P6000 EVA with SAS disk drives—Ready for the Enterprise

Technical white paper

Table of contents

Overview

The new P6000 Enterprise Virtual Array (EVA) Systems support 6 Gb/s SAS disk drives. As customers may be more familiar with Fibre Channel (FC) drives for midrange enterprise storage, this white paper aims to describe the SAS technology and why it is appropriate to use with business critical applications.

This whitepaper also discusses differences in 2.5" SAS Small Form Factor (SFF) drives and 3.5" SAS Large Form Factor (LFF) drives and suggests some factors to consider when deciding which drives best match a customer's requirements.

SAS drives

Maturity

SAS was first introduced in 2004, with volume shipments beginning in 2005. Since then it has evolved with a number of changes. In 2008, the SAS-2, or SAS 2.0, specification was released by the T10 Committee. The new specification includes the following improvements over the initial release.

- Double the transfer rate, up to 6 Gb/s interface speeds
- Greater signal integrity, which also helps increase the cable interconnect length from 5 meters to 10 meters
- Spread Spectrum Clocking (SSC), which helps reduce EMI
- Mini-SAS connectors, which improve cabling and connector options

The 6 Gb/s SAS drives used in the P6000 EVA family are all SAS-2 drives.

Interface speeds

SAS disk drives have been available for some time with 3 Gb/s serial transfer rates, slower than the 4 Gb/s rates available with FC disk drives. In 2009, drive vendors introduced 6 Gb/s SAS transfer rates, surpassing the currently available 4 Gb/s FC disk transfer rates. So SAS drive interface speeds are faster than FC interface speeds. Faster interface speeds mean that more drives and/or larger block sizes can be supported more easily.

Functionality

The functionality of SAS-2.0 disks and FC disks is very similar. The same SCSI-2 command sets are implemented on both drive types. The controllers in the drives have the same basic architecture, processors, and memory. Both have dual I/O ports for redundant paths to the disk. The physical interfaces—the I/O chips and connectors—are different, but approximately 98% of the drive is same.

Workload utilization

Workload utilization,[1] referring to the portion of time a disk drive is reading or writing data (including seek activity), is an important factor in determining which disk drives to use for a particular application. Higher workload means that the disk drive has to do more reading and writing (and moving the read/write heads). This creates more stress on the disk drive as more mechanical work is being done. With more work being done, more heat is generated, which elevates the drive temperature and in turn decreases the life of the disk drive.

Enterprise 10k rpm and 15k rpm disk drives, both LFF and SFF, are designed for high workload utilization. They are capable of sustaining continuous read/write operations.

Midline drives are suitable for applications that do not require high workload utilization (i.e., not exceeding 40% workload), such as backup, archival, and image storage. For applications where a large amount of data needs to be stored, but the data is not accessed frequently and speed is not of high importance, midline SAS drives are very economical, while meeting the high reliability, but lower workload utilization requirements of the enterprise.

The reliability figures below assume that the drives are being used within their workload utilization specification range. Using midline drives beyond the workload utilization specification severely degrades their reliability of the midline drives.

SAS reliability

Drive vendors produce families of drives with the choice of several different interfaces on the same drive mechanism. The drive mechanism, with all of the moving parts—heads, actuator, rotating media, and spindle motor—is a major factor in the disk drive reliability specification. The interfaces/controllers of SAS and FC drives have about the same level of complexity—both have all solid state electronics. So, as long as disk drives have the same drive mechanism and such similar electronics, the SAS and the FC versions have the same reliability specification.

There are, however, reliability specification differences between drive types. Enterprise SAS drives, like enterprise FC drives, are designed and specified with high (100%) workloads and high reliability specifications, about 1.6 million hour Mean-Time-Between-Failures (MTBF), depending on the drive vendor. Midline SAS drives, on the other hand, like Fibre Attached Technology Adapted (FATA) disk drives, are not designed for high-workload environments and have a lower reliability specification than enterprise drives. Currently, the reliability specification for midline SAS (or FATA) drives is 1.2 million hours MTBF, when workload utilization does not exceed 40%.

