Performance factors for HP ProLiant Serial Attached Storage (SAS)

technology brief

Abstract

Enterprise-class hard disk drives (HDD) must meet the maximum reliability and scalable performance goals of the 24/7 enterprise environment. Serial Attached SCSI (SAS) has become the defacto HDD standard for mission-critical applications. This paper explains the emergence of SAS and the key parameters of SAS drive technology. It also includes technical data and comparison information of the latest small form factor SAS drives available from HP at time of publication.

Introduction

The mission-critical, 24/7 enterprise environment places stringent demands on storage technology. Reliability and performance are of paramount importance. While server downtime can be costly, even slow server performance can increase operating costs. Hard disk controllers, interfaces, and drives have evolved to address the reliability and performance issues of the enterprise storage system.

HDD controller technologies

To enhance the operating efficiency of multiple drives, HDD controllers use logic-based solutions including:

- Data buffering and read/write caching
- Queuing control and read/write reordering
- Error management and pre-failure warning
- Redundant Array of Independent Drives (RAID)

HDD controllers employ on-board data buffering and caching techniques to avoid the use of slower system memory. Queuing techniques such as Tagged Command Queuing (TCQ) and Native Command Queuing (NCQ) allow controllers and compatible HDDs to take advantage of the read/write head position for more efficient drive operations. RAID is another logic-based solution that places data in stripes across multiple drives to enhance reliability, performance, and data integrity. A variety of RAID strategies[1] are in use:

- RAID 0 – Striping to two or more disks (no parity) for performance improvement (no redundancy)
- RAID 1 – Mirroring data on two disks (no parity) for redundancy, slight performance improvement
- RAID 0+1 – Mirroring and striping for redundancy and performance improvement
- RAID 1+0 (10) – Mirroring and striping for redundancy and performance improvement
- RAID 3 – Striping (byte level) with parity for improved performance and fault tolerance
- RAID 4 – Striping (block level) with parity for improved performance and fault tolerance
- RAID 5 – Striping with distributed parity for improved performance and fault tolerance
- RAID 6 – Striping with dual parity for improved performance and fault tolerance[2]

The choice of RAID strategy depends on the desired balance of protection and performance, along with the number of hard drives available.

---

[1] For more information about disk drive technology and RAID refer to the HP technology brief titled "Disk Drive Technology Overview" available at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01071496/c01071496.pdf .

[2] For more information about RAID 6 technology refer to the HP technology brief titled "RAID 6 with HP ADG Technology" available at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00386950/c00386950.pdf

HDD interface technologies

Since the days of Advanced Technology Attachment (ATA, also known as Integrated Drive Electronics or IDE), the server industry has transitioned through several HDD interface technologies:

- Small computer system interconnect (SCSI)
- Serial Attached ATA (SATA)
- Serial Attached SCSI (SAS)

Table 1 lists the key characteristics of these HDD interface technologies.

Table 1. Comparison of HDD interface technologies

	SCSI	SATA	SAS
Transfer/connection type	Parallel/shared bus	Serial/point-to-point	Serial/point-to-point
Current bandwidth [1]	320 MB/s	3.0 Gb/s (300 MB/s)	3 GB/s (300 MB/s)
Future bandwidth growth?	No	Yes, to 6 Gb/s	Yes, to 6 Gb/s and 12 Gb/s
# of devices supported per interface per connection	16	15 [2]	16, 256 [3]
HDD type supported	SCSI	SATA	SAS and SATA
Relative reliability	Good	Adequate	Very good
Best suited for	Enterprise servers (replaced by SAS)	Entry level servers	Enterprise servers

NOTES:
[1] Actual data rates are slightly lower due to protocol overhead.
[2] Through the use of a SATA port multiplier
[3] Through the use of SAS port expanders

As Table 1 suggests, the SAS interface offers the best solution for the enterprise environment and has emerged as the preferred choice for high input/out applications.

Key HDD design parameters for enterprise environments

Enterprise-class HDDs must provide maximum performance under a 100 percent duty cycle and continuous I/O workload in a high-vibration environment. While Mean Time Before Failure (MTBF) is used to express the length of HDD life in general, a more meaningful benchmark, the Annual Failure Rate (AFR), better defines the estimated life of an HDD in the enterprise environment.

The AFR is the relation (in percent) between the MTBF and the number of hours that the device is expected to run per year (100 percent duty cycle = 8760 hours per year). For example, for an enterprise HDD with an MTBF of 1,200,000 hours, the AFR is calculated as follows:

1,200,000 hours/8760 hours = 136.9863 years, *then* (1 failure/136.9863 years) x 100% = AFR of 0.73%

An AFR of 0.73 percent means that 0.73 percent of the population of HDDs can be expected to fail in the average year. In other words, in a system of 100,000 drives, 730 could be expected to fail. The AFR is less applicable for smaller systems but is meaningful for high-density infrastructures with thousands or hundreds of thousands of drives. SAS drives, particularly small form factor (SFF) SAS drives, typically have lowest AFR[3].

