An Oracle White Paper
May 2012

Oracle's Pillar Axiom Storage — A Better Choice Than EMC

A Comparison of Oracle's Pillar Axiom 600 Storage System with EMC's VNX Family

ORACLE

Executive Overview

Oracle’s purchase of Pillar Data Systems in July 2011 brought together a valuable Storage Area Network (SAN) solution with an industry-leading software company. The Oracle Pillar Axiom 600 storage system addresses some key challenges that competing storage vendors have yet to overcome. For organizations running Oracle software, the Pillar Axiom 600 is an ideal fit for consolidating storage for multiple Oracle Databases or Oracle Applications. It is also a good choice for EMC CLARiiON customers nearing the end of a lease agreement. Rather than upgrading to the EMC VNX product family, this juncture is the perfect time to make the switch to Oracle's Pillar Axiom 600.

This paper compares Oracle's Pillar Axiom 600 storage system with the EMC VNX family of storage products and describes how the Pillar Axiom 600 system can help organizations reduce costs and increase agility, while meeting the demanding performance requirements of today's virtualized application and server infrastructures.

The Pillar Axiom 600 system was designed from the ground up to deliver deterministic Quality of Service (QoS) consistent with the business value of the application and its data. It eliminates the requirement to overbuy and underutilize storage to buffer against application interference in a multi-tenancy shared storage system. This enables organizations to buy the capacity they require today, scale up when needed, and conserve capital in the process — all while delivering consistent performance under load.

When compared directly with the EMC VNX family of storage products, the Pillar Axiom 600 system offers the following key advantages:

- **Investment protection through flexible multidimensional scalability**
  The Pillar Axiom 600 storage system is designed to scale smoothly in both capacity and performance without the disruption of a platform migration or the cost of increased software license fees. The EMC VNX product family, on the other hand, requires that customers scale in fixed increments. This can result in costly, highly disruptive, and risky replacement efforts when storage requirements surpass the capabilities of a given storage system model. The modular architecture of the Pillar Axiom 600 enables customers to start small and scale across the entire family of EMC VNX systems, thus protecting investments and avoiding the disruption of replacing outgrown storage systems.

- **Up to 2x better storage efficiency with performance**
  Industry averages for storage capacity utilization, also called storage efficiency, hover around 40%[1] as companies use performance solutions that require excess capacity. The Pillar Axiom 600 system is based on a unique storage architecture that enables performance to scale along with capacity. Many Pillar Axiom 600 customers see storage utilization rates up to and often exceeding 80% with no loss of I/O performance. This is approximately double the industry average for efficiency with traditional storage arrays, such as EMC's VNX family.

- **Up to 30x compression for better storage efficiency with Oracle Database**
  Oracle Database 11*g* includes a high-efficiency data compression feature called Hybrid Columnar Compression that works only with Oracle storage. EMC customers who are currently running read-oriented Oracle Database environments with EMC VNX or CLARiiON systems can see an immediate 3x to 5x reduction in storage footprint and related costs by migrating to the Pillar Axiom 600. The example scenario cited in this paper shows more than a 70% reduction in required storage capacity when using Hybrid Columnar Compression.

- **Up to 63% cost savings for Oracle Database and up to 41% savings via consolidation**
  The Pillar Axiom 600 offers as much as 63% cost savings over EMC for a large Oracle Database deployment through support for Hybrid Columnar Compression. The ability to consolidate storage workloads that would otherwise have been isolated on multiple EMC VNX systems also provides as much as a 41% cost savings over EMC.

Consistent SLA performance in consolidated environments

The application-aware Pillar Axiom 600 storage system enables deterministic performance in consolidated environments, even under heavy loads. EMC customers are often forced to spend more by purchasing separate storage systems for each business application so that performance will not be impacted by I/O contention from other applications. By contrast, the Pillar Axiom 600 with QoS has been designed to effectively minimize I/O contention, enabling multiple storage workloads to be consolidated on the same storage system without undue risk of breaching application service level agreements (SLAs). This consolidation enables greater utilization of the storage system and significant cost savings.

---

[1] Source: "Wasted space: IT aims to fill disks," Computerworld, August 9, 2010, *http://www.computerworld.com/s/article/351035/Wasted_Space_IT_Aims_to_Fill_Disks*.

Investment Protection Through Flexible Multidimensional Scalability

Today's enterprise servers have incredible processing power and can sort through queries against massive data warehouses faster than most storage systems can keep up. The combined trends of increasingly powerful servers and continually dropping disk drive prices have naturally resulted in an increase in the amount of data that is stored and processed. An important effect of these trends is that today's storage systems need to scale in multiple dimensions. They must be able to scale capacity to hold more data and also scale performance to avoid storage I/O bottlenecks that can slow down performance of enterprise applications.

Many storage vendors, including EMC, still design their storage product families in such a way that customers must scale in fixed increments that typically involve a model upgrade. This results in costly, highly disruptive, and risky replacement efforts when storage requirements surpass the capabilities of a given storage system model.

The Pillar Axiom 600 storage system is designed to adapt and accommodate changing business needs, thus protecting investments in storage technology. The Pillar Axiom 600 provides a long-term, flexible platform that can be easily expanded over time to accommodate increased requirements for both capacity and performance.

Where Traditional Storage Architectures Fall Short

Traditional storage architectures were not designed to support the level of I/O throughput that is being required in today's data centers. According to industry analysts, most storage systems only reach about 40% of their total capacity.[2] The practices of adding extra spindles and placing data at the outer edges of drives to boost performance means that excess capacity is required.

Because of the performance limitations of these traditional storage systems, many organizations use work-around solutions to help boost I/O performance. Common approaches include increasing the spindle count even though additional capacity is not needed and placing data on the outer edges of disk drives where drive performance is greater. While these approaches do indeed boost performance, they have a very damaging effect on storage efficiency. They result in disk drives that are only partially utilized, so organizations have to purchase more capacity than would otherwise be needed, thus reducing return on investment. There is also a limit to how far performance can be stretched using these methods. Because of the way that these traditional storage systems are architected, their performance will not scale beyond a certain point even when more drives are added.

As shown in Figure 1, traditional storage array architectures have been based on a pair of storage controllers and matching redundant disk trays. Each storage controller typically has one RAID controller and the architecture prevents adding more RAID controllers or storage controllers when

[2] Source: "Wasted space," *http://www.computerworld.com/s/article/351035/Wasted_Space_IT_Aims_to_Fill_Disks*.

more disk drives are added. EMC VNX systems, for example, support a maximum of two RAID controllers, one in each storage controller.

These architectures become stressed when their disk trays are fully populated and capacity utilization begins to reach the practical limit of around 40% as mentioned above. When this point is reached, storage I/O performance begins to be limited by the storage controllers or RAID controllers.

