Fujitsu Server PRIMERGY
Performance Report
PRIMERGY RX1330 M5

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMERGY RX1330 M5.

Explaines PRIMERGY RX1330 M5 performance data in comparison to other PRIMERGY models. In addition to the benchmark results, the explanation for each benchmark and benchmark environment are also included.

Version
1.3
2023-10-03

Contents

Technical data

PRIMERGY PRIMERGY RX1330 M5

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 RX1330 M5
Form factor	Rack server
Chipset	Intel C256
Number of sockets	1
Number of processors orderable	1
Processor type	Intel Xeon E-2300 processor family / Intel Pentium Gold G6405
Number of memory slots	4
Maximum memory configuration	128 GB
Onboard HDD controller	Controller with RAID (0/1, or 5/6) function
PCI slots	PCI-Express 4.0 (x8 lane): 2 (Low Profile) PCI-Express 4.0 (x4 lane): 1 (Low Profile)
Max. number of internal storage	10 x 2.5 inches or 8 x 2.5 inches or 4 x 2.5 inches or 4 x 3.5 inches

Processor model	Number of cores	Number of threads	Cache	Rated frequency	Maximum turbo frequency	Maximum memory frequency	TDP
			[MB]	[GHz]	[GHz]	[MHz]	[W]
Xeon E-2388G	8	16	12	3.2	5.1	3,200	95
Xeon E-2386G	6	12	12	3.5	5.1	3,200	95
Xeon E-2378G	8	16	16	2.8	5.1	3,200	80
Xeon E-2378	8	16	16	2.6	4.8	3,200	65
Xeon E-2374G	4	8	8	3.7	5.0	3,200	80
Xeon E-2356G	6	12	12	3.2	5.0	3,200	80
Xeon E-2336	6	12	12	2.9	4.8	3,200	65
Xeon E-2334	4	8	8	3.4	4.8	3,200	65
Xeon E-2324G	4	4	8	3.1	4.6	3,200	65
Xeon E-2314	4	4	8	2.8	4.5	3,200	65
Pentium Gold G6405	2	4	4	4.1	-	2,666	58

All processors that can be ordered with PRIMERGY RX1330 M5 support Intel Turbo Boost Technology 2.0.

This technology allows you to operate the processor with higher frequencies than the rated frequency. The "maximum turbo frequency" listed in the processor list above is the theoretical maximum frequency when there is only one active core per processor. The maximum frequency that can actually be achieved depends on the number of active cores, current consumption, power consumption, and processor temperature.

As a general rule, Intel does not guarantee that maximum turbo frequencies will be achieved. This is related to manufacturing tolerances, and the performance of each individual processor model varies from each other.

The range of difference covers the range including all of the rated frequency and the maximum turbo frequency.

The turbo function can be set in the BIOS option. Generally, Fujitsu always recommends leaving the [Turbo Mode] option set at the standard setting [Enabled], as performance is substantially increased by the higher frequencies. However, the Turbo Mode frequency depends on the operating conditions mentioned above and is not always guaranteed. The turbo frequency fluctuates in applications where AVX instructions are used intensively and the number of instructions per clock is large. If you need stable performance or want to reduce power consumption, it may be beneficial to set the [Turbo Mode] option to [Disabled] to disable the turbo function.

Memory modules									
Type	Capacity	Number of ranks	Bit width of the memory chips	Frequency	3DS	Load Reduced	Registered	NVDIMM	ECC
	[GB]			[MHz]					
8 GB (1x8 GB) 1Rx8 DDR4-3200 U ECC	8	1	8	3,200					✓
16 GB (1x16 GB) 1Rx8 DDR4-3200 U ECC	16	1	8	3,200					✓
16 GB (1x16 GB) 2Rx8 DDR4-3200 U ECC	16	2	8	3,200					✓
32 GB (1x 32 GB) 2Rx8 DDR4-3200 U ECC	32	2	8	3,200					✓

Power supplies	Maximum number
Standard PSU 300W	1
Modular PSU 500W platinum	2
Modular PSU 900W platinum	2

Includes components that will be supported after the system release. Also, some components may not be available in all countries or sales regions.

Detailed technical information is available in the data sheet of PRIMERGY RX1330 M5.

SPEC CPU2017

Benchmark description

SPEC CPU2017 is a benchmark which measures the system efficiency with integer and floating-point operations. It consists of an integer test suite (SPECrate 2017 Integer, SPECspeed 2017 Integer) containing 10 applications and a floating-point test suite (SPECrate 2017 Floating Point, SPECspeed 2017 Floating Point) containing 14 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.

SPEC CPU2017 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.

SPEC CPU2017 contains two different performance measurement methods. The first method (SPECspeed 2017 Integer or SPECspeed 2017 Floating Point) determines the time which is required to process a single task. The second method (SPECrate 2017 Integer or SPECrate 2017 Floating Point) 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." They differ in the use of compiler optimization. When publishing the results, the base values are always used and the peak values are optional.

Benchmark	Number of single benchmarks	Arithmetics	Type	Compiler optimization	Measurement result
SPECspeed2017_int_peak	10	Integer	Aggressive (peak)	Speed	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_base	10		Conservative (base)		Performance
SPECspeed2017_int_energy_base					Power efficiency
SPECspeed2017_int_peak	10		Aggressive (peak)	Throughput	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_base	10		Conservative (base)		Performance
SPECspeed2017_int_energy_base					Power efficiency
SPECspeed2017_int_peak	10	Floating point	Aggressive (peak)	Speed	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_base	10		Conservative (base)		Performance
SPECspeed2017_int_energy_base					Power efficiency
SPECspeed2017_int_peak	13		Aggressive (peak)	Throughput	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_base	13		Conservative (base)		Performance
SPECspeed2017_int_energy_base					Power efficiency

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 favor of the lower individual results. "Normalized" means that the measurement is how fast is the test system compared to a reference system. For example, value "1" was defined for the SPECspeed2017_int_base, SPECrate2017_int_base, SPECspeed2017_fp_base, and SPECrate2017_fp_base results of the reference system. A SPECspeed2017_int_base value of 2 means that the measuring system has handled this benchmark twice as fast as the reference system. A SPECrate2017_fp_base value of 4 means that the measuring system has handled this benchmark about 4/[# base copies] times faster than the reference system. "# base copies" specifies how many parallel instances of the benchmark have been executed.

Not every SPEC CPU2017 measurement is submitted by Fujitsu for publication at SPEC. This is why the SPEC web pages do not have every result. As Fujitsu archives the log files for all measurements, it is possible to prove the correct implementation of the measurements at any time.

Benchmark environment

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY RX1330 M5
• Processor	Intel Xeon E-2300 processor family / Intel Pentium Gold G6405
• Memory	2 x 16 GB 2Rx8 PC4-3200AA-E
Software	
• BIOS settings	SPECspeed2017_int_base: • Energy Efficient Turbo = Disabled • SA GV High Gear = Gear1 • FAN Control = Full SPECSpeed2017_fp_base: • Hyper Threading = Disabled • Package C-State un-demotion = Enabled • REFRESH_2X_MODE = 2- Enabled HOT only • FAN Control = Full SPECrate2017_int_base: • Adjacent Cache Line Prefetch = Disabled • Package C-State limit = C6 • Per Core P State OS control mode = Disabled • FAN Control = Full SPECrate2017_fp_base: • Hyper Threading = Disabled • C-States Auto Demotion = Disabled • C-States Un Demotion = Disabled • DDR Speed Control = Auto • DMI Gen3 ASPM = ASPM L0s • FAN Control = Full
• Operating system	SUSE Linux Enterprise Server 15 SP3 5.3.18-57-default
• Operating system settings	Stack size set to unlimited using "ulimit -s unlimited" SPECspeed2017_int, SPECrate2017_int: cpupower -c all frequency-set -g performance SPECrate2017_fp: echo madvise > /sys/kernel/mm/transparent_hugepage/enabled
• Compiler	Fortran: Version 2021.1 of Intel Fortran Compiler for Linux SPECspeed2017_fp: C/C++: Version 2021.1 of Intel C/C++ Compiler for Linux SPECspeed2017_int, SPECrate: C/C++: Version 2021.1 of Intel oneAPI DPC++/C++ Compiler 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 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.

The results with "est." are the estimated values.

Processor	Number of cores	SPECrate2017 int_base	SPECrate2017 fp_base
Xeon E-2388G	8	**69.7**	**59.8**
Xeon E-2386G	6	58.3 est.	53.6 est.
Xeon E-2378G	8	65.8 est.	57.6 est.
Xeon E-2378	8	57.5 est.	53.3 est.
Xeon E-2374G	4	42.6 est.	43.7 est.
Xeon E-2356G	6	56.1 est.	52.3 est.
Xeon E-2336	6	52.9 est.	50.5 est.
Xeon E-2334	4	39.6 est.	41.7 est.
Xeon E-2324G	4	34.5 est.	42.1 est.
Xeon E-2314	4	29.2 est.	38.4 est.
Pentium Gold G6405	2	16.3 est.	17.1 est.

Processor	Number of cores	SPECspeed2017 int_base	SPECspeed2017 fp_base
Xeon E-2388G	8	**15.8**	**44.8**
Xeon E-2386G	6	15.2 est.	41.1 est.
Xeon E-2378G	8	15.6 est.	43.2 est.
Xeon E-2378	8	14.8 est.	40.6 est.
Xeon E-2374G	4	14.6 est.	34.4 est.
Xeon E-2356G	6	14.8 est.	40.3 est.
Xeon E-2336	6	14.4 est.	39.0 est.
Xeon E-2334	4	14.0 est.	33.2 est.
Xeon E-2324G	4	13.1 est.	33.5 est.
Xeon E-2314	4	12.7 est.	30.4 est.
Pentium Gold G6405	2	8.85 est.	14.3 est.

