White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY TX1320 M4

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY TX1320 M4.

The PRIMERGY TX1320 M4 performance data are compared with the data of other PRIMERGY models and discussed. In addition to the benchmark results, an explanation has been included for each benchmark and for the benchmark environment.

Version
1.4
2020/04/17

Contents

Document history

Version 1.0 (2018/11/27)

New:

- Technical data
- SPECcpu2017
  Measurements with Celeron G4900, Pentium Gold G5400 and Intel® Xeon® Processor E-2100 Product Family
- SPECjbb2015
  Measurement with Intel® Xeon® E-2186G
- STREAM
  Measurements with Celeron G4900, Pentium Gold G5400, Core i3-8100 and Intel® Xeon® Processor E-2100 Product Family

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

Version 1.2 (2019/08/05)

Updated:

- SPECpower_ssj2008
  Added new measurement results on SLES12SP4

Version 1.3 (2019/12/27)

Updated:

- 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

Version 1.4 (2020/04/17)

Updated:

- SPECcpu2017
  Measured with Celeron G4930, calculated with Pentium Gold G5420, Core i3-9100 and Intel® Xeon® Processor E-2200 Product Family
- STREAM
  Measured with Celeron G4930, calculated with Pentium Gold G5420, Core i3-9100 and Intel® Xeon® Processor E-2200 Product Family

Technical data

PRIMERGY TX1320 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 TX1320 M4
Model versions	PY TX1320 M4 LFF PY TX1320 M4 SFF
Form factor	Tower server
Chipset	Intel® C246
Number of sockets	1
Number of processors orderable	1
Processor type	Intel® Celeron® G4900 Intel® Celeron® G4930 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	PCI-Express 3.0 x 8 × 2 PCI-Express 3.0 x 4 × 1 PCI-Express 3.0 x 1 × 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							
Celeron G4900	2	2	2	3.1		2400	54
Pentium Gold G5400	2	4	4	3.7		2400	58
Core i3-8100	4	4	6	3.6		2400	65
Xeon E-2124	4	4	8	3.3	4.3	2666	71
Xeon E-2124G	4	4	8	3.4	4.5	2666	71
Xeon E-2126G	6	6	12	3.3	4.5	2666	80
Xeon E-2134	4	8	8	3.5	4.5	2666	71
Xeon E-2136	6	12	12	3.3	4.5	2666	80
Xeon E-2144G	4	8	8	3.6	4.5	2666	71
Xeon E-2146G	6	12	12	3.5	4.5	2666	80
Xeon E-2174G	4	8	8	3.8	4.7	2666	71
Xeon E-2176G	6	12	12	3.7	4.7	2666	80
Xeon E-2186G	6	12	12	3.8	4.7	2666	95
November 2019 released							
Celeron G4930	2	2	2	3.2		2400	54
Pentium Gold G5420	2	4	4	3.8		2400	54
Core i3-9100	4	4	6	3.6	4.2	2400	65
Xeon E-2224	4	4	8	3.4	4.6	2666	71
Xeon E-2224G	4	4	8	3.5	4.7	2666	71
Xeon E-2226G	6	6	12	3.4	4.7	2666	80
Xeon E-2234	4	8	8	3.6	4.8	2666	71
Xeon E-2236	6	12	12	3.4	4.8	2666	80
Xeon E-2244G	4	8	8	3.8	4.8	2666	71
Xeon E-2246G	6	12	12	3.6	4.8	2666	80
Xeon E-2274G	4	8	8	4.0	4.9	2666	83
Xeon E-2276G	6	12	12	3.8	4.9	2666	80
Xeon E-2278G	8	16	16	3.4	5.0	2666	80
Xeon E-2286G	6	12	12	4.0	4.9	2666	95
Xeon E-2288G	8	16	16	3.7	5.0	2666	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 TX1320 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	2666				✓
8 GB (1 x 8 GB) 1Rx8 DDR4-2666 U ECC	8	1	8	2666				✓
16 GB (1 x 16 GB) 2Rx8 DDR4-2666 U ECC	16	2	8	2666				✓
32 GB (1 x 32 GB) 2Rx8 DDR4-2666 U ECC	32	2	8	2666				✓

Power supplies (since system release)	Max. number
Standard PSU 250W	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 TX1320 M4.

SPECcpu2017

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 result.

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)	Throughput	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)	Throughput	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 TX1320 M4
Processor	Intel® Celeron® G4900 Intel® Celeron® G4930 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 × 16GB (1 x 16GB) 2Rx8 DDR4-2666 U ECC
Software	
BIOS settings	Celeron® G4900: SPECrate2017_int: Fan Control = Full SPECrate2017_fp Fan Control = Full Xeon E-2186G: 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 Hyper-Threading = Disabled Software Guard Extensions (SGX) = Disabled 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

BIOS settings (cont'd)	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 E-2186G 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 E-2186G 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