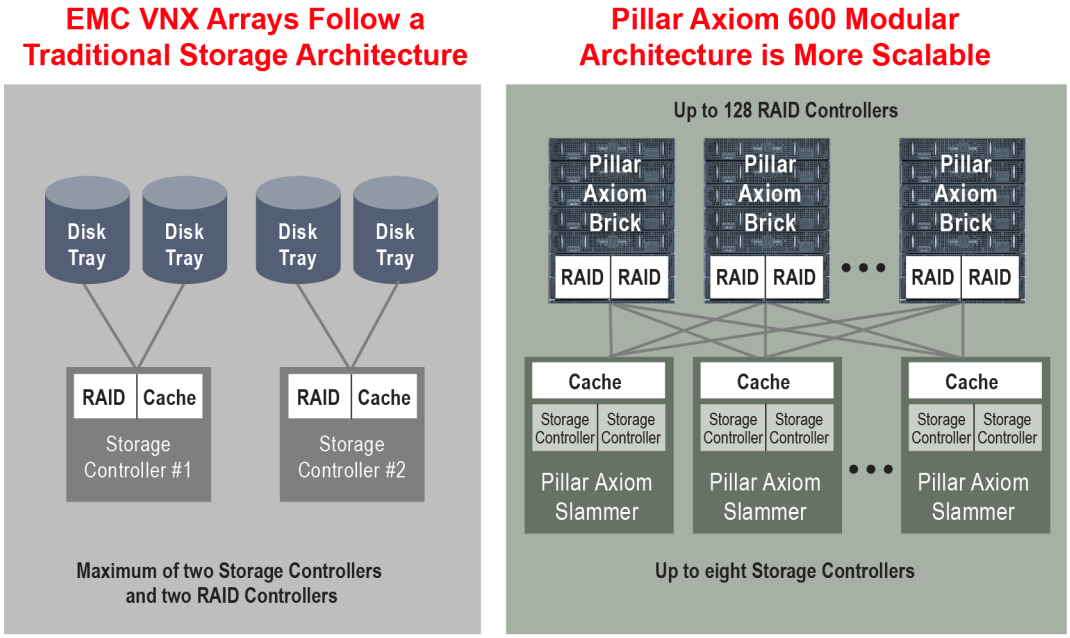

Figure 1. Architecture comparison of a traditional storage array such as EMC VNX versus Pillar Axiom 600.

The right-hand side of Figure 1 shows a very different modular architecture in the Pillar Axiom 600. This architecture supports scalability across multiple dimensions, enabling the core I/O infrastructure to scale along with capacity expansion. The following subsections describe in more detail the Pillar Axiom 600 system architecture and the advantages it has over EMC VNX systems, including the ability to offer linear performance scalability all the way from small departmental systems to large enterprise solutions without replacing the storage system.

Pillar Axiom 600 Modular Architecture

The Pillar Axiom 600 has no single point of failure. The architecture is based on a set of modular building blocks as follows:

- **Pillar Axiom Brick** — A Pillar Axiom Brick is a fully redundant storage enclosure that is used as a modular building block to create a shared storage pool. The Pillar Axiom 600 supports up to 64 Bricks. Two embedded active/active RAID controllers are combined with a tray of disk drives, enabling distributed RAID processing and offloading this processing from the storage controllers. The Pillar Axiom Brick comes in three flavors to support Fibre Channel drives, SATA drives, or solid-state drives (SSDs). Each Brick also includes a hot-spare drive.

- **Pillar Axiom Slammer** — The Pillar Axiom Slammer includes two active/active storage controllers and the Pillar Axiom 600 supports up to four Slammers (for a total of eight storage controllers). The Slammer's primary function is to move data across the SAN (or NAS) connections to and from the Pillar Axiom Bricks. There are different versions of the Pillar Axiom Slammer to support 4 Gbit/s or 8 Gbit/s Fibre Channel SAN, 1 Gbit/s iSCSI SAN, a combination of Fibre Channel and iSCSI SAN, or 10 Gbit/s Ethernet NAS connectivity. . The Pillar Axiom Slammer uses a private switched SAN as the back-end network to connect the Slammers to the Bricks
- **Pillar Axiom Pilot** — The Pillar Axiom Pilot controls all system configuration, management, and monitoring within the Pillar Axiom 600 system. It consists of two independent control units that operate in an active/standby mode and support automated failover. Communication to the Pillar Axiom Slammers is achieved over an Ethernet-based, out-of-band private management interface (PMI). Management of the Pillar Axiom is accomplished via a Java-based GUI that connects to the Pilot over Ethernet or a full-featured command line interpreter (CLI).

This modular architecture enables organizations to mix and match building block components to create configurations that closely match both their performance and capacity requirements. More importantly, it enables easy scalability so that organizations can meet their ever-changing storage needs without having to do a rip and replace upgrade as would be required with EMC VNX systems. As storage requirements grow, simply add more Pillar Axiom Slammers and Pillar Axiom Bricks to meet the desired performance and capacity.

Distributed RAID Architecture Offers Superior Performance Scalability

EMC integrates its RAID controllers into the VNX storage controller units. Thus, additional RAID controllers cannot be added within the EMC VNX platform no matter how many disk drives are added. Even the VNX 7500, EMC's largest model, has only two RAID controllers and two storage controllers to serve up to 1,000 disk drives.

The distributed RAID architecture of the Pillar Axiom 600 provides a superior approach. With up to 128 RAID controllers (two per Brick) and up to eight storage controllers (two per Slammer), I/O performance can be scaled linearly along with capacity. This means that the ratio of RAID controllers to disk drives and capacity can remain constant from small configurations up to very large configurations. Similarly, the number of disk drives that must be managed by a single pair of storage controllers never has to exceed 208. By contrast, the maximum capacity the EMC VNX 7500 configuration supports is 1,000 disk drives with a single pair of storage controllers, nearly a 5x difference.

With multiple storage controllers and many RAID controllers all operating in parallel, the Pillar Axiom 600 system has a robust I/O infrastructure that can scale I/O performance so that disk drives can operate at about 2X the utilization of traditional storage architectures, such as the EMC VNX family.

Performance Advantages of Large DRAM Cache

One of the key differences between the Pillar Axiom 600 and the EMC VNX family is the cache architecture. The Pillar Axiom 600 system supports a very large DRAM cache rather than using a

secondary cache based on solid-state drives (SSDs) like EMC's FAST cache. The Pillar Axiom 600 supports a DRAM cache up to 192 GB, which is 4x the memory capacity of the largest EMC VNX configuration.

For almost all workloads, a large DRAM cache will outperform an architecture based on a more limited DRAM cache coupled with a secondary SSD cache. This is because the I/O write performance of DRAM is approximately 1,000 times that of enterprise SSDs. In addition, a secondary cache essentially doubles the amount of data that must be written to cache whenever there is a cache miss. With a large enough DRAM cache, the number of cache misses is usually very limited, thus enabling extremely fast performance.

