Bill Mottram and Robin Harris
November, 2008

Research Study

Improving the Economics of Large-Scale Data Storage

Abstract:

Executives seeking to reduce the costs associated with data storage while improving service levels can do both by investing in innovative storage architectures. However, a break from traditional thinking is required if meaningful cost reductions and operational efficiencies are to be realized. Scale out storage is such an approach and is gaining momentum in enterprise data centers that are struggling under the pressure of dramatic data growth and increasing business demands for rapid access to both current and historical enterprise data.

This paper examines the latest entry into the marketplace, the HP StorageWorks 9100 Extreme Data Storage System (ExDS9100) and compares the key cost of ownership factors to three competitive approaches from Isilon, NetApp and Sun to determine the most effective, efficient and economical choice.

The results of Data Mobility Group's research are striking. The HP ExDS9100 has significantly lower cost of acquisition and lower operational costs than comparable alternatives. With a total solution list price below $2/gigabyte for raw storage, combined with a high storage density, the ExDS9100 offers the most attractive economics for fixed content, active archive data storage.

Introduction

In an era of economic stringency, IT and financial executives must look for ways to reduce information technology costs. Data storage, which typically makes up 50% or more of enterprise data center hardware budgets, is a prime candidate.

Despite the current economic turmoil, enterprise data continues to grow at rates that can vary from 50% to over 100% annually. The challenge facing storage administrators is how to manage the growth, maintain increasing service level expectations, and keep within a budget that is flat at best, but probably declining.

Recent research[1] has found that most enterprise file data, once created, quickly becomes inactive or persistent. The term "persistent data" describes fixed data that is accessed less frequently over time, but still considered a valuable digital corporate asset. Persistent data sits dormant on the disk waiting in anticipation of some future requirement that could create an access request. Observation shows that most data within a data center becomes persistent after 30 days and industry estimates suggest that the volume of persistent data in a data center ranges from 70% to 85% of the total data—with much of this data lying dormant for months if not years.

Traditionally, as data aged beyond its initial usefulness it was, by default, left occupying expensive tier 1 and 2 storage just in case it was needed. However the economics of traditional storage made this practice prohibitive. Consequently, much of the historical data within an enterprise was archived on tape and shuffled off into an inactive archive with poor access.

Persistent data is the digital memory of an enterprise and takes many forms. It could be aged accounting information, historical customer data, product data, emails, images, videos, PowerPoint's, schematics, manufacturing data, etc. The possibilities are endless. In today's increasingly competitive world, availability of this data—for reasons of compliance, governance or competitive advantage—is critical.

The modern knowledge worker's requirements for rapid access to historical corporate data have generated new and creative thinking in storage architectures. End user demand coupled with the rapidly declining cost of disk storage makes online data storage an attractive option for persistent

data. Historically the barrier to widespread adoption of disk technologies for mass data storage has been cost, but recent innovations in large scale storage architectures have minimized the cost challenge and enabled the positive attributes of online disk storage to be optimized and applied to the massive volumes of enterprise, persistent data.

An illustration of the growing demand for rapid access to large volumes of historical corporate data is the impressive growth in data warehouses. A recent report by David Winter[2] in Information Week presented data that clearly illustrated the dramatic growth in the size of data warehouses. Admittedly, Winter referenced the largest installations, but data illustrates that enterprises are retaining more data, not just for corporate governance and regulatory compliance but for the business intelligence they gain through data mining. This allows them to drive increased competitive advantage.

Recent changes to the Federal Rules of Civil Procedure that apply to Electronically Stored Information (ESI) have placed a premium on ready access to business records during discovery. These benefits have two key elements:

- Liability reduction through complete discovery and avoidance of damages and adverse interference; and
- Reduced litigation costs through automation of discovery processes. Online storage of business records are critical to managing both issues.

With these increased economic and service level pressures on enterprise IT and a growing awareness of the true access patterns of corporate file data, Data Mobility Group predicts that storage architectures specifically designed to meet the needs of long term retention of persistent data (e.g., scale-out architectures) shall become much more common in the next five years.

