

## Pillar Data's Application Aware Storage

**Date:** February, 2008

**Author:** Mark Peters, Analyst and Steve Duplessie, Senior Analyst

**Abstract:** Pillar Data Systems has spent a great deal of time and money designing a storage architecture that can mesh with the specific needs of varying applications—and is now set to unveil the first proof points.

### Background: How and Why Pillar Arrived at This Point

One of Pillar's assertions from the start has been that it is on a path to 'make storage less exciting.' Of course, that's marketing-speak, but it's actually not a bad encapsulation of what the company has been doing—as long as you understand that being less exciting has nothing to do with being less capable. Indeed, it has everything to do with adding so much automation and built in 'smarts' that the system only *looks* dumb because it doesn't need a cadre of specialists to tune and tweak it in order to optimize it.

It also has a lot to do with the very essence of why Pillar exists. The company began as Digital Appliance about 10 years ago and it was (and is) funded by Oracle's Larry Ellison, who wanted to bring "thin computing" capabilities to the mass market. What's the relevance to the super-smart and flexible Pillar products? Well, from the perspective of a dominant database provider, 'thin computing' equals a huge back-end and a very thin client—which makes perfect sense since a huge scalable backend can run an equally huge, and growing, database. Make it simple for that huge database to work, to be optimized and to grow on an affordable and easy-to-use server with a dedicated storage platform and there's a major business opportunity.

The 'skunk work' R&D project toiled around for several years without any significant results as the market headed the other way—towards more distributed computing. When it once again caught Ellison's attention, he brought in some friendly and trusted industry veterans to evaluate what had been achieved, and discovered that part of the goal—that of a seamlessly scalable commodity based back-end super server—was becoming less likely...*but* also found that the storage initially created to accompany and support that server had considerable promise. There followed a period of more focused R&D—and Pillar was born. Launching its Axiom line of storage systems in the second half of 2005, Pillar hit 100 customers in under a year and has been growing market traction and acceptance ever since—today, its base exceeds 300 customers and well over 500 systems.

The Axiom system has found market favor—and a loyal customer base—based on all the expected, impressive basic speeds and feeds which have been supported from the outset by a solid set of practical capabilities (see sidebar: Pillar: The Pedigree and the Product).

Pillar has never been short on features and functions, and continues to add and improve. This has helped assure customers that Pillar is in the game for the long haul; strong financial backing and customer growth support that. Indeed, in so many respects, Pillar could never be viewed as 'your standard storage 'start-up.' It had mega-backing from the start, which has been combined with the patience required to get things right. For example, Pillar built a world class test and support structure *in advance* of an installed base. We can't think of any venture-backed entity that has ever had that luxury. Patience and deep pockets have allowed the company to focus on putting the pieces in place to fight the battle on the next frontier—a Web 2.0 world where all the high-end requirements of old are mixed with the new attributes of the day, such as infinite effective scale, heavy automation, massive intelligence and ultra-simplistic management. Supporting NAS and SAN in one device was hardly the norm and yet that 'complex simplicity' was what was built into 'Slammers' (storage controllers), 'Bricks' and 'Pilot' (the management system). Pillar's AxiomONE software suite allowed ILM to become more than just a promising three letter acronym, providing high-end functions in order to dramatically improve internal IT service offerings while simultaneously lowering cost and complexity.

With a pedigree that proves its ability to deliver impressive functionality, the company is now turning its attention to the wider compute ecosystem and is driving its original concept of ‘Application Aware Storage’

### ‘Application Aware’ Defined

Perhaps the best way to explain this is to first establish the decades old standard approach of “Application Unaware Storage;” CPUs are CPUs, applications are applications and storage is storage. The disk subsystem is storage...period. It stores and delivers whatever is sent its way or requested, at whatever the service level it has available at that particular moment dictates. It is the ‘peanut butter’ approach—an equal spread for every one and every application. As data processing demands have become more dynamic, very often we set up autonomous stovepipe systems—with their own storage—to handle certain applications. In terms of data storage, such an approach can be effective in getting the job done, perhaps, but is certainly not the most efficient in terms of resources or management. It also tends to be over-provisioned and rarely changes regardless of the effects of eventual data growth. Even if all of the infrastructure were perfectly tuned to an applications requirements on day one, by day 1000 those requirements will be different.

