

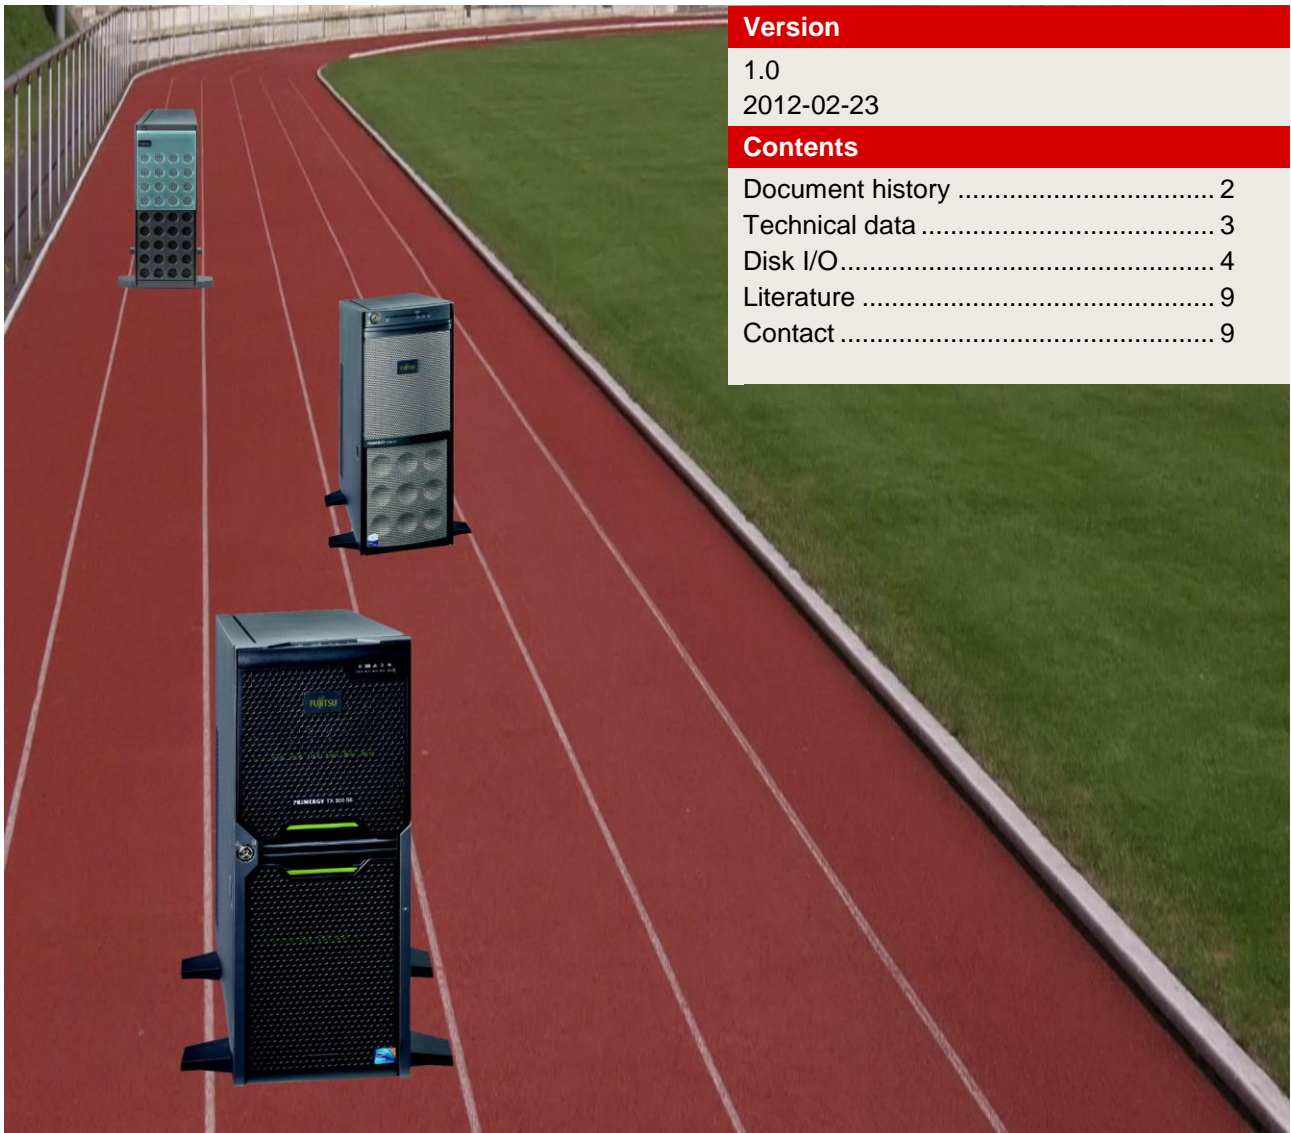
WHITE PAPER

FUJITSU PRIMERGY SERVERS

PERFORMANCE REPORT PRIMERGY SX960 S1

This document contains a summary of the benchmarks executed for the PRIMERGY SX960 S1.

In addition to the benchmark results, an explanation has been included for each benchmark and for the benchmark environment.



Document history

Version 1.0

- Disk I/O
Measurements with “RAID Ctrl SAS 5/6 512MB (D2616)” Controller

Technical data

The PRIMERGY SX960 S1 Storage Blade can be used to expand the internal disk capacity of a PRIMERGY Server Blade by up to 10 TB. It has room for ten 2.5" SATA-SSDs, SATA-HDDs or SAS-HDDs, which are run via one integrated SAS-RAID controller and can be used by one adjacent server blade. Up to two PRIMERGY SX960 S1 can be plugged into the PRIMERGY BX400 S1 Blade Server or in the PRIMERGY BX900 S1 Blade Server. Each one occupies two slots in the Blade Server - the slots are on the right next to the appropriate server blade.



Detailed technical information is available in the

- [data sheet PRIMERGY SX960 S1](#)
- [data sheet PRIMERGY BX900 S1](#)
- [data sheet PRIMERGY BX400 S1](#)

Disk I/O

Benchmark description

Performance measurements of disk subsystems for PRIMERGY servers are used to assess their performance and enable a comparison of the different storage connections for PRIMERGY servers. As standard, these performance measurements are carried out with a defined measurement method, which models the hard disk accesses of real application scenarios on the basis of specifications.

The essential specifications are:

- Share of random accesses / sequential accesses
- Share of read / write access types
- Block size (kB)
- Number of parallel accesses (# of outstanding I/Os)

A given value combination of these specifications is known as "load profile". The following five standard load profiles can be allocated to typical application scenarios:

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