Scale-out storage is gaining acceptance as the most effective and economic approach for the long-term storage of large volumes of persistent data that must be online and must deliver the response times characteristic of traditional online data storage. While the benefits of scale-out storage will be discussed in a later section suffice it to say that it differs from monolithic or modular storage systems in that the underlying cluster technology enables cost effective scalability well beyond that of a single array. Scalability extends to management

with complexity minimized and management enabled through a “single pain of glass”. This consolidated view empowers an administrator to comfortably manage data volumes measured in the hundreds of terabytes if not petabytes.

In the remainder of this white paper the Data Mobility Group (DMG) looks at three different approaches to delivering large data repositories. The solutions—from respected storage vendors NetApp, Isilon and Sun—are contrasted to the HP ExDS9100 Extreme Data Storage System.

Scaling out storage for long-term retention of persistent data.

Designers of data storage have for many years been held hostage by traditional thinking. Disk solutions evolved from a transaction environment and storage solution were designed to meet the characteristics for access, availability and reliability for transactional data. For these primary data applications monolithic and modular architectures have dominated but for secondary data applications they tend to lack scalability, have configurations that are inflexible and they are much too costly. Transactional data storage (primary storage) typically requires large and expensive mirrored caches to overcome disk latency and I/O limitations. Disk fan-out is limited to maximize bandwidth and minimize latency and I/O bus contention and special low-latency interconnects are often required to maintain cluster coherency. These components raise the cost of the storage solution.

Secondary data storage has a different set of challenges. Long term, fixed content, historical data, also referred to as persistent data, has unique access, availability and reliability characteristics. Persistent data is largely read-only data. The requesting servers will cache frequently accessed files, reducing the need for storage system caching. Data can be laid out across multiple disk controllers with larger fan-out to maximize total bandwidth without excessive cost and persistent data requires much less cluster intercommunication, reducing cluster interconnects costs. The practical impact of these differences will be apparent when we review the pricing of the competitive systems.

The volume of persistent data is massive dominating the enterprise datastore, creating an economic challenge that separates long term, persistent data storage, access and retrieval requirements from traditional transactional storage. Yet until recently there was no obvious alternative.

However, Google's revolutionary demonstration of high performance storage using ordinary consumer disk drives in large clusters changed the picture. Google proved that modular and monolithic storage arrays were not the only option for mission critical storage infrastructures. The entry of the world's largest IT company, Hewlett-Packard, which ships 1 of every 2 disk drives manufactured, validates the enthusiasm for scale-out storage as a viable solution for long term retention of persistent data.

Scale-out storage architectures are designed to economically meet the scalability, flexibility, on-line response times and data accessibility requirements for long term data storage. Scale-out solutions are characterized by integrating multiple smaller standard, components (server, controllers, storage) to deliver the aggregated effect of a single larger solution. They use clustering technology to create these scalable and upgradeable-in-place storage infrastructures that tend to support massive capacities, be extremely dense, highly flexible and very cost effective. These large scale storage solutions are finding their place in the "clouds" as evidenced by the role Hewlett Packard's, ExDS9100 is expected to play in the open sourced test bed being created for the advancement of cloud computing research. This effort is sponsored by HP, Yahoo and Intel and will provide an internet-scale testing environment.

While clustering technology is well-understood, what CIOs and CFOs want to understand is the economic benefits it confers and it is the objective of this paper to highlight some specific areas where these cost benefits will be realized.

Understanding the trade-offs

Scale-out storage systems typically share several characteristics:

- High density disk drive packaging to ensure space efficient data storage.
- Large capacity disk drives, today typically 1 TB to drive down cost.
- Cluster file systems to enable an easily expandable and upgradeable storage infrastructure.
- High volume hardware components to reduce cost and facilitate upgrades.
- Gigabit or 10 Gb Ethernet access.

However, within these parameters there is a great deal of variation among different vendor's implementations of scale-out storage. Chief among these differences is acquisition cost and storage density.