And, while distributed processing has been found to make sense in many situations, invariably distributed data storage—i.e., data—has not. As storage is

consolidated, we have had to develop new and increasingly intricate means to ensure that the varying demands of different users and applications are met. Storage administrators, if they can ever find the time, are forced to create different pools, perform manual tiering, over-provision capacity to ensure service levels, and so on. All we’ve really achieved is to move the problem and superficially change the symptoms—the cause is unchanged and the results are often inevitably ineffective. In fact, one could argue that the situation is deteriorating, as manifested by the widening disconnect between the applications owners and storage owners in many organizations. One promising source of improvement—to help close that gap—is the concept of application awareness.

‘Application awareness’ means that the storage system, and some day other areas of the infrastructure stack, recognizes the application that is using it and reacts (dynamically, ideally) by presenting and delivering different capabilities to that application, thus providing the best possible match between the applications’ needs and the storage devices’ capabilities. This is valuable in any situation, but is especially crucial in a mixed workload environment where multiple application types are concurrently accessing the storage system. This is currently the norm in data centers and will only increase in heterogeneity and cause more variable outcomes as virtualization and consolidation efforts multiply.

### Pillar: The Pedigree and the Product

- **Ease of Management:** in an early realization of its importance, ease of use included the obvious - installation and volume management as well as virtually eliminating performance tuning and capacity optimization
- **SAN and NAS in one:** can only improve resource optimization and consolidation
- **Clustered Architecture:** just-in-time scalability to flex in line with business demands
- **Large File System:** simplifies management and improves workflow
- **Virtualized Capacity:** the single pool helps with everything from performance and utilization to volume expansion and data transfer
- **Quality of Service:** this was Pillar’s initial foray into addressing the fact that not all data is created equal and therefore, different data requires different storage attributes. QoS permits specific values to be assigned to specific volumes in order to optimize known application performance needs
- **Advanced Data Management:** such as differential snapshots, file replication and volume-based remote mirroring
- **Rapid RAID Rebuilds:** achieved because rebuilds are localized in ‘Bricks’ (storage enclosures with disk drives and RAID controllers), and thus reducing users’ potential exposure to data loss
- **Axiom Tiered Storage:** beyond merely supporting different drive types, the system also uses knowledge of the disk geometry to place data depending on its QoS
- **AxiomONE:** a software suite that makes the promise of ILM a simplified reality—despite the elaborate nature of the tools themselves, they remain easy to use. Standard dynamic performance and capacity management can be enhanced with WORM, replication, CDP and VTL capabilities.

By definition, 'Application Aware' anything is hard to deliver upon— which is why almost everyone side-steps the problem and generalizes storage service levels in order to try to homogenize the results—'peanut-butter,' spread evenly so that it's roughly good enough for everyone and every application. Once again, it can be effective, but it's not likely to be efficient and almost never is it as good as it can be.

Given its architecture and approach, Pillar had always intended to marry the abilities of the storage to the specific requirements of the application itself. Once again, Pillar is not following the crowd. Many storage players have been creating very generic devices; some—it must be said—with very impressive raw capabilities in terms of performance, capacity or functionality. However, other than "qualifying" that such systems work, there hasn't been a great deal of focused attention paid to specific performance-requirement-matching. Nonetheless, it remains a simple truth that not all applications are the same—they write, read and manipulate data differently. For example, some are random access and small block, and some are sequential large block applications. Moreover, just about all storage uses 512k blocks, but not all applications are designed for that block size. Being generic is logical if you design for the mass market, as the cost and complexity of designing "one off" systems has proven economically disastrous historically. Pillar seeks to offer universally accepted storage systems that also provide the ability to customize all the way down to the specific requirements of individual applications—within a mixed application environment.