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

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				
Celeron G4900	*5.69 est.*		**8.73**	**9.37**
Pentium Gold G5400	*7.06 est.*		*12.7 est.*	*13.6 est.*
Core i3-8100	*7.97 est.*		*22.0 est.*	*23.4 est.*
Xeon E-2124	*9.39 est.*		*22.8 est.*	*24.2 est.*
Xeon E-2124G	*9.78 est.*		*23.8 est.*	*25.2 est.*
Xeon E-2126G	*10.1 est.*		*34.7 est.*	*36.9 est.*
Xeon E-2134	*10.0 est.*		*29.7 est.*	*31.8 est.*
Xeon E-2136	*10.3 est.*		*43.0 est.*	*46.1 est.*
Xeon E-2144G	*10.1 est.*		*29.9 est.*	*32.0 est.*
Xeon E-2146G	*10.3 est.*		*41.6 est.*	*44.6 est.*
Xeon E-2174G	*10.4 est.*		*30.4 est.*	*32.5 est.*
Xeon E-2176G	*10.7 est.*		*43.1 est.*	*46.2 est.*
Xeon E-2186G	**10.7**		**43.5**	**46.8**
November 2019 released				
Celeron G4930	**5.81**	**6.10**	**8.87**	**9.51**
Pentium Gold G5420	*7.29 est.*	*7.57 est.*	*12.9 est.*	*13.9 est.*
Core i3-9100	*9.76 est.*	*9.88 est.*	*26.1 est.*	*27.2 est.*
Xeon E-2224	*10.7 est.*	*10.9 est.*	*28.5 est.*	*29.1 est.*
Xeon E-2224G	*10.9 est.*	*11.1 est.*	*29.1 est.*	*30.1 est.*
Xeon E-2226G	*11.3 est.*	*11.5 est.*	*40.7 est.*	*42.4 est.*
Xeon E-2234	*11.4 est.*	*11.6 est.*	*35.8 est.*	*36.6 est.*
Xeon E-2236	*11.7 est.*	*11.9 est.*	*48.1 est.*	*48.1 est.*
Xeon E-2244G	*11.4 est.*	*11.6 est.*	*35.1 est.*	*36.9 est.*
Xeon E-2246G	*11.7 est.*	*11.9 est.*	*46.2 est.*	*48.9 est.*
Xeon E-2274G	*11.5 est.*	*11.8 est.*	*34.6 est.*	*36.4 est.*
Xeon E-2276G	*11.9 est.*	*12.0 est.*	*46.5 est.*	*49.1 est.*
Xeon E-2278G	*12.2 est.*	*12.4 est.*	*57.6 est.*	*61.0 est.*
Xeon E-2286G	*11.9 est.*	*12.0 est.*	*47.9 est.*	*50.5 est.*
Xeon E-2288G	**12.1**	*12.4 est.*	**60.5**	*63.9 est.*

Processor	SPECspeed2017 int_energy_base	SPECrate2017 int_energy_base
Xeon E-2288G	**219**	**435**

On November 13, 2018, the PRIMERGY TX1320 M4 with a Xeon E-2186G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECrate2017 int_base.

On November 13, 2018, the PRIMERGY TX1320 M4 with a Xeon E-2186G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECrate2017 int_peak.

On November 13, 2018, the PRIMERGY TX1320 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 TX1320 M4 with a Xeon E-2288G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECspeed2017 int_energy_base.

On November 1, 2019, the PRIMERGY TX1320 M4 with a Xeon E-2288G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECrate2017 int_energy_base.

Processor	SPECspeed2017 fp_base	SPECspeed2017 fp_peak	SPECrate2017 fp_base	SPECrate2017 fp_peak
November 2018 released				
Celeron G4900	*11.7 est.*		**11.6**	**11.8**
Pentium Gold G5400	*13.5 est.*		*15.5 est.*	*15.6 est.*
Core i3-8100	*23.2 est.*		*26.5 est.*	*26.9 est.*
Xeon E-2124	*26.0 est.*		*29.7 est.*	*30.1 est.*
Xeon E-2124G	*26.6 est.*		*30.6 est.*	*31.1 est.*
Xeon E-2126G	*31.5 est.*		*37.6 est.*	*38.2 est.*
Xeon E-2134	*26.8 est.*		*30.7 est.*	*31.2 est.*
Xeon E-2136	*31.9 est.*		*38.2 est.*	*38.8 est.*
Xeon E-2144G	*26.7 est.*		*30.9 est.*	*31.4 est.*
Xeon E-2146G	*31.7 est.*		*37.3 est.*	*38.0 est.*
Xeon E-2174G	*27.2 est.*		*31.2 est.*	*31.7 est.*
Xeon E-2176G	*32.1 est.*		*38.4 est.*	*39.1 est.*
Xeon E-2186G	**32.1**		**38.7**	**39.3**
November 2019 released				
Celeron G4930	**11.9**	**12.1**	**11.8**	**12.1**
Pentium Gold G5420	*13.8 est.*	*14.1 est.*	*15.7 est.*	*16.0 est.*
Core I3-9100	*24.3 est.*	*24.8 est.*	*27.8 est.*	*28.5 est.*
Xeon E-2224	*26.7 est.*	*27.4 est.*	*30.7 est.*	*31.4 est.*
Xeon E-2224G	*27.1 est.*	*27.8 est.*	*31.4 est.*	*32.2 est.*
Xeon E-2226G	*32.0 est.*	*32.7 est.*	*37.5 est.*	*38.5 est.*
Xeon E-2234	*27.4 est.*	*28.0 est.*	*31.5 est.*	*32.3 est.*
Xeon E-2236	*32.0 est.*	*32.7 est.*	*37.7 est.*	*38.7 est.*
Xeon E-2244G	*27.5 est.*	*28.2 est.*	*31.8 est.*	*32.5 est.*
Xeon E-2246G	*32.2 est.*	*32.9 est.*	*37.9 est.*	*38.9 est.*
Xeon E-2274G	*27.3 est.*	*27.9 est.*	*31.5 est.*	*32.1 est.*
Xeon E-2276G	*32.8 est.*	*33.5 est.*	*38.6 est.*	*39.5 est.*
Xeon E-2278G	*36.0 est.*	*36.8 est.*	*43.2 est.*	*44.3 est.*
Xeon E-2286G	*33.0 est.*	*33.6 est.*	*39.1 est.*	*40.0 est.*
Xeon E-2288G	**36.5**	*37.3 est.*	**44.6**	*45.7 est.*

Processor	SPECspeed2017 fp_energy_base	SPECrate2017 fp_energy_base
Xeon E-2288G	**312**	**363**

On November 13, 2018, the PRIMERGY TX1320 M4 with a Xeon E-2186G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECrate2017 fp_base.

On November 13, 2018, the PRIMERGY TX1320 M4 with a Xeon E-2186G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECrate2017 fp_peak.

On November 13, Nov 2018, the PRIMERGY TX1320 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 TX1320 M4 with a Xeon E-2288G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECspeed2017 fp_energy_base.

On November 1, 2019, the PRIMERGY TX1320 M4 with a Xeon E-2288G processor was ranked first in the 1-socket Xeon E server systems category for the benchmark SPECrate2017 fp_energy_base.