Other configuration choices, such as the use of blade servers instead of larger rack servers affect flexibility, storage density and file system performance. A blade server chassis can support 8 full height blades or 16 half height blades in the same space as 10 1U or 5 2U servers. Given the speed of price decline and/or performance improvements in storage and CPUs, a flexible modular architecture is a TCO benefit. Storage and processor elements that can be added when needed contribute to lower capital and operating costs.

Over-provisioning is the common response to inflexible system architecture. The ability to easily add arbitrary units of CPU, network and capacity to a system is critical to maximizing TCO.

A key issue in architecting and evaluating scale-out storage is the trade-off between network bandwidth and storage capacity. Network bandwidth, which includes the internal network in scale-out storage, is expensive per gigabit compared to the cost of a gigabyte of disk capacity. For example, at today's prices an 8 port 10Gb switch prices out at about $1500/port. $1500 will buy about 5TB of SAS disk capacity.

As a result, the fan-out ratio - network bandwidth in Gb/s divided by storage capacity in GB - is a key determinate of overall system cost. Note that prices of both network ports and disk capacity are dropping - the precise ratio is less important than the relative costs of bandwidth and capacity.

Thus, while advocates of scale-out storage promote scale-out storage for general-purpose file serving, they do not factor in the costs of the higher bandwidth required for this purpose. This cost differential makes it economically advantageous for IT architects to differentiate between transactional storage and storage optimized for the long-term retention of rapidly accessible persistent data.

The important issue is for planners to be realistic about their data access and bandwidth requirements. It is the Data Mobility Group's opinion that many file servers are significantly over-provisioned with costly bandwidth at the expense of abundant capacity.

Flexibility also extends to CPU provisioning. Different workloads place different demands on a storage system. An uncompressed HD video archive will have many terabytes of large files. File accesses will be infrequent with relatively little NFS or CIFS CPU overhead. Cost-saving fan-out ratios may be large. But as time goes on, more short files - 15 second promos, for example - may be stored and frequently accessed. Adding CPU resources to handle higher request rates is a cost-effective strategy.

Methodology

Total cost of ownership (TCO) studies are difficult to credibly perform and usually even more difficult for users to evaluate. Implicit assumptions can skew the results, especially when looking at soft costs. In DMG's experience many financial people discount claims of operational expense savings, particularly manpower, reasoning that even if the operational savings materialize, they will rarely result in a hard dollar expense reduction. Instead the resources are shifted elsewhere.

In DMG's view, user experience with products is the best guide to operational strengths and weaknesses. A TCO study is best suited to looking at hard dollar costs and for that reason this TCO study focuses its attention on those factors that have a quantifiable effect on life cycle cost.

The one obvious variable that we did not try to monetize was the value of the vendors support infrastructure and the worldwide reach of the vendor under consideration. Support includes break fix and installation to post sales support through to broad and complimentary professional services. Three of the four companies under consideration, HP, Sun and NetApp have extensive support capabilities whose suitability, reach and quality is best evaluated by the prospective user. These variables make differentiation within the context and scope of this paper difficult and most certainly contentious. However users of complex technology must be sure that whichever vendor is selected they are able to get the depth of support they need.

Preparing for the study we began with similar sized configurations, roughly 820 TB, and using published list pricing we developed a comparative pricing evaluation. To calculate user available storage we discount the raw capacity by subtracting the overhead associated with implementing RAID 6 (or RAID 6 equivalent) and the file system. We then look at the power, cooling and

floor space requirements of each configuration to generate hard dollar operational costs.

When compiling each of the configurations used in the comparison we tried to establish equivalency wherever possible. For example as the ExDS9100 can support up to 16 servers we applied the same criteria to Sun and Isilon. However configuration limitations dictated some compromises. In the case of NetApp we focused on a FAS6080 configuration with two filers. To compare a GX cluster configuration would require multiple FAS6080 storage nodes, (up to 8 HA pairs), which would unreasonably skew the comparisons. This choice allowed a comparison against a more traditional architecture.

