THE POSSIBILITIES ARE INFINITE

WHITE PAPER

Performance Report PRIMERGY RX100 S6

Version 2.0
March 2010

Pages 31

Abstract

This document contains a summary of the benchmarks executed for the PRIMERGY RX100 S6.

The PRIMERGY RX100 S6 performance data are compared with the data of other PRIMERGY models and discussed. In addition to the benchmark results, an explanation has been included for each benchmark and for the benchmark environment.

Contents

Document history

Version 1.0

First report version including the benchmark chapters

- SPECcpu2006
  Measurements with
  - Pentium G6950
  - Core i3-540
  - Xeon X3430, X3440, X3450, X3460 and X3470
- SPECjbb2005
  Measurement with Xeon X3470
- SPECpower_ssj2008
  Measurement with Intel Xeon X3470 and 1 x SATA 2.5“ 5400rpm
- StorageBench
  Measurements with Onboard SATA Controller
  “RAID 0/1 SAS based on LSI MegaRAID 4Port” Controller
  “RAID Ctrl SAS 5/6 512MB (D2616)” Controller
- OLTP-2
  Measurement with
  - Celeron G1101
  - Pentium G6950
  - Core i3-530 and i3-540
  - Xeon L3426, X3430, X3440, X3450, X3460 and X3470

Version 2.0

Updated benchmark chapters:

- SPECcpu2006
  Measurements with
  - Celeron G1101
  - Core i3-530
  - Xeon L3406 and L3426

Technical Data

Like its predecessor the PRIMERGY RX100 S5, the PRIMERGY RX100 S6 is a mono socket rack server with one height unit. It includes the Intel 3420 chip set, an Intel Celeron, Pentium, Core i3 or Xeon processor, up to 32 GB memory, depending on the processor used a 1067 MHz or 1333 MHz bus, an Intel 82574 and an Intel 82578 1-GBit LAN controller and three PCI slots (1 x PCI-Express Gen2 x4, 2 x PCI-Express Gen2 x8).

The PRIMERGY RX100 S6 is offered in two version with either up to two integrated 3.5“ hard disks or up to four integrated 2.5“ hard disks. A 6-port SATA controller with RAID 0, 1 and 10 functionality is available for the use of SATA hard disks. Either an 8-port SAS controller with RAID 0, 1 and RAID 1E functionality or an 8-port SAS controller with RAID 0, 1, 10, 5, 50, 6 and RAID 60 functionality can be used for SAS hard disks.




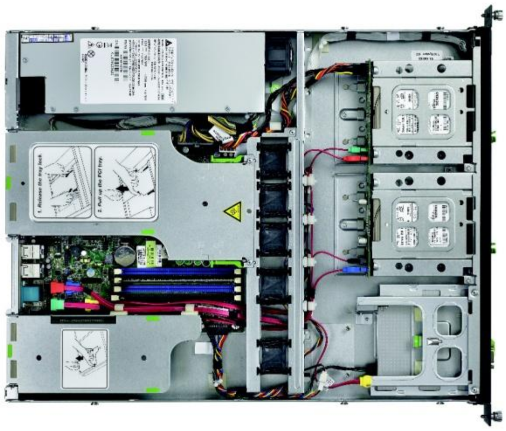

See Data sheet PRIMERGY RX100 S6 for detailed technical information.

SPECcpu2006*

Benchmark description

SPECcpu2006 is a benchmark to measure system efficiency during integer and floating point operations. It consists of an integer test suite containing 12 applications and a floating point test suite containing 17 applications which are extremely computing-intensive and concentrate on the CPU and memory. Other components, such as disk I/O and network, are not measured by this benchmark.

SPECcpu2006 is not bound to a specific operating system. The benchmark is available as source code and is compiled before the actual benchmark. Therefore, the compiler version used and its optimization settings have an influence on the measurement result.

SPECcpu2006 contains two different methods of performance measurement: The first method (SPECint2006 and SPECfp2006) determines the time required to complete a single task. The second method (SPECint_rate2006 and SPECfp_rate2006) determines the throughput, i.e. how many tasks can be completed in parallel. Both methods are additionally subdivided into two measuring runs, "base" and "peak", which differ in the way the compiler optimization is used. The "base" values are always used when results are published, the "peak" values are optional.

Benchmark	Arithmetic	Type	Compiler optimization	Measuring result	Application
SPECint2006	integer	peak	aggressive	speed	single threaded
SPECint_base2006	integer	base	conservative	speed	single threaded
SPECint_rate2006	integer	peak	aggressive	throughput	multithreaded
SPECint_rate_base2006	integer	base	conservative	throughput	multithreaded
SPECfp2006	floating point	peak	aggressive	speed	single threaded
SPECfp_base2006	floating point	base	conservative	speed	single threaded
SPECfp_rate2006	floating point	peak	aggressive	throughput	multithreaded
SPECfp_rate_base2006	floating point	base	conservative	throughput	multithreaded

The results represent the geometric mean of normalized ratios determined for the individual benchmarks. Compared with the arithmetic mean, the geometric mean results in the event of differingly high single results in a weighting in favor of the lower single results. "Normalized" means measuring how fast the test system runs in comparison to a reference system. The value of "1" was determined for the SPECint_base2006, SPECint_rate_base2006, SPECfp_base2006 and SPECfp_rate_base2006 results of the reference system. Thus a SPECint_base2006 value of 2 means for example that the measuring system has executed this benchmark approximately twice as fast as the reference system. A SPECfp_rate_base2006 value of 4 means that the measuring system has executed this benchmark about 4/[# base copies] times as fast as the reference system. "# base copies" here specifies how many parallel instances of the benchmark have been executed.

We do not submit all SPECcpu2006 measurements for publication at SPEC. So not all results appear on SPEC's web sites. As we archive the log data for all measurements, we are able to prove the correct realization of the measurements any time.

* SPEC®, SPECint®, SPECfp® and the SPEC logo are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).

Benchmark results

The PRIMERGY RX100 S6 was measured with four different processor versions:

- Celeron G1101
- Pentium G6950
- Core i3-530 and i3-540
- Xeon L3406, L3426, X3430, X3440, X3450, X3460 and X3470

On January 5, 2010 the PRIMERGY RX100 S6 achieved first place in the 1-socket server category in the benchmarks SPECint_base2006 and SPECint2006.[1]

On January 6, 2010 the PRIMERGY RX100 S6 achieved first place in the 1-socket server category in the benchmarks SPECint_rate_base2006 and SPECint_rate2006.[2]

The following tables show results, in which all benchmark programs were compiled with the Intel C++/Fortran compiler 11.1 and run under SUSE Linux Enterprise Server 11 (64-bit). Bold values are published at http://www.spec.org.

[1] Competitive benchmark results stated above reflect results published as of January 5, 2010. The comparison presented above is based on the best performing 1-socket servers. For the latest SPECint_base2006 and SPECint2006 benchmark results, visit http://www.spec.org/cpu2006/results.

[2] Competitive benchmark results stated above reflect results published as of January 6, 2010. The comparison presented above is based on the best performing 1-socket servers. For the latest SPECint_rate_base2006 and SPECint_rate2006 benchmark results, visit http://www.spec.org/cpu2006/results.

Processor	Cores	GHz	L3 cache	FSB	TDP	SPECint_base2006	SPECint2006
Celeron G1101	2	2.27	2 MB	1067 MHz	73 watt	**18.5**	**20.1**
Pentium G6950	2	2.80	3 MB	1067 MHz	73 watt	**22.4**	**24.4**
Core i3-530	2	2.93	4 MB	1333 MHz	73 watt	**25.5**	**27.8**
Core i3-540	2	3.07	4 MB	1333 MHz	73 watt	**26.3**	**28.7**
Xeon L3406	2	2.27	4 MB	1333 MHz	30 watt	21.6	24.0
Xeon L3426	4	1.87	8 MB	1333 MHz	45 watt	**28.2**	**32.8**
Xeon X3430	4	2.40	8 MB	1333 MHz	95 watt	**27.4**	**30.0**
Xeon X3440	4	2.53	8 MB	1333 MHz	95 watt	**28.5**	**31.0**
Xeon X3450	4	2.67	8 MB	1333 MHz	95 watt	**30.3**	**33.3**
Xeon X3460	4	2.80	8 MB	1333 MHz	95 watt	**32.1**	**35.4**
Xeon X3470	4	2.93	8 MB	1333 MHz	95 watt	**33.4**	**36.8**

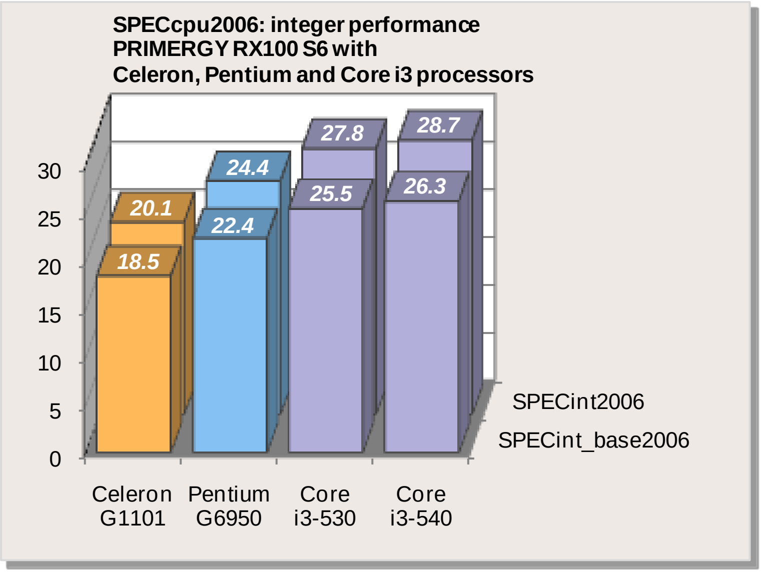

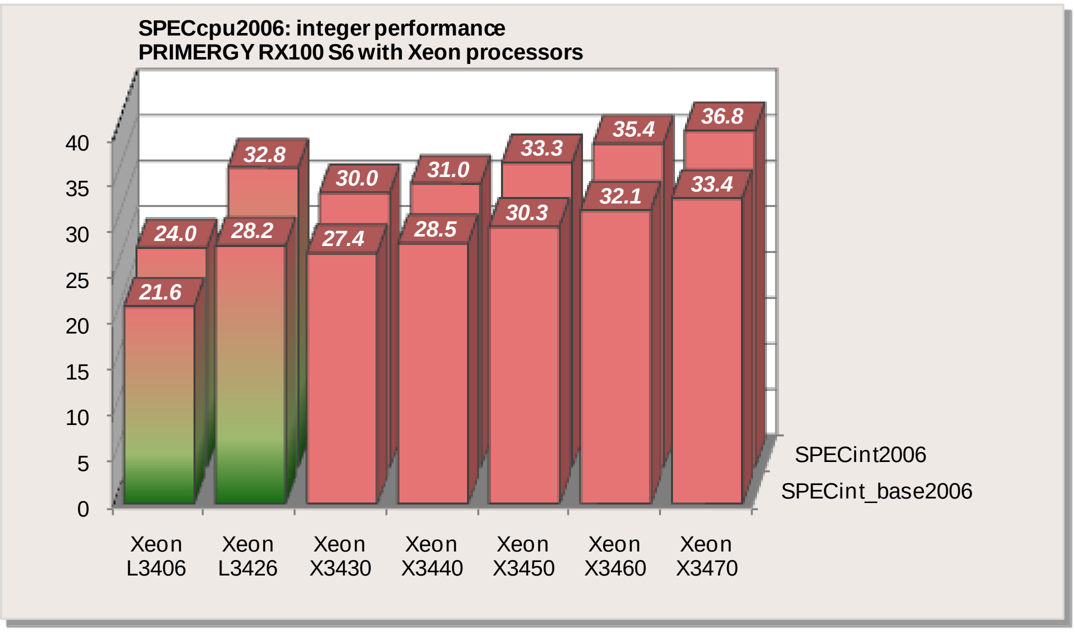

Processor	Cores	GHz	L3 cache	FSB	TDP	SPECint_rate_base2006	SPECint_rate2006
Celeron G1101	2	2.27	2 MB	1067 MHz	73 watt	**32.3**	**35.0**
Pentium G6950	2	2.80	3 MB	1067 MHz	73 watt	**39.3**	**42.6**
Core i3-530	2	2.93	4 MB	1333 MHz	73 watt	**56.7**	**60.1**
Core i3-540	2	3.07	4 MB	1333 MHz	73 watt	**58.4**	**61.8**
Xeon L3406	2	2.27	4 MB	1333 MHz	30 watt	45.4	48.3
Xeon L3426	4	1.87	8 MB	1333 MHz	45 watt	**82.9**	**89.0**
Xeon X3430	4	2.40	8 MB	1333 MHz	95 watt	**85.7**	**92.3**
Xeon X3440	4	2.53	8 MB	1333 MHz	95 watt	**105**	**112**
Xeon X3450	4	2.67	8 MB	1333 MHz	95 watt	**109**	**115**
Xeon X3460	4	2.80	8 MB	1333 MHz	95 watt	**114**	**120**
Xeon X3470	4	2.93	8 MB	1333 MHz	95 watt	**120**	**127**