The large DRAM cache or memory capacity in the Pillar Axiom 600 thus offers a good architecture for performance-hungry applications and is part of what enables the Pillar Axiom 600 to achieve twice the capacity utilization of traditional storage architectures. As discussed in the section titled "Cost Savings — Up to 2x Better Efficiency," this utilization advantage also helps saves costs. Thus, the DRAM cache of the Pillar Axiom 600 can be more cost-effective than the seemingly less expensive SSD-based FAST cache option in EMC VNX systems.

Measured Performance Scalability with Pillar Axiom 600

Oracle's testing of various Pillar Axiom 600 configurations with SATA drives in a RAID 5 configuration showed linear performance scalability of both random reads and random writes. The results are illustrated in Figure 2 and demonstrate that the Pillar Axiom 600 provides an excellent growth path from small systems that compare to the EMC VNX 5300 all the way up to large enterprise systems that approach the capacity of EMC's largest configuration, the VNX 7500.

An important note about these performance tests is that random read performance for the largest configuration showed 114 read IOPS per drive, which is essentially the limit of the drive itself. This proves that the modular architecture of the Pillar Axiom 600 provides a scalable I/O infrastructure that will not become a performance bottleneck as in traditional storage architectures, such as the EMC VNX product family.

The Pillar Axiom 600 system has also achieved outstanding results on industry standard benchmarks. A half rack Pillar Axiom 600 storage system achieved a performance result of 70,102.27 SPC-1 IOPS and an SPC-1 price-performance rating of $7.32/SPC-1 IOPS.[3] EMC does not publish benchmark data so the reader is urged to ask EMC if they can provide comparative performance data for a similarly priced configuration to enable an apples-to-apples comparison.

[3] Pillar Axiom 600 SPC-1 benchmark results as of April 19, 2011, are available at http://www.storageperformance.org/results/. (U.S. list price as of April 19, 2011. Pricing subject to change without notice.)

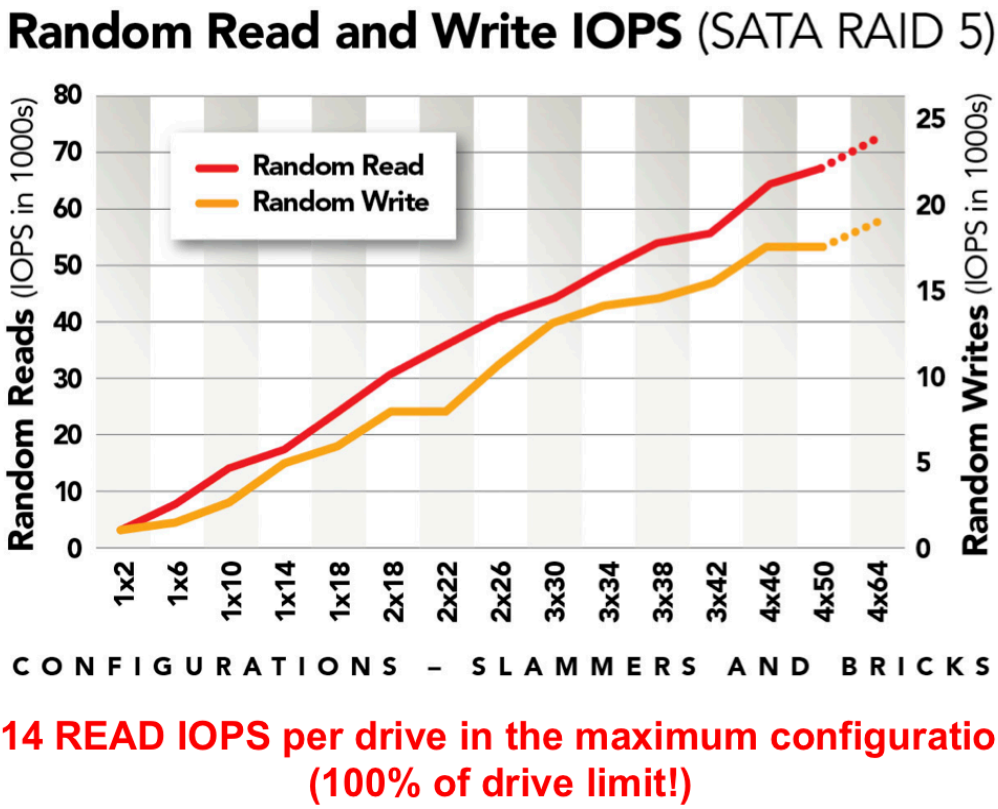

Figure 2. The Pillar Axiom 600 shows linear scalability across the entire range of configurations.

Summary of Configuration Options on EMC VNX and Pillar Axiom 600

Table 1 provides a summary of the configuration options across the EMC VNX product family in comparison with the Pillar Axiom 600. The modular architecture of the Pillar Axiom 600 enables this single system to offer configurations that compare against the entire family of EMC VNX systems. This enables organizations to protect their storage investments as they scale from small to large configurations. In contrast, when an organization has outgrown an EMC VNX 5300 or VNX 5500 model, they must disrupt their storage environment to rip out the old system and replace it with the next larger EMC VNX model.

As noted above, the most dramatic advantages for the Pillar Axiom 600 include:

- Up to 4x the number of storage controllers
- Up to 64x the number of RAID controllers
- Nearly 4x the memory or DRAM cache size of the largest EMC VNX configuration

TABLE 1. EMC VNX VERSUS PILLAR AXIOM 600 CONFIGURATION OPTIONS

	EMC VNX 5300	EMC VNX 5500	EMC VNX 5700	EMC VNX 7500	PILLAR AXIOM 600
MAXIMUM NUMBER OF STORAGE CONTROLLERS	2	2	2	2	8
MAXIMUM NUMBER OF RAID CONTROLLERS	2	2	2	2	128
MAXIMUM DISK DRIVES	125	250	500	1,000	832
MEMORY	16 GB	24 GB	36 GB	48 GB	Up to 192 GB (48 GB per Pillar Axiom Slammer)
MAXIMUM RAW CAPACITY	360 TB	720 TB	1.5 PB	1,974 TB	1.66 PB (Requires two or more Pillar Axiom Slammers)

Cost Savings — Up to 2x Better Storage Efficiency

A recent study conducted by the Enterprise Strategy Group (ESG) revealed that — for the fourth consecutive year — reducing costs is the business initiative most frequently impacting technology purchases. Specifically, 42% of organizations indicated that cost reduction would be among the key business initiatives guiding IT spending decisions over the course of 2012.[4]

Cost is an obvious business driver and today's IT managers are looking for storage solutions that can help them drive down storage costs. The Pillar Axiom 600 offers the following cost saving advantages over the EMC VNX family:

- **Twice the efficiency with higher performance** — The same SLA performance targets can be met with lower cost systems because the Pillar Axiom 600 enables up to twice the storage efficiency while also delivering high performance. This means that more data can be stored on the Pillar Axiom 600 without having to purchase additional unneeded capacity just to meet performance targets.
- **Lower OPEX** — Up to 2x better storage efficiency means that half as many disk drives are required. This helps reduce administration and management costs as well as power and cooling costs.
- **Consolidated storage** — The Pillar Axiom 600 provides assurance that business-critical applications get high priority so that SLAs can be met in a cost-effective consolidated environment.

