FUJITSU

WHITE PAPER
FUJITSU PRIMERGY SERVER
SINGLE DISK PERFORMANCE

This technical documentation is aimed at people who are involved in selecting hard disk drives and solid state drives for operating disk subsystems on PRIMERGY servers. The document is intended to help find at an early stage the appropriate data media, – usually for RAID configurations, – for a planned solution. The various data media and their performance data are presented.

Version

1.0a
2011-08-31

Content

© Fujitsu Technology Solutions 2011

Document history

Version 1.0a

© Fujitsu Technology Solutions 2011

Introduction

Hard disks are a security factor as well as critical performance components in the server environment. It is thus important to bundle the performance of such components via intelligent organization so that they do not cause a system bottleneck. They should simultaneously compensate for any failure of an individual component. Methods exist for arranging several hard disks in an array so that any hard disk failure can be compensated. This is known as a "Redundant Array of Independent Disks" or in short RAID. Special RAID controllers are normally used.

The PRIMERGY servers are available in a wide range of internal configuration versions with different RAID controller and hard disk configurations. The "Modular RAID" concept that is offered as a standard for all servers in the PRIMERGY family consists of a modular controller family and standardized management via the Fujitsu RAID Manager software known as "ServerView RAID". The comprehensive offer of RAID solutions enables the user to select the appropriate controller for a particular application scenario. The performance of the disk subsystem is defined by the controller, the selected hard disks and the features of the RAID level.

Several documents have been created in the PRIMERGY white paper series which illustrate all aspects of "Modular RAID" regarding performance:

- We recommend - as a comprehensive introduction to disk I/O performance - the white paper "Basics of Disk I/O Performance".
- This document "Single Disk Performance" presents the hard disks currently available for PRIMERGY and their performance in different application scenarios.
- The document "RAID Controller Performance" covers all the RAID controllers currently on offer for PRIMERGY and their performance.
- The document "RAID Performance" compares the individual RAID levels and provides information about performance and the optimal configuration.

In the past the terms "Hard Disk" and "Hard Disk Drive" (HDD) were used for a hard magnetic-coated, rotating, digital, non-volatile storage medium that could be directly addressed. Technical development has now seen new "hard disk" versions introduced as storage media; they use the same interface to the server and as far as the server is concerned are handled accordingly as hard disks. A typical example is a flash memory (Solid State Drive = SSD), which is an electronic storage medium without any movable parts; but which is still referred to as a hard disk in everyday talk. This document uses the term "hard disk" throughout. If necessary, for example, to illustrate construction features, the terms "SSD" and "HDD" will be used.

This document specifies hard disk capacities on a basis of 10 (1 TB = 1 billion bytes), while all other capacities as well as throughput are specified on a basis of 2 (1 MB/s = 2^{20} Byte/s).

Basics of hard disks

As hard disks are usually the only directly addressable, non-volatile storage media in a server and they have a very high access time in comparison with server components, such as processors or main memory, special attention should be paid regarding their selection for sizing and configuration of disk subsystems. The manner in which hard disks are used in servers is basically different when compared to those in a desktop or notebook. Server hard disks are usually faced with a much greater workload due to the higher number of users. Data transfer requests can arrive suddenly at any time and must be processed as quickly as possible. High transaction rates (IO/s) and fast response times (latency) are essential. Furthermore, data logging and backup/restore functions require a high data throughput rate (MB/s). At the same time, server hard disks usually run over a longer period, around the clock in an extreme situation. Server hard disks are normally used in a RAID array (Redundant Array of Independent Disks) in order to ensure as high data security levels as possible despite such workloads and possible failures of single hard disks. A further intention of using RAID arrays is of achieving throughput rates higher than those attained by just an individual hard disk. Various RAID levels plus their security and performance-specific features are described in detail in the document "RAID Performance". The lifespan requirements for such an individual hard disk are still much greater than for a disk used in the desktop environment. And last of all because the disk subsystem of a server can be very large, its hard disks often play a major role in the power consumption of a server.

The high level of quality requirements facing server hard disks can be described in more detail, as they depend on the specific task that a server has to execute. Here are some examples:

- **File servers** are mainly used for file storage. Not only do they usually have internal hard disks, but are also connected with external hard disk arrays. Fail-safety is a number 1 priority for such a file server. The performance level for file server hard disks depends on the tasks: hard disks are probably not accessed quite so much if the server is just used for archiving purposes. In such situations, speed does not play a significant role. The behavior is different if the file server is used as a central part of daily business. In this case, a high level of hard disk performance is essential.
- **Database servers** house databases which provide data within a company network or even via the Internet. From the user's viewpoint access to a database server is usually via an application server or web server. As database servers usually have to handle many parallel accesses, data security and fast response times are also essential.
- **Mail servers** are responsible for sending, receiving and storing e-mails and are thus of great significance in the communications world. Such a scenario requires reliable hard disks. Fast response times are also required.
- **Application servers** are used to execute application programs in a network. These include for example web applications, authentication services; in short: all services running on a server which are essential for a transaction between client and server. This can be limited to a restricted number of users or openly available for all Internet users. Such usage requires hard disks which are the fastest and most secure.
- **Streaming servers** provide the user with multimedia files. Usually in 7×24 operations. Both fail-safety and high throughput rates are the main factors when selecting the hard disks.
- **Virtualization servers** provide virtual machines in which different operating systems and applications work. The fastest and most reliable hard disks are essential in such a scenario.

Indeed, any combination of the above different server activities could run on one piece of server hardware. In this case, the hard disks should be grouped according to tasks and if possible should not be used in common for all tasks.

 © Fujitsu Technology Solutions 2011

Quality & reliability

The reliability requirements are an important aspect when selecting the right hard disk.

The MBTF value of a hard disk type is used to determine the approximate lifespan of a hard disk. **MTBF** stands for **M**ean **T**ime **B**etween **F**ailures and specifies in hours the average operating time between two failures. As this is simply a statistic, there is no guarantee that the MTBF value is achieved in a particular situation.

Dual porting

SAS hard disks have two ports and can thus simultaneously be operated via two different controllers. This redundancy also enables hard disk access if one of the two controllers should fail.

SATA hard disks require a port selector for the connection to two controllers. No simultaneous access is then possible. The second controller can only access the SATA hard disk if the first controller is not active.

Transfer protocols

SATA and SAS protocols are used during data exchange at the interface between controllers and hard disks. The transfer rates specified below for both protocols must not be confused with hard disk throughputs. The actual throughput of a hard disk is lower than the throughput, for which your interface is designed. Downwards compatibility applies for the interfaces named below. This means that for example a hard disk with a SATA 6.0 Gbit/s interface can be operated on a server with a SATA 3.0 Gbit/s interface, but it can be only with the lower transfer speed. To ensure that an interface does not become a bottleneck, its transfer rate must be designed to be higher than the maximum possible throughput of a hard disk.

SATA interface

SATA stands for **S**erial **A**dvanced **T**echnology **A**ttachment and is the successor of the ATA interface for driving storage media. Whereas 16 bit can be transferred in parallel with ATA, transfer is serial with SATA. There are in the meantime three generations of the SATA standard. The theoretically usable transfer rate, specified in the following list in MB/s, is about 80% of the transfer rate specified for the respective SATA generation.

- SATA 1.5 Gbit/s enables a throughput of up to 143 MB/s.
- SATA 3.0 Gbit/s offers a maximum throughput of 286 MB/s. Hot-plug, staggered spinup, NCQ and longer connecting cables have also been introduced here.
- The maximum throughput is 572 MB/s for SATA 6.0 Gbit/s. Improved NCQ (Streaming Command) and energy-saving features have also been implemented and special connectors (LIF, ZIF) introduced.

SAS interface

SAS stands for **S**erial **A**ttached **S**CSI and is the successor to the SCSI interface for driving storage media. Whereas the transfer was 16 bit in parallel with the last SCSI generation "Ultra-320 SCSI", transfer is serial with SAS. There are in the meantime two generations of the SAS standard. The theoretically usable transfer rate is about 80% of the transfer rate specified for the respective SAS generation.

- SAS 3 Gbit/s offers a maximum throughput of 286 MB/s.
- SAS 6 Gbit/s enables a throughput of up to 572 MB/s.
- SAS 12 Gbit/s for a maximum throughput of 1144 MB/s is being planned.

Hard disks with SATA interfaces can also be operated on controllers with SAS interfaces.

Solid state drives vs. conventional hard disks

Solid state drives (SSDs) have also been offered for some time now. Contrary to conventional hard disks, SSDs do not have any moving parts, but are based on flash memory. Outer design and interface correspond to those of conventional hard disks. In comparison with conventional hard disks, SSDs stand out for very low access times, mechanical robustness, low noise and very low energy consumption, but still have considerably lower capacity today and a higher price.

SSDs store information in so-called NAND flash memory **s**ingle **l**evel **c**ells (SLC) or in **m**ulti **l**evel **c**ells (MLC). In comparison with MLC-based drives, the capacity of SLC-based SSDs is lower. In contrast, they are considerably more durable and offer a clearly higher write performance. Although SSDs are generally indefinitely readable, they are not indefinitely writable. However, the controller that exists in the SSD ensures that write operations are equally distributed over all flash memory cells thus ensuring even wear and tear. As with conventional hard disks, it is also possible to monitor SSDs via S.M.A.R.T. (**S**elf-**M**onitoring, **A**nalysis and **R**eporting **T**echnology). The lifespan of an SSD can be compared with that of conventional hard disks of the Enterprise class (see "Drive classes").

SSDs generally provide a higher throughput than conventional hard disks. However, this greatly depends on when an SSD is accessed during its lifecycle. The first write accesses to an SSD take place at the highest possible speed, because only empty cells are written. After a period of time the throughput settles at a lower level, because for each write operation it is first necessary to delete any old contents. Only these are throughputs that are really relevant to practice. The SSDs were used in this condition in our measurements. This settling effect does not exist for the read performance. However, an SSD still achieves a much higher throughput than conventional hard disks - at least with random access - at this lower throughput level.

© Fujitsu Technology Solutions 2011

Drive classes

For PRIMERGY servers Fujitsu offers a large number of hard disks with different performance and availability features for all conceivable scenarios - from the small department server that is only needed during working hours up to the high-availability database server. To facilitate making a selection from this wide range Fujitsu has introduced three hard disk classes:

- Hard disks of the **Economic class (ECO)** are the most cost-effective variant. These disks are chosen if performance and fail-safety do not have to meet the highest requirements. They should only be used in non-critical application areas, in which high transaction rates are not to be expected and in which 7×24-hour operation is not necessary. ECO hard disks have a SATA interface and work at a rotational speed of 5400 or 7200 rpm.
- Hard disks of the **Business-Critical class (BC)** are suited for 7×24-hour operation and thus meet the highest requirements for fail-safety. Due to their high capacities the lowest costs per GB are incurred here. BC hard disks either have a SATA interface and work at a rotational speed of 7200 rpm or have a SAS interface and work at a rotational speed of 10000 rpm.
- Hard disks of the **Enterprise class (EP)** offer the best performance and simultaneously meet the highest security requirements. They were specially developed for scenarios, in which the highest requirements for throughput also have to be met under high loads. The EP hard disks include ones with a SAS interface and a rotational speed of 15000 rpm as well as SATA SSDs.

The following overview is a snapshot taken at the time the white paper was written. New PRIMERGY models and hard disk types can be added. A current compilation of the available hard disks can be found in the data sheet or in the configurator of the appropriate PRIMERGY server.

The following hard disks are on offer in the economic class for the current PRIMERGY servers:

Drive class	Data medium type	Interface	Form factor	rpm	Capacity [GB]	PRIMERGY servers
ECO	HDD	SATA 3 Gbit/s	3.5″	7200	250	All tower servers except TX120, TX140
ECO	HDD	SATA 3 Gbit/s	2.5″	5400	500	All tower servers except TX100, TX120 CX122 BX920 SX940, SX960
ECO	HDD	SATA 3 Gbit/s	2.5″	5400	320	All tower servers except TX100 All scale-out servers RX100 S6 BX620, BX920, BX922 All storage blades
ECO	HDD	SATA 3 Gbit/s	2.5″	5400	160	All tower servers except TX100 All scale-out servers RX100 S6 BX620, BX920, BX922 All storage blades

All tower servers: *PRIMERGY TX100 S2, TX120 S3, TX140 S1, TX150 S6, TX200 S6, TX300 S6*
All rack servers: *PRIMERGY RX100 S6, RX100 S7, RX200 S6, RX300 S6, RX600 S5, RX600 S6, RX900 S1, RX900 S2*
All scale-out servers: *PRIMERGY CX120 S1, CX122 S1*
All server blades: *PRIMERGY BX620 S6, BX920 S2, BX922 S2, BX924 S2, BX960 S1*
All storage blades: *PRIMERGY SX650, SX940 S1, SX960 S1*

The following hard disks are on offer in the business-critical class for the current PRIMERGY servers:

Drive class	Data medium type	Interface	Form factor	rpm	Capacity [GB]	PRIMERGY servers
BC	HDD	SAS 6 Gbit/s	2.5″	10000	600	All tower servers except TX100 All rack servers BX920 SX940, SX960
BC	HDD	SAS 6 Gbit/s	2.5″	10000	450	All tower servers except TX100 All rack servers BX920 SX940, SX960
BC	HDD	SAS 6 Gbit/s	2.5″	10000	300	All tower servers except TX100 All rack servers BX920 SX940, SX960
BC	HDD	SAS 6 Gbit/s	2.5″	10000	146	All tower servers except TX100 All rack servers BX920 SX940, SX960
BC	HDD	SATA 3 Gbit/s	3.5″	7200	2000	All tower servers except TX100 RX100, RX300
BC	HDD	SATA 3 Gbit/s	3.5″	7200	1000	All tower servers except TX120 RX100, RX300
BC	HDD	SATA 3 Gbit/s	3.5″	7200	500	All tower servers except TX120 RX100, RX300
BC	HDD	SATA 3 Gbit/s	3.5″	7200	250	All tower servers except TX120 RX100, RX300
BC	HDD	SATA 3 Gbit/s	2.5″	7200	1000	All tower servers except TX100 RX100, RX200, RX300, RX600 S5 BX620, BX920 All storage blades
BC	HDD	SATA 3 Gbit/s	2.5″	7200	500	All tower servers except TX100 All scale-out servers RX100, RX200 RX300, RX600 S5 BX620, BX920 All storage blades
BC	HDD	SATA 3 Gbit/s	2.5″	7200	160	All tower servers except TX100 All scale-out servers RX100, RX200 RX300, RX600 S5 BX620, BX920 All storage blades

All tower servers: *PRIMERGY TX100 S2, TX120 S3, TX140 S1, TX150 S6, TX200 S6, TX300 S6*
All rack servers: *PRIMERGY RX100 S6, RX100 S7, RX200 S6, RX300 S6, RX600 S5, RX600 S6, RX900 S1, RX900 S2*
All scale-out servers: *PRIMERGY CX120 S1, CX122 S1*
All server blades: *PRIMERGY BX620 S6, BX920 S2, BX922 S2, BX924 S2, BX960 S1*
All storage blades: *PRIMERGY SX650, SX940 S1, SX960 S1*

© Fujitsu Technology Solutions 2011

The following hard disks are on offer in the enterprise class for the current PRIMERGY servers:

Drive class	Data medium type	Interface	Form factor	rpm	Capacity [GB]	PRIMERGY servers
EP	SSD	SATA 3 Gbit/s	2.5″	-	64	All tower servers except TX100 All rack servers All server blades except BX620 SX940, SX960
EP	SSD	SATA 3 Gbit/s	2.5″	-	32	All tower servers except TX100, TX140 All rack servers All server blades except BX620 SX940, SX960
EP	HDD	SAS 6 Gbit/s	3.5″	15000	600	All tower servers except TX100, TX120 RX100, RX300
EP	HDD	SAS 6 Gbit/s	3.5″	15000	450	All tower servers except TX100, TX120 RX100, RX300
EP	HDD	SAS 6 Gbit/s	3.5″	15000	300	All tower servers except TX100, TX120 RX100, RX300
EP	HDD	SAS 3 Gbit/s	3.5″	15000	146	All tower servers except TX100, TX120 RX100, RX300
EP	HDD	SAS 6 Gbit/s	2.5″	15000	146	All tower servers except TX100 All rack servers BX920 SX940, SX960
EP	HDD	SAS 6 Gbit/s	2.5″	15000	73	All tower servers except TX100 All rack servers BX920 SX940, SX960

All tower servers: *PRIMERGY TX100 S2, TX120 S3, TX140 S1, TX150 S6, TX200 S6, TX300 S6*
All rack servers: *PRIMERGY RX100 S6, RX100 S7, RX200 S6, RX300 S6, RX600 S5, RX600 S6, RX900 S1, RX900 S2*
All scale-out servers: *PRIMERGY CX120 S1, CX122 S1*
All server blades: *PRIMERGY BX620 S6, BX920 S2, BX922 S2, BX924 S2, BX960 S1*
All storage blades: *PRIMERGY SX650, SX940 S1, SX960 S1*

Disk I/O performance measurements

At Fujitsu disk I/O performance measurements are performed for all the hard disks released for PRIMERGY servers in the PRIMERGY Performance Lab. In contrast to the application benchmarks, it is in general not the performance of an entire server including disk subsystem that is tested, but merely the performance of the disk subsystem, in other words the storage media in themselves as well as their controllers. In the single disk measurements presented in this document the disk subsystem has been sized in such a way that server components such as the processor or main memory do not constitute a bottleneck.

The Iometer measuring tool was used for the disk I/O performance measurements. Iometer makes it possible to reproduce the behavior of real applications with regard to accesses to disk subsystems. In this way, the most varied application scenarios can be reproduced - be it with sequential read or write, random read or write and also combinations thereof when using different block sizes as well as a variably adjustable number of simultaneous accesses. This method permits the performance of various hard disks with certain access patterns to be compared.

As a result Iometer provides the following important indicators:

- Throughput [MB/s] Throughput in megabytes per second
- Transactions [IO/s] Transaction rate in I/O operations per second
- Latency [ms] Average response time in ms

Data throughput has generally established itself as the usual measurement variable for sequential load profiles, whereas the measurement variable "transaction rate" is normally used for the random load profiles with their small block sizes, as is also the case in this document.