The following two diagrams illustrate the comparison of SPECrate2017 on Xeon E-2186G and Xeon E-2288G measured with PRIMERGY TX1320 M4, in their respective most performance configurations.

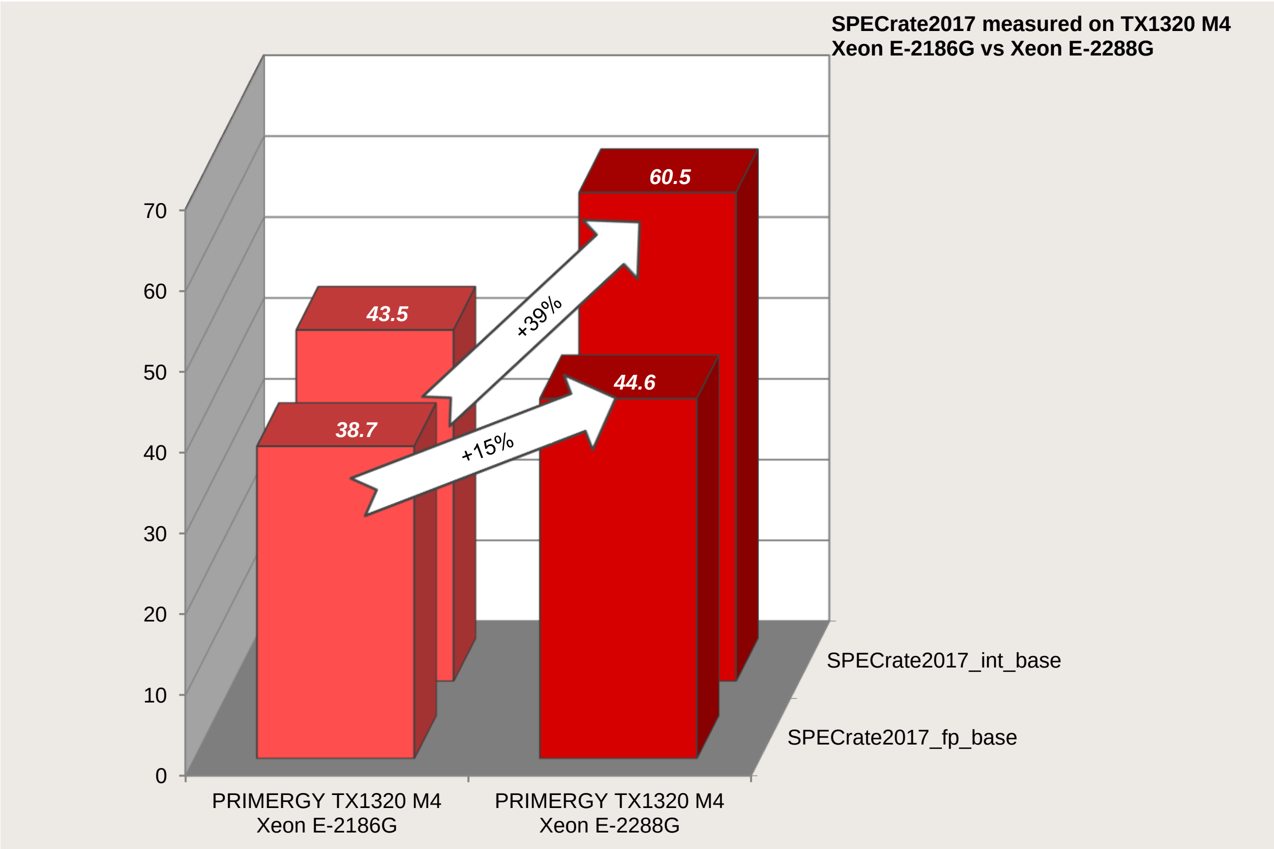

SPECpower_ssj2008

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 as they have 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 the 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.

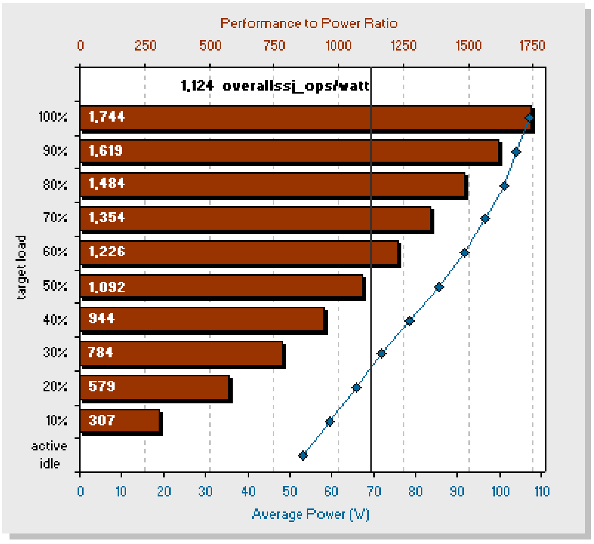

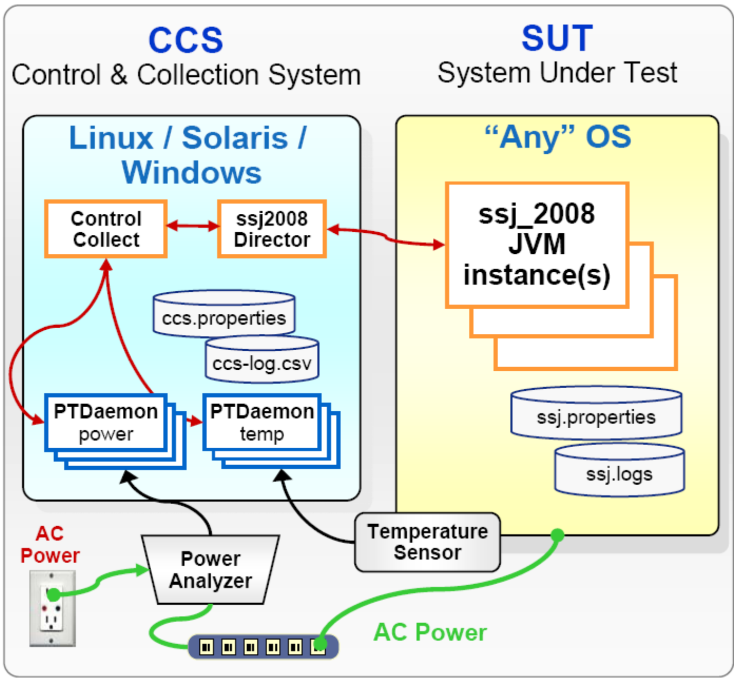

The benchmark runs on a wide variety of operating systems and hardware architectures, and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System 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 Windows OS measurement	
Hardware	
Model	PRIMERGY TX1320 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 250 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

