White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY TX1330 M2

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY TX1330 M2.

The PRIMERGY TX1330 M2 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.

Version
1.1
2016-02-29

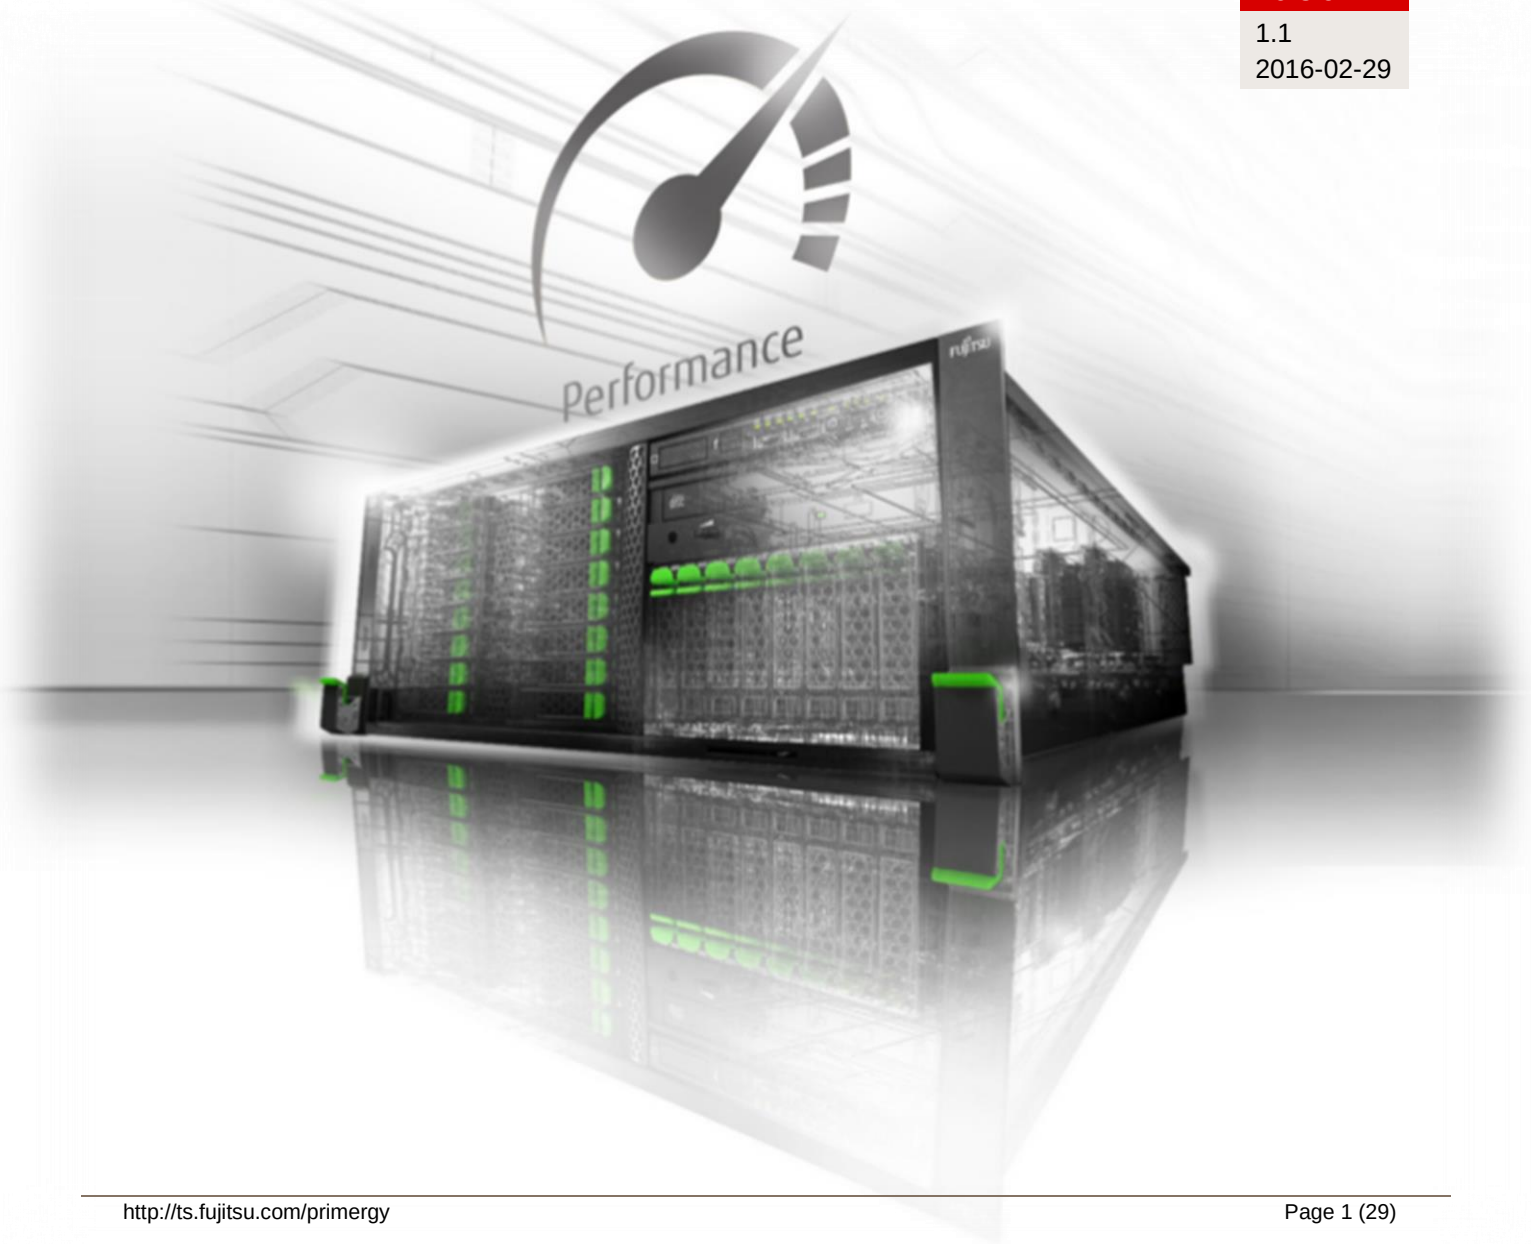

Contents

Document history

Version 1.0 (2016-01-14)

New:

- Technical data
- SPECcpu2006
  Measurements with Pentium G4400, Core i3-6100 and Intel® Xeon® Processor E3-1200 v5 Product Family
- SPECpower_ssj2008
  Measurement with Xeon E3-1240L v5
- STREAM
  Measurements with Pentium G4400, Core i3-6100 and Intel® Xeon® Processor E3-1200 v5 Product Family

Version 1.1 (2016-02-29)

New:

- Disk I/O: Performance of storage media
  Results for 2.5" and 3.5" storage media
- Disk I/O: Performance of RAID controllers
  Measurements with “LSI SW RAID on Intel C236 (Onboard SATA)”, “PRAID CP400i”, “PRAID EP400i” and “PRAID EP420i” controllers

Updated:

- Technical data
  Xeon E3-1225 v5 added
- SPECcpu2006
  Measurements with Xeon E3-1225 v5
- STREAM
  Measurement with Xeon E3-1225 v5

Technical data

PRIMERGY TX1330 M2

Decimal prefixes according to the SI standard are used for measurement units in this white paper (e.g. 1 GB = 10^9 bytes). In contrast, these prefixes should be interpreted as binary prefixes (e.g. 1 GB = 2^{30} bytes) for the capacities of caches and memory modules. Separate reference will be made to any further exceptions where applicable.