Load profiles

As standard, a series of load profiles, which are characterized by the following mass storage accesses, is used for disk I/O performance measurements:

Access	Type of access		Transfer request size [kB] (block size)	Number of outstanding I/Os
	read	write		
sequential	100%	0%	1, 4, 8, 64, 128, 512, 1024	1, 3, 8, 16, 32, 64, 128, 256, 512
sequential	0%	100%	1, 4, 8, 64, 128, 512, 1024	1, 3, 8, 16, 32, 64, 128, 256, 512
random	100%	0%	1, 4, 8, 64, 256, 1024	1, 3, 8, 16, 32, 64, 128, 256, 512
random	0%	100%	1, 4, 8, 64, 256, 1024	1, 3, 8, 16, 32, 64, 128, 256, 512
random	67%	33%	1, 4, 8, 16, 32, 64, 128	1, 3, 8, 16, 32, 64, 128, 256, 512
random	50%	50%	64	1, 3, 8, 16, 32, 64, 128, 256, 512

Some of these load profiles can be assigned with typical applications:

Standard load profile	Access	Type of access		Block size [kB]	Application
		read	write		
File copy	random	50%	50%	64	Copying of files
File server	random	67%	33%	64	File server
Database	random	67%	33%	8	Database (data transfer) Mail server
Streaming	sequential	100%	0%	64	Database (log file), Data backup; Video streaming (partial)
Restore	sequential	0%	100%	64	Restoring of files

Apart from a number of measurement results, which clarify the influence of the block size on hard disk performance (see section "Essential parameters", subsection "Block size"), all the measurement results presented in this document can be allocated to these five practice-relevant standard load profiles.

 © Fujitsu Technology Solutions 2011

Measurement environment

All the measurement results presented in this document were determined in a standardized measurement environment with the following hardware and software components.

Components	Details						
Operating system	Windows Server 2008 Enterprise Edition						
File system	NTFS						
Measuring tool	Iometer 2006.07.27						
Measurement data	32 GB measurement file						
Controller	LSI MegaRAID SAS 1068						
Drives	ECO	HDD	SATA	3 Gbit/s	2.5″	5400 rpm	160 GB
	ECO	HDD	SATA	3 Gbit/s	2.5″	5400 rpm	320 GB
	ECO	HDD	SATA	3 Gbit/s	2.5″	5400 rpm	500 GB
	ECO	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	250 GB
	BC	HDD	SATA	3 Gbit/s	2.5″	7200 rpm	160 GB
	BC	HDD	SATA	3 Gbit/s	2.5″	7200 rpm	500 GB
	BC	HDD	SATA	3 Gbit/s	2.5″	7200 rpm	1 TB
	BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	250 GB
	BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	500 GB
	BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	1 TB
	BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	2 TB
	BC	HDD	SAS	6 Gbit/s	2.5″	10000 rpm	146 GB
	BC	HDD	SAS	6 Gbit/s	2.5″	10000 rpm	300 GB
	BC	HDD	SAS	6 Gbit/s	2.5″	10000 rpm	450 GB
	BC	HDD	SAS	6 Gbit/s	2.5″	10000 rpm	600 GB
	EP	HDD	SAS	6 Gbit/s	2.5″	15000 rpm	73 GB
	EP	HDD	SAS	6 Gbit/s	2.5″	15000 rpm	146 GB
	EP	HDD	SAS	6 Gbit/s	3.5″	15000 rpm	146 GB
	EP	HDD	SAS	6 Gbit/s	3.5″	15000 rpm	300 GB
	EP	HDD	SAS	6 Gbit/s	3.5″	15000 rpm	450 GB
	EP	HDD	SAS	6 Gbit/s	3.5″	15000 rpm	600 GB
	EP	SSD	SATA	3 Gbit/s	2.5″		32 GB
	EP	SSD	SATA	3 Gbit/s	2.5″		64 GB
RAID levels	JBOD						

Essential parameters

To understand the performance data presented in the following sections the influence that the variation of two important parameters has is dealt with here: block size and load intensity.

Block size

Data transfer always takes place block-by-block when accessing a disk subsystem. The size of the transferred data blocks depends on features of the operating system and/or the application and cannot be influenced by the user. The throughput rates that can be achieved therefore depend very much on the access pattern that an application generates on the disk subsystem.

The following table shows examples of typical access patterns of various applications:

Application	Access pattern
Operating system	random, 40% read, 60% write, blocks ≥ 4 kB
File copy (SMB)	random, 50% read, 50% write, 64 kB blocks
File server (SMB)	random, 67% read, 33% write, 64 kB blocks
Mail server	random, 67% read, 33% write, 8 kB blocks
Database (transaction processing)	random, 67% read, 33% write, 8 kB blocks
Web server	random, 100% write, 64 kB blocks
Database (log file)	sequential, 100% write, 64 kB blocks
Video streaming	sequential, 100% read, blocks ≥ 64 kB
Backup	sequential, 100% read, blocks ≥ 64 kB
Restore	sequential, 100% write, blocks ≥ 64 kB

As the implemented block size is an application feature that cannot be modified, comparing the hard disk performance when varying the block size amounts to comparing the I/O behavior of different applications (file server, database server, etc.) or even a specific application with an exotic or even completely irrelevant access pattern. In comparison with the relevant access patterns, others may exist with which a higher hard disk throughput can be achieved. However, this is always at the expense of poorer response times. Therefore, this type of measurement data is of theoretical interest only and is not practice-relevant.

Below we look at the influence of the block size on the throughputs, transactions and latencies of a SATA-HDD, a SAS-HDD and a SATA-SSD. To this end, measurement results were selected which show this influence with random access and a read/write ratio of 67% to 33%, as well as with sequential read and sequential write access. The load intensity (outstanding I/Os) was kept the same in all cases.

 © Fujitsu Technology Solutions 2011

Random access, 67% read, 33% write, 2 outstanding IOs

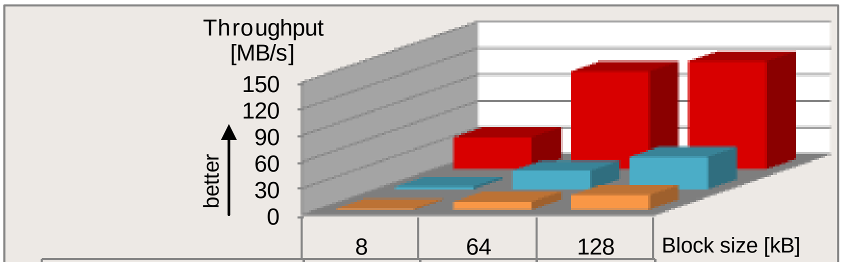

	8	64	128
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	1.26	9.18	16.56
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	2.98	21.23	36.91
EP SSD SATA 3 Gbit/s 2.5" 64 GB	36.19	110.65	121.97

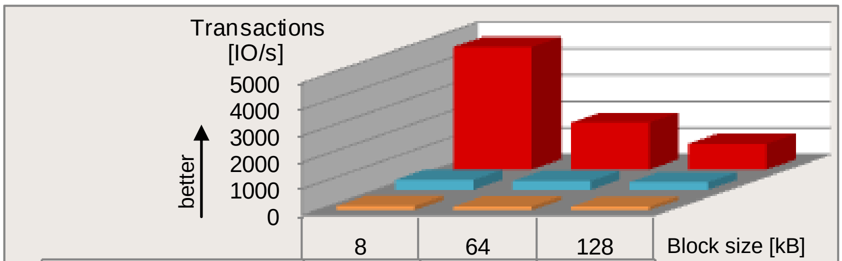

	8	64	128
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	161	147	132
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	382	340	295
EP SSD SATA 3 Gbit/s 2.5" 64 GB	4633	1770	976

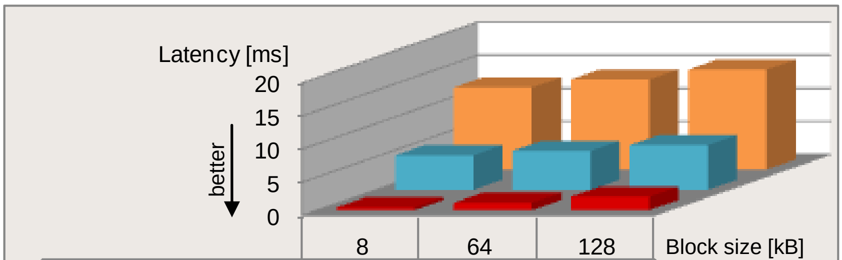

	8	64	128
EP SSD SATA 3 Gbit/s 2.5" 64 GB	0.4	1.1	2.0
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	5.2	5.9	6.8
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	12.4	13.6	15.1

A reduction in the transaction rates can be seen (diagram 2) for random access with 67% read share and increasing block size. However, in comparison with the increase in block size the decrease is so minor that this results in a significantly growing throughput (diagram 1). But at the same time higher average response times must also be expected with higher block sizes (diagram 3). The measurement results with a block size of 8 kB correspond to those of typical transaction processing in databases, whereas the measurement results with a block size of 64 kB reflect that of a typical file server. As can be seen in the right data column in the top diagrams, applications that used even larger blocks would be able to generate even higher throughputs, whereby this would of course be at the expense of further sinking transaction rates and further increasing response times.

Sequential access, 100% read, 2 outstanding IOs

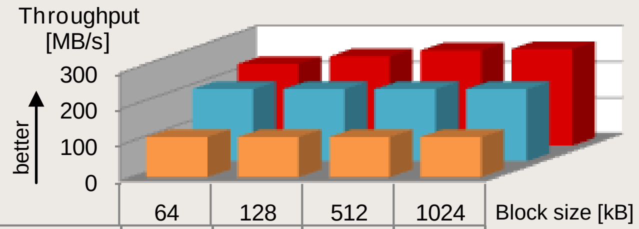

	64	128	512	1024
BC HDD SATA 3 Gbit/s 2.5” 7200 rpm 500 GB	111.16	111.13	111.15	111.10
EP HDD SAS 6 Gbit/s 2.5” 15000 rpm 146 GB	200.15	200.15	200.20	200.17
EP SSD SATA 3 Gbit/s 2.5" 64 GB	226.40	247.50	264.62	268.63

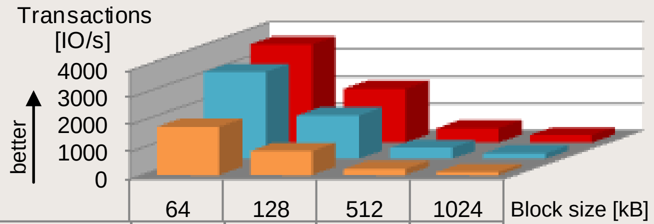

	64	128	512	1024
BC HDD SATA 3 Gbit/s 2.5” 7200 rpm 500 GB	1779	889	222	111
EP HDD SAS 6 Gbit/s 2.5” 15000 rpm 146 GB	3202	1601	400	200
EP SSD SATA 3 Gbit/s 2.5" 64 GB	3622	1980	529	269

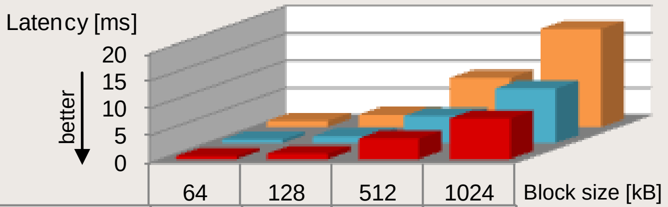

	64	128	512	1024
EP SSD SATA 3 Gbit/s 2.5" 64 GB	0.6	1.0	3.8	7.4
EP HDD SAS 6 Gbit/s 2.5” 15000 rpm 146 GB	0.6	1.2	5.0	10.0
BC HDD SATA 3 Gbit/s 2.5” 7200 rpm 500 GB	1.1	2.2	9.0	18.0

The throughput of SATA and SAS hard disks with sequential read access does not depend on the block size. Thus, twice as high average response times and halved I/O rates are recorded with double the block size. This is not the case with SATA-SSDs, because with block sizes of more than 64 kB they still offer minor increases in throughput. As only throughput - and not the transaction rates and response times - plays a role during backup, large blocks are used here, whereas smaller block sizes are generally used during streaming.

 © Fujitsu Technology Solutions 2011

Sequential access, 100% read, 2 outstanding IOs

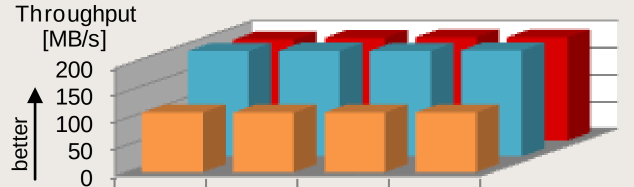

	64	128	512	1024
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	110.72	110.65	110.76	110.85
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	195.32	196.09	195.94	196.83
EP SSD SATA 3 Gbit/s 2.5" 64 GB	187.30	190.81	191.61	191.64

Block size [kB]

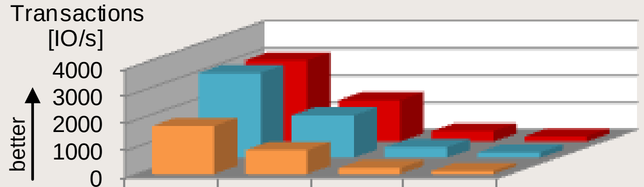

	64	128	512	1024
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	1772	885	222	111
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	3125	1569	392	197
EP SSD SATA 3 Gbit/s 2.5" 64 GB	2997	1526	383	192

Block size [kB]

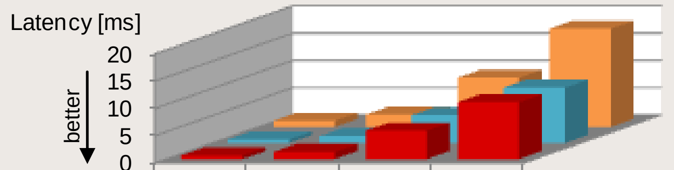

	64	128	512	1024
EP SSD SATA 3 Gbit/s 2.5" 64 GB	0.7	1.3	5.2	10.4
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	0.6	1.3	5.1	10.2
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	1.1	2.3	9.0	18.0

Block size [kB]

The throughput of all hard disks does not depend on block size with sequential write access. Thus, twice as high average response times and halved transaction rates are recorded with double the block size. As long as response times are of no interest, it does not matter whether an application runs with a block size of 64 kB or even larger blocks.

© Fujitsu Technology Solutions 2011

Load intensity

Numerous clients normally access a server simultaneously. It is also possible for a client to place several requests with one server without waiting for responses in the meantime. Both can result in simultaneous accesses to a single controller or even to a single data medium. Most controllers and data media therefore have queuing mechanisms, which when processing several parallel requests ensure under certain circumstances a higher performance than when processing fewer parallel or even only single requests. This, however, is always done at the price of higher response times. If many parallel requests only entail an increase in response time, but no longer in throughput, the disk subsystem is overloaded. If the load is increased further, throughput generally remains constant and does not slump. If faced with an application scenario in which the response behavior of the server cannot be disregarded, a choice must be made between optimizing the throughput and optimizing the response behavior. The server must then be sized and configured in such a way that it can manage the processing of parallel requests in accordance with individual requirements.

Later on this document the diagrams for load profiles with random access show the maximum transaction rates, whereas the diagrams for load profiles with sequential access report the maximum throughput achieved. With regard to throughputs or transaction rates with the highest and second highest load intensity there is no longer any difference worth mentioning for any of the hard disks (less than 5%) in all the diagrams.

For the standard load profiles "File server", "Database", "File copy", "Streaming" and "Restore" the following ten diagrams show the performance of a SATA-HDD, a SAS-HDD and a SATA-SSD depending on load intensity (outstanding IO) in an exemplary way on the basis of the following hard disks:

BC	HDD	SATA	3 Gbit/s	2.5″	7200 rpm	500 GB
EP	HDD	SAS	6 Gbit/s	2.5″	15000 rpm	146 GB
EP	SSD	SATA	3 Gbit/s	2.5″		64 GB

Standard load profile "Database" (random access, 67% read, 33% write, 8 kB block size)

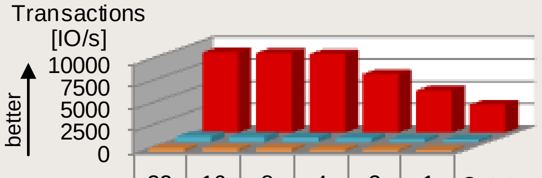

Outstanding IO	32	16	8	4	2	1
BC HDD SATA 3 Gbit/s 2.5” 7200 rpm 500 GB	312	316	312	267	226	188
EP HDD SAS 6 Gbit/s 2.5” 15000 rpm 146 GB	632	624	557	470	382	284
EP SSD SATA 3 Gbit/s 2.5" 64 GB	9049	8998	8856	6603	4752	3179

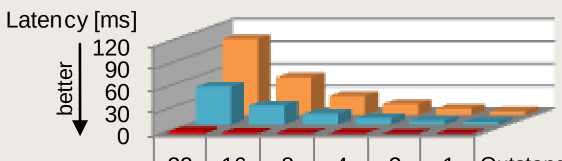

Outstanding IO	32	16	8	4	2	1
EP SSD SATA 3 Gbit/s 2.5" 64 GB	3.5	1.8	0.9	0.6	0.4	0.3
EP HDD SAS 6 Gbit/s 2.5” 15000 rpm 146 GB	50.6	25.6	14.4	8.5	5.2	3.5
BC HDD SATA 3 Gbit/s 2.5” 7200 rpm 500 GB	102.6	50.7	25.7	15.0	8.9	5.3

© Fujitsu Technology Solutions 2011

Standard load profile "File server" (random access, 67% read, 33% write, 64 kB block size)

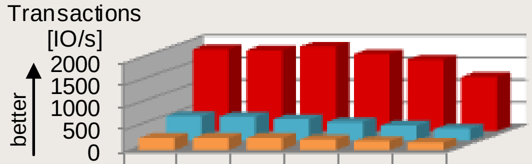

Outstanding IO	32	16	8	4	2	1
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	273	273	273	240	198	171
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	552	546	483	419	340	267
EP SSD SATA 3 Gbit/s 2.5" 64 GB	1846	1825	1908	1741	1614	1222

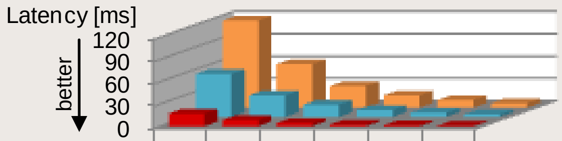

Outstanding IO	32	16	8	4	2	1
EP SSD SATA 3 Gbit/s 2.5" 64 GB	17.3	8.8	4.2	2.3	1.2	0.8
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	58.0	29.3	16.6	9.6	5.9	3.7
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	117.1	58.5	29.3	16.6	10.1	5.9

Standard load profile "File copy" (random access, 50% read, 50% write, 64 kB block size)

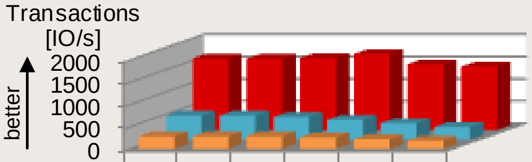

Outstanding IO	32	16	8	4	2	1
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	284	284	289	258	220	189
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	541	543	508	443	369	286
EP SSD SATA 3 Gbit/s 2.5" 64 GB	1614	1606	1627	1711	1486	1431

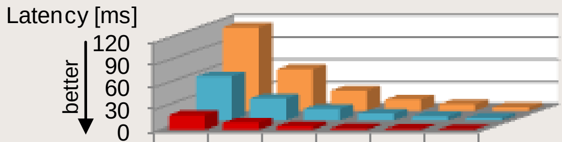

Outstanding IO	32	16	8	4	2	1
EP SSD SATA 3 Gbit/s 2.5" 64 GB	19.8	10.0	4.9	2.3	1.3	0.7
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	59.2	29.4	15.7	9.0	5.4	3.5
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	112.8	56.3	27.7	15.5	9.1	5.3

© Fujitsu Technology Solutions 2011

With a low load the SAS-HDD offers 50% higher transaction rates for the three standard load profiles with random access than the SATA-HDD. With higher loads the SAS-HDD provides about twice the performance of the SATA-HDD. If the load on the SAS-HDD is twice as high as that on the SATA-HDD, the response times are approximately equal.