In order to model applications that access in parallel with a different load intensity, the "# of Outstanding I/Os" is increased, starting with 1, 3, 8 and going up to 512 (from 8 onwards in increments to the power of two).

The measurements of this document are based on these standard load profiles.

The main results of a measurement are:

- 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

The throughput rate has established itself as the normal measurement variable for sequential load profiles, whereas the measurement variable "transaction rate" is mostly used for random load profiles with their small block sizes. Data throughput and transaction rate are directly proportional to each other and can be transferred to each other according to the formula

<i>Data throughput [MB/s]</i>	$= \text{Transaction rate [IO/s]} \times \text{Block size [MB]}$
<i>Transaction rate [IO/s]</i>	$= \text{Data throughput [MB/s]} / \text{Block size [MB]}$

All the details of the measurement method and the basics of disk I/O performance are described in the white paper "[Basics of Disk I/O Performance](#)".

Benchmark environment

All results were determined by way of example on a PRIMERGY BX920 S2.

System Under Test (SUT)	
Hardware (Shared)	
Enclosure	PRIMERGY BX400 S1
Hardware	
Model	PRIMERGY BX920 S2
Disk subsystem	PRIMERGY SX960 S1
Controller	1 x RAID Ctrl SAS 6G 5/6 512 MB (D2616)
Drive	EP HDD SAS 6 Gbit/s 2.5" 15000 rpm 146 GB EP SSD SATA 3 Gbit/s 2.5" 64 GB
Software	
Operating system	Microsoft Windows Server 2008 R2 Enterprise
Administration software	ServerView RAID Manager 5.0.2
Initialization of RAID arrays	RAID arrays are initialized before the measurement with an elementary block size of 64 kB ("stripe size")
File system	NTFS
Measuring tool	Iometer 27.07.2006
Measurement data	Measurement files of 32 GB with 1 – 8 hard disks; 64 GB with 9 – 16 hard disks

Some components may not be available in all countries / sales regions.