For Linux OS measurement	
Hardware	
Model	PRIMERGY TX1320 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 250 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	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 rcu_nocbs=1-111 nohz_full=1-111 isolcpus=1-111 modprobe cpufreq_conservative cpupower frequency-set --governor conservative echo -n 98 > /sys/devices/system/cpu/cpufreq/conservative/up_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 97 > /sys/devices/system/cpu/cpufreq/conservative/down_threshold 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 powertop --auto-tune echo 0 > /proc/sys/kernel/nmi_watchdog cpupower frequency-set -u 2500MHz 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 -Xmn1700m -Xms1950m -Xmx1950m -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

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

Benchmark results(Windows)

The PRIMERGY TX1320 M4 in Microsoft Windows Server 2016 Standard achieved the following result:

SPECpower_ssj2008 = 9,435 overall ssj_ops/watt

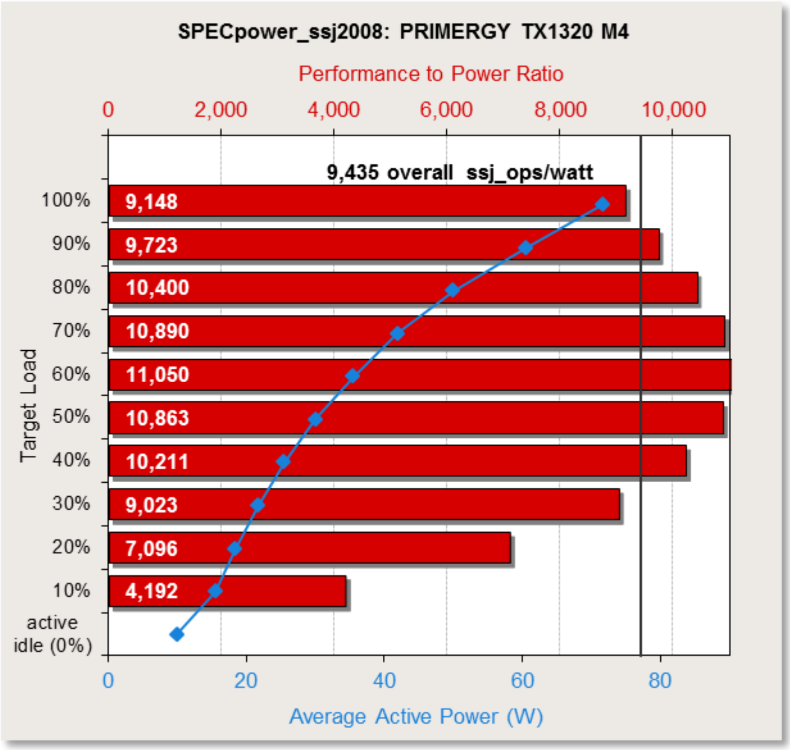

The adjoining diagram shows the results for the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (top x-axis) for each target load level on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhombus. The black vertical line shows the benchmark result of 9,435 overall ssj_ops/watt for the PRIMERGY TX1320 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%	655,699	71.7	9,148
90%	589,566	60.6	9,723
80%	522,263	50.0	10,400
70%	455,895	41.9	10,890
60%	391,846	35.5	11,050
50%	325,677	30.0	10,863
40%	260,292	25.5	10,211
30%	196,623	21.8	9,023
20%	130,916	18.4	7,096
10%	65,014	15.5	4,192
Active Idle	0	10.0	0
			∑ssj_ops / ∑power = 9,435

Benchmark results(Linux)

The PRIMERGY TX1320 M4 in SUSE Linux Enterprise Server 12 SP4 achieved the following result:

SPECpower_ssj2008 = 11,089 overall ssj_ops/watt

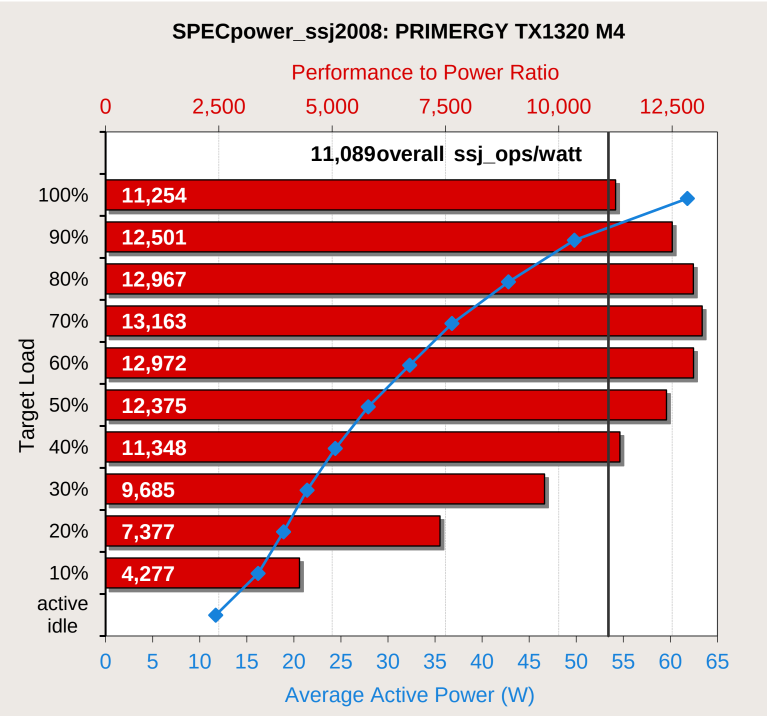

The adjoining diagram shows the results for the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (top x-axis) for each target load level on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhombus. The black vertical line shows the benchmark result of 11,089 overall ssj_ops/watt for the PRIMERGY TX1320 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%	695,687	61.8	11,254
90%	622,632	49.8	12,501
80%	555,470	42.8	12,967
70%	484,278	36.8	13,163
60%	418,655	32.3	12,972
50%	345,565	27.9	12,375
40%	277,162	24.4	11,348
30%	207,569	21.4	9,685
20%	139,347	18.9	7,377
10%	69,264	16.2	4,277
Active Idle	0	11.7	0
			∑ssj_ops / ∑power = 11,089

The PRIMERGY TX1320 M4 achieved a new class record with this result (date: June 5, 2019). Thus, the PRIMERGY TX1320 M4 proves itself to be the most energy-efficient 1-socket Xeon E-2100 server in the world. For the latest SPECpower_ssj2008 benchmark results, visit: http://www.spec.org/power_ssj2008/results.

The following diagram shows for each load level the power consumption (on the right y-axis) and the throughput (on the left y-axis) of the PRIMERGY TX1320 M4 compared to the predecessor PRIMERGY TX1320 M3.

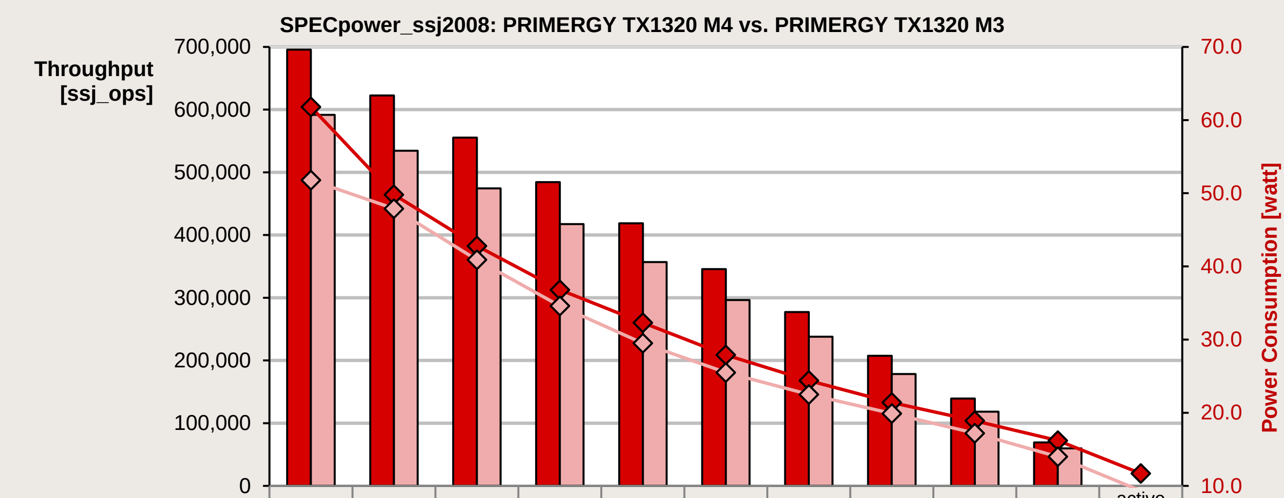

Load Level	100%	90%	80%	70%	60%	50%	40%	30%	20%	10%	active idle
TX1320 M4 ssj_ops	695,687	622,632	555,470	484,278	418,655	345,565	277,162	207,569	139,347	69,264	0
TX1320 M3 ssj_ops	591,742	534,498	474,350	417,545	356,872	296,352	237,875	178,292	118,311	59,729	0
TX1320 M4 watt	61.8	49.8	42.8	36.8	32.3	27.9	24.4	21.4	18.9	16.2	11.7
TX1320 M3 watt	51.8	47.9	40.9	34.6	29.5	25.5	22.5	19.9	17.2	14.0	9.3

Thanks to the latest Xeon E processors, the PRIMERGY TX1320 M4 has a higher throughput. This results in an overall 6.35% increase in energy efficiency in the PRIMERGY TX1320 M4.

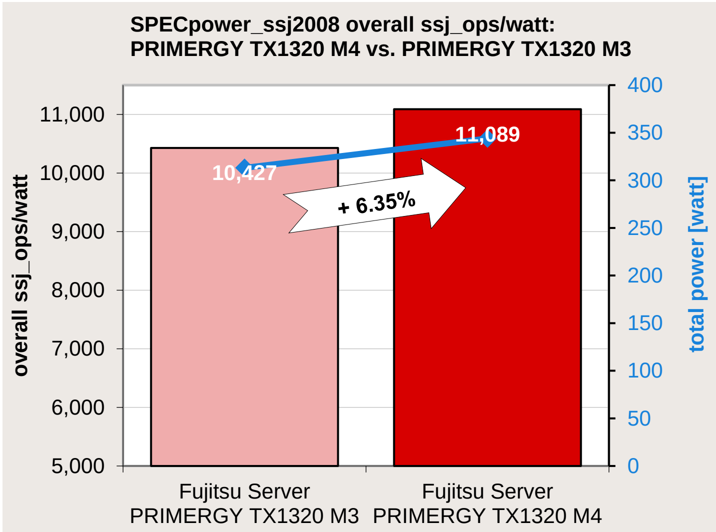

Difference of score by OS & JVM version

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.