Model	PRIMERGY TX1330 M2
Model versions	PY TX1330M2/f/Standard PSU PY TX1330M2/f/Red. PSU PY TX1330M2/r/Red. PSU
Form factor	Tower server
Chipset	Intel® C236
Number of sockets	1
Processor type	Intel® Pentium® G4400 Intel® Core™ i3-6100 Intel® Xeon® Processor E3-1200 v5 Product Family
Number of memory slots	4
Maximum memory configuration	64 GB
Onboard LAN controller	2 × 1 Gbit/s
Onboard HDD controller	Controller with RAID 0, RAID 1 or RAID 10 for up to 4 SATA HDDs
PCI slots	2 × PCI-Express 3.0 x8 1 × PCI-Express 3.0 x4 1 × PCI-Express 3.0 x1 (mech. x4)
Max. number of internal hard disks	24

Processors (since system release)							
Processor	Cores	Threads	Cache [MB]	Rated Frequency [Ghz]	Max. Turbo Frequency [Ghz]	Max. Memory Frequency [MHz]	TDP [Watt]
Pentium G4400	2	2	3	3.30	n/a	2133	54
Core i3-6100	2	4	3	3.70	n/a	2133	51
Xeon E3-1220 v5	4	4	8	3.00	3.50	2133	80
Xeon E3-1225 v5	4	4	8	3.30	3.70	2133	80
Xeon E3-1240L v5	4	8	8	2.10	3.20	2133	25
Xeon E3-1230 v5	4	8	8	3.40	3.80	2133	80
Xeon E3-1240 v5	4	8	8	3.50	3.90	2133	80
Xeon E3-1270 v5	4	8	8	3.60	4.00	2133	80
Xeon E3-1280 v5	4	8	8	3.70	4.00	2133	80

All the processors of Intel® Xeon® Processor E3-1200 v5 Product Family that can be ordered with the PRIMERGY TX1330 M2 support Intel® Turbo Boost Technology 2.0. This technology allows you to operate the processor with higher frequencies than the nominal frequency. Listed in the processor table is "Max. Turbo Frequency" for the theoretical frequency maximum with only one active core per processor. The maximum frequency that can actually be achieved depends on the number of active cores, the current consumption, electrical power consumption and the temperature of the processor.

As a matter of principle Intel does not guarantee that the maximum turbo frequency will be reached. This is related to manufacturing tolerances, which result in a variance regarding the performance of various examples of a processor model. The range of the variance covers the entire scope between the nominal frequency and the maximum turbo frequency.

The turbo functionality can be set via BIOS option. Fujitsu generally recommends leaving the "Turbo Mode" option set at the standard setting "Enabled", as performance is substantially increased by the higher frequencies. However, since the higher frequencies depend on general conditions and are not always guaranteed, it can be advantageous to disable the "Turbo Mode" option for application scenarios with intensive use of AVX instructions and a high number of instructions per clock unit, as well as for those that require constant performance or lower electrical power consumption.

Memory modules (since system release)								
Memory module	Capacity [GB]	Ranks	Bit width of the memory chips	Frequency [MHz]	Low voltage	Load reduced	Registered	ECC
4GB (1x4GB) 1Rx8 DDR4-2133 U ECC	4	1	8	2133				✓
8GB (1x8GB) 2Rx8 DDR4-2133 U ECC	8	2	8	2133				✓
16GB (1x16GB) 2Rx8 DDR4-2133 U ECC	16	2	8	2133				✓

Power supplies (since system release)	Max. number
Standard PSU 300W	1
Modular PSU 450W platinum hp	2

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

Detailed technical information is available in the data sheet PRIMERGY TX1330 M2.

SPECcpu2006

Benchmark description

SPECcpu2006 is a benchmark which measures the system efficiency with integer and floating-point operations. It consists of an integer test suite (SPECint2006) containing 12 applications and a floating-point test suite (SPECfp2006) containing 17 applications. Both test suites are extremely computing-intensive and concentrate on the CPU and the memory. Other components, such as Disk I/O and network, are not measured by this benchmark.

SPECcpu2006 is not tied to a special operating system. The benchmark is available as source code and is compiled before the actual measurement. The used compiler version and their optimization settings also affect the measurement result.

SPECcpu2006 contains two different performance measurement methods: the first method (SPECint2006 or SPECfp2006) determines the time which is required to process single task. The second method (SPECint_rate2006 or SPECfp_rate2006) determines the throughput, i.e. the number of tasks that can be handled in parallel. Both methods are also divided into two measurement runs, “base” and “peak” which differ in the use of compiler optimization. When publishing the results the base values are always used; the peak values are optional.

Benchmark	Arithmetics	Type	Compiler optimization	Measurement result	Application
SPECint2006	integer	peak	aggressive	Speed	single-threaded
SPECint_base2006	integer	base	conservative		
SPECint_rate2006	integer	peak	aggressive	Throughput	multi-threaded
SPECint_rate_base2006	integer	base	conservative		
SPECfp2006	floating point	peak	aggressive	Speed	single-threaded
SPECfp_base2006	floating point	base	conservative		
SPECfp_rate2006	floating point	peak	aggressive	Throughput	multi-threaded
SPECfp_rate_base2006	floating point	base	conservative		

The measurement results are the geometric average from normalized ratio values which have been determined for individual benchmarks. The geometric average - in contrast to the arithmetic average - means that there is a weighting in favour of the lower individual results. Normalized means that the measurement is how fast is the test system compared to a reference system. Value “1” was defined for the SPECint_base2006-, SPECint_rate_base2006, SPECfp_base2006 and SPECfp_rate_base2006 results of the reference system. For example, a SPECint_base2006 value of 2 means that the measuring system has handled this benchmark twice as fast as the reference system. A SPECfp_rate_base2006 value of 4 means that the measuring system has handled this benchmark some 4/[# base copies] times faster than the reference system. “# base copies” specify how many parallel instances of the benchmark have been executed.

Not every SPECcpu2006 measurement is submitted by us for publication at SPEC. This is why the SPEC web pages do not have every result. As we archive the log files for all measurements, we can prove the correct implementation of the measurements at any time.