The following graph compares the throughput of PRIMERGY RX1330 M5 and its older model, PRIMERGY RX1330 M4, with maximum performance configurations.

SPECrate2017: Comparison of PRIMERGY RX1330 M5 and PRIMERGY RX1330 M4,

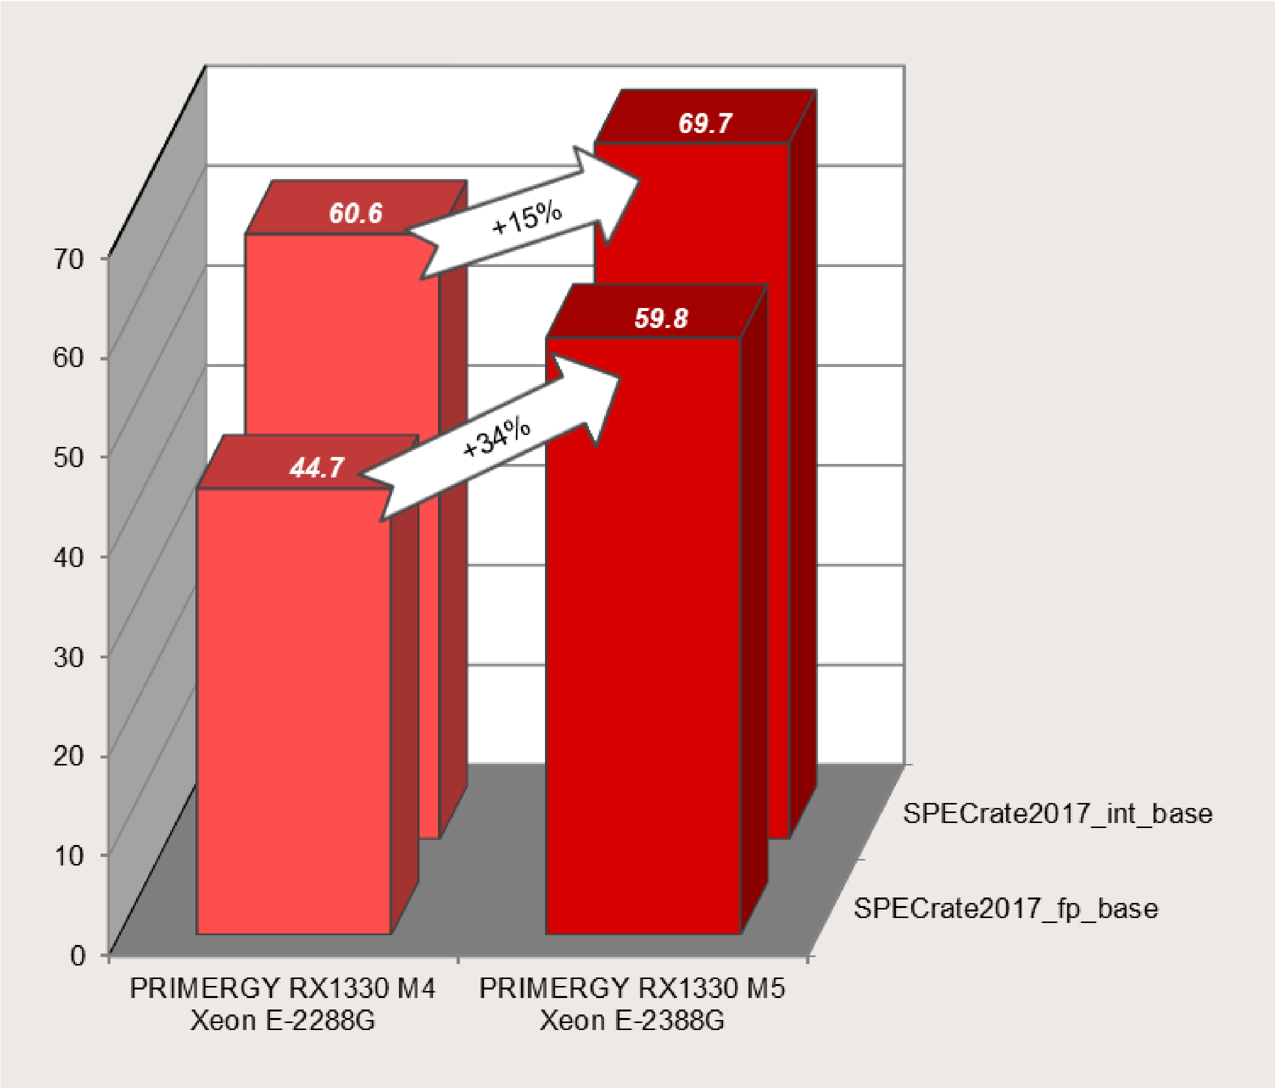

Measurement results of SPECcpu2017 (January 5, 2022)

On January 5th, 2022, PRIMERGY RX1330 M5 with an Intel Xeon E-2388G processor won first place in the 1-socket Intel Xeon category of the SPECspeed2017_int_base benchmark.

STREAM

Benchmark description

STREAM is a synthetic benchmark that has been used for many years to determine memory throughput and 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. This provides optimal load distribution for the available processor cores.

In the STREAM benchmark, a data area consisting of 8-byte elements is continuously copied to four operation types. Arithmetic operations are also performed on operation types other than COPY.

Arithmetics type	Arithmetics	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.

In this chapter, throughputs are indicated as a power of 10. (1 GB/s = 10^9 Byte/s)

Benchmark environment

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY RX1330 M5
• Processor	Intel Xeon E-2300 processor family / Intel Pentium Gold G6405
• Memory	2 x 32 GB 2Rx8 PC4-3200AA-E
Software	
• BIOS settings	• FAN Control = Full
• Operating system	SUSE Linux Enterprise Server 15 SP3 5.3.18-57-default
• Operating system settings	Default
• Compiler	C/C++: Version 2021.1 of Intel C/C++ Compiler for Linux
• Benchmark	STREAM Version 5.10

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

Benchmark results

The results with "est." are the estimated values.

Processor	Memory frequency [MHz]	Maximum memory bandwidth [GB/s]	Number of cores	Rated frequency [GHz]	TRIAD [GB/s]	
Xeon E-2388G	3,200	51.2	8	3.2	43.8	est.
Xeon E-2386G	3,200	51.2	6	3.5	44.1	est.
Xeon E-2378G	3,200	51.2	8	2.8	43.8	est.
Xeon E-2378	3,200	51.2	8	2.6	43.7	est.
Xeon E-2374G	3,200	51.2	4	3.7	43.2	est.
Xeon E-2356G	3,200	51.2	6	3.2	44.2	est.
Xeon E-2336	3,200	51.2	6	2.9	44.1	est.
Xeon E-2334	3,200	51.2	4	3.4	42.9	est.
Xeon E-2324G	2,933	51.2	4	3.1	43.0	est.
Xeon E-2314	3,200	51.2	4	2.8	43.1	est.
Pentium Gold G6405	3,200	42.7	2	4.1	35.3	est.

The following diagram illustrates the throughput of the PRIMERGY RX1330 M5 in comparison to its predecessor, the RX1330 M4,.