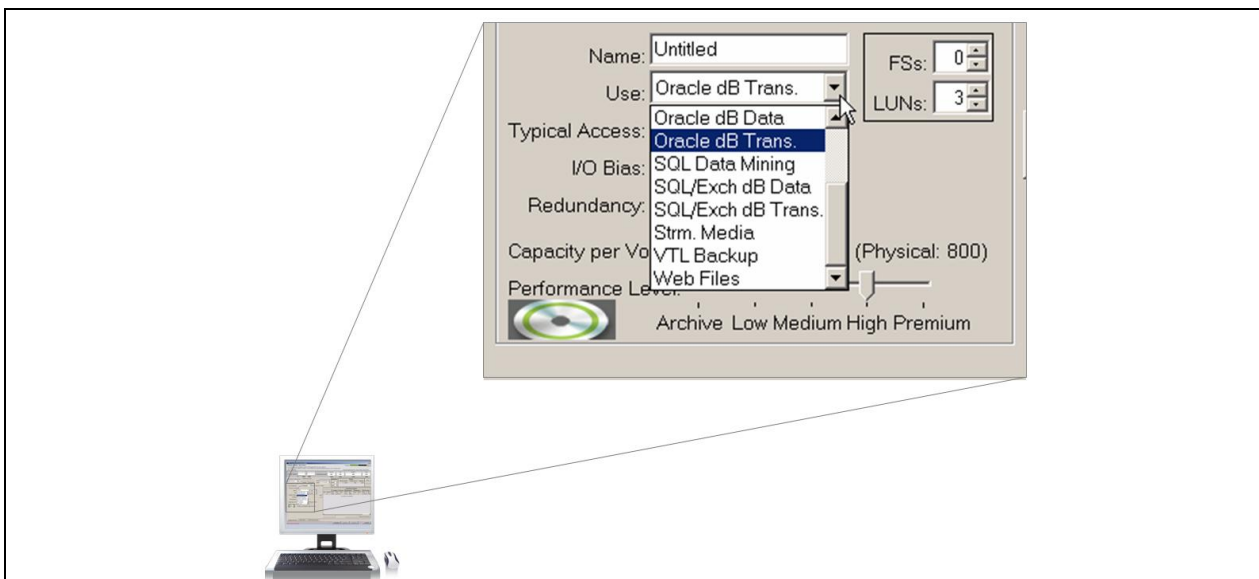
### The Pillar Strategy

To take advantage of the opportunity, Pillar designed its systems with the flexibility to meet the specific requirements of individual application types. The company realized that databases do not act in the same way—or have the same needs—as, for example, e-mail systems or streaming media systems. Additionally, Pillar recognized that specific performance requirements would need to be combined with point-in-time value—understanding that nothing should be forced to be treated the same way forever. Pillar has always had the ability to alter overall performance and cost metrics based on value by tiering to lower cost, higher density devices. This is the QoS (Quality of Service) metric. Now Pillar intends to add specific performance and availability attributes as dictated by the application itself.

### How is Application Awareness Applied?

Application awareness sounds good, but how does it work? While there's a lot of magic going on behind the scenes, at a high level it's actually rather simple. Administrators now have the capability to define storage volumes based on an application profile as shown in Figure 1. As you can see, Pillar has created profiles for a number of common applications including Oracle, SQL Server, Exchange and VTL backups. Profiles define the behavior of the application based on attributes like Typical Access, I/O Bias, Redundancy, Capacity per Volume, Performance Level, etc.

