Fujitsu Server PRIMERGY
Performance Report
PRIMERGY TX1330 M4

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMERGY TX1330 M4.

Explaines PRIMERGY TX1330 M4 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
1.5 2023-10-03

Contents

Technical data

PRIMERGY TX1330 M4

Decimal prefixes according to the SI standard are used for storage capacity 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 provided for any further exceptions where applicable.

Model	PRIMERGY TX1330 M4
Model versions	PY TX1330 M4/ Floorstand /Standard PSU PY TX1330 M4/ Floorstand /Red. PSU PY TX1330 M4/ Rack/Red. PSU
Form factor	Tower server
Chipset	Intel 246
Number of sockets	1
Number of processors orderable	1
Processor type	Intel Pentium Gold G5400 Intel Pentium Gold G5420 Intel Core i3-8100 Intel Core i3-9100 Intel Xeon Processor E-2100 Product Family Intel Xeon Processor E-2200 Product Family
Number of memory slots	4
Maximum memory configuration	128 GB
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 x 8 1 × PCI-Express 3.0 x 4 1 × PCI-Express 3.0 x 1
Max. number of internal hard disks	PY TX1320 M4 LFF : 3.5" × 2 PY TX1320 M4 SFF : 2.5" × 8

Processors (since system release)							
Processor	Cores	Threads	Cache [MB]	Rated Frequency [Ghz]	Max. Turbo Frequency [Ghz]	Max. Memory Frequency [MHz]	TDP [Watt]
November 2018 released							
Pentium Gold G5400	2	4	4	3.7		2,400	58
Core i3-8100	4	4	6	3.6		2,400	65
Xeon E-2124	4	4	8	3.3	4.3	2,666	71
Xeon E-2124G	4	4	8	3.4	4.5	2,666	71
Xeon E-2126G	6	6	12	3.3	4.5	2,666	80
Xeon E-2134	4	8	8	3.5	4.5	2,666	71
Xeon E-2136	6	12	12	3.3	4.5	2,666	80
Xeon E-2144G	4	8	8	3.6	4.5	2,666	71
Xeon E-2146G	6	12	12	3.5	4.5	2,666	80
Xeon E-2174G	4	8	8	3.8	4.7	2,666	71
Xeon E-2176G	6	12	12	3.7	4.7	2,666	80
Xeon E-2186G	6	12	12	3.8	4.7	2,666	95
November 2019 released							
Pentium Gold G5420	2	4	4	3.8		2,400	54
Core i3-9100	4	4	6	3.6	4.2	2,400	65
Xeon E-2224	4	4	8	3.4	4.6	2,666	71
Xeon E-2224G	4	4	8	3.5	4.7	2,666	71
Xeon E-2226G	6	6	12	3.4	4.7	2,666	80
Xeon E-2234	4	8	8	3.6	4.8	2,666	71
Xeon E-2236	6	12	12	3.4	4.8	2,666	80
Xeon E-2244G	4	8	8	3.8	4.8	2,666	71
Xeon E-2246G	6	12	12	3.6	4.8	2,666	80
Xeon E-2274G	4	8	8	4.0	4.9	2,666	83
Xeon E-2276G	6	12	12	3.8	4.9	2,666	80
Xeon E-2278G	8	16	16	3.4	5.0	2,666	80
Xeon E-2286G	6	12	12	4.0	4.9	2,666	95
Xeon E-2288G	8	16	16	3.7	5.0	2,666	95

All the processors of the Intel Xeon Processor E-2100 Product Family and the Intel Xeon Processor E-2200 Product Family, and Core i3-9100 processor, that can be ordered with the PRIMERGY TX1330 M4 support Intel® Turbo Boost Technology 2.0. This technology allows you to operate the processor with higher frequencies than the nominal frequency. "Max. Turbo Frequency" listed in the processor table is the theoretical maximum frequency 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 can 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 a BIOS option. Fujitsu generally recommends leaving the "Turbo Mode" option set at the standard setting of "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
4 GB (1 x 4 GB) 1Rx8 DDR4-2666 U ECC	4	1	8	2,666				✓
8 GB (1 x 8 GB) 1Rx8 DDR4-2666 U ECC	8	1	8	2,666				✓
16 GB (1 x 16 GB) 2Rx8 DDR4-2666 U ECC	16	2	8	2,666				✓
32 GB (1 x 32 GB) 2Rx8 DDR4-2666 U ECC	32	2	8	2,666				✓

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 M4.

SPEC CPU2017

Benchmark description

SPECcpu2017 is a benchmark which measures the system efficiency with integer and floating-point operations. It consists of an integer test suite (SPECrate 2017 Integer or SPECspeed 2017 Integer) containing 10 applications and a floating-point test suite (SPECrate 2017 Floating Point or SPECspeed 2017 Floating Point) containing 14 applications. Both test suites are extremely computation-intensive and concentrate on the CPU and the memory. Other components, such as disk I/O and network, are not measured by this benchmark.

SPECcpu2017 is not tied to a specific operating system. The benchmark is available as source code and is compiled before the actual measurement. The compiler version used and its optimization settings also affect the measurement results.

SPECcpu2017 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. The performance per power can also be measured during performance measurements, by measuring the system power with a power meter.

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 and the peak values are optional.

Benchmark	Number of single benchmarks	Arithmetic	Compiler optimization	Measurement result	
SPECspeed2017_int_peak	10	Integer	Aggressive (peak)	Speed	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	10		Conservative (base)		Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	10		Aggressive (peak)	Through put	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	10		Conservative (base)		Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	10	Floating point	Aggressive (peak)	Speed	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	10		Conservative (base)		Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	13		Aggressive (peak)	Through put	Performance
SPECspeed2017_int_energy_peak					Power efficiency
SPECspeed2017_int_peak	13		Conservative (base)		Performance
SPECspeed2017_int_energy_peak					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 the test system is compared to a reference system. A value of "1" has been defined to be the SPECspeed2017_int_base, SPECrate2017_int_base, SPECspeed2017_fp_base, and SPECrate2017_fp_base results of the reference system. For example, 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 SPECcpu2017 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 M4
• Processor	Intel Pentium Gold G5400 Intel Pentium Gold G5420 Intel Core i3-8100 Intel Core i3-9100 Intel Xeon Processor E-2100 Product Family Intel Xeon Processor E-2200 Product Family
• Memory	4 × 16 GB (1 x 16 GB) 2Rx8 DDR4-2666 U ECC
Software	
• BIOS settings	Pentium Gold G5400, Core i3-8100: SPECspeed2017_int: • DCU Streamer Prefetcher = Disabled • DDR PowerDown and idle counter = PCODE • CState Pre-Wake = Disabled • Package C-State Un-demotion = Enabled • REFRESH_2X_MODE = 1- Enabled for WARM or HOT SPECrate2017_int: Fan Control = Full SPECrate2017_fp • Fan Control = Full Xeon Processor E-2100 Product Family: SPECspeed2017_int: • DCU Streamer Prefetcher = Disabled • DDR PowerDown and idle counter = PCODE • CState Pre-Wake = Disabled • Package C-State Un-demotion = Enabled • REFRESH_2X_MODE = 1- Enabled for WARM or HOT SPECspeed2017_fp • Energy Efficient Turbo = Disabled SPECrate2017_int: • Hardware Prefetcher = Disabled • Adjacent Cache Line Prefetch = Disabled • VT-d = Disabled • Fan Control = Full • Race To Halt (RTH) = Disabled • DMI Link ASPM Control = L0s • REFRESH_2X_MODE = 2- Enabled HOT only SPECrate2017_fp: • Fan Control = Full