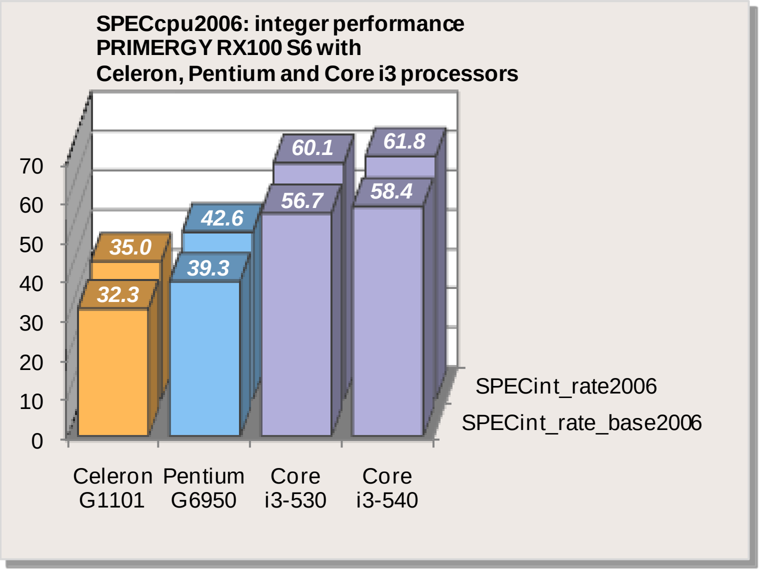

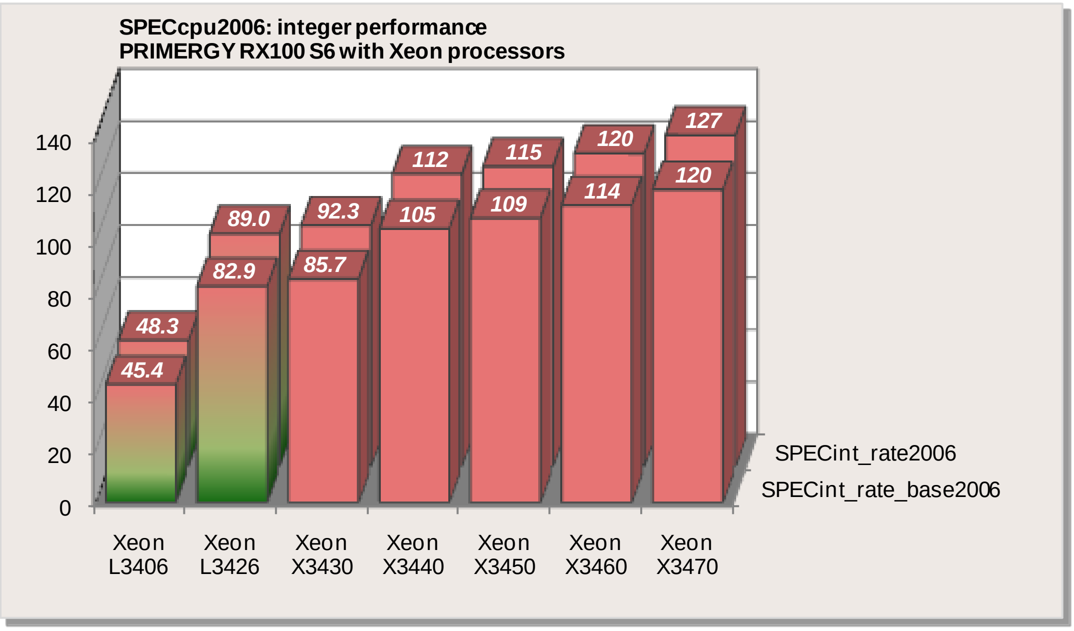

Processor	Cores	GHz	L3 cache	FSB	TDP	SPECfp_base2006	SPECfp2006
Celeron G1101	2	2.27	2 MB	1067 MHz	73 watt	**21.8**	**22.7**
Pentium G6950	2	2.80	3 MB	1067 MHz	73 watt	**25.4**	**26.5**
Core i3-530	2	2.93	4 MB	1333 MHz	73 watt	**28.4**	**29.9**
Core i3-540	2	3.07	4 MB	1333 MHz	73 watt	**29.3**	**30.7**
Xeon L3406	2	2.27	4 MB	1333 MHz	30 watt	24.0	25.5
Xeon L3426	4	1.87	8 MB	1333 MHz	45 watt	**32.2**	**35.5**
Xeon X3430	4	2.40	8 MB	1333 MHz	95 watt	**31.7**	**33.8**
Xeon X3440	4	2.53	8 MB	1333 MHz	95 watt	**32.8**	**34.9**
Xeon X3450	4	2.67	8 MB	1333 MHz	95 watt	**34.7**	**37.0**
Xeon X3460	4	2.80	8 MB	1333 MHz	95 watt	**36.4**	**39.0**
Xeon X3470	4	2.93	8 MB	1333 MHz	95 watt	**37.6**	**39.9**

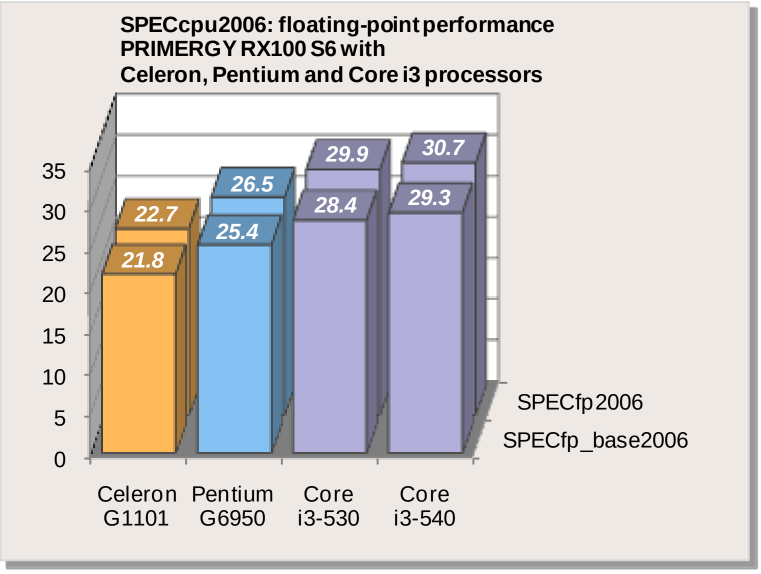

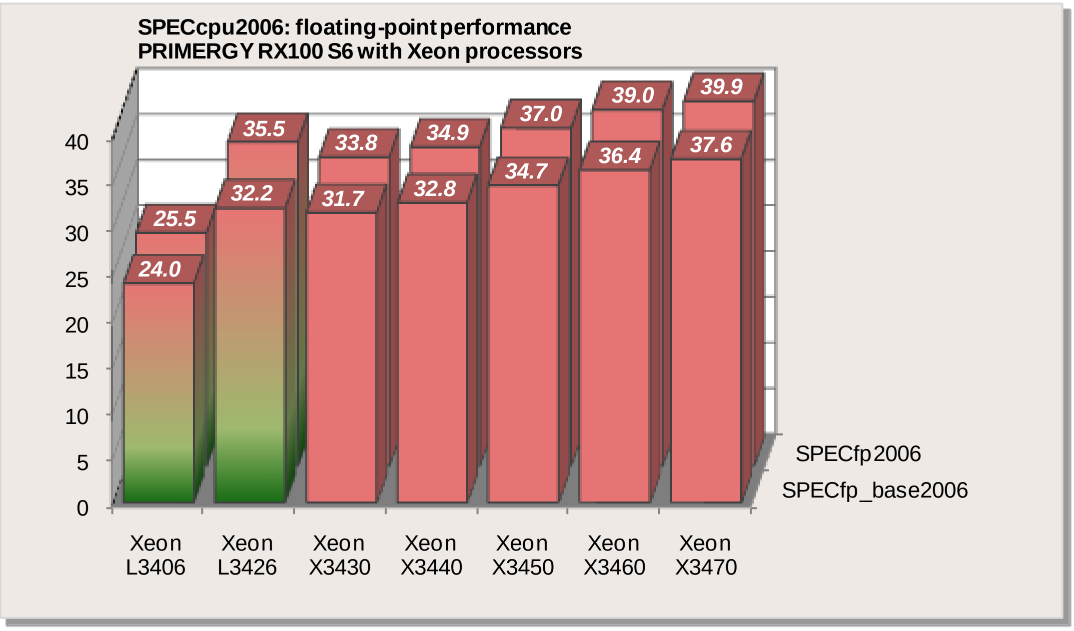

Processor	Cores	GHz	L3 cache	FSB	TDP	SPECfp_rate_base2006	SPECfp_rate2006
Celeron G1101	2	2.27	2 MB	1067 MHz	73 watt	**34.1**	**35.4**
Pentium G6950	2	2.80	3 MB	1067 MHz	73 watt	**39.1**	**40.7**
Core i3-530	2	2.93	4 MB	1333 MHz	73 watt	**48.5**	**50.5**
Core i3-540	2	3.07	4 MB	1333 MHz	73 watt	**49.5**	**51.7**
Xeon L3406	2	2.27	4 MB	1333 MHz	30 watt	39.8	41.3
Xeon L3426	4	1.87	8 MB	1333 MHz	45 watt	**68.4**	**70.8**
Xeon X3430	4	2.40	8 MB	1333 MHz	95 watt	**73.2**	**75.7**
Xeon X3440	4	2.53	8 MB	1333 MHz	95 watt	**80.5**	**83.2**
Xeon X3450	4	2.67	8 MB	1333 MHz	95 watt	**82.1**	**85.0**
Xeon X3460	4	2.80	8 MB	1333 MHz	95 watt	**84.6**	**87.6**
Xeon X3470	4	2.93	8 MB	1333 MHz	95 watt	**87.2**	**90.4**

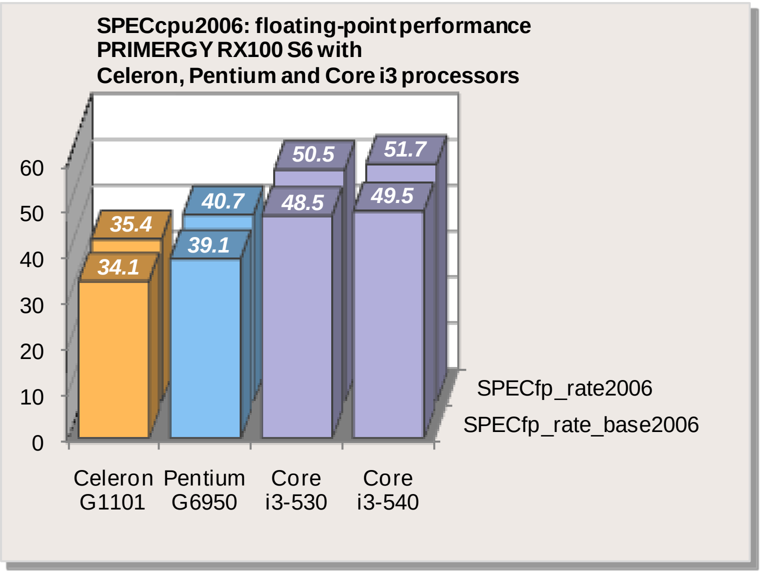

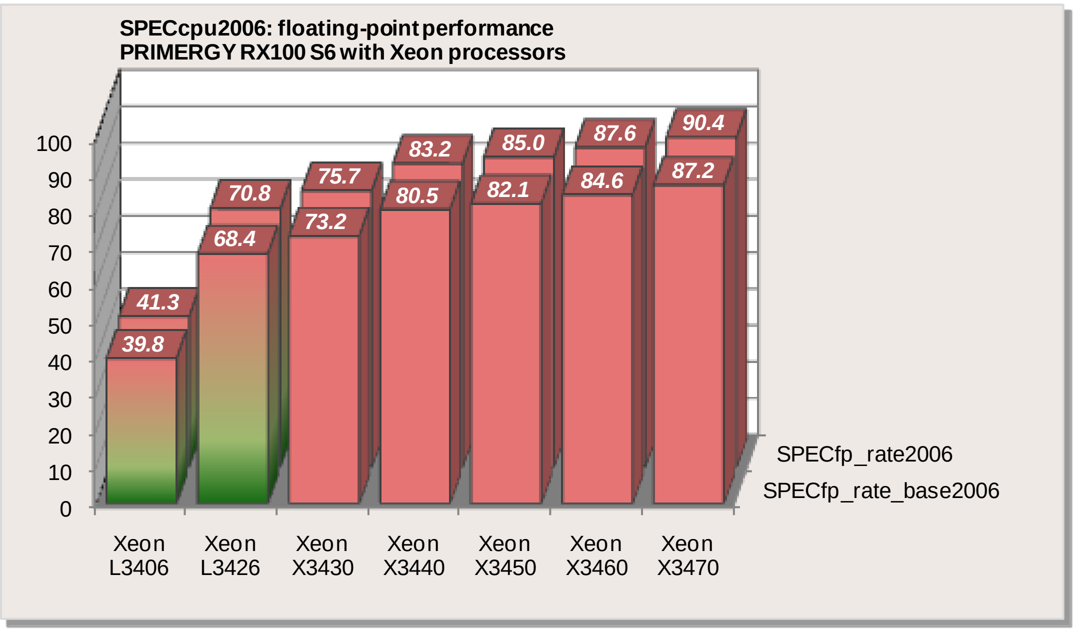

The diagrams below illustrate the performance of the PRIMERGY RX100 S6 compared with its predecessor PRIMERGY RX100 S5, both in their highest performance configurations.