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY TX1330 M2
Processor	Pentium G4400 Core i3-6100 Intel® Xeon® Processor E3-1200 v5 Product Family
Memory	4 × 16GB (1x16GB) 2Rx8 DDR4-2133 U ECC
Software	
Operating system	SUSE Linux Enterprise Server 12 (x86_64)
Operating system settings	echo always > /sys/kernel/mm/transparent_hugepage/enabled
Compiler	C/C++: Version 16.0.0.101 of Intel C++ Studio XE for Linux Fortran: Version 16.0.0.101 of Intel Fortran Studio XE for Linux

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

Benchmark results

In terms of processors the benchmark result depends primarily on the size of the processor cache, the support for Hyper-Threading, the number of processor cores and on the processor frequency. In the case of processors with Turbo mode the number of cores, which are loaded by the benchmark, determines the maximum processor frequency that can be achieved. In the case of single-threaded benchmarks, which largely load one core only, the maximum processor frequency that can be achieved is higher than with multi-threaded benchmarks.

Processor	SPECint_base2006	SPECint2006	SPECint_rate_base2006	SPECint_rate2006
Pentium G4400	53.1	54.9	95.9	99.3
Core i3-6100	61.9	64.1	133	139
Xeon E3-1220 v5	65.3	67.5	196	202
Xeon E3-1225 v5	68.2	70.7	201	208
Xeon E3-1240L v5	59.7	61.6	201	208
Xeon E3-1230 v5	69.5	71.8	244	253
Xeon E3-1240 v5	71.3	73.8	249	258
Xeon E3-1270 v5	73.0	75.6	254	262
Xeon E3-1280 v5	73.4	75.8	250	259

Processor	SPECfp_base2006	SPECfp2006	SPECfp_rate_base2006	SPECfp_rate2006
Pentium G4400	70.6	71.3	104	105
Core i3-6100	85.6	87.0	131	134
Xeon E3-1220 v5	90.6	92.7	174	176
Xeon E3-1225 v5	92.6	94.9	176	178
Xeon E3-1240L v5	84.0	86.1	165	169
Xeon E3-1230 v5	94.8	96.9	188	192
Xeon E3-1240 v5	96.4	98.5	189	193
Xeon E3-1270 v5	98.0	100	191	197
Xeon E3-1280 v5	98.1	100	190	195

SPECpower_ssj2008

Benchmark description

SPECpower_ssj2008 is the first industry-standard SPEC benchmark that evaluates the power and performance characteristics of a server. With SPECpower_ssj2008 SPEC has defined standards for server power measurements 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 platforms and easy to run. The benchmark tests CPUs, caches, the memory hierarchy and scalability of symmetric multiprocessor systems (SMPs), as well as the implementation of Java Virtual Machine (JVM), Just In Time (JIT) compilers, garbage collection, threads and some aspects of the operating system.

SPECpower_ssj2008 reports power consumption for servers at different performance levels — from 100% to "active idle" in 10% 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 provides information about the energy efficiency of the measured server. The defined measurement standard enables customers to compare it with other configurations and servers measured with SPECpower_ssj2008. The diagram shows a typical graph of a SPECpower_ssj2008 result.

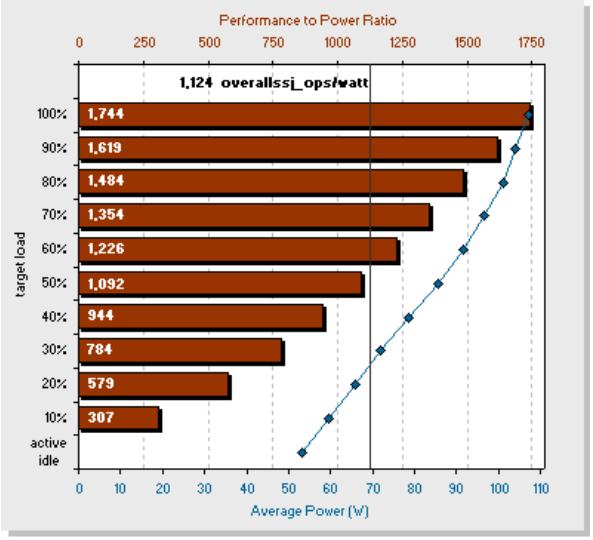

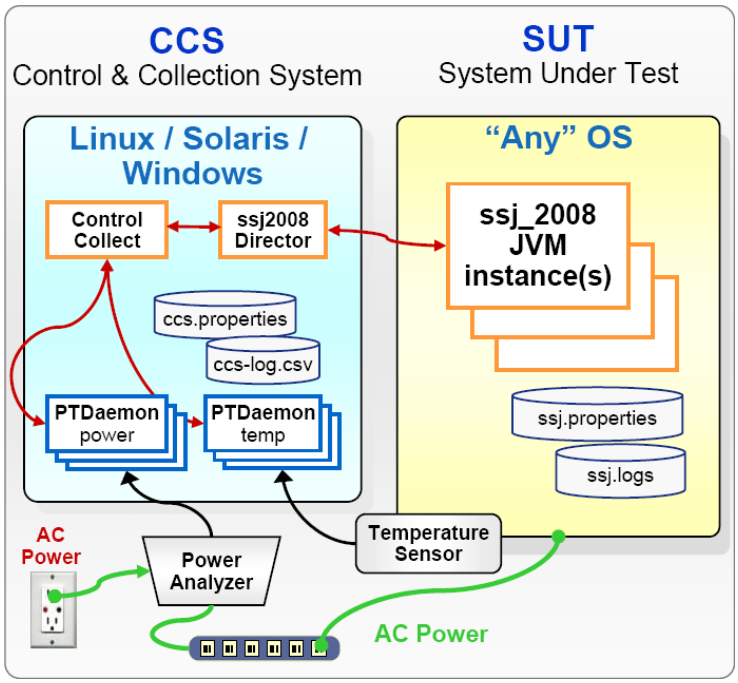

The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System Under Test (SUT) which runs one of the supported operating systems and the JVM. The JVM provides the environment required to run the SPECpower_ssj2008 workload which is implemented in Java. The other computer is a "Control & Collection System" (CCS) which controls the operation of the benchmark and captures the power, performance and temperature readings for reporting. The diagram provides an overview of the basic structure of the benchmark configuration and the various components.

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY TX1330 M2
Model version	PY TX1330M2/f/Standard PSU
Processor	Xeon E3-1240L v5
Memory	2 × 8GB (1x8GB) 2Rx8 DDR4-2133 U ECC
Network-Interface	Onboard LAN-Controller (1 port used)
Disk-Subsystem	Onboard HDD controller 1 × SSD SATA 6G 64GB DOM N H-P
Power Supply Unit	1 × Standard PSU 300W
Software	
BIOS	R1.6.0
BIOS settings	Hardware Prefetcher = Disabled Adjacent Cache Line Prefetch = Disabled DCU Streamer Prefetcher = Disabled ASPM Support = Auto DMI Control = Gen1 Intel Virtualization Technology = Disabled LAN 2 Controller = Disabled
Firmware	8.10F
Operating system	Microsoft Windows Server 2012 R2 Standard
Operating system settings	Using the local security settings console, “lock pages in memory” was enabled for the user running the benchmark. Power Management: Enabled (“Fujitsu Enhanced Power Settings” power plan) Set “Turn off hard disk after = 1 Minute” in OS. Benchmark was started via Windows Remote Desktop Connection.
JVM	Oracle Java HotSpot(TM) 64-Bit Server VM (build 24.80-b11, mixed mode), version 1.7.0_80
JVM settings	-server -Xmn10g -Xms12g -Xmx12g -XX:SurvivorRatio=60 -XX:TargetSurvivorRatio=90 -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:LoopUnrollLimit=45 -XX:InitialTenuringThreshold=12 -XX:MaxTenuringThreshold=15 -XX:ParallelGCThreads=8 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:FreqInlineSize=2500 -XX:+AggressiveOpts -XX:+UseLargePages -XX:+UseParallelOldGC -XX:-UseAdaptiveSizePolicy