	• Race To Halt (RTH) = Disabled • Energy Efficient Turbo = Disabled • DMI Link ASPM Control = Disabled • Package C-State Un-demotion = Enabled • Native PCIE Enable = Disabled CPUs equipped with HT: • Hyper-Threading = Disabled E-2174G, E-2176G, E-2186G: • Software Guard Extensions (SGX) = Disabled Xeon E-2288G: SPECspeed2017_int: • Adjacent Cache Line Prefetch = Disabled • CState Pre-Wake = Disabled • DCU Streamer Prefetcher = Disabled • DDR PowerDown and idle counter = PCODE • Energy Efficient Turbo = Disabled • Enhanced C-states = Disabled • Intel Virtualization Technology = Disabled • Native ASPM = Disabled • Package C-State Un-demotion = Enabled • REFRESH_2X_MODE = 1- Enabled for WARM or HOT SPECspeed2017_fp • Energy Efficient Turbo = Disabled • Fan Control = Full • Hyper Threading = Disabled • SW Gurad Extension (SGX) = Enabled SPECrate2017_int: • Adjacent Cache Line Prefetch = Disabled • C states = Disabled • Fan Control = Full • Hardware Prefetcher = Disabled • Intel Virtualization Technology = Disabled • Intel Speed Shift Technology = Disabled SPECrate2017_fp • AES = Disabled • DCU Streamer Prefetcher = Dsiabled • Fan Control = Full • Hyper-Threading = Disabled • Package C State Limit = C0
• Operating system	SPECspeed2017: Red Hat Enterprise Linux Server release 7.5 3.10.0-862.el7.x86_64 SPECrate2017:SUSE Linux Enterprise Server 15 4.12.14-23-default
• Operating system settings	Stack size set to unlimited using "ulimit -s unlimited" Xeon® Processor E-2100 Product Family: SPECrate2017

	echo 500000 > /proc/sys/kernel/sched_cfs_bandwidth_slice_us Xeon E-2288G: SPECrate2017_int、SPECspeed2017_int nohz_full = 1-15 SPECspeed2017_fp sched_min_granularity_ns = 100,000,000 sched_wakeup_granularity_ns = 150,000,000 sched_latency_ns = 240,000,000
• Compiler	Xeon Processor E-2100 Product Family: C/C++: Version 19.0.0.117 of Intel C/C++ Compiler for Linux Fortran: Version 19.0.0.117 of Intel Fortran Compiler for Linux Xeon E-2288G SPECrate2017 C/C++: Version 19.0.4.227 of Intel C/C++ Compiler for Linux Fortran: Version 19.0.4.227 of Intel Fortran Compiler for Linux SPECspeed2017 C/C++: Version 19.0.5.281 of Intel C/C++ Compiler for Linux Fortran: Version 19.0.5.281 of Intel Fortran Compiler for Linux

Benchmark results

For processors, the benchmark results depend 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 mainly load only one core, the maximum processor frequency that can be achieved is higher than with multi-threaded benchmarks.

Results with "est." are estimated values.

Processor	SPECspeed2017 int_base	SPECspeed2017 int_peak	SPECrate2017 int_base	SPECrate2017 int_peak
November 2018 released				
Pentium Gold G5400	7.06		12.7	13.6
Core i3-8100	7.97		15.4	15.6
Xeon E-2124	9.39		22.7	24.1
Xeon E-2124G	9.78		23.7	25.1
Xeon E-2126G	10.1		34.5	36.8
Xeon E-2134	10.0		29.6	31.7
Xeon E-2136	10.3		42.8	46.0
Xeon E-2144G	10.1		29.8	31.9
Xeon E-2146G	10.3		41.4	44.5
Xeon E-2174G	10.4		30.3	32.4
Xeon E-2176G	10.7		42.9	46.1
Xeon E-2186G	10.7		43.3	46.7
November 2019 released				
Pentium Gold G5420	7.22	7.57	13.0	13.9
Core i3-9100	9.67	9.88	26.3	27.2
Xeon E-2224	10.6	10.9	28.5	29.1
Xeon E-2224G	10.8	11.1	29.1	30.1
Xeon E-2226G	11.2	11.5	40.7	42.4
Xeon E-2234	11.3	11.6	35.8	36.6
Xeon E-2236	11.6	11.9	48.1	48.1
Xeon E-2244G	11.3	11.6	35.1	36.9
Xeon E-2246G	11.6	11.9	46.2	48.9
Xeon E-2274G	11.4	11.8	34.6	36.4
Xeon E-2276G	11.8	12.0	46.5	49.1
Xeon E-2278G	12.1	12.4	57.6	61.0
Xeon E-2286G	11.8	12.0	47.9	50.5
Xeon E-2288G	12.0	12.4	60.5	63.9

On November 13, 2018, the PRIMERGY TX1330 M4 with a Xeon E-2186G processor was ranked first (tie) in the 1-socket Xeon E server systems category for the benchmark SPECspeed2017 int_base.

On November 1, 2019 the PRIMERGY TX1330 M4 with a Xeon E-2288G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECrate2017 int_base.

Processor	SPECspeed2017 fp_base	SPECspeed2017 fp_peak	SPECrate2017 fp_base	SPECrate2017 fp_peak
November 2018 released				
Pentium Gold G5400	**13.5**		**21.9**	**23.3**
Core i3-8100	**23.2**		**26.4**	**26.8**
Xeon E-2124	**26.0**		**29.5**	**30.0**
Xeon E-2124G	**26.6**		**30.4**	**31.0**
Xeon E-2126G	**31.5**		**37.4**	**38.1**
Xeon E-2134	**26.8**		**30.5**	**31.1**
Xeon E-2136	**31.9**		**38.0**	**38.7**
Xeon E-2144G	**26.7**		**30.7**	**31.3**
Xeon E-2146G	**31.7**		**37.1**	**37.9**
Xeon E-2174G	**27.2**		**31.0**	**29.5**
Xeon E-2176G	**32.1**		**38.2**	**39.0**
Xeon E-2186G	**32.1**		**38.5**	**39.2**
November 2019 released				
Pentium Gold G5420	**13.8**	**14.1**	**15.7**	**16.0**
Core i3-9100	**24.5**	**24.8**	**27.9**	**28.5**
Xeon E-2224	**27.0**	**27.4**	**30.8**	**31.4**
Xeon E-2224G	**27.4**	**27.8**	**31.6**	**32.2**
Xeon E-2226G	**32.3**	**32.7**	**37.7**	**38.5**
Xeon E-2234	**27.6**	**28.0**	**31.7**	**32.3**
Xeon E-2236	**32.3**	**32.7**	**37.9**	**38.7**
Xeon E-2244G	**27.8**	**28.2**	**32.0**	**32.5**
Xeon E-2246G	**32.5**	**32.9**	**38.1**	**38.9**
Xeon E-2274G	**27.5**	**27.9**	**31.6**	**32.1**
Xeon E-2276G	**33.1**	**33.5**	**38.8**	**39.5**
Xeon E-2278G	**36.3**	**36.8**	**43.4**	**44.3**
Xeon E-2286G	**33.3**	**33.6**	**39.3**	**40.0**
Xeon E-2288G	**36.8**	**37.3**	**44.8**	**45.7**

On November 13, 2018, the PRIMERGY TX1330 M4 with a Xeon E-2186G processor was ranked first (tie) in the 1-socket Xeon E server systems category for the benchmark SPECspeed2017 fp_base.

On November 1, 2019, the PRIMERGY TX1330 M4 with a Xeon E-2288G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECspeed2017 fp_base.

On November 1, 2019, the PRIMERGY TX1330 M4 with a Xeon E-2288G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECrate2017 fp_base.