[3] The AFR calculations given are for illustration purposes only. The actual failure rate experienced could vary depending on manufacturing deviations, material quality, and the actual application environment, among other factors.

Another important factor to HDD performance is seek time, described as the time from when a read or write action is initiated until the data transfer from or to the disk actually begins. The smaller platters of SFF SAS drives inherently yield lower seek times, an advantage in file servers with frequent random accesses.

SAS use for enterprise systems

This section describes SAS design features that affect the performance and operating efficiency of SAS HDDs for enterprise-class systems. For a general discussion of SAS technology, refer to the HP white paper "Serial Attached SCSI Technology" available at the following HP web page: http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00302340/c00302340.pdf

Since RAID strategies and AFR numbers favor high-density systems, SFF SAS drives are an attractive solution for large enterprise systems. Compared to 3.5-inch large form factor (LFF) SAS drives, 2.5-inch SFF SAS drives have these advantages:

- Lower AFR
- 15 to 20 percent higher Input/output Operations Per Second (IOPS)
- 30 percent smaller physical size
- Up to half the power consumption

Increasing HDD density in a RAID-based system invariably enhances the overall performance of the storage system. Increasing the drive density also expands the choice of RAID strategies as indicated in Table 2.

Table 2. Comparison of drive support in two 1U HP ProLiant servers

Server	**DL320 G5**	**DL360 G5**
Drive type	LFF SAS [1] or SATA	SFF SAS or SATA
# of drives	4	6
RAID support	RAID 0, 1 [1]	RAID 0, 1, 10, 5, 6 [2]

NOTES:
[1] Requires optional SAS controller
[2] RAID 6 requires P400i with 512MB Battery Backed Write Cache option

At the drive level, smaller/faster is usually better. SFF drives offer a spindle-per-U advantage that yields better performance in most systems.

Traditionally, LFF drives have a cost-per-GB advantage over SFF types. For workstation and small-to-medium business customers, getting the most storage capacity for the money is a key factor in setting up a system. For the enterprise IT architect, however, the goal should be to get the most work (that is, more IOPS) out of the system, and this is best achieved with high density storage systems.

HDD performance comparisons

This section provides performance results of comparisons between HDD interface types, form factors, and platter speeds. Other factors such as operating system, application environment, controller type, RAID strategy, and general workload also play a part in system performance and should be considered when evaluating systems.

Performance comparison between SATA and SAS drives

Systems that support SAS can be configured with a mix of SAS and SATA drives[4], although performance goals are more easily achieved using SAS drives. Figures 1 through 3 are plots of performance data from tests of SATA drives (non-NCQ and NCQ) and SAS drives using the following configuration:

- System: HP StorageWorks 60 Modular Smart Array
- SAS Controller: HP Smart Array P400, firmware version 4.12, 512 MB cache, RAID 1

Figure 1. Relative performance: sequential operations of non-NCQ SATA, NCQ SATA, and SAS 2-drive array with HP Smart Array P400 Controller, RAID 1