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

Benchmark results

The PRIMERGY TX1330 M2 achieved the following result:

SPECpower_ssj2008 = 8,187 overall ssj_ops/watt

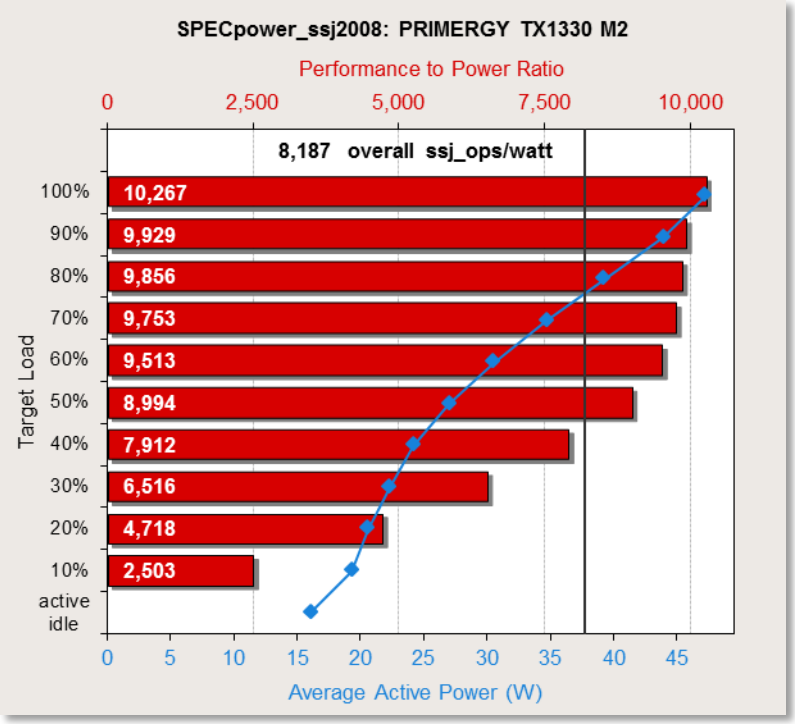

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 8,187 overall ssj_ops/watt for the PRIMERGY TX1330 M2. This is the quotient of the sum of the transaction throughputs for each load level and the sum of the average power consumed for each measurement interval.

The following table shows the benchmark results for the throughput in ssj_ops, the power consumption in watts and the resulting energy efficiency for each load level.

Performance		Power	Energy Efficiency
Target Load	ssj_ops	Average Power (W)	ssj_ops/watt
100%	484,122	47.2	10,267
90%	435,666	43.9	9,929
80%	386,466	39.2	9,856
70%	338,638	34.7	9,753
60%	289,910	30.5	9,513
50%	242,418	27.0	8,994
40%	191,117	24.2	7,912
30%	145,498	22.3	6,516
20%	96,975	20.6	4,718
10%	48,224	19.3	2,503
Active Idle	0	16.1	0
			∑ssj_ops / ∑power = 8,187

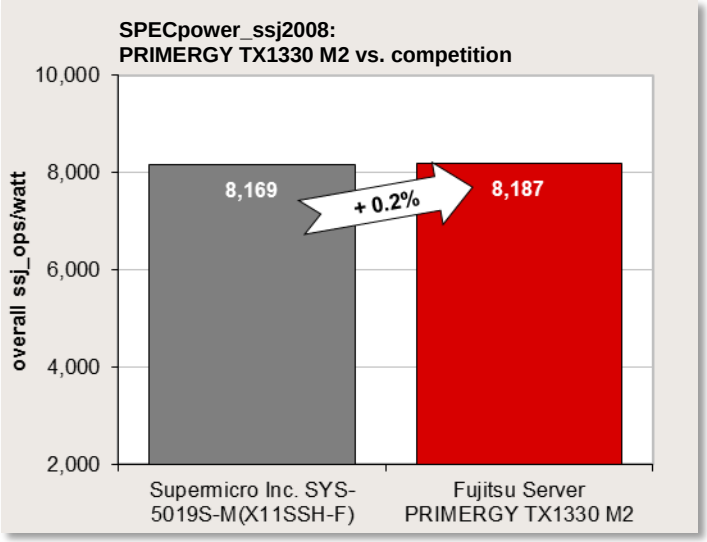

The comparison with the competition makes the advantage of the PRIMERGY TX1330 M2 in the field of energy efficiency evident. Compared to the best result of the competition, the Supermicro Inc. SYS-5019S-M(X11SSH-F) server, the PRIMERGY TX1330 M2 achieves a result with 0.2% higher energy efficiency (as of January 11th, 2016).

Disk I/O: Performance of storage media

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 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 from 1 to 512 (in steps 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 data throughput 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

Data throughput [MB/s]	*= Transaction rate [IO/s] × Block size [MB]*
Transaction rate [IO/s]	*= Data throughput [MB/s] / Block size [MB]*

This section specifies capacities of storage media on a basis of 10 (1 TB = 10^{12} bytes) while all other capacities, file sizes, block sizes and throughputs are specified on a basis of 2 (1 MB/s = 2^{20} bytes/s).

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 the measurement results discussed in this section apply for the hardware and software components listed below:

System Under Test (SUT)		
Hardware		
Model	PRIMERGY TX1330 M2	
Controller	1 × PRAID CP400i	
Storage media	SSD	HDD
	Intel SSDSC2BA100G3C Intel SSDSC2BA200G3C Intel SSDSC2BB120G4C Intel SSDSC2BB240G4C Intel SSDSC2BB480G4C Intel SSDSC2BB800G4C	HGST HUC101812CS4204 HGST HUC101812CSS204 HGST HUC101818CS4204 HGST HUC101890CS4204 HGST HUC101890CSS204 HGST HUC101830CSS204 HGST HUC101860CS4204 HGST HUC101860CSS204 HGST HUC156030CSS204 HGST HUC156045CSS204 HGST HUC156060CSS204 Seagate ST1000NM0033 Seagate ST1000NX0313 Seagate ST1000NX0333 Seagate ST2000NM0024 Seagate ST2000NM0033 Seagate ST2000NM0034 Seagate ST2000NX0253 Seagate ST2000NX0273 Seagate ST4000NM0024 Seagate ST4000NM0033 Seagate ST4000NM0034 Seagate ST6000NM0024 Seagate ST6000NM0034 Seagate ST91000640NS Seagate ST9146853SS Seagate ST9500620NS Seagate ST500DM002 Toshiba MK1401GRRB Western Digital WD1003FBYX Western Digital WD2000FYYZ Western Digital WD4000FYYZ Western Digital WD5003ABYX
Software		
Operating system	Microsoft Windows Server 2012 Standard R2	
Administration software	ServerView RAID Manager 6.2.1	
Benchmark version	3.0	
RAID type	Logical drive of type RAID 0 consisting of 1 hard disk	
Stripe size	Controller default (here 64 kB)	
Measuring tool	Iometer 1.1.0	
Measurement area	The first 10% of the usable LBA area is used for sequential accesses; the next 25% for random accesses.	
File system	raw	
Total number of Iometer workers	1	