- The HP ExDS9100 is configured with 820 TB of raw capacity and HP's PolyServe cluster file system. Each implementation can be configured with up to 16 HP blade servers and 10 individual capacity blocks. This separation of performance and capacity modules is key to its flexible scalability.
- The NetApp FAS6080 is configured with 826 TB of raw capacity and configured with ONTAP 7G This configuration is a traditional modular architecture. There was a second reason why we decided not to use ONTAP GX. Our understanding is that the GX configuration is primarily suited for high performance computing and not the active archive space. The FAS6080 can expand to 1008 spindles and managed by two FAS6000 series filers in a fault tolerant configuration. To grow beyond this capacity (up to 14PB) and still enjoy a global namespace across multiple filers the system will require an upgrade to ONTAP GX. The strength of ONTAP GX is to increase compute power when I/O is the issue not capacity.
- The Sun Fire X4540-based system is configured with 816 TB using multiple X4540 storage servers running the ZFS file system locally and the IBRIX Fusion file system to provide the scale-out cluster capability. ZFS does not currently offer a native cluster capability. One of the disadvantages of the Sun solution as configured is the fact the only four Sun Fire units can populate a single rack which reduces storage and wastes data center floor space.
- The Isilon X-Series IQ 12000 is configured with 820 TB of raw storage and uses their OneFS cluster file system. The configuration did include some storage only units but Isilon configuration rules support only one storage extension node per platform node.

DMG chose to focus on vendor delivered systems rather than customer integrated commodity hardware and open source software. While technically astute (and resource rich) customers, such as Google, have achieved good results for specialized applications, enterprises with more general requirements and a greater sensitivity to risk may prefer vendor integrated and supported solutions.

System Configuration Comparison				
	Hewlett Packard ExDS9100	NetApp FAS6080	Sun Fire X4540	Isilon X-Series IQ 12000
File System	PolyServe	Data ONTAP 7G	IBRIX Fusion	OneFS
System Configuration[1]				
Raw Storage Capacity	820	826	816	820
Cost as Configured	$1,506,146	$2,866,560	$1,894,265	$2,871,440
Racks Required	2	5	5	4
Performance Resource Scalability	individual blade servers	individual filers	integrated servers and storage	integrated servers and storage
Storage Units Required	10	59	17	68
Storage Units Per Rack[2]	4-6	10-14	4	21
Spindles Per Storage Unit	82	14	48	12
Maximum Storage Per Storage Unit[3]	82	14	48	12
Total Spindles	820	826	816	820
Supported RAID Levels	6	1, 4 , 6	0, 1, 0+1, 5, 6	6 equivalent
Disk Drive Type	SAS Midline	FC or SATA	SATA	SATA
Storage Efficiency[4]	78%	83%[5]	78%	78%
Available Storage[6]	640	644	636	640
File System Overhead	6%	20%	10%	12%
Overall Efficiency	72%	63%	68%	66%
Total Usable Capacity	590	520	555	541

Table 1: System Configuration Comparison

[1]For the purposes of this comparison a system is defined as a configuration that can deliver approx 820 TB Raw

[2]Rack storage densities vary depending on number of servers/filers installed

[3]1TB drives and TB units used throughout

[4]Raw storage less allocations for RAID, sparing, etc.

[5]Size of NetApp RAID group was assumed to be 12+2, other configurations sized at 8+2.

[6]Before file system overhead

An analysis was completed on each on the competitive solution using data extracted from available published documentation and field data. For the purposes of this comparison a system is defined as a configuration that can deliver approx 820 TB Raw.

To compare customer usable storage the overhead associated with RAID 6 or equivalent was subtracted from the raw capacity number. Subtracting the file system overhead to arrive at the total usable storage for each solution further diminished this number. The adjustments used for each file system overhead was 6% for HP, 20% for NetApp, 10% for Sun and 12% for Isilon. These percentages were based on vendor input, observation and field reports.