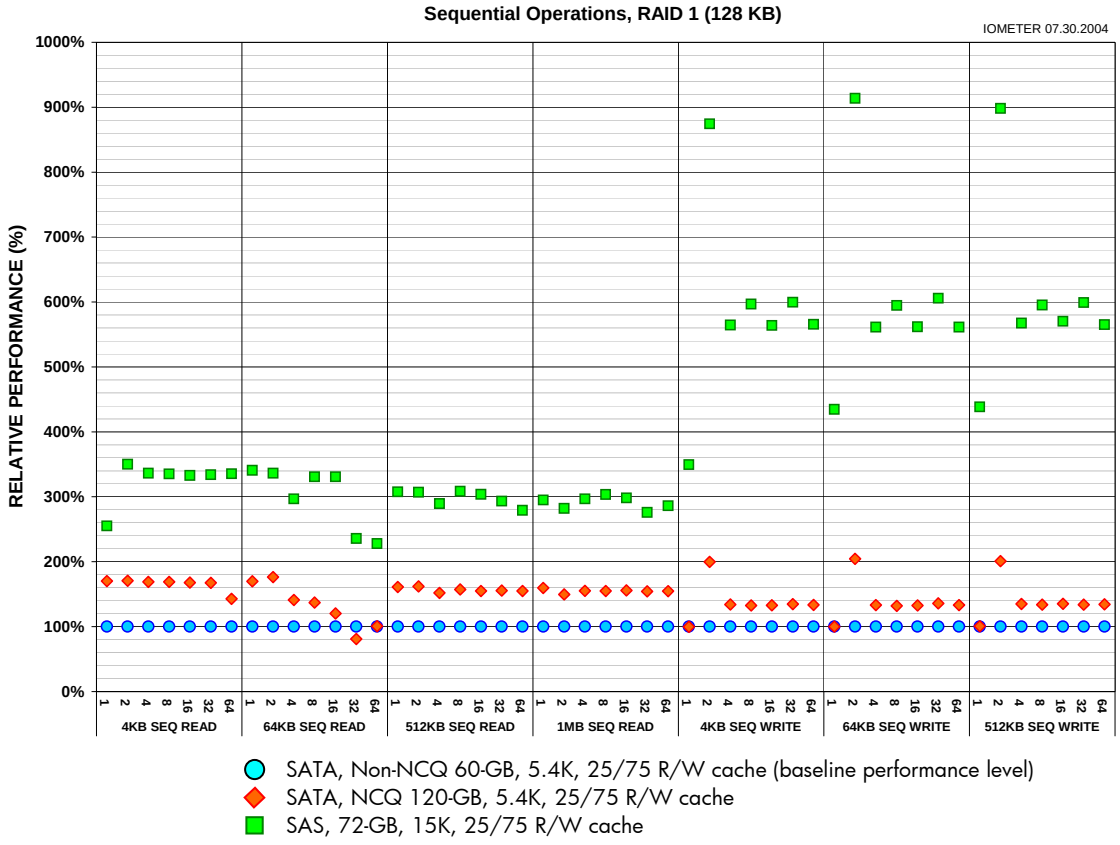

[4] Mixing SATA and SAS drives within a logical drive is NOT allowed.

Figure 2. Relative performance: random operations of non-NCQ SATA, NCQ SATA, and SAS 2-drive array with HP Smart Array Controller P400 Controller, RAID 1

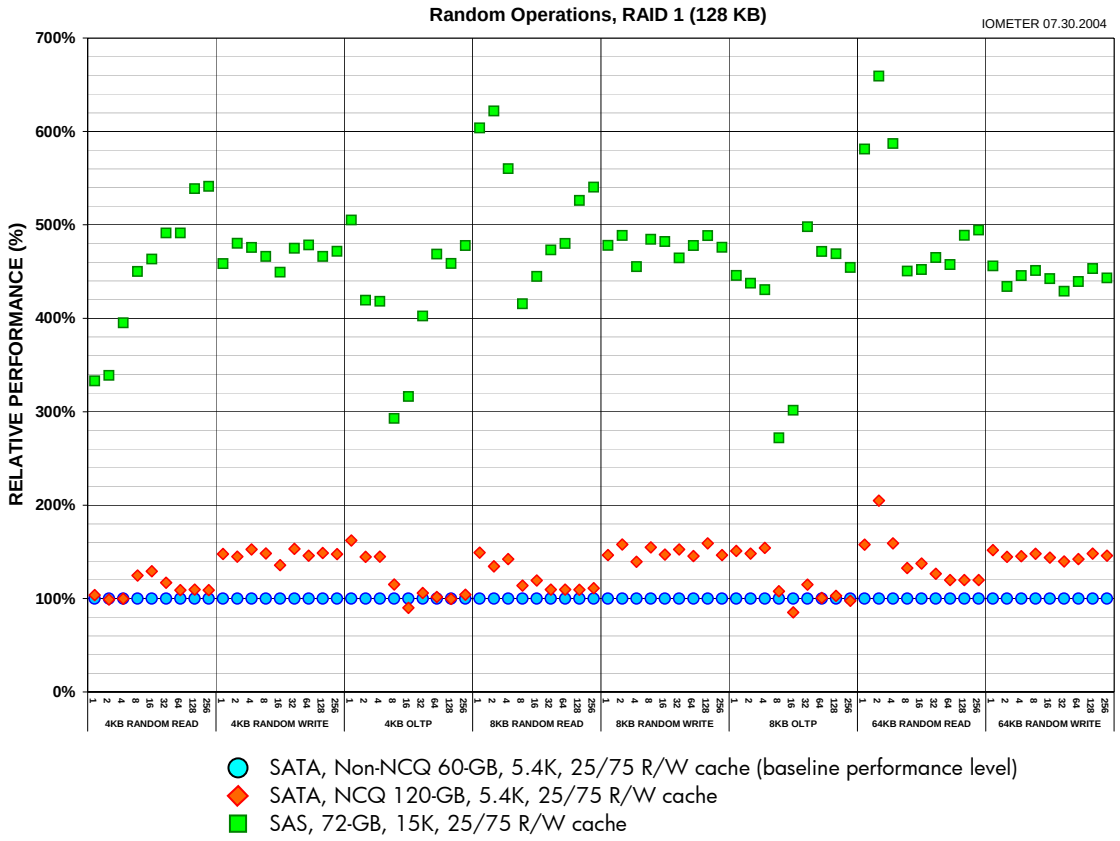

The HP P400 controller supports NCQ, so the NCQ SATA drives provide an increase in performance over non-NCQ SATA drives. The SAS arrays, however, are much faster than either of the SATA arrays and exhibit a substantial performance increase for most operations.