[4] Source: ESG Research Report, *2012 IT Spending Intentions Survey*, January 2012.

- **Ease of Management** — The Pillar Axiom 600 is a very easy-to-use system that further decreases the management cost because fewer administrators are needed to manage a given environment.

Storage Efficiency Is Only Meaningful When Considered with Performance

Nearly every storage vendor claims to offer better storage efficiency. However, only the Pillar Axiom 600 can claim up to 80% storage efficiency **with no decrease in I/O performance**. This is double the industry norm with traditional storage architectures such as EMC VNX systems. Twice the efficiency with performance means that Pillar Axiom 600 customers can purchase systems that have only half as many disk drives. In essence, the Pillar Axiom 600 customers get to use more of what they buy.

EMC emphasizes compression and deduplication technologies as a means for their customers to improve storage efficiency by storing more data in each GB of disk space. However, these technologies are intended for situations where performance is not a critical requirement such as when storing user desktop data or operating a disk-to-disk backup environment. For systems that are already under-performing in terms of IOPS, customers cannot take advantage of compression or deduplication because that would just make matters worse. These technologies put another strain on already strained and limited storage processors. Furthermore, even if more data is packed into each GB of storage, performance bottlenecks are still likely to occur when the system reaches 40% of total capacity. Thus EMC's deduplication and compression don't extend the 40% of usable capacity; they just enable more data to be stored in that 40%.

Lower OPEX Costs

In many cases, storage purchase decisions put too much emphasis on the initial purchase outlay, or CAPEX, and not enough emphasis on ongoing OPEX, partly because OPEX is harder to quantify. Yet OPEX can be a significant component of overall TCO. It is therefore important to consider the following ongoing costs with as much accuracy as possible:

- Data center floor space
- Energy consumption
- Administration costs for provisioning and managing storage

The 80% utilization rate often achieved with the Pillar Axiom 600 has a big effect on operating expenses in addition to initial purchase costs. Disk drives consume the same amount of power whether they are 40% utilized or 80% utilized. So having half as many drives that are 80% utilized requires only half the power and cooling cost. Even service and support costs are lower when there are only half as many disk drives under warranty and support.

Lower Administration Costs

As storage systems have gotten larger and more complex, administration costs are on the rise. Highly trained storage administrators are often required and it can take longer to find performance bottlenecks in today's large systems.

The Pillar Axiom 600 offers automated tuning and simplified management to help reduce administration costs. Application profiles allow a new LUN to be provisioned in less than one minute. Simple questions and choices are presented to administrators, masking a lot of the underlying technical complexity of the system and focusing instead on the business needs of the application. Pre-built application profiles also provide a starting point and can be customized as needed.

When using an application profile to provision a LUN, the administrator can complete the task in **just three clicks**: selecting the application profile, choosing the LUN size, and then naming the LUN. Even without an application profile, the provisioning task can be completed in **just six clicks**. Faster provisioning also means greater business agility because new applications or databases can be up and running with a fully optimized storage environment in very little time.

Pillar Axiom 600 systems are often used to consolidate storage, achieving additional management savings because the consolidated environment requires less management effort than multiple separate storage systems. Multiple Pillar Axiom systems can also be managed from a single pane of glass with Pillar Axiom MaxMan. Multiple local or geographically dispersed Pillar Axiom 600 systems can be managed using Pillar Axiom MaxMan.

Perhaps most importantly, Pillar Axiom 600 systems typically have half as many drives to manage because the systems deliver twice the utilization of traditional storage systems, such as the EMC VNX family.

Storage Consolidation Cost Savings

The Pillar Axiom supports prioritized storage I/O, which enables organizations to consolidate more applications on the same storage system and still meet SLAs. A description of how Pillar Axiom Quality of Service (QoS) helps meet SLAs can be found in the later section titled "Consistent SLA Performance in Consolidated Environments."

Storage consolidation enables organizations to meet their storage needs without having to purchase and manage separate storage systems for each application. Rather than deploying multiple separate storage systems that are each partially utilized, organizations can deploy a single consolidated storage platform and achieve big savings. A consolidated storage platform not only costs less to purchase but also requires less floor space, power, and cooling.

Software license costs will also be significantly lower when consolidating storage on the Pillar Axiom 600. The purchase of any Pillar Axiom 600 system includes all of the software licenses needed to operate the system regardless of how large the system grows. No additional software licenses are required even when upgrading to the maximum number of Slammers and Bricks. For EMC VNX systems, some software licenses are based on total storage capacity. Thus, when certain capacity thresholds are hit, EMC VNX customers may need to purchase additional software licenses. When consolidating multiple EMC VNX storage systems onto a Pillar Axiom 600, the system can then grow as needed without further license upgrades.

Price Comparison Scenarios

To demonstrate the cost savings that can be achieved with the Pillar Axiom 600, two scenarios are provided below. Each scenario includes a description of requirements based on a characterized application environment. The Pillar Axiom 600 and EMC VNX configurations that would be able to meet the needs of this environment are then defined along with list pricing to allow for competitive comparison.

Up to 63% Savings for OLTP Oracle Database Scenario

An organization running Oracle E-Business Suite has 40 TB of data in their Oracle Database and is anticipating rapid growth in the future so they want a scalable storage platform. The 40 TB database consists of five years of sales and customer data. The most recent 12 months of data requires immediate read and write access, while the previous four years of data must simply be maintained for compliance purposes and occasional searches. This older data, which represents 75% of the total data, does not require write access because the data is historical, and thus no new transactions will update the data.

This scenario provides a good example of when Oracle Partitioning and Hybrid Columnar Compression functionality in the Oracle Database can be very helpful. Oracle Partitioning allows a table, index, or index-organized table to be subdivided into smaller pieces where each partition has its own name and may optionally have its own storage characteristics.

Thus, the data for years two through five can be managed in a separate partition and can reside in a separate LUN on the Pillar Axiom 600 system, where it can take full advantage of Hybrid Columnar Compression. The most recent 12 months of data cannot take advantage of Hybrid Columnar Compression because there is still write activity against this data and read activity is still quite heavy.