The following two diagrams illustrate the comparison of SPECrate2017 on Xeon E-2186G and Xeon E-2288G measured with PRIMERGY TX1330 M4, in their respective most performance configurations.

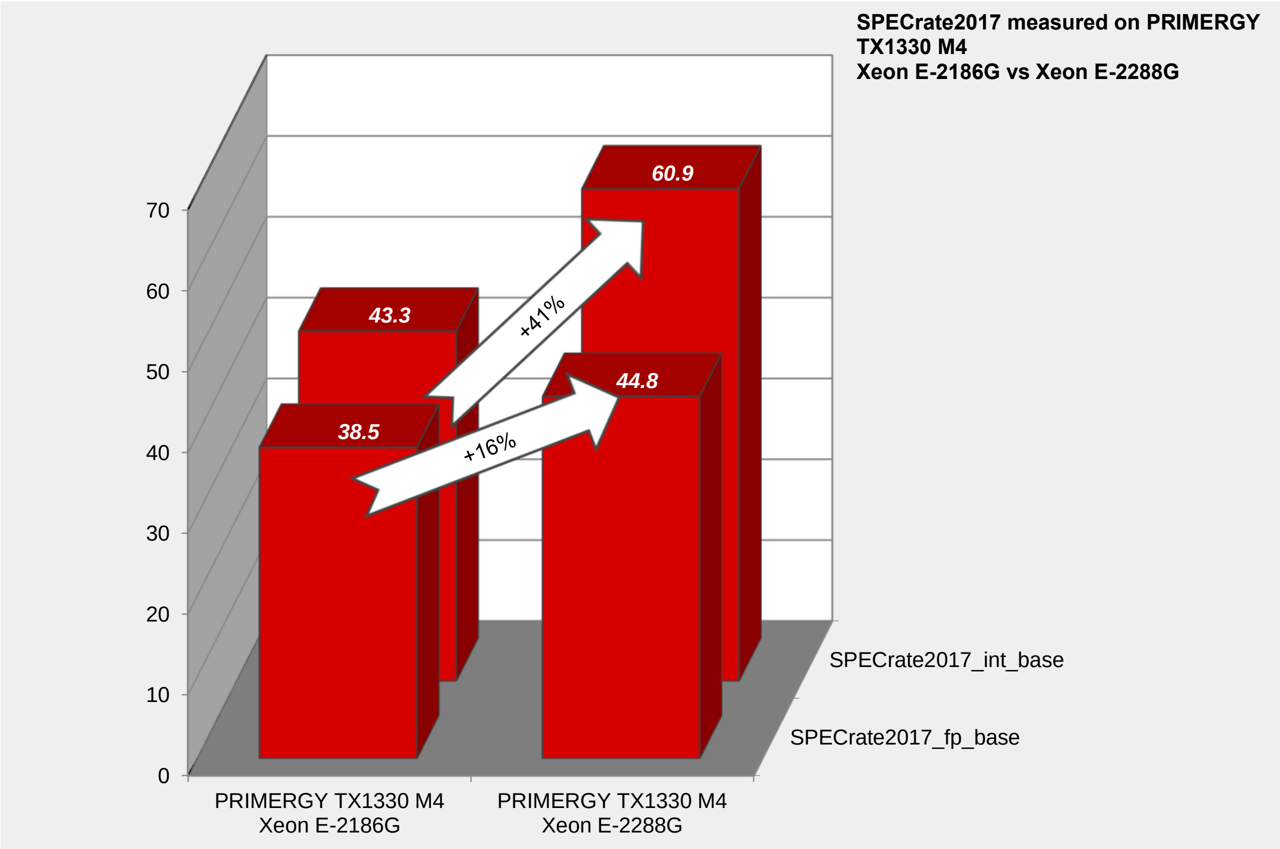

STREAM

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 TX1330 M4
• Processor	Intel Pentium Gold G5400 Intel Pentium Gold G5420 Intel Core i3-8100 Intel Core i3-9100 Intel Xeon Processor E-2100 Product Family Intel Xeon Processor E-2200 Product Family
• Memory	4 × 16 GB (1 x 16 GB) 2Rx8 DDR4-2666 U ECC
Software	
• BIOS settings	• Fan Control = Full
• Operating system	SUSE Linux Enterprise Server 15 (x86_64)
• Compiler	Version 18.0.2.199 of Intel C/C++ Compiler for Linux
• Benchmark	Stream.c Version 5.10

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

Benchmark results

Results with "est." are estimated values.

Processor	Memory Frequency	Max. Memory Bandwidth	Cores	Processor Frequency	TRIAD
	[MHz]	[GB/s]		[GHz]	[GB/s]
November 2018 released					
Pentium Gold G5400	2,400	37.5	2	3.7	33.8
Core i3-8100	2,400	37.5	4	3.6	33.0
Xeon E-2124	2,666	41.6	4	3.3	36.7
Xeon E-2124G	2,666	41.6	4	3.4	37.0
Xeon E-2126G	2,666	41.6	6	3.3	36.2
Xeon E-2134	2,666	41.6	4	3.5	36.7
Xeon E-2136	2,666	41.6	6	3.3	35.8
Xeon E-2144G	2,666	41.6	4	3.6	36.6
Xeon E-2146G	2,666	41.6	6	3.5	35.9
Xeon E-2174G	2,666	41.6	4	3.8	36.9
Xeon E-2176G	2,666	41.6	6	3.7	36.2
Xeon E-2186G	2,666	41.6	6	3.8	35.9
November 2019 released					
Pentium Gold G5420	2,400	37.5	2	3.8	34.2
Core i3-9100	2,400	37.5	4	3.6	33.2
Xeon E-2224	2,666	41.6	4	3.4	36.8
Xeon E-2224G	2,666	41.6	4	3.5	36.6
Xeon E-2226G	2,666	41.6	6	3.4	35.4
Xeon E-2234	2,666	41.6	4	3.6	36.8
Xeon E-2236	2,666	41.6	6	3.4	35.4
Xeon E-2244G	2,666	41.6	4	3.8	36.8
Xeon E-2246G	2,666	41.6	6	3.6	35.4
Xeon E-2274G	2,666	41.6	4	4.0	36.8
Xeon E-2276G	2,666	41.6	6	3.8	35.5
Xeon E-2278G	2,666	41.6	8	3.4	34.8
Xeon E-2286G	2,666	41.6	6	4.0	35.4
Xeon E-2288G	2,666	41.6	8	3.7	34.9

The following diagram shows the comparison of STREAM TRIAD on the E-2200 Product Family and its predecessor, the E-2100 Product Family measured with PRIMERGY TX1330 M4.