STREAM TRIAD: Comparison of PRIMERGY RX1330 M5 and PRIMERGY RX1330 M4,

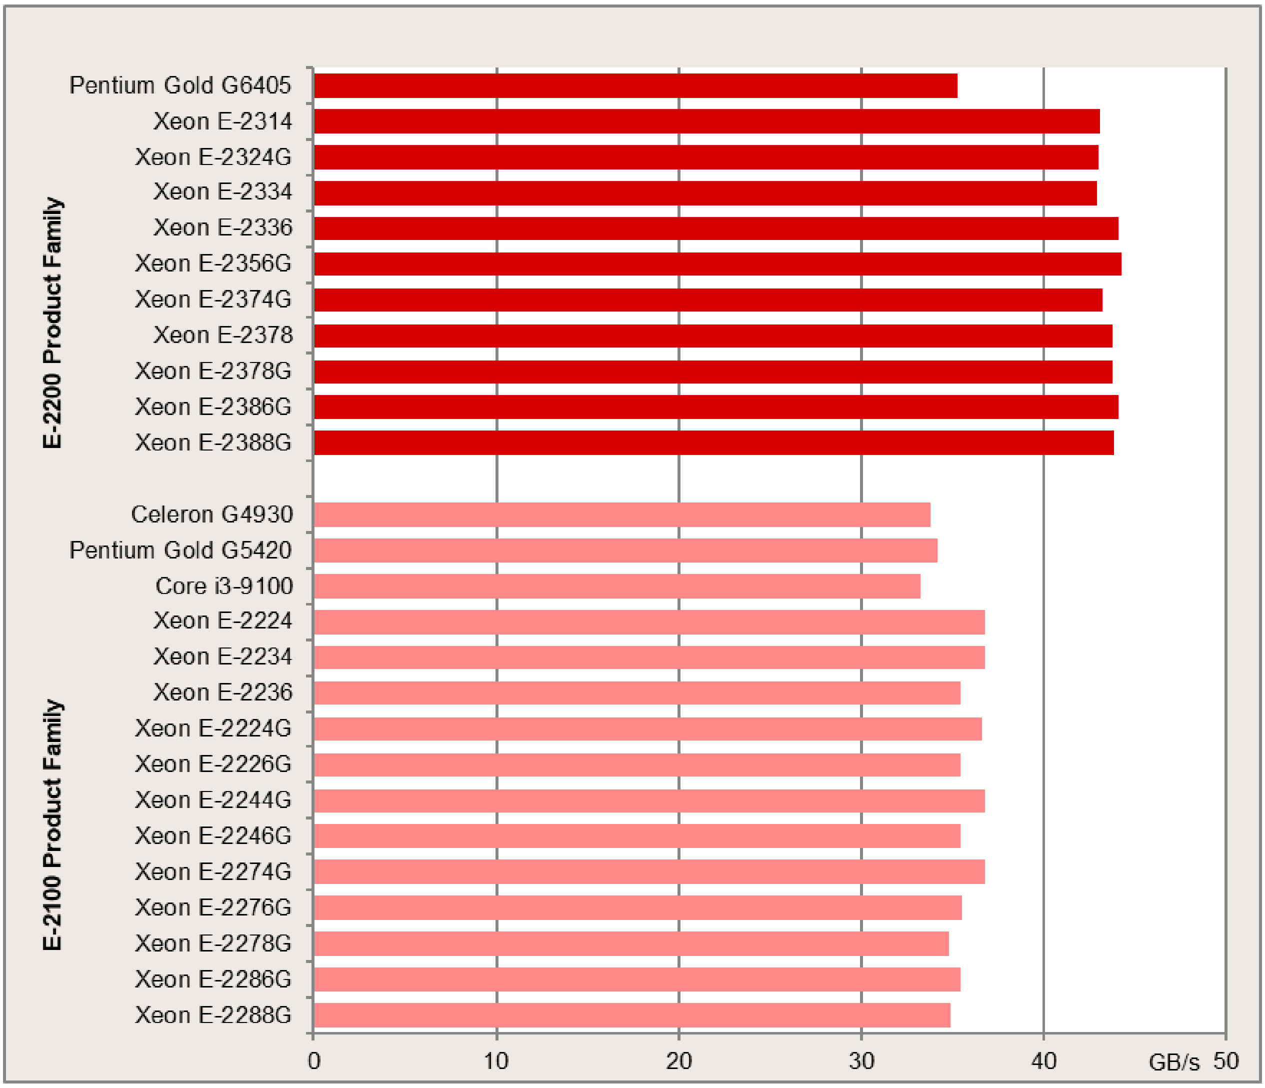

SPECpower_ssj2008

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.

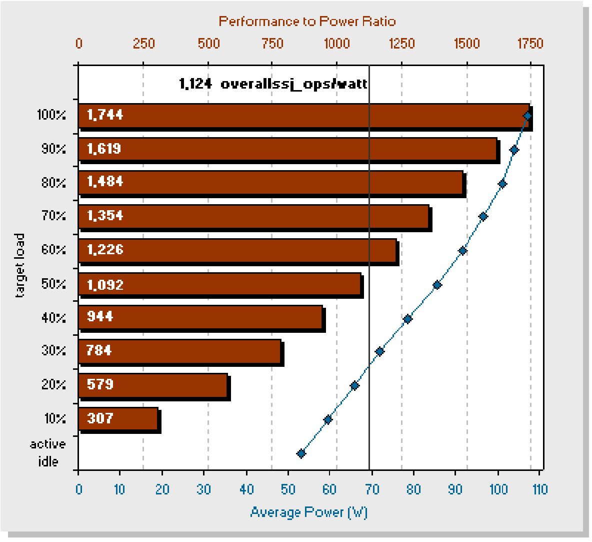

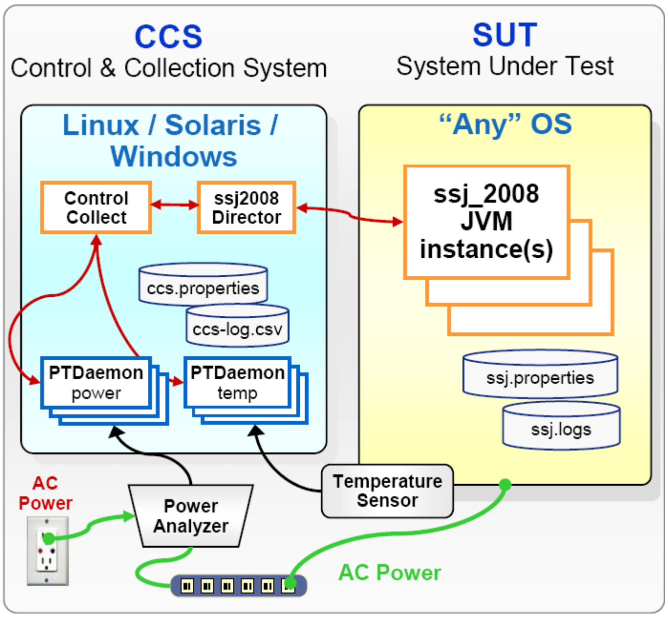

The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System 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)	
For Linux OS measurement	
Hardware	
• Model	PRIMERGY RX1330 M5
• Processor	1 x Intel Xeon E-2388G
• Memory	2 x 8 GB 1Rx8 PC4-3200AA-ED2-11
• Network interface	2 x Intel I210 Gigabit Network Connection (onboard)
• Disk subsystem	1 x SSD M.2 240GB, S26361-F5787-E240
• Power Supply Unit	1 x 300 W, S26113-E614-V70-1
Software	
• BIOS	R1.30.0
• BIOS settings	ASPM Support = Auto Hardware Prefetcher = Disabled Adjacent Cache Line Prefetch = Disabled Intel Virtualization Technology = Disabled DDR Speed Control = Auto DMI Gen3 ASPM = Auto DMI Link ASPM Control = Auto SATA Controller Port0/1/2/3/4/6 = Disabled Serial Port = Disabled LAN2 Controller = Disabled
• iRMC Firmware	1.01S
• Operating system	SUSE Linux Enterprise Server 15 SP3, 5.3.18-57-default
• Operating system settings	kernel parameter: pcie_aspm=force pcie_aspm.policy=powersave intel_pstate=passive Benchmark started via ssh. modprobe cpufreq_conservative cpupower frequency-set -g conservative echo 3000000 > /sys/devices/system/cpu/cpufreq/conservative/sampling_rate echo 92 > /sys/devices/system/cpu/cpufreq/conservative/up_threshold echo 1 > /sys/devices/system/cpu/cpufreq/conservative/freq_step echo 91 > /sys/devices/system/cpu/cpufreq/conservative/down_threshold echo always > /sys/kernel/mm/transparent_hugepage/enabled sysctl -w kernel.nmi_watchdog=0
• JVM	Oracle Java HotSpot 64-Bit Server VM 18.9 (build 11.0.9+7-LTS, mixed mode)
• JVM settings	-server -Xmn9500m -Xms11000m -Xmx11000m -XX:+UseHugeTLBFS -XX:+UseLargePages -XX:+UseTransparentHugePages -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchInstr=0 -XX:AllocatePrefetchLines=4 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:ParallelGCThreads=8 -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:+UseParallelOldGC -XX:FreqInlineSize=2500 -XX:MinJumpTableSize=18 -XX:UseAVX=0 -XX:+UseBiasedLocking -XX:-ThreadLocalHandshakes

For Windows OS measurement	
Hardware	
• Model	PRIMERGY RX1330 M5
• Processor	1 x Intel Xeon E-2388G
• Memory	2 x 8 GB 1Rx8 PC4-3200AA-ED2-11
• Network interface	2 x Intel I210 Gigabit Network Connection (onboard)
• Disk subsystem	1 x SSD M.2 240GB, S26361-F5787-E240
• Power Supply Unit	1 x 300 W, S26113-E614-V70-1
Software	
• BIOS	R1.30.0
• BIOS settings	ASPM Support = Auto Hardware Prefetcher = Disabled Adjacent Cache Line Prefetch = Disabled Intel Virtualization Technology = Disabled DMI Max Link Speed = Gen1 DDR Speed Control = Auto DMI Gen3 ASPM = Auto DMI Link ASPM Control = Auto SATA Controller Port0/1/2/3/4/6 = Disabled Serial Port = Disabled LAN2 Controller = Disabled
• iRMC Firmware	1.01S
• Operating system	Microsoft Windows Server 2019 Standard
• Operating system settings	Turn off hard disk after = 1 Minute Turn off display after = 1 Minute Minimum processor state = 0% Maximum processor state = 100% Using the local security settings console, "lock pages in memory" was enabled for the user running the benchmark. Benchmark was started via Windows Remote Desktop Connection.
• JVM	Oracle Java HotSpot 64-Bit Server VM 18.9 (build 11.0.9+7-LTS, mixed mode)
• JVM settings	-server -Xmn9500m -Xms11000m -Xmx11000m -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:ParallelGCThreads=2 - XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:LoopUnrollLimit=45 - XX:InitialTenuringThreshold=12 -XX:MaxTenuringThreshold=15 -XX:InlineSmallCode=3900 - XX:MaxInlineSize=270 -XX:FreqInlineSize=2500 -XX:+UseLargePages -XX:+UseParallelOldGC - XX:UseAVX=0 -XX:-UseAdaptiveSizePolicy -XX:-ThreadLocalHandshakes

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

Benchmark results (Linux)

The PRIMERGY RX1330 M5 in SUSE Linux Enterprise Server 15 SP3 achieved the following result:

SPECpower_ssj2008 = 10,767 overall ssj_ops/watt

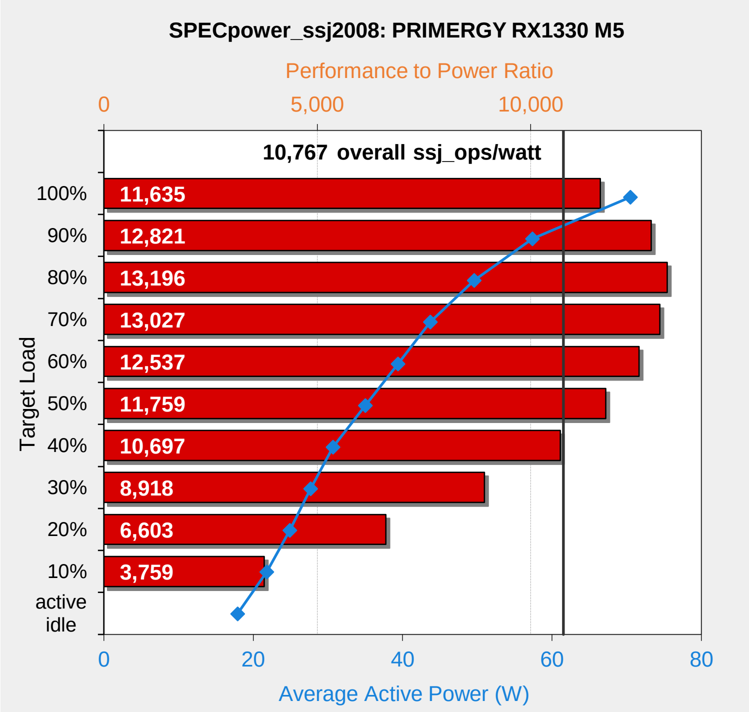

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 10,767 overall ssj_ops/watt for the PRIMERGY RX1330 M5. 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%	820,751	70.5	11,635
90%	735,856	57.4	12,821
80%	655,157	49.6	13,196
70%	569,227	43.7	13,027
60%	493,752	39.4	12,537
50%	411,493	35.0	11,759
40%	328,344	30.7	10,697
30%	247,043	27.7	8,918
20%	164,504	24.9	6,603
10%	82,121	21.8	3,759
Active Idle	0	17.9	0
			Σ ssj_ops / Σ power = 10,767

Benchmark results (Windows)

The PRIMERGY RX1330 M5 in Microsoft Windows Server 2019 Standard achieved the following result:

SPECpower_ssj2008 = 8,700 overall ssj_ops/watt

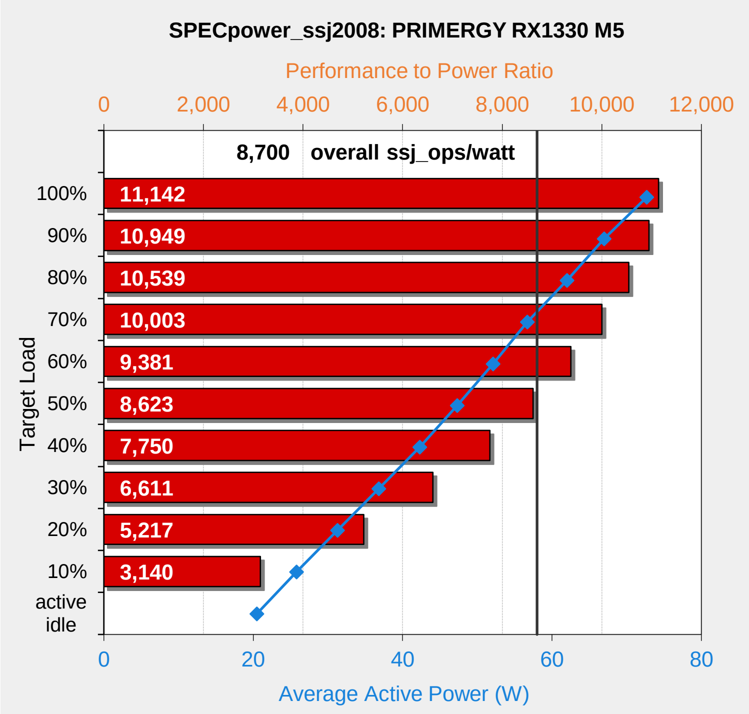

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 8,700 overall ssj_ops/watt for the PRIMERGY RX1330 M5. 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%	810,100	72.7	11,142
90%	733,433	67.0	10,949
80%	653,806	62.0	10,539
70%	567,629	56.7	10,003
60%	488,352	52.1	9,381
50%	407,466	47.3	8,623
40%	328,071	42.3	7,750
30%	243,614	36.8	6,611
20%	163,483	31.3	5,217
10%	81,006	25.8	3,140
Active Idle	0	20.5	0
			Σ ssj_ops / Σ power = 8,700

Comparison with the predecessor

The following diagram shows for each load level (on the x-axis) the throughput (on the left y-axis) and the power consumption (on the right y-axis) of the PRIMERGY RX1330 M5 compared to the predecessor PRIMERGY RX1330 M4, with the JVM versions that affect the SPECpower_ssj2008 benchmark.

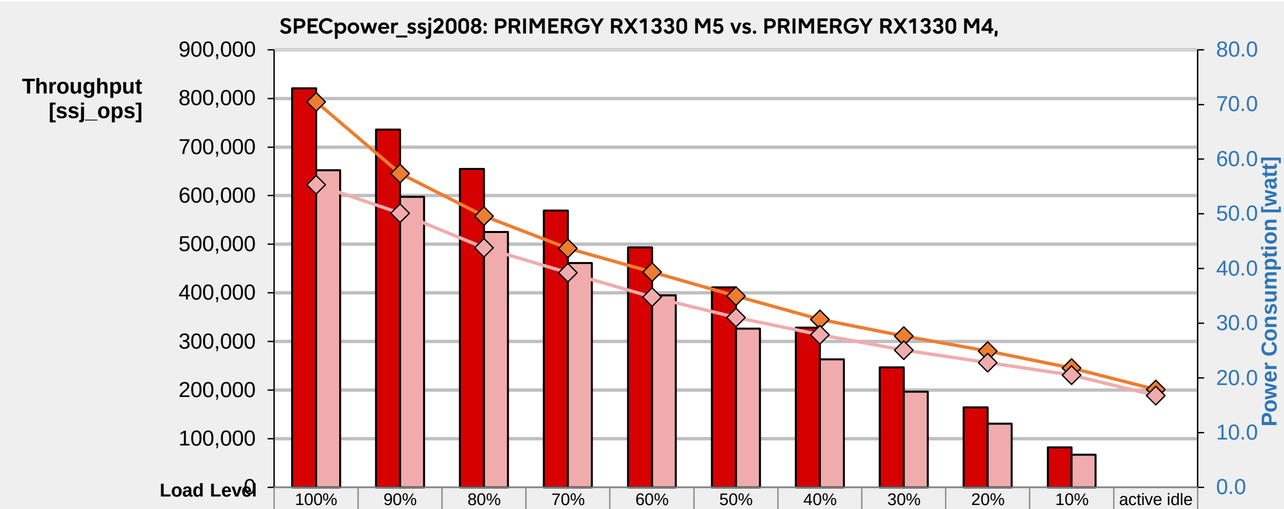

Load Level	100%	90%	80%	70%	60%	50%	40%	30%	20%	10%	active idle
Xeon E-2388G ssj_ops	820,751	735,856	655,157	569,227	493,752	411,493	328,344	247,043	164,504	82,121	0
Xeon E-2288G ssj_ops	652,394	597,828	525,290	461,597	395,067	326,690	263,437	196,841	131,230	66,841	0
Xeon E-2388G watt	70.5	57.4	49.6	43.7	39.4	35.0	30.7	27.7	24.9	21.8	17.9
Xeon E-2288G watt	55.4	50.1	43.8	39.3	34.8	31.1	27.9	25.1	22.8	20.5	16.8

The total average throughput of the PRIMERGY RX1330 M5 is 4,508,248 ssj_ops, an improvement of 24.6% over the 3,617,215 ssj_ops of the PRIMERGY RX1330 M4,.

On the other hand, the total average power consumption of the PRIMERGY RX1330 M5 is 419 W, which is 13.9% higher than the 368 W of the PRIMERGY RX1330 M4,. The reason for the increased power consumption is the increased power consumption when idle. It was 16.8W on the PRIMERGY RX1330 M4,, but it increased 1.07 times to 17.9W on the PRIMERGY RX1330 M5.

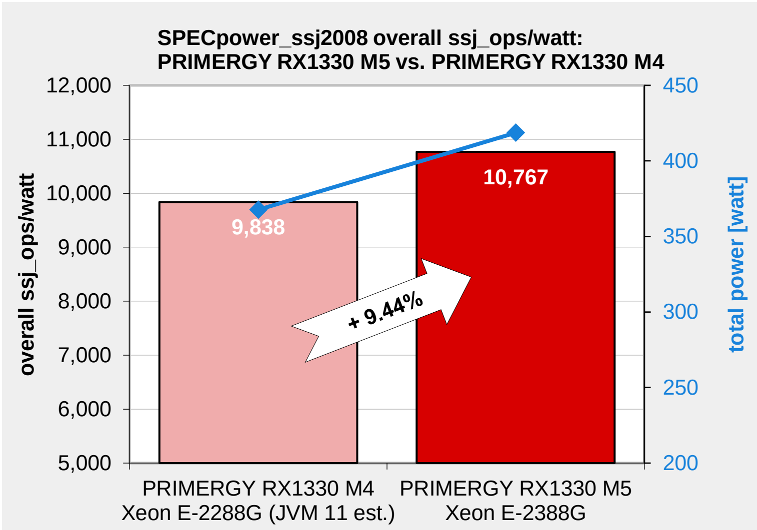

The overall energy efficiency of the PRIMERGY RX1330 M5 has improved by 9.44% due to a 13.9% increase in power consumption but a 24.6% improvement in performance.