Alignment of Iometer accesses	Aligned to whole multiples of 4096 bytes

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

Benchmark results

The results shown here are intended to help you select the appropriate storage media for the PRIMERGY TX1330 M2 under the aspect of disk-I/O performance. For this purpose, a single storage medium was measured in the configuration specified in the subsection Benchmark environment. The measurements were made using controllers which have their main features listed in the table below:

Storage medium	Controller name	Cache	Supported interfaces		RAID levels
SSD/HDD	PRAID CP400i	-	SATA 6G SAS 12G	PCIe 3.0 x8	0, 1, 1E, 10, 5, 50

Storage media

When selecting the type and number of storage media you can move the weighting in the direction of storage capacity, performance, security or price. The following types of HDD and SSD storage media can be used in the PRIMERGY TX1330 M2:

Storage medium type	Interface	Form factor
HDD	SATA 6G	2.5"
HDD	SATA 6G	3.5"
HDD	SAS 6G	2.5"
HDD	SAS 12G	2.5"
HDD	SAS 12G	3.5"
SSD	SATA 6G	2.5"

HDDs and SSDs are operated via host bus adapters, usually RAID controllers, with a SATA or SAS interface. The interface of the RAID controller to the chipset of the systemboard is typically PCIe or, in the case of the integrated onboard controllers, an internal bus interface of the systemboard.

Of all the storage medium types SSDs offer by far the highest transaction rates for random load profiles as well as the shortest access times. In return, however, the price per gigabyte of storage capacity is substantially higher.

Cache settings

In most cases, the cache of HDDs 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. To prevent data loss in case of power failure you are recommended to equip the system with a UPS.

For the purpose of easy and reliable handling of the settings for RAID controllers and hard disks it is advisable to use the RAID-Manager software “ServerView RAID” that is supplied for PRIMERGY servers. All the cache settings for controllers and hard disks can usually be made en bloc – specifically for the application – by using the pre-defined modi “Performance” or “Data Protection”. The “Performance” mode ensures the best possible performance settings for the majority of the application scenarios.

Performance values

The performance values of the PRIMERGY TX1330 M2 are summarized in the following tables, in each case specifically for a single storage medium and with various access types and block sizes. The established measurement variables, as already mentioned in the subsection Benchmark description, are used here. Thus, transaction rate is specified for random accesses and data throughput for sequential accesses. To

avoid any confusion among the measurement units the tables have been separated for the two access types.

The table cells contain the maximum achievable values. This means that each value is the maximum achievable value of the whole range of load intensities (# of Outstanding I/Os). In order to also visualize the numerical values each table cell is highlighted with a horizontal bar, the length of which is proportional to the numerical value in the table cell. All bars shown in the same scale of length have the same color. In other words, a visual comparison only makes sense for table cells with the same colored bars. Since the horizontal bars in the table cells depict the maximum achievable performance values, they are shown by the color getting lighter as you move from left to right. The light shade of color at the right end of the bar tells you that the value is a maximum value and can only be achieved under optimal prerequisites. The darker the shade becomes as you move to the left, the more frequently it will be possible to achieve the corresponding value in practice.

SSDs in comparison with the most powerful HDD

Random accesses (maximum performance values in IO/s):

PRIMERGY TX1330 M2						
Capacity [GB]	Storage device	Interface	Form factor	Transactions [IO/s]		
				Database	Fileserver	Filecopy
800	Intel SSDSC2BB800G4C	SATA 6G	2.5"	30495	3117	3960
600	*HGST HUC156060CSS204 (HDD)*	*SAS 12G*	*2.5"*	*768*	*653*	*670*
480	Intel SSDSC2BB480G4C	SATA 6G	2.5"	29575	4067	3521
240	Intel SSDSC2BB240G4C	SATA 6G	2.5"	23354	2870	2432
200	Intel SSDSC2BA200G3C	SATA 6G	2.5"	42572	5421	5147
120	Intel SSDSC2BB120G4C	SATA 6G	2.5"	15950	2026	1554
100	Intel SSDSC2BA100G3C	SATA 6G	2.5"	37420	4224	3980

Sequential accesses (maximum performance values in MB/s):

PRIMERGY TX1330 M2					
Capacity [GB]	Storage device	Interface	Form factor	Throughput [MB/s]	
				Streaming	Restore
800	Intel SSDSC2BB800G4C	SATA 6G	2.5"	478	437
600	*HGST HUC156060CSS204 (HDD)*	*SAS 12G*	2.5"	*237*	*237*
480	Intel SSDSC2BB480G4C	SATA 6G	2.5"	478	434
240	Intel SSDSC2BB240G4C	SATA 6G	2.5"	478	267
200	Intel SSDSC2BA200G3C	SATA 6G	2.5"	469	370
120	Intel SSDSC2BB120G4C	SATA 6G	2.5"	443	140
100	Intel SSDSC2BA100G3C	SATA 6G	2.5"	471	199

HDDs

Random accesses (maximum performance values in IO/s):

PRIMERGY TX1330 M2						
Capacity [GB]	Storage device	Interface	Form factor	Transactions [IO/s]		
				Database	Fileserver	Filecopy
6000	Seagate ST6000NM0034	SAS 12G	3.5"	348	313	318
6000	Seagate ST6000NM0024	SATA 6G	3.5"	301	277	287
4000	Seagate ST4000NM0034	SAS 12G	3.5"	356	320	319
4000	Seagate ST4000NM0024	SATA 6G	3.5"	310	285	295
4000	Seagate ST4000NM0033	SATA 6G	3.5"	265	245	242
4000	Western Digital WD4000FYYZ	SATA 6G	3.5"	187	176	187
2000	Seagate ST2000NM0034	SAS 12G	3.5"	352	317	316
2000	Seagate ST2000NX0273	SAS 12G	2.5"	356	308	297
2000	Seagate ST2000NM0024	SATA 6G	3.5"	311	286	295
2000	Seagate ST2000NM0033	SATA 6G	3.5"	264	240	243
2000	Western Digital WD2000FYYZ	SATA 6G	3.5"	191	177	185
2000	Seagate ST2000NX0253	SATA 6G	2.5"	329	285	290
1800	HGST HUC101818CS4204	SAS 12G	2.5"	605	517	534
1200	HGST HUC101812CS4204	SAS 12G	2.5"	651	533	543
1200	HGST HUC101812CSS204	SAS 12G	2.5"	596	495	510
1000	Seagate ST1000NX0333	SAS 12G	2.5"	363	312	301
1000	Seagate ST1000NM0033	SATA 6G	3.5"	274	249	248
1000	Western Digital WD1003FBYX	SATA 6G	3.5"	163	151	158
1000	Seagate ST1000NX0313	SATA 6G	2.5"	328	281	290
1000	Seagate ST91000640NS	SATA 6G	2.5"	302	258	243
900	HGST HUC101890CS4204	SAS 12G	2.5"	651	533	545
900	HGST HUC101890CSS204	SAS 12G	2.5"	639	539	552
600	HGST HUC156060CSS204	SAS 12G	2.5"	768	653	670
600	HGST HUC101860CS4204	SAS 12G	2.5"	678	558	567
600	HGST HUC101860CSS204	SAS 12G	2.5"	632	535	545
500	Seagate ST500DM002	SATA 6G	3.5"	214	196	195
500	Western Digital WD5003ABYX	SATA 6G	3.5"	199	182	191
500	Seagate ST9500620NS	SATA 6G	2.5"	300	257	249
450	HGST HUC156045CSS204	SAS 12G	2.5"	744	608	631
300	HGST HUC156030CSS204	SAS 12G	2.5"	758	625	636
300	HGST HUC101830CSS204	SAS 12G	2.5"	670	567	577
146	Seagate ST9146853SS	SAS 6G	2.5"	655	545	522
146	Toshiba MK1401GRRB	SAS 6G	2.5"	603	506	492