Figure 3 compares the throughput performance of the SATA and SAS drives in the same configuration, and illustrates the substantial improvement of SAS drives.

Figure 3. Throughput performance: random operations of non-NCQ SATA, NCQ SATA, and SAS 2-drive array with HP Smart Array P400 Controller, RAID 1

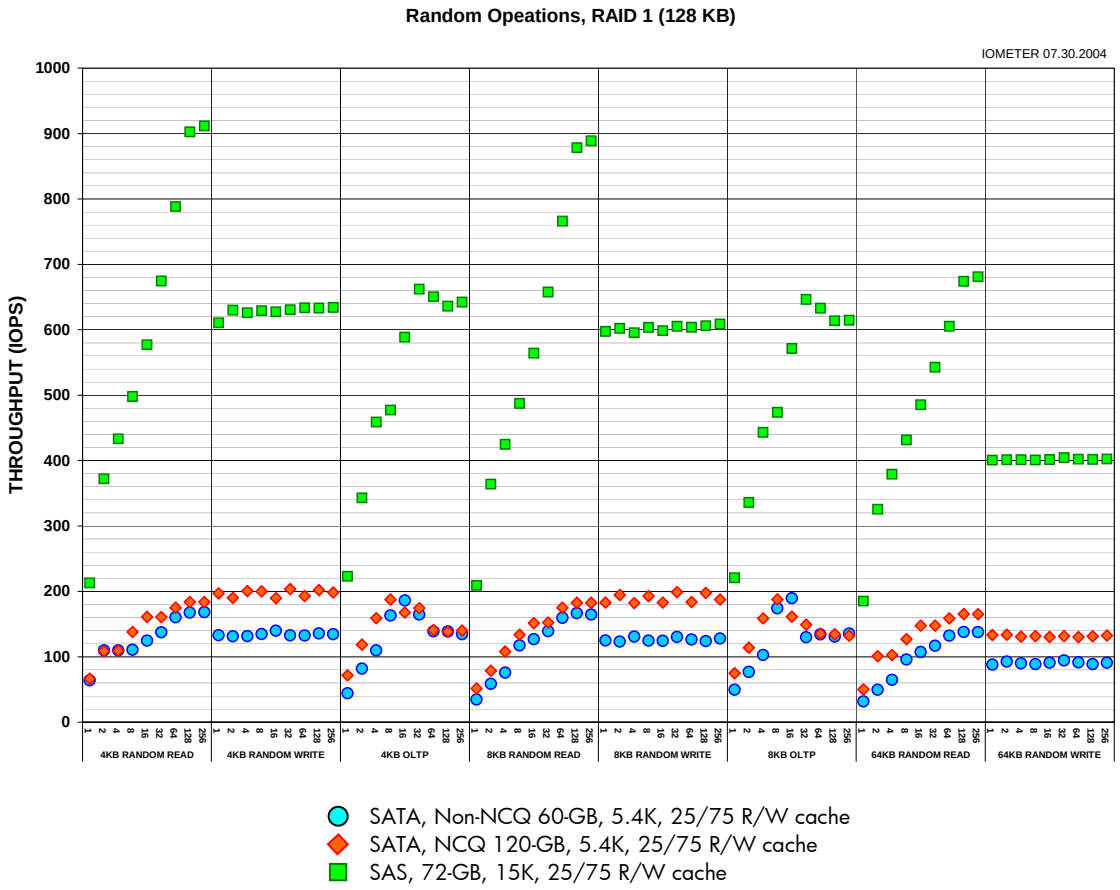

Performance comparisons between SAS drives

This section includes comparisons of SAS drives based on form factor and platter speed, along with differences in RAID strategies.