In comparison with the SAS-HDD, the SATA-SSD offers many times the transaction rates with significantly lower response times. It can be loaded twice as much as a SAS-HDD without supplying poorer response times than the latter. Therefore, in a RAID array the performance of SAS-HDDs can be achieved with considerably fewer SATA-SSDs, unless capacity requirements oppose this.

Standard load profile "Streaming" (sequential access, 100% read, 64 kB block size)

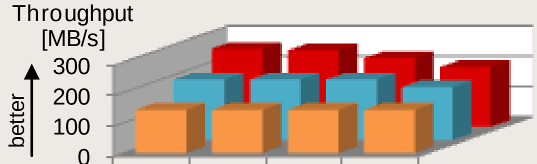

Outstanding IO	8	4	2	1
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	143.15	143.15	143.11	143.07
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	200.15	200.15	200.15	175.15
EP SSD SATA 3 Gbit/s 2.5" 64 GB	257.20	252.24	226.40	197.53

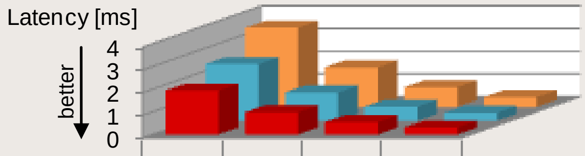

Outstanding IO	8	4	2	1
EP SSD SATA 3 Gbit/s 2.5" 64 GB	1.9	1.0	0.6	0.3
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	2.5	1.2	0.6	0.4
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	3.5	1.7	0.9	0.4

With streaming the throughput of the SAS-HDD is about 40% higher than that of the SATA-HDD. Compared with the SAS-HDD, the SATA-SSD offers a higher throughput of about 30%. The response times are reduced by about 20% - 30% in each case.

Standard load profile "Restore" (sequential access, 100% write, 64 kB block size)

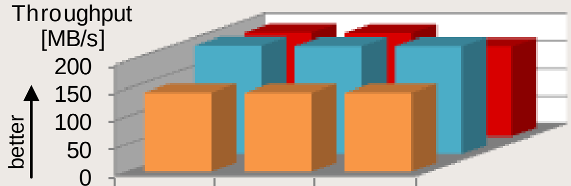

Outstanding IO	4	2	1
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	142.78	142.57	142.53
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	195.99	195.32	195.04
EP SSD SATA 3 Gbit/s 2.5" 64 GB	188.00	187.30	163.12

© Fujitsu Technology Solutions 2011

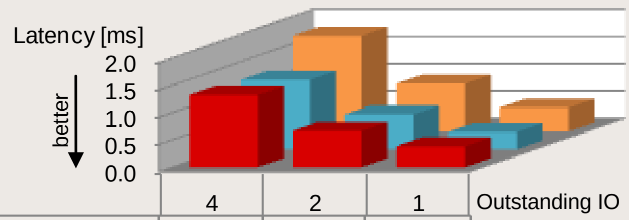

Outstanding IO	4	2	1
EP SSD SATA 3 Gbit/s 2.5" 64 GB	1.3	0.7	0.4
EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB	1.3	0.6	0.3
BC HDD SATA 3 Gbit/s 2.5" 7200 rpm 500 GB	1.8	0.9	0.4

With restore the throughput of the SAS-HDD is just under 40% higher than that of the SATA-HDD, the response times are just under 30% lower. In comparison with the SAS-HDD, the throughput and response times of the SATA-SSD are slightly worse.

Performance data of all hard disks

Below is an overview of the performance data of all of the hard disks examined, sorted according to drive classes.

Hard disk of the ECO class

The following ten diagrams show performance values (transactions, throughput rates and response times) for the standard load profiles "File server", "Database", "File copy", "Streaming" and "Restore" with increasing loads through the number of parallel requests (outstanding IO). The following actual hard disks are considered

ECO HDD SATA 3 Gbit/s 2.5″ 5400 rpm 160 GB

ECO HDD SATA 3 Gbit/s 2.5″ 5400 rpm 320 GB

ECO HDD SATA 3 Gbit/s 2.5″ 5400 rpm 500 GB

ECO HDD SATA 3 Gbit/s 3.5″ 7200 rpm 250 GB

Standard load profile "Database" (random access, 67% read, 33% write, 8 kB block size)

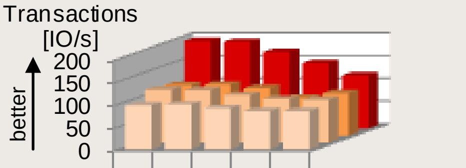

Outstanding IO	16	8	4	2	1
2.5" 5400 rpm 160 GB	97	99	91	85	85
2.5" 5400 rpm 320 GB	117	116	105	97	95
2.5" 5400 rpm 500 GB	111	112	105	83	93
3.5" 7200 rpm 250 GB	194	192	169	145	118

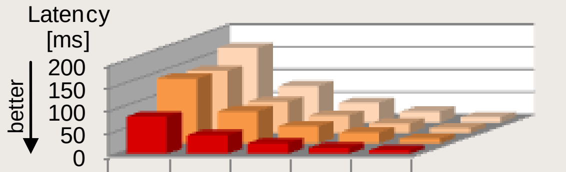

Outstanding IO	16	8	4	2	1
3.5" 7200 rpm 250 GB	82.7	41.7	23.6	13.8	8.5
2.5" 5400 rpm 500 GB	143.8	71.1	38.2	24.0	10.7
2.5" 5400 rpm 320 GB	137.1	69.1	38.2	20.7	10.6
2.5" 5400 rpm 160 GB	164.8	80.6	43.9	23.6	11.8

Standard load profile "File server" (random access, 67% read, 33% write, 64 kB block size)

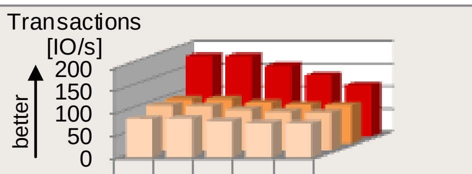

Outstanding IO	16	8	4	2	1
2.5" 5400 rpm 160 GB	85	86	80	76	75
2.5" 5400 rpm 320 GB	100	98	90	85	84
2.5" 5400 rpm 500 GB	98	98	90	86	85
3.5" 7200 rpm 250 GB	178	178	156	135	113

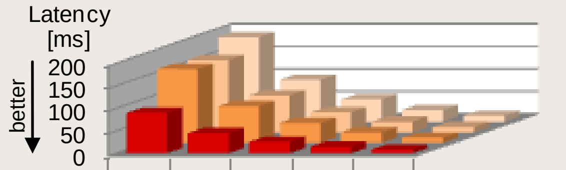

Outstanding IO	16	8	4	2	1
3.5" 7200 rpm 250 GB	90.1	45.1	25.6	14.8	8.8
2.5" 5400 rpm 500 GB	163.6	81.5	44.4	23.3	11.8
2.5" 5400 rpm 320 GB	160.7	81.9	44.3	23.4	11.9
2.5" 5400 rpm 160 GB	187.9	93.4	49.7	26.4	13.3

© Fujitsu Technology Solutions 2011

Standard load profile "File copy" (random access, 50% read, 50% write, 64 kB block size)

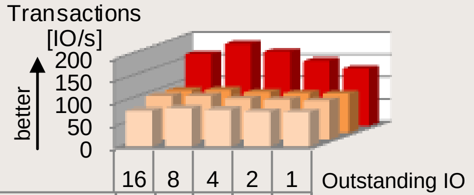

	16	8	4	2	1
2.5" 5400 rpm 160 GB	81	86	83	79	78
2.5" 5400 rpm 320 GB	98	99	93	90	88
2.5" 5400 rpm 500 GB	95	97	91	89	88
3.5" 7200 rpm 250 GB	162	185	168	147	129

Outstanding IO

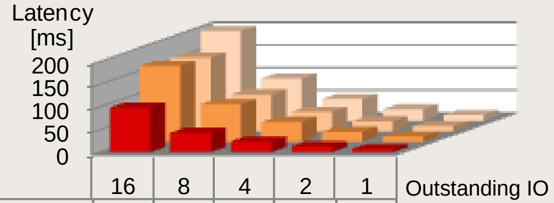

	16	8	4	2	1
3.5" 7200 rpm 250 GB	98.9	43.3	23.8	13.6	7.8
2.5" 5400 rpm 500 GB	167.7	82.8	43.8	22.4	11.3
2.5" 5400 rpm 320 GB	162.5	81.2	43.1	22.2	11.3
2.5" 5400 rpm 160 GB	198.3	93.4	48.3	25.4	12.9

Outstanding IO

Standard load profile "Streaming" (sequential access, 100% read, 64 kB block size)

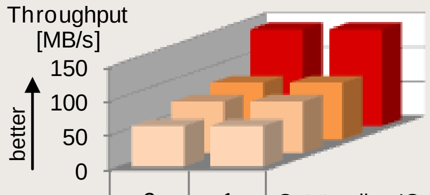

	2	1
2.5" 5400 rpm 160 GB	59.63	59.62
2.5" 5400 rpm 320 GB	74.91	74.90
2.5" 5400 rpm 500 GB	82.47	82.47
3.5" 7200 rpm 250 GB	140.67	140.79

Outstanding IO

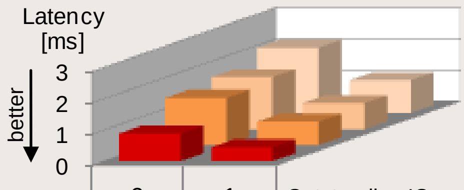

	2	1
3.5" 7200 rpm 250 GB	0.9	0.4
2.5" 5400 rpm 500 GB	1.5	0.8
2.5" 5400 rpm 320 GB	1.7	0.8
2.5" 5400 rpm 160 GB	2.1	1.0

Outstanding IO

Standard load profile "Restore" (sequential access, 100% write, 64 kB block size)

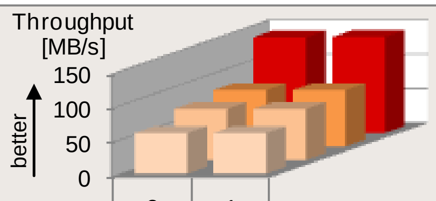

	2	1
2.5" 5400 rpm 160 GB	59.35	59.43
2.5" 5400 rpm 320 GB	74.61	74.72
2.5" 5400 rpm 500 GB	82.24	82.19
3.5" 7200 rpm 250 GB	139.04	139.54

Outstanding IO

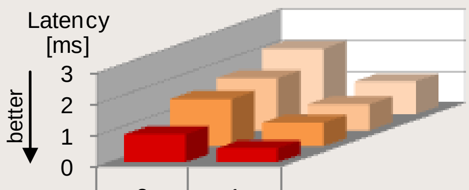

	2	1
3.5" 7200 rpm 250 GB	0.9	0.4
2.5" 5400 rpm 500 GB	1.5	0.8
2.5" 5400 rpm 320 GB	1.7	0.8
2.5" 5400 rpm 160 GB	2.1	1.1

Outstanding IO

© Fujitsu Technology Solutions 2011

Hard disks of the BC class

The following ten diagrams show performance values (transactions, throughput rates and response times) for the standard load profiles "File server", "Database", "File copy", "Streaming" and "Restore" with increasing loads through the number of parallel requests (outstanding IO). The following actual hard disks are considered

BC	HDD	SATA	3 Gbit/s	2.5″	7200 rpm	160 GB
BC	HDD	SATA	3 Gbit/s	2.5″	7200 rpm	500 GB
BC	HDD	SATA	3 Gbit/s	2.5″	7200 rpm	1 TB
BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	250 GB
BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	500 GB
BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	1 TB
BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	2 TB
BC	HDD	SAS	6 Gbit/s	2.5″	10000 rpm	146 GB
BC	HDD	SAS	6 Gbit/s	2.5″	10000 rpm	300 GB
BC	HDD	SAS	6 Gbit/s	2.5″	10000 rpm	450 GB
BC	HDD	SAS	6 Gbit/s	2.5″	10000 rpm	600 GB

© Fujitsu Technology Solutions 2011

Standard load profile "Database" (random access, 67% read, 33% write, 8 kB block size)

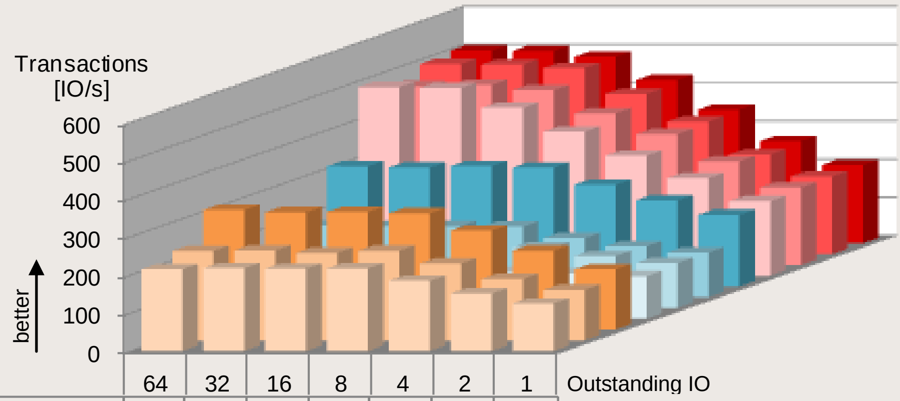

	64	32	16	8	4	2	1
SATA 2.5" 160 GB	215	217	216	217	184	151	125
SATA 2.5" 500 GB	234	235	229	235	202	161	132
SATA 2.5" 1 TB	314	307	308	305	260	207	158
SATA 3.5" 250 GB	136	133	133	136	122	119	112
SATA 3.5" 500 GB	192	187	188	190	159	136	118
SATA 3.5" 1 TB	188	187	186	185	155	133	117
SATA 3.5" 2 TB	315	312	316	312	267	226	188
SAS 2.5" 146 GB	494	492	440	378	315	256	195
SAS 2.5" 300 GB	469	472	458	397	343	272	203
SAS 2.5" 450 GB	499	497	488	420	348	261	204
SAS 2.5" 600 GB	505	504	489	429	350	268	206

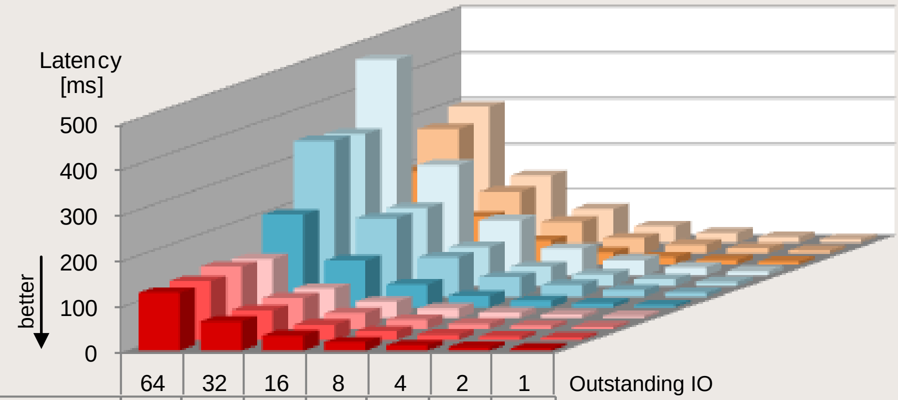

	64	32	16	8	4	2	1
SAS 2.5" 600 GB	126.6	63.4	32.7	18.7	11.4	7.5	4.9
SAS 2.5" 450 GB	128.4	64.4	32.8	19.0	11.5	7.6	4.9
SAS 2.5" 300 GB	136.6	67.8	34.9	20.1	11.7	7.4	4.9
SAS 2.5" 146 GB	129.5	65.0	36.4	21.2	12.7	7.8	5.1
SATA 3.5" 2 TB	203.2	102.6	50.7	25.7	15.0	8.9	5.3
SATA 3.5" 1 TB	340.8	171.5	86.0	43.3	25.8	15.0	8.5
SATA 3.5" 500 GB	332.9	170.7	85.1	42.0	25.2	14.7	8.5
SATA 3.5" 250 GB	471.3	240.9	120.5	58.8	32.7	16.9	9.0
SATA 2.5" 1 TB	204.0	104.1	51.9	26.2	15.4	9.7	6.3
SATA 2.5" 500 GB	273.0	136.0	69.8	34.1	19.8	12.4	7.6
SATA 2.5" 160 GB	298.4	147.3	74.1	36.9	21.7	13.2	8.0

© Fujitsu Technology Solutions 2011

Standard load profile "File server" (random access, 67% read, 33% write, 64 kB block size)

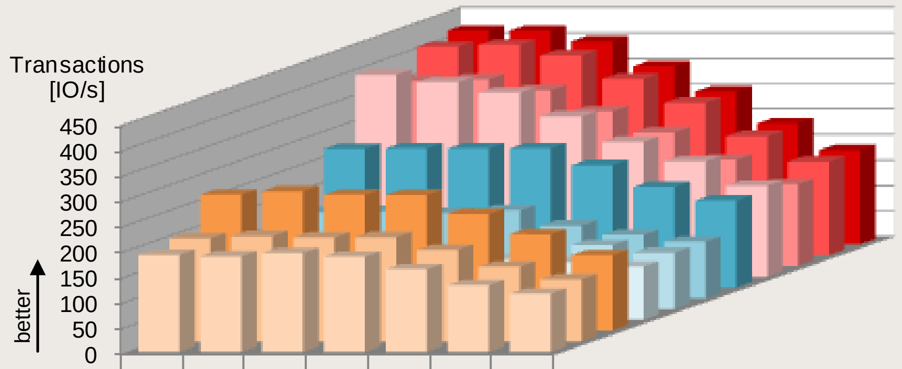

Outstanding IO	64	32	16	8	4	2	1
SATA 2.5” 160 GB	191	188	195	187	163	132	115
SATA 2.5” 500 GB	203	207	204	205	180	147	122
SATA 2.5” 1 TB	267	274	267	266	229	189	148
SATA 3.5” 250 GB	128	125	123	126	119	113	105
SATA 3.5” 500 GB	173	172	174	172	147	126	111
SATA 3.5” 1 TB	169	171	168	173	143	124	111
SATA 3.5” 2 TB	272	273	273	273	240	198	171
SAS 2.5” 146 GB	397	382	362	315	265	227	179
SAS 2.5” 300 GB	364	367	345	304	262	209	160
SAS 2.5” 450 GB	410	413	393	347	299	234	184
SAS 2.5” 600 GB	420	420	399	350	301	238	185