Benchmark results

This section is designed to help you choose the right solution from the various configuration options of the PRIMERGY SX960 S1 in the light of disk-I/O performance. The determining factors here are the suitable components and their correct parameter settings. These two aspects should therefore be dealt with as preparation for the discussion of the performance values.

Components

The hard disks are the first essential component. If there is a reference below to "hard disks", this is meant as the generic term for HDDs ("hard disk drives", in other words conventional hard disks) and SSDs ("solid state drives", i.e. non-volatile electronic storage media). When selecting the type of hard disk and number of hard disks you can move the weighting in the direction of storage capacity, performance, security or price. In order to enable a pre-selection of the hard disk types – depending on the required weighting – the hard disk types for PRIMERGY servers are divided into three classes:

- "Economic" (ECO): low-priced hard disks
- "Business Critical" (BC): very failsafe hard disks
- "Enterprise" (EP): very failsafe and very high-performance hard disks

The following table is a list of the hard disk types that have been available for the PRIMERGY SX960 S1 since system release.

Drive class	Data medium type	Interface	Form factor	krpm
Economic	HDD	SATA 3G	2.5"	5.4
Business Critical	HDD	SATA 3G	2.5"	7.2
Enterprise	HDD	SAS 3G/6G	2.5"	10,15
Enterprise	SSD	SATA 3G	2.5"	-

Mixed drive configurations of SAS and SATA hard disks in one system are permitted, unless they are excluded in the configurator for special hard disk types.

The SATA-HDDs offer high capacities right up into the terabyte range at a very low cost. 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). SAS-HDDs with a rotational speed of 15 krpm have better access times and throughputs than comparable HDDs with a rotational speed of 10 krpm. And the standard 6G has in the meantime established itself among the SAS-HDDs.

Of all the hard disk types SSDs offer on the one hand by far the highest transaction rates for random load profiles, and on the other hand the shortest access times. In return, however, the price per gigabyte of storage capacity is substantially higher.

More detailed performance statements about hard disk types are available in the white paper "[Single Disk Performance](#)".

After the hard disks the RAID controller is the second performance-determining key component.

The following table summarizes the most important features of the available RAID controllers of the system. A short alias is specified here for each controller, which is used in the subsequent list of the performance values.

Controller name	Alias	Cache	Supported interfaces	Max. # disks in the system	RAID levels	BBU/FBU	
RAID Ctrl SAS 6G 5/6 512 MB (D2616)	LSI2108	512 MB	SATA 3G/6G SAS 3G/6G	PCIe 2.0 x8	10 × 2.5"	0, 1, 5, 6, 10, 50, 60	✓/✓

System-specific interfaces

The interfaces of a controller to the motherboard and to the hard disks have in each case specific maximum values for data throughput. These maximum values are listed in the following table. The minimum of these two values is a definite limit, which cannot be exceeded when using the respective controller in the system. This value is highlighted in bold in the following table.

Controller alias	Effective in the system				Maximum throughput of PCIe interface	Expander in the system?
	# Disk channels	Maximum throughput of disk interface	PCIe version	PCIe width		
LSI2108	8 x SAS 6G	3890 MB/s	2.0	x4	1716 MB/s	✓

An expander makes it possible to connect more hard disks in a system than the SAS channels that the controller has. An expander cannot increase the possible maximum throughput of a controller, but makes the overall throughput of all connected hard disks available. More details about the RAID controllers of the PRIMERGY systems are available in the white paper "[RAID Controller Performance](#)".

Specific to the PRIMERGY SX960 S1 is the fact that (in contrast to the PRIMERGY SX980 S1) the RAID controller is in the storage blade (and not in the server blade). As a result, this storage blade has the same kind of connection from the controller to the hard disks as a server with internal hard disks and can from the viewpoint of disk-I/O performance be treated identically.

Settings

In most cases, the cache of the hard disks has a great influence on disk-I/O performance. It is frequently regarded as a security problem in case of power failure and is thus switched off. On the other hand, it was integrated by hard disk manufacturers for the good reason of increasing the write performance. For performance reasons it is therefore advisable to enable the hard disk cache. This is particular valid for SATA-HDDs. The performance can as a result increase more than tenfold for specific access patterns and hard disk types. More information about the performance impact of the hard disk cache is available in the document "[Single Disk Performance](#)". To prevent data loss in case of power failure you are recommended to equip the system with a UPS.