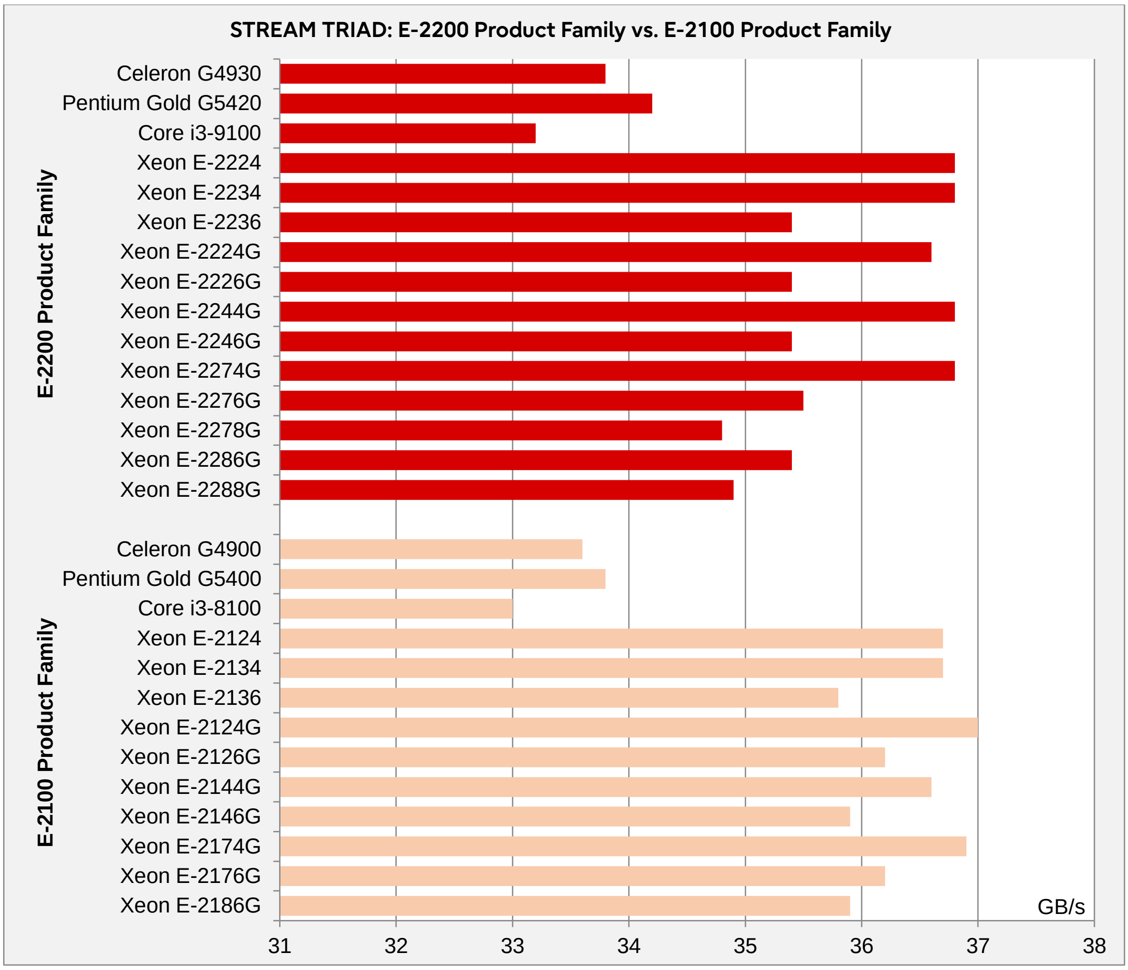

SPECpower_ssj2008

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.

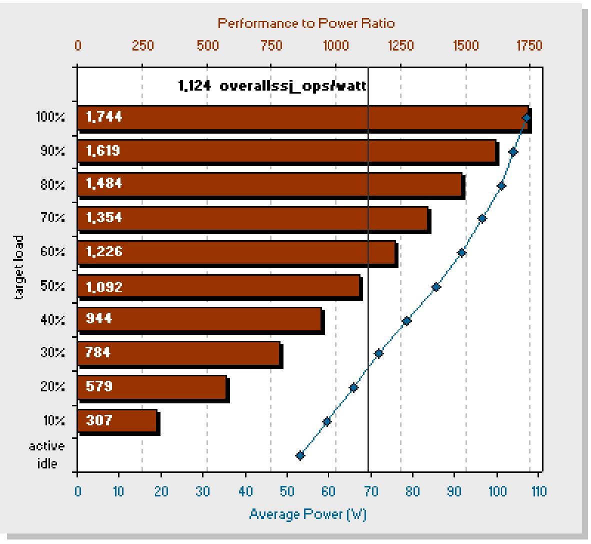

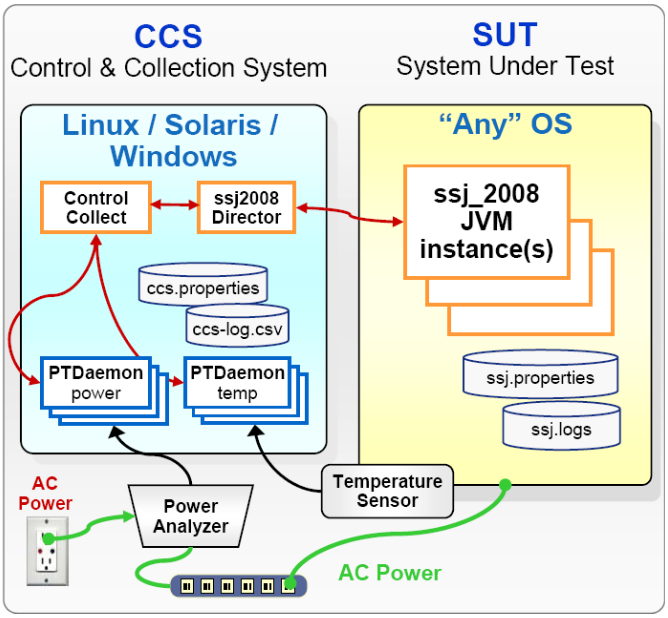

The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System 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)	
Hardware	
• Model	PRIMERGY TX1330 M4
• Processor	Intel Xeon Platinum E-2176G
• Memory	2 × 8 GB (1 x 8 GB) 1Rx8 DDR4-2666 U ECC
• Network interface	2 × Intel I210 Gigabit Network Connection
• Disk subsystem	Onboard SATA. controller 1 × M.2 SSD 240 GB, S26361-F5706-E240
• Power Supply Unit	1 × Standard PSU 300 W
Software	
• BIOS	R1.1.0
• BIOS settings	ASPM Support = Force F0s Adjacent Cache Line Prefetch = Disabled Hardware Prefecher = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled SATA Port 1/2/3/4/5/7 = Disabled Turbo = Disabled Serial port = Disabled LAN2 = Disabled USB Port Control = Enable internal ports only Software Guard Extensions = Disabled Network Stack = Disabled
• Firmware	1.60h
• Operating system	Microsoft Windows Server 2012 R2 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. The benchmark was started via Windows Remote Desktop Connection. SPECpower_ssj.props input.load_level.number_warehouses set to 12.
• JVM	Oracle Java HotSpot(TM) 64-Bit Server VM 18.9(build 11+28, mixed mode), version 11
• JVM settings	-server -Xmn10500m -Xms12000m -Xmx12000m -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:ParallelGCThreads=2 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:FreqInlineSize=2500 -XX:+UseLargePages -XX:+UseParallelOldGC -XX:AllocatePrefetchInstr=0 -XX:MinJumpTableSize=18 -XX:UseAVX=0 -XX:TenuredGenerationSizeSupplement=40 -XX:-UseFastStosb

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

Benchmark results

The PRIMERGY TX1330 M4 achieved the following result:

SPECpower_ssj2008 = 9,031 overall ssj_ops/watt

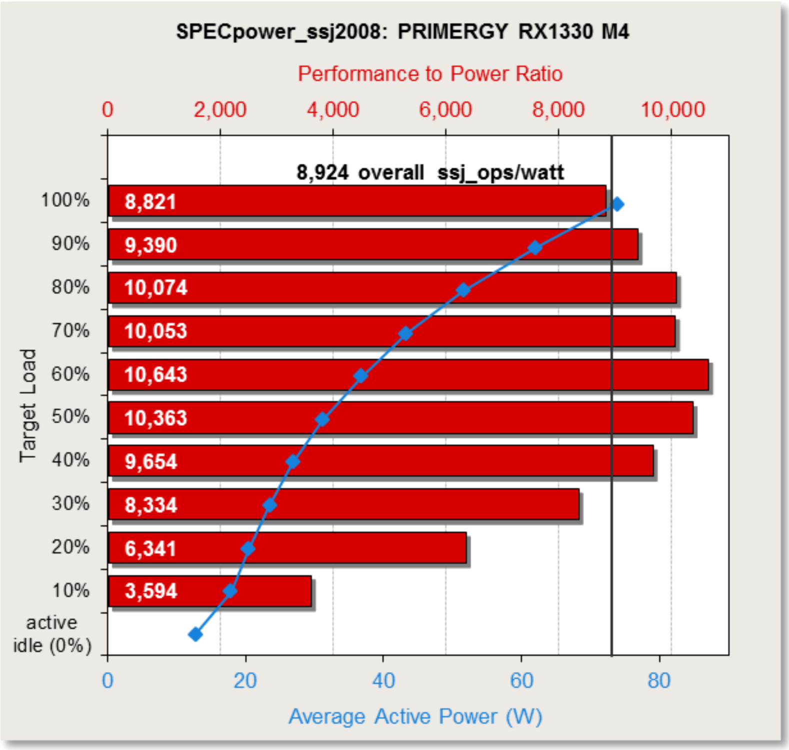

The adjoining diagram shows the results for the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (top x-axis) for each target load level on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhombus. The black vertical line shows the benchmark result of 9,031 overall ssj_ops/watt for the PRIMERGY TX1330 M4. 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%	651,586	73.9	8,821
90%	582,666	62.1	9,390
80%	519,605	51.6	10,074
70%	455,039	43.3	10,053
60%	391,263	36.8	10,643
50%	322,940	31.2	10,363
40%	259,677	26.9	9,654
30%	195,508	23.5	8,334
20%	130,148	20.5	6,341
10%	64,423	17.9	3,594
Active Idle	0	12.8	0
			Σ ssj_ops / Σ power = 9,031

Difference of score by OS & JVM version

The SPECpower_ssj2008 score differs about 10% depending on the OS used in the system. The OS itself has an influence on performance. This means that the usable JVM version will be different depending on the OS type. Currently, the combinations Windows Server 2012 R2 & JVM7, Windows Server 2016 & JVM11, and Linux & JVM7 are used in Fujitsu and other vendors' submission results.
With appropriate OS settings and JVM options, the score increases in the order Linux & JVM7 ≥ Windows Server 2012 R2 & JVM7 > Windows Server2016 & JVM11.
There are very few differences between Linux & JVM7 and Windows Server 2012. On the other hand, the score for the combination Windows Server 2016 & JVM11 is about 10% lower than the other two combinations' scores.
Under the SPECpower_ssj2008 rules, Windows Server 2016 (a relatively new OS) is not allowed to be measured with JVM7. Therefore, a later JVM version has to be used. Alt-rt.jar, a module included in JVM7, is related to accelerated collection type HashMap. However, the module has been deleted in JVM11. This is the main reason that the SPECpower_ssj2008 score measured with JVM11 is lower.

For this reason, the SPECpower_ssj2008 score measured with TX1330 M4 and JVM 11 is lower than the score measured with TX1330 M3 and JVM 7. However, this does not mean that TX1330 M4 is inferior to the predecessor system TX1330 M4 in terms of power efficiency. Fujitsu has verified that the SPECpower_ssj2008 score is higher than that for TX1330 M3 if the same JVM version is used.

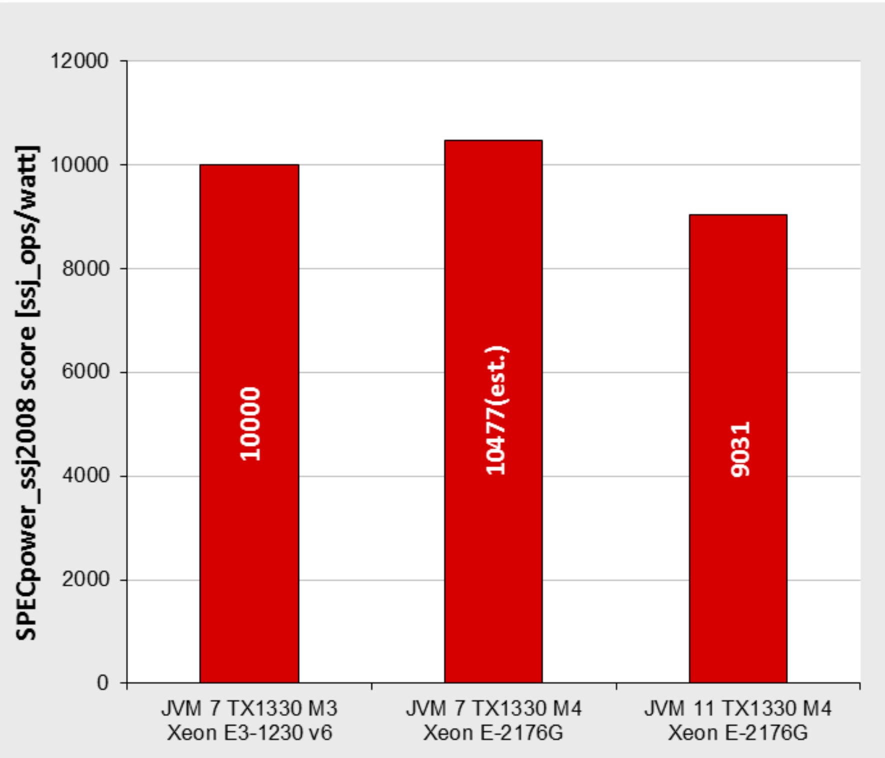

Benchmark environment (E-2200 Product Family)

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY TX1330 M4
• Processor	Intel Xeon E-2288G
• Memory	2 × 8 GB (1 x 8 GB) 1Rx8 DDR4-2666V-E
• Network interface	2 × Intel I210 Gigabit Network Connection (on board)
• Disk subsystem	Onboard SATA. controller 1 × M.2 SSD 240 GB, S26361-F5706-E240
• Power Supply Unit	1 × Standard power supply part of base unit S26361-K1639-V101
Software	
• BIOS	R1.12.0
• BIOS settings	ASPM Support = Force F0s Adjacent Cache Line Prefetch = Disabled Hardware Prefecher = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled SATA Port 1/2/3/4/5/7 = Disabled Turbo = Disabled Serial port = Disabled LAN2 = Disabled Management LAN = Disabled USB Port Control = Enable internal ports only Software Guard Extensions = Disabled Network Stack = Disabled DMI Max Link Speed = Gen1 Enabled ACPI Auto Configuration = Enabled Native_ASPM=Disabled DDR Performance = Power balanced C States=Disabled
• Firmware	2.50P
• Operating system	SUSE Linux Enterprise Server 12 SP4 4.12.14-94.41-default

• Operating system settings	kernel parameter:pcie_aspm=force pcie_aspm.policy=powersave intel_pstate=disable Benchmark started via ssh modprobe cpufreq_conservative cpupower frequency-set --governor conservative echo -n 94 > /sys/devices/system/cpu/cpufreq/conservative/up_threshold echo -n 93 > /sys/devices/system/cpu/cpufreq/conservative/down_threshold echo -n 1 > /sys/devices/system/cpu/cpufreq/conservative/freq_step echo -n 1000000 > /sys/devices/system/cpu/cpufreq/conservative/sampling_rate echo -n 0 > /sys/devices/system/cpu/cpufreq/conservative/ignore_nice_load sysctl -w kernel.sched_migration_cost_ns=6000 echo -n 1 > /sys/devices/system/cpu/cpufreq/conservative/sampling_down_factor sysctl -w kernel.sched_min_granularity_ns=10000000 echo always > /sys/kernel/mm/transparent_hugepage/enabled cpupower frequency-set -u 2800MHz powertop --auto-tune echo 0 > /proc/sys/kernel/nmi_watchdog sysctl -w vm.swappiness=50 sysctl -w vm.laptop_mode=5
• JVM	Oracle Java HotSpot(TM) 64-Bit Server VM (build 24.80-b11, mixed mode), version 1.7.0_80
• JVM settings	-server -Xmn10500m -Xms12000m -Xmx12000m -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -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:+UseHugeTLBFS -XX:+UseTransparentHugePages -XX:UseAVX=0