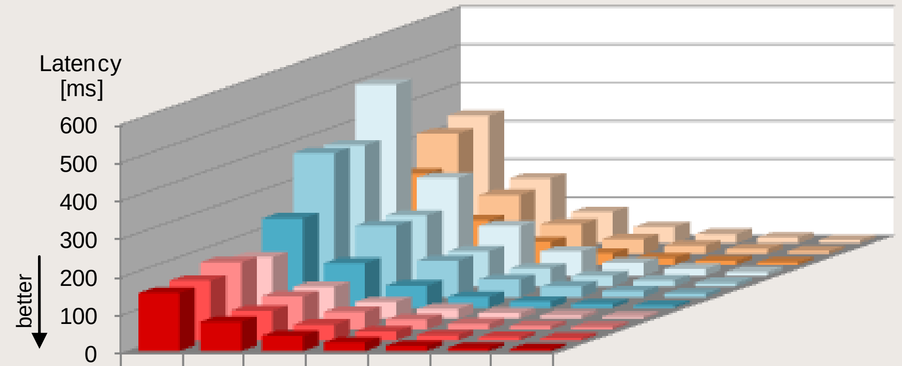

Outstanding IO	64	32	16	8	4	2	1
SAS 2.5” 600 GB	152.3	76.1	40.1	22.8	13.3	8.4	5.4
SAS 2.5” 450 GB	156.2	77.4	40.8	23.1	13.4	8.6	5.4
SAS 2.5” 300 GB	176.0	87.3	46.3	26.3	15.3	9.6	6.2
SAS 2.5” 146 GB	161.2	83.7	44.2	25.4	15.1	8.8	5.6
SATA 3.5” 2 TB	234.9	117.1	58.5	29.3	16.6	10.1	5.9
SATA 3.5” 1 TB	377.9	186.9	95.4	46.3	28.1	16.1	9.0
SATA 3.5” 500 GB	369.9	185.5	92.0	46.5	27.2	15.9	9.0
SATA 3.5” 250 GB	500.9	256.1	129.8	63.2	33.7	17.7	9.6
SATA 2.5” 1 TB	239.3	116.9	60.0	30.0	17.5	10.6	6.7
SATA 2.5” 500 GB	316.0	154.8	78.4	39.1	22.2	13.6	8.2
SATA 2.5” 160 GB	334.6	170.0	82.2	42.7	24.5	15.2	8.7

 © Fujitsu Technology Solutions 2011

Standard load profile "File copy" (random access, 50% read, 50% write, 64 kB block size)

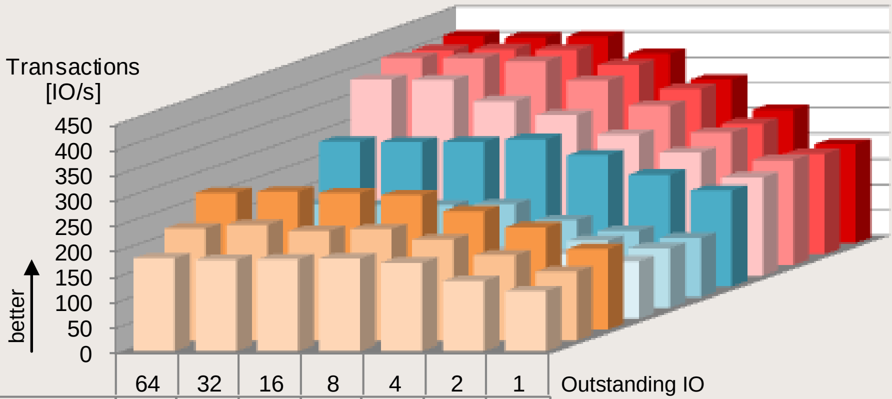

	64	32	16	8	4	2	1
SATA 2.5" 160 GB	182	179	180	182	174	138	118
SATA 2.5" 500 GB	220	227	215	219	198	167	135
SATA 2.5" 1 TB	269	271	268	264	233	201	158
SATA 3.5" 250 GB	139	137	137	137	132	127	113
SATA 3.5" 500 GB	184	185	186	184	155	132	118
SATA 3.5" 1 TB	182	181	181	183	152	130	117
SATA 3.5" 2 TB	284	284	284	289	258	220	189
SAS 2.5" 146 GB	385	385	343	316	277	242	193
SAS 2.5" 300 GB	407	406	400	362	313	260	206
SAS 2.5" 450 GB	401	403	401	372	325	257	197
SAS 2.5" 600 GB	408	405	407	373	322	262	195

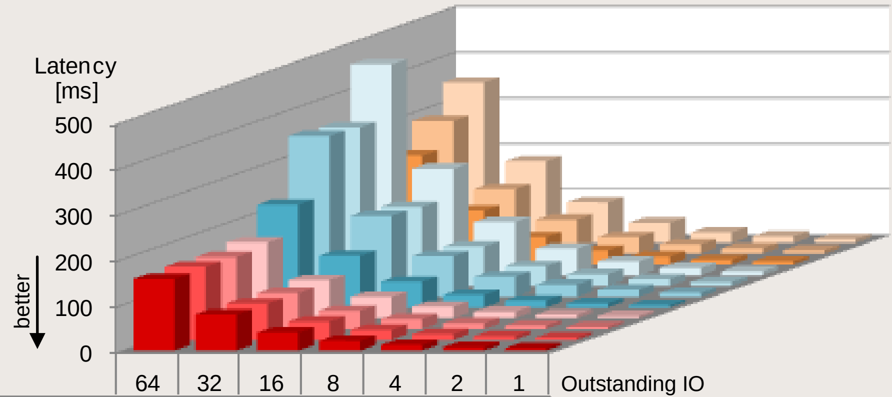

	64	32	16	8	4	2	1
SAS 2.5" 600 GB	156.8	79.0	39.3	21.4	12.4	7.6	5.1
SAS 2.5" 450 GB	159.6	79.4	39.9	21.5	12.3	7.8	5.1
SAS 2.5" 300 GB	157.3	78.8	40.0	22.1	12.8	7.7	4.9
SAS 2.5" 146 GB	166.0	83.2	46.6	25.3	14.4	8.3	5.2
SATA 3.5" 2 TB	225.0	112.8	56.3	27.7	15.5	9.1	5.3
SATA 3.5" 1 TB	351.7	176.6	88.5	43.8	26.4	15.4	8.6
SATA 3.5" 500 GB	347.2	173.2	86.2	43.4	25.8	15.1	8.5
SATA 3.5" 250 GB	461.2	233.3	116.5	58.3	30.2	15.7	8.8
SATA 2.5" 1 TB	238.2	118.0	59.7	30.3	17.2	10.0	6.3
SATA 2.5" 500 GB	290.7	140.9	74.4	36.6	20.2	12.0	7.4
SATA 2.5" 160 GB	351.1	178.6	88.7	44.0	23.0	14.5	8.5

Standard load profile "Streaming" (sequential access, 100% read, 64 kB block size)

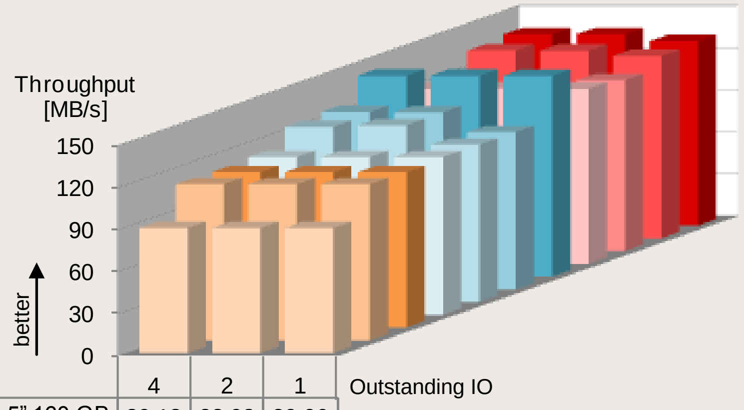

	4	2	1
SATA 2.5" 160 GB	89.13	88.98	89.06
SATA 2.5" 500 GB	111.20	111.16	111.13
SATA 2.5" 1 TB	110.68	110.58	110.58
SATA 3.5" 250 GB	112.42	112.38	112.37
SATA 3.5" 500 GB	124.83	125.86	112.01
SATA 3.5" 1 TB	126.28	126.16	112.07
SATA 3.5" 2 TB	143.15	143.11	143.07
SAS 2.5" 146 GB	124.77	124.82	124.77
SAS 2.5" 300 GB	121.70	121.77	121.74
SAS 2.5" 450 GB	133.74	133.71	129.98
SAS 2.5" 600 GB	136.18	136.18	131.33

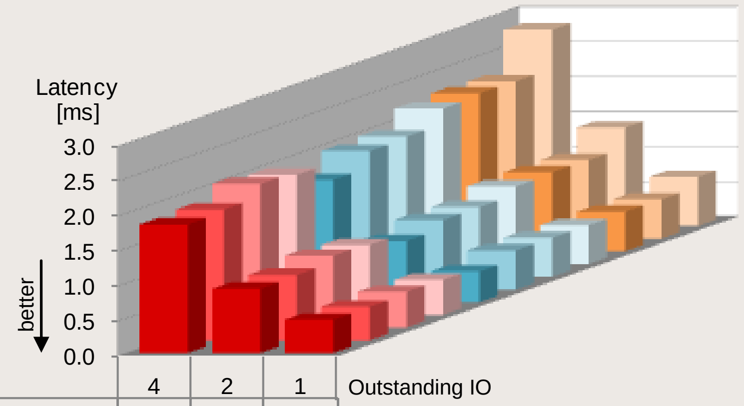

	4	2	1
SAS 2.5" 600 GB	1.8	0.9	0.5
SAS 2.5" 450 GB	1.9	0.9	0.5
SAS 2.5" 300 GB	2.1	1.0	0.5
SAS 2.5" 146 GB	2.0	1.0	0.5
SATA 3.5" 2 TB	1.7	0.9	0.4
SATA 3.5" 1 TB	2.0	1.0	0.6
SATA 3.5" 500 GB	2.0	1.0	0.6
SATA 3.5" 250 GB	2.2	1.1	0.6
SATA 2.5" 1 TB	2.3	1.1	0.6
SATA 2.5" 500 GB	2.2	1.1	0.6
SATA 2.5" 160 GB	2.8	1.4	0.7

 © Fujitsu Technology Solutions 2011

Standard load profile "Restore" (sequential access, 100% write, 64 kB block size)

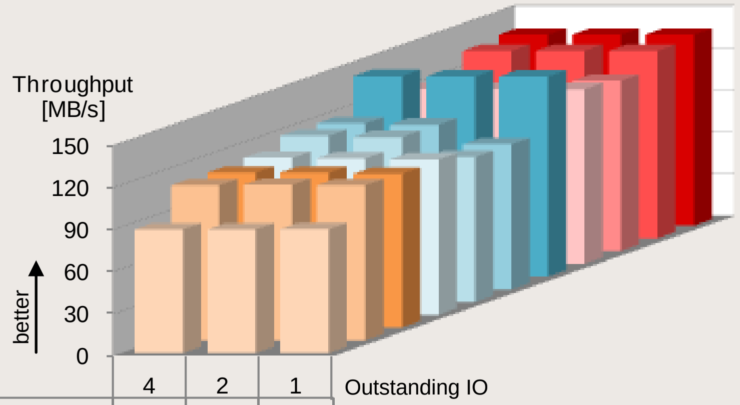

	4	2	1
SATA 2.5" 160 GB	88.04	88.04	88.23
SATA 2.5" 500 GB	110.38	110.72	110.43
SATA 2.5" 1 TB	110.35	110.23	109.43
SATA 3.5" 250 GB	111.78	111.33	110.61
SATA 3.5" 500 GB	118.33	117.48	103.36
SATA 3.5" 1 TB	118.14	117.27	103.30
SATA 3.5" 2 TB	142.78	142.57	142.53
SAS 2.5" 146 GB	124.63	124.44	124.42
SAS 2.5" 300 GB	121.46	121.34	121.32
SAS 2.5" 450 GB	133.56	133.57	133.54
SAS 2.5" 600 GB	136.01	135.98	135.93

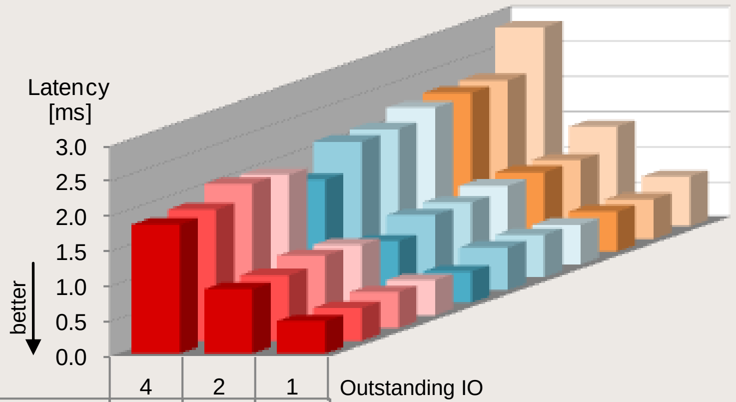

	4	2	1
SAS 2.5" 600 GB	1.8	0.9	0.5
SAS 2.5" 450 GB	1.9	0.9	0.5
SAS 2.5" 300 GB	2.1	1.0	0.5
SAS 2.5" 146 GB	2.0	1.0	0.5
SATA 3.5" 2 TB	1.8	0.9	0.4
SATA 3.5" 1 TB	2.1	1.1	0.6
SATA 3.5" 500 GB	2.1	1.1	0.6
SATA 3.5" 250 GB	2.2	1.1	0.6
SATA 2.5" 1 TB	2.3	1.1	0.6
SATA 2.5" 500 GB	2.3	1.1	0.6
SATA 2.5" 160 GB	2.8	1.4	0.7

© Fujitsu Technology Solutions 2011

Hard disks of the EP class

The following ten diagrams show performance values (transactions, throughput rates and response times) for the standard load profiles "File server", "Database", "File copy", "Streaming" and "Restore" with increasing loads through the number of parallel requests (outstanding IO). The following actual hard disks are considered

EP	HDD	SAS	6 Gbit/s	2.5″	15000 rpm	73 GB
EP	HDD	SAS	6 Gbit/s	2.5″	15000 rpm	146 GB
EP	HDD	SAS	6 Gbit/s	3.5″	15000 rpm	146 GB
EP	HDD	SAS	6 Gbit/s	3.5″	15000 rpm	300 GB
EP	HDD	SAS	6 Gbit/s	3.5″	15000 rpm	450 GB
EP	HDD	SAS	6 Gbit/s	3.5″	15000 rpm	600 GB
EP	SSD	SATA	3 Gbit/s	2.5″		32 GB
EP	SSD	SATA	3 Gbit/s	2.5″		64 GB

Standard load profile "Database" (random access, 67% read, 33% write, 8 kB block size)

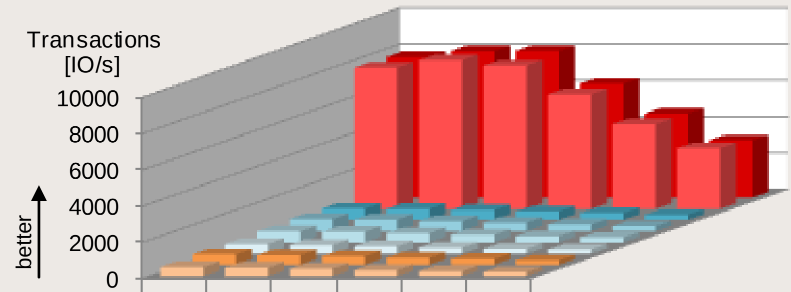

Outstanding IO	32	16	8	4	2	1
HDD SAS 2.5" 73 GB	505	467	408	343	286	236
HDD SAS 2.5" 146 GB	589	551	489	406	335	258
HDD SAS 3.5" 146 GB	539	481	410	336	273	232
HDD SAS 3.5" 300 GB	617	592	516	429	346	268
HDD SAS 3.5" 450 GB	669	641	553	468	371	281
HDD SAS 3.5" 600 GB	632	624	557	470	382	284
SSD SATA 2.5" 32 GB	7783	8198	7898	6305	4671	3343
SSD SATA 2.5" 64 GB	7754	8077	8083	6279	4633	3153

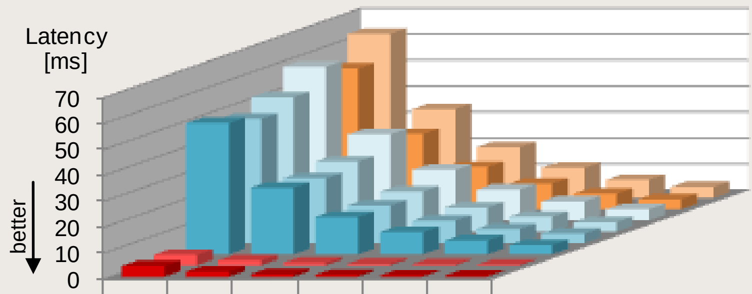

Outstanding IO	32	16	8	4	2	1
SSD SATA 2.5" 64 GB	4.1	2.0	1.0	0.6	0.4	0.3
SSD SATA 2.5" 32 GB	4.1	2.0	1.0	0.6	0.4	0.3
HDD SAS 3.5" 600 GB	50.6	25.6	14.4	8.5	5.2	3.5
HDD SAS 3.5" 450 GB	47.9	25.0	14.5	8.5	5.4	3.6
HDD SAS 3.5" 300 GB	51.9	27.0	15.5	9.3	5.8	3.7
HDD SAS 3.5" 146 GB	59.4	33.2	19.5	11.9	7.3	4.3
HDD SAS 2.5" 146 GB	54.3	29.0	16.3	9.9	6.0	3.9
HDD SAS 2.5" 73 GB	63.3	34.3	19.6	11.7	7.0	4.2

© Fujitsu Technology Solutions 2011

Standard load profile "File server" (random access, 67% read, 33% write, 64 kB block size)