One observation from the table is the significantly fewer cabinets required for the HP solution, which is a tangible, end user benefit of their very high storage density.

Physical Attribute Comparison				
	Hewlett Packard ExDS9100	NetApp FAS6080	Sun Fire X4540	Isilon X-Series IQ 12000
Specifications				
Drives Per Storage Unit	82	14	48	12
Fully Populated Rack Weight (lbs)	2027-2415	1529-4059	2176	1622
Single Rack Footprint (sq.ft.)	6.12	6.00	6.43	6.43
Racks Per System	2	5	5	4
Total System Footprint (sq.ft.)	12.24	30	32.15	25.72

Table 2: Physical Attribute Comparison

Table 2 looks at the physical metrics of each of the solutions. The major stand out in the comparison is the required floor space for each full system. The HP ExDS9100 is by far the most space efficient solution with a 110% to 160% advantage over the other solutions. This translates to real operating expense savings as will be discussed later.

Environmental Metric Comparison				
	Hewlett Packard ExDS9100	NetApp FAS6080	Sun Fire X4540	Isilon X-Series IQ 12000
Specifications				
Base Power Consumption (KW) per rack*	11.97-12.42	6.74-7.56	7.58	7.35
Thermal Rating (BTU)*	40860-42378	23010-25800	14971	26250
Air Conditioning Load (tons)*	3.50-3.63	1.97-2.21	1.28	2.25
Air Conditioning Load (KW)*	11.97-12.42	5.56-6.30	4.39	7.69
Total Power Consumption (KW) per rack*	23.95-24.84	12.34-13.86	11.97	15.04
Total Power Consumption (KW) per system	48.79	60.85	54.23	49.42
Energy Efficiency (**Raw** TB/KW)	16.80	13.60	15.90	16.80
Space Efficiency (**Raw** TB/sq.ft.)	67.00	27.50	26.90	31.90
Space Efficiency (**Usable** TB/sq.ft.)	48.20	17.30	18.30	21.00
Cost Per KW (in USD)	$0.11			
BTUs Per System	83238	113105	68031	85312
Energy Cost Per System/Hr	$5.37	$6.69	$5.97	$5.44
Annual Power Cost (8760 hrs)	$47,014	$58,635	$52,257	$47,621

Table 3: Environmental Metric Comparison

*All racks fully populated, with and without servers where indicated.

Table 3 compares the power and cooling requirements for each option. It is no surprise that the power needs for each scale-out solution is very comparable. The exception being the monolithic (scale-up) NetApp configuration which shows a somewhat higher power requirement and lower Tb/KW efficiency ratio. Power needs of a storage array are primarily driven by the disks within an array. It is safe to assume that 80% or more of the power consumed by an enterprise array is due to the rotating disks—the more they are active the more power they draw. Because none of the solutions presented offer any significant power saving techniques that impact disk power requirements, the power draw for each will be similar.

Data Mobility Group

www.datamobilitygroup.com

Scalability

Scalability is one of the most abused words in the computer industry-marketing lexicon. It can mean anything from increasing the capacity of a single storage array to clustering 1000 or more servers.

Traditionally scalability in the storage world has been myopically viewed in terms of capacity growth and addressed by simply adding more storage. However dramatic data growth plus greater demands for online data availability, accessibility and increased operational efficiencies (power, space, management), are all trends that question this lazy man's approach to managing storage growth. Unfortunately however, this tradition still lingers and remains standard operating procedures in many data rich enterprises.

Scalability is a multidimensional challenge and while just one of many selection criteria when making purchase decisions its importance requires deeper thinking. Data growth, increasing application demand, service level expectations and vendor product roadmaps should be understood and balanced when considering the complex question of scalability.

A pragmatic definition for scalability would be; "A system can grow large enough to meet foreseeable application needs without compromising performance and at a reasonable cost - where cost includes hardware,

Snapfish - a study in successful large scale- out storage deployment.