Measurement results of SPECpower_ssj2008 (January 19, 2022)

8,700 SPECpower_ssj2008

On January 19, 2022, PRIMERGY RX1330 M5 with an Intel Xeon E-2388G processor achieved a performance value of 8,700 on the Windows Server 2019 Standard in the SPECpower_ssj2008 benchmark, in the Windows division of Intel Xeon E-2300 processor family category and won first place in SPECpower_ssj2008 performance.

For the latest results of the SPECpower_ssj2008, see https://www.spec.org/power_ssj2008/results/.

Disk I/O: Performance of storage media

Benchmark description

Performance measurements of disk subsystems for PRIMERGY servers 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 as follows.

- Random access / sequential access ratio
- Read / write access ratio
- Block size (kiB)
- Queue Depth (number of IO requests to issue at one time)

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 [kiB]	Application
		read	write		
Filecopy	Random	50%	50%	64	Copying files
Fileserver	Random	67%	33%	64	Fileserver
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 files

In order to model applications that access in parallel with a different load intensity the Queue Depth 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 measurement items are as follows.

- Throughput [MiB/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 following formula.

Data throughput [MiB/s]	= Transaction rate [IO/s] x Block size [MiB]
Transaction rate [IO/s]	= Data throughput [MiB/s] / Block size [MiB]

In this section, a power of 10 (1 TB = 10^{12} bytes) is used to indicate the capacity of the hard storage medium, and a power of 2 (1 MiB / s = 2^{20} bytes) is used to indicate the capacity of other media, file size, block size, and throughput.

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		
3.5 inch model		
Controller: PRAID EP540i		
Storage media	Category	Drive name
HDD	SAS HDD (SAS 12Gbps, 15k rpm) [512n]	ST300MP0006 ST600MP0006
	NL-SAS HDD (SAS 12Gbps, 7.2k rpm) [512e]	ST6000NM029A ST8000NM001A
	NL-SAS HDD (SAS 12Gbps, 7.2k rpm) [512n]	ST2000NM003A ST4000NM003A
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512e]	ST6000NM021A ST8000NM000A
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512n]	ST2000NM000A ST4000NM000A
SSD	SATA SSD (SATA 6Gbps, Mixed Use)	MTFDDAK240TDT MTFDDAK480TDT MTFDDAK960TDT MTFDDAK1T9TDT MTFDDAK3T8TDT
	SATA SSD (SATA 6Gbps, Read Intensive)	MTFDDAK240TDS MTFDDAK480TDS MTFDDAK960TDS MTFDDAK1T9TDS MTFDDAK3T8TDS MTFDDAK7T6TDS

Controller: Intel C256 Standard SATA AHCI controller		
Storage media	Category	Drive name
SSD	M.2 Flash module	MTFDDAV240TDS MTFDDAV480TDS

2.5 inch model

Controller: PRAID EP540i

Storage media	Category	Drive name
HDD	SAS HDD (SAS 12Gbps, 15k rpm) [512n]	ST300MP0006 ST600MP0006
	SAS HDD (SAS 12Gbps, 10k rpm) [512e]	AL15SEB18EQ AL15SEB24EQ
	SAS HDD (SAS 12Gbps, 10k rpm) [512n]	AL15SEB060N AL15SEB120N
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512e]	ST1000NX0313
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512n]	ST1000NX0423
SSD	SAS SSD (SAS 12Gbps, Write Intensive)	XS400ME70084 XS800ME70084 XS1600ME70084
	SAS SSD (SAS 12Gbps, Mixed Use)	XS800LE70084 XS1600LE70084 XS3200LE70084
	SAS SSD (SAS 12Gbps, Read Intensive)	XS960SE70084 XS1920SE70084 XS3840SE70084 XS7680SE70084
	SATA SSD (SATA 6Gbps, Mixed Use)	MTFDDAK240TDT MTFDDAK480TDT MTFDDAK960TDT MTFDDAK1T9TDT MTFDDAK3T8TDT
	SATA SSD (SATA 6Gbps, Read Intensive)	MTFDDAK240TDS MTFDDAK480TDS MTFDDAK960TDS MTFDDAK1T9TDS MTFDDAK3T8TDS MTFDDAK7T6TDS
	PCIe SSD (Mixed Use)	SSDPE2KE016T8 SSDPE2KE032T8 SSDPE2KE064T8
	PCIe SSD (Read intensive)	SSDPE2KX010T8 SSDPE2KX020T8 SSDPE2KX040T8

Controller: Intel C256 Standard SATA AHCI controller

Storage media	Category	Drive name
SSD	M.2 Flash module	MTFDDAV240TDS MTFDDAV480TDS

Software		
Operating system		Microsoft Windows Server 2016 Standard
Benchmark version		3.0
RAID type		Type RAID 0 logical drive consisting of 1 hard disk
Stripe size		HDD: 256KiB, SSD: 64 KiB
Measuring tool		Iometer 1.1.0
Measurement area	HDD,SSD (other than M.2, E1.S)	RAW file system is used. The first 32GiB of available LBA space is used for sequential access. The following 64GiB is used for random access.
	SSD (M.2, E1.S)	NTFS file system is used. The first 32GiB of available LBA space is used for sequential access. The following 64GiB is used for random access.
Total number of Iometer worker		1
Alignment of Iometer accesses		Aligned to access block size

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

Benchmark results

The results shown here are intended to help you select the appropriate storage media 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."

Controller

The measurements were made using controllers in the table below.

Storage media	Controller name	Cache	Supported interfaces		RAID levels
			host	drive	
SSD/HDD PCIe SSD 2.5”	PRAID EP540i	-	8x PCIe 3.0	SATA 6G SAS 12G 16x PCIe	0, 1, 1E, 10, 5, 50
M.2 Flash	C620 Standard SATA AHCI controller	-	4x DMI 3.0	SATA 6G	-

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 for PRIMERGY servers.

model	Storage media type	Interface	Form factor
3.5 inch model	HDD	SAS 12G	3.5 inch or 2.5 inch [1)]
		SATA 6G	3.5 inch
	SSD	SATA 6G	2.5 inch [1)] or M.2
2.5 inch model	HDD	SAS 12G	2.5 inch
		SATA 6G	2.5 inch
	SSD	SAS 12G	2.5 inch
		SATA 6G	2.5 inch or M.2
		PCIe 3.0	2.5 inch

1) *It is available with a 3.5 inch cage.*

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 system board is typically PCIe or, in the case of the integrated onboard controllers, an internal bus interface of the system board.

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 mode "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 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.

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 (number 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.

Storage media performance

3.5 inch model

HDDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
300	ST300MP0006	SAS 12G	790	696	666	304	304
600	ST600MP0006	SAS 12G	736	651	601	301	300
6,000	ST6000NM029A	SAS 12G	369	333	323	252	252
8,000	ST8000NM001A	SAS 12G	354	310	310	255	255
2,000	ST2000NM003A	SAS 12G	378	343	336	237	237
4,000	ST4000NM003A	SAS 12G	369	333	330	214	215
6,000	ST6000NM021A	SATA 6G	326	293	302	253	253
8,000	ST8000NM000A	SATA 6G	325	290	301	249	248
2,000	ST2000NM000A	SATA 6G	331	304	313	230	207
4,000	ST4000NM000A	SATA 6G	313	290	297	211	210

SSDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
240	MTFDDAK240TDT	SATA 6G	46,406	5,989	6,121	508	370
480	MTFDDAK480TDT	SATA 6G	49,138	6,383	6,600	508	437
960	MTFDDAK960TDT	SATA 6G	50,488	6,970	7,136	508	486
1,920	MTFDDAK1T9TDT	SATA 6G	50,669	7,183	7,336	508	487
3,840	MTFDDAK3T8TDT	SATA 6G	49,490	7,115	7,208	493	474
240	MTFDDAK240TDS	SATA 6G	42,594	5,435	5,510	508	301
480	MTFDDAK480TDS	SATA 6G	47,577	6,109	6,310	508	401
960	MTFDDAK960TDS	SATA 6G	50,134	6,633	6,852	506	480
1,920	MTFDDAK1T9TDS	SATA 6G	50,638	7,078	7,286	508	488
3,840	MTFDDAK3T8TDS	SATA 6G	49,542	7,097	7,196	495	477
7,680	MTFDDAK7T6TDS	SATA 6G	47,200	7,134	7,563	508	487
240	MTFDDAV240TDS	SATA 6G	32,138	5,494	5,506	506	298
480	MTFDDAV480TDS	SATA 6G	39,970	6,373	6,541	508	395

2.5 inch model

HDDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
300	ST300MP0006	SAS 12G	790	696	666	304	304
600	ST600MP0006	SAS 12G	736	651	601	301	300
1,800	AL15SEB18EQ	SAS 12G	767	631	624	255	249
2,400	AL15SEB24EQ	SAS 12G	754	620	617	264	260
600	AL15SEB060N	SAS 12G	698	586	600	232	232
1,200	AL15SEB120N	SAS 12G	732	604	615	230	226
1,000	ST1000NX0313	SATA 6G	393	340	337	134	134
1,000	ST1000NX0423	SATA 6G	399	344	346	134	134

SSDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
400	XS400ME70084	SAS 12G	122,956	22,969	19,438	1,052	872
800	XS800ME70084	SAS 12G	123,848	23,784	19,435	1,052	874
1,600	XS1600ME70084	SAS 12G	123,277	23,725	19,270	1,051	884
800	XS800LE70084	SAS 12G	121,914	23,707	19,257	1,052	871
1,600	XS1600LE70084	SAS 12G	122,949	23,771	19,455	1,052	874
3,200	XS3200LE70084	SAS 12G	123,090	22,816	19,418	1,051	872
960	XS960SE70084	SAS 12G	123,014	23,678	19,424	1,052	870
1,920	XS1920SE70084	SAS 12G	123,093	23,760	19,423	1,052	874
3,840	XS3840SE70084	SAS 12G	122,810	22,949	19,406	1,051	871
7,680	XS7680SE70084	SAS 12G	123,461	22,899	19,516	1,051	880
240	MTFDDAK240TDT	SATA 6G	46,406	5,989	6,121	508	370
480	MTFDDAK480TDT	SATA 6G	49,138	6,383	6,600	508	437
960	MTFDDAK960TDT	SATA 6G	50,488	6,970	7,136	508	486
1,920	MTFDDAK1T9TDT	SATA 6G	50,669	7,183	7,336	508	487
3,840	MTFDDAK3T8TDT	SATA 6G	49,490	7,115	7,208	493	474
240	MTFDDAK240TDS	SATA 6G	42,594	5,435	5,510	508	301
480	MTFDDAK480TDS	SATA 6G	47,577	6,109	6,310	508	401
960	MTFDDAK960TDS	SATA 6G	50,134	6,633	6,852	506	480
1,920	MTFDDAK1T9TDS	SATA 6G	50,638	7,078	7,286	508	488
3,840	MTFDDAK3T8TDS	SATA 6G	49,542	7,097	7,196	495	477
7,680	MTFDDAK7T6TDS	SATA 6G	47,200	7,134	7,563	508	487
1,600	SSDPE2KE016T8	PCIe3 x4	276,785	45,739	40,923	3,214	1,972
3,200	SSDPE2KE032T8	PCIe3 x4	306,446	53,059	50,093	3,220	2,461
6,400	SSDPE2KE064T8	PCIe3 x4	297,505	56,338	56,632	3,219	2,499
1,000	SSDPE2KX01	PCIe3 x4	153,263	25,891	21,942	2,799	1,109
2,000	SSDPE2KX02	PCIe3 x4	237,530	38,336	34,740	3,181	1,979
4,000	SSDPE2KX04	PCIe3 x4	242,546	39,242	38,151	2,905	2,417
240	MTFDDAV240TDS	SATA 6G	32,138	5,494	5,506	506	298
480	MTFDDAV480TDS	SATA 6G	39,970	6,373	6,541	508	395

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 system.

In contrast to benchmarks such as SPEC CPU 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. OLTP-2 is based on the well-known database benchmark TPC-E. OLTP-2 was designed in such a way that a wide range of configurations can be measured to present the scaling of a system with regard to the CPU and memory configuration.

Even if the two benchmarks OLTP-2 and TPC-E simulate similar application scenarios using the same load profiles, 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.

Further information can be found in the document Benchmark Overview OLTP-2.

Benchmark environment

The typical measurement set-up is illustrated below:

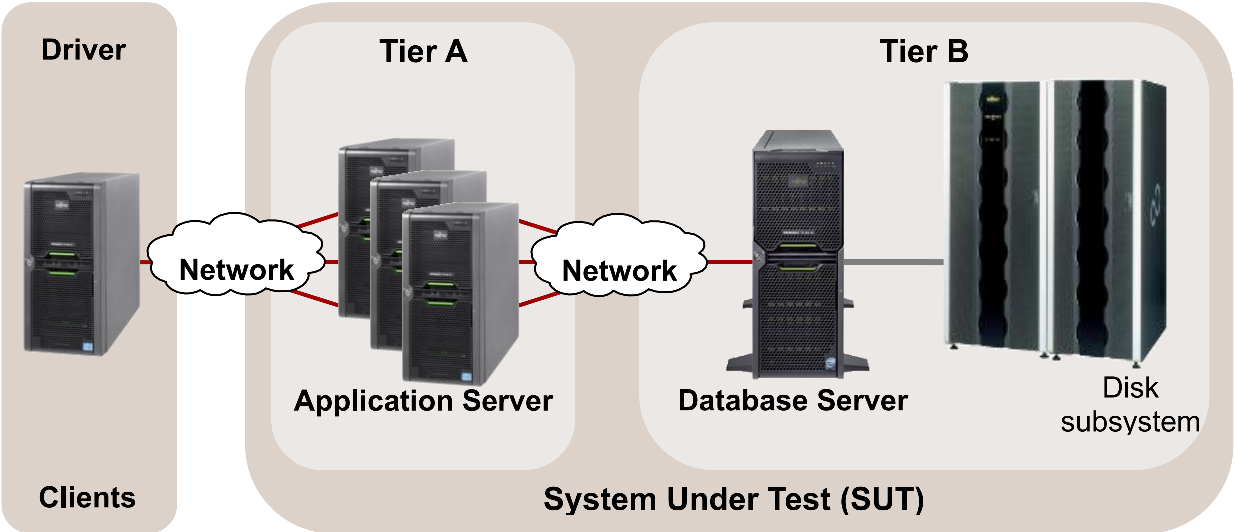

All OLTP-2 results were measured with the configuration of the next following pages of PRIMERGY RX1330 M5.

Database Server (Tier B)	
Hardware	
• Model	PRIMERGY RX1330 M5
• Processor	Intel Xeon E-2300 processor family Intel Pentium Gold G6405
• Memory	1 processors:4 x 32 GB (1x32 GB) 2Rx8 DDR4-3200 ECC
• Network interface	2 x Onboard LAN 1 Gbps
• Disk subsystem	PRIMERGY RX1330 M5: Operating system and database applications, RAID 0 (OS) Sequential access, optimized to reduce response time, RAID 1 (LOG) Write, optimize for response time priority, RAID 0 (temp) Random access, optimize throughput, RAID 0 (data)
Software	
• BIOS	Version R1.13.0
• Operating system	Microsoft Windows Server 2019 Standard + KB4577069
• Database	Microsoft SQL Server 2017 Enterprise + KB4341265

Application Server (Tier A)	
Hardware	
• Model	1 x PRIMERGY RX2530 M2
• Processor	2 x Intel Xeon 2699 v4
• Memory	128 GB, 2400 MHz Registered ECC DDR4
• Network interface	1 x Dual port LAN 10 Gbps (at 1Gbps) 1 x Dual port onboard LAN 1 Gbps
• Disk subsystem	2 x 300 GB 10k rpm SAS drive
Software	
• Operating system	Microsoft Windows Server 2012 R2 Standard

Client	
Hardware	
• Model	1 x PRIMERGY RX2530 M2
• Processor	2 x Intel Xeon E5-2667 v4
• Memory	128 GB, 2400 MHz Registered ECC DDR4
• Network interface	1 x Quad port onboard LAN 1 Gbps
• Disk subsystem	1 x 300 GB 10k rpm SAS drive
Software	
• Operating system	Microsoft Windows Server 2012 R2 Standard
• Benchmark	OLTP-2 Software EGen version 1.14.0

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

Benchmark results

Database performance greatly depends on the configuration options with CPU, memory and on the connectivity of an adequate disk subsystem for the database. In the following scaling considerations for the processors we assume that both the memory and the disk subsystem has been adequately chosen and is not a bottleneck.

A guideline in the database environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses. This why a configuration with a total memory of 128 GB was considered for the measurements. The memory configurations have memory access of 2933 MHz at the measurements.

The result with "est." are the estimated values.

Processor	Cores	Threads	Score
Xeon E-2388G	8	16	1,537
Xeon E-2386G	6	12	1,261
Xeon E-2378G	8	16	1,438
Xeon E-2378	8	16	1,248
Xeon E-2374G	4	8	888
Xeon E-2356G	6	12	1,193
Xeon E-2336	6	12	1,104
Xeon E-2334	4	8	819
Xeon E-2324G	4	4	623
Xeon E-2314	4	4	527
Pentium Gold G6405	2	4	378

The following graph shows the OLTP-2 transaction rates obtained with Intel Xeon E-2300 processor family and Intel Pentium Gold G6405.