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

Benchmark results (E-2200 Product Family)

By applying latest E-2200 Product Family processor, the PRIMERGY TX1330 M4 in SUSE Linux Enterprise Server 12 SP4 achieved the following result:

SPECpower_ssj2008 = 12,631 overall ssj_ops/watt

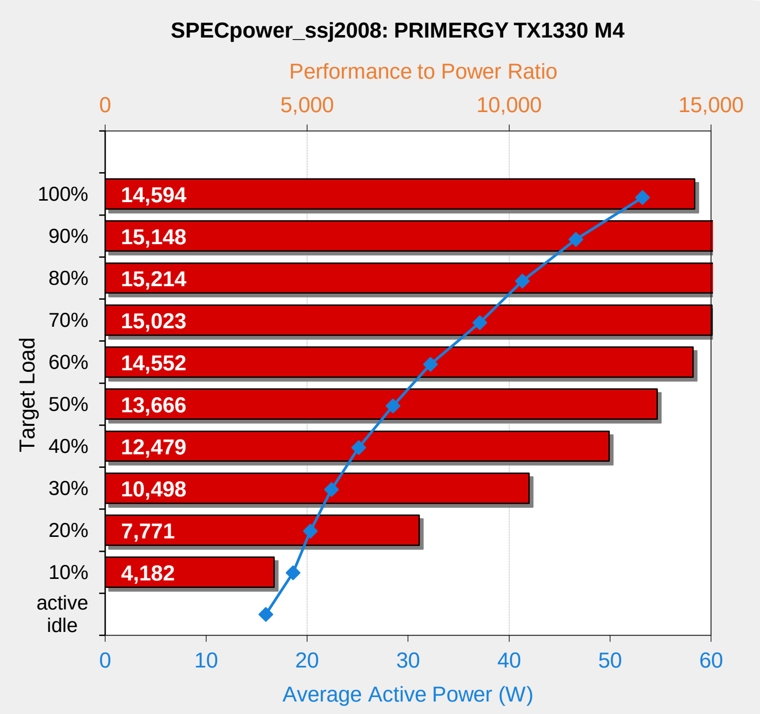

The adjoining diagram shows the results for the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (top x-axis) for each target load level on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhombus. The black vertical line shows the benchmark result of 12,631 overall ssj_ops/watt for the PRIMERGY TX1330 M4. 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%	776,812	53.2	14,594
90%	705,646	46.6	15,148
80%	628,060	41.3	15,214
70%	556,871	37.1	15,023
60%	469,204	32.2	14,552
50%	389,284	28.5	13,666
40%	312,897	25.1	12,479
30%	235,087	22.4	10,498
20%	157,397	20.3	7,771
10%	77,883	18.6	4,182
Active Idle	0	15.9	0
			Σ ssj_ops / Σ power = 12,631

In this way, latest E-2200 Product Family processor enables to achieve better power efficiency.

SPECjbb2015

Benchmark description

The SPECjbb2015 benchmark is the latest version of a series of Java benchmarks following SPECjbb2000, SPECjbb2005 and SPECjbb2013. “jbb” stands for Java Business Benchmark. It evaluates the performance and the scalability of the Java business application environment.

The SPECjbb2015 is a benchmark modeled on the business activity of a world-wide supermarket company’s IT infrastructure. The company has some supermarket stores, headquarters which manage them and suppliers who replenish their inventory. The following processing is performed based on the requests from customers and inside the company.

- POS (Point Of Sales) processing in supermarkets and processing of online purchases
- Issuing and managing coupons and discounts and customer payments management
- Managing receipts, invoices and customer databases
- Interaction with suppliers to replenish inventory
- Data mining operations to identify sale patterns and generate quarterly business reports

The SPECjbb2015 benchmark has two metrics:

- max-jOPS: This is the maximum transaction rate that can be achieved by the system under test while meeting the benchmark constraints. That is, it is a metric of the maximum processing throughput of the system.
- critical-jOPS: This is the geometric mean of the maximum transaction rates that can be achieved while meeting the constraints on the response time of 10, 25, 50, 75, and 100 milliseconds. In other words, it is a metric of the maximum processing throughput of the system under response time constraints.

The SPECjbb2015 benchmark consists of the following three components: Backends (BE), which contains the business logic and data, Transaction Injector (TxI), which issues transaction requests, and Controller (Ctr), which controls them. Through the configuration of these components, the benchmark is divided into the following three categories:

- SPECjbb2015 Composite

  All components run on one JVM running on one host.
- SPECjbb2015 MultiJVM

  All components exist on one host, but each runs on a separate JVM.
- SPECjbb2015 Distributed

  The Backends exists on a separate host from those on which the other components are running. The Backends and the other components are connected by networks.

Results can not be compared between categories.

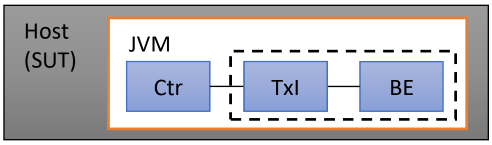

(a) Example of SPECjbb2015 Composite configuration

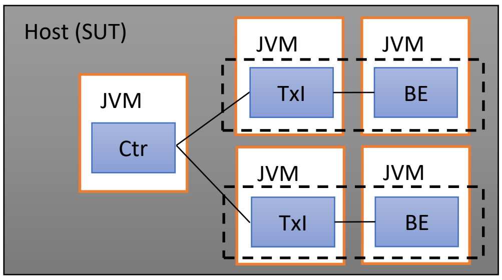

(b) Example of SPECjbb2015 MultiJVM configuration

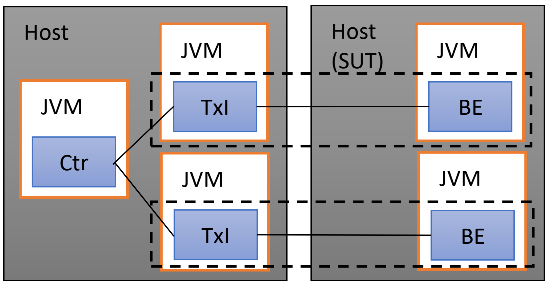

(c) Example of SPECjbb2015 Distributed configuration

The results of the SPECjbb2015 benchmark reflects not only the performance of the Java runtime environment (JRE) but also the performance of the operating system and the hardware underneath it. For the JRE, factors like the Java Virtual Machine (JVM), the Just-in-time (JIT) Compiler, garbage collection, and user thread affect the performance score, and for the hardware, it is affected by the performance of the processors, memory subsystem, and network. The SPECjbb2015 benchmark does not cover disk I/O performance.

The detailed specifications of the benchmark can be found at https://www.spec.org/jbb2015/.

Benchmark environment

PRIMERGY RX1330 M4 was configured for the SPECjbb2015 MultiJVM benchmark measurement.