Sequential accesses (maximum performance values in MB/s):

PRIMERGY TX1330 M2					
Capacity [GB]	Storage device	Interface	Form factor	Throughput [MB/s]	
				Streaming	Restore
6000	Seagate ST6000NM0034	SAS 12G	3.5"	210	210
6000	Seagate ST6000NM0024	SATA 6G	3.5"	209	209
4000	Seagate ST4000NM0034	SAS 12G	3.5"	211	211
4000	Seagate ST4000NM0024	SATA 6G	3.5"	220	219
4000	Seagate ST4000NM0033	SATA 6G	3.5"	183	183
4000	Western Digital WD4000FYYZ	SATA 6G	3.5"	162	161
2000	Seagate ST2000NM0034	SAS 12G	3.5"	213	212
2000	Seagate ST2000NX0273	SAS 12G	2.5"	133	134
2000	Seagate ST2000NM0024	SATA 6G	3.5"	204	204
2000	Seagate ST2000NM0033	SATA 6G	3.5"	179	179
2000	Western Digital WD2000FYYZ	SATA 6G	3.5"	157	157
2000	Seagate ST2000NX0253	SATA 6G	2.5"	133	134
1800	HGST HUC101818CS4204	SAS 12G	2.5"	242	242
1200	HGST HUC101812CS4204	SAS 12G	2.5"	238	238
1200	HGST HUC101812CSS204	SAS 12G	2.5"	211	211
1000	Seagate ST1000NX0333	SAS 12G	2.5"	133	134
1000	Seagate ST1000NM0033	SATA 6G	3.5"	184	184
1000	Western Digital WD1003FBYX	SATA 6G	3.5"	130	129
1000	Seagate ST1000NX0313	SATA 6G	2.5"	131	131
1000	Seagate ST91000640NS	SATA 6G	2.5"	108	108
900	HGST HUC101890CS4204	SAS 12G	2.5"	233	233
900	HGST HUC101890CSS204	SAS 12G	2.5"	212	211
600	HGST HUC156060CSS204	SAS 12G	2.5"	237	237
600	HGST HUC101860CS4204	SAS 12G	2.5"	230	230
600	HGST HUC101860CSS204	SAS 12G	2.5"	210	210
500	Seagate ST500DM002	SATA 6G	3.5"	138	138
500	Western Digital WD5003ABYX	SATA 6G	3.5"	130	127
500	Seagate ST9500620NS	SATA 6G	2.5"	111	111
450	HGST HUC156045CSS204	SAS 12G	2.5"	237	237
300	HGST HUC156030CSS204	SAS 12G	2.5"	236	236
300	HGST HUC101830CSS204	SAS 12G	2.5"	210	210
146	Seagate ST9146853SS	SAS 6G	2.5"	195	195
146	Toshiba MK1401GRRB	SAS 6G	2.5"	199	198

Disk I/O: Performance of RAID controllers

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 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 from 1 to 512 (in steps 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 data throughput 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

Data throughput [MB/s]	*= Transaction rate [IO/s] × Block size [MB]*
Transaction rate [IO/s]	*= Data throughput [MB/s] / Block size [MB]*

This section specifies capacities of storage media on a basis of 10 (1 TB = 10^{12} bytes) while all other capacities, file sizes, block sizes and throughputs are specified on a basis of 2 (1 MB/s = 2^{20} bytes/s).

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 the measurement results discussed in this chapter were determined using the hardware and software components listed below:

System Under Test (SUT)	
Hardware	
Controller	1 × “LSI SW RAID on Intel C236 (Onboard SATA)” Intel C236 PCH, Code name Sunrise Point Driver name: megasr1.sys, Driver version: 17.01.2015.0716 BIOS version: A.15.08211538R 1 × „PRAID CP400i“ 1 × „PRAID EP400i“ 1 × „PRAID EP420i“ Driver name: megasas2.sys, Driver version: 6.706.06 Firmware package: 24.7.0-0061
Drive	4 × 3.5" SATA SSD Intel SSDSC2BA400G3 8 × 3.5" SATA SSD Intel SSDSC2BB480G4C 4 × 3.5" SATA HDD Seagate ST1000NM0033 4 × 3.5" SAS HDD HGST HUC156045CSS204 4 × 2.5" SATA SSD Intel SSDSC2BA400G3 24 × 2.5" SATA SSD Intel SSDSC2BB480G4C 4 × 2.5" SATA HDD Seagate ST91000640NS 24 × 2.5" SAS HDD HGST HUC156045CSS204
Software	
Operating system	Microsoft Windows Server 2012 Standard R2
Administration software	ServerView RAID Manager 6.2.1
Benchmark version	3.0
Stripe size	Controller default
Measuring tool	Iometer 1.1.0
Measurement area	The first 10% of the usable LBA area is used for sequential accesses; the next 25% for random accesses.
File system	raw
Total number of Iometer workers	1
Alignment of Iometer accesses	Aligned to whole multiples of 4096 bytes

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

Benchmark results

The results presented here are designed to help you choose the right solution from the various configuration options of the PRIMERGY TX1330 M2 in the light of disk-I/O performance. Various combinations of RAID controllers and storage media will be analyzed below. Information on the selection of storage media themselves is to be found in the section “Disk I/O: Performance of storage media”.

Hard disks

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).

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.

More hard disks per system are possible as a result of using 2.5" hard disks instead of 3.5" hard disks. Consequently, the load that each individual hard disk has to overcome decreases and the maximum overall performance of the system increases.

More detailed performance statements about hard disk types are available in the section “Disk I/O: Performance of storage media” in this performance report.

Model versions

The maximum number of hard disks in the system depends on the system configuration. The following table lists the essential cases.