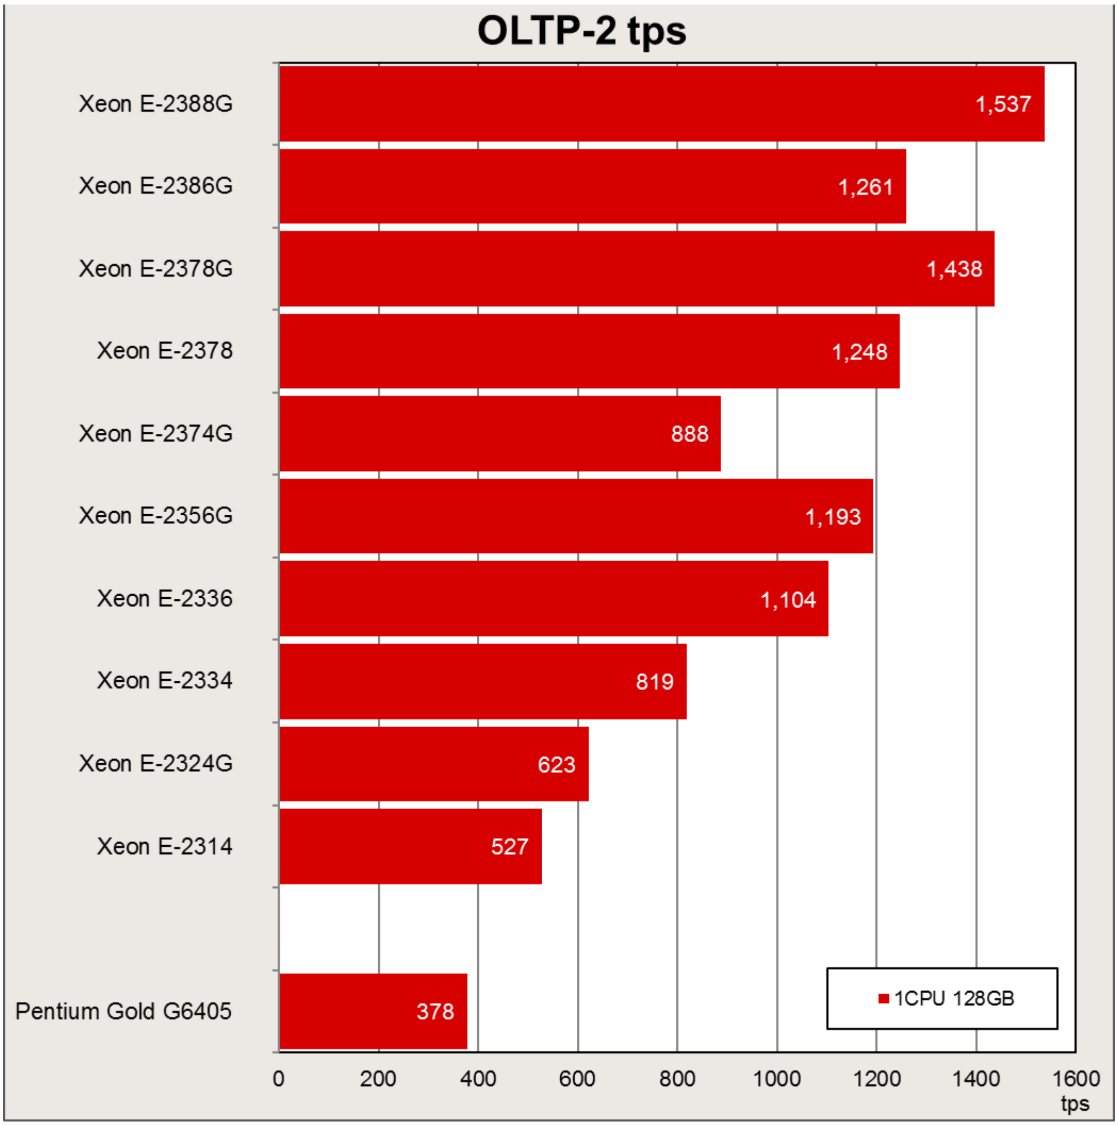

In general, the relatively large performance differences between the processors can be explained by their features. The values scale on the basis of the number of cores, the size of the L3 cache and the CPU clock frequency and as a result of the features of Hyper-Threading and turbo mode, which are available in most processor types. Furthermore, the data transfer rate between processors ("UPI Speed") also determines the performance.

Within a group of processors with the same number of cores, scaling can be seen via the CPU clock frequency.

The highest value for OLTP-2 on the current PRIMERGY model is about 17% higher than the highest value on the previous model.

It cannot be said exactly because the measurement conditions such as OS are different, but approximately 70% of the performance improvement is due to the increase in system memory.

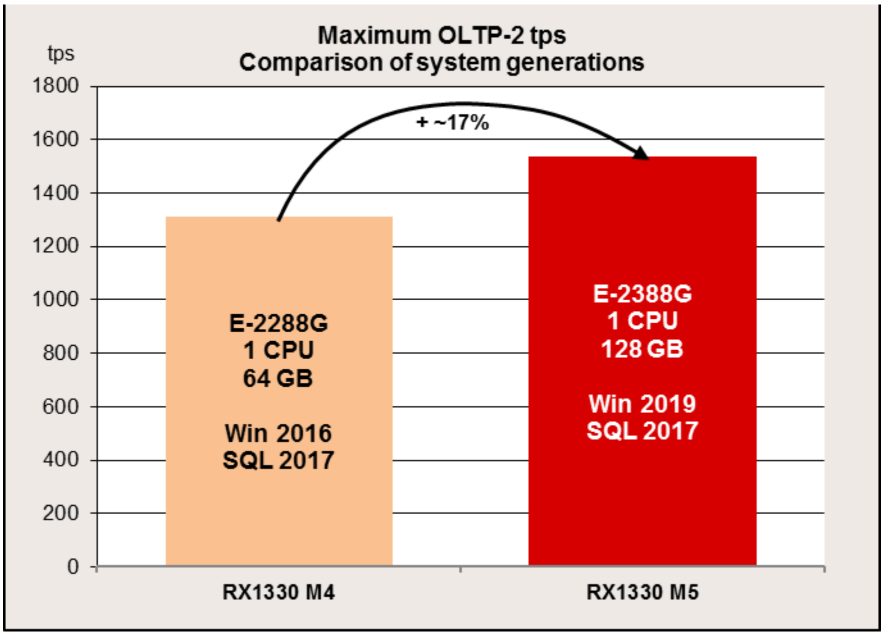

VMmark V3

Benchmark description

VMmark V3 is a benchmark developed by VMware to compare server configurations with hypervisor solutions from VMware regarding their suitability for server consolidation. In addition to the software for load generation, the benchmark consists of a defined load profile and binding regulations. The benchmark results can be submitted to VMware and are published on their Internet site after a successful review process. After the discontinuation of the proven benchmark “VMmark V2” in September 2017, it has been succeeded by “VMmark V3”. VMmark V2 required a cluster of at least two servers and covers data center functions, like Cloning and Deployment of virtual machines (VMs), Load Balancing, as well as the moving of VMs with vMotion and also Storage vMotion. VMmark V3 covers the moving of VMs with XvMotion in addition to VMmark V2. Also, changes application architecture to more scalable workloads.

In addition to the “Performance Only” result, alternatively measure the electrical power consumption and publish it as a “Performance with Server Power” result (power consumption of server systems only) and/or “Performance with Server and Storage Power” result (power consumption of server systems and all storage components).

Application scenario	Load tool	# VMs
Scalable web system	Weathervane	14
E-commerce system	DVD Store 3 client	4
Standby system		1

VMmark V3 is not a new benchmark in the actual sense. It is in fact a framework that consolidates already established benchmarks, as workloads in order to simulate the load of a virtualized consolidated server environment. Two proven benchmarks, which cover the application scenarios Scalable web system and E commerce system were integrated in VMmark V3.

Each of the three application scenarios is assigned to a total of 18 dedicated virtual machines. Then add to these an 19th VM called the “standby server”. These 19 VMs form a “tile”. Because of the performance capability of the underlying server hardware, it is usually necessary to have started several identical tiles in parallel as part of a measurement in order to achieve a maximum overall performance.

In VMmark V3 there is an an infrastructure component, which is present once for every two hosts. It measures the efficiency levels of data center consolidation through VM Cloning and Deployment, vMotion, XvMotion and Storage vMotion. The Load Balancing capacity of the data center is also used (DRS, Distributed Resource Scheduler).

The result of VMmark V3 for test type “Performance Only” is a number, known as a “score”, which provides information about the performance of the measured virtualization solution. The score is the maximum sum of the benefits of server aggregation and is used as a comparison criterion for different hardware platforms.

This score is determined from the individual results of the VMs and an infrastructure components result. Each of the five VMmark V3 application or front-end VMs provides a specific benchmark result in the form of application-specific transaction rates for each VM. In order to derive a normalized score, the individual benchmark result for each tile is put in relation to the respective results of a reference system. The resulting dimensionless performance values are then averaged geometrically and finally added up for all VMs. This value is included in the overall score with a weighting of 80%. The infrastructure workload is only present in the benchmark once for every two hosts; it determines 20% of the result. The number of transactions per hour and the average duration in seconds respectively are determined for the score of the infrastructure components workload.

In addition to the actual score, the number of VMmark V3 tiles is always specified with each VMmark V3 score. The result is thus as follows: “Score@Number of Tiles”, for example “8.11@8 tiles”.

In the case of the two test types “Performance with Server Power” and “Performance with Server and Storage Power”, a so-called “Server PPKW Score” and “Server and Storage PPKW Score” are determined. These are the performance scores divided by the average power consumption in kilowatts (PPKW = performance per kilowatt (kW)).

The results of the three test types should not be compared with each other.

A detailed description of VMmark V3 is available in the document Benchmark Overview VMmark V3.