Performance comparisons based on form factor

Figures 4 through 8 compare the performance of SFF SAS drives and LFF SAS drives. The following configuration was used for testing the SAS drives:

- System: HP StorageWorks 70 Modular Smart Array
- SAS Controller: HP Smart Array P800, firmware version 4.06, 512 MB cache

Figure 4. Performance comparison: SFF to LFF 15K 72-GB SAS drives with HP Smart Array P800 Controller, RAID 0 (128K)

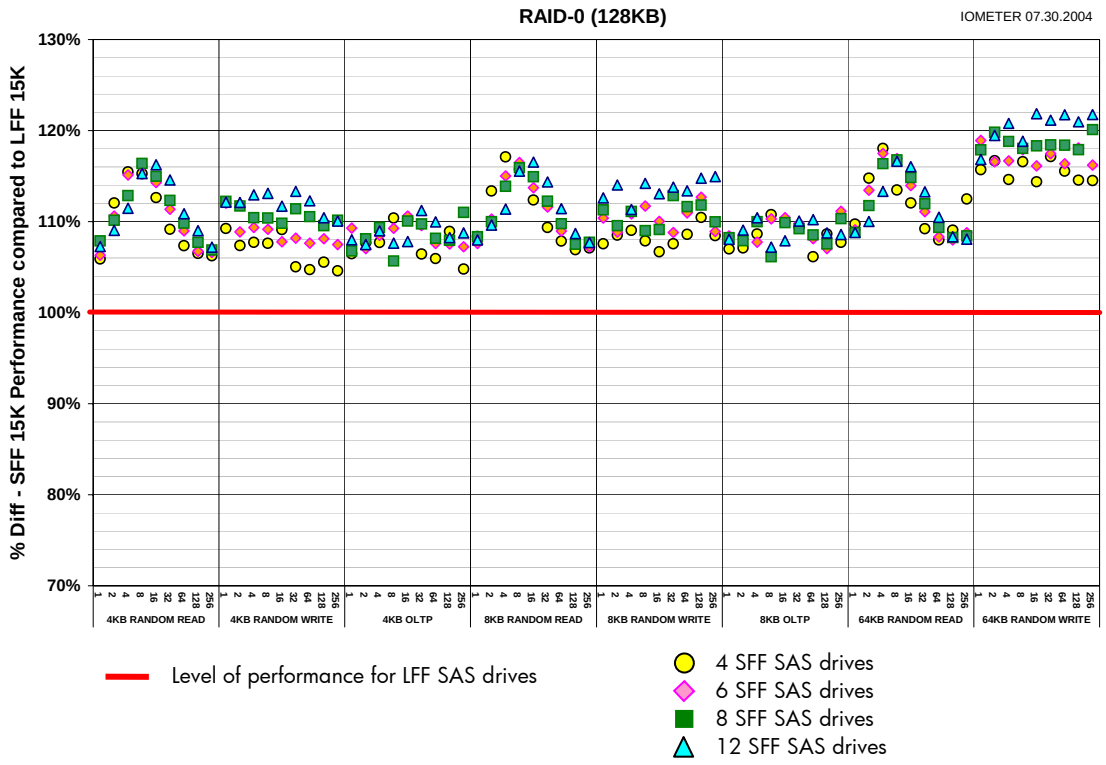

As indicated in Figure 4, SFF SAS drive arrays yield better performance in a RAID 0 solution than LFF SAS drive arrays. While providing maximum storage capacity and performance, RAID 0 does not provide any disk redundancy.

Figure 5. Performance comparison: SFF to LFF 15K 72-GB SAS drives with HP Smart Array P800 Controller, RAID 1 (128K)