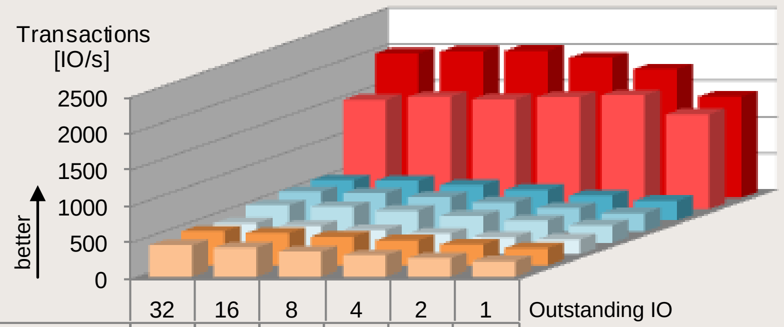

	32	16	8	4	2	1
HDD SAS 2.5" 73 GB	433	398	345	294	255	213
HDD SAS 2.5" 146 GB	470	443	395	344	286	234
HDD SAS 3.5" 146 GB	424	386	332	282	235	203
HDD SAS 3.5" 300 GB	519	496	445	372	306	243
HDD SAS 3.5" 450 GB	550	528	472	403	325	254
HDD SAS 3.5" 600 GB	552	546	483	419	340	267
SSD SATA 2.5" 32 GB	1504	1544	1510	1543	1566	1307
SSD SATA 2.5" 64 GB	1990	2015	2023	1928	1770	1391

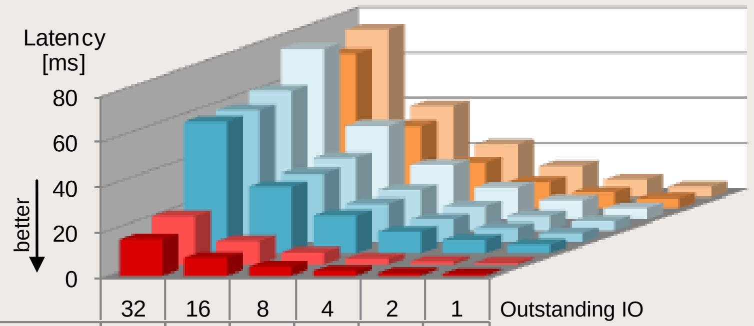

	32	16	8	4	2	1
SSD SATA 2.5" 64 GB	16.1	7.9	4.0	2.1	1.1	0.7
SSD SATA 2.5" 32 GB	21.3	10.4	5.3	2.6	1.3	0.8
HDD SAS 3.5" 600 GB	58.0	29.3	16.6	9.6	5.9	3.7
HDD SAS 3.5" 450 GB	58.1	30.3	17.0	9.9	6.1	3.9
HDD SAS 3.5" 300 GB	61.7	32.2	18.0	10.7	6.5	4.1
HDD SAS 3.5" 146 GB	75.4	41.4	24.1	14.2	8.5	4.9
HDD SAS 2.5" 146 GB	68.1	36.2	20.3	11.6	7.0	4.3
HDD SAS 2.5" 73 GB	73.9	40.2	23.2	13.6	7.9	4.7

Standard load profile "File copy" (random access, 50% read, 50% write, 64 kB block size)

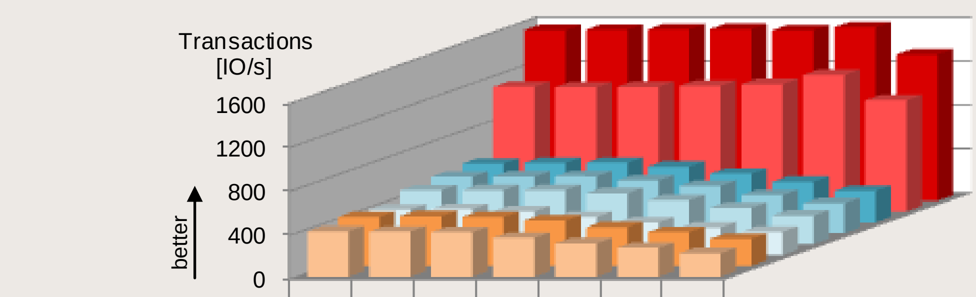

Outstanding IO	64	32	16	8	4	2	1
HDD SAS 2.5” 73 GB	415	412	407	356	305	265	214
HDD SAS 2.5” 146 GB	447	454	446	415	353	308	248
HDD SAS 3.5” 146 GB	414	418	397	348	297	250	211
HDD SAS 3.5” 300 GB	491	491	493	464	402	330	259
HDD SAS 3.5” 450 GB	514	516	518	481	432	348	274
HDD SAS 3.5” 600 GB	534	541	543	508	443	369	286
SSD SATA 2.5" 32 GB	1137	1132	1136	1145	1159	1245	1014
SSD SATA 2.5" 64 GB	1555	1557	1561	1564	1552	1580	1331

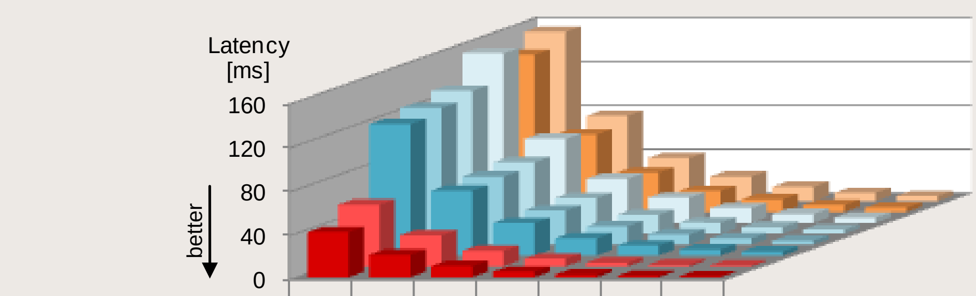

Outstanding IO	64	32	16	8	4	2	1
SSD SATA 2.5" 64 GB	41.2	20.5	10.2	5.1	2.6	1.3	0.8
SSD SATA 2.5" 32 GB	56.3	28.3	14.1	7.0	3.5	1.6	1.0
HDD SAS 3.5” 600 GB	119.8	59.2	29.4	15.7	9.0	5.4	3.5
HDD SAS 3.5” 450 GB	124.4	62.0	30.9	16.6	9.3	5.7	3.7
HDD SAS 3.5” 300 GB	130.3	65.1	32.5	17.3	10.0	6.1	3.9
HDD SAS 3.5” 146 GB	154.5	76.6	40.3	23.0	13.5	8.0	4.7
HDD SAS 2.5” 146 GB	143.3	70.5	35.9	19.3	11.3	6.5	4.0
HDD SAS 2.5” 73 GB	154.3	77.7	39.4	22.5	13.1	7.6	4.7

© Fujitsu Technology Solutions 2011

Standard load profile "Streaming" (sequential access, 100% read, 64 kB block size)

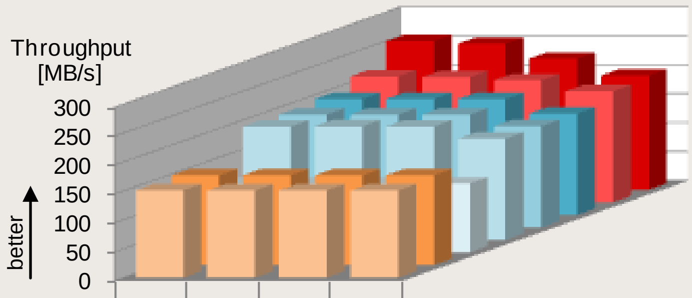

Outstanding IO	8	4	2	1
HDD SAS 2.5" 73 GB	150.60	150.60	150.60	150.58
HDD SAS 2.5" 146 GB	154.95	155.04	155.00	155.02
HDD SAS 3.5" 146 GB	119.30	119.27	119.29	119.24
HDD SAS 3.5" 300 GB	196.20	196.28	196.23	175.13
HDD SAS 3.5" 450 GB	195.10	195.18	195.12	175.13
HDD SAS 3.5" 600 GB	200.15	200.15	200.15	175.15
SSD SATA 2.5" 32 GB	217.72	216.02	211.07	191.74
SSD SATA 2.5" 64 GB	257.20	252.24	226.40	197.53

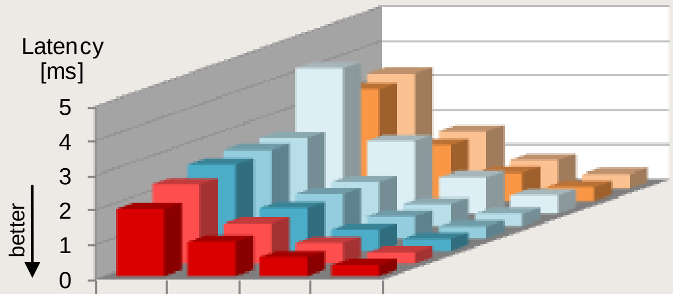

Outstanding IO	8	4	2	1
SSD SATA 2.5" 64 GB	1.9	1.0	0.6	0.3
SSD SATA 2.5" 32 GB	2.3	1.2	0.6	0.3
HDD SAS 3.5" 600 GB	2.5	1.2	0.6	0.4
HDD SAS 3.5" 450 GB	2.6	1.3	0.6	0.4
HDD SAS 3.5" 300 GB	2.5	1.3	0.6	0.4
HDD SAS 3.5" 146 GB	4.2	2.1	1.0	0.5
HDD SAS 2.5" 146 GB	3.2	1.6	0.8	0.4
HDD SAS 2.5" 73 GB	3.3	1.7	0.8	0.4

Standard load profile "Restore" (sequential access, 100% write, 64 kB block size)

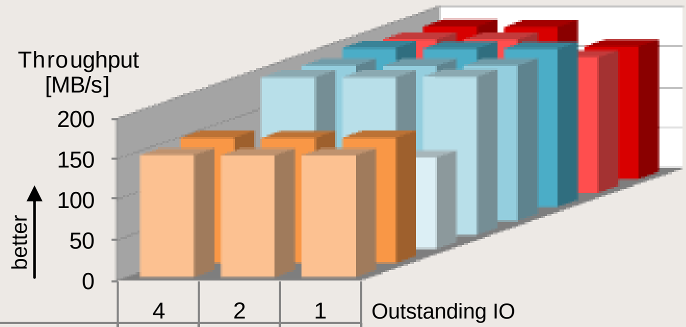

	4	2	1
HDD SAS 2.5” 73 GB	150.44	150.31	150.33
HDD SAS 2.5” 146 GB	154.71	154.49	154.59
HDD SAS 3.5” 146 GB	113.65	113.88	113.10
HDD SAS 3.5” 300 GB	194.60	194.65	195.41
HDD SAS 3.5” 450 GB	191.65	191.28	191.62
HDD SAS 3.5” 600 GB	195.99	195.32	195.04
SSD SATA 2.5" 32 GB	189.19	189.16	167.17
SSD SATA 2.5" 64 GB	188.00	187.30	163.12

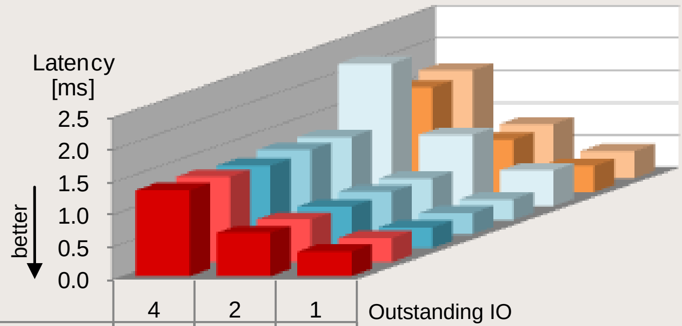

	4	2	1
SSD SATA 2.5" 64 GB	1.3	0.7	0.4
SSD SATA 2.5" 32 GB	1.3	0.7	0.4
HDD SAS 3.5” 600 GB	1.3	0.6	0.3
HDD SAS 3.5” 450 GB	1.3	0.7	0.3
HDD SAS 3.5” 300 GB	1.3	0.6	0.3
HDD SAS 3.5” 146 GB	2.2	1.1	0.6
HDD SAS 2.5” 146 GB	1.6	0.8	0.4
HDD SAS 2.5” 73 GB	1.7	0.8	0.4

 © Fujitsu Technology Solutions 2011

Performance data in comparison

The individual hard disk features, which are significant for the influence on their performance, are handled below. One of these features is a switchable parameter: the enabling/disabling of the disk cache has a great impact on hard disk performance and simultaneously touches upon security aspects that have to be observed when configuring a server. The three other hard disk features "form factor", "rotational speed" and "capacity" must be taken into account when selecting the suitable hard disk for a specific server and do not only have an impact on server performance. They are of central importance when it comes to server sizing.

Server sizing

The minimum requirements for the total capacity must first be observed when selecting a server and its hard disks. Especially the number of hard disks that can be installed in a server plays a role here and also in relation to the individual hard disk

- form factor
- rotational speed
- capacity

The following table provides an overview of the PRIMERGY servers that are available today and their expandability with hard disks (form factor 2.5″ hard disks or 3.5″ hard disks) and the maximum total capacity achievable:

Server blades	Maximum configuration 3.5″	Maximum overall capacity 3.5″	Maximum configuration 2.5″	Maximum overall capacity 2.5″
PRIMERGY BX620 S6			2	SAS: 1.2 TB SATA: 2 TB
PRIMERGY BX920 S2			2	SAS: 1.2 TB SATA: 2 TB
PRIMERGY BX922 S2			2	SATA: 0.3 TB
PRIMERGY BX924 S2			2	SATA: 0.1 TB
PRIMERGY BX960 S1			2	SATA: 0.1 TB
Storage blades	**Maximum configuration 3.5″**	**Maximum overall capacity 3.5″**	**Maximum configuration 2.5″**	**Maximum overall capacity 2.5″**
PRIMERGY SX650			5	SAS: 1.5 TB SATA: 5 TB
PRIMERGY SX940 S1			4	SAS: 2.4 TB SATA: 4 TB
PRIMERGY SX960 S1			10	SAS: 6 TB SATA: 10 TB
Scale-out servers	**Maximum configuration 3.5″**	**Maximum overall capacity 3.5″**	**Maximum configuration 2.5″**	**Maximum overall capacity 2.5″**
PRIMERGY CX120 S1			2	SATA: 1 TB
PRIMERGY CX122 S1			2	SATA: 1 TB

© Fujitsu Technology Solutions 2011

Rack servers	Maximum configuration 3.5″	Maximum overall capacity 3.5″	Maximum configuration 2.5″	Maximum overall capacity 2.5″
PRIMERGY RX100 S6	2	SAS: 1.2 TB SATA: 4 TB	4	SAS: 2.4 TB SATA: 4 TB
PRIMERGY RX100 S7	2	SAS: 1.2 TB SATA: 4 TB	4	SAS: 2.4 TB SATA: 4 TB
PRIMERGY RX200 S6			8	SAS: 4.8 TB SATA: 8 TB
PRIMERGY RX300 S6	6	SAS: 3.6 TB SATA: 12 TB	12	SAS: 7.2 TB SATA: 12 TB
PRIMERGY RX600 S5			8	SAS: 4.8 TB SATA: 8 TB
PRIMERGY RX600 S6			8	SAS: 4.8 TB SATA: 0.5 TB
PRIMERGY RX900 S1			8	SAS: 4.8 TB SATA: 0.5 TB
PRIMERGY RX900 S2			8	SAS: 4.8 TB SATA: 0.5 TB
Tower servers	**Maximum configuration 3.5″**	**Maximum overall capacity 3.5″**	**Maximum configuration 2.5″**	**Maximum overall capacity 2.5″**
PRIMERGY TX100 S2	4	SATA: 8 TB		
PRIMERGY TX120 S3	2	SATA: 4 TB	4	SAS: 2.4 TB SATA: 4 TB
PRIMERGY TX140 S1	4	SAS: 2.4 TB SATA: 8 TB	8	SAS: 4.8 TB SATA: 8 TB
PRIMERGY TX150 S7	4	SAS: 2.4 TB SATA: 8 TB	8	SAS: 4.8 TB SATA: 8 TB
PRIMERGY TX200 S6	6	SAS: 3.6 TB SATA: 12 TB	16	SAS: 9.2 TB SATA: 16 TB
PRIMERGY TX300 S6	8	SAS: 4.8 TB SATA: 16 TB	20	SAS: 12 TBSATA : 20 TB

The maximum total capacity can be achieved with the 2.5″ hard disk "BC HDD SATA 3 Gbit/s 2.5″ 7200 rpm 1 TB" with the PRIMERGY TX200 S6 and the PRIMERGY TX300 S6. With the other servers the maximum total capacity is equally high with 2.5″ and 3.5″ hard disks. When using hard disks with a SAS interface, the maximum total capacity can be achieved with the 2.5″ hard disk "BC HDD SAS 6 Gbit/s 2.5″ 10000 rpm 600 GB".

Up to twelve hard disks can be connected to a single controller in the PRIMERGY TX200 S6 and the PRIMERGY TX300 S6. A second controller is required for further hard disks.

© Fujitsu Technology Solutions 2011

Form factor

The hard disks with a 3.5″ form factor were widely used in servers for a long time. Compared with the 2.5″ hard disks, they offered higher capacity and/or better performance over a long period. Today, this only applies with regard to the capacity of SATA hard disks, and now differs from case to case as far as performance is concerned.

The advantages of 2.5″ hard disks are lower power consumption, less heat generation and less space required. As more 2.5″ hard disks can be housed in a server than 3.5″ hard disks, the capacity and performance differences are relativized in RAID arrays. The RAID performance of a server which is fully configured with 2.5″ hard disks is higher than a server fully configured with 3.5″ hard disks. Furthermore, the same high total capacity levels can be achieved in all PRIMERGY servers with 2.5″ hard disks and with 3.5″ SAS hard disks. And there are even considerably higher total capacity levels in some servers.