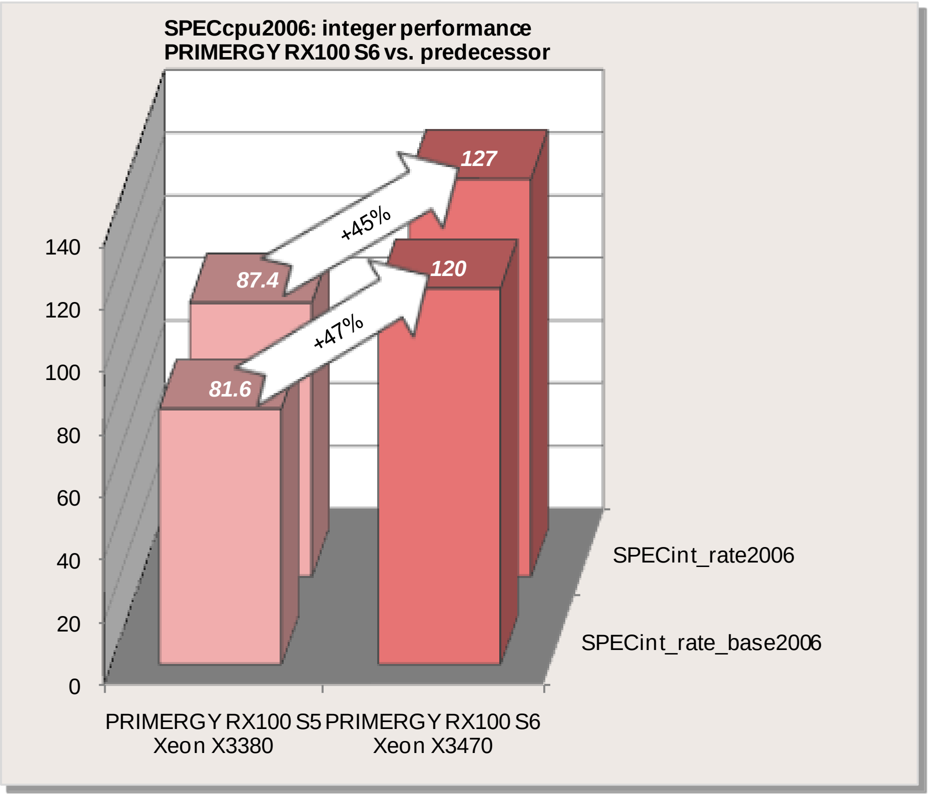

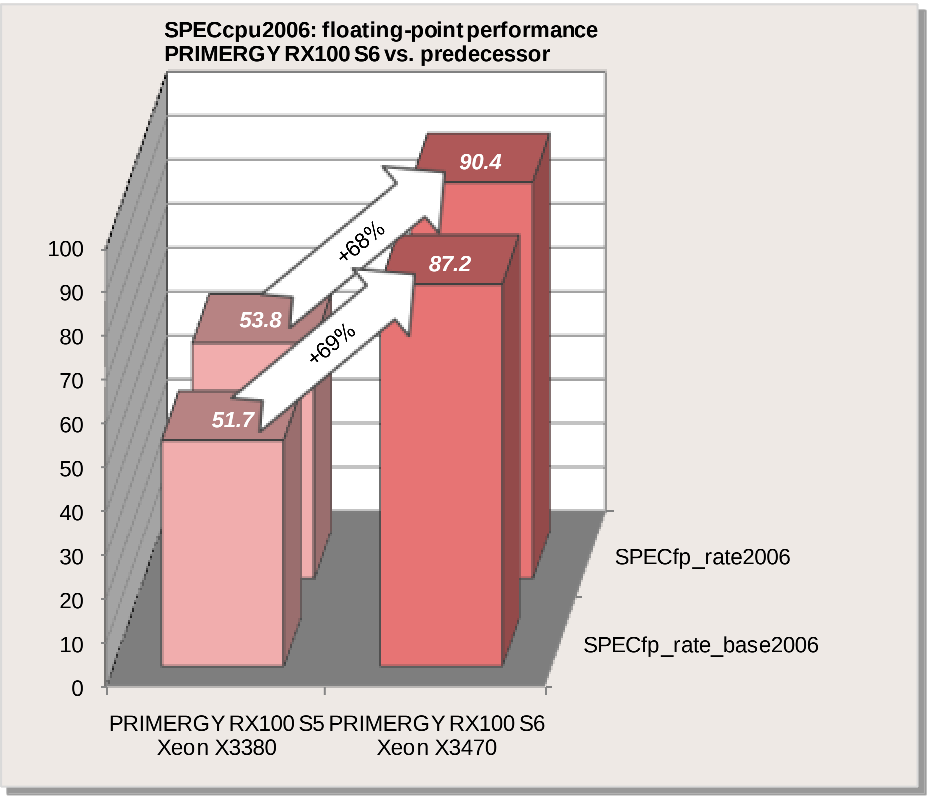

Benchmark environment*

All SPECcpu2006 measurements were performed on a PRIMERGY RX100 S6 with the following hardware and software configuration:

Hardware	
Model	PRIMERGY RX100 S6
CPU	Celeron G1101 Pentium G6950 Core i3-530 and i3-540 Xeon L3406, L3426, X3430, X3440, X3450, X3460 and X3470
Number of chips	Celeron G1101, Pentium G6950, Core i3-530 and i3-540, Xeon L3406: 1 chip, 2 cores, 2 cores per chip Xeon L3426, X3430, X3440, X3450, X3460 and X3470: 1 chip, 4 cores, 4 cores per chip
Primary Cache	32 kB instruction + 32 kB data on chip, per core
Secondary Cache	256 kB on chip, per core
Other Cache	Celeron G1101: 2 MB (I+D) on chip, per chip Pentium G6950: 3 MB (I+D) on chip, per chip Core i3-530 and i3-540, Xeon L3406: 4 MB (I+D) on chip, per chip Xeon L3426, X3430, X3440, X3450, X3460 and X3470: 8 MB (I+D) on chip, per chip
Memory	Dual-Core processors: 2 x 4 GB Quad-Core processors: 4 x 4 GB
Software	
Operating System	SUSE Linux Enterprise Server 11 (64-bit)
Compiler	Intel C++/Fortran Compiler 11.1

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

SPECjbb2005*

Benchmark description

SPECjbb2005 is a Java business benchmark that focuses on the performance of Java server platforms. It is essentially a modernized version of SPECjbb2000 with the main differences being:

- The transactions have become more complex in order to cover a greater functional scope.
- The working set of the benchmark has been enlarged to the extent that the total system load has increased.
- SPECjbb2000 allows only one active Java Virtual Machine instance (JVM), whereas SPECjbb2005 permits several instances, which in turn achieves greater closeness to reality, particularly with large systems.

On the software side SPECjbb2005 measures the implementations of the JVM, JIT (Just-In-Time) compiler, garbage collection, threads and some aspects of the operating system. As far as hardware is concerned, it measures the efficiency of the CPUs and caches, the memory subsystem and the scalability of shared memory systems (SMP). Disk and network I/O are irrelevant.

SPECjbb2005 emulates a 3-tier client/server system that is typical for modern business process applications with emphasis on the middle tier system:

- Clients generate the load, consisting of driver threads, which on the basis of the TPC-C benchmark generate OLTP accesses to a database without thinking times.
- The middle-tier system implements the business processes and the updating of the database.
- The database takes on the data management and is emulated by Java objects that are in the memory. Transaction logging is implemented on an XML basis.

The major advantage of this benchmark is that it includes all three tiers that run together on a single host. The performance of the middle tier is measured, thus avoiding large-scale hardware installations and making direct comparisons possible between SPECjbb2005 results of different systems. Client and database emulation are also written in Java.

SPECjbb2005 only needs the operating system as well as a Java Virtual Machine with J2SE 5.0 features.

The scaling unit is a warehouse with approx. 25 MB Java objects. Precisely one Java thread per warehouse executes the operations on these objects. The business operations are assumed by TPC-C:

- New Order Entry
- Payment
- Order Status Inquiry
- Delivery
- Stock Level Supervision
- Customer Report

However, these are the only features SPECjbb2005 and TPC-C have in common. The results of the two benchmarks are not comparable.

SPECjbb2005 has 2 performance metrics:

- bops (business operations per second) is the overall rate of all business operations performed per second.
- bops/JVM is the ratio of the first metrics and the number of active JVM instances.

In comparisons of various SPECjbb2005 results it is necessary to state both metrics.

The following rules, according to which a compliant benchmark run has to be performed, are the basis for these metrics:

A compliant benchmark run consists of a sequence of measuring points with an increasing number of warehouses (and thus of threads) with the number in each case being increased by one warehouse. The run is started at one warehouse up through 2*MaxWh but not less than 8 warehouses. MaxWh is the number of

* SPEC®, SPECjbb® and the SPEC logo are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).

warehouses with the highest operation rate per second the benchmark expects. Per default the benchmark equates MaxWh with the number of CPUs visible by the operating system.

The metrics bops is the arithmetic average of all measured operation rates with between MaxWh warehouses and 2*MaxWh warehouses.

Benchmark results

In December 2009 the PRIMERGY RX100 S6 was measured with a memory configuration with 16 GB PC3-10600 DDR3-SDRAM with one Xeon X3470 processor. The measurement was taken under Windows Server 2008 Enterprise x64 Edition SP2. Two instances of JRockit(R) 6 P28.0.0 from Oracle were used as JVM for the measurement. All measured values between 4 and 8 warehouses were incorporated in the benchmark result. When comparing the PRIMERGY RX100 S6 and its predecessor the PRIMERGY RX100 S5 all in their highest performance configurations, a throughput increase of +31% is noted.

