

WHITE PAPER

Bringing Clarity to Hard Disk Drive Choices for Enterprise Storage Systems

Sponsored by: Hewlett-Packard

John Rydning
February 2009

INTRODUCTION

Technology migrations for IT organizations are often challenging. But choosing the right storage hardware technology wisely, and implementing it appropriately for a given storage application and its associated workload, can lead to improved performance and lower operating costs. The keys to realizing these benefits are understanding emerging hard disk drive (HDD) technologies and evolving HDD form factors and selecting those technologies that will have longevity and broad adoption and that are backed by an entire ecosystem of suppliers and standards organizations. Awareness of these HDD technology trends is important for IT professionals who are making decisions about future enterprise server and storage system architectures.

IN THIS WHITE PAPER

In this white paper, IDC highlights current enterprise server and storage system trends influencing the design of HDDs used in storage systems and discusses several HDD products and technologies entering the market that will likely become mainstream in the near future.


SITUATION OVERVIEW

Data Growth Continues

Digital content growth is a long-term phenomenon with which organizations will continue to cope for the foreseeable future. In fact, disk drive shipments, as measured in terabytes (TB), for enterprise storage system applications will increase more than threefold from 2008 to 2012 (see Figure 1).

FIGURE 1

Cumulative Worldwide HDD Shipments in Terabytes for Enterprise Storage Applications, 2008–2012



Source: IDC, 2009

Fortunately, the disk drive industry has done a commendable job of delivering disk drives with steadily increasing storage densities. From 1997 to 2008, the average growth rate of the capacity of performance-optimized HDDs designed for transaction-intensive workloads used in enterprise server and storage applications increased at a compound annual growth rate of 38%. While this advancement is impressive, it was not good enough to keep up with the pace of enterprise storage requirements. As a result, IT managers were confronted with the dilemma of economically finding a way to meet TB storage growth requirements of greater than 50% per year with disk drive capacities that were increasing at a rate of less than 40% per year.

The Emergence of the Capacity-Optimized Enterprise Disk Drive

Rather than simply adding more disk drives per array, IT managers realized that not all data is created equal and that certain data could be stored on lower-cost, less frequently accessed disk drives. What followed was the development of the capacity-optimized HDD to store low-intensity workload data (optimized for enterprise storage). Identifying lower "tiers" of data to store on capacity-optimized HDDs continues to be an important strategy for IT organizations.

The capacity-optimized HDD platform evolved from an HDD design used initially for desktop PCs. With this HDD platform, a low-cost Serial ATA (SATA) interface with "good enough" performance generally is used for most PC and CE applications. Early adoption of these HDDs by server and storage system platforms exposed some of the weaknesses inherent with HDDs designed for relatively low workloads found in PC and CE applications as opposed to 24 x 7 enterprise datacenter operating environments. However, once HDD OEMs foresaw the inevitable use of these drives in server and storage applications, they responded by tailoring the design of the capacity-optimized drive to deliver a combination of capacity and reliability in storage arrays affordably.

Today, capacity-optimized HDDs with a SATA interface are widely used for appropriate enterprise storage workload applications, but other native physical interfaces are available as well, including Serial Attached SCSI (SAS) and Fibre Channel (FC). One example is HP's EVA storage array. In this system, capacity-optimized HDDs equipped with a native FC interface are used to deliver a combination of capacity, reliability, low power, and improved overall storage system performance.

Fibre Channel: The Interface Notorious for Performance

The FC interface, first launched in 1996, was designed specifically for enterprise storage environments to provide new levels of redundancy, scalability, and performance. This interface is a point-to-point topology where HDDs are typically deployed in an arbitrated loop. The FC protocol allows for a large number of HDDs to be attached to a fabric topology, yet still allows for efficient and reliable communication from the host with each drive. The FC interface offers fast serial data transfer rates of up to 4Gb/sec with built-in redundancy via dual-port access to each disk drive.

