


**WHITEPAPER:**

# Understanding Pillar Axiom Data Protection Options

## Introduction

This document gives an overview of the Pillar Data System® Axiom RAID protection schemas. It does not delve into corner cases or exceptions on functionality outside of what is considered to be “normal” operations. On the Pillar Axiom® storage system, data sets are protected with hardware RAID and striped across multiple disk drives for performance. These data sets can also be doubled for additional levels of data protection.



## Axiom Hardware Components

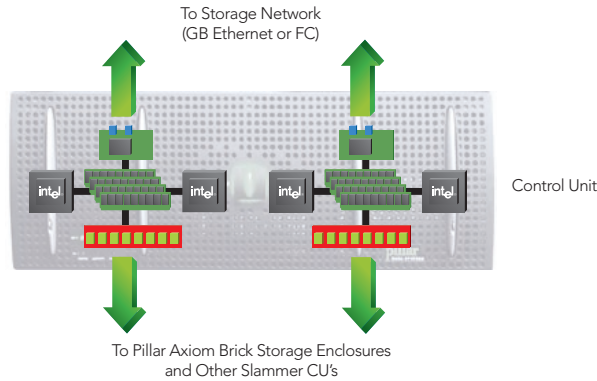
This section defines the various hardware components of Pillar Data Systems Axiom storage system.

### Pillar Axiom Pilot Policy Controller

The Pilot is an out-of-band management system that directs and manages all system activity. It has no connection to the data path. It communicates with the rest of the Axiom system over the private Ethernet network. Each Axiom system has one Pilot, which includes a Graphical User Interface (GUI) for user configuration and management of the Axiom system; a Command Line Interface (CLI) used in scripts for user configuration and management of the Axiom system; two independent Control Units (CUs) that operate in active/passive mode. While the pilot contains redundant components, it is not considered a critical piece of the data protection puzzle.

### Pillar Axiom Slammer Storage Controller

The Slammer provides an external interface to a storage network. It processes and manages every Input/Output (I/O) request. An Axiom system can include NAS and SAN Slammers. SAN slammers can be either Fibre Channel (FC), iSCSI, or both. The current architecture allows up to two Slammers per system with future plans for expansion up to four.




A Slammer contains two CUs that function as an active-active pair, and are configured identically. The primary components of a CU are a pair of network interfaces (GbE for NAS, iSCSI SAN, or FC for SAN), CPU and memory resources, and an integrated FC switch.

Slammers are connected to the system storage pool through eight Fibre Channel ports on imbedded Fibre Channel loop switches. Bricks connect to Slammers through a switched Fibre Channel fabric. This is a private Fibre Channel fabric called the System Storage Fabric or SSF. The Slammers then access and virtualize the storage pool, so you can easily increase the number and size of file systems and LUNs.

A primary job of the Slammer is to store I/Os from hosts in cache and then de-stage them to the backend disk pool which is composed of multiple Bricks. Each segment of data written by the Slammer is striped across multiple drives in the intelligent storage enclosure called Bricks.

**Pillar Axiom Brick Storage Enclosure**

There are two types of Bricks: SATA Bricks and FC Bricks. For the purpose of this paper we will discuss the SATA bricks, however FC bricks have similar properties. SATA Bricks have 13 disk drives where the 13th drive is used as a hot spare for automatic failover. The remaining 12 disks are pre-configured as two 6 drive RAID-5 groups (or disk arrays) with each controlled by the primary controller as well as the secondary backup controller. The Brick protects data from drive failures. Drive rebuilds for current drive technologies are under 4 hours, eliminating the need for more costly RAID protection like RAID 6 or RAID 1+0 (also called RAID 10).



The Brick is fully redundant internally. Each hot-swappable RAID controller acts as the primary controller for one of the two RAID-5 arrays and has a path to each one of the disk drives inside the Brick. If one controller fails, the surviving controller continues to process I/Os for both RAID arrays. The current architecture supports up to 64 Bricks per system.


## Understanding Pillar's RAID Options

Now that we've defined the basic components of an Axiom, we can discuss how Pillar's Standard and Double Protection levels protect customers' critical information against data loss.

### Standard Protection

Standard Protection, Pillar's default level of protection, is RAID 5+0 (also known as RAID 50). RAID 5 includes both parity and disk striping across multiple drives. The Disk Array is RAID 5. A Pillar virtual LUN is striped (RAID 0) across multiple disk arrays. RAID 50 provides increased write performance, improved data protection, and faster rebuild times even in the event of a drive failure. During a drive rebuild, system performance remains high because the other distributed RAID 5 arrays are functioning fully. Other architectures place a strain on the I/O controller to support RAID rebuilds, thus impacting overall performance and lengthening rebuild times. The figure below shows how data sent from the Slammer is distributed across two bricks which are comprised of four RAID 5 disk arrays for a total of 24 drives. Standard protection is typically striped over four disk arrays, so two SATA Bricks are required.


A LUN or Volume that uses Standard Protection across two bricks can survive the simultaneous loss of four drives (one per RAID group) with no loss of data. Two of the failed drives (one per Brick) will rebuild, in parallel, in less than 4 hours to the dedicated hot-spare in each Brick. The disk arrays that successfully completed the rebuild process could then lose another drive each and still have no data loss. It would take a seventh drive failure to have data loss. A double drive failure in any of the disk arrays in a 4 hour window or a catastrophic failure of a Brick can cause data loss; both of these scenarios are extremely unlikely to happen.



**Double Protection**

For enhanced data protection, Pillar offers Double Protection also known as RAID 50+1. With Double Protection, Pillar mirrors each stripe of data written to a disk array to another disk array on another Brick. Double protection requires a minimum of two Bricks.

The figure below shows how the data from the Slammer will be written using Double Protection. Note how a copy of Stripe A is written to disk array 1 on Brick 1 while its mirrored copy is written to disk array 3 on Brick 2. Also note that with Double Protection, an entire Brick can fail and no data will be lost.



## Conclusion

Losing two drives in the same disk array in less than 4 hours (the amount of time needed to rebuild a failed 500GB SATA drive) is extremely rare. With drives becoming more and more reliable and with Pillar continuing to decrease rebuild times, the likelihood of a double drive failure is becoming even more remote. The same is true with losing a complete Brick: It rarely, if ever, happens. Pillar allows end-users to set the protection level of a virtual LUN or Volume to survive a multiple drive failure or even a double Brick failure. Other vendors require the use of host-based RAID protection in combination with their controller-based protection, introducing more points of failure, more complexity, and more cost to provide these levels of protection. Pillar is the clear leader when it comes to protecting user's data.

Pillar Data Systems takes a sensible, customer-centric approach to networked storage. We started with a simple, yet powerful idea: Build a successful storage company by creating value that others had promised, but never produced. At Pillar, we're delivering the most cost-effective, highly available networked storage solutions on the market. We build reliable, flexible solutions that, for the first time, seamlessly unite SAN with NAS and enable multiple tiers of storage on a single platform. In the end, we created an entirely new class of storage. [www.pillardata.com](http://www.pillardata.com)