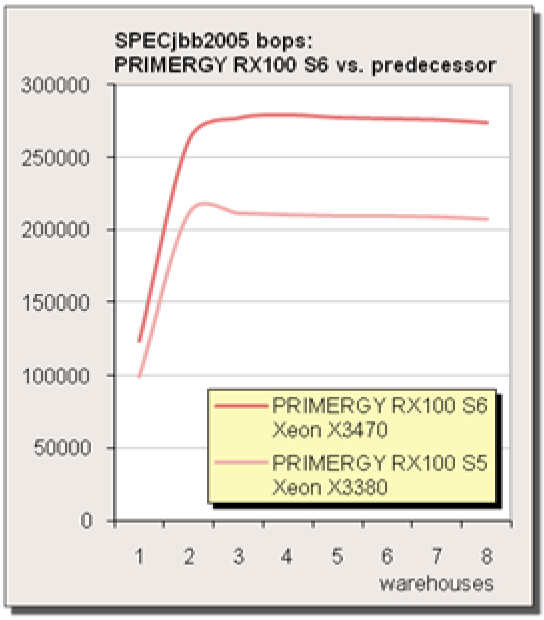

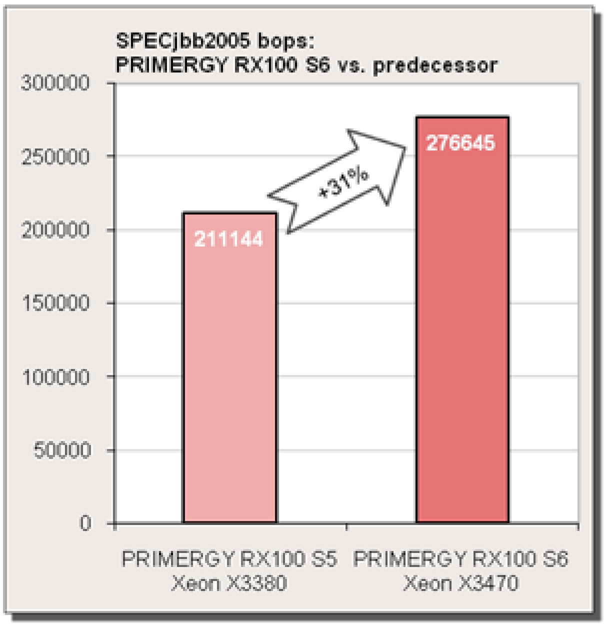

Benchmark environment*

The SPECjbb2005 measurement was performed on a PRIMERGY RX100 S6 with the following hardware and software configuration:

Hardware	
Model	PRIMERGY RX100 S6
CPU	Xeon X3470
Number of chips	1 chip, 4 cores, 4 cores per chip
Primary Cache	32 kB instruction + 32 kB data on chip, per core
Secondary Cache	256 kB (I+D) on chip, per chip
Other Cache	8 MB on chip, per chip
Memory	4 x 4 GB PC3-10600R DDR3-SDRAM
Software	
Operating System	Windows Server 2008 Enterprise x64 Edition SP2
JVM Version	Oracle JRockit(R) 6 P28.0.0 (build P28.0.0-29-114096-1.6.0_11-20090427-1759-windows-x86_64)

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

SPECpower_ssj2008*

Benchmark description

SPECpower_ssj2008 is the first industry-standard SPEC benchmark that evaluates the power and performance characteristics of server class computers. With SPECpower_ssj2008, SPEC has defined server power measurement standards in the same way they have done for performance.

The benchmark workload represents typical server-side Java business applications. The workload is scalable, multi-threaded, portable across a wide range of operating environments, and economical to run. It exercises CPUs, caches, memory hierarchy, and the scalability of symmetric multiprocessor systems (SMPs), as well as implementations of the Java Virtual Machine (JVM), Just In Time (JIT) compiler, garbage collection, threads, and some aspects of the operating system.

SPECpower_ssj2008 reports power consumption for servers at different performance levels — from 100-percent to active idle in 10-percent segments — over a set period of time. The graduated workload recognizes the fact that processing loads and power consumption on servers vary substantially over the course of days or weeks. To compute a power-performance metric across all levels, measured transaction throughputs for each segment are added together, and then divided by the sum of the average power consumed for each segment. The result is a figure of merit called "overall ssj_ops/watt." This ratio gives information about the energy efficiency of the measured server. Because of its defined measurement standard it allows the customers to compare it to other configurations and servers measured with SPECpower_ssj2008. The adjoining diagram shows a typical graph of a SPECpower_ssj2008 result.

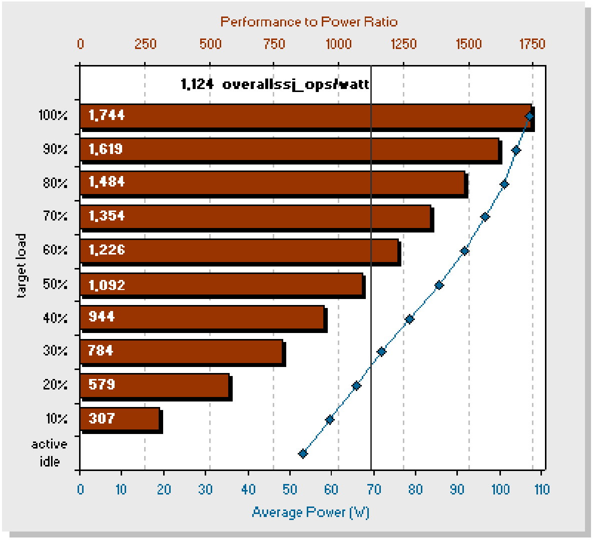

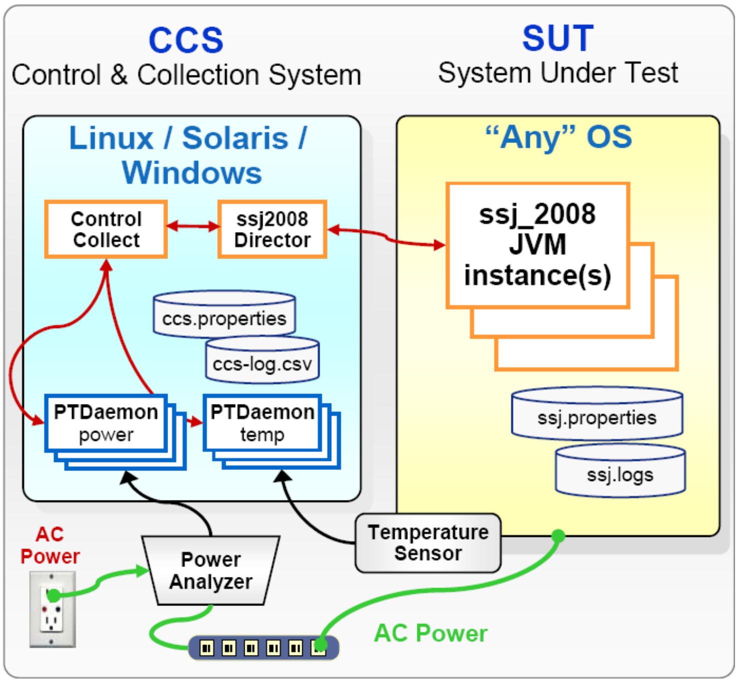

The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System Under Test (SUT) running any of the supported operating systems along with the JVM installed. The JVM provides the environment required to run the SPECpower_ssj2008 workload which is implemented in Java. The other computer is a Collect and Control System (CCS) which controls the operation of the benchmark and captures the power, performance and temperature readings for reporting. The adjoining diagram gives an overview of the different components of this framework.

* SPEC®, SPECpower_ssj2008® and the SPEC logo are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).

Benchmark results

In January 2010 the PRIMERGY RX100 S6 was measured with an Intel Xeon X3470 processor and 4 GB of PC3-10600E DDR3-SDRAM memory. The measurement was taken under Windows Server 2008 Enterprise x64 Edition with SP2 and a JRockit(R) 6 P28.0.0 JVM by Oracle was used.

With the Intel Xeon X3470 processor the PRIMERGY RX100 S6 achieved a result of **1,842 overall ssj_ops/watt**.

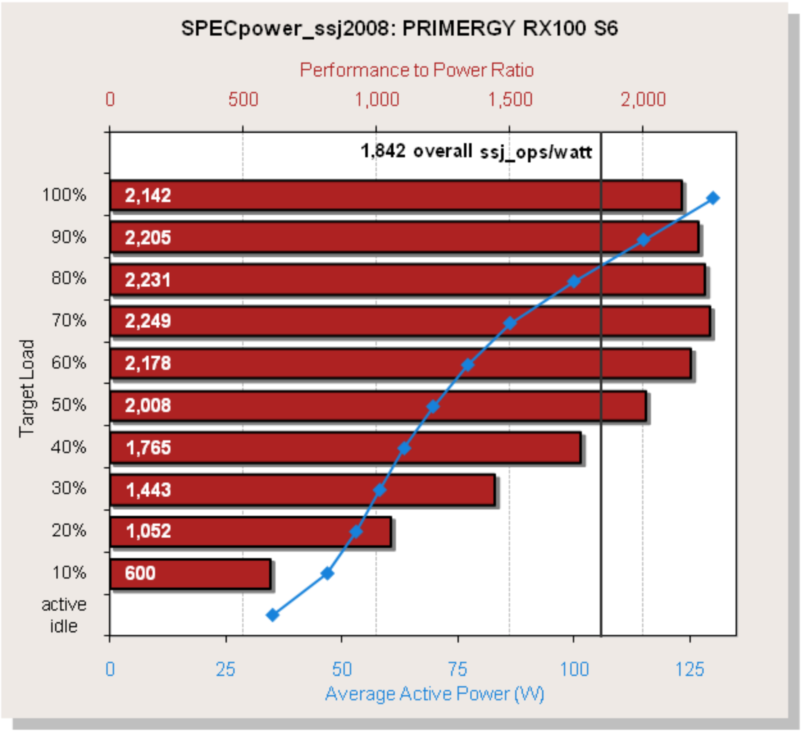

The adjoining diagram shows the result graph of the configuration described above, measured with the PRIMERGY RX100 S6. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level which are tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 1,842 overall ssj_ops/watt for the PRIMERGY RX100 S6. This is the quotient of the sum of the transaction throughputs for each measurement interval and the sum of the average power consumed for each measurement interval.

The diagram shows that the highest energy efficiency of the server is reached at 70% load. This is different to former SPECpower_ssj2008 results on other PRIMERGY servers and already published results on www.spec.org, which have the peak in efficiency at 100% load. The reason behind is that the PRIMERGY RX100 S6 is one of the first servers which has been measured with the “Fujitsu Enhanced Power Settings” power plan. This power plan has been developed for the Microsoft Windows Server 2008 SP2 operating system and is installed on supported PRIMERGY servers via Server View Installation Manager by default. Cunstomers may choose this power plan in the power options of the control panel in Windows as alternative to the default “Balanced” power plan. With the development of the “Fujitsu Enhanced Power Settings” power plan Fujitsu counteracts the fact that the typical server utilization is significantly less than 100% in almost all the time. The “Fujitsu Enhanced Power Settings” power plan optimizes the power management to increase the energy efficiency in the typical load range of a server.

The following table shows the detailed information for the throughput in ssj_ops, the average power consumption in watts and the resulting energy efficiency for each load level of the measured result.

Performance		Power	Energy Efficiency
Target Load	ssj_ops	Average Power (W)	ssj_ops/watt
100%	278,441	130	2,142
90%	252,747	115	2,205
80%	223,063	100	2,231
70%	193,931	86.2	2,249
60%	167,911	77.1	2,178
50%	139,758	69.6	2,008
40%	111,916	63.4	1,765
30%	83,882	58.1	1,443
20%	55,771	53.0	1,052
10%	28,045	46.8	600
Active Idle	0	35.0	0
		∑ssj_ops / ∑power =	**1,842**

The configuration was tuned to get the best possible result for this server in terms of performance per watt. The memory configuration with 2 x 2 GB was selected to meet the criteria of best performance at lowest power consumption by populating only one slot of each available memory channel. This configuration enables the benchmark to use the full potential of the available memory bandwidth and at the same time consumes less power than a comparable performance equivalent 4 DIMMs configuration. The most important factor in the hardware configuration is the right choice of the processor. Processors are the part of a server which consumes the most power beside the memory subsystem. For the PRIMERGY RX100 S6 the quad-core Intel Xeon X3470 processor with a Thermal Design Power (TDP) of 95 watt showed the best efficiency score.

Compared to the predecessor PRIMERGY RX100 S5 which was measured with the Intel Xeon L3360 processor the PRIMERGY RX100 S6 is 64% more energy efficient. The advance of the SPECpower_ssj2008 result of the PRIMERGY RX100 S6 is explained by the new Intel Nehalem microarchitecture which results in lower power consumption and improved throughput in ssj_ops.