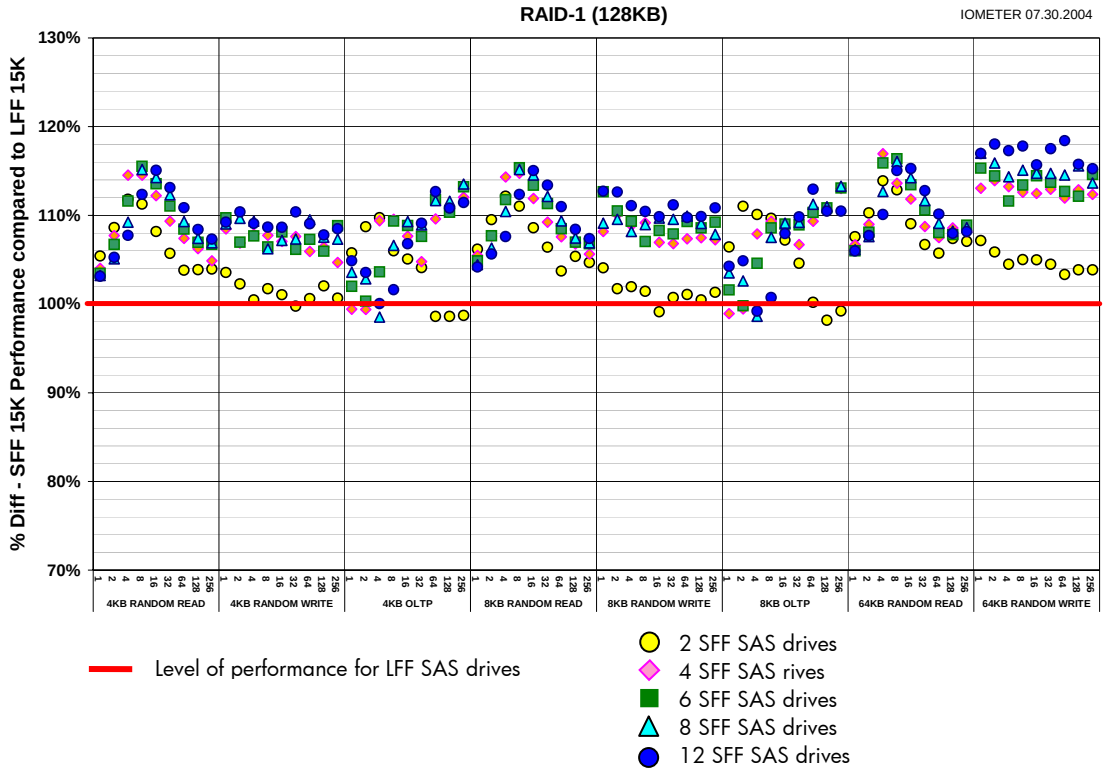

RAID 1 is a strategy for achieving redundancy rather than enhanced performance. Nevertheless, as indicated in Figure 5, SFF SAS drives provide better performance than LFF SAS drives for most RAID 1 configurations.

Figure 6 illustrates the IOPS performance of various array sizes of SFF SAS drives with the HP Smart Array P800 Controller using RAID 5 (64K).

Figure 6. RAID5 performance of SFF 15K 72-GB SAS drive arrays

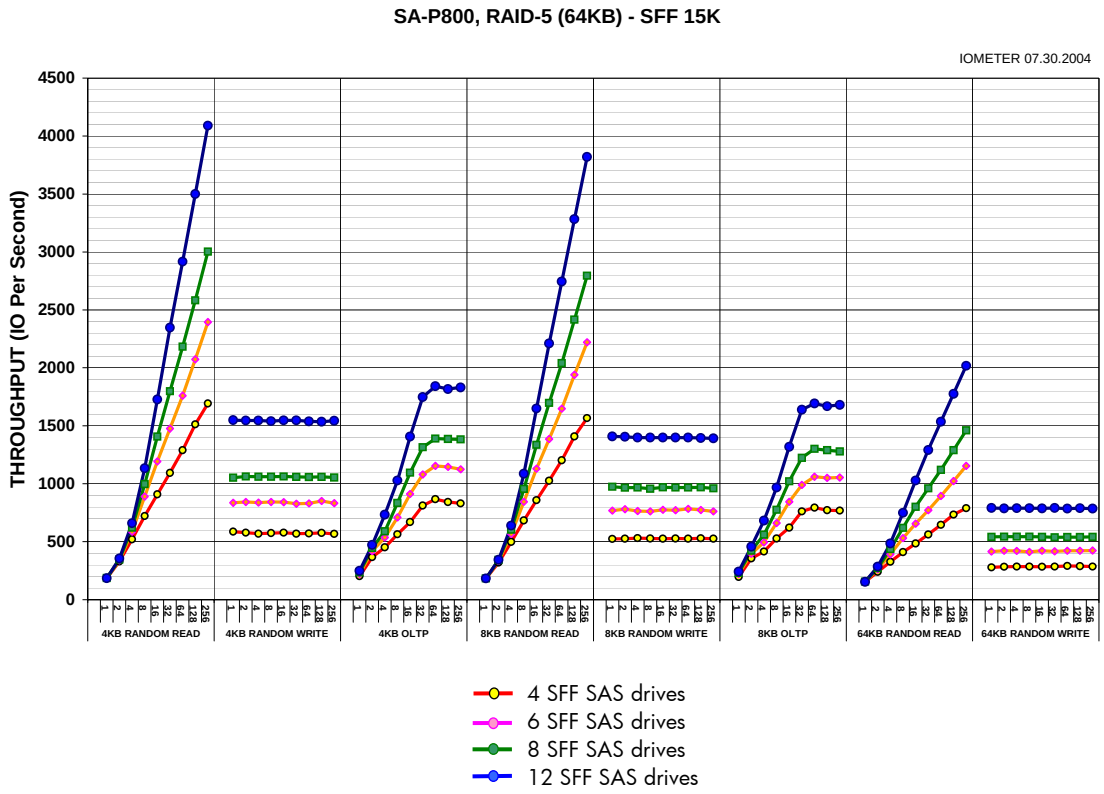

Performance comparisons based on platter speed

The charts in Figure 7 compare sequential read performance between 10K (rpm) and 15K SFF SAS drives operating as a 6-drive array with an HP Smart Array P800 controller using RAID 5. Charts 7A through 7E show the drive array performance for sequential reads per block size at a specific queue (Q) depth. Queue depth is the number of commands waiting to be serviced by the storage system. In workstations and small server systems queue depth rarely exceeds 1, but in an enterprise server environment a large number of users generating requests can result in a queue depth of tens or hundreds. Chart 7F shows random read IOPS performance.