Form factor	Interface	Connection type	Number of PCIe controllers	Maximum number of hard disks
2.5", 3.5"	SATA 6G	direct	0	4
2.5", 3.5"	SATA 6G, SAS 12G	direct	1	8
3.5"	SATA 6G, SAS 12G	Expander	1	12
2.5"	SATA 6G, SAS 12G	Expander	1	24

RAID controller

In addition to the hard disks the RAID controller is the second performance-determining key component. In the case of these controllers the “modular RAID” concept of the PRIMERGY servers offers a plethora of options to meet the various requirements of a wide range of different application scenarios.

The following table summarizes the most important features of the available RAID controllers of the PRIMERGY TX1330 M2. 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		In the system		FBU
					Max. # disks per controller	RAID levels	
LSI SW RAID on Intel C236 (Onboard SATA)	Sunrise Point	-	SATA 6G	-	4 × 2.5" 4 × 3.5"	0, 1, 10	-
PRAID CP400i	PRAID CP400i	-	SATA 6G SAS 12G	PCIe 3.0 x8	8 × 2.5" 8 × 3.5"	0, 1, 1E, 5, 10, 50	-
PRAID EP400i	PRAID EP400i	1 GB	SATA 6G SAS 12G	PCIe 3.0 x8	24 × 2.5" 12 × 3.5"	0, 1, 1E, 5, 6, 10, 50, 60	✓
PRAID EP420i	PRAID EP420i	2 GB	SATA 6G SAS 12G	PCIe 3.0 x8	24 × 2.5" 12 × 3.5"	0, 1, 1E, 5, 6, 10, 50, 60	✓

The onboard RAID controller is implemented in the chipset Intel C236 on the system board of the server and uses the CPU of the server for the RAID functionality. This controller is a simple solution that does not require a PCIe slot.

System-specific interfaces

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

Controller alias	Effective in the configuration				Connection via expander
	# Disk-side data channels	Limit for throughput of disk interface	# CPU-side data channels	Limit for throughput of CPU-side interface	
Sunrise Point	4 × SATA 6G	**2060 MB/s**	4 × DMI 3.0	3433 MB/s	-
PRAID CP400i	8 × SAS 12G	8240 MB/s	8 × PCIe 3.0	**6761 MB/s**	-
PRAID EP400i	8 × SAS 12G	8240 MB/s	8 × PCIe 3.0	**6761 MB/s**	-/✓
PRAID EP420i	8 × SAS 12G	8240 MB/s	8 × PCIe 3.0	**6761 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 it available in total to all connected hard disks.

More details about the RAID controllers of the PRIMERGY systems are available in the white paper “RAID Controller Performance”.

Settings

In most cases, the cache of HDDs 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. 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, the data temporarily stored in the cache should be safeguarded against loss in case of power failure. Suitable accessories are available for this purpose (e.g. a FBU).

For the purpose of easy and reliable handling of the settings for RAID controllers and hard disks it is advisable to use the software “ServerView RAID Manager” that is supplied for PRIMERGY servers. All the cache settings for controllers and hard disks can usually be made en bloc – specifically for the application – by using the pre-defined modi “Performance”, “Data Protection” or “Fast Path optimum”. The “Performance” mode ensures the best possible performance settings for the majority of the application scenarios. In connection with the “FastPath” RAID controller option, the “Fast Path optimum” mode should be selected if maximum transaction rates are to be achieved with SSDs for random accesses with small blocks (≤ 8 kB, e. g. OLTP operation of databases).

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 TX1330 M2 if the configurations measured there are also supported by this system.

The performance values of the PRIMERGY TX1330 M2 are listed in table form below, specifically for different RAID levels, access types and block sizes. Substantially different configuration versions are dealt with separately. The established measurement variables, as already mentioned in the subsection Benchmark description, are used here. Thus, transaction rate is specified for random accesses and data throughput for sequential accesses. To avoid any confusion among the measurement units the tables have been separated for the two access types.

The table cells contain the maximum achievable values. This has three implications: On the one hand hard disks with optimal performance were used (the components used are described in more detail in the subsection 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. And ultimately each value is the maximum value for the entire load intensity range (# of outstanding I/Os).

In order to also visualize the numerical values each table cell is highlighted with a horizontal bar, the length of which is proportional to the numerical value in the table cell. All bars shown in the same scale of length have the same color. In other words, a visual comparison only makes sense for table cells with the same colored bars.

Since the horizontal bars in the table cells depict the maximum achievable performance values, they are shown by the color getting lighter as you move from left to right. The light shade of color at the right end of the bar tells you that the value is a maximum value and can only be achieved under optimal prerequisites. The darker the shade becomes as you move to the left, the more frequently it will be possible to achieve the corresponding value in practice.

2.5" - Random accesses (maximum performance values in IO/s):

PRIMERGY TX1330 M2							
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	Hard disk type	#Disks					
C236	ST91000640NS SATA HDD SSDSC2BA400G3 SATA SSD	2	1	380	332	42226	6879
		4	0	683	390	77940	14370
		4	10	562	323	61093	10684
PRAID CP400i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	1153	971	37739	4685
		8	10	4294	2229	115104	20736
		8	0	4784	2531	142935	25238
		8	5	2435	1382	28848	12990
PRAID EP400i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	1590	888	38230	4953
		8	10	4623	3265	162574	20132
		8	0	5291	3784	185132	28118
		8	5	3008	2097	73854	10405
		16	10	9166	6377	N/A	N/A
		24	0	14086	10367	N/A	N/A
		24	5	8017	5582	N/A	N/A
PRAID EP420i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	1544	994	38253	4813
		8	10	4616	3213	163936	19886
		8	0	5230	3729	187813	28664
		8	5	2970	2039	73612	10381
		16	10	9360	6461	N/A	N/A
		24	0	14115	10381	N/A	N/A
		24	5	7884	5622	N/A	N/A

(*cursive*: calculated)

2.5" - Sequential accesses (maximum performance values in MB/s):

PRIMERGY TX1330 M2							
Configuration version				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	Hard disk type	#Disks	RAID level				
C236	ST91000640NS SATA HDD SSDSC2BA400G3 SATA SSD	2	1	113	108	746	428
		4	0	427	423	1361	1620
		4	10	224	212	1047	784
PRAID CP400i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	327	230	914	397
		8	10	1051	892	3397	1274
		8	0	1795	1779	3377	2534
		8	5	1579	1558	3375	1734
PRAID EP400i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	350	232	914	403
		8	10	1148	945	3605	1284
		8	0	1874	1892	3602	2538
		8	5	1648	1658	3584	2211
		16	10	2097	1888	N/A	N/A
		24	0	5395	5356	N/A	N/A
		24	5	5178	3169	N/A	N/A
PRAID EP420i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	375	232	910	389
		8	10	1174	942	3589	1271
		8	0	1871	1892	3589	2548
		8	5	1648	1650	3589	2131
		16	10	2114	1889	N/A	N/A
		24	0	5396	5386	N/A	N/A
		24	5	5172	3093	N/A	N/A

(*cursive*: calculated)

3.5" - Random accesses (maximum performance values in IO/s):