Large, external enterprise storage systems consume the vast majority of FC HDDs that ship today. These storage systems leverage the performance and reliability benefits associated with the FC protocol, coupled with a rugged mechanical HDD platform used for 10,000rpm and 15,000rpm performance-optimized HDDs.

However, a challenge associated with FC storage systems is integrating the capacity-optimized HDDs into the storage array. The FC interface lacks the capability to directly control the SATA interface protocol commonly found on capacity-optimized HDDs.

One solution is to use a native FC interface in place of a native SATA interface on the HDD. HP, for example, uses a native FC interface on the capacity-optimized HDDs in its current generation of EVA storage arrays. The drawback with a native FC interface on a capacity-optimized HDD is a potentially higher-cost HDD given that more sophisticated FC electronic (read *expensive*) components must be incorporated into the HDD design. Further, there is a lack of multiple sources for capacity-optimized HDDs that have a native FC interface.

Another solution for incorporating capacity-optimized SATA HDDs into an FC storage array is to use an interposer card that physically sits between the HDD and the storage system backplane. This card interfaces with the SATA protocol on the HDD and emulates the FC protocol to the storage system host. This intermediary device also provides an opportunity for interposer card vendors to add drive diagnostic functionality to a capacity-optimized HDD (not a native feature of the SATA protocol). The drawbacks of using interposer cards are a potentially higher cost for the overall storage system, increased power consumption, and an additional potential point of failure in the system.

A "Hot" Challenge Within Datacenters: Power and Cooling

The growth of digital content has brought a correlated increase in the number of disk drives in enterprise server and storage systems. The adoption of capacity-optimized HDDs into enterprise servers and storage systems has alleviated some of the pain associated with digital content growth, but it has not completely eradicated a number of ongoing challenges.

IT managers must still deal with limited square footage in datacenters, limited power and cooling supplies, and limited budgets. The tremendous growth in the capacity of disk drives has helped to reduce the price per GB of new HDD storage, and it also has helped to reduce the watts per GB to power and cool the drives. Granted, the cost to power and cool storage is not as severe as the cost to power and cool servers. Nevertheless, it continues to increase, and organizations are looking for ways to reduce power consumption wherever they can.

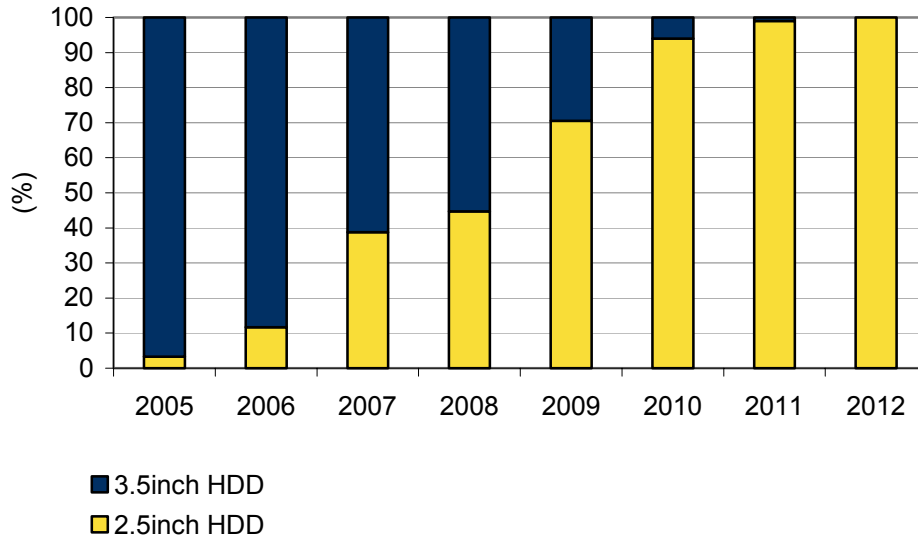
Getting Smaller: The Transition from 3.5inch to 2.5inch Disk Drives

One way to reduce the power consumption per disk drive is to transition to smaller form factors. A move from a 3.5inch to a 2.5inch HDD can cut the watts consumed per drive in half. To address the needs of power, density, and performance, the HDD industry introduced the performance-optimized 2.5inch disk drive, also referred to as a small form factor (SFF) enterprise-class drive, in 2004.