The following ten diagrams show performance values (transactions, throughput rates and response times) for the standard load profiles "File server", "Database", "File copy", "Streaming" and "Restore" with increasing loads through the number of parallel requests (outstanding IO). The following actual hard disks are considered

BC HDD SATA 3 Gbit/s 2.5″ 7200 rpm 500 GB

BC HDD SATA 3 Gbit/s 3.5″ 7200 rpm 500 GB

Standard load profile "Database" (random access, 67% read, 33% write, 8 kB block size)

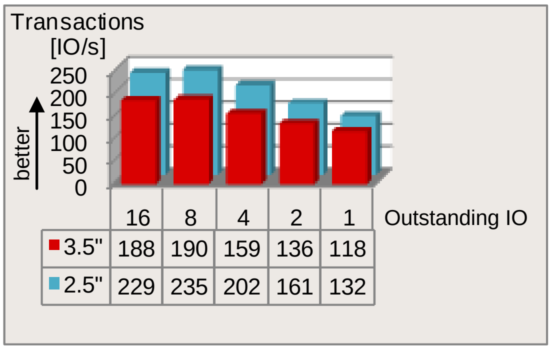

Outstanding IO	16	8	4	2	1
3.5"	188	190	159	136	118
2.5"	229	235	202	161	132

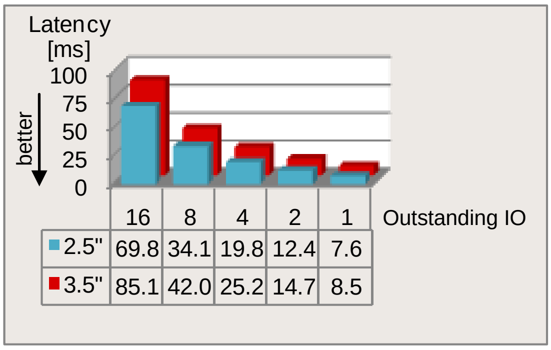

Outstanding IO	16	8	4	2	1
2.5"	69.8	34.1	19.8	12.4	7.6
3.5"	85.1	42.0	25.2	14.7	8.5

Standard load profile "File server" (random access, 67% read, 33% write, 64 kB block size)

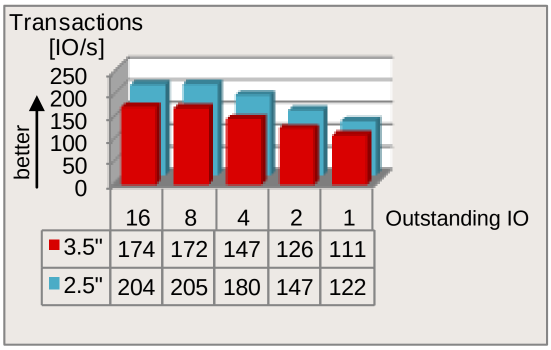

Outstanding IO	16	8	4	2	1
3.5"	174	172	147	126	111
2.5"	204	205	180	147	122

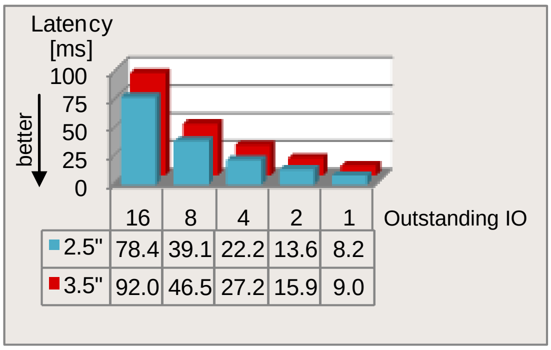

Outstanding IO	16	8	4	2	1
2.5"	78.4	39.1	22.2	13.6	8.2
3.5"	92.0	46.5	27.2	15.9	9.0

Standard load profile "File copy" (random access, 50% read, 50% write, 64 kB block size)

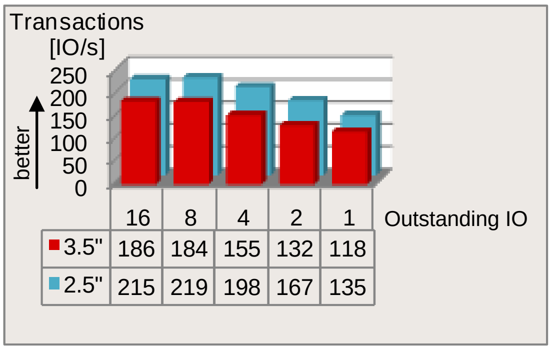

Outstanding IO	16	8	4	2	1
3.5"	186	184	155	132	118
2.5"	215	219	198	167	135

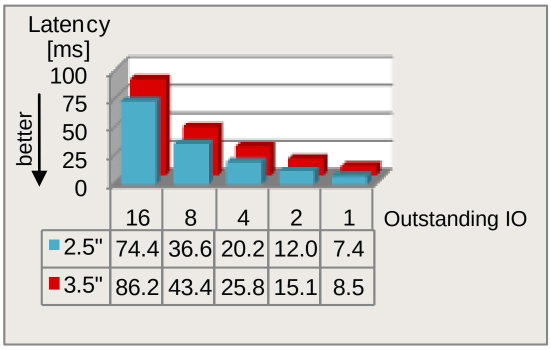

Outstanding IO	16	8	4	2	1
2.5"	74.4	36.6	20.2	12.0	7.4
3.5"	86.2	43.4	25.8	15.1	8.5

© Fujitsu Technology Solutions 2011

Standard load profile "Streaming" (sequential access, 100% read, 64 kB block size)

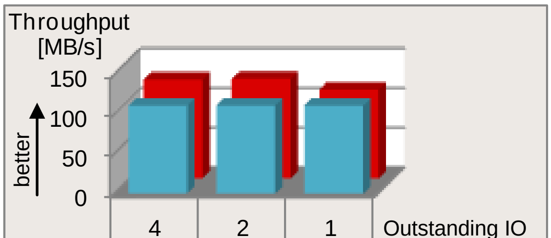

Outstanding IO	4	2	1
2.5"	111.20	111.16	111.13
3.5"	124.83	125.86	112.01

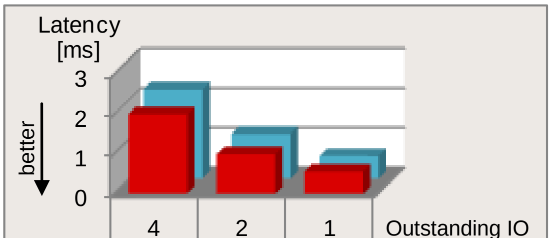

Outstanding IO	4	2	1
3.5"	2.0	1.0	0.6
2.5"	2.2	1.1	0.6

Standard load profile "Restore" (sequential access, 100% write, 64 kB block size)

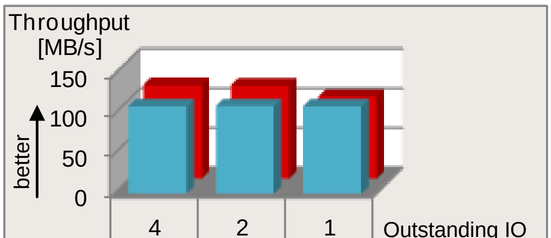

Outstanding IO	4	2	1
2.5"	110.38	110.72	110.43
3.5"	118.33	117.48	103.36

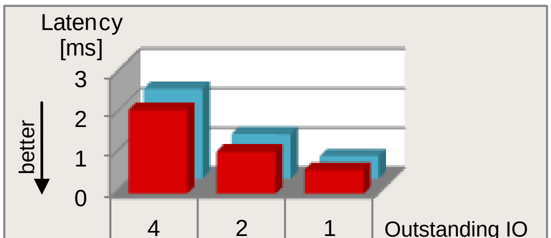

Outstanding IO	4	2	1
3.5"	2.1	1.1	0.6
2.5"	2.3	1.1	0.6

Whereas the 3.5″ hard disk with the standard load profiles and with sequential access provides a higher throughput than the 2.5″ hard disk, the 2.5″ hard disk with standard load profiles and with random access is at the forefront. However, this example cannot be generalized.

Rotational speed

In contrast to SSDs, data is saved to rotating disks with conventional hard disks. These revolve at a constant speed. The number of revolutions is normally not specified in the international SI system unit of 1/s, but in the English unit rpm for "revolutions per minute".

In high-performance scenarios either SSDs or hard disks with a rotational speed of 15000 rpm are used for servers. Hard disks with 10000 or 7200 rpm are used in the middle of the range performance segment. Where performance only plays a subordinate role or even no role at all, hard disks with 7200 rpm or even with only 5400 rpm can be used.

The faster the rotational speed of a hard disk, the higher its power consumption and thus also its heat generation, which means higher requirements for the cooling mechanisms in the server. Another aspect that has to be observed is the hard disk vibration that accompanies the rotation. As hard disks with different rotational speeds can disrupt each other, all the hard disks that are located in a common hard disk cage should be of the same type.

If one of the two hard disk types has a 50% higher rotational speed, this does not mean that all performance values are 50% higher. There are several reasons for this. On the one hand the two hard disk types can differ in the number of sectors per track, which has an effect on the amount of data that can be processed per revolution. It is typical for current 10000 rpm hard disks to be offered with higher capacities than the current 15000 rpm hard disks. This difference usually reduces the advantage of the higher revolution speed of the faster disk. On the other hand, the response time for a hard disk access does not only comprise the rotation-defined latency time, but also includes additional components, such as the relevant "track-to-track seek time" of random accesses.

Due to the constant technical development of hard disks the change to a newer hard disk generation can provide better performance, i.e. a modern hard disk with 10000 rpm can be faster than an older hard disk with 15000 rpm.

 © Fujitsu Technology Solutions 2011

The influence of the rotational speed on performance is now to be explained on the basis of two hard disk types, whose order-relevant technical data only differs in this aspect. The hard disk cache was enabled.

BC	HDD	SAS	6 Gbit/s	2.5″	10000 rpm	146 GB
EP	HDD	SAS	6 Gbit/s	2.5″	15000 rpm	146 GB

Standard load profile "Database" (random access, 67% read, 33% write, 8 kB block size)

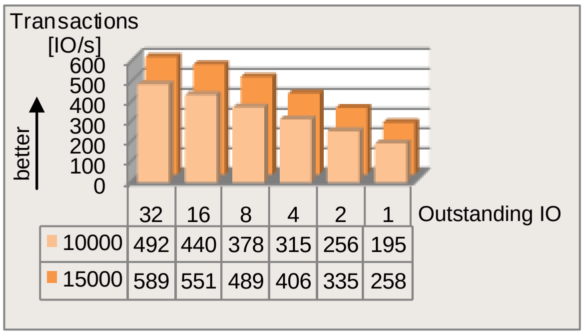

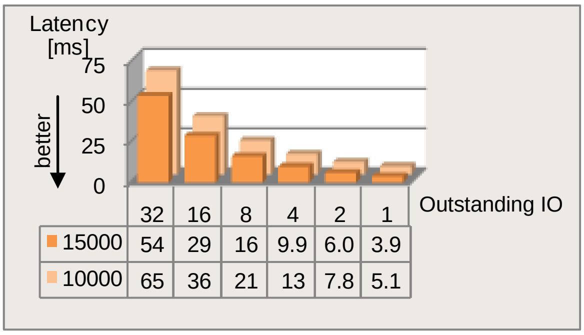

With the random load profile "Database" the faster rotating hard disk for low workload intensities (1 outstanding IO) has a performance advantage of 32%. The advantage is reduced to 20% for higher load intensities (32 outstanding IOs).

Standard load profile "File server" (random access, 67% read, 33% write, 64 kB block size)

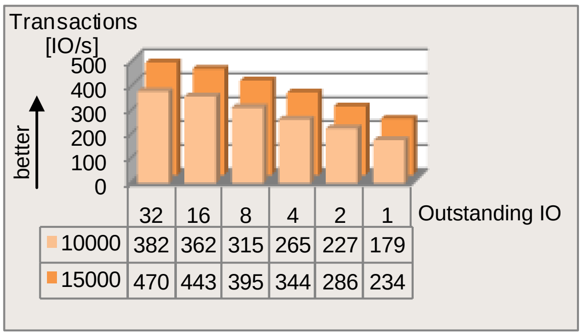

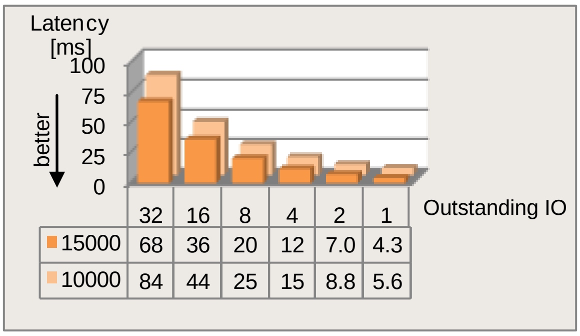

With the random load profile "File Server" the faster rotating hard disk for low load intensities (1 outstanding IO) has a performance advantage of 31%. The advantage is reduced to 23% for higher load intensities (32 outstanding IOs).

Standard load profile "File copy" (random access, 50% read, 50% write, 64 kB block size)

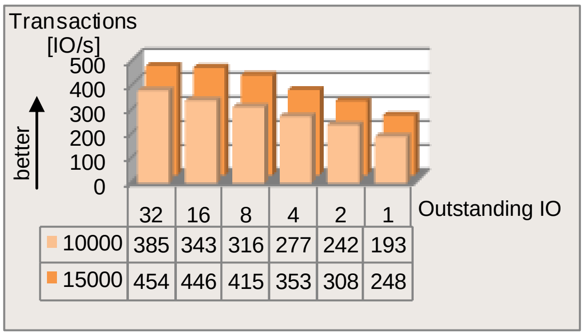

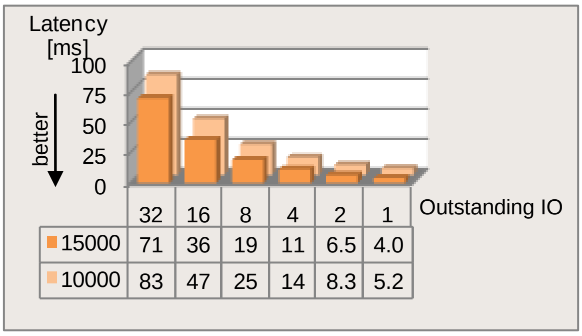

With the random load profile "File copy", the faster rotating hard disk for lower load intensities (1 outstanding IO) has a performance advantage of 28%. The advantage is reduced to 18% for higher load intensities (32 outstanding IOs).

Standard load profile "Streaming" (sequential access, 100% read, 64 kB block size)

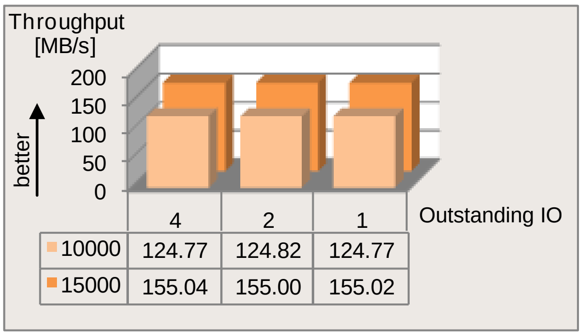

	4	2	1
10000	124.77	124.82	124.77
15000	155.04	155.00	155.02

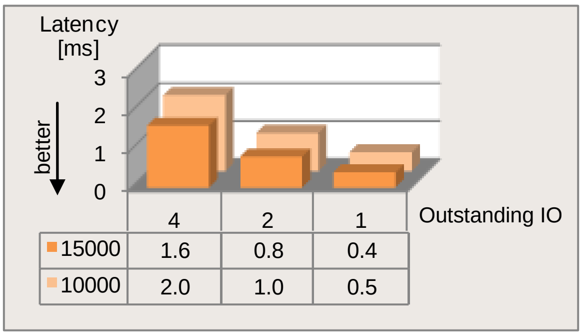

	4	2	1
15000	1.6	0.8	0.4
10000	2.0	1.0	0.5

With the sequential load profile "Streaming" the faster rotating hard disk remains very consistent with 24% more performance than the more slowly rotating hard disk.

Standard load profile "Restore" (sequential access, 100% write, 64 kB block size)

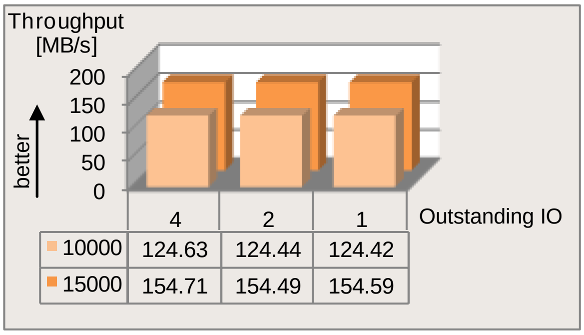

	4	2	1
10000	124.63	124.44	124.42
15000	154.71	154.49	154.59

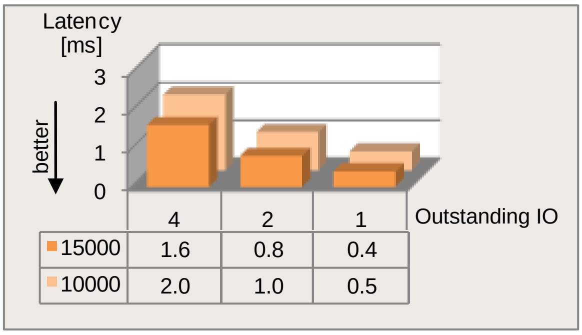

	4	2	1
15000	1.6	0.8	0.4
10000	2.0	1.0	0.5

With the sequential load profile "Restore" the faster rotating hard disk remains very consistent with 24% more performance than the more slowly rotating hard disk.

The time required to reposition the read/write head is important for performance with random load profiles, as it takes place frequently. The time required to reposition only plays a comparably minor role with sequential load profiles. The decisive factors here are the rotational speed and the number of sectors per track. It is therefore no wonder that the two sequential load profiles show a very constant performance difference here, whereas the advantages of the faster rotating hard disk with the three random load profiles vary more, depending on the type of load profile and the load intensity, but remain on average within the same dimension.

The example makes no claim to universal validity, but is certainly not untypical. In individual cases, it depends on the exact technical features of two hard disk types that are to be compared.

 © Fujitsu Technology Solutions 2011

Capacity

One of the main criteria when selecting the appropriate hard disk is its capacity. 2.5″ SATA-HDDs are offered with a capacity of up to 1 TB, 3.5″ SATA-HDDs with up to 2 TB. SAS-HDDs have capacities of up to 600 GB. SATA-SSDs currently have a maximum capacity of 64 GB.

