



Lowering the TCO of Storage Growth — Dell Deploys SSDs in Virtualized iSCSI Arrays

Analyst: David Reine

Management Summary

For the average consumer, shopping for a new automobile has become an adventure, trying to determine the most important variables, and their value to you, in trying to determine the total cost of ownership against the family budget. If we exclude some of the more superficial categories, such as paint color or the inclusion of a GPS, we have to look at what we consider the mission-critical features: how much performance we require, what capacity we need, and which safety features are required and which are optional. In terms of performance, how much horsepower do we need to drive around town as opposed to how much is required to get up to 60 MPH to merge into freeway traffic? Do we purchase a vehicle that can seat four? What happens when we need to drive the carpool that includes more than four? In terms of safety, do we settle for the front airbags or do we insist upon side airbags also? All of these options include a cost factor. Going for the six-cylinder engine may provide the pickup you need for the highway but it may be over-provisioning for 90% of normal driving requirements. You also need to factor in the on-going gas economy being wasted 90% of the time in order to have the extra “giddy-up”. The same is true if you need to configure for towing: extra hp, over-sized radiator, and heavy-duty battery for the rare occasion that you need to tow your yacht to the lake. Then there are seating and storage requirements: do you configure for the occasional peak load or for the normal workload. In all of these cases, you can configure for the peak requirements, but you are wasting money and resources in terms of acquisition cost, operational cost, and fuel economy. Sometimes you have no choice; you have to go with the over-provisioned model, unless you can afford a second vehicle.

These same issues confront the CIO of every enterprise trying to configure the data center for both mission-critical and business-critical high-performance applications, as well as high-capacity archiving functions. How can you configure storage requirements for high-performance requirements without being forced to over-provision the data center with higher cost storage for applications that do not require it? There are some applications that require literally hundreds of thousands of I/Os per second (IOPS). In order to meet this requirement, the SAN is configured with hundreds, or thousands, of high-cost Fibre Channel (FC) devices, each configured for short-stroking, i.e., writing to only the fewest outside cylinders to order to improve performance. This wastes not only 90% of your storage capacity, but also 90% of the acquisition budget, power requirements, and floor space.

In order to improve performance and increase storage and power utilization, and reduce floor space consumption, many data centers are looking to solid state disks (SSDs). One company that has recognized this growing problem is Dell. Dell has now configured its *EqualLogic PS6000* storage array with an optional set of SSD devices, creating the *PS6000S*. To learn more about the *PS6000S*, please read on.

IN THIS ISSUE

- **High-Performance Data Center Needs..** 2
- **Solid State Storage Options** 2
- **Dell’s EqualLogic PS6000S** 4
- **Conclusion** 5

High-Performance Data Center Needs

There has always been a requirement for a wide variety of storage devices in the enterprise data center. Mission-critical applications have always demanded high-performance, high-availability devices to ensure rapid response with high availability. These Tier-1 devices are more expensive than the high-capacity Tier-2 drives needed to satisfy the scalability demands of a rapidly growing enterprise. The typical data center also maintains a Tier-3 storage requirement to satisfy its archive and long-term storage needs. This Tier-3 requirement has usually been satisfied by tape, due to its low cost structure and low energy demands.

Tier-1 requirements can be further sub-divided in terms of performance and capacity, with 15K RPM Fibre Channel (FC) hard disk devices (HDDs) carrying a premium over 10K RPM drives. HDDs, however, have failed to maintain parity with the increased capabilities of server I/O controllers, especially when considering the rotational capability, which has remained static since going to 15,000 RPM last century! Tier 2 media also has multiple performance points, with 7200RPM and 5400RPM SATA drives. These low-cost drives often can provide more value to the data center in terms of business-critical application needs for high-growth data such as email. In terms of capacity, SATA drives have maintained a high rate of improvement with 1TB drives currently ruling in the high capacity space, with 2TB drives on the near horizon.

We have also seen a significant increase in both the capabilities of, and the demand for, iSCSI based disk arrays, in lieu of the FC arrays that are prevalent in the higher end enterprise data center. With an increase in the throughput of Ethernet-based controllers from 1GbE to 10GbE, iSCSI-based arrays have made substantial gains, especially in the SMB data center and in departmental arrays in the larger enterprise. It is important to consider the functionality that is available with these iSCSI arrays; features such as thin provisioning and snapshots are critical to the success of any iSCSI SAN.