Figure 7. Performance comparison: 10K versus 15K RPM, RAID 5 controller and 6 SFF SAS drives

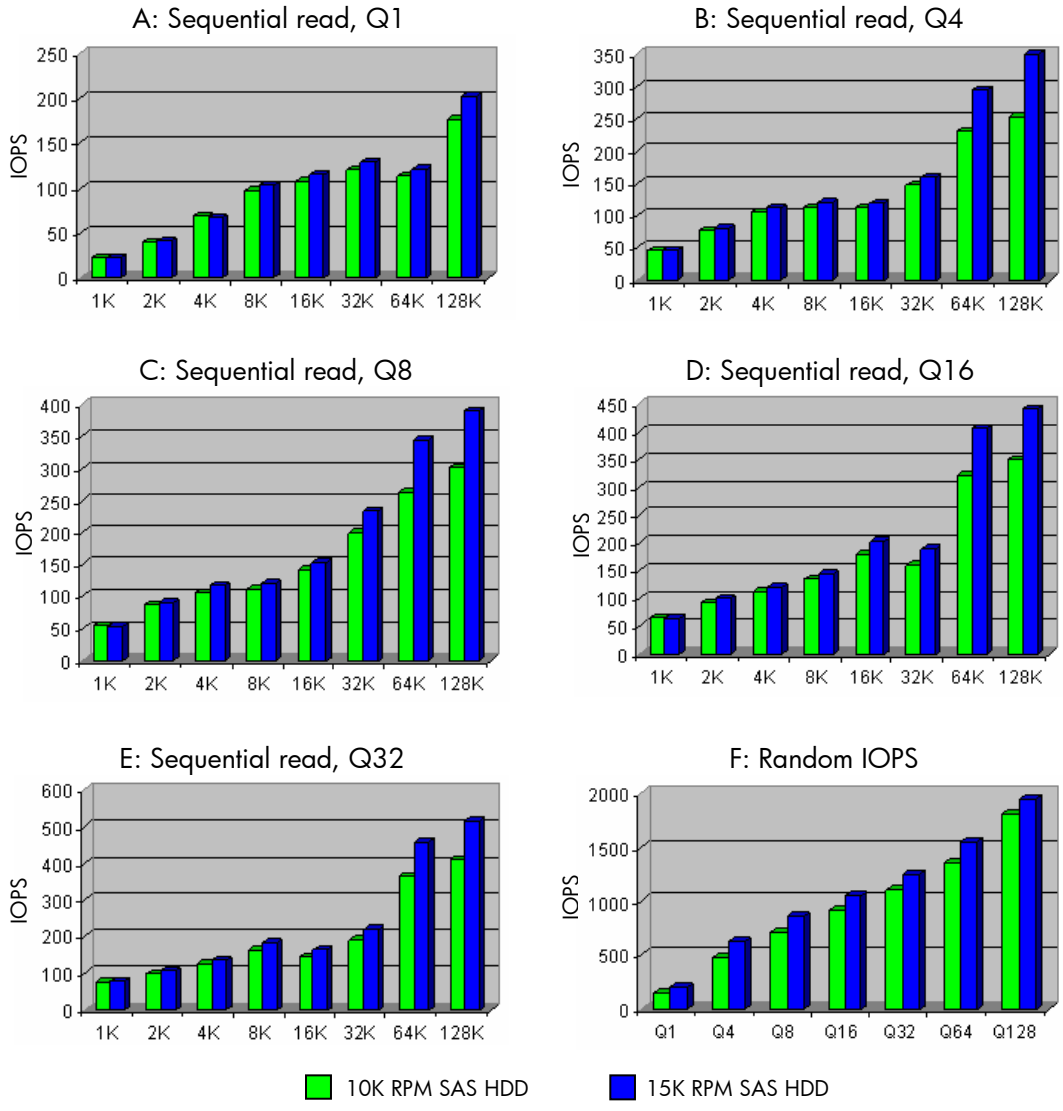

As indicated by the charts, 15K drives show a significant advantage where the requested data blocks are stripe size (64 KB) or larger. Although the IOPS rating can increase with greater queue depths, user interaction response of the system is best when the queue depth remains below 32.

Figure 8 compares the random read performance of 10K and 15K SFF SAS drives for arrays of three to eight drives, using RAID 5 with the block size set to 128 KB.