30x Compression Using Hybrid Columnar Compression

For archive data such as this, compression ratios of 30x to 50x are possible with Hybrid Columnar Compression. This feature is only available on Oracle storage platforms, so the EMC VNX platform cannot achieve anywhere near this level of efficiency. Other compression techniques and deduplication approaches typically don't provide a significant efficiency gain for database tables. For the purposes of this scenario, a conservative 30x compression ratio is assumed for the Pillar Axiom 600 partition housing the archive data.

Figure 3 shows how dramatically this impacts total storage requirements — more than a 70% savings in total storage for all five years of data. The data for years two through five requires 30 TB of storage uncompressed, but can be compressed to just 1 TB of storage with Hybrid Columnar Compression. Since the most recent 12 months of data requires 10 TB of storage, the total storage requirement is reduced from 40 TB to 11 TB.

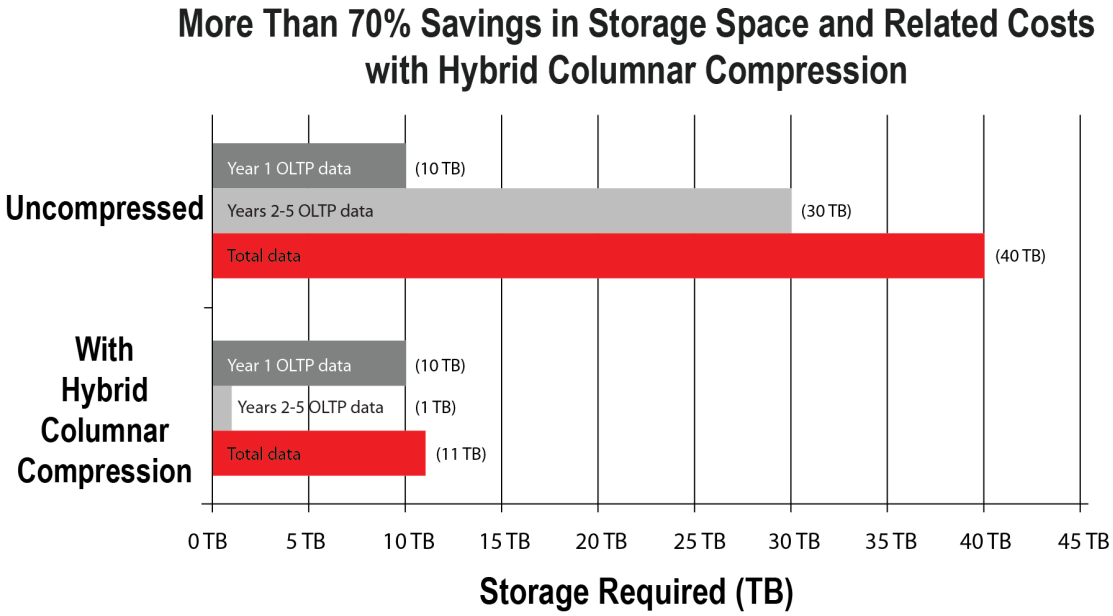

Figure 3. Using Hybrid Columnar Compression on the data set for years two through five enables more than 70% savings in total storage space and related costs.

Table 2 shows the Oracle and EMC configurations that can meet the needs of this scenario. The systems were configured with a little more than 70 TB of raw capacity to provide at least 40 TB of useable capacity after allowance for RAID parity data. For simplicity sake, the systems were configured to deploy the entire data set (years one through five) on 7200 RPM drives. The EMC configuration uses 2 TB 7200 RPM SAS drive shelves while the Pillar Axiom 600 systems use 2 TB 7200 RPM SATA Bricks. Both types of drives should offer similar performance. In reality, the first year of data would likely be deployed on faster Fibre Channel drives or perhaps even SSDs, but since that cost would be similar in both systems, a simplified configuration consisting entirely of 7200 RPM drives has been used for this scenario.

The EMC VNX 5700 model was chosen because it is the smallest of the VNX family that comes closest to the Pillar Axiom 600 in terms of scalability for capacity and performance. The EMC VNX 5700 can support up to 500 drives while the Pillar Axiom 600 can support up to 832 drives. The recommended Pillar Axiom 600 configuration has one Slammer, which contains two controllers. Two EMC VNX 5700 systems are configured for this scenario so that the EMC offering will also have two storage controllers. It was not considered necessary to compare the Pillar Axiom 600 to EMC's largest system, the VNX 7500, which supports up to 1,000 drives because the cost differential is even greater.

The table shows two different Pillar Axiom 600 configurations, one that assumes the use of Hybrid Columnar Compression and one that does not. The Pillar Axiom 600 configuration that uses no compression offers an impressive 43.7% savings over EMC, and the configuration with Hybrid Columnar Compression offers an even more dramatic 63.1% cost savings over EMC.

TABLE 2. LIST PRICE COMPARISON FOR OLTP ORACLE DATABASE SCENARIO

	HARDWARE AND SOFTWARE CONFIGURATION	HARDWARE AND SOFTWARE LIST PRICE*	PERCENT SAVINGS OVER EMC*
EMC VNX	• 2x EMC VNX 5700 controllers • 72TB raw storage (36x 2TB 7200 RPM SAS disk drives) • VNX Operating Environment • EMC Unisphere Unified Storage Management Solution	$233,318.	N/A
PILLAR AXIOM 600 (UNCOMPRESSED DATA)	• Pillar Axiom Fibre Channel SAN Slammer • 72TB raw storage (3x Pillar Axiom SATA Brick with thirteen 2TB 7200 RPM SATA drives per Brick, including one spare drive per Brick) • Pillar Axiom Pilot with 4GB memory • All software included	$131,393.	43.7%
PILLAR AXIOM 600 (HYBRID COLUMNAR COMPRESSION)	• Pillar Axiom Fibre Channel SAN Slammer • 24TB raw storage (1x Pillar Axiom SATA Brick with thirteen 2TB 7200 RPM SATA drives per Brick, including one spare drive per Brick) • Pillar Axiom Pilot with 4GB memory • All software included	$86,159.	63.1%

* EMC pricing obtained from publicly available sources. Pillar Axiom 600 pricing is based on Oracle list prices as of March 19, 2012 The prices shown do not include the cost of support.

Up to 41% Savings When Consolidating Storage for Multiple Applications

An organization would like to consolidate the storage for three application data sets that are currently stored on three separate EMC VNX systems. The applications include:

- Microsoft Exchange email application requiring 10 TB of storage
- 20 TB data warehouse and reporting tool used to monitor sales and inventory
- Product Data Management (PDM) application that contains important engineering designs and 3D models and requires 50 TB of storage

These applications would typically be deployed with dedicated storage rather than consolidating the storage for all three applications on a single EMC VNX system. Microsoft best practices state that

dedicated storage should be used for Microsoft Exchange.[5] If the Microsoft Exchange storage were consolidated with the data warehouse, a single table scan on the data warehouse could cause a service interruption for Microsoft Exchange users. The table scan would make a long sequential run through the data and would disrupt I/O to the Microsoft Exchange application.