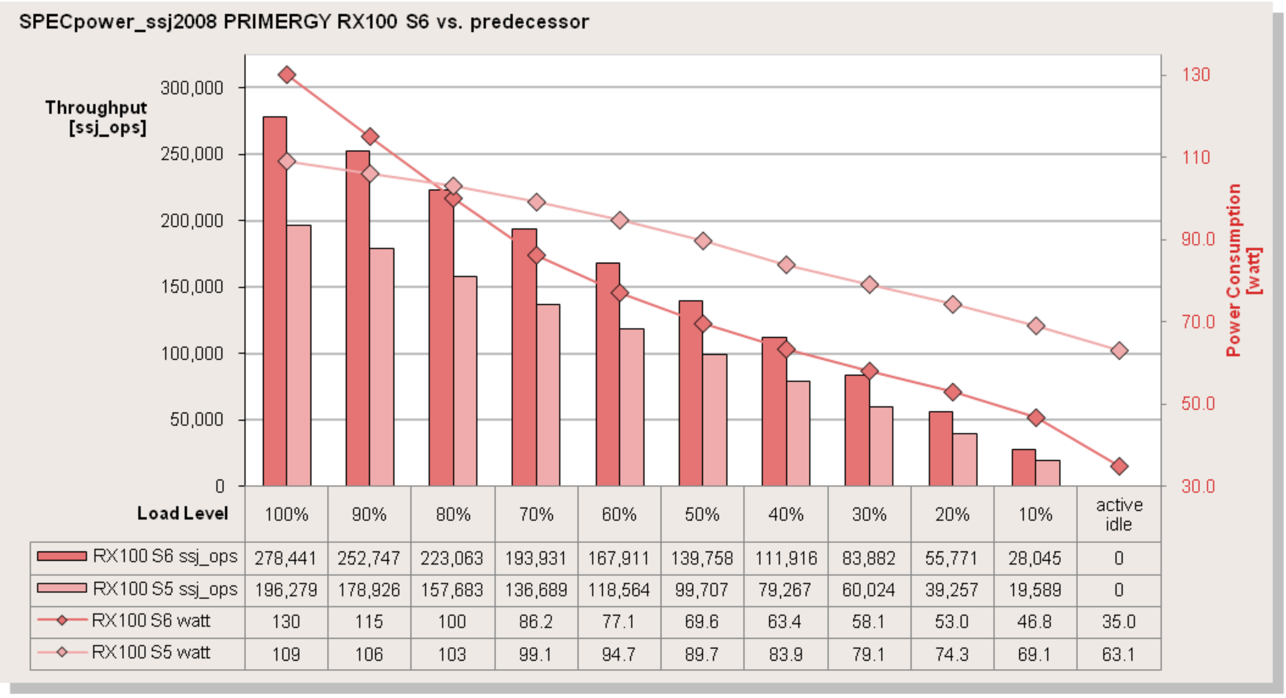

Benchmark environment*

The SPECpower_ssj2008 measurement presented here was performed on a PRIMERGY RX100 S6 with the following hardware and software configuration using the Yokogawa WT210 power analyzer:

Hardware	
Model	PRIMERGY RX100 S6
Processor (TDP)	Intel Xeon X3470 (95 W)
Number of chips	1 chip, 4 cores per chip
Primary Cache	32 KB instruction + 32 KB data on chip, per core
Secondary Cache	256 KB (I+D) on chip, per core
Tertiary Cache	8 MB (I+D) on chip, per chip
Memory	2 x 2 GB PC3-10600E DDR3-SDRAM
Network Interface	1 x Intel 82578DM and 1 x Intel 82574L Gigabit Network Connection (onboard)
Disk Subsystem	1 × Integrated SATA controller 1 × 2.5" SATA HDD, 160 GB, JBOD
Power Supply Unit	1 x 350 W Fujitsu Technology Solutions DPS-350YB A
Software	
Operating System	Windows Server 2008 Enterprise x64 Edition + SP2
JVM Version	Oracle JRockit(R) 6 P28.0.0 (build P28.0.0-29-114096-1.6.0_11-20090427-1759-windows-x86_64)
JVM affinity	start /affinity [0x0F,0xF0]
JVM options	-Xms1625m -Xmx1625m -Xns1400m -XXaggressive -Xlargepages -Xgc:genpar -XXcallprofiling -XXgcthreads=4 -XXtlasize:min=4k,preferred=1024k

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

StorageBench

Benchmark description

To estimate the capability of disk subsystems Fujitsu Technology Solutions defined a benchmark called StorageBench to compare the different storage systems connected to a system. To do this StorageBench makes use of the Iometer measuring tool developed by Intel combined with a defined set of load profiles that occur in real customer applications and a defined measuring scenario.

Measuring tool

Since the end of 2001 Iometer has been a project at http://SourceForge.net and is ported to various platforms and enhanced by a group of international developers. Iometer consists of a user interface for Windows systems and the so-called “dynamo” which is available for various platforms. For some years now it has been possible to download these two components under “Intel Open Source License” from http://www.iometer.org/ or http://sourceforge.net/projects/iometer.

Iometer gives you the opportunity to reproduce the behavior of real applications as far as accesses to IO subsystems are concerned. For this purpose, you can among other things configure the block sizes to be used, the type of access, such as sequential read or write, random read or write and also combinations of these. You can also configure the number of simultaneous accesses ("Outstanding IOs"). As a result Iometer provides a text file with comma separated values (.csv) containing basic parameters, such as throughput per second, transactions per second and average response time for the respective access pattern. This method permits the efficiency of various subsystems with certain access patterns to be compared. Iometer is in a position to access not only subsystems with a file system, but also so-called raw devices.

With Iometer it is possible to simulate and measure the access patterns of various applications, but the file cache of the operating system remains disregarded and operation is in blocks on a single test file.

Load profile

The manner in which applications access the mass storage system considerably influences the performance of a storage system. Examples of various access patterns of a number of applications:

Application	Access pattern
Database (data transfer)	random, 67% read, 33% write, 8 KB (SQL Server)
Database (log file)	sequential, 100% write, 64 KB blocks
Backup	sequential, 100% read, 64 KB blocks
Restore	sequential, 100% write, 64 KB blocks
Video streaming	sequential, 100% read, blocks ≥ 64 KB
File server	random, 67% read, 33% write, 64 KB blocks
Web server	random, 100% read, 64 KB blocks
Operating system	random, 40% read, 60% write, blocks ≥ 4 KB
File copy	random, 50% read, 50% write, 64 KB blocks

From this four distinctive profiles were derived:

Load profile	Access	Access pattern		Block size	Outstanding IOs	Load tool
		read	write			
Streaming	sequential	100%		64 KB	3	Iometer
Restore	sequential		100%	64 KB	3	Iometer
Database	random	67%	33%	8 KB	3	Iometer
File server	random	67%	33%	64 KB	3	Iometer

All four profiles were generated with Iometer.

Measurement scenario

In order to obtain comparable measurement results it is important to perform all the measurements in identical, reproducible environments. This is why StorageBench is based, in addition to the load profile described above, on the following regulations:

- Since real-life customer configurations work only in exceptional situations with raw devices, performance measurements of internal disks are always conducted on disks containing file systems. NTFS is used for Windows and ext3 for Linux, even if higher performance could possibly be achieved with other file systems or raw devices.
- Hard disks are among the most error-prone components of a computer system. This is why RAID controllers are used in server systems in order to prevent data loss through hard disk failure. Here several hard disks are put together to form a "Redundant Array of Independent Disks", known as RAID in short – with the data being spread over several hard disks in such a way that all the data is retained even if one hard disk fails – except with RAID 0. The most usual methods of organizing hard disks in arrays are the RAID levels RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50 and RAID 60. Information about the basics of various RAID arrays is to be found in the paper Performance Report - Modular RAID for PRIMERGY.

  Depending on the number of disks and the installed controller, the possible RAID configurations are used for the StorageBench analyses of the PRIMERGY servers. For systems with two hard disks we use RAID 1 and RAID 0, for three and more hard disks we also use RAID 1E and RAID 5 and, where applicable, further RAID levels – provided that the controller supports these RAID levels.
- Regardless of the size of the hard disk, a measurement file with the size of 8 GB is always used for the measurement.
- In the evaluation of the efficiency of I/O subsystems, processor performance and memory configuration do not play a significant role in today's systems - a possible bottleneck usually affects the hard disks and the RAID controller, and not CPU and memory. Therefore, various configuration alternatives with CPU and memory need not be analyzed under StorageBench.

Measurement results

For each load profile StorageBench provides various key indicators: e.g. "data throughput" in megabytes per second, in short MB/s, "transaction rate" in I/O operations per second, in short IO/s, and "latency time" or also "mean access time" in ms. For sequential load profiles data throughput is the normal indicator, whereas for random load profiles with their small block sizes the transaction rate is normally used. Throughput and transaction rate are directly proportional to each other and can be calculated according to the formula

Data throughput [MB/s]	*= Transaction rate [Disk-I/O s^{-1}] × Block size [MB]*
Transaction rate [Disk-I/O s^{-1}]	*= Data throughput [MB/s] / Block size [MB]*

Benchmark results

The PRIMERGY RX100 S6 is equipped with controllers from the “Modular RAID” family. The variety of the RAID solutions enables the user to choose the right controller for his application scenario.

The PRIMERGY RX100 S6 has the following RAID solutions to offer:

1. SATA RAID onboard controller

   The controller is implemented directly on the motherboard of the server in the Intel Ibex Peak chipset and the RAID stack is realized by the server CPU. This RAID solution is only foreseen for the connection of SATA hard disks. Support is provided for RAID levels 0, 1 and 10. This controller does not have a cache.

2. “RAID 0/1 SAS based on LSI MegaRAID 4Port” Controller (LSI MegaRAID SAS 1064)

   The controller is supplied as a PCI Express card. The maximum number of SATA, SAS and SAS-2.0 hard disks that can be connected to the controller is four. Support is provided for RAID levels 0, 1 and 1E. The controller does not have a cache.

3. “RAID Ctrl SAS 5/6 512MB (D2616)” Controller (LSI MegaRAID SAS 2108)

   The controller supports the 6 Gb/s SAS-2.0 interface for the connection of SAS-2.0 hard disks. It is supplied as a PCI Express card and offers the user a complete RAID solution. Support is provided for RAID levels 0, 1, 5, 6, 10, 50 and 60. This controller is on offer with a 512 MB controller cache. The controller cache can be protected against power failure by an optional battery backup unit (BBU). The controller supports up to 240 hard disks.

Various SATA, SAS and SAS-2.0 hard disks can be connected to these controllers. Depending on the performance required, it is possible to select the appropriate disk subsystem. And depending on the model version the PRIMERGY RX100 S6 offers two hot-plug bays for 3½" hard disks or four hot-plug bays for 2½" hard disks.

The following hard disks can be chosen for the PRIMERGY RX100 S6:

- 3½" SATA hard disks with a capacity of 160 GB, 250 GB, 500 GB, 750 GB and 1 TB (7.2 krpm)
- 2½" SAS-2.0 hard disks with a capacity of 146 GB and 300 GB (10 krpm)
- 2½" SAS-2.0 hard disks with a capacity of 73 GB and 146 GB (15 krpm)
- 3½" SAS hard disks with a capacity of 146 GB (15 krpm)
- 3½" SAS-2.0 hard disks with a capacity of 300 GB, 450 GB and 600 GB (15 krpm)

SATA RAID Onboard Controller

Based on 3½" SATA hard disks, the following diagram shows the throughput differences of a single hard disk (Single Disk, SD) and of two hard disks that are configured in a RAID 1.

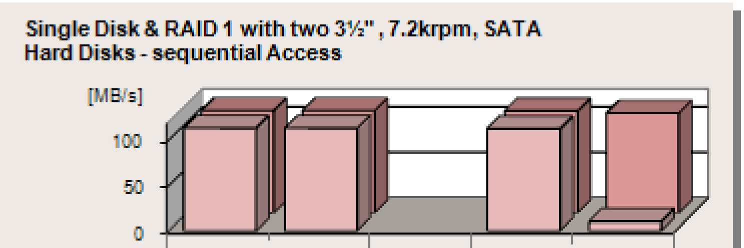

	64 KB sequential reads			64 KB sequential writes	
	SD	RAID 1		SD	RAID 1
HD Cache off	112.60	112.60		111.95	10.26
HD Cache on	112.57	112.60		112.47	109.84

Onboard SATA Controller

Read throughput for sequential reading of 64 KB blocks is not dependent on the cache settings. In RAID 1 the same throughput values are achieved as in the single disk configuration, although RAID 1 offers the benefit of data redundancy.

The throughput for sequential write with 64 KB blocks to a single hard disk is also not dependent on the cache settings.

In contrast, write throughput with sequential access with 64 KB blocks in RAID 1 largely depends on the cache settings.

To ensure data redundancy in RAID 1 it is necessary with write to perform two accesses to two different hard disks. If the disk cache is not enabled and as a result optimized write access is not possible, write throughput in the RAID 1 decreases considerably.

Enabling the disk cache improves the write throughput by a factor of about 11 in RAID 1. The much higher write throughput is explained by the optimized write accesses to the hard disks and the shorter latency times.

Enabling the disk cache leads also to an increase in throughput during random access in RAID 1. However this increase is not as noticeable as with sequential writing. With a random access in RAID 1 the increase in throughput is about 27%.

With a random access with 8 KB respectively 64 KB blocks and a single disk configuration the throughput is not dependant on the cache settings.