Recently, however, high-performance applications with extremely high IOPS requirements have been placing a tremendous strain on the performance limits of the HDD array, as well as the data center storage budget. In fact, some applications with the need for speed are demanding hundreds of thousands of IOPS to satisfy

“C” level management requirements, far beyond the capability of “normal” spinning media. In order to meet this demand, the data center has been deploying mission-critical Tier-1 resources with a “short-stroke” architecture, only accessing the outermost cylinders of the disks in order to accelerate the response time. In order to achieve the required level of IOPS, many enterprises are deploying thousands of expensive, short-stroked FC devices, filling banks of racks with energy-consuming spinning media, at utilization rates of less than 10%. This is a waste of resources and budget dollars, both CapEx and OpEx.

In order to reduce the total cost of ownership (TCO) of the IT infrastructure, many data centers are implementing a new Tier-0 storage class for applications with the highest IOPS performance requirements. These applications include ERP solutions such as SAP and Oracle, high-transaction OLTP servers, email servers, and virtualized appliances that reduce the predictability of the next I/O. This premium Tier-0 level is being deployed with solid state disks (SSDs) configured with Flash memory¹. However, in order to utilize the full capability of the SSD, the I/O controller in the server must have sufficient throughput to be able to scale commensurate with the storage.

Solid State Disk Options

Solid-state flash drives can accelerate operational performance by improving application response times. They can also increase storage efficiency by doubling server throughput while reducing space and energy requirements. Surprisingly, SSDs can also *reduce* the TCO of the IT infrastructure by eliminating the necessity to add additional servers and HDDs to boost performance. SSDs will also increase reliability because they contain no moving parts.

Where do Flash SSDs belong? In the server, as direct attached storage (DAS), or in the storage array (SAN)? The answer is both, as the amount of SSD storage required will determine the deployment. While the cost per GB for SSDs is higher than HDDs, the \$/IOPS is much lower. Flash, however, is much faster than HDDs but slower than the DRAM often found within the server. While DRAM speed is measured in nano-

¹ See the issue of *Clipper Notes* dated February 10, 2009, entitled *A New Tier of Storage Appears – Faster, Solid-State Drives State Their Case*, and available at <http://www.clipper.com/research/TCG2009006.pdf>.

seconds, SSD speed is only measured in micro-seconds. SSDs, however, are significantly faster than HDDs, with speed measured in milliseconds. DRAM, unfortunately, is still very expensive, has limited capacity, and is volatile². This makes SSDs more attractive to the data center for applications with very high IOPS requirements while HDDs retain the interest of the enterprise in terms of high capacity.

An SSD can replace a drawer of HDDs; a drawer of SSDs can replace a rack of HDDs. The relatively higher cost per GB of SSDs, however, often leads to the deployment of a hybrid environment with SSDs being used for high performance, and SATA drives deployed for high capacity. Some mission-critical applications can still be satisfied with a modest amount of traditional FC drives, however, leading to the deployment of a multi-tiered architecture of SSD and FC/SATA HDDs.

SSDs bring a new level of performance to the storage array, delivering over 100 times more IOPS to operational performance than HDDs. With no spinning media, SSDs reduce latency, lower energy consumption, and eliminate head movement, improving reliability and performance, while lowering TCO. In addition, with a familiar format, the IT staff can rapidly deploy SSDs into existing environments with the same form factor as traditional HDDs. With data stored on integrated circuits that can withstand shock and vibration, SSDs last longer, reducing diagnostic time and minimizing the number of devices that need to be replaced.

There are two types of flash memory – *NOR* and *NAND* – that form the circuitry of flash devices. With long erase and write times and large voltage requirements, NOR flash is well-suited for random access to data that needs to be updated infrequently. NAND flash, on the other hand, provides block access to data and supports greater storage densities, provides greater endurance, and costs less per unit of storage, making NAND flash the more commonly implemented variety.

There are two types of NAND flash technology: *single-level cell (SLC)* and *multi-level cell (MLC)*. SLC devices store a single binary value in each memory cell, while MLC supports

up to four or eight values per cell. **Because of its longer lifespan and better performance SLC NAND flash devices are preferred for enterprise applications.** However, there are some very attractive use cases for MLC flash when coupled with new designs in the Flash Memory Controller firmware for enterprise applications. Many of these are read-intensive applications, such as caching and predominately read databases, like those found in Telco. Capacity, as opposed to performance, is the main reason to consider MLC.