The PDM application is also typically deployed with dedicated storage because it serves a large engineering community (potentially a few thousand engineers) and the engineers need immediate response when they request access to a drawing or their productivity will suffer. Engineering design files can be quite large, often hundreds of MBytes or even tens of GBytes of data. When several engineers request a file at the same time, the storage system must be able to respond quickly and not get held up waiting for I/O for other applications.

For these reasons, these applications would be configured on three different EMC VNX systems as follows:

- Microsoft Exchange — Microsoft Exchange would run on a pair of EMC VNX 5500 controllers with a total of 16 TB of raw storage to meet the requirement for 10 TB of usable storage. While the smaller EMC VNX 5300 system could also be configured with this much storage capacity, Microsoft Exchange requires the I/O bandwidth of the EMC VNX 5500 model.
- Data Warehouse — The data warehouse would also be deployed on a pair of EMC VNX 5500 controllers, which would be configured with a total of 40 TB of raw storage to meet the requirement for 20 TB of usable storage with a little room to grow.
- Product Data Management — The PDM application would be deployed on a pair of EMC VNX 5700 controllers with a total of 80 TB of raw storage to meet the need for 50 TB of usable storage.

As discussed earlier, the Pillar Axiom 600 system offers a unique QoS capability that enables these applications to be consolidated onto a single storage system while maintaining appropriate I/O bandwidth for each application. Storage administrators have the controls they need on the Pillar Axiom 600 to meet the service level requirements for each application without the extra cost of dedicated storage.

Table 3 shows the configurations and list prices for the EMC and Oracle configurations that would house the three different data sets. The Pillar Axiom 600 system that offers consolidated storage is configured with two Slammers and thus four storage controllers. It has 168 TB of disk, slightly more than the combined total of 136 TB for the three EMC VNX configurations. The cost savings from consolidating the three different EMC VNX environments onto this single Pillar Axiom 600 system is $186,860 or 41.3%.

[5] http://technet.microsoft.com/en-us/library/ee832792.aspx.

TABLE 3. LIST PRICE COMPARISON FOR CONSOLIDATION SCENARIO

	HARDWARE AND SOFTWARE CONFIGURATION	HARDWARE AND SOFTWARE LIST PRICE*
EMC VNX SYSTEM FOR MICROSOFT EXCHANGE APPLICATION	• 2x EMC VNX 5500 controllers • 16TB raw storage (8x 2TB 7200 RPM SAS disk drives) • VNX Operating Environment • EMC Unisphere Unified Storage Management Solution	$87,811.
EMC VNX SYSTEM FOR DATA WAREHOUSE APPLICATION	• 2x EMC VNX 5500 controllers • 40TB raw storage (20x 2TB 7200 RPM SAS disk drives) • VNX Operating Environment • EMC Unisphere Unified Storage Management Solution	$ 120,547.
EMC VNX SYSTEM FOR PRODUCT DATA MANAGEMENT APPLICATION	• 2x EMC VNX 5700 controllers • 80TB raw storage (40x 2TB 7200 RPM SAS disk drives) • VNX Operating Environment • EMC Unisphere Unified Storage Management Solution	$244,230.
TOTAL FOR ALL THREE EMC VNX ENVIRONMENTS	• Two pairs of EMC VNX 5500 systems and one pair of EMC VNX 5700 systems with a combined total of 136TB	$ 452,588.
PILLAR AXIOM 600 FOR CONSOLIDATION OF APPLICATIONS A, B, AND C	• 2 x Pillar Axiom Fibre Channel SAN Slammer • 168TB raw storage (7x Pillar Axiom SATA Brick with thirteen 2TB 7200 RPM SATA drives per Brick, including one spare drive per Brick) • Pillar Axiom Pilot with 4GB memory • All software included	$265,728.
Savings with a single Pillar Axiom 600 compared to three EMC VNX systems		**41.3%**

* EMC pricing obtained from publicly available sources. Pillar Axiom 600 pricing is based on Oracle list prices as of March 19, 2012. The prices shown do not include the cost of support.

Consistent SLA Performance in Consolidated Environments

Applications are the lifeblood of the business, and if application performance is jeopardized by storage I/O performance, there is an immediate business impact. The company may miss business opportunities or lose customers because employees cannot get access to information in a timely manner. For this reason, organizations have traditionally purchased separate storage systems for each

of their business applications, thus minimizing the risk that other applications could interfere with storage performance.

Although isolated storage islands result in poor utilization, they have been commonly deployed because of the potential performance risk when consolidating storage. Today's IT budget pressures and continued growth in storage requirements are now causing IT managers to reconsider this old approach and evaluate storage consolidation as a means to save on storage costs. However, storage consolidation is only practical if there is a means to ensure that SLA targets can be met in the consolidated environment.

Traditional storage arrays such as the EMC VNX product family use a first-in-first-out queue for all I/O requests regardless of their business value. As shown in the left-hand side of Figure 4, this can cause business-critical I/O requests to get stuck waiting for those from other less critical applications. By contrast, the Pillar Axiom 600 system uses priority queues that enable high priority I/O requests to be executed first. As shown on the right-hand side of the figure, this means that lower priority applications may sometimes have to wait while some I/O requests from the higher priority queues are serviced. However, by prioritizing I/O requests according to their business value, the Pillar Axiom 600 enables organizations to achieve deterministic performance for business-critical applications in a consolidated environment while minimizing I/O contention. This, in turn, enables increased utilization because other applications can share the same storage system without increasing business risk.