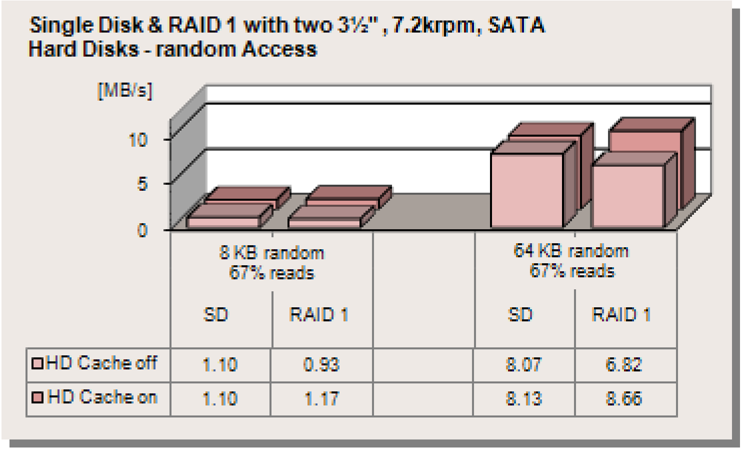

	8 KB random 67% reads			64 KB random 67% reads	
	SD	RAID 1		SD	RAID 1
HD Cache off	1.10	0.93		8.07	6.82
HD Cache on	1.10	1.17		8.13	8.66

Onboard SATA Controller

LSI MegaRAID SAS 1064

The performance of the available hard disk types on the LSI MegaRAID SAS 1064 controller is compared below. This controller does not have a controller cache. This is why only the impact of the disk cache parameters was examined in the measurements and the measurements for the hard disk comparison were in each case performed with and without a disk cache.

The hard disk cache has influence on disk I/O performance. This is frequently seen as a safety problem in the event of a power failure and is therefore disabled. Nevertheless, it was for a good reason integrated by the hard disk manufacturers to increase write performance. For performance reasons it is advisable to enable the disk cache for the SATA hard disks in particular, which in comparison with the SAS hard disks rotate slowly. The by far larger cache for I/O accesses and thus a potential safety risk (data loss) in the event of a power failure is situated in the main memory and is administered by the operating system. To prevent data losses it is advisable to equip the system with an uninterruptible power supply (UPS).

In the test setup two hard disks were connected to the controller and configured as a RAID 1. In the measurements all hard disk types currently available for the PRIMERGY RX100 S6 were analyzed. The throughputs of the individual hard disk types in RAID 1 are compared below with different access patterns.

The throughputs that were achieved with sequential read and write in RAID 1 with a block size of 64 KB are shown in the diagram.

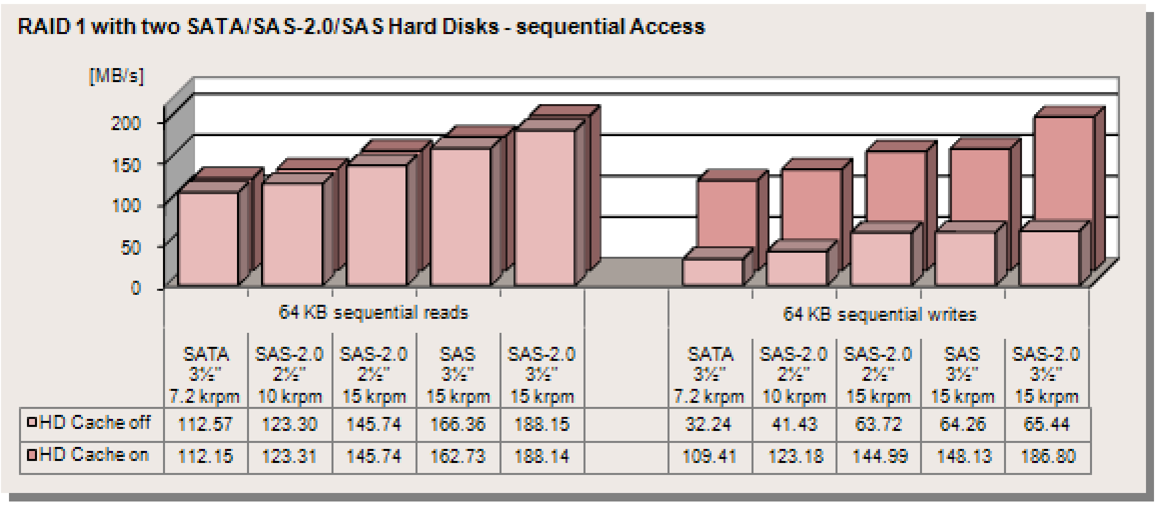

	64 KB sequential reads						64 KB sequential writes				
	SATA 3½" 7.2 krpm	SAS-2.0 2½" 10 krpm	SAS-2.0 2½" 15 krpm	SAS 3½" 15 krpm	SAS-2.0 3½" 15 krpm		SATA 3½" 7.2 krpm	SAS-2.0 2½" 10 krpm	SAS-2.0 2½" 15 krpm	SAS 3½" 15 krpm	SAS-2.0 3½" 15 krpm
HD Cache off	112.57	123.30	145.74	166.36	188.15		32.24	41.43	63.72	64.26	65.44
HD Cache on	112.15	123.31	145.74	162.73	188.14		109.41	123.18	144.99	148.13	186.80

LSI MegaRAID SAS 1064

If for sequential read in RAID 1 with enabled disk cache a 2½" SAS-2.0 hard disk with 15 krpm is used instead of one 2½" SAS-2.0 hard disk with a rotational speed of 10 krpm, the result is an increase in throughput of about 18%. If you compare the throughputs of the 3½" 15 krpm SAS and 3½" 7.2 krpm SATA hard disks, you can then see that the throughput in RAID 1 for sequential read, with the SAS hard disks is about 45% higher than for the SATA hard disks.

If you compare the throughputs of the 2½" and 3½" SAS-2.0 hard disks, all of them with a rotational speed of 15 krpm, you can see that the throughput for sequential access with the 3½" hard disks in RAID 1 is about 29% higher than for the 2½" hard disks.

If for sequential write in RAID 1 with enabled disk cache a 2½" SAS-2.0 hard disk with a rotational speed of 15 krpm is used instead of one 2½" SAS-2.0 hard disk with 10 krpm the result is an increase in throughput of about 18%. If for sequential write in RAID 1 with enabled disk cache a 3½" SAS hard disk with a rotational speed of 15 krpm is used instead of one 3½" SATA hard disk with 7.2 krpm the result is an increase in throughput of about 35%.

A special increase in throughput for sequential write, up to 4.4-fold, can be achieved with the SATA hard disk by enabling the disk cache. The increase in throughput gained with SAS hard disks through enabling the disk cache is not so pronounced as with the SATA hard disks, but it is still significant. For the 2½" SAS-2.0 hard disks with 10 krpm throughput increases by a factor of 3 and 2.3 for the 2½" SAS-2.0 hard disks with 15 krpm. For the 3½" SAS hard disks with 15 krpm the throughput increases by a factor of 2.3 for the 3½" SAS-2.0 hard disks with 15 krpm the throughput increases by a factor of 2.9.

The following diagram shows that for random access with 67% reads the disk cache also plays an important role in improving throughput.

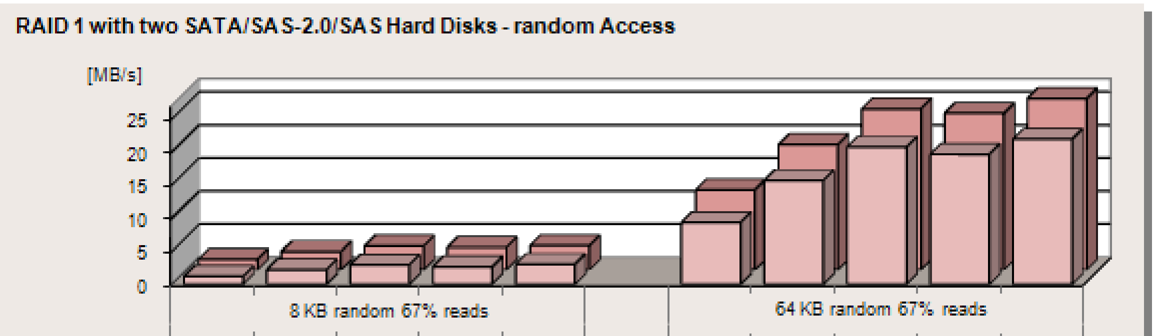

	8 KB random 67% reads						64 KB random 67% reads				
	SATA 3½" 7.2 krpm	SAS-2.0 2½" 10 krpm	SAS-2.0 2½" 15 krpm	SAS 3½" 15 krpm	SAS-2.0 3½" 15 krpm		SATA 3½" 7.2 krpm	SAS-2.0 2½" 10 krpm	SAS-2.0 2½" 15 krpm	SAS 3½" 15 krpm	SAS-2.0 3½" 15 krpm
HD Cache off	1.21	2.21	2.93	2.66	3.01		9.27	15.62	20.68	19.58	21.88
HD Cache on	1.59	2.73	3.57	3.37	3.61		12.03	18.90	24.28	23.62	25.85

LSI MegaRAID SAS 1064

Through enabling the disk cache an increase in throughput of about 22% was achieved for random access with 8 KB blocks in RAID 1 with two 2½" 10 krpm or 15 krpm SAS-2.0 hard disks. If you compare the throughputs of the 2½" 10 krpm and 15 krpm SAS-2.0 hard disks in RAID 1 it is then evident that the throughput of the 15 krpm hard disks for random access with 8 KB blocks and enabled disk cache is about 31% higher than that of the 10 krpm hard disks. If you compare the throughputs of the 3½" 15 krpm SAS respectively SAS-2.0 hard disks and 3½" 7.2 krpm SATA hard disks in RAID 1 it is then evident that the throughput of the SAS espectively SAS-2.0 hard disks for random access with 8 KB blocks is about 2 times higher than that of the SATA hard disks. If you compare throughputs in RAID 1 for random access with 8 KB blocks and enabled disk cache of the 2½" and 3½" SAS-2.0 hard disks, all of them with a rotational speed of 15 krpm, you can see that the throughputs are about within the same performance range.

Through enabling the disk cache an increase in throughput of about 20% respectively 18% was achieved for random access with 8 KB respectively 64 KB blocks in RAID 1 with two 3½" 15 krpm SAS-2.0 hard disks. Through enabling the disk cache an increase in throughput of about 22% respectively 17% was achieved for random access with 8 KB respectively 64 KB blocks in RAID 1 with two 2½" 15 krpm SAS-2.0 hard disks. Through enabling the disk cache an increase in throughput of about 30% was achieved for random access with 64 KB blocks in RAID 1 with two 3½" 7.2 krpm SATA hard disks.

If you compare the throughputs of the 3½" and 2½" SAS-2.0 hard disks, all of them with a rotational speed of 15 krpm, it is then evident that the throughput of the 3½" hard disks for random access with 64 KB blocks in RAID 1 and enabled disk cache is about 6% higher than that of the 2½" hard disks. If you compare the throughputs of the 3½" 15 krpm SAS and 3½" 7.2 krpm SATA hard disks, it is then evident that the throughput of the SAS hard disks for random access with 64 KB blocks in RAID 1 and enabled disk cache is about 96% higher than that of the SATA hard disks.

LSI MegaRAID SAS 2108

The RAID array defines the way in which data is treated as regards availability. How quickly the data is transferred in the respective RAID array context depends largely on the data throughput of the hard disks. The number of hard disks configured for the measurements in a RAID array was defined depending on the RAID level. Between two and three hard disks were used. To ensure that the hard disks do not represent a bottleneck when determining the efficiency of the controller under various cache settings, the measurements were performed with 2½" SAS-2.0 hard disks with a rotational speed of 15 krpm.

The throughput can in certain cases be considerably increased through the cache settings. However, these increases in throughput differ – depending on the data structure and access pattern. For the measurements the controller cache option "Read-Mode" is always set to "No Read-ahead" and the option "I/O cache" is always set to "I/O direct". The options "Write-Mode" and "Disk cache" were varied.

The following diagram shows the throughputs for sequential read and write with 64 KB blocks and for different cache settings in RAID 1 with two and in RAID 5 with three 2½" SAS-2.0 hard disks.

The sequential read throughput achieves very good values and does not depend on the cache settings.

In contrast, the write throughput depends on the cache settings. In order to achieve optimal performance with RAID 1, it is necessary to use the "Disk cache enabled" option as the optimal cache setting. In our case the throughput was improved by a factor 2.3 using sequential write.