PRIMERGY TX1330 M2							
Configuration version							
RAID Controller	Hard disk type	#Disks	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]
C236	ST1000NM0033 SATA HDD SSDSC2BA400G3 SATA SSD	2	1	351	321	42226	6879
		4	0	634	364	77940	14370
		4	10	544	315	61093	10684
PRAID CP400i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	1153	971	37739	4685
		8	10	4294	2229	115104	20736
		8	0	4784	2531	142935	25238
		8	5	2435	1382	28848	12990
PRAID EP400i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	1590	888	38230	4953
		8	10	4623	3265	162574	20132
		8	0	5291	3784	185132	28118
		8	5	3008	2097	73854	10405
PRAID EP420i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	1544	994	38253	4813
		8	10	4616	3213	163936	19886
		8	0	5230	3729	187813	28664
		8	5	2970	2039	73612	10381

(*cursive*: calculated)

3.5" - Sequential accesses (maximum performance values in MB/s):

PRIMERGY TX1330 M2							
Configuration version				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	Hard disk type	#Disks	RAID level				
C236	ST1000NM0033 SATA HDD SSDSC2BA400G3 SATA SSD	2	1	187	176	746	428
		4	0	693	693	1361	1620
		4	10	365	344	1047	784
PRAID CP400i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	327	230	914	397
		8	10	1051	892	3397	1274
		8	0	1795	1779	3377	2534
		8	5	1579	1558	3375	1734
PRAID EP400i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	350	232	914	403
		8	10	1148	945	3605	1284
		8	0	1874	1892	3602	2538
		8	5	1648	1658	3584	2211
PRAID EP420i	HUC156045CSS204 SAS HDD SSDSC2BB480G4C SATA SSD	2	1	375	232	910	389
		8	10	1174	942	3589	1271
		8	0	1871	1892	3589	2548
		8	5	1648	1650	3589	2131

(*cursive*: calculated)

Conclusion

At full configuration with powerful hard disks the PRIMERGY TX1330 M2 achieves a throughput of up to 5396 MB/s for sequential load profiles and a transaction rate of up to 187813 IO/s for typical, random application scenarios.

For best possible performance we recommend one of the plug-in PRAID controllers. To operate SSDs within the maximum performance range the PRAID CP400i is already suited for the simpler RAID levels 0, 1 and 10, and a PRAID controller with cache is to be preferred for RAID 5.

In the event of HDDs the controller cache for random load profiles with a significant write share has performance advantages for all RAID levels.

STREAM

Benchmark description

STREAM is a synthetic benchmark that has been used for many years to determine memory throughput and which was developed by John McCalpin during his professorship at the University of Delaware. Today STREAM is supported at the University of Virginia, where the source code can be downloaded in either Fortran or C. STREAM continues to play an important role in the HPC environment in particular. It is for example an integral part of the HPC Challenge benchmark suite.

The benchmark is designed in such a way that it can be used both on PCs and on server systems. The unit of measurement of the benchmark is GB/s, i.e. the number of gigabytes that can be read and written per second.

STREAM measures the memory throughput for sequential accesses. These can generally be performed more efficiently than accesses that are randomly distributed on the memory, because the processor caches are used for sequential access.

Before execution the source code is adapted to the environment to be measured. Therefore, the size of the data area must be at least 12 times larger than the total of all last-level processor caches so that these have as little influence as possible on the result. The OpenMP program library is used to enable selected parts of the program to be executed in parallel during the runtime of the benchmark, consequently achieving optimal load distribution to the available processor cores.

During implementation the defined data area, consisting of 8-byte elements, is successively copied to four types, and arithmetic calculations are also performed to some extent.

Type	Execution	Bytes per step	Floating-point calculation per step
COPY	a(i) = b(i)	16	0
SCALE	a(i) = q × b(i)	16	1
SUM	a(i) = b(i) + c(i)	24	1
TRIAD	a(i) = b(i) + q × c(i)	24	2

The throughput is output in GB/s for each type of calculation. The differences between the various values are usually only minor on modern systems. In general, only the determined TRIAD value is used as a comparison.

The measured results primarily depend on the clock frequency of the memory modules; the processors influence the arithmetic calculations.

This chapter specifies throughputs on a basis of 10 (1 GB/s = 10^9 Byte/s).

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY TX1330 M2
Processor	Pentium G4400 Core i3-6100 Intel® Xeon® Processor E3-1200 v5 Product Family
Memory	4 × 16GB (1x16GB) 2Rx8 DDR4-2133 U ECC
Software	
Operating system	SUSE Linux Enterprise Server 12 (x86_64)
Operating system settings	Transparent Huge Pages inactivated
Compiler	Intel C++ Composer XE 2015 for Linux
Benchmark	Stream.c Version 5.10

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

Benchmark results

Processor	Memory Frequency [MHz]	Max. Memory Bandwidth [GB/s]	Cores	Processor Frequency [GHz]	TRIAD [GB/s]
Pentium G4400	2133	34.1	2	3.30	30.2
Core i3-6100	2133	34.1	2	3.70	30.3
Xeon E3-1240L v3	2133	34.1	4	2.10	29.2
Xeon E3-1220 v3	2133	34.1	4	3.00	29.8
Xeon E3-1225 v3	2133	34.1	4	3.30	29.9
Xeon E3-1230 v3	2133	34.1	4	3.40	29.9
Xeon E3-1240 v3	2133	34.1	4	3.50	30.0
Xeon E3-1270 v5	2133	34.1	4	3.60	30.0
Xeon E3-1280 v5	2133	34.1	4	3.70	30.0

Literature

PRIMERGY Servers

http://primergy.com/

PRIMERGY TX1330 M2

This White Paper:
http://docs.ts.fujitsu.com/dl.aspx?id=0aeedb9d-cf47-4463-9d67-1db099d60cfb
http://docs.ts.fujitsu.com/dl.aspx?id=ada803f2-5e08-433f-ad1d-687bdfde37ea
http://docs.ts.fujitsu.com/dl.aspx?id=0200cd8e-4fb7-46f5-a8b1-f89469b3051d

Data sheet
http://docs.ts.fujitsu.com/dl.aspx?id=29bfbf8a-2991-4e6b-8177-0ba62f2ad1bb

PRIMERGY Performance

http://www.fujitsu.com/fts/x86-server-benchmarks

Performance of Server Components

http://www.fujitsu.com/fts/products/computing/servers/mission-critical/benchmarks/x86-components.html

RAID Controller Performance 2013
http://docs.ts.fujitsu.com/dl.aspx?id=e2489893-cab7-44f6-bff2-7aeea97c5aef

RAID Controller Performance 2016
http://docs.ts.fujitsu.com/dl.aspx?id=9845be50-7d4f-4ef7-ac61-bbde399c1014

Disk I/O: Performance of storage media and RAID controllers

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

Information about Iometer
http://www.iometer.org

SPECcpu2006

http://www.spec.org/osg/cpu2006

Benchmark overview SPECcpu2006
http://docs.ts.fujitsu.com/dl.aspx?id=1a427c16-12bf-41b0-9ca3-4cc360ef14ce

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

STREAM

http://www.cs.virginia.edu/stream/

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

2016-02-29 WW EN