As a result of using a purely electronic solution, as opposed to mechanical devices, the data center can gain additional advantages by using SSD devices. For one thing, the data center will experience lower energy consumption, improving the TCO, improving the “green” image of the enterprise, and improving the bottom line. Secondly, with fewer drives required, the data center will reduce floor space consumption. Combined with the reduction in energy consumption, perhaps this may extend the life of the data center for a number of years, eliminating the necessity of a large capital CapEx for a new facility.

There are cautions, however, such as the durability and performance of write operations. There is a write performance penalty when the target is an SSD. The SSD cells must be erased before writing to them, creating a latency delay, although not nearly the delay encountered with spinning media. In addition, SSD cells degrade over time. In general, an SLC SSD will hold up for up to 100,000 writes, with MLC SSDs experiencing a significantly lower rate, up to 10,000 writes. This requires an intelligent controller to manage the provisioning of the SSDs to ensure a long life.

In order to take advantage of this latest innovation in technology, the data center needs to find an SSD solution that is not only performant, but also easy to use and scalable to help solve the I.T. issues of today and tomorrow. One company that is addressing these exact issues is Dell, with the addition of two new models to their *Equal Logic PS6000* family³, the *PS6000S* and the *PS6010S*.

² Volatile means that data is lost when power is removed, as in a power failure. Non-volatile memory doesn't require constant power to preserve its contents.

³ See [The Clipper Group Navigator](http://www.clipper.com/research/TCG2007087.pdf) dated September 10, 2007, entitled *Dell Expands Storage Tiers for the SMB – Introduces Low-Cost, Extensible Storage*, which is available at <http://www.clipper.com/research/TCG2007087.pdf>.

Dell's EqualLogic PS60x0S

The Dell PS60x0S is an intelligent, virtualized iSCSI^{4,5} SAN platform that is dedicated to solid-state disks and combines intelligence and automation, along with fault tolerance, to provide any enterprise with the performance, simplified administration, high reliability, and seamless scalability that it requires. Dell EqualLogic SSD arrays are offered to the SMB community at a significantly more reasonable price, compared to competitive SSD offerings, to enable more enterprises with the opportunity to afford the benefits of solid-state technology. Instead of using the highest performance SSD devices, carrying the highest cost, Dell has configured the PS60x0S with more moderately priced 50GB and 100GB high performance SSD drives. This makes the benefits of SSD available to more data centers.

However, SSDs remain more expensive than HDDs. As such, the IT staff is faced with the challenge to ensure that only certain high-priority or niche applications run on SSD devices, as it is important to provide these applications with the highest performance. Some SSD offerings are included as a small RAID set, a second stage cache, or as another tier of storage virtualization behind a controller head. These solutions accentuate the problem of controller resource contention. SSD drives in a frame architecture have to share the same controllers with the rest of the storage in the array, a problem made worse by the enhanced performance capabilities of SSDs. This could force the IT staff to still run separate SANs, defeating some of the benefits of networked storage. This problem does not exist with the Dell EqualLogic PS60x0S family.

The PS6000S has been designed as a dedicated pool or tier within the SAN, not sharing controller resources with lower priority workloads, enabling the data center to unlock the potential for higher performance from the SAN, avoiding the linked problems of controller or network port resource contention, and providing an outstanding random IOPS capability.

With significantly higher IOPS than a 15K RPM SAS array, the PS6000S is ideal for intense,

moderate capacity workloads, such as OLTP or consolidated, virtualized environments. As with the other members of Dell's EqualLogic family, the PS6000S provides simplified network storage, enabling the IT staff to work smarter, including all of the software and applications expected from a Tier-1, enterprise-level SAN – and that functionality is included at no extra charge, simplifying the acquisition process.

Each PS60x0S can be deployed with either 8 or 16 hot-pluggable SSD devices per array, for a capacity of 800GB, or 1.6TB per drawer⁶. Each array is configured with dual controllers with 4 GbE network interfaces per controller, and a total of 4GB of battery-backed cache memory to provide up to 72 hours of data protection. Each PS60x0S also has hot-swappable power supplies/cooling fans for additional reliability, making the PS60x0S virtually bulletproof.

It is highly scalable, on-demand, with the capability to consolidate up to 12 PS60x0S arrays into a single group, or pool, of high-performance data storage, up to 19.2TBs. This enables the data center to minimize initial CapEx, operating on a seamless, pay-as-you-grow basis. Moreover, the data center can create a consolidated, multi-tiered environment, combining PS60x0S arrays with other members of the PS6000 family, utilizing SSD, SAS, and SATA drives in a single pool. Any data can be migrated between PS6000 arrays in order to balance workloads, maintain high availability, or simply move data from SSD to SAS or SATA, and back again, as the value of the information changes. The PS60x0S supports up to 1024 volumes with 512 snapshots per volume, up to a total of 10,000 snapshots.