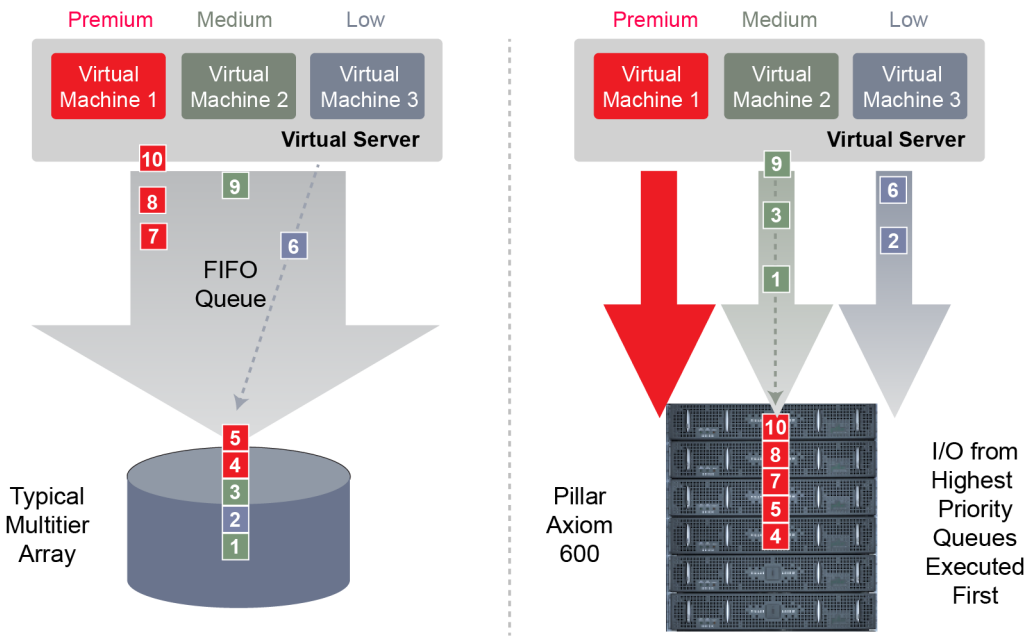

Figure 4. Pillar Axiom 600 systems prioritize I/O requests so that business-critical applications don't have to wait.

Pillar Axiom 600 Quality of Service (QoS) Technology

Many storage vendors use disk tiering alone to address the need for different levels of service in a multiple application workload. This usually meets with limited success because disk tiering does not control the I/O queue to the drives. Even though high priority data may be stored on faster disk drives or striped to the outer edges of disk platters, read or write requests for high priority data can still get queued up behind less critical requests at the storage controller level. For optimal performance, the disk head movement needs to be minimized by understanding the relative priorities of I/O requests to the drive.

The Pillar Axiom 600 was designed from the ground up with Quality of Service (QoS) in mind. Unlike systems that bolt on storage classification as an afterthought, the Pillar Axiom 600 has native QoS technology. The result is a carefully balanced, automatically tunable storage system that optimizes performance of five classes of data within the same array. The Pillar Axiom 600 implements QoS throughout its modular components. All layers work in concert to deliver predictable performance to every application, contributing to deterministic SLA compliance. This is patented technology that no other storage vendor can deliver today.

How Pillar Axiom QoS Technology Works

The QoS functionality in the Pillar Axiom 600 provides controls that determine how the file system or LUN is laid out physically on the disk drives within the Bricks and how application I/O requests are managed by the Slammers. Each file system or LUN is assigned a priority of premium, high, medium, low, or archive, depending on the business priority of the data to be stored on that LUN. I/O requests to a prioritized LUN are handled according to the priority level chosen when the LUN was provisioned. For mixed QoS workloads, the Slammer will maintain a relative QoS queuing of 38% for premium, 25% for high, 19% for medium, 12% for low, and 6% for archive.

For example, an I/O request that comes in on a premium priority LUN will get approximately six times (38% vs. 6%) the compute cycles on the Slammer control unit compared to I/O requests that come in on an archive priority LUN. Additionally, higher priority LUNs will be striped across a greater number of drives than lower priority LUNs. Higher priority LUNs are also placed closer to the outer edge of the disk platters. With a heavily loaded I/O environment, QoS will ensure that lower priority I/O operations are still serviced, preventing I/O starvation.

Figure 5 shows the various QoS functions within the Pillar Axiom 600. The vertical columns in red represent the five queues to be managed and show the default weighting for each queue. When an I/O request arrives, the Queue Assignment function assigns it a priority to match the target LUN and places the request in the appropriate queue. The QoS Prioritized I/O Scheduler will then schedule high priority requests before low priority requests. It also balances queue weightings so as to avoid I/O starvation for lower priority queues.

The QoS I/O Cache Management function manages how the cache is used for different types of I/O. For example, all write requests get copied to the write cache of the control unit and to a copy-of-write cache on the peer control unit within the same Slammer. Additionally, QoS I/O Cache Management utilizes an Aggressive Read Ahead feature for read-intensive sequential I/Os, thus filling up cache with

data that is expected to be used. If the I/O requests are random in nature and write-intensive, cached data will be flushed as soon as possible to make room for other I/O requests.

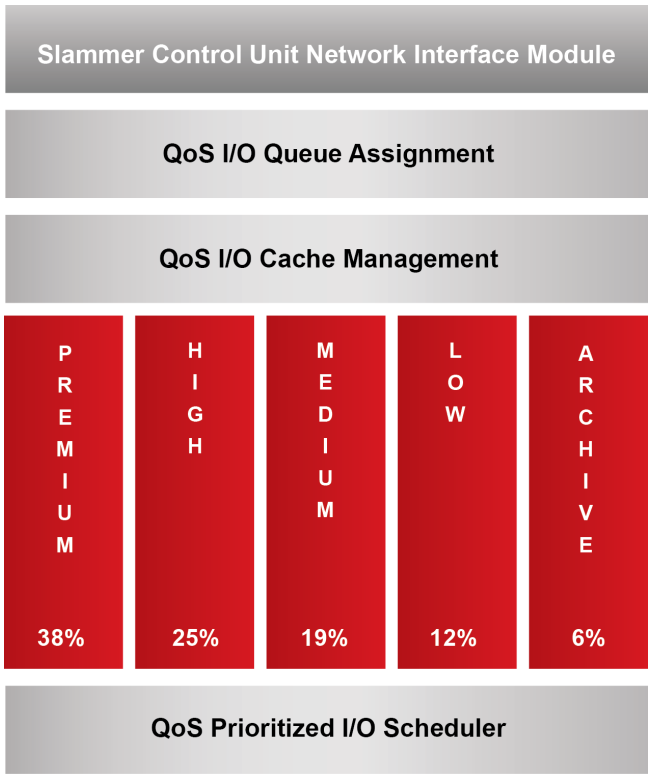

Figure 5. Data is prioritized according to its business value, with the highest priority data being given the largest percentage of queue service time.

Establishing Application Profiles for QoS

The Pillar Axiom 600 comes with approximately 30 pre-defined application profiles (such as for Oracle Database) that can be used immediately. These profiles can also be tweaked by administrators as they see fit. The administrator can also establish custom application profiles for each LUN or file system by answering six simple questions in a configuration wizard with popup menus. The information that QoS needs to provision a LUN is as follows:

- LUN name
- LUN size in GBytes
- Drive type (SSD, SATA, or FC)
- Access bias
- I/O bias
- Thin provisioning

Figure 6 illustrates the type of information that can be defined for each LUN. Administrators can choose the appropriate I/O priority and storage class (SSD, SATA or FC) as well as the typical I/O profile or storage characteristics for applications that use the LUN.