60 million members in 20 + countries and over 5 billion unique photographs stored

Storage grew from 0.5PB to 7PB and expected to grow to 20PB by the end of 2009

Snapfish receive 10 million internet requests for prints on a peak day.
Key requirements are

- Highly scalable storage to handle the millions of new customers that join Snapfish each year, doubling storage.
- Reduced data latency to rapidly display pictures prevented customer frustration and loss.
- Improved storage management allowing five storage administrators to manage 7PB of data.
- Extreme reliability is a critical requirement; downtime means loss of customers.
- Affordable meaning that high performance and high scalability needs to be balanced with low cost.

Snapfish is a Hewlett Packard customer but is an excellent example where the scalability and flexibility of scale-out storage drives significant end user benefit.

 76 Northeastern Blvd. Suite 29A, Nashua NH 03062 Phone: 603.835.6141 Fax: 877.254.4886

integration, floor space, power, cooling and management."

- **Capacity**: Capacity comes in two flavors, raw or usable storage. In the case of the HP ExDS9100, the capacity modules (storage) are physically separated from the compute, performance modules. This level of granularity simplifies the task of increasing capacity transparently without unnecessary performance increases. Sun and Isilon are more tightly integrated architectures and do not enjoy the degree of flexibility presented by HP. NetApp, while storage modules are independent from the filers, the challenges presented to integrate storage is significantly greater than the clustered examples.

- **Performance:** Storage performance normally includes measures of both data latency and bandwidth. However, in the case of scale-out, large volume storage for active archive data, the Data Mobility Group believes that performance considerations should look beyond these traditional measures and consider the ease and transparency of IOP scalability. This key variable drives higher access densities, improving the ability to mange peak query volumes and variable query complexity. A particularly critical metric in today's internet driven world.

  Snapfish is an good example of the challenges facing today's internet driven enterprises.

  In any mass storage implementation application specific requirements have to be considered. Data latency and I/O performance are critical when servicing users dealing with files common in enterprise and e-commerce and when requests are of varying volume and complexity. In contrast are those applications where bandwidth dominates such as video post processing where intensive rendering and nonlinear editing create demanding requirements.

  It is worth underlining that the goodness factor of available performance such as bandwidth is determined by the application and where excessive bandwidth is available but not used results in unnecessary expense.

  Highly flexible architectures are needed to meet these requirements. The HP ExDS9100 has the advantage of employing individual blade servers, called Performance Blocks, based on HP BL460c blade servers. This allows the ExDS9100 to easily increase compute capacity to address both I/O, latency and bandwidth with a degree of granularity

uncommon in comparative solutions.

- **Scale through Time:** Simply put, are expensive investments protected by an architecture that is easy and non-disruptively upgradeable. Can new technologies be easily integrated into the existing architecture extending its useful life or will it take a major, disruptive event to introduce these new technologies. This is a dimension of scalability that receives little attention but considering the amount of data that will reside in these massive storage solutions, the idea of migrating petabytes of data every 3 or so years to the latest and greatest technology is no longer viable. Data repositories must have architectures that allow gradual migration of technology without forklift upgrades. The ExDS9100 is such an architecture that has the flexibility to upgrade transparently assimilate technology as it matures.

  The obvious examples are faster CPUs to increase performance, drive densities to increase capacity, evolving SSD technologies or perhaps holographic recording, interconnect hardware and protocols etc.

  The ExDS9100 has the advantage in that it is highly granular so as each innovation becomes available the task of transparent integration is simplified.

Time to first customer data

Time to production is the measure of ease in installation, integration and vendor support. The massive capacity of scale-out storage creates its own operational difficulties. Loading petabytes of data is a significant challenge that will be an ongoing process that could take weeks or months.

Taking additional time to assemble, expand and integrate scale-out storage is therefore a particular burden. The best scale-out storage systems install and configure quickly.

The HP ExDS9100 meets the challenge by shipping each configuration factory assembled and installs in mere hours. A worthy note is that HP wraps installation cost and support into the system purchase price, an interesting turnkey approach and one that is thoughtful to their customers.