Adoption of the performance-optimized 2.5inch form factor has been relatively quick. In examining the mix of HDDs shipped in 2008 for the storage shipped within servers, IDC found that 45% of the performance-optimized HDDs used were 2.5inch form factor products. As shown in Figure 2, IDC expects that virtually all performance-optimized HDDs that ship with servers will consist of the 2.5inch form factor HDD by 2011. Propelling the crossover in volume from 3.5inch to 2.5inch form factor HDDs will be the launch of 300GB performance-optimized 2.5inch HDDs in 2009.

FIGURE 2

Worldwide Share of Performance-Optimized HDD Shipments for Storage Internal to Servers by Form Factor, 2005–2012



Source: IDC, 2009

Table 1 illustrates the various types of HDDs available today for use in enterprise server and storage systems.

TABLE 1

Current Hard Disk Drive Options for Enterprise Server and Storage Applications

Drive Classification	Form Factor	RPM	Interface	Current Maximum Capacity	MTBF
Performance optimized	3.5inch	10,000 and 15,000	FC, SCSI, SAS	600GB	1.6 million hours
Performance optimized	2.5inch	10,000 and 15,000	FC, SCSI, SAS	300GB	1.6 million hours
Capacity optimized	3.5inch	5,400 and 7,200*	SATA, FC, SAS	2.0TB	≥1.2 million hours
Capacity optimized	2.5inch	5,400 and 7,200	SATA, FC, SAS	500GB	<1.0 million hours

* One HDD program is available with a 10,000rpm spin speed.

Note: Capacity-optimized HDD MTBF ratings are based on lower workloads than performance-optimized HDDs.

Source: IDC, 2009

HDDs used for entry servers are not included in Table 1. An entry server, generally consisting of one or two HDDs, continues to leverage HDD platforms designed mainly for PCs.

External Storage: Poised to Transition to 2.5inch Performance-Optimized HDDs

Today, the vast majority of external enterprise storage systems utilize 3.5inch HDDs given the capacity advantage of the 3.5inch form factor versus the 2.5inch form factor and the commensurate lower cost per GB. However, there is a growing realization that the acquisition cost per GB is not the only cost consideration for storage system customers. As mentioned earlier, the cost to power and cool storage is also a growing concern.

Performance and Cost: The Power of the HDD Form Factor

The 2.5inch SFF performance-optimized HDD is already well-recognized for providing greater storage density and higher IOPS per U in server and storage systems as well as consuming less power. But a 2.5inch form factor HDD carries a capacity penalty of roughly half that of a similar-generation 3.5inch HDD. The reason is simple: 3.5inch HDDs can have a maximum of four platters per drive, while current performance-optimized 2.5inch HDDs have a maximum of two platters per drive. However, this is not a fixed rule, and HDD configurations are about to change.

By 2010, the HDD industry is expected to increase the maximum number of platters per 2.5inch performance-optimized HDD from two to three, enabling them to accelerate delivering a doubling of capacity per drive, and subsequently achieving 50% capacity increases per drive over a shorter time frame. The additional capacity per drive will again reduce watts per GB and will make 2.5inch HDDs an even more compelling option from a power standpoint compared with 3.5inch performance-optimized HDDs. Tables 2 and 3 provide an overview of IDC's expectations for new generations of performance-optimized HDDs through 2012. Note that IDC expects that the HDD industry's last generation of 3.5inch performance-optimized HDDs will be launched in 2009.