Note that reliability differs from availability. Using best RAID redundancy practices, the availability of even a large number of drives can be very high.

Shipment history

HP began shipping SAS drives in server and storage products in 2006. Since 2007, HP has shipped about nine times more SAS drives than FC drives, demonstrating expertise in handling SAS drives.

With the introduction of the 6 Gb/s SAS interface speeds, with our extensive experience with SAS in server and entry storage products, and with the availability of new drive capacities in both SFF and LFF SAS drive sizes, HP is confident that SAS drives are ready for midrange enterprise arrays.

[1] Workload utilization has often been referred to as duty cycle. However, drive vendors are now referring to duty cycle as the power-on time. So, if a drive is powered on 100% of the time, but it is not accessed frequently, it can have a low workload, but a 100% duty cycle, using the drive vendor definition.

Midline SAS vs. SATA drives

While this paper is focused on SAS drives, a short discussion regarding SATA drives will be helpful. The P6000 EVA family does not support SATA drives, primarily because SATA drives do not offer the same level of redundancy as SAS or FC. For example, there is only one I/O port on the drive. And, while array controllers can make it appear that there are redundant paths to the drives, all I/O ultimately has to pass through the one I/O port on the drive, which is a potential single point of failure. In addition, SAS drive error-recovery and error-reporting uses SCSI commands which have more functionality than the ATA SMART commands used by SATA drives.

Tips on selecting SAS SFF and LFF drives

The P6000 EVA family offers a variety of choices for Enterprise SAS and midline SAS drives:

- Two physical sizes (SFF or LFF)
- A variety of speeds (10k and 15k for enterprise and 7.2k for midline)
- A variety of capacities (from 146 GB to 2 TB)

These choices result in different performances (depending on the speed), different workload capabilities (100% for enterprise and 40% for midline), different reliability specifications, and different prices. This section is intended to review the choice points and provide some general guidance on how to select the right disk drive(s) to meet customer requirements.

SAS drives available in SFF and LFF sizes

Physical size, power draw, and cooling are important factors that have driven drive vendors to develop SFF SAS drives, also described as 2.5-inch disks. LFF drives, also described as 3.5-inch drives, have been very popular in disk arrays, especially in higher 15k rpm speeds and also in high capacity and lower performance drives.

The following table compares the physical specifications of SFF and LFF drives (without drive carriers):

	SFF	LFF	Ratio
Height	14.8 mm (0.583 in)	26.11 mm (1.028 in)	.567 ~1/2
Width	70 mm (2.76 in)	101.85 mm (4.010 in)	.687~ 2/3
Length	100.5 mm (3.957 in)	146.99 mm (5.787 in)	.684 ~ 2/3
Weight*	227 grams (0.50 lbs)	694 grams (1.53 lbs) for 10k/15k; 610 grams (1.345 lbs) for 7.2k	.327 ~ 1/3
Frontal area (H x W)	1,036 mm^2 (1.61 in^2)	2,659.3 mm^2 (4.15 in^2)	.389 ~ 1/3
Volume (H x W x L)	104,118 mm^2 (6.37 in^2)	390,890.5 mm^2 (24.02 in^2)	.266 ~ 1/4

*Weight will vary by model of drive. These are representative examples.

The SFF drives are about 1/4th the volume and about 1/3rd the weight of an LFF drive.

Figure 1, contains pictures of a 600 GB 10k rpm SFF drive, the largest capacity SFF 10k drive currently available, and a 600 GB 15k rpm LFF drive, the largest capacity LFF 15k drive currently available, showing their relative sizes. The SFF drive has the same capacity as the LFF drive, but in a smaller package. However, the SFF drive has a 10k rpm rotational speed versus the 15k rpm rotational speed of the LFF drive.

Figure 1: LFF and SFF drive comparison

Capacity ranges for SFF and LFF drives

Drive capacities have increased from year to year for both SFF and LFF drives. SFF drives now match the capacities of LFF drives, but at different enterprise speeds (10k rpm for SFF; 15k rpm for LFF). In contrast, SFF drive capacities continue to lag LFF drive capacities in midline drive speeds (7.2k rpm).