Once installed the challenge of adding addition resources has also been simplified. Once the performance module (blade server) or the capacity module has been physically installed a single

command line enables the transparent integration of these additional resources.
While some vendors factory install the IBRIX Fusion file system, it is not clear if that is the case with Sun. If not then extended installation and setup can be expected. As with all solutions that integrate technology from multiple companies the strength of support, ownership and responsibilities should be established.

The NetApp option is a traditional, modular scale-up solution, which implies that none of the ease of use features associated with scale-out architectures will be available. Not being a cluster architecture it is subjected to the traditional architectural limitation in connectivity, scalability, flexibility and manageability.

Isilon's OneFS Systems are designed to be easy to install. Once powered on and connected to the network a new node will automatically join the local cluster and began performing I/O. Ease of installation and ease of use is a trademark of the Isilon system.

Affordability

As the chart opposite illustrates, there is a major difference in the cost of capacity across the different systems.

The raw cost is the list price of the hardware and software for each configuration divided by the raw capacity. From this DMG derived a cost per user GB by assuming RAID 6 equivalent data protection to protect against data loss due to a disk failure and subsequent unrecoverable read error during a rebuild.

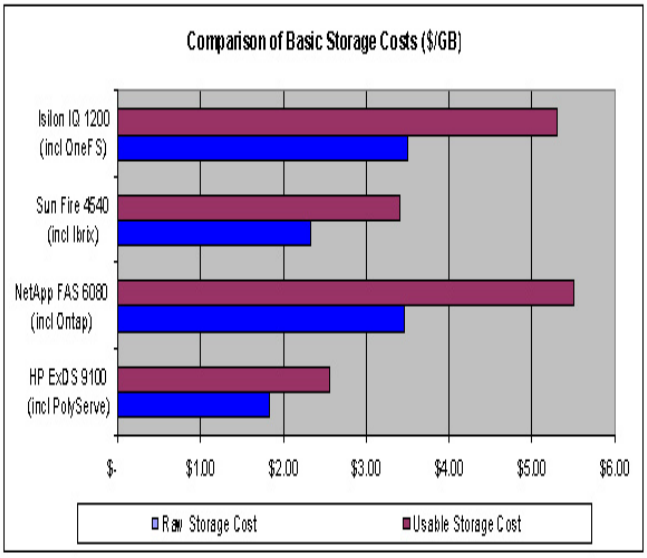

We then adjusted for file system overhead 6% for HP, 20% for NetApp, 10% for Sun and 12% for Isilon, to arrive at the final storage cost per user GB of capacity.

Environmental Efficiencies

Power is an issue of topical interest. Disk drives dominate storage system power consumption and since the vendors buy standard disk drives our study did not find any significant differences in power consumption. Note, no solution employed any meaningful power management techniques such as MAID.