**FIGURE 1. CONFIGURING APPLICATION AWARENESS**



These profiles provide the awareness that Pillar needs to optimize the performance and availability of applications. To illustrate how this works, consider the difference between the access patterns of a Microsoft Exchange database compared to a VTL backup volume. The I/O for an Exchange database is mostly random reads that are accessed in relatively small chunks (4 KB or 8 KB depending on the version of Exchange). A VTL backup volume sees mostly writes that are relatively large (e.g., 256 KB). Based on this deep knowledge of how the application does I/O, application awareness can now be used to tune and prioritize the performance and availability of applications to meet the needs of the business.

Some expectations of where this might work well are obvious: right back to Pillar's beginnings, one has to expect an Oracle story here, especially given the wide variations of storage demands that can be associated with databases. Others are as much to do with the changing face of processing as that of storage; thus, for example, the ability to mix, move and migrate applications within and across virtual servers (using VMware and the like) cries out for an exceedingly, and dynamically, tuned storage system. Infrastructure virtualization employed to create higher utilization and asset consolidation also brings great unknowns. If you collapse 59 physical machines onto 7 in a virtual world, the overall makeup of the application environment changes dramatically. By default, bandwidth requirements to and from storage will change, but so will the overall consolidated stream of requests. Fifty-nine low-use database applications spread across fifty-nine machines can turn into an ultra-high transaction system when collapsed onto 7 machines. Exchange, CRM, finance, code management and all other corporate application environments will act differently when virtualized and concentrated—just *how* they will act is an unknown.

That is exactly what Pillar is counting on—if the company can adapt to changing environments and maintain optimized data flow, it can take a huge risk out of IT, accelerate time to deployment and enable faster adoption of consolidation technologies such as server virtualization. By providing multiple protocol support and the ability to morph as requirements dictate, Pillar can show that it is a company ready for the next wave of IT requirements.

## The Bottom Line

With its introduction of Application Aware Storage, Pillar is challenging a couple of notions that have been around for a long time. No longer would data storage systems have to be merely 'equal-opportunity-infrastructures.' And no longer would data storage devices be 'just' peripherals. Storage would be an integral part of the entire infrastructure and application solution because anything that increases the ability to apply the right tool to the right job in a data center is not only a good thing—it is a vital thing. For a storage system to 'adjust' flexibly in order to do that based on an awareness of the 'job' is a significant forward step. Indeed, it's one of those attributes that will allow IT to willingly say 'yes' more often, knowing that their infrastructure has the flexibility to cope with unpredictable application and workload changes and alterations (even transient changes, which are much harder to deal with) as well as pure growth.

Of course, with anything new there is risk, but Pillar's pedigree and backing are unlike anything that has come before. Mr. Ellison has shown exceptional patience—it isn't as though he needs the money—in order to let Pillar try to solve stunningly hard problems. We have no guarantee that the company will successfully execute on the strategy, but early proof points are encouraging, and its market success with Axiom to date has shown that Pillar has the wherewithal to build and support products in this market. Years of development with this goal embedded at each step also lends confidence. The company was set up to change the playing field and 'to make storage less exciting.' It does not take much to envision just why that would be important to someone such as Ellison—Oracle is the major management tool of most of the world's mission critical data and it most likely irritates him that some "dumb box" at the end of the wire garners such a large piece of the budgetary pie. If Pillar can succeed, not only is it a win, but the balance of "value" in the chain moves firmly back to that data management layer.

Finally, the giants of the storage business clearly recognize the potential threat of any company backed by Ellison, and watch Pillar making inroads. The company is on the hit list of most of the major players—for obvious reasons. The reality is that the market dynamics are such that Pillar doesn't need to take share from traditional players to ultimately be successful. This new world of digital content growth—and all the new characteristics it brings with it—is so large and so rapidly growing that for the foreseeable future, everyone can participate and prosper. Eventually, the new characteristics of the Web 2.0 world, applied to commercial enterprises, will "trifurcate" and a new class will emerge—a class that dwarfs both traditional transaction-based computing as well as the distributed

processing based computing data generations. It's the next next thing in our world of data, and that's what Pillar has set its sights on.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.