The following table shows the drives available at the introduction of the P6000 EVA family.

P6000 EVA family disk drive list

Capacity and RPM	SFF	LFF
146 GB 15k	ü	
300 GB 10k	ü	
450 GB 10k	ü	
600 GB 10k	ü	
500 GB 7.2k Midline	ü	
300 GB 15k		ü
450 GB 15k		ü
600 GB 15k		ü
2 TB 7.2k Midline		ü

NOTE:
Additional drive offerings will be added to the P6000 EVA family over time, so this chart is subject to change.

There are several points to note regarding the table above:

- For 15k rpm drives, the maximum SFF drive capacity today is 146 GB, while the maximum LFF drive capacity today is 600 GB
- For 10k rpm drives, only available in SFF drives, the SFF 10k drives match the capacities of the LFF 15k drives, with a maximum of 600 GB, but only require half of the rack space
- For 7.2k rpm midline drives, the maximum SFF drive today is 500 GB; the maximum LFF drive today is 2 TB
- The 500 GB 7.2k SFF midline drive may be a good low performance mate to the 10k SFF enterprise drives, as they will fit in the same SFF drive enclosure
- The 2 TB 7.2k midline drive has the highest capacity per drive, >3x that of a 10k or 15k enterprise drive, and may be the drive of choice for low performance, large capacity configurations

The following table shows the maximum capacities by each drive type and capacity for the P6000 EVA family:

P6000 EVA family maximum capacities by drive types

Capacity and RPM	P6300	P6500
	250 SFF	450 SFF
146 GB 15k	36.5 TB	65.7 TB
300 GB 10k	75 TB	135 TB
450 GB 10k	112.5 TB	202.5 TB
600 GB 10k	150 TB	270 TB
500 GB 7.2k Midline	125 TB	225 TB
	120 LFF	240 LFF
300 GB 15k	36 TB	72 TB
450 GB 15k	54 TB	108 TB
600 GB 15k	72 TB	144 TB
2 TB 7.2k Midline	240 TB	480 TB

As the table above shows, there are differences in the maximum P6000 EVA capacity that can be achieved with each size, LFF or SFF, of disk:

- For enterprise drives, 10k rpm SFF drives can deliver twice the capacity as 15k rpm LFF drives, in the same rack space.
- The 2 TB 7.2k rpm LFF midline drives achieve the highest capacity points.

The following tables show the number of drives and drive enclosures required for each P6000 model for selected capacity points:

P6300

Drive capacity (GB)	Drive type	Max number of drives		Total capacity 25 (TB)	50 (TB)	75 (TB)	100 (TB)	125 (TB)	150 (TB)	175 (TB)	200 (TB)	225 (TB)
146	15k SFF	250	# drives	172								
			# enclos	7								
300	10k SFF	250	# drives	84	167	250						
			# enclos	4	7	10						
450	10k SFF	250	# drives	56	112	167	223					
			# enclos	3	5	7	9					
600	10k SFF	250	# drives	42	84	125	167	209	250			
			# enclos	2	4	5	7	9	10			
500	7.2k SFF MDL	250	# drives	50	100	150	200	250				
			# enclos	2	4	6	8	10				
300	15k LFF	120	# drives	84								
			# enclos	7								
450	15k LFF	120	# drives	56	112							
			# enclos	5	10							
600	15K LFF	120	# drives	42	84							
			# enclos	4	7							
2000	7.2k LFF MDL	120	# drives	13	25	38	50	63	75	88	100	113
			# enclos	2	3	4	5	6	7	8	9	10