The following ten diagrams show performance values (transactions, throughput rates and response times) for the standard load profiles "File server", "Database", "File copy", "Streaming" and "Restore" with increasing loads through the number of parallel requests (outstanding IO). The following actual hard disks are considered

BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	250 GB
BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	500 GB
BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	1 TB
BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	2 TB

Standard load profile "Database" (random access, 67% read, 33% write, 8 kB block size)

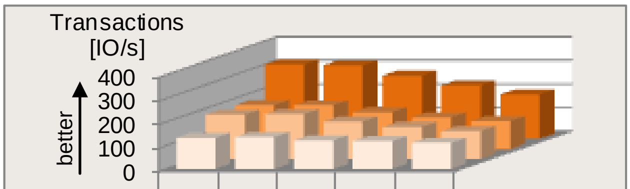

Outstanding IO	16	8	4	2	1
250 GB	133	136	122	119	112
500 GB	188	190	159	136	118
1000 GB	186	185	155	133	117
2000 GB	316	312	267	226	188

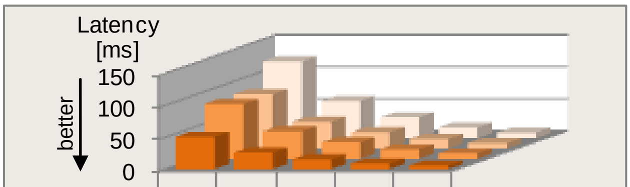

Outstanding IO	16	8	4	2	1
2000 GB	50.7	25.7	15.0	8.9	5.3
1000 GB	86.0	43.3	25.8	15.0	8.5
500 GB	85.1	42.0	25.2	14.7	8.5
250 GB	120.5	58.8	32.7	16.9	9.0

Standard load profile "File server" (random access, 67% read, 33% write, 64 kB block size)

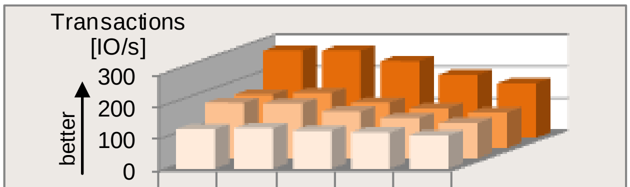

Outstanding IO	16	8	4	2	1
250 GB	123	126	119	113	105
500 GB	174	172	147	126	111
1000 GB	168	173	143	124	111
2000 GB	273	273	240	198	171

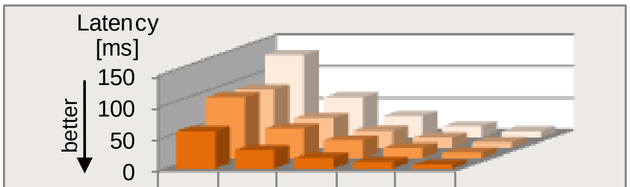

Outstanding IO	16	8	4	2	1
2000 GB	58.5	29.3	16.6	10.1	5.9
1000 GB	95.4	46.3	28.1	16.1	9.0
500 GB	92.0	46.5	27.2	15.9	9.0
250 GB	129.8	63.2	33.7	17.7	9.6

Standard load profile "File copy" (random access, 50% read, 50% write, 64 kB block size)

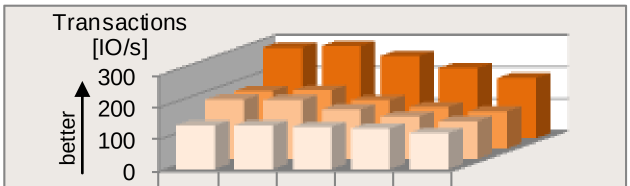

Outstanding IO	16	8	4	2	1
250 GB	137	137	132	127	113
500 GB	186	184	155	132	118
1000 GB	181	183	152	130	117
2000 GB	284	289	258	220	189

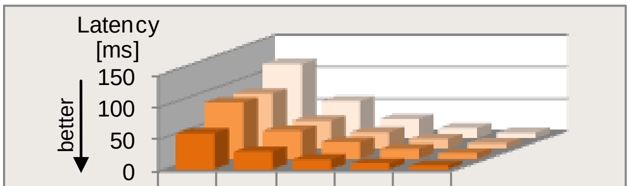

Outstanding IO	16	8	4	2	1
2000 GB	56.3	27.7	15.5	9.1	5.3
1000 GB	88.5	43.8	26.4	15.4	8.6
500 GB	86.2	43.4	25.8	15.1	8.5
250 GB	116.5	58.3	30.2	15.7	8.8

Standard load profile "Streaming" (sequential access, 100% read, 64 kB block size)

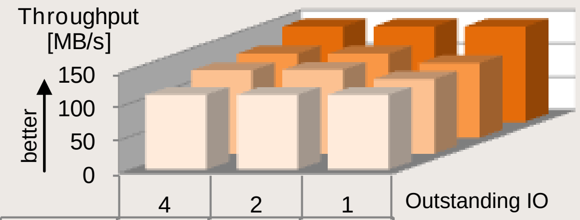

	4	2	1
250 GB	112.42	112.38	112.37
500 GB	124.83	125.86	112.01
1000 GB	126.28	126.16	112.07
2000 GB	143.15	143.11	143.07

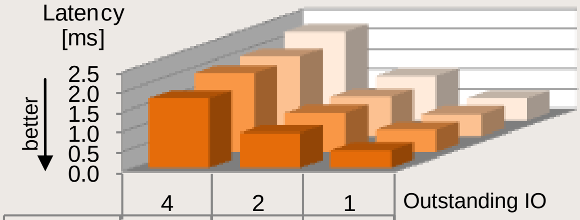

	4	2	1
2000 GB	1.7	0.9	0.4
1000 GB	2.0	1.0	0.6
500 GB	2.0	1.0	0.6
250 GB	2.2	1.1	0.6

Standard load profile "Restore" (sequential access, 100% write, 64 kB block size)

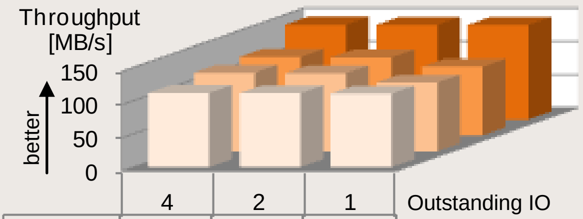

	4	2	1
250 GB	111.78	111.33	110.61
500 GB	118.33	117.48	103.36
1000 GB	118.14	117.27	103.30
2000 GB	142.78	142.57	142.53

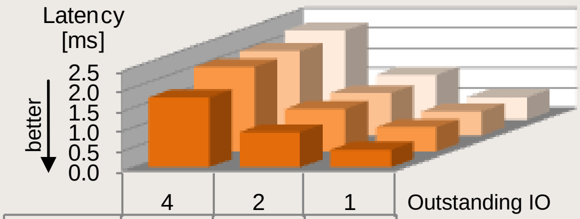

	4	2	1
2000 GB	1.8	0.9	0.4
1000 GB	2.1	1.1	0.6
500 GB	2.1	1.1	0.6
250 GB	2.2	1.1	0.6

The following ten diagrams show performance values (transactions, throughput rates and response times) for the standard load profiles "File server", "Database", "File copy", "Streaming" and "Restore" with increasing loads through the number of parallel requests (outstanding IO). The following actual hard disks are considered

EP HDD SAS 6 Gbit/s 3.5″ 15000 rpm 300 GB
EP HDD SAS 6 Gbit/s 3.5″ 15000 rpm 450 GB
EP HDD SAS 6 Gbit/s 3.5″ 15000 rpm 600 GB

Standard load profile "Database" (random access, 67% read, 33% write, 8 kB block size)

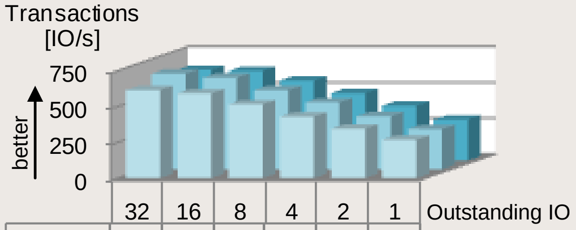

	32	16	8	4	2	1
300 GB	617	592	516	429	346	268
450 GB	669	641	553	468	371	281
600 GB	632	624	557	470	382	284

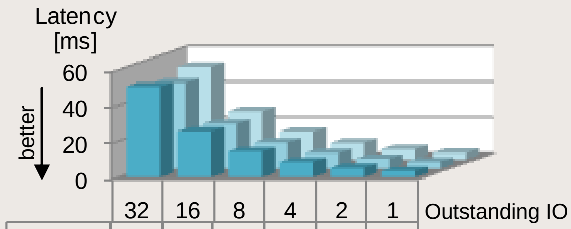

	32	16	8	4	2	1
600 GB	50.6	25.6	14.4	8.5	5.2	3.5
450 GB	47.9	25.0	14.5	8.5	5.4	3.6
300 GB	51.9	27.0	15.5	9.3	5.8	3.7

© Fujitsu Technology Solutions 2011

Standard load profile "File server" (random access, 67% read, 33% write, 64 kB block size)

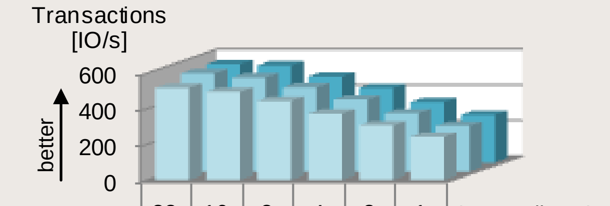

	32	16	8	4	2	1
300 GB	519	496	445	372	306	243
450 GB	550	528	472	403	325	254
600 GB	552	546	483	419	340	267

Outstanding IO

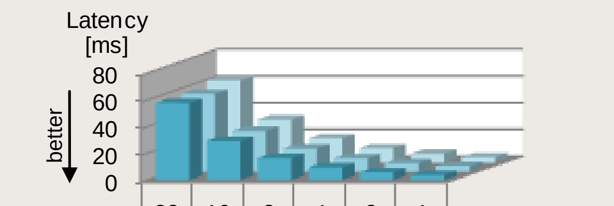

	32	16	8	4	2	1
600 GB	58.0	29.3	16.6	9.6	5.9	3.7
450 GB	58.1	30.3	17.0	9.9	6.1	3.9
300 GB	61.7	32.2	18.0	10.7	6.5	4.1

Outstanding IO

Standard load profile "File copy" (random access, 50% read, 50% write, 64 kB block size)

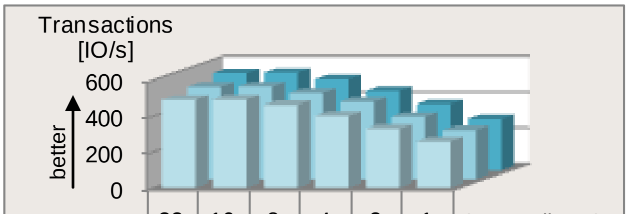

	32	16	8	4	2	1
300 GB	491	493	464	402	330	259
450 GB	516	518	481	432	348	274
600 GB	541	543	508	443	369	286

Outstanding IO

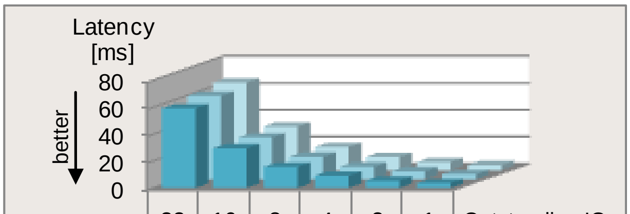

	32	16	8	4	2	1
600 GB	59.2	29.4	15.7	9.0	5.4	3.5
450 GB	62.0	30.9	16.6	9.3	5.7	3.7
300 GB	65.1	32.5	17.3	10.0	6.1	3.9

Outstanding IO

Standard load profile "Streaming" (sequential access, 100% read, 64 kB block size)

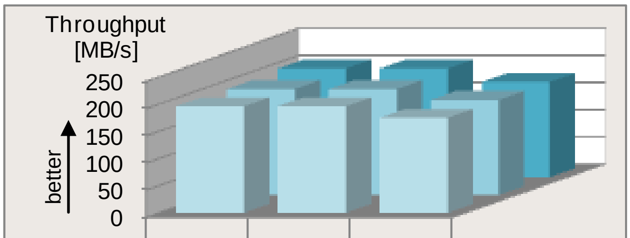

	4	2	1
300 GB	196.28	196.23	175.13
450 GB	195.18	195.12	175.13
600 GB	200.15	200.15	175.15

Outstanding IO

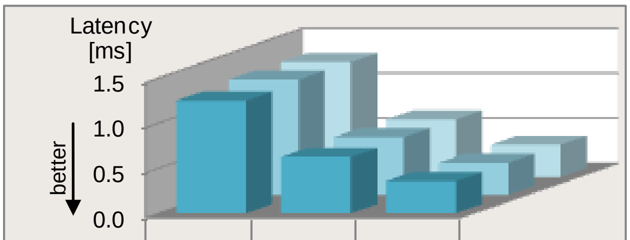

	4	2	1
600 GB	1.2	0.6	0.4
450 GB	1.3	0.6	0.4
300 GB	1.3	0.6	0.4

Outstanding IO

Standard load profile "Restore" (sequential access, 100% write, 64 kB block size)

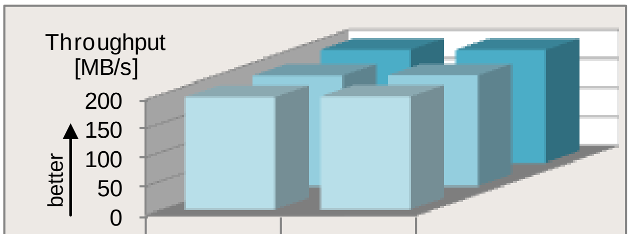

	2	1
300 GB	194.65	195.41
450 GB	191.28	191.62
600 GB	195.32	195.04

Outstanding IO

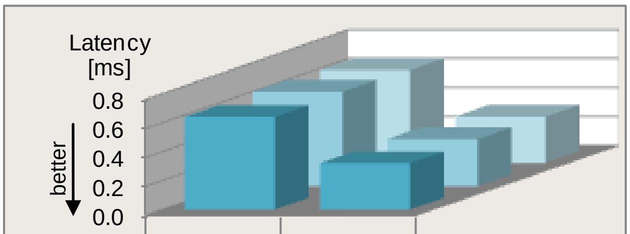

	2	1
600 GB	0.6	0.3
450 GB	0.7	0.3
300 GB	0.6	0.3

Outstanding IO

The following ten diagrams show performance values (transactions, throughput rates and response times) for the standard load profiles "File server", "Database", "File copy", "Streaming" and "Restore" with increasing loads through the number of parallel requests (outstanding IO). The following actual hard disks are considered

EP SSD SATA 3 Gbit/s 2.5″ 32 GB

EP SSD SATA 3 Gbit/s 2.5″ 64 GB

Standard load profile "Database" (random access, 67% read, 33% write, 8 kB block size)

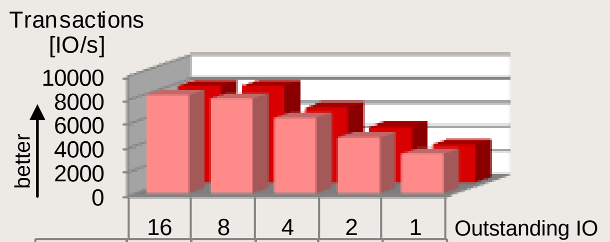

Outstanding IO	16	8	4	2	1
32 GB	8198	7898	6305	4671	3343
64 GB	8077	8083	6279	4633	3153

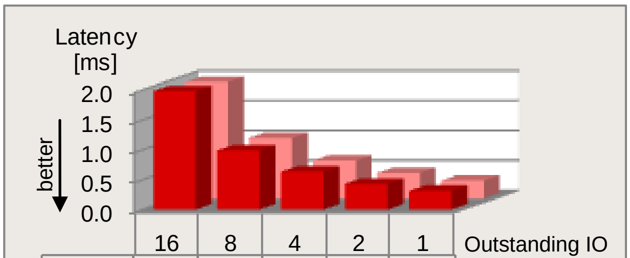

Outstanding IO	16	8	4	2	1
64 GB	2.0	1.0	0.6	0.4	0.3
32 GB	2.0	1.0	0.6	0.4	0.3

Standard load profile "File server" (random access, 67% read, 33% write, 64 kB block size)

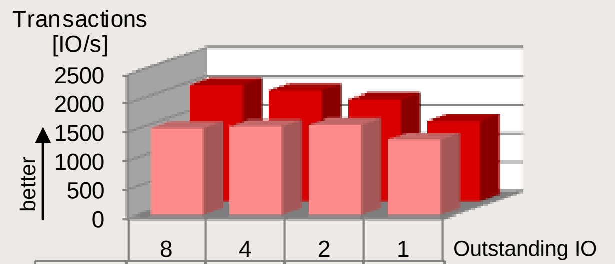

Outstanding IO	8	4	2	1
32 GB	1510	1543	1566	1307
64 GB	2023	1928	1770	1391

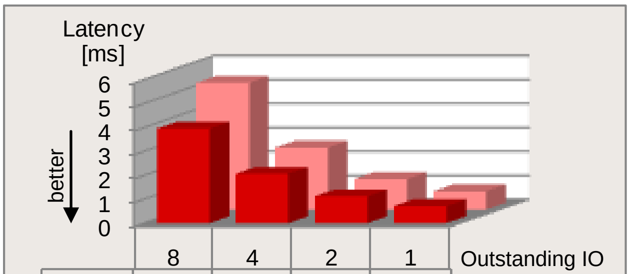

Outstanding IO	8	4	2	1
64 GB	4.0	2.1	1.1	0.7
32 GB	5.3	2.6	1.3	0.8

Standard load profile "File copy" (random access, 50% read, 50% write, 64 kB block size)

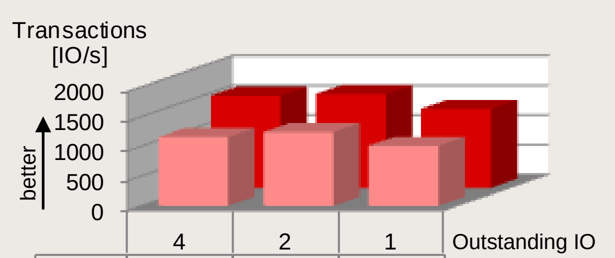

Outstanding IO	4	2	1
32 GB	1159	1245	1014
64 GB	1552	1580	1331

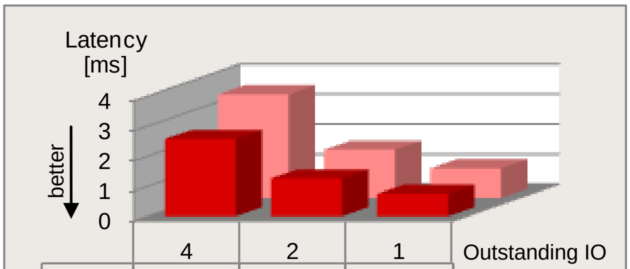

Outstanding IO	4	2	1
64 GB	2.6	1.3	0.8
32 GB	3.5	1.6	1.0

 © Fujitsu Technology Solutions 2011

Standard load profile "Streaming" (sequential access, 100% read, 64 kB block size)

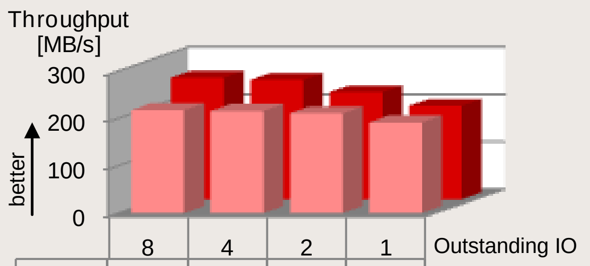

Outstanding IO	8	4	2	1
32 GB	217.72	216.02	211.07	191.74
64 GB	257.20	252.24	226.40	197.53

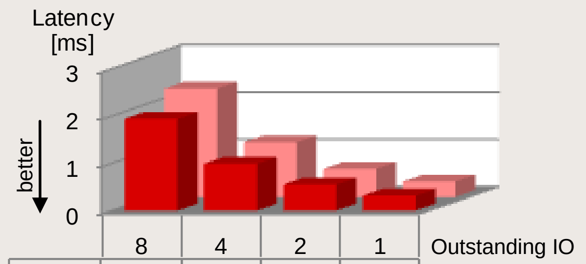

Outstanding IO	8	4	2	1
64 GB	1.9	1.0	0.6	0.3
32 GB	2.3	1.2	0.6	0.3

Standard load profile "Restore" (sequential access, 100% write, 64 kB block size)

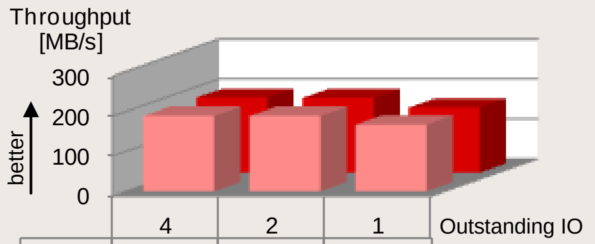

Outstanding IO	4	2	1
32 GB	189.19	189.16	167.17
64 GB	188.00	187.30	163.12

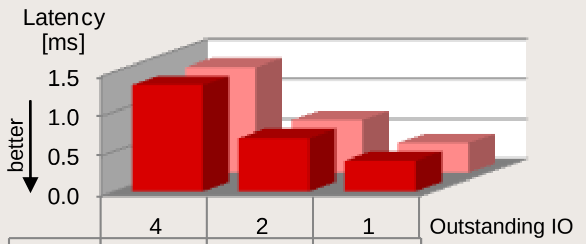

Outstanding IO	4	2	1
64 GB	1.3	0.7	0.4
32 GB	1.3	0.7	0.4

Conclusion of the review on capacity:

Of two hard disks that only differ in their capacity, the one with the higher capacity also often offers a somewhat higher performance. However, hard disks are not organized into RAID arrays just for security reasons. From a performance point of view, it is also best for the required total capacity to be distributed over as many hard disks as possible. In general, this also applies if the one with less capacity is the one with less performance. This performance disadvantage is more than compensated by the performance advantage associated with the larger number of hard disks in the RAID array.