The importance of optimal cache settings for a good performance can be seen particularly clearly with RAID 5. The diagram shows that sequential write throughput increases considerably, by a factor of 22.7, as a result of enabling the controller cache with the option "Write-back" and the disk cache with the option "enabled".

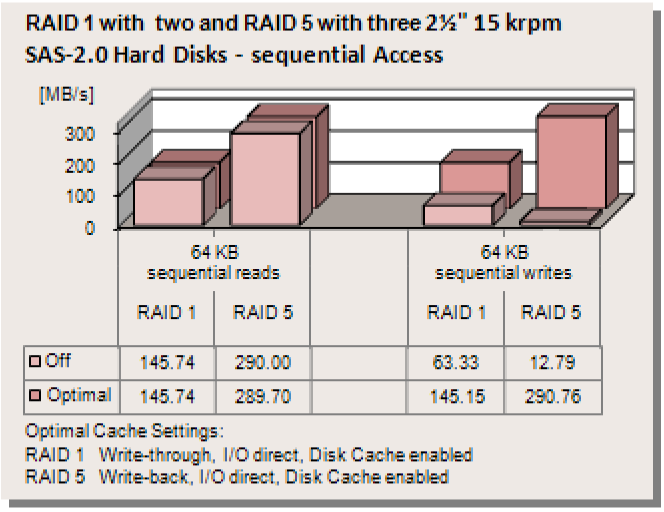

LSI MegaRAID SAS 2108 with 512 MB Cache

To achieve optimal throughput for random access with RAID 1 it is important to set the write mode option of the controller cache to "Write-through" and to enable the disk cache of the hard disk. As a result of these optimal cache settings, improvements in throughput of 28% and 23% are achieved depending on whether blocks of 8 KB or 64 KB are used for random access.

To achieve optimal throughput for random access with RAID 5 it is important to set the write mode option of the controller cache to "Write-back" and to enable the disk cache of the hard disk. Due to these optimal cache settings, improvements in throughput of 61% and 42% are achieved depending on block size.

More detailed information about this topic is available in the paper Performance Report - Modular RAID for PRIMERGY.

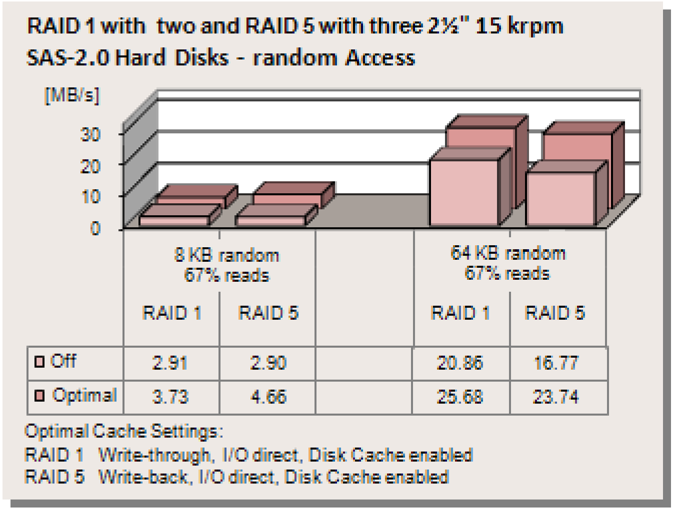

LSI MegaRAID SAS 2108 with 512 MB Cache

Controller comparison

The following comparison depicts the throughputs of two controllers. The measurements were made with the same hard disk types in the same RAID 1 array. The diagram shows the throughputs achieved with disabled caches (Off) and with optimal cache settings (Optimal).

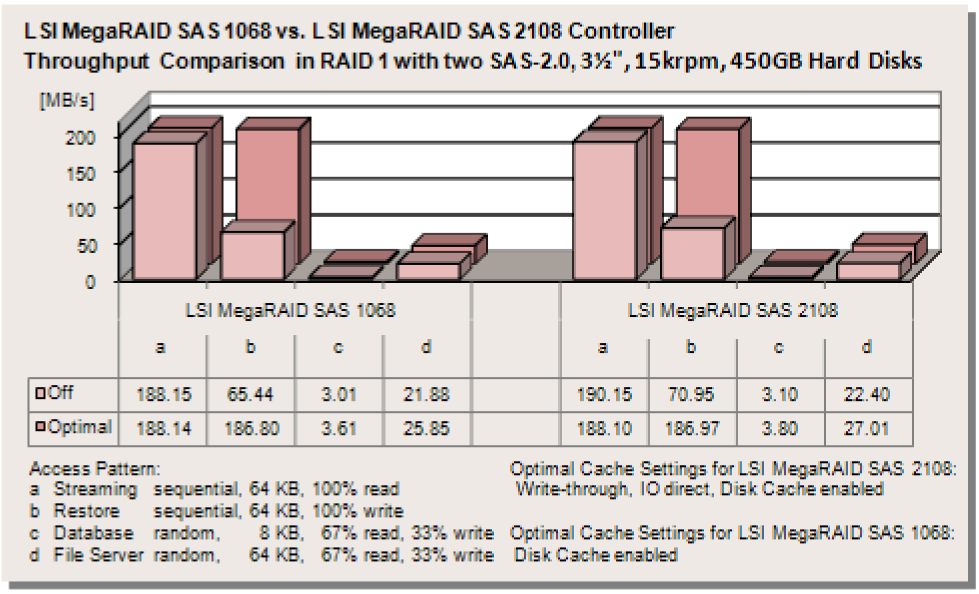

The LSI MegaRAID SAS 1064 or the LSI MegaRAID SAS 2108 controller can be used to connect SAS-2.0 hard disks. The differences in performance of the controllers used are minimal with pure sequential access and optimal cache settings. With sequential read all the controllers achieve the range of the maximum throughput rate regardless of the cache settings. Even for sequential write both controllers are within the same performance range.

The LSI MegaRAID SAS 2108 controller, which with its controller cache and extended function is optimally equipped for a more heavily loaded application profile and higher RAID levels, also performs well here for random access in RAID 1 with this load profile.

Conclusion

With the “Modular RAID” concept, the PRIMERGY RX100 S6 offers a plethora of opportunities to meet the various requirements of different application scenarios.

For a configuration with SATA hard disks the onboard SATA RAID controller offers the user good solution options with a very good price/performance ratio.

The entry-level controller, represented by the LSI MegaRAID SAS 1064 controller, offers the basic RAID solutions RAID 0, RAID 1 and RAID 1E and supports these RAID levels with a very good performance.

The “high-end” controller, represented by the LSI MegaRAID SAS 2108 controller, offers all today’s current RAID solutions; for the PRIMERGY RX100 S6, which can be expanded with up to four internal hard disks, this can be RAID levels 0, 1, 5, 6 and 10. This controller is supplied with a 512 MB controller cache and can as an optional extra be secured with a BBU. Various options for setting the use of the cache enable controller performance to be flexibly adapted to suit the RAID levels used.

Use of RAID 5 or RAID 6 enables the existing hard disk capacity to be utilized economically for a good performance. However, we recommend a RAID 10 for optimal performance and security.

The PRIMERGY RX100 S6 offers a choice between SATA, SAS and SAS-2.0 hard disks. 3½" hard disks with a rotational speed of 7.2 krpm are available with the SATA hard disks.

2½" and 3½" hard disks are available with the SAS-2.0 hard disks. The 2½" hard disks are offered with a rotational speed of 10 krpm or 15 krpm and the 3½" hard disks with a rotational speed of 15 krpm.

The 3½" SAS hard disks are offered with a rotational speed of 15 krpm.

Depending on the performance required, a decision must be taken as to which hard disk type with which rotational speed is to be used. The 2½" hard disks with a rotational speed of 15 krpm generally provide higher throughputs than hard disks that rotate more slowly. However, throughput very much depends on the access pattern. During our analyses throughput differences of up to 53% were determined, especially with random access. As a result of using 2½" hard disks it is possible – depending on the RAID level – to achieve higher parallelism through the use of more hard disks in the RAID array.

For maximum performance it is advisable, particularly with SATA hard disks or when using a controller without a controller cache, to enable the hard disk cache. Depending on the disk type used, the increase in performance is 22.7-fold. When the hard disk cache is enabled we recommend the use of a UPS.

Benchmark environment*

All the measurements presented here were performed with the hardware and software components listed below.

Component	Details
Server	PRIMERGY RX100 S6
Operating system	Windows Server 2008, Enterprise Edition Version: 6.0.6001 Service Pack 1 Build 6001
File system	NTFS
Measuring tool	Iometer 27.07.2006
Measurement data	Measurement file of 32 GB
Onboard SATA Controller	Intel Ibex Peak BIOS: 6.00.1.05 SATA RAID mode
Controller “RAID 0/1 SAS based on LSI MegaRAID 4Port” (LSI MegaRAID SAS 1064)	Driver Name: lsi_sas.sys, Driver Version: 1.29.03.00 Firmware version: 1.27.00.00 BIOS version: 06.26.00.00
Controller “RAID Ctrl SAS 5/6 512MB (D2616)” (LSI MegaRAID SAS 2108)	Driver name: megasys2.sys, Driver version: 4.18.0.64 Firmware package: 12.0.1-0057 Firmware version: 2.0.03-0673 BIOS version: 3.07.00 Controller cache: 512 MB
Hard Disk SATA, 3½", 7.2 krpm	Western Digital WD5002ABYS, 500 GB
Hard Disk SAS-2.0, 2½", 10 krpm	Fujitsu MBD2147RC, 147 GB
Hard Disk SAS-2.0, 2½", 15 krpm	Fujitsu MBE2147RC, 147 GB
Hard Disk SAS, 3½", 15 krpm	Seagate ST3146356SS, 147 GB
Hard Disk SAS-2.0, 3½", 15 krpm	Hitachi HUS156045VLS600, 450 GB

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

OLTP-2

Benchmark description

OLTP stands for Online Transaction Processing. The OLTP-2 benchmark is based on the typical application scenario of a database solution. In OLTP-2 database access is simulated and the number of transactions achieved per second (tps) determined as the unit of measurement for the performance of the system measured.

In contrast to benchmarks such as SPECint and TPC-E, which were standardized by independent bodies and for which adherence to the respective rules and regulations are monitored, OLTP-2 is an internal benchmark of Fujitsu Technology Solutions. The partially enormous hardware and time expenditure for standardized benchmarks has been reduced to a reasonable degree in OLTP-2 so that a variety of configurations can be measured within an acceptable period of time.

Even if the two benchmarks OLTP-2 and TPC-E simulate similar application scenarios using the same workload, the results cannot be compared or even treated as equal, as the two benchmarks use different methods to simulate user load. OLTP-2 values are typically similar to TPC-E values. A direct comparison, or even referring to the OLTP-2 result as TPC-E, is not permitted, especially because there is no price-performance calculation.

Benchmark results

There are a couple of processors released for PRIMERGY RX100 S6. The following table show the processors used in OLTP-2 measurement.

Processor	Cores/ Chip	HT	TM	GHz	L3-cache	Memory	TDP	tps
Celeron G1101	2			2.27	2 MB per chip	1067 MHz	73 watt	91.67
Pentium G6950	2			2.80	3 MB per chip	1067 MHz	73 watt	106.85
Core i3-530	2	✓		2.93	4 MB per chip	1333 MHz	73 watt	166.54
Core i3-540	2	✓		3.07	4 MB per chip	1333 MHz	73 watt	171.04
Xeon L3426	4	✓	✓	1.87	8 MB per chip	1333 MHz	45 watt	289.08
Xeon X3430	4		✓	2.40	8 MB per chip	1333 MHz	95 watt	272.48
Xeon X3440	4	✓	✓	2.53	8 MB per chip	1333 MHz	95 watt	350.76
Xeon X3450	4	✓	✓	2.67	8 MB per chip	1333 MHz	95 watt	358.50
Xeon X3460	4	✓	✓	2.80	8 MB per chip	1333 MHz	95 watt	374.83
Xeon X3470	4	✓	✓	2.93	8 MB per chip	1333 MHz	95 watt	386.18

HT = Hyper-Threading, TM = Turbo Mode, TDP = Thermal Design Power

All results were determined on the basis of the operating system Microsoft Windows Server 2008 Enterprise x64 Edition SP1 and the database SQL Server 2008 Enterprise x64 Edition SP1. OLTP-2 benchmark results depend to a great degree on the configuration options of a system with hard disks and their controllers. Therefore, the system was equipped with two 2-port SAS RAID controllers that were connected to 12 PRIMERGY SX40 and a total of 144 SAS-hard disks. The disk subsystem was dimensioned to be no bottleneck within the measurements. Comparable results may also be achievable with other disk subsystems being no bottleneck. See the Benchmark environment section for further information on the system configuration.