TABLE 2

Performance-Optimized 3.5inch Hard Disk Drive Products, 2007–2012

	2007	2008	2009	2010	2011	2012
Maximum capacity	300GB	450GB	600GB	600GB	600GB	600GB
Maximum number of platters per drive	4	4	4	4	4	4
Spin speed (rpm)	10,000, 15,000	10,000, 15,000	10,000, 15,000	10,000, 15,000	10,000, 15,000	10,000, 15,000
Interfaces	FC, SCSI, SAS 1.0	FC, SCSI, SAS 1.0	FC, SAS 1.0, and SAS 2.0	SAS 2.0	SAS 2.0	SAS 2.0
Average price per GB	\$1.06	\$0.84	\$0.58	\$0.43	\$0.32	\$0.25
Average watts per GB 15,000rpm	0.048	0.034	0.025	0.025	0.025	0.025
Average watts per GB 10,000rpm	0.046	0.036	0.023	0.023	0.023	0.023

Note: Watts calculation is weighted by capacity and interface and workload assumptions.

Source: IDC, 2009

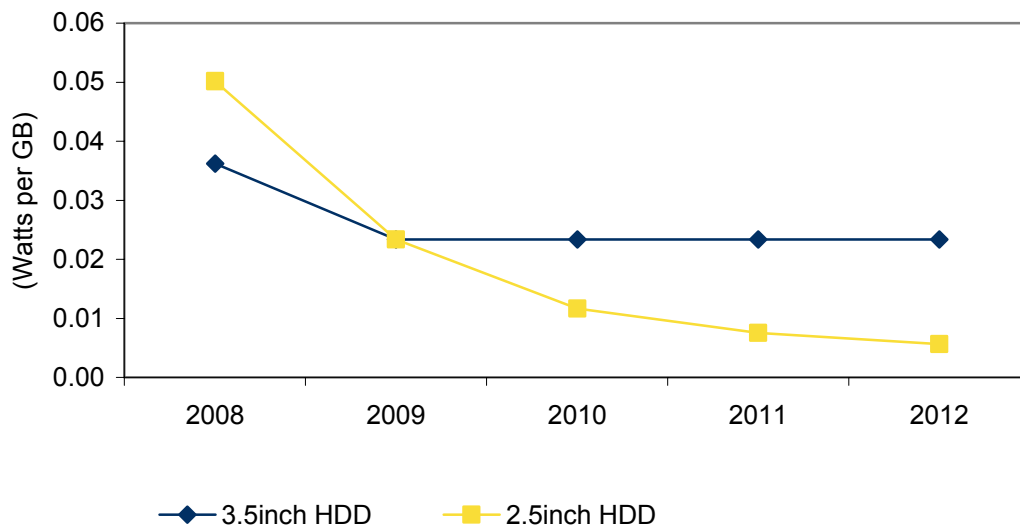
TABLE 3**Performance-Optimized 2.5inch Hard Disk Drive Products, 2007–2012**

	2007	2008	2009	2010	2011	2012
Maximum capacity	147GB	147GB	300GB	600GB	900GB	1.2TB
Maximum number of platters per drive	2	2	2	3	3	3
Spin speed (rpm)	10,000, 15,000	10,000, 15,000	10,000, 15,000	10,000, 15,000	10,000, 15,000	10,000, 15,000
Interfaces	FC, SAS 1.0	FC, SAS 1.0	FC, SAS 1.0, SAS 2.0	SAS 2.0	SAS 2.0	SAS 2.0
Average price per GB	\$1.54	\$1.06	\$0.70	\$0.45	\$0.32	\$0.24
Average watts per GB 15,000 rpm	0.050	0.050	0.025	0.012	0.008	0.006
Average watts per GB 10,000 rpm	0.051	0.050	0.023	0.012	0.008	0.006

Note: Watts calculation is weighted by capacity and interface and workload assumptions.

Source: IDC, 2009

Note that 2.5inch performance-optimized HDDs will nearly reach parity with 3.5inch performance-optimized HDDs by 2010. Moreover, given that the average watts per GB for a 2.5inch performance-optimized HDD will be roughly half that of an equivalent-capacity 3.5inch HDD by 2010 (see Figure 3), 2.5inch performance-optimized HDDs will offer the best combination of price and watts per GB. IDC anticipates that enterprise storage systems will begin to incorporate 2.5inch HDDs into system designs.