Figure 8. Random read performance: 10K versus 15K RPM SFF SAS drives, RAID 5

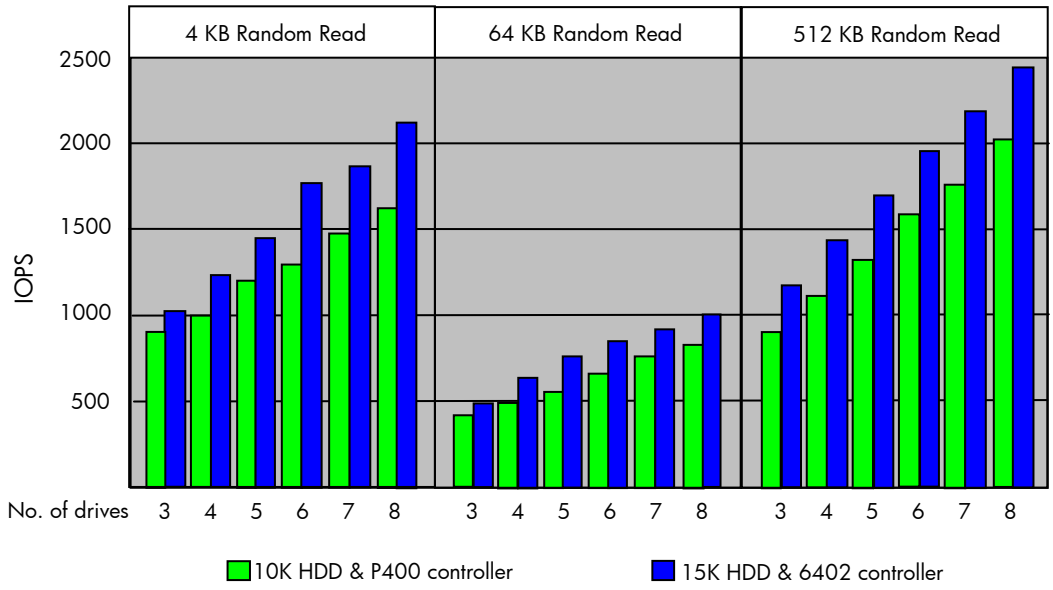

As indicated in Figure 8, 15K drive arrays consistently provide higher random read IOPS than 10K drive arrays. Also, overall performance scales with array size (number of drives). Note that performance is best when retrieving data blocks that are larger than the RAID 5 stripe size (128 KB in this case). For optimum performance, the stripe size should be configured for the specific application.

SAS drive specifications

Table 3 compares the seek time and power consumption of SFF (2.5-inch) SAS drives available from HP at time of publication.

NOTE:
Drive types and specifications are subject to change without notice. Refer to the following site for a listing and specifications of currently available SAS drives:
http://h18004.www1.hp.com/products/servers/proliantstorage/serial/sas/index.html

Table 3. Comparison of seek time and power consumption of HP SFF SAS drives

Capacity/RPM	36 GB/15K	72 GB/10K	72 GB/15K	146 GB/10K
Seek time:				
Single track	0.20 ms	0.20 ms	0.20 ms	0.60 ms
Average	3.0 ms	4.0 ms	3.0 ms	4.1 ms
Full stroke	7.0 ms	8.1 ms	7.0 ms	8.0 ms
Power consumption:				
Idle	5.74 watts	5.96 watts	6.11 watts	6.31 watts
Maximum	7.25 watts	8.42 watts	8.29 watts	8.73 watts

NOTE: All drives listed In Table 3 have a cache/buffer size of 16 MB and a transfer time of 3 Gb/s.

Conclusion

A variety of hard drive types are used in today's server industry. The right HDD type is the one that provides the best return on investment (ROI) for a particular server application. Small-to-medium businesses (SMBs) or environments with less-than-demanding applications can be well served with SCSI or SATA drives that offer economical service and good dollar-per-gigabyte value.

In the mission-critical, performance-oriented, enterprise server environment where hard drive duty cycle is 100 percent, SFF SAS drives excel in performance and reliability. Since SFF drives require only 70 percent of the space and half the power of 3.5-inch drives, higher drive densities per U space can be achieved *without* a significant increase in power consumption. Higher drive densities provide better overall performance, reliability, and lower operating costs.

For more information

For additional information, refer to the resources listed below.

Resource description	Web address
QuickSpecs for HP SAS drives	http://h18004.www1.hp.com/products/quickspecs/12244_na/12244_na.html
"Serial ATA technology" technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00301688/c00301688.pdf
"Serial-Attached SCSI technology" technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00302340/c00302340.pdf

Call to action

Send comments about this paper to TechCom@HP.com.

TC080502TB, May 2008