Any volume in that pool can be dynamically moved, transparently, between storage tiers to take advantage of the most price/performance tier of storage for that specific data. Because of the innovative EqualLogic iSCSI technology, as the data center adds storage, they also add controllers, to marry additional performance to storage growth, providing the enterprise with full system scalability. The PS6000S supports a wide variety of operating environments, including multiple Windows iterations, VMware and Citrix virtualization engines, Linux, Solaris, and a variety of UNIX implementations.

⁴ See the issue of *Captain's Log* dated November 15, 2007, entitled *iSCSI Train Gains Steam – Dell Drives Ahead on the Fast Track*, and available at <http://www.clipper.com/research/TCG2007100.pdf>.

⁵ See the issue of *Clipper Notes* dated March 5, 2007, entitled *iSCSI SANs – Panacea or Placebo?*, and available at <http://www.clipper.com/research/TCG2007037.pdf>.

⁶ Each 64GB SSD is over-provisioned down to 50GB in order to minimize the write penalty of NAND SSDs.

Exhibit 1 — PS6000S Functionality

- **Data Management** – includes SAN HQ to provide the data center with a unique multi-group performance and event monitoring tool, along with rapid provisioning and a remote setup wizard;
- **Data Reliability Features** – including snapshots, replication, volume cloning, failover, and multi-path I/O;
- **Storage Virtualization** – with storage pools, automatic load balancing, volume migration, storage tiering, and thin provisioning; and
- **Server Management & Protection Integration** – to automate Snapshot Manager, MPIO, Smart Copy, and VMware Site Recovery Manager.

Source: Dell

The PS60x0S provides full RAID⁷ support and includes all of the basic and advanced management features, providing comprehensive data management services to the IT staff, once available only to Tier-1 data centers. (See Exhibit 1, above.) One highlight is that the PS60x0S provides a simplified virtual data protection with *Auto-Snapshot Manager/VMware Edition* (ASM/VE). ASM/VE provides a high-performance, space-efficient, hypervisor-aware, SAN-based snapshot and replication functionality to automate a consolidated, virtual backup and recovery or disaster recovery scenario.

The PS60x0S also comes with a full set of Dell Data Services available, including *Dell ProSupport* for the SMB or department with limited or no dedicated IT staff and *Dell ProSupport for IT*, providing tech-to-tech support for IT professionals, database administrators, and internal help desks. Dell also offers a full set of *Server and Storage Deployment Services* for on-site and remote installation services to help reduce the time required to get your system up and running.

Conclusion

In order to reduce the total cost of ownership of the IT infrastructure, the data center must gain control of a rapidly growing storage network – and it must do so without negatively impacting

performance. Dell has invested in innovative technology to do just that. With the addition of solid state drives in the Dell EqualLogic PS6000 family, the data center can reduce the number of drives required to satisfy mission-critical SLA requirements of their enterprise applications deployed in a virtualized environment. With the PS6000S and the PS6010S, the data center can:

- Improve performance, and performance tuning, for both sequential and transactional data, with automated, real-time load balancing across storage tiers;
- Centralize and simplify storage and data management;
- Simplify ease of use of storage provisioning with intelligent, automated management to minimize otherwise tedious tasks;
- Increase reliability with a fault tolerant, fully redundant solution, designed for no-single-point-of-failure with no moving parts;
- Facilitate a seamless linear scalability with a modular design enabling the data center to grow storage and performance as needed and optimize capacity utilization; and
- Lower the TCO for the data center with an all-inclusive offering of features and functions, without additional licensing, using an affordable iSCSI connection to keep costs down.

The EqualLogic family makes it easy for the SMB to deploy solid state disks, whether they are first-time EqualLogic users or have already deployed an EqualLogic SAN. If your data center has deployed an under-utilized storage infrastructure in order to achieve an IOPS rate commensurate with the SLA required by your user base, then it may be time to investigate the advantages of replacing those mechanical drives with solid-state disks. It is simple and fast to migrate workloads to a PS60x0S array online with no downtime. If you require affordability, scalability, and ease of use to go along with the high performance, then solid state from Dell may be the answer you have been seeking.



⁷ Automatic RAID 5, RAID 6, RAID 10, and RAID 50.

About The Clipper Group, Inc.

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

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

About the Author

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.)***

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.