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY TX1330 M4
• Processor	1 × Intel Xeon E-2186G
• Memory	4 × 16 GB (1 x 16 GB) 2Rx8 DDR4-2666 U ECC
• Network interface	1 Gbit/s LAN
• Disk subsystem	Disk : 1 × 2TB SATA HDD
Software	
For measurement result (1)	
• BIOS settings	Intel(VMX) Virtualization Technology set to Disabled C states set to Disabled
• Operating system	SUSE Linux Enterprise Server 15 4.12.14-23-default
• Operating system settings	cpupower -c all frequency-set -g performance echo 16000000 > /proc/sys/kernel/sched_latency_ns echo 1500 > /proc/sys/vm/dirty_writeback_centisecs
• JVM	Oracle Java SE 10.0.2
• JVM settings	-server -Xms27g -Xmx27g -Xmn26g -XX:SurvivorRatio=22 --add-modules=java.xml.bind -XX:MaxTenuringThreshold=15 -XX:+UseParallelOldGC -Xnoclassgc -XX:+UseNUMA -XX:-UseBiasedLocking -XX:-UseAdaptiveSizePolicy -XX:TargetSurvivorRatio=90 -XX:ParallelGCThreads=12 -verbose:gc -XX:+UseHugeTLBFS -XX:ActiveProcessorCount=0 -XX:+UseRTMDeopt -XX:MaxGCPauseMillis=300
For measurement result (2)	
• BIOS settings	Intel(VMX) Virtualization Technology set to Disabled C states set to Disabled Hardware Prefetcher set to Disabled VT-d set to Disabled
• Operating system	SUSE Linux Enterprise Server 15 4.12.14-23-default
• Operating system settings	cpupower -c all frequency-set -g performance echo 16000000 > /proc/sys/kernel/sched_latency_ns echo 1500 > /proc/sys/vm/dirty_writeback_centisecs
• JVM	Oracle Java SE 10.0.2
• JVM settings	-server -Xms27g -Xmx27g -Xmn26g -XX:SurvivorRatio=22 --add-modules=java.xml.bind -XX:MaxTenuringThreshold=15 -XX:+UseParallelOldGC -Xnoclassgc -XX:+UseNUMA -XX:-UseBiasedLocking -XX:-UseAdaptiveSizePolicy -XX:TargetSurvivorRatio=90 -XX:ParallelGCThreads=12 -verbose:gc -XX:+UseHugeTLBFS -XX:ActiveProcessorCount=0 -XX:+UseCondCardMark -XX:+UseRTMDeopt -XX:InlineSmallCode=10k
For measurement result (3)	
• BIOS settings	Intel(VMX) Virtualization Technology set to Disabled C states set to Disabled Hardware Prefetcher set to Disabled

	VT-d set to Disabled VT-d set to Disabled
• Operating system	SUSE Linux Enterprise Server 15 4.12.14-23-default
• Operating system settings	cpupower -c all frequency-set -g performance echo 16000000 > /proc/sys/kernel/sched_latency_ns echo 1500 > /proc/sys/vm/dirty_writeback_centisecs
• JVM	Oracle Java SE 10.0.2
• JVM settings	-server -Xms27g -Xmx27g -Xmn26g -XX:SurvivorRatio=22 --add-modules=java.xml.bind -XX:MaxTenuringThreshold=15 -XX:+UseParallelOldGC -Xnoclassgc -XX:+UseNUMA -XX:-UseBiasedLocking -XX:-UseAdaptiveSizePolicy -XX:TargetSurvivorRatio=90 -XX:ParallelGCThreads=12 -verbose:gc -XX:+UseHugeTLBFS -XX:ActiveProcessorCount=0 -XX:+UseCondCardMark -XX:+UseRTMDeopt -XX:+UseParallelGC

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

Benchmark results

"SPECjbb2015 Composite" measurement result (1) (November 6, 2018)
26,012 SPECjbb2015-Composite max-jOPS
13,098 SPECjbb2015-Composite critical-jOPS

"SPECjbb2015 Composite" measurement result (2) (November 8, 2018)
27,083 SPECjbb2015-Composite max-jOPS
12,653 SPECjbb2015-Composite critical-jOPS

On November 6, 2018 PRIMERGY TX1330 M4 with a Xeon E-2186G processor achieved the scores of 27,083 SPECjbb2015-Composite max-jOPS. With this result, it ranked first in the 1-socket Xeon E server category for SPECjbb2015-Composite max-jOPS.

"SPECjbb2015 Composite" measurement result (3) (November 8, 2018)
26,397 SPECjbb2015-Composite max-jOPS
13,426 SPECjbb2015-Composite critical-jOPS

On November 8, 2018 PRIMERGY TX1330 M4 with a Xeon E-2186G processor achieved the scores of 13,426 SPECjbb2015-Composite critical-jOPS. With this result, it ranked first in the 1-socket Xeon E server category for SPECjbb2015-Composite critical-jOPS.

The latest results of the SPECjbb2015 benchmark can be found at https://www.spec.org/jbb2015/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: 1x PRAID CP400i

Storage media	Category	Drive Name
HDD	SAS HDD (SAS 12 Gbps, 10k rpm) [512e]	AL15SEB18EQ *2 *3
	SAS HDD (SAS 12 Gbps, 10k rpm) [512n]	AL15SEB030N *2 *3
	SAS HDD (SAS 12 Gbps, 15k rpm) [512n]	ST300MP0006 *1 *3
	NL-SAS HDD (SAS 12 Gbps, 7.2k rpm) [512e]	HUH721212AL5204 *2 *3
	NL-SAS HDD (SAS 12 Gbps, 7.2k rpm) [512n]	ST2000NM0045 *1 *3
	BC-SATA HDD (SATA 6 Gbps, 7.2k rpm) [512e]	ST6000NM0115 *1 *3 HUH721212ALE604 *2 *3
	BC-SATA HDD (SATA 6 Gbps, 7.2k rpm) [512n]	HUS722T1TALA604 *2 *3 ST2000NM0055 *1 *3
	SATA HDD (SATA 6 Gbps, 7.2 krpm) [512e]	ST1000DM003-1SB *1 *3
SSD	SATA SSD (SATA 6 Gbps, Mixed Use)	MZ7KH240HAHQ *2 *3 MZ7KH480HAHQ *2 *3 MZ7KH960HAJR *2 *3 MZ7KH1T9HAJR *2 *3 MZ7KH3T8HALS *2 *3
	SATA SSD (SATA 6 Gbps, Read Intensive)	MTFDDAK240TCB *2 *3 MTFDDAK480TDC *2 *3 MTFDDAK960TDC *2 *3 MTFDDAK1T9TDC *2 *3 MTFDDAK3T8TDC *2 *3 MTFDDAK7T6TDC *2 *3

Controller: Intel® C620 Standard SATA AHCI controller

Storage media	Category	Drive Name
SSD	M.2 Flash Module	MTFDDAV240TCB *2 *4 MTFDDAV480TCB *2 *4

2.5 inch Model:

Controller: 1x PRAID CP400i

Storage media	Category	Drive Name
HDD	SAS HDD (SAS 12 Gbps, 10k rpm) [512e]	AL15SEB06EQ *2 *3
	SAS HDD (SAS 12 Gbps, 10k rpm) [512n]	AL15SEB030N *2 *3
	SAS HDD (SAS 12 Gbps, 15k rpm) [512n]	ST300MP0006 *1 *3
	BC-SATA HDD (SATA 6 Gbps, 7.2k rpm) [512e]	ST1000NX0313 *1 *3
	BC-SATA HDD (SATA 6 Gbps, 7.2k rpm) [512n]	ST2000NX0403 *1 *3
SSD	SATA SSD (SATA 6 Gbps, Mixed Use)	MZ7KH240HAHQ *2 *3 MZ7KH480HAHQ *2 *3 MZ7KH960HAJR *2 *3 MZ7KH1T9HAJR *2 *3 MZ7KH3T8HALS *2 *3
	SATA SSD (SATA 6 Gbps, Read Intensive)	MTFDDAK240TCB *2 *3 MTFDDAK480TDC *2 *3 MTFDDAK960TDC *2 *3 MTFDDAK1T9TDC *2 *3 MTFDDAK3T8TDC *2 *3 MTFDDAK7T6TDC *2 *3

Controller: Intel® C620 Standard SATA AHCI controller

Storage media	Category	Drive Name
SSD	M.2 Flash Module	MTFDDAV240TCB *2 *4 MTFDDAV480TCB *2 *4

*1) Operating system used Microsoft Windows Server 2012 Standard R2.