Benchmark environment (E-2200 Product Family)

System Under Test (SUT)	
Hardware	
Model	PRIMERGY TX1320 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 ×S26113-F575-E13
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 TX1320 M4 in SUSE Linux Enterprise Server 12 SP4 achieved the following result:

SPECpower_ssj2008 = 12,364 overall ssj_ops/watt

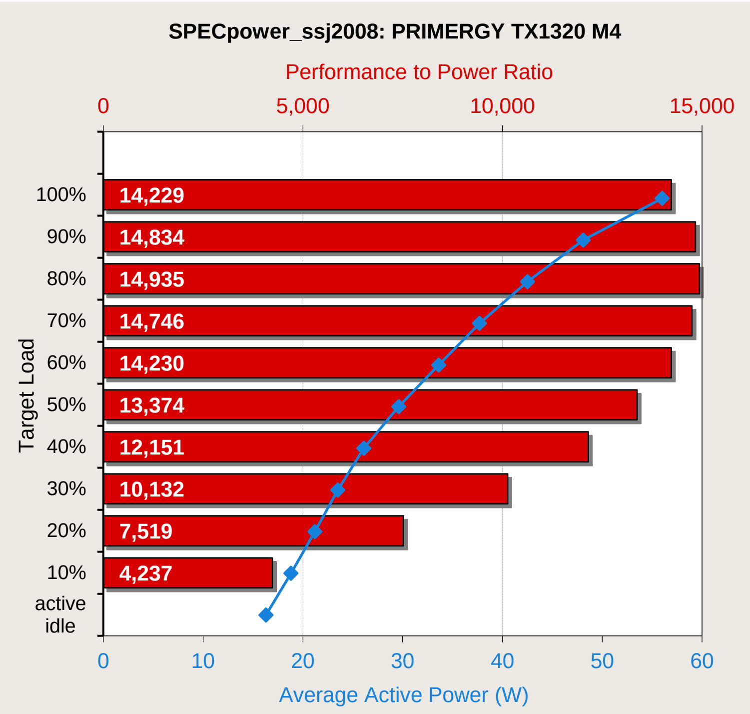

The adjoining diagram shows the results for the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (top x-axis) for each target load level on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhombus. The black vertical line shows the benchmark result of 12,364 overall ssj_ops/watt for the PRIMERGY TX1320 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%	797,225	56.0	14,229
90%	712,828	48.1	14,834
80%	635,321	42.5	14,935
70%	555,428	37.7	14,746
60%	477,700	33.6	14,230
50%	395,875	29.6	13,374
40%	317,519	26.1	12,151
30%	237,991	23.5	10,132
20%	159,095	21.2	7,519
10%	79,687	18.8	4,237
Active Idle	0	16.3	0
			∑ssj_ops / ∑power = 12,364

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

The following diagram shows for each load level the power consumption (on the right y-axis) and the throughput (on the left y-axis) of the E-2288G compared to the predecessor E-2176G on PRIMERGY TX1320 M4.

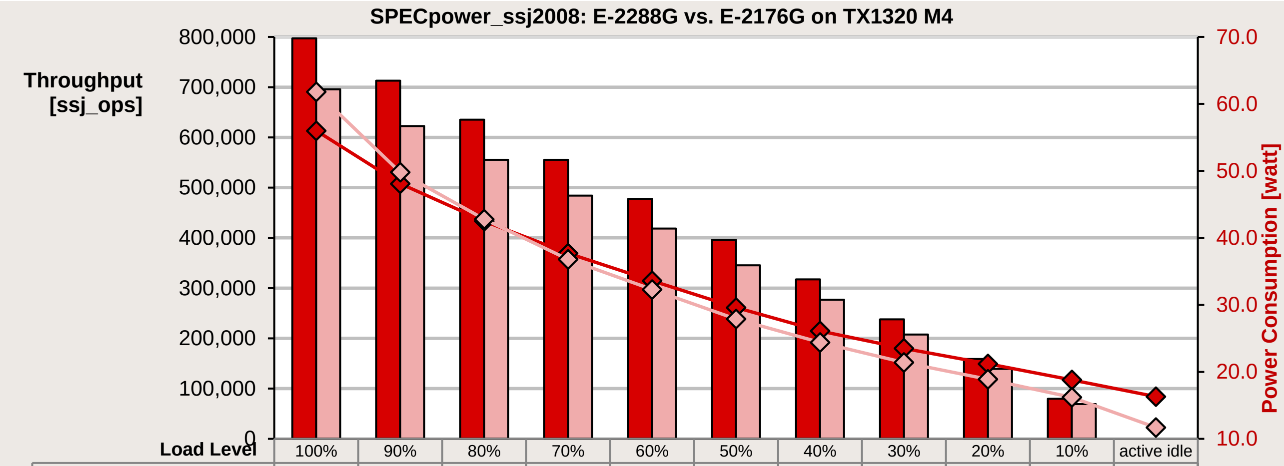

Load Level	100%	90%	80%	70%	60%	50%	40%	30%	20%	10%	active idle
Xeon E-2288G ssj_ops	797,225	712,828	635,321	555,428	477,700	395,875	317,519	237,991	159,095	79,687	0
Xeon E-2176G ssj_ops	695,687	622,632	555,470	484,278	418,655	345,565	277,162	207,569	139,347	69,264	0
Xeon E-2288G watt	56	48	43	38	34	30	26	24	21.2	18.8	16.3
Xeon E-2176G watt	62	50	43	37	32	28	24	21	19	16.2	11.7

Thanks to the latest Xeon E processors, the PRIMERGY TX1320 M4 has a higher throughput. This results in an overall 7.3% increase in energy efficiency in the E-2288G.