Benchmark environment

The typical measurement set-up is illustrated below:

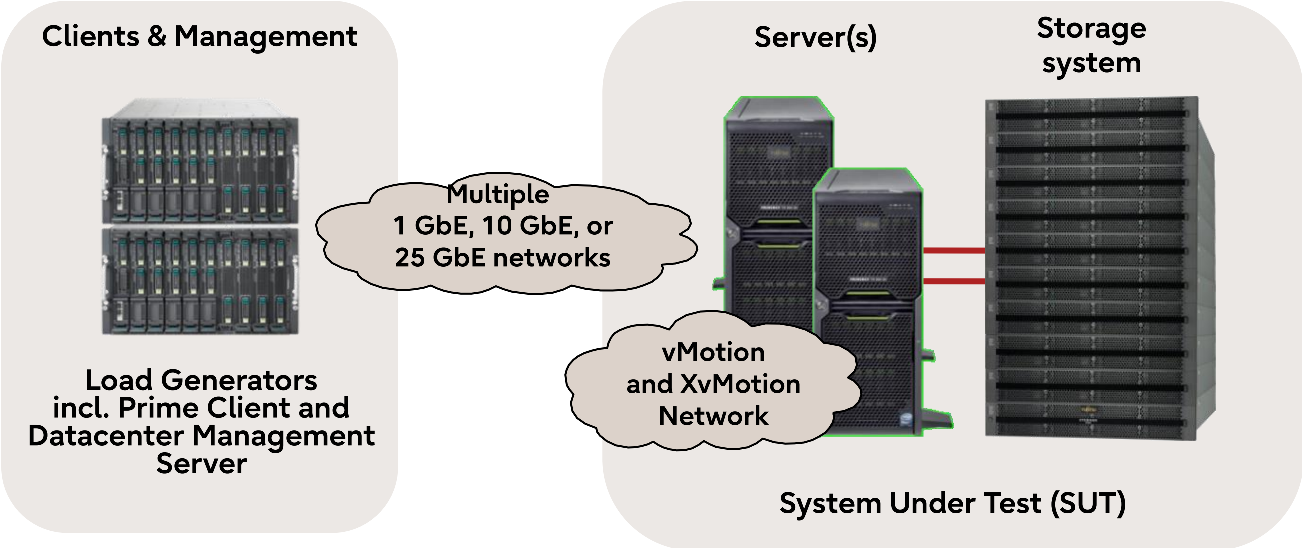

System Under Test (SUT)	
Hardware	
• Number of servers	2
• Model	PRIMERGY RX1330 M5
• Processor	Intel Xeon E-2388G
• Memory	128 GB: 4 × 32GB (1x32GB) 2Rx8 DDR4-3200 U ECC
• Network interface	1 × Mellanox MCX4121A-ACAT dual port 25Gb SFP28 PCIe Adapter 1 × Intel I210 1Gb dual port onboard
• Disk subsystem	1 × Emulex LPe35002 dual port 32Gb PCIe Adapter 1 × PRIMERGY RX2540 M4 configured as Fibre Channel target: 2 × Micron MTFDDAK480TDC SATA SSD (480 GB) 1 × Intel P46800 PCIe SSD (4 TB) 2 × Intel P4800X PCIe SSD (750 GB)
Software	
• BIOS	R1.33.0
• BIOS settings	See details
• Operating system	VMware ESXi 7.0 U3 GA, Build 19193900
• Operating system settings	ESX settings: see details

Detail	
See disclosure	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2022-05-03-Fujitsu-PRIMERGY-RX1330M5.pdf https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2022-05-03-Fujitsu-PRIMERGY-RX1330M5-serverPPKW.pdf https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2022-05-03-Fujitsu-PRIMERGY-RX1330M5-serverstoragePPKW.pdf

Datacenter Management Server (DMS)	
Hardware	
• Model	1 × PRIMERGY RX2530 M2
• Processor	2 × Intel Xeon E5-2698 v4
• Memory	64 GB
• Network interface	1 × Emulex One Connect Oce14000 1 GbE Dual Port Adapter
Software	
• Operating system	Hypervisor: VMware ESXi 6.7 EP 02a Build 9214924
Datacenter Management Server (DMS) VM	
Hardware	
• Processor	4 × logical CPU
• Memory	19 GB
• Network interface	1 × 1 Gbit/s LAN
Software	
• Operating system	VMware vCenter Server Appliance 7.0 U2 Build 17694817

Load generator	
Hardware	
• Model	2 × PRIMERGY RX2530 M2
• Processor	2 × Intel Xeon E5-2699 v4
• Memory	256 GB
• Network interface	1 × Emulex One Connect Oce14000 1GbE Dual Port Adapter 1 × Emulex One Connect Oce14000 10GbE Dual Port Adapter
Software	
• Operating system	VMware ESXi 6.7 EP 08 Build 13473784

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

Benchmark results

“Performance Only” measurement result (May 3, 2022)

On May 3, 2022 Fujitsu achieved with a PRIMERGY RX1330 M5 with an Intel Xeon E-2388G processor and VMware ESXi 7.0 U3 a VMmark V3 score of “1.46@1 tile” in a system configuration with a total of 16 processor cores and when using two identical servers in the “System under Test” (SUT). With this result the PRIMERGY RX1330 M5 is in the official VMmark V3 “Performance Only” ranking the most powerful Intel Xeon E processor server in a “Matched Pair” configuration consisting of two identical hosts (valid as of benchmark results publication date).

All comparisons for the competitor products reflect the status of May 3, 2022.

The current VMmark V3 “Performance Only” results as well as the detailed results and configuration data are available at https://www.vmware.com/products/vmmark/results3x.html .

All VMs, their application data, the host operating system, and any additional data needed were stored on a powerful Fibre Channel disk subsystem. This disk subsystem used fast PCIe SSDs such as Intel Optane to improve the response time of storage media. The host-side network connectivity to the load generators and infrastructure load connectivity between hosts were implemented using 25GbE LAN ports.

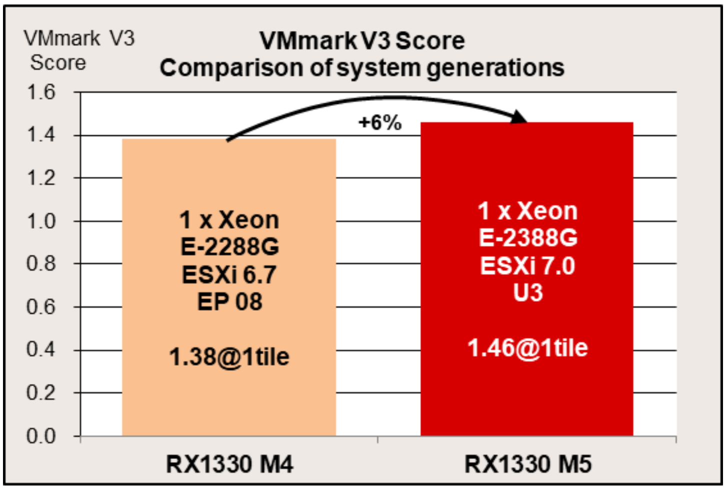

The graph on the left compares the VMmark V3 scores of the PRIMEGY RX1330 M5 and the previous generation PRIMEGY RX1330 M4. The PRIMERGY RX1330 M5 achieved a 6% improvement in score compared to the previous generation PRIMERGY RX1330 M4. This is due to the improved performance of the Intel Xeon E processor and the effective use of the capabilities of the VMware ESXi hypervisor.

“Performance with Server Power” measurement result (May 3, 2022)

On May 3, 2022 Fujitsu achieved with a PRIMERGY RX1330 M5 with an Intel Xeon E-2388G processor and VMware ESXi 7.0 U3 a VMmark V3 “Server PPKW Score” of “6.3858@1 tile” in a system configuration with a total of 16 processor cores and when using two identical servers in the “System under Test” (SUT). With this result the PRIMERGY RX1330 M5 is in the official VMmark V3 “Performance with Server Power” ranking the most energy-efficient virtualization server worldwide (valid as of benchmark results publication date).

The current VMmark V3 “Performance with Server Power” results as well as the detailed results and configuration data are available at https://www.vmware.com/products/vmmark/results3x.html .

Literature

PRIMERGY Servers

https://www.fujitsu.com/global/products/computing/servers/primergy/

PRIMERGY RX1330 M5

This Whitepaper

- https://docs.ts.fujitsu.com/dl.aspx?id=3af7ee68-4e6c-4291-b5fa-eb5acb0eb0b2
- https://docs.ts.fujitsu.com/dl.aspx?id=c37849fa-4b65-4b1b-a4cd-8b992fc6abb4

Data sheet

https://docs.ts.fujitsu.com/dl.aspx?id=bb81ac54-e78b-40a4-8e2d-bb8f0078e487

PRIMERGY Performance

https://www.fujitsu.com/global/products/computing/servers/primergy/benchmarks/

SPEC CPU2017

https://www.spec.org/osg/cpu2017

Benchmark Overview SPECcpu2017

https://docs.ts.fujitsu.com/dl.aspx?id=20f1f4e2-5b3c-454a-947f-c169fca51eb1

STREAM

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

LINPACK

The LINPACK Benchmark: Past, Present, and Future

https://www.netlib.org/utk/people/JackDongarra/PAPERS/hplpaper.pdf

TOP500

https://www.top500.org/

HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers

https://www.netlib.org/benchmark/hpl/

Intel Math Kernel Library – LINPACK Download

https://www.intel.com/content/www/us/en/developer/articles/technical/onemkl-benchmarks-suite.html

SPECpower_ssj2008

https://www.spec.org/power_ssj2008

Benchmark Overview SPECpower_ssj2008

https://docs.ts.fujitsu.com/dl.aspx?id=166f8497-4bf0-4190-91a1-884b90850ee0

OLTP-2

Benchmark Overview OLTP-2

https://docs.ts.fujitsu.com/dl.aspx?id=e6f7a4c9-aff6-4598-b199-836053214d3f

VMmark V3

https://www.vmware.com/products/vmmark.html

Benchmark Overview VMmark V3

https://docs.ts.fujitsu.com/dl.aspx?id=e6f9973c-90d6-47c6-b317-e388a978bfb7

Document change history

Version	Date	Description
1.3	2023-10-03	Update: • New Visual Identity format
1.2	2022-07-05	Update: • Minor collections
1.1	2022-05-23	New: • VMmark • Measured with Intel Xeon E-2388G Update: • Minor collections
1.0	2022-02-18	New: • Technical data • SPECcpu2017, OLTP-2, STREAM Measured and calculated with Intel Xeon E-2300 processor family / Intel Pentium Gold G6405 • SPECpower_ssj2008 Measured with Intel Xeon E-2388G • Disk I/O Measured with 2.5" and 3.5" storage media

Contact

Fujitsu

Web site: https://www.fujitsu.com

PRIMERGY Performance and Benchmarks

mailto:fj-benchmark@dl.jp.fujitsu.com