P6500

Drive capacity (GB)	Drive type	Max number of drives		Total capacity 50 (TB)	100 (TB)	150 (TB)	200 (TB)	250 (TB)	300 (TB)	350 (TB)	400 (TB)	450 (TB)
146	15k SFF	450	# drives	343								
			# enclos	14								
300	10k SFF	450	# drives	167	334							
			# enclos	7	14							
450	10k SFF	450	# drives	112	223	334	445					
			# enclos	5	9	14	18					
600	10k SFF	450	# drives	84	167	250	334	417				
			# enclos	4	7	10	14	17				
500	7.2k SFF MDL	450	# drives	100	200	300	400					
			# enclos	4	8	12	16					
300	15k LFF	250	# drives	167								
			# enclos	14								
450	15K LFF	250	# drives	112	223							
			# enclos	10	19							
600	15k LFF	250	# drives	84	167							
			# enclos	7	14							
2000	7.2K LFF MDL	250	# drives	25	50	75	100	125	150	175	200	225
			# enclos	3	5	7	9	11	13	15	17	19

Notice that the maximum limits of each model are applied. For example, for the P6300, 50 TB cannot be achieved using 146 GB 15k SFF drives. It would require 343 drives, which is above the maximum number of 250 SFF drives that can be supported by the P6300. Similarly, 50 TB cannot be achieved using 300 GB 15k LFF drives in the P6300. It would require 167 drives, which is above the maximum number of 120 LFF drives that can be supported by the P6300.

Notice also that the 10k SFF drives can reach higher total capacities than the 15k LFF drives. For example, for the P6300, 250 of the 600 GB 10k SFF drives can provide 150 TB. But only 120 of the 600 GB 15k LFF drives can be supported, for a capacity of only 72 TB.

Performance of SFF and LFF drives

Performance numbers for SFF and LFF drives, on a drive-per-drive basis, are very similar for the same drive type. That is, a 10k rpm SFF drive versus a 10k LFF rpm drive or a 15k SFF drive versus a 15k LFF drive have similar drive performance specification numbers.

However, if different drive capacities are being configured to meet the same capacity requirement, then more small capacity drives can outperform fewer large capacity drives, particularly for random I/O operations. For example, (100) 300 GB 10k SFF drives will outperform (50) 600 GB 10k SFF drives.

Rotational speed is also important when estimating performance. A 15k drive (SFF or LFF) will outperform a 10k drive (SFF or LFF). But today, the 300/400/600 GB capacities are available in 10k rpm for SFF and 15k rpm for LFF. (There is a 146 GB 15k SFF that is available for high performance at lower capacity configurations.) So, the drive decision will likely come down to selecting some number of 10k SFF drives versus some different number of 15k LFF drives.

The following table shows drive performance parameters for the SAS drives supported in the P6000 EVA family:

SAS performance parameters

	SFF					LFF			
	146 GB 15k	300 GB 10k	450 GB 10k	600 GB 10k	500 GB 7.2k	300 GB 15k	450 GB 15k	600 GB 15k	2 TB 7.2k
Average rotational latency (ms)	2.0	3.0	3.0	3.0	4.2	2.0	2.0	2.0	4.16
Seek Time— Avg. Read (ms) Avg. Write (ms)	 2.9 3.3	 3.6 4.2	 3.6 4.2	 3.6 4.2	 8.5 9.5	 3.4 3.9	 3.4 3.9	 3.4 3.9	 8.5 9.5
Track-to-Track— Read (ms) Write (ms)	 0.2 0.4	 0.2 0.3	 0.2 0.3	 0.2 0.3	 0.8 1.0	 0.2 0.4	 0.2 0.4	 0.2 0.4	 0.8 1.0

Note:
These are representative performance parameters. Because multiple vendors may be used for each drive type and because drive improvements and changes are made over time, the values may vary, but not significantly.

Some things to note in the table above:

- The performance parameters are the same for drives with the same rpm and form factor. For example the 300 GB, 450 GB and 600 GB 10k drives all perform identically.
- 15k rpm drives' average rotational latency time is less than 10k rpm drives' rotational latency time, which is less than 7.2k drives' rotational latency time. That translates to faster performance for the higher speed 15k rpm drives.
- For smaller capacities with high performance needs, the 146 GB 15k SFF drive may be a good fit.
- For larger capacities with high performance needs, the 300/450/600 GB 15k LFF drives may be the right choice.
- For medium performance needs, the 300/450/600 GB 10k SFF drives may be the right choice.
- Seek times, both read and write, vary by drive family. For example, the 300/450/600 GB LFF 15k drives' seek times are slightly faster than the 300/450/600 GB SFF 10k drives.
  - (While SFF drives have smaller distances to move the data heads because of their smaller size, the positioner mechanisms in the SFF drives are not as strong as the positioner mechanisms of the LFF drives. So, even though the LFF drives have longer distances to move the data heads, they are able to move faster than the SFF drives.)
- Track-to-track (seek) times for reads are the same between the SFF and LFF drives, except for the 7.2k midline drive, which is slower, The track-to-track times for writes are slightly slower for the 15k drives and even slower for the 7.2k midline drives.
- The 7.2k midline drives are considerably slower in all parameters, compared to either the 10k or 15k drives.

While the drive performance parameters help explain the performance differences for the various drives, it is really more important to know how the drives perform in the P6000 EVAs under different application parameters. See the P6000 EVA family performance white paper, at the P6000 web page, at www.hp.com/go/P6000, under Resource Library, for more information on array performance.

Power draw of SFF and LFF drives

The power draw of the different drives is also an important factor to consider when evaluating the different drive sources.

The following table shows the power figures for the various drives under idle and heavy load conditions.

SAS power parameters

	SFF					LFF			
	146 GB 15k	300 GB 10k	450 GB 10k	600 GB 10k	500 GB 7.2k	300 GB 15k	450 GB 15k	600 GB 15k	2 TB 7.2k
Power idle (Watts)	7.0	3.4	3.8	3.8	3.5	9.0	10.3	12.0	8.1
Power heavy load (Watts)	8.5	5.4	5.7	6.0	4.1	12.5	14.2	16.0	9.8
Power per TB—Idle (Watts/TB)	47.9	11.3	8.4	6.3	7.0	30.0	22.9	20.0	4.0
Power per TB—Heavy load (Watts/TB)	58.2	18.0	12.7	10.0	8.2	41.7	31.5	26.7	4.9

Note:
These are representative power numbers. Because multiple vendors may be used for each drive type and because drive improvements and changes are made over time, the values may vary, but not significantly.

Some things to note in the table above:

- The 10k SFF drives draw less than half of the power than the 15k LFF drives under heavy loads. At the same time, they can provide the same capacity per drive as the 15k LFF drives.
- The midline drives draw less power than the 10k SFF drives or 15k SFF drives and generally have larger capacities.
- On a watt/TB basis (idle and heavy load), the midline drives are the most efficient, followed by the 10k drives, then the 15k drives.
- The 2 TB 7.2k midline drive has the lowest watt/TB rating of all the drives, 4.9 watts/TB for heavy load and 4.0 watts for idle. Note, however, that 100% workload utilization (heavy load) is not recommended.

To help compare the power draw of different P6000 EVA configurations with the different drives, including drive enclosures and array controllers, under idle or heavy load conditions, a power calculator tool is available at: http://www.hp.com/servers/powercalculator.

The following tables show the power required for each P6000 model for selected capacity points, using the power calculator referenced above:

P6300

Drive Capacity (GB)	Drive Type	Max Number of Drives		Total Power (Watts) / Capacity								
				25 (TB)	50 (TB)	75 (TB)	100 (TB)	125 (TB)	150 (TB)	175 (TB)	200 (TB)	225 (TB)
146	15k SFF	250	Power	2110								
			# drives	172								
300	10k SFF	250	Power	840	1504	2159						
			# drives	84	167	250						
450	10k SFF	250	Power	632	1095	1540	2030					
			# drives	56	112	167	223					
600	10k SFF	250	Power	493	891	1214	1608	1939	2332			
			# drives	42	84	125	167	209	250			
500	7.2k SFF MDL	250	Power	443	790	1137	1484	1831				
			# drives	50	100	150	200	250				
300	15k LFF	120	Power	1661								
			# drives	84								
450	15k LFF	120	Power	1264	2432							
			# drives	56	112							
600	15k LFF	120	Power	1066	1984							
			# drives	42	84							
2000	7.2k LFF MDL	120	Power	363	550	753	942	1145	1337	1534	1725	1919
			# drives	13	25	38	50	63	75	88	100	113