*2) Operating system used Microsoft Windows Server 2016 Standard.

*3) Measurement area is type 1.

*4) Measurement area is type 2.

.

Software		
Operating system		Microsoft Windows Server 2012 Standard R2 Microsoft Windows Server 2016 Standard
Benchmark version		3.0
RAID type		Logical drive of type RAID 0 consisting of 1 hard disk
Stripe size		Controller default (here, 64 kiB)
Measuring tool		Iometer 1.1.0
Measurement area	Type 1	RAW file system is used. The first 10% of the usable LBA area is used for sequential accesses; the next 25% is used for random accesses.
	Type 2	NTFS file system is used. A 32 GiB area is secured at the top of the target drive and is used for sequential accesses and random accesses.
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 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	PRAID CP400i	-	8x PCIe 3.0	SATA 6G SAS 12G	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 type	Storage medium type	Interface	Form factor
3.5 inch model	HDD	SAS 12G	3.5 inch or 2.5 inch1)
		SATA 6G	3.5 inch
	SSD	SATA 6G	M.2
2.5 inch model	HDD	SAS 12G	2.5 inch
		SATA 6G	2.5 inch
	SSD	SATA 6G	2.5 inch or M.2

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
1,800	AL15SEB18EQ	SAS 12G	600	512	547	258	255
300	AL15SEB030N	SAS 12G	645	546	568	231	230
300	ST300MP0006	SAS 12G	768	662	472	304	304
12,000	HUH721212AL5204	SAS 12G	396	339	364	245	244
2,000	ST2000NM0045	SAS 12G	376	336	343	206	206
6,000	ST6000NM0115	SATA 6G	392	362	371	213	208
12,000	HUH721212ALE604	SATA 6G	350	313	341	246	246
1,000	HUS722T1TALA604	SATA 6G	287	264	269	201	201
2,000	ST2000NM0055	SATA 6G	339	301	314	196	195
1,000	ST1000DM003-1SB	SATA 6G	222	210	204	208	203

SSDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
240	MZ7KH240HAHQ	SATA 6G	49,159	7,313	7,431	526	486
480	MZ7KH480HAHQ	SATA 6G	50,558	7,774	7,810	526	485
960	MZ7KH960HAJR	SATA 6G	50,647	7,793	7,916	525	485
1,920	MZ7KH1T9HAJR	SATA 6G	50,702	8,040	7,960	526	485
3,840	MZ7KH3T8HALS	SATA 6G	50,766	8,039	7,936	526	485
240	MTFDDAK240TCB	SATA 6G	18,959	3,367	4,516	487	258
480	MTFDDAK480TDC	SATA 6G	24,710	3,799	5,006	507	362
960	MTFDDAK960TDC	SATA 6G	30,152	4,625	5,553	507	440
1,920	MTFDDAK1T9TDC	SATA 6G	37,234	5,606	5,566	507	483
3,840	MTFDDAK3T8TDC	SATA 6G	41,711	6,429	6,133	504	481
7,680	MTFDDAK7T6TDC	SATA 6G	40,683	6,874	6,672	469	482
240	MTFDDAV240TCB	SATA 6G	20,113	3,936	5,021	510	271
480	MTFDDAV480TCB	SATA 6G	22,596	4,993	6,331	509	403

2.5 inch model

HDDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
1,200	AL15SEB12EQ	SAS 12G	594	520	546	260	259
300	AL15SEB030N	SAS 12G	645	546	568	231	230
300	ST300MP0006	SAS 12G	768	662	472	304	304
1,000	ST1000NX0313	SATA 6G	324	281	288	131	131
2,000	ST2000NX0403	SATA 6G	326	286	294	133	133

SSDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
240	MZ7KH240HAHQ	SATA 6G	49,159	7,313	7,431	526	486
480	MZ7KH480HAHQ	SATA 6G	50,558	7,774	7,810	526	485
960	MZ7KH960HAJR	SATA 6G	50,647	7,793	7,916	525	485
1,920	MZ7KH1T9HAJR	SATA 6G	50,702	8,040	7,960	526	485
3,840	MZ7KH3T8HALS	SATA 6G	50,766	8,039	7,936	526	485
240	MTFDDAK240TCB	SATA 6G	18,959	3,367	4,516	487	258
480	MTFDDAK480TDC	SATA 6G	24,710	3,799	5,006	507	362
960	MTFDDAK960TDC	SATA 6G	30,152	4,625	5,553	507	440
1,920	MTFDDAK1T9TDC	SATA 6G	37,234	5,606	5,566	507	483
3,840	MTFDDAK3T8TDC	SATA 6G	41,711	6,429	6,133	504	481
7,680	MTFDDAK7T6TDC	SATA 6G	40,683	6,874	6,672	469	482
240	MTFDDAV240TCB	SATA 6G	20,113	3,936	5,021	510	271
480	MTFDDAV480TCB	SATA 6G	22,596	4,993	6,331	509	403

Literature

PRIMERGY Servers

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

PRIMERGY TX1330 M4

This Whitepaper

http://docs.ts.fujitsu.com/dl.aspx?id=b7e38640-3198-4c56-9f0e-f2b6a5887baa

http://docs.ts.fujitsu.com/dl.aspx?id=773f5e9a-1d20-44d6-9714-59346ab3d307

Data sheet

http://docs.ts.fujitsu.com/dl.aspx?id=043850d3-9e2f-4897-96fe-a7fa7ead8af6

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/

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

SPECjbb2015

https://www.spec.org/jbb2015/

Document change history

Version	Date	Description
1.5	2023-10-03	Update: • New Visual Identity format
1.4	2020-04-17	Update: • SPECcpu2017 Measured with Pentium Gold G5440, Core i3-9100 and Intel® Xeon® Processor E-2200 Product Family • STREAM Measured with Pentium Gold G5440, Core i3-9100 and Intel® Xeon® Processor E-2200 Product Family
1.3	2019-12-27	Update: • Technical data Added Intel® Xeon® Processor E-2200 Product Family • SPECcpu2017 Added measurement with Intel® Xeon® Processor E-2200 Product Family • SPECpower_ssj2008 Added measurement with Intel® Xeon® Processor E-2200 Product Family • Disk I/O: Performance of storage media Results for 2.5" and 3.5" storage media
1.2	2019-08-05	Update: • SPECpower_ssj2008 Changed the description about the OS and JVM version
1.1	2018-12-25	New: • SPECpower_ssj2008 Measurement with Intel® Xeon® E-2176G • Disk I/O: Performance of storage media Results for 2.5" and 3.5" storage media

Document change history

Version	Date	Description
1.0	2018-11-27	New: • Technical data • SPECcpu2017 Measurements with Pentium Gold G5400, Core i3-8100 and Intel Xeon Processor E-2100 Product Family • SPECjbb2015 Measurement with Intel® Xeon® E-2186G • STREAM Measurements with Pentium Gold G5400, Core i3-8100 and Intel® Xeon® Processor E-2100 Product Family

Contact
Fujitsu
Web site: https://www.fujitsu.com
PRIMERGY Performance and Benchmarks
mailto:fj-benchmark@dl.jp.fujitsu.com