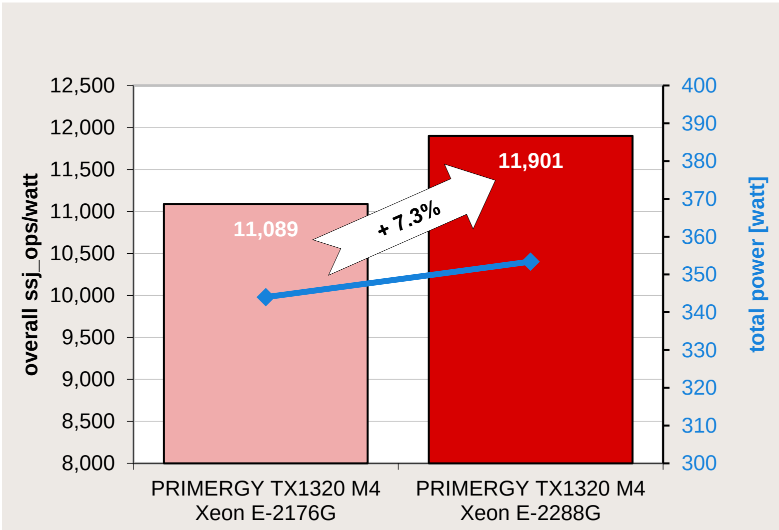

SPECjbb2015

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 the replenish inventory
- Data mining operations to identify sale patterns and generate quarterly business reports

The SPECjbb2015 benchmark has a two performance 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 constraint 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 constraint.

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

Results can not be compared between categories.

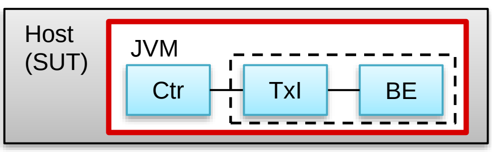

(a) Example of SPECjbb2015 Composite configuration

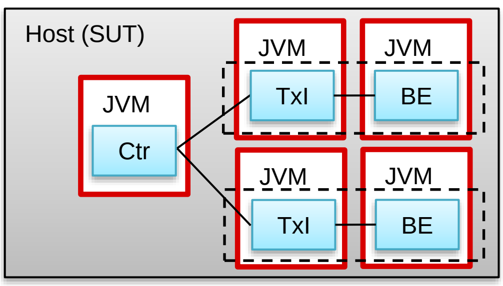

(b) Example of SPECjbb2015 MultiJVM configuration

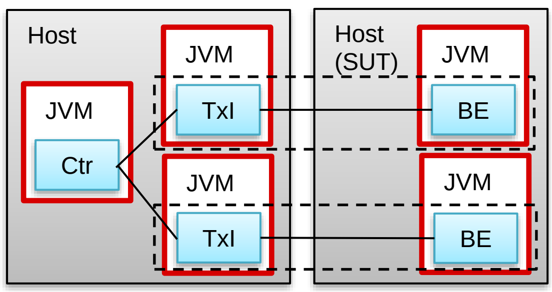

(c) Example of SPECjbb2015 Distributed configuration

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 specification of the benchmark can be found at https://www.spec.org/jbb2015/ .

Benchmark environment

PRIMERGY TX1320 M4 was configured for the SPECjbb2015 Composite benchmark measurement.

System Under Test (SUT)	
Hardware	
Model	PRIMERGY TX1320 M4
Processor	1 × Intel® Xeon® E-2186G
Memory	4 × 16 GB (1x16 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
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 TX1320 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 TX1320 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 used to assess their performance and enable a comparison of the different storage connections for PRIMERGY servers. In the standard implementation, these performance measurements are carried out using a defined measurement method which models the accesses in real application scenarios on the basis of specifications.

The essential specifications are:

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

A given combination of values for these specifications is known as a “load profile”. The following five standard load profiles can be allocated to typical application scenarios:

Standard load profile	Access	Type of access		Block size [kB]	Application
		Read	Write		
File copy	Random	50%	50%	64	Copying files
File server	Random	67%	33%	64	File server
Database	Random	67%	33%	8	Database (data transfer), mail server
Streaming	Sequential	100%	0%	64	Database (log file), data backup, video streaming (partial)
Restore	Sequential	0%	100%	64	Restoring files

In order to model applications that access in parallel with a different load intensity the "# of outstanding I/Os" is increased from 1 to 512 (in steps of powers of two).

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

The main results for a measurement are:

- Throughput [MB/s]   Throughput in megabytes per second
- Transactions [IO/s]   Transaction rate in I/O operations per second
- Latency [ms]   Average response time in ms

The data throughput has established itself as the normal measurement variable for sequential load profiles, whereas the measurement variable “transaction rate” is mostly used for random load profiles with their small block sizes. Data throughput and transaction rate are directly proportional to each other and can be converted between each other using the formulas:

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

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

All the details of the measurement method and the basics of disk I/O performance are described in the white paper “Basics of Disk I/O Performance”.

Benchmark environment

All the measurement results discussed in this section apply with regard to 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	BC-SATA HDD (SATA 6 Gbps, 7.2 krpm) [512e]	ST6000NM0115 *1 *3
		BC-SATA HDD (SATA 6 Gbps, 7.2 krpm) [512n]	HUS722T1TALA604 *2 *3 ST2000NM0055 *1 *3
		SATA HDD (SATA 6 Gbps, 7.2 krpm) [512e]	ST1000DM003-1SB *1 *3
Controller: Intel(R) 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, 10 krpm) [512e]	AL15SEB12EQ *2 *3
		SAS HDD (SAS 12 Gbps, 10 krpm) [512n]	AL15SEB030N *2 *3
		SAS HDD (SAS 12 Gbps, 15 krpm) [512n]	ST300MP0006 *1 *3
		BC-SATA HDD (SATA 6 Gbps, 7.2 krpm) [512e]	ST1000NX0313 *1 *3
		BC-SATA HDD (SATA 6 Gbps, 7.2 krpm) [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 kB)
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 GB 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 most appropriate storage media in terms of disk-I/O performance. For this purpose, a single storage medium was measured using the configuration specified in the subsection "Benchmark environment".

Controller

The measurements were made using the controllers in the table below.

Storage medium	Storage medium	Cache	Supported interfaces		RAID levels
			Host	Drive	
SSD/HDD	PRAID CP400i	-	PCIe 3.0 x8	SATA 6G SAS 12G	0, 1, 1E, 10, 5, 50
M.2 Flash	C620 Standard SATA AHCI controller	-	DMI 3.0 x4	SATA 6G	-

Storage media

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

Model type	Storage medium type	Interface	Form factor
3.5 inch model	HDD	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

HDDs and SSDs are operated via host bus adapters, usually RAID controllers, with a SATA or SAS interface. The interface for the RAID controller to the chipset of the system board is typically PCIe or, in the case of 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. The trade-off, however, is that the price per gigabyte of storage capacity is substantially higher.