@230V and heavy workload

P6500

Drive Capacity (GB)	Drive Type	Max Number of Drives		Total Power (Watts) / Capacity								
				50 (TB)	100 (TB)	150 (TB)	200 (TB)	250 (TB)	300 (TB)	350 (TB)	400 (TB)	450 (TB)
146	15k SFF	450	Power	4129								
			# drives	343								
300	10k SFF	450	Power	1504	2914							
			# drives	167	334							
450	10k SFF	450	Power	1095	1942	2914	3901					
			# drives	112	223	334	445					
600	10k SFF	450	Power	891	1608	2332	3124	3857				
			# drives	84	167	250	334	417				
500	7.2k SFF MDL	450	Power	790	790	1137	1484					
			# drives	100	200	300	400					
300	15k LFF	240	Power	3209								
			# drives	167								
450	15k LFF	240	Power	2434	4683							
			# drives	112	223							
600	15k LFF	240	Power	1984	3853							
			# drives	84	167							
2000	7.2k LFF MDL	240	Power	550	942	1337	1726	2117	2500	2889	3282	3680
			# drives	25	50	75	100	125	150	175	200	225

@230V and heavy workload

Summary

HP believes SAS drives are ready for Enterprise storage arrays. SAS performance, functionality and reliability equal that of FC drives. SAS drives are proven in servers and business class arrays with unit shipments of SAS drives in the past three years far exceeding the shipments of FC drives.

When choosing SAS drives, all of the user's requirements need to be considered: their capacity needs, including future growth, their workload utilization requirements, their performance needs, their price range, and their power target. Because there are more drive choices available with the SAS drives, customer requirements are better met. While it is tempting to make a drive choice or recommendation on capacity and/or price alone, it is best to invest a little time to consider all of the other factors to achieve optimum customer satisfaction.

Some points to remember when selecting drives:

- 10k SFF drives have the same capacities today as 15k LFF drives, but take half the rack space.
- 10k SFF drives draw half the power than comparable 15k LFF drives.
- 15k rpm drives have the best performance per drive.
- For smaller capacities with high performance needs, the 146 GB 15k SFF drive may be a good fit.
- For larger capacities with high performance needs, the 300/450/600 GB 15k drives may be the right choice.
- For medium performance needs, the 300/450/600 GB 10k SFF drives may be the right choice.
- 7.2k midline drives should only be used when the workload utilization requirements are light, 40% or less.
- The 500 GB 7.2k SFF midline drive may be a good low performance mate to the 10k SFF enterprise drives, as they will fit in the same SFF drive enclosure.
- The 2 TB 7.2k midline drive has the highest capacity per drive, >3x that of a 10k or 15k enterprise drive, and may be the low performance and low workload drive of choice for large capacity configurations.
- The 2 TB 7.2k midline drive has the lowest watt/TB rating of all the drives, 4.9 watts/TB for heavy load (note that heavy load is 100%, but midline drive workload utilization should be limited to 40%) and 4.0 watts/TB for idle.

For more information

The SCSI Trade Association (STA), a member-run industry association established to support and promote SCSI technology: http://www.scsita.org/

T10, a Technical Committee of the International Committee on Information Technology Standards (INCITS, pronounced "insights"): http://www.t10.org/

"Enterprise-Optimized 6Gb/s SAS Rivals Fibre Channel Performance and Scalability at Lower Cost," Seagate Technology Paper: http://www.seagate.com/docs/pdf/whitepaper/tp_enterprise_otimized_6gbs.pdf

"6Gb/s Serial-Attached SCSI (SAS)," Hitachi white paper: http://www.hitachigst.com/tech/techlib.nsf/techdocs/EBDE428D64C507BA8625757F0080CDAF/$file/SAS6Gbps_WP_final.pdf

To know more about HP P6000 Enterprise Virtual Array, visit: www.hp.com/go/P6000

4AA3-2517ENW, Created June 2011