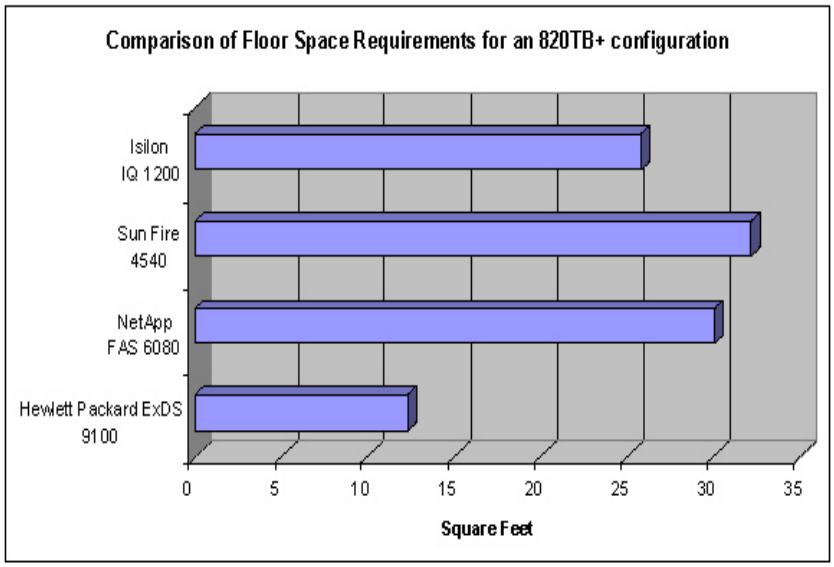

Footprint is the other area that affects TCO. The HP ExDS9100 achieves an admirable density of 67 TB per square foot, more than twice that of the nearest competitor. In urban areas where floor space can be a critical concern the HP ExDS9100 has a clear advantage.

The Impact of Operational Expense On Storage Costs

Impact on Storage Costs				
	Hewlett Packard ExDS9100	NetApp FAS6080	Sun Fire X4540	Isilon X-Series IQ 12000
Annual Energy Costs	$47,014	$46,744	$52,257	$47,621
Annual Floor Space Costs	$23,032	$54,000	$57,870	$46,296
Annual OPEX	$69,046	$100,744	$110,127	$93,917
Impact on $/GB Raw Cost	0.08	0.12	0.13	0.11
Impact on $/GB Usable Cost	0.12	0.19	0.20	0.17
3-Year Impact on Raw Cost	0.25	0.37	0.40	0.34
3-Year Impact on Usable Cost	0.35	0.58	0.60	0.52
Compared to HP	—	+66%	+70%	+48%

Table 4: Impact on Storage Costs

Table 4 provides a quick look at how operational expense, power, cooling and floor space can impact the real cost of storage. While not significant in relative terms, it reflects the impact on the bottom line and is worth considering in these days of constricting budgets and challenged margins.

Note that over a three year period operational expense will add 35¢ to the $/GB cost from HP while adding 58¢ to NetApp, 60¢ to Sun/IBRIX and 52¢ to Isilon. These costs are hard and will increase with time.

Conclusion

It is Data Mobility Group's experience that it is rare for a vendor to have a significant cost advantage with a system of equivalent goodness, but in this case Hewlett Packard has a clear and evidently sustainable advantage. The initial purchase cost is very competitive, its highly dense storage footprint and its competitive management framework keep operating costs very competitive.

It is clear that the scale-out architecture is the solution of choice for large data repositories. It delivers the massive scalability in capacity and performance that is required to meet today's storage challenge. The HP ExDS9100 enhances this scalability attribute with a granularity that simplifies independent growth in either capacity or performance - a major differentiating advantage.

It would appear that the ExDS9100 is particularly suited for those massive data repositories that populate large corporations, research institutions and those data rich, growing internet based enterprises. The example of Snapfish is a clear example of how this technology has evolved to meet the demands of internet commerce.

The scalability of both Sun/IBRIX and Isilon configurations are hurt because the configuration flexibility is limited. Each tightly bind their servers and storage limiting flexibility and Sun limits the number of Sun Fires that can be hosted in a single rack. Isilon offers accelerator nodes and capacity nodes, but only one capacity node can be attached to an accelerator node. Sun has a

similar limitation. As a result they cannot get the fan-out, with its resulting efficiency in network provisioning, that the HP ExDS9100 achieves. In fairness, the Isilon system is not aimed at the scale out, persistent data market. Instead they are optimized for the video production space where large scale, very high-bandwidth is usually required. The HP ExDS9100 is well-suited for serving video content but its design is not optimized to handle bandwidth intensive rendering and nonlinear editing applications.

The big win for the HP ExDS9100 is in the day-to-day world of real life file serving. As research has demonstrated most business files are rarely accessed after they've been created and reviewed. This means that file servers that serve humans rather than other computer systems can be configured for very large capacity without excessive bandwidth deman As the HP ExDS9100 demonstrates there are significant cost advantages to this architecture.

Endnotes

[1] "Measurement and analysis of Large Scale Network File System" Workloads; Andrew Leung, Shanker Pasupathy,Garth Goodson and Ethan Miller; presented at USENIX 08, Technical Conference