The maximum memory size of PRIMERGY RX100 S6 with 4 memory modules depends on the processor type. Processors Celeron, Pentium and Core i3 can be used with a maximum of 16 GB UDIMM at 1067 MHz or 1333 MHz. The Xeon Processors can be used with a maximum of 32 GB RDIMM. Although these Processors can access memory with 1333 MHz, the access for 32 GB total is only 1067 MHz.

The next diagrams show the OLTP-2 performance data for the PRIMERGY RX100 S6 with processors Celeron G1101, Pentium G6950, Core i3-530, Core i3-540 at 16 GB memory and processors Xeon L3426, X3430, X3440, X3450, X3460, X3470 at 32 GB memory.

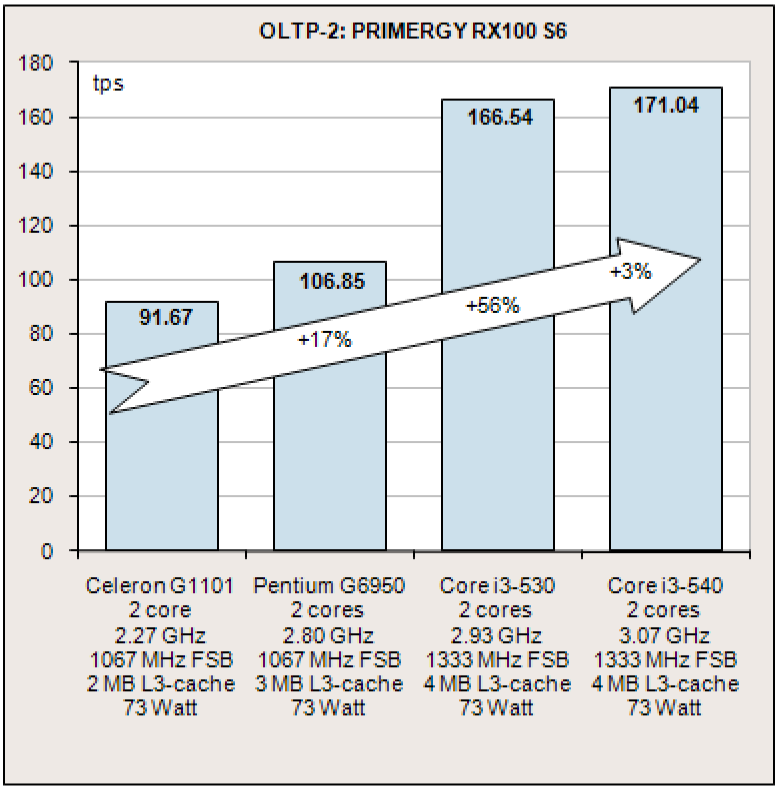

The largest scaling in the first graphic with +56% is at Pentium to Core-i3. Higher frequency, larger cache, higher memory access frequency and Hyper-Threading are responsible for the increase in throughput.

The largest scaling over all Xeon processor types is at X3430 to X3440 with +29%. The reason is Hyper-Treading. In all cases higher frequency increases the throughput.

Comparing Dual-Core i3-540 and Quad-Core X3430 the performance increase is +59% by doubling the number of cores and memory size.

OLTP-2 values RX100 S6 with processors Celeron G1101, Pentium G6950, Core i3-530, Core i3-540 and 16 GB Memory

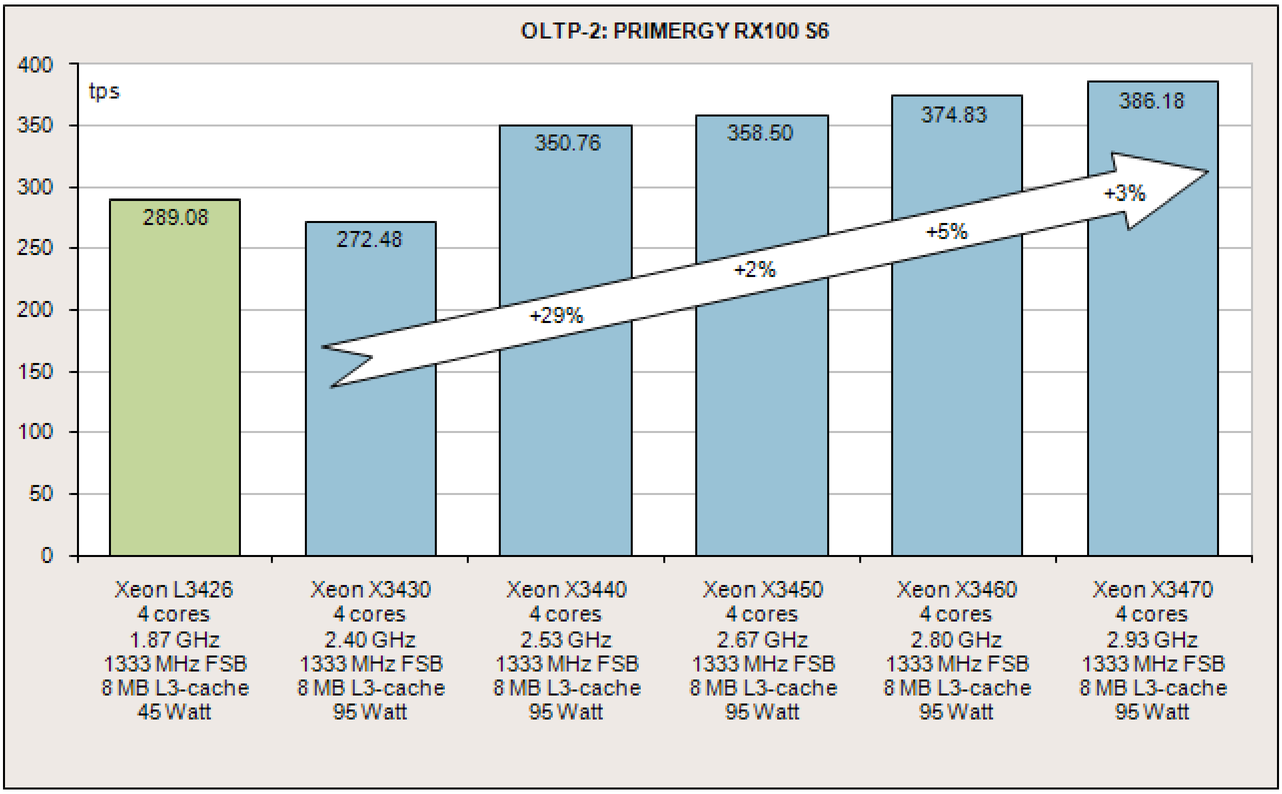

OLTP-2 values RX100 S6 with Xeon processors and 32 GB Memory

Benchmark environment*

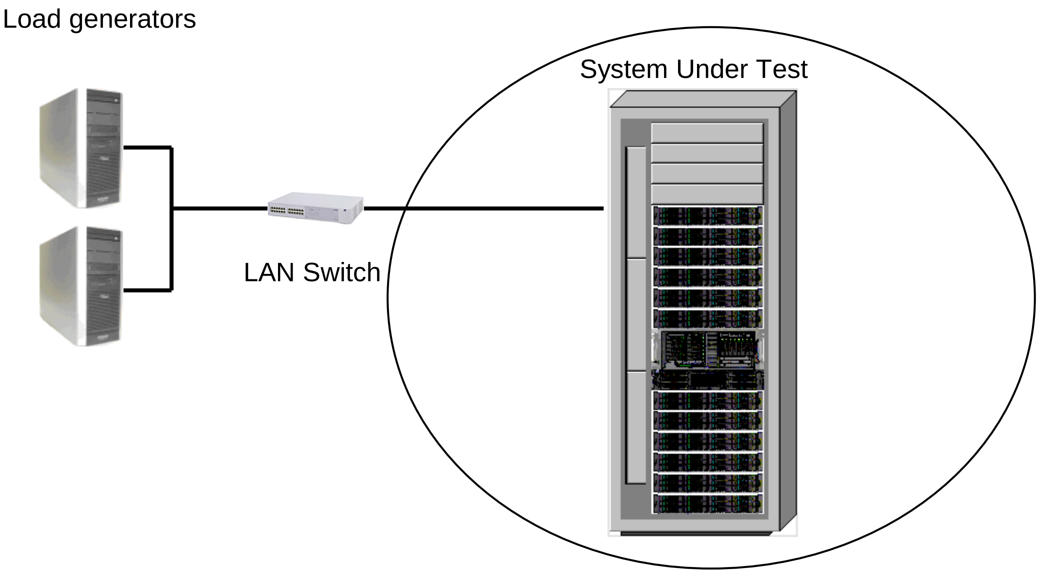

System Under Test (Tier B)	
Hardware	
Server	PRIMERGY RX100 S6
Processor	Celeron G1101, Pentium G6950, Core i3-530, Core i3-540, Xeon L3426, X3430, X3440, X3450, X3460, X3470
Memory	4 x 4 GB DDR3 PC3-10600E or 4 x 8 GB DDR3 PC3-8500R
Settings (default)	Turbo Mode enabled, Hyper-Threading enabled
Network Interface	2 x 1-GBit LAN Intel (onboard)
Disk Subsystem	PRIMERGY RX100 S6: 1 x RAID Ctrl SAS 6G 5/6 512MB 1 x HD SAS 6G 146GB 15K 2.5" RAID-0, OS 3 x HD SAS 6G 146GB 15K 2.5" RAID-0, log 2 x LSI SAS MegaRAID 9280-8e 12 x FibreCAT SX40: 144 x Seagate 146 GB 15 krpm, RAID-0, data
Software	
Operating System	Windows Server 2008 Enterprise x64 Edition SP1
Database	SQL Server 2008 Enterprise x64 Edition SP1

Front End (Tier A)	
Hardware	
Model	1 x PRIMERGY RX300 S4
Processor	2 x Xeon E5420 2.50 GHz, 2 MB L2 cache
Memory	4 GB FBD667 PC2-5300F
Network Interface	2 x 1-GBit LAN (onboard), dual port LAN 1-GBit
Software	
Operating System	Windows Server 2003 R2 Standard x64 Edition

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

Load Generators	
Hardware	
Model	2 x PRIMERGY Econel 200
Processor	2 x Xeon 3.40 GHz, 2 MB L2 cache
Memory	2 GB DDR-SDRAM PC2700
Network Interface	1 x 1-GBit LAN (onboard)
Software	
Operating System	Windows Server 2003 Standard Edition SP1 (x86)
OLTP-2 Software	EGen version 1.8.0-1015

Literature

PRIMERGY Systems	
	http://ts.fujitsu.com/primergy
PRIMERGY RX100 S6	
	Data sheet http://docs.ts.fujitsu.com/dl.aspx?id=4cfe0d12-d479-44c3-8667-7184ddf58bd5
PRIMERGY Performance	
	http://ts.fujitsu.com/products/standard_servers/primergy_bov.html
OLTP-2	
	Benchmark Overview OLTP-2 http://docs.ts.fujitsu.com/dl.aspx?id=e6f7a4c9-aff6-4598-b199-836053214d3f
SPECcpu2006	
	http://www.spec.org/osg/cpu2006
	Benchmark Overview-SPECcpu2006 http://docs.ts.fujitsu.com/dl.aspx?id=1a427c16-12bf-41b0-9ca3-4cc360ef14ce
SPECjbb2005	
	http://www.spec.org/jbb2005
	Benchmark Overview SPECjbb2005 http://docs.ts.fujitsu.com/dl.aspx?id=5411e8f9-8c56-4ee9-9b3b-98981ab3e820
SPECpower_ssj2008	
	http://www.spec.org/power_ssj2008
	Benchmark Overview SPECpower_ssj2008 http://docs.ts.fujitsu.com/dl.aspx?id=166f8497-4bf0-4190-91a1-884b90850ee0
StorageBench	
	Performance Report - Modular RAID for PRIMERGY http://docs.ts.fujitsu.com/dl.aspx?id=8f6d5779-2405-4cdd-8268-1f948ba050e6
	iometer http://www.iometer.org

Contact

PRIMERGY Hardware

PRIMERGY Product Marketing
mailto:Primergy-PM@ts.fujitsu.com

PRIMERGY Performance and Benchmarks

mailto:primergy.benchmark@ts.fujitsu.com

Published by department:

Enterprise Products
PRIMERGY Server
PRIMERGY Performance Lab
mailto:primergy.benchmark@ts.fujitsu.com

Internet:
http://ts.fujitsu.com/primergy
Extranet:
http://partners.ts.fujitsu.com/com/products/servers/primergy