In the case of controllers with a cache there are several parameters that can be set. The optimal settings can depend on the RAID level, the application scenario and the type of data medium. In the case of RAID levels 5 and 6 in particular (and the more complex RAID level combinations 50 and 60) it is obligatory to enable the controller cache for application scenarios with write share. If the controller cache is enabled, it is advisable to have a BBU or FBU to safeguard the data temporarily stored in the cache against loss in case of power failure. More information about the setting options of the controller cache is available in the white paper "[RAID Controller Performance](#)".

Performance values

In general, disk-I/O performance of a RAID array depends on the type and number of hard disks, on the RAID level and on the RAID controller. If the limits of the [system-specific interfaces](#) are not exceeded, the statements on disk-I/O performance are therefore valid for all PRIMERGY systems. This is why all the performance statements of the document "[RAID Controller Performance](#)" also apply for the PRIMERGY SX960 S1 if the configurations measured there are also supported by this system.

The following list now concentrates on the maximum achievable performance values of the system. These values are listed in table form, specifically for different RAID levels, access types and block sizes. The values can be achieved with hard disks that have been selected for optimal performance (the components used are described in more detail in the section [Benchmark environment](#)). Furthermore, cache settings of controllers and hard disks, which are optimal for the respective access scenario and the RAID level, are used as a basis. More information about the selection of the RAID level and the cache settings for an access scenario is available in the document "[RAID Performance](#)".

The first table deals with random accesses.

Configuration version		RAID level	HDDs random 8 kB blocks 67% read [IO/s]	HDDs random 64 kB blocks 67% read [IO/s]	SSDs random 8 kB blocks 67% read [IO/s]	SSDs random 64 kB blocks 67% read [IO/s]
RAID Controller	#Disks					
1 x LSI2108	2	RAID 1	767	400	4855	2033
	10	RAID 5	3178	1671	13644	5031
	10	RAID 10	4311	2228	16859	8716
	10	RAID 0	5704	2899	24760	10460

The next table deals with sequential accesses.

Configuration version		RAID level	HDDs sequential 64 kB blocks 100% read [MB/s]	HDDs sequential 64 kB blocks 100% write [MB/s]	SSDs sequential 64 kB blocks 100% read [MB/s]	SSDs sequential 64 kB blocks 100% write [MB/s]
RAID Controller	#Disks					
1 x LSI2108	2	RAID 1	287	151	368	193
	10	RAID 5	1341	1177	1227	1002
	10	RAID 10	1148	756	1254	625
	10	RAID 0	1449	1472	1511	1151

At full configuration with powerful hard disks (configured as RAID 0) the PRIMERGY SX960 S1 achieves a throughput of up to 1511 MB/s for sequential load profiles and a transaction rate of up to 24760 IO/s for typical, random application scenarios.

Literature

PRIMERGY Systems

<http://primergy.com/>

PRIMERGY SX960 S1

Data sheet

<http://docs.ts.fujitsu.com/dl.aspx?id=489d2a0c-81f1-4d3e-aaa3-df55ac446930>

PRIMERGY BX900 S1

Data sheet

<http://docs.ts.fujitsu.com/dl.aspx?id=0a5dcae5-f5a2-42dc-9039-7f887182bc5e>

PRIMERGY BX400 S1

Data sheet

<http://docs.ts.fujitsu.com/dl.aspx?id=fab1fbb5-5d63-4b2d-816d-3def6c7cb6da>

PRIMERGY Performance

<http://www.fujitsu.com/fts/products/computing/servers/primergy/benchmarks/>

Disk I/O

Basics of Disk I/O Performance

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

Single Disk Performance

<http://docs.ts.fujitsu.com/dl.aspx?id=0e30cb69-44db-4cd5-92a7-d38bacec6a99>

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

Website: <http://www.fujitsu.com/>

PRIMERGY Product Marketing

<mailto:Primergy-PM@ts.fujitsu.com>

PRIMERGY Performance and Benchmarks

<mailto:primergy.benchmark@ts.fujitsu.com>