Disk cache

All hard disks, even SSDs, have a cache.

The influence of the cache on performance is achieved by the buffering of data. Write requests are temporarily stored in the hard disk cache, because their processing is generally not as time-critical as that of read requests that the requesting application is waiting for. This also applies analogously for solid state drives.

The influence of the hard disk cache on disk I/O performance is quite high. Unfortunately, this is frequently seen as a security problem during a power failure and is therefore disabled. On the other hand, it was integrated by hard disk manufacturers for a good reason to increase the write performance. For performance reasons it is particularly advisable to enable the disk cache for the SATA hard disks, which rotate more slowly in comparison with SAS hard disks. The much larger cache for I/O accesses, and therefore a potential security risk for data loss during power failure, is already located in the main memory and administered by the operating system. To prevent data loss, you are recommended to equip the system with a UPS. The hard disk comparison measurements were in each case performed with and without a disk cache.

The following five diagrams deal with the influence of the hard disks cache for the standard load profiles "File server", "Database", "File copy", "Streaming" and "Restore" with increasing load intensity (outstanding IO) in an exemplary way for the hard disks:

ECO	HDD	SATA	3 Gbit/s	2.5″	5400 rpm	320 GB
BC	HDD	SATA	3 Gbit/s	3.5″	7200 rpm	2 TB
BC	HDD	SAS	6 Gbit/s	2.5″	10000 rpm	600 GB
EP	HDD	SAS	6 Gbit/s	3.5″	15000 rpm	600 GB
EP	SSD	SATA	3 Gbit/s	2.5″		64 GB

The respective performance growth of a hard disk with an enabled cache is shown in comparison with a hard disk with the cache disabled. The growth of the throughput corresponds to the growth of the transaction rate. The reduction in the average response time is proportional to the increase in the transaction rate. Measurement accuracy is approximately 95%.

Standard load profile "Database" (random access, 67% read, 33% write, 8 kB block size)

A 15% to 25% higher transaction rate can be achieved with an enabled cache with SATA-HDDs. With SAS-HDDs the gain in performance depends on the load intensity. With a medium load it is slightly less than with the SATA-HDDs. With low and very high loads the gain in performance is practically zero. The strongest effect is obtained when the cache is enabled for the SATA-SSD. The growth of the transaction rate increases by up to 175%.

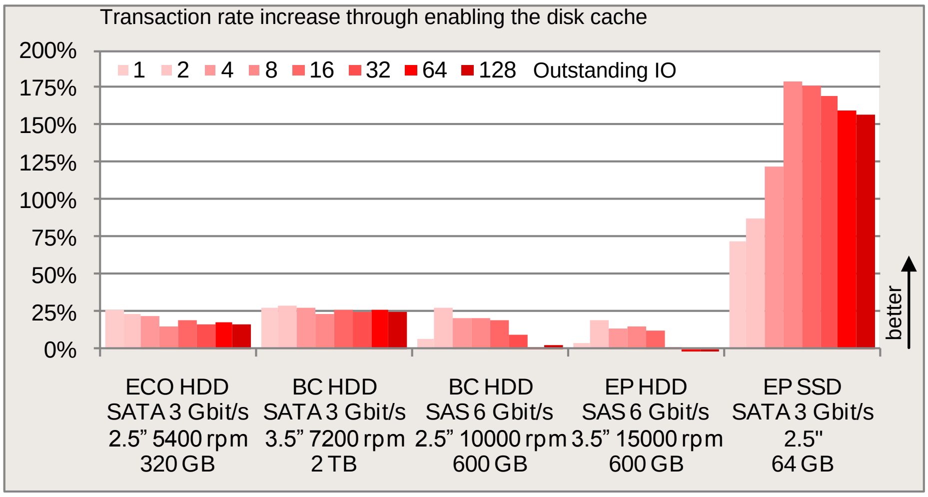

 © Fujitsu Technology Solutions 2011

Standard load profile "File server" (random access, 67% read, 33% write, 64 kB block size)

With an enabled cache it is possible to achieve an at least 20% higher transaction rate with the BC SATA-HDD. With SAS-HDDs and with the ECO SATA-HDD the gain in performance depends on the load intensity. This is mainly between 5% and 20%. The same performance values with an enabled cache can also only be achieved with a disabled cache under very high loads. The strongest effect is obtained when the cache is enabled for the SATA-SSD. The growth of the transaction rate increases by up to 50%.

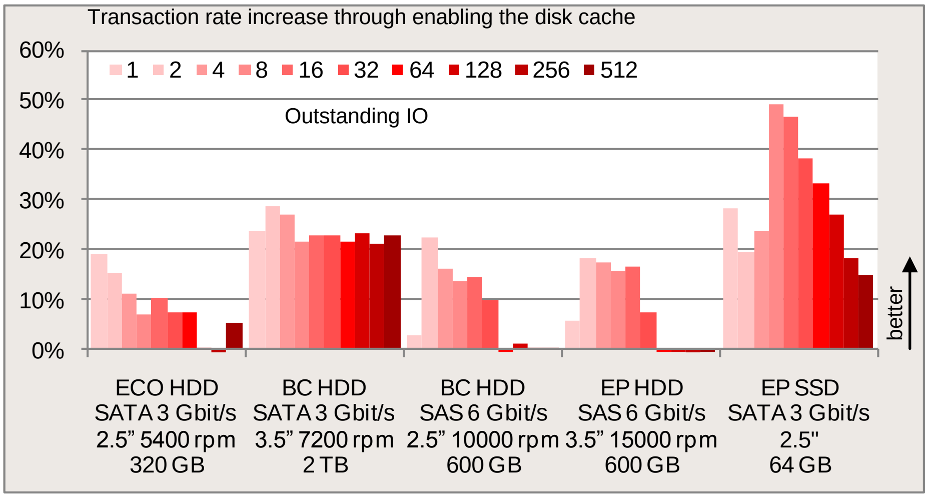

Standard load profile "File copy" (random access, 50% read, 50% write, 64 kB block size)

With an enabled cache a 13% to 27% higher transaction rate can be achieved with the ECO SATA-HDD. Growth is even between 31% and 53% with the BC SATA-HDD. With the SAS-HDDs the gain in performance depends on the load intensity. It increases to more than 30% and with very high loads sinks to practically zero. Performance growth is between 12% and 29% with the SATA-SSD.

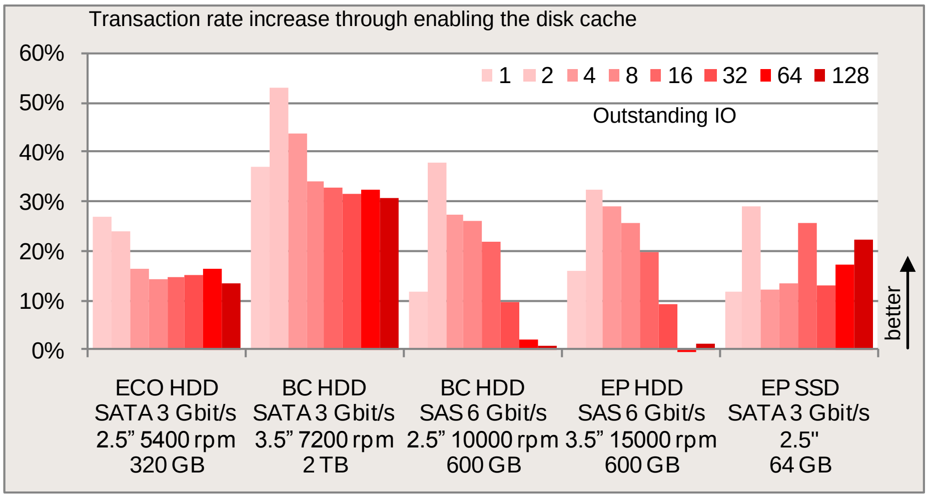

Standard load profile "Streaming" (sequential access, 100% read, 64 kB block size)

The use of the hard disk cache generally has no advantage here.

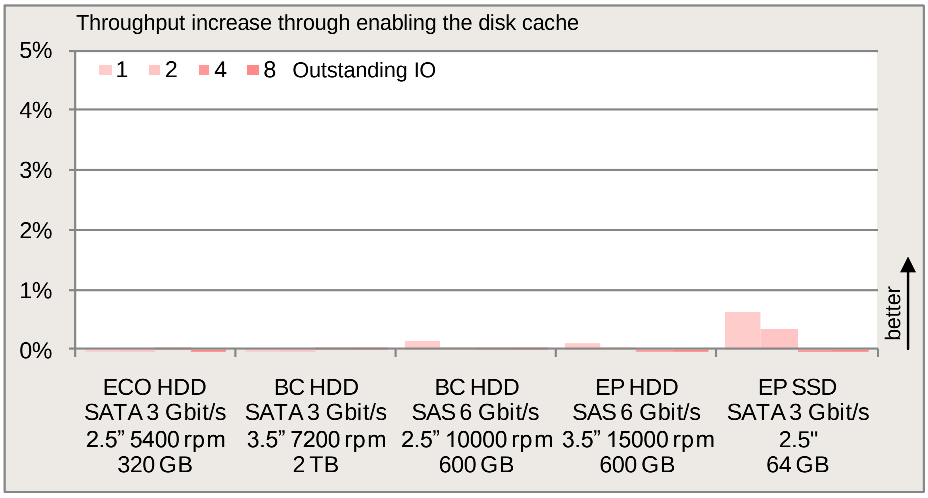

Standard load profile "Restore" (sequential access, 100% write, 64 kB block size)

By enabling the cache it is possible to achieve a fourteen-fold increase in throughput with the ECO SATA-HDD. And throughput is at least twice as high with higher loads. With the other HDDs it is possible to greatly increase the throughput by enabling the cache with minor loads. If the load increases, this high throughput is also achieved with a disabled cache. With the SATA-SSD an almost four-fold increase in throughput can be achieved by enabling the cache. And throughput is at least twice as high with higher loads.

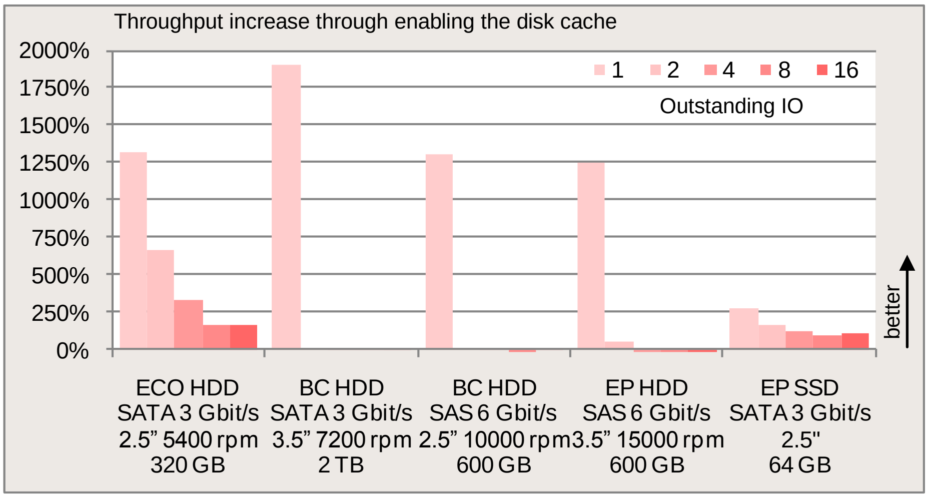

© Fujitsu Technology Solutions 2011

Conclusion

For PRIMERGY servers Fujitsu offers a large number of hard disks with different performance and availability features for all conceivable scenarios - from the small department server that is only needed during working hours up to the high-availability database server. The three defined hard disk classes Economic (ECO), Business-Critical (BC) and Enterprise (EP) provide orientation as to how the individual hard disk types should be classified.

- Hard disks of the Economic class (ECO) are the most cost-effective variant. These disks are chosen if performance and fail-safety do not have to meet the highest requirements. They should only be used in non-critical application areas, in which high transaction rates are not to be expected and in which 7×24-hour operation is not necessary. ECO hard disks have a SATA interface and work at a rotational speed of 5400 or 7200 rpm.
- Hard disks of the Business-Critical class (BC) are suited for 7×24-hour operation and thus meet the highest requirements for fail-safety. Due to their high capacities the lowest costs per GB are incurred here. BC hard disks either have a SATA interface and work at a rotational speed of 7200 rpm or have a SAS interface and work at a rotational speed of 10000 rpm.
- Hard disks of the Enterprise class (EP) offer the best performance and simultaneously meet the highest security requirements. They were specially developed for scenarios, in which the highest requirements for throughput also have to be met under high loads. The EP hard disks include ones with a SAS interface and a rotational speed of 15000 rpm as well as SATA-SSDs.

In addition to this classification, the hard disks differ in their technology (HDD or SSD), in the interface (SAS or SATA), in the form factor (2.5″ or 3.5″) and in the rotational speed and capacity. All of these factors play a role if the disk subsystem is to be sized with regard to capacity or performance.

In comparison with conventional hard disks, SSDs stand out for very low access times, mechanical robustness, low noise and very low energy consumption, but still have considerably lower capacity today and a higher price. Especially with random accesses, such as with database applications, SSDs offer a much higher transaction rate than conventional hard disks.

The SAS-HDDs have shorter access times and achieve higher throughputs due to the higher rotational speed of the SAS-HDDs (in comparison with the SATA-HDDs). Furthermore, the standard SAS-2.0 with the 6 Gbit/s interface has in the meantime established itself among the SAS-HDDs. Due to the lower rotational speed of the SATA-HDDs in comparison with the SAS-HDDs, the SATA-HDDs have higher access times and achieve lower throughputs. With SATA the hard disks today are for the most part still connected with the 3 Gbit/s interface. However, the SATA-HDDs offer high capacities right up in the terabyte range and at a very low cost. They are above all suited for use in secondary storage and backup systems.

More hard disks per server are possible thanks to the use of 2.5″ hard disks. Depending on the RAID level, this enables higher parallelism to be achieved due to the use of more hard disks in the RAID array and thus a better overall performance.

Depending on the required performance, a decision must be taken as to which hard disk type should be used with which rotational speed. Current SAS hard disks with 15000 rpm offer a better performance of between 20% and 30% than the current SAS-HDDs with 10000 rpm.

Of two hard disks that only differ in their capacity, the one with the higher capacity also often offers a somewhat higher performance. From a performance point of view, it is best for the required total capacity to be distributed over as many hard disks as possible. In general, this also applies if the one with less capacity is the one with less performance. This performance disadvantage is more than compensated by the performance advantage associated with the larger number of hard disks in the RAID array.

To ensure maximum performance it is advisable to enable the hard disk cache, especially with SATA-HDDs. When the hard disk cache is activated we recommend the use of a UPS.

© Fujitsu Technology Solutions 2011

Literature

PRIMERGY Systems

http://ts.fujitsu.com/primergy

PRIMERGY Performance

http://ts.fujitsu.com/products/standard_servers/primergy_bov.html

Basics of Disk I/O Performance

http://docs.ts.fujitsu.com/dl.aspx?id=65781a00-556f-4a98-90a7-7022feacc602

RAID Controller Performance

http://docs.ts.fujitsu.com/dl.aspx?id=ada7c1bf-74e3-4953-b783-839cdeec790b

RAID Performance

http://docs.ts.fujitsu.com/dl.aspx?id=c55404e9-69c6-4d1e-a556-6a322b27da5b

Information about Iometer

http://www.iometer.org

Contact

FUJITSU Technology Solutions

Website: http://ts.fujitsu.com

PRIMERGY Product Marketing

mailto:PRIMERGY-PM@ts.fujitsu.com

PRIMERGY Performance and Benchmarks

mailto:primergy.benchmark@ts.fujitsu.com

All rights reserved, including industrial property rights. Delivery subject to availability; right of technical modifications reserved. No liability or warranty assumed for completeness, validity and accuracy of the specified data and illustrations. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owners.
Further details are available under http://ts.fujitsu.com/terms_of_use.html

2011-08-31 WW EN Copyright © Fujitsu Technology Solutions GmbH 2011