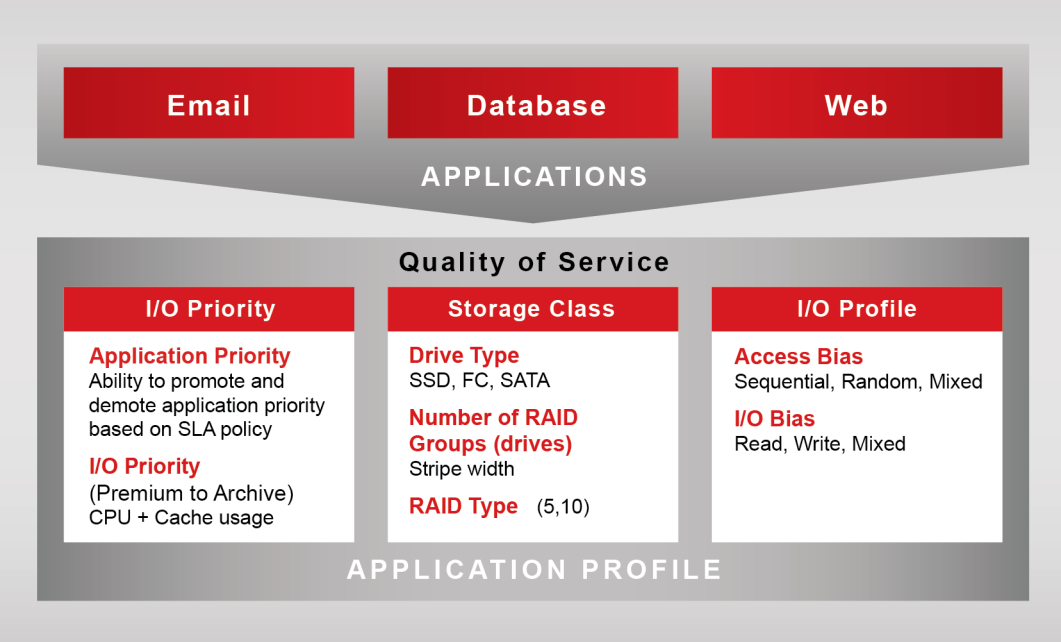

Figure 6. Application profiles enable administrators to quickly and easily define performance.

Dynamic Resource Allocation

The Pillar Axiom 600 dynamically adapts to changing application requirements to help achieve consistent performance in a consolidated environment. Storage resources are initially allocated according to business value, and then unused resources can be dynamically reallocated to meet immediate demands. This is achieved while ensuring high service levels by prioritizing I/O requests according to their business value.

Secure Isolation with Storage Domains

Another requirement for storage consolidation is that data from different applications be well segregated so that one application cannot overwrite another's data if something goes wrong. In the event of a security breach, the security threat must be quarantined so that it cannot easily expand from one application's data to another.

In the Pillar Axiom 600, each storage domain consists of specific Bricks or disk trays. This means that there are separate RAID controllers and physical separation of data for each storage domain. Each storage domain can also have different QoS levels within the domain so that multiple applications can share the same domain and have different priorities if desired.

More Consistent Performance Through Faster Disk Drive Rebuild Times

Disk drives are mechanical devices that are known to fail, and today's storage systems automatically rebuild a failed drive onto a spare drive under the direction of the RAID controller. From a performance perspective, the duration of time that it takes to rebuild a drive affects overall system throughput and thus can impact SLAs.

In traditional systems, such as an EMC VNX, a centralized storage controller is used to rebuild failed drives, so this takes away resources from the rest of the system. Therefore, performance for the entire array degrades significantly during drive rebuilds, which can take up to twice as long on a busy LUN compared with an idle drive.

A 600 GB Fibre Channel drive can be rebuilt in approximately 4 hours on a Pillar Axiom 600 system whereas a typical drive rebuild on other platforms can be significantly longer (2-4x longer on a busy system). Pillar Axiom 600 systems have two RAID controllers and a spare drive in each Brick so that when a drive is being rebuilt, it does not affect any part of the storage system outside of that particular disk tray. The performance impact is localized, and the rebuild is much faster because data does not need to travel outside of the disk tray.

Performance Synergy with Oracle Database

The performance benefits of Pillar Axiom 600 systems are even greater when used with Oracle software environments. The storage is engineered for Oracle software and includes a number of unique integration points that are available only when deployed with Oracle software. For example, Oracle Database includes a high-efficiency data compression feature called Hybrid Columnar Compression that works only with storage developed by Oracle. The ability to leverage Hybrid Columnar Compression, integration with Oracle VM, and other unique benefits of the Oracle software environment further speed the time to value for the Pillar Axiom 600, as well as lowering business risk for companies that deploy it with Oracle Applications.

Data warehouses are the most common use case for Hybrid Columnar Compression. OLTP and mixed-use environments where database archives are required for business or regulatory reasons will want to take advantage of Oracle's Advanced Compression Option (ACO). Hybrid Columnar Compression typically achieves between 10x and 50x compression ratio. The compressed data set can be sent to the storage system from the Oracle Database, so there can also be a 10x to 50x reduction in I/O traffic, which translates to much faster storage performance as well.

Since EMC VNX systems do not support Hybrid Columnar Compression, customers who are currently running Oracle Database environments with EMC VNX or CLARiiON systems can see up to 3x to 5x reduction in their storage footprint by migrating to a Pillar Axiom 600 system. Or, stated differently, EMC VNX users will require 3-5x more disk capacity to house an uncompressed database compared with users running a Pillar Axiom 600 using Hybrid Columnar Compression.

Conclusion

Today's demanding data center environments require scalable SAN solutions that can expand without disruption and maintain SLAs for critical business applications. Yet growing storage capacity requirements are also making cost an increasingly important criterion.

Oracle's Pillar Axiom 600 delivers better scalability, better service levels, and better TCO than EMC's VNX family of products. By choosing the Pillar Axiom 600 over EMC VMX, organizations can achieve the following benefits:

- Eliminate costly, inefficient data silos.
- Meet or exceed SLA targets while consolidating storage on a single shared platform.
- Drive down costs with up to 2x greater storage efficiency and simplified provisioning.
- Protect storage investments and increase business agility with storage that scales across the entire EMC VNX product family.
- Achieve up to 3x to 5x reduction in storage footprint for Oracle Database environments.
- Optimize performance and integration with Oracle software environments.

For More Information

For more information on Oracle's Pillar Axiom 600, call +1.800.ORACLE1 to speak to an Oracle representative, or visit the Web links listed in Table 4.

TABLE 4. WEB RESOURCES FOR FURTHER INFORMATION

WEB RESOURCE DESCRIPTION	WEB RESOURCE URL
Oracle storage offerings	http://oracle.com/storage/
Oracle's Pillar Axiom storage system	http://oracle.com/us/products/servers-storage/storage/san/overview/

Oracle's Pillar Axiom Storage — A Better Choice Than EMC
May 2012

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

Hardware and Software, Engineered to Work Together