FIGURE 3**Average Watts per GB for Performance-Optimized HDDs:
3.5inch Versus 2.5inch, 2008–2012**

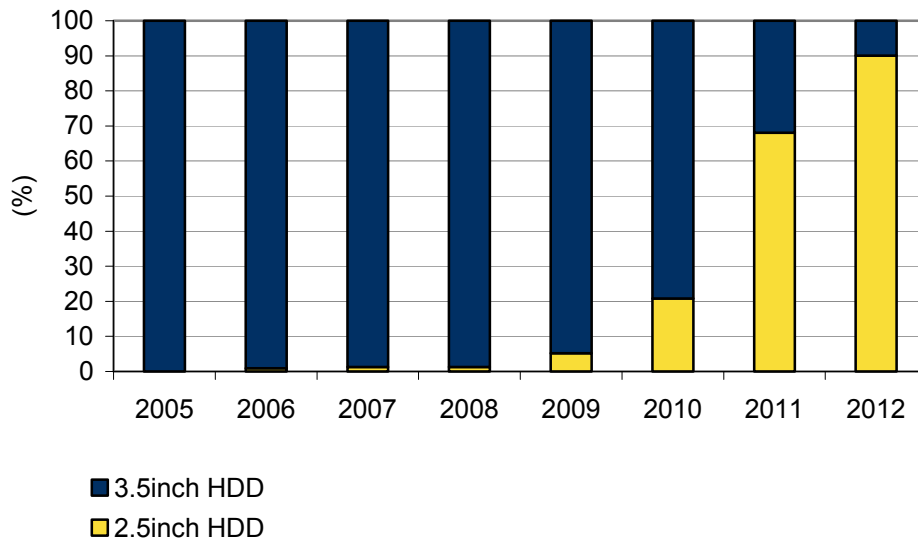
Note: Comparison is for 10,000rpm products.

Source: IDC, 2009

IDC expects that the transition to 2.5inch performance-optimized HDDs for external storage will happen quickly given the assumption that the HDD industry will not launch another generation of 3.5inch performance-optimized HDD with more than 600GB of capacity (see Figure 4).

FIGURE 4

Worldwide Share of Performance-Optimized HDD Shipments for Storage External to Servers by Form Factor, 2005–2012



Source: IDC, 2009

SAS: Emerging as a Viable Disruption to FC Inside Storage Systems

There is certainly a large installed base of enterprise storage systems with FC HDD slots and FC technology expertise. FC technology is mature and has a reputation of being best in class for performance and reliability. Furthermore, the FC interface historically provided, until just recently, a data transfer rate advantage.

Disruption of an existing technology is difficult, but it is possible when the new technology is "good enough" to meet the essential requirements for storage systems at a substantial cost savings for users and solution providers. The SAS 2.0 interface and protocol, developed by an entire ecosystem of suppliers, is backward compatible with SAS 1.0 and will begin to ship in storage arrays in 2009.

The SAS 2.0 interface boasts a 6Gb/sec data transfer rate, which is faster than the current 4Gb/sec data transfer rate available for the FC interface. IDC believes that all of the essential storage system components from most suppliers, HDD vendors included, are prepared to support the transition to SAS 2.0 technology.

Current SAS 1.0 HDD products fell just short of being "good enough" for external storage systems. But the new SAS 2.0 specification will make possible large, rich topologies that allow for larger-scale storage system designs that can compete with storage systems built around an FC infrastructure.

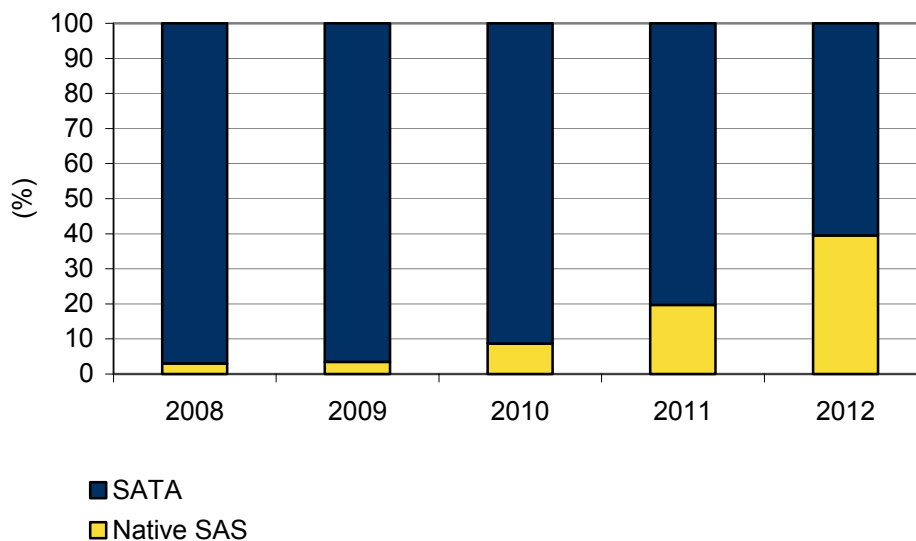
One advantage of migrating to SAS 2.0 for an external storage system OEM is the opportunity to leverage the same interface that will be used within internal storage systems and to achieve better economies of scale for all components, hence lower costs. For system builders and suppliers, it reduces the number of backplanes and components to develop and validate. For end users, it simplifies technology choices and helps to future proof investments in the SAS architecture. Additionally, lower ongoing maintenance and support costs will be made possible given the lower number of replacement or upgrade HDD SKUs that will be needed in the future. This will allow for simpler and leaner inventory management throughout the supply chain.

SAS Support of the SATA Protocol

One benefit associated with the SAS interconnect technology is its ability to support both the SAS and SATA interface protocols and to physically connect each type of drive directly to a common SAS backplane. This makes it possible to use a capacity-optimized SATA HDD when dual-port access is not a performance requirement, yet have the option to use a native SAS interface capacity-optimized HDD when a given application or task requires redundant data paths. IDC's outlook for the mix of capacity-optimized HDDs shipping with either a SATA or SAS interface is provided in Figure 5.

FIGURE 5

Worldwide Capacity-Optimized HDD Shipment Share: SAS Versus SATA, 2008–2012



Source: IDC, 2009

In short, the SAS architecture enables customers to incorporate a mix of both performance- and capacity-optimized HDDs in a server or storage array to achieve an optimal balance between capacity, cost, reliability, and availability.

HP's HDD Solutions for Enterprise Server and Storage

Server and storage system architectures continue to evolve to support a wide variety of applications and installation environments. Whether now or in the future, HP believes it is prepared to help its customers transition to the most advanced HDD technologies needed in storage systems to realize their overall cost, performance, power consumption, availability, and reliability goals for data storage.

A brief overview of the HDD products that HP will be offering for major enterprise storage system segments leveraging many of the new technologies mentioned earlier includes:

- ☒ Enterprise SFF 2.5inch 10,000rpm performance-optimized HDDs with up to 300GB of capacity and the new SAS 2.0 interface that allows for up to 6Gb/sec data transfer rates. This new product will help to accelerate the transition to 2.5inch HDDs in enterprise server storage systems. Moreover, the combination of capacity, power, and performance in this product will act as a catalyst that starts the adoption of 2.5inch performance-optimized HDDs in external enterprise storage systems.
- ☒ Enterprise SFF 2.5inch 15,000rpm performance-optimized HDDs with up to 146GB of capacity and the new SAS 2.0 interface. This drive will provide customers with an economical I/O storage performance boost in enterprise servers.
- ☒ Enterprise large form factor (LFF) 3.5inch 10,000rpm and 15,000rpm performance-optimized HDDs with up to 600GB of capacity and either the new SAS 2.0 interface or an FC interface with 4Gb/sec data transfer rates. Current storage system architectures designed around this form factor will be able to realize a boost in both performance and capacity.
- ☒ Midline 3.5inch capacity-optimized HDDs with a native SAS 2.0 interface for storage applications that have less intensive workloads but need the availability and reliability associated with dual paths to drives in the storage system architecture. Capacity-optimized HDDs with a SAS 2.0 interface will be available to HP's customers with up to 2TB of storage capacity in 2009.
- ☒ Midline 3.5inch capacity-optimized HDDs with a SATA interface designed for the enterprise environment aimed at providing cost-effective storage for less intensive workloads. The HDD industry, working cooperatively with storage system companies such as HP, will continue to increase the capacity of these drives each year, realizing that HDD reliability must not be sacrificed with each new product generation. HP will offer its customers capacity-optimized HDDs with a SATA interface with up to 2TB of storage capacity in 2009.

- ☒ Midline 2.5inch capacity-optimized HDDs designed to fit into the same slot as an enterprise SFF 2.5inch HDD with either a SAS 2.0 or SATA interface. This emerging HDD form factor will give customers new flexibility to optimize the capacity of a given storage array cost-effectively.

In addition to these HDD products, HP will continue to provide its entry server customers with relatively low-cost and reliable HDDs at capacities similar to those generally available for most PCs in the market today.

Future Outlook

Enterprise users, solution providers, and manufacturers are not quick to move off of legacy HDD technology. IDC has tracked several prior HDD technology transitions for enterprise applications closely and has found it can take several years to complete a migration.

Nevertheless, given the growing concern with the cost to power and cool storage, IDC believes that real cost savings can be realized when making the transition to smaller HDD form factors. Legacy systems will require ongoing support and will continue to consume 3.5inch form factor products well into the future. But the momentum is clear: Performance-optimized HDD shipments are shifting quickly to the 2.5inch form factor. IDC expects that more 2.5inch than 3.5inch performance-optimized HDDs will ship in 2010 and that the HDD industry's R&D efforts will be focused on future generations of this form factor.

In addition, IDC believes that new server and storage system architectures and platforms will lead the transition to 6Gb/sec SAS, including:

- ☒ Virtualized servers that effectively reduce costs by consolidating multiple logical servers into one physical server. A consolidated server is much more I/O and storage intensive and demands higher throughput I/O from the server and storage subsystems.
- ☒ High-end servers and storage systems running storage-intensive applications such as mission-critical databases and online transaction processing that require quick data access and retrieval.

Challenges for HP and Its Customers

There are more enterprise HDD choices for servers and storage than ever before. This variety equates to more flexibility for HP's customers, as well as the emergence of new server and storage system architectures. At one time, the HDD choices for enterprise server and storage used to be simple. But today they are growing more complex given all of the new HDD form factor and interface options available.

Most enterprise customers have stable datacenter system architectures in place and have no tolerance for disruption. Many times, these systems are responsible for generating significant revenue for the company, as well as sustaining current business processes. Hence, there will be some reluctance to move off of these stable platforms to newer technologies (especially if the costs of the newer systems

are higher). HP must illustrate the long-term advantage of migrating to these new platforms and articulate the clear advantages from various perspectives (e.g., green, performance, consolidation).

HP will also need to manage the transition of its own product portfolio. Long-tail platforms will need to be supported, but sales of these platforms must be brought to end of life fairly quickly to prevent a burgeoning and perplexing plethora of storage solution options. Product choices must be clear and intuitive, with logical migration paths from one family to another, or else there is the potential for customer confusion and complexity.

HP will have an important role to play to help its customers understand the underlying HDD technologies used within datacenters and the pros and cons of each technology. HP will need to be prepared to educate its customers about the new HDD technologies and form factors that are now available. With proper guidance, HP's customers will be able to select the best combination of enterprise HDDs for servers and storage that will deliver many years of reliable service.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2009 IDC. Reproduction without written permission is completely forbidden.