Cache settings

In most cases, the cache in HDDs has a significant influence on disk-I/O performance. It is frequently regarded as a security problem in the event of a power failure, and is therefore switched off. On the other hand, HDD manufacturers integrate it into their hard disks for a good reason: it increases the write performance. For performance reasons, it is therefore advisable to enable the hard disk cache. To prevent data loss in the event of a power failure, equipping the system with a UPS is recommended.

To be able to handle the settings for RAID controllers and hard disks easily and reliably, it is advisable to use the RAID-Manager software "ServerView RAID" that is supplied with PRIMERGY servers. All the cache settings for controllers and hard disks can usually be made in one go – specifically for the application – by using the pre-defined modes "Performance" or "Data Protection". The "Performance" mode ensures you will get the best possible performance settings for the majority of 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 already mentioned in the subsection "Benchmark description", are used here. Thus, transaction rate is specified for random accesses and data throughput for sequential accesses. To avoid any confusion among the measurement units, the two access types have been given separate tables.

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

3.5 inch model storage media

HDDs

Random accesses (units: IO/s):

Capacity [GB]	Storage device	Interface	Transactions [IO/s]		
			Database	Fileserver	filecopy
6,000	ST6000NM0115	SATA 6G	392	362	371
1,000	HUS722T1TALA604	SATA 6G	287	264	269
2,000	ST2000NM0055	SATA 6G	339	301	314
1,000	ST1000DM003-1SB	SATA 6G	222	210	204

Sequential accesses (units: MB/s):

Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
6,000	ST6000NM0115	SATA 6G	213	208
1,000	HUS722T1TALA604	SATA 6G	201	201
2,000	ST2000NM0055	SATA 6G	196	195
1,000	ST1000DM003-1SB	SATA 6G	208	203

SSDs

Random accesses (units: IO/s):

Capacity [GB]	Storage device	Interface	Transactions [IO/s]		
			Database	Fileserver	filecopy
240	MTFDDAV240TCB	SATA 6G	20,113	3,936	5,021
480	MTFDDAV480TCB	SATA 6G	22,596	4,993	6,331

Sequential accesses (units: MB/s):

Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
240	MTFDDAV240TCB	SATA 6G	510	271
480	MTFDDAV480TCB	SATA 6G	509	403

2.5 inch model storage media

HDDs

Random accesses (units: IO/s):

Capacity [GB]	Storage device	Interface	Transactions [IO/s] Database	Fileserver	filecopy
1,200	AL15SEB12EQ	SAS 12G	594	520	546
300	AL15SEB030N	SAS 12G	645	546	568
300	ST300MP0006	SAS 12G	768	662	472
1,000	ST1000NX0313	SATA 6G	324	281	288
2,000	ST2000NX0403	SATA 6G	326	286	294

Sequential accesses (units: MB/s):

Capacity [GB]	Storage device	Interface	Throughput [MB/s] Streaming	Restore
1,200	AL15SEB12EQ	SAS 12G	260	259
300	AL15SEB030N	SAS 12G	231	230
300	ST300MP0006	SAS 12G	304	304
1,000	ST1000NX0313	SATA 6G	131	131
2,000	ST2000NX0403	SATA 6G	133	133

SSDs

Random accesses (units: IO/s):

Capacity [GB]	Storage device	Interface	Transactions [IO/s] Database	Fileserver	filecopy
240	MZ7KH240HAHQ	SATA 6G	49,159	7,313	7,431
480	MZ7KH480HAHQ	SATA 6G	50,558	7,774	7,810
960	MZ7KH960HAJR	SATA 6G	50,647	7,793	7,916
1,920	MZ7KH1T9HAJR	SATA 6G	50,702	8,040	7,960
3,840	MZ7KH3T8HALS	SATA 6G	50,766	8,039	7,936
240	MTFDDAK240TCB	SATA 6G	18,959	3,367	4,516
480	MTFDDAK480TDC	SATA 6G	24,710	3,799	5,006
960	MTFDDAK960TDC	SATA 6G	30,152	4,625	5,553
1,920	MTFDDAK1T9TDC	SATA 6G	37,234	5,606	5,566
3,840	MTFDDAK3T8TDC	SATA 6G	41,711	6,429	6,133
7,680	MTFDDAK7T6TDC	SATA 6G	40,683	6,874	6,672
240	MTFDDAV240TCB	SATA 6G	20,113	3,936	5,021
480	MTFDDAV480TCB	SATA 6G	22,596	4,993	6,331

Sequential accesses (units: MB/s):

Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
240	MZ7KH240HAHQ	SATA 6G	526	486
480	MZ7KH480HAHQ	SATA 6G	526	485
960	MZ7KH960HAJR	SATA 6G	525	485
1,920	MZ7KH1T9HAJR	SATA 6G	526	485
3,840	MZ7KH3T8HALS	SATA 6G	526	485
240	MTFDDAK240TCB	SATA 6G	487	258
480	MTFDDAK480TDC	SATA 6G	507	362
960	MTFDDAK960TDC	SATA 6G	507	440
1,920	MTFDDAK1T9TDC	SATA 6G	507	483
3,840	MTFDDAK3T8TDC	SATA 6G	504	481
7,680	MTFDDAK7T6TDC	SATA 6G	469	482
240	MTFDDAV240TCB	SATA 6G	510	271
480	MTFDDAV480TCB	SATA 6G	509	403

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 field 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 for 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 in the memory, because the processor caches are used for sequential access.

Before execution, the source code is adjusted to suit 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 they will have as little influence as possible on the result. The OpenMP program library is used to enable selected parts of the program to be executed in parallel during the runtime of the benchmark, consequently achieving optimal load distribution to the available processor cores.

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

Type	Execution	Bytes per step	Floating-point calculations 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 for comparison.

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

This chapter specifies throughputs in base 10 (1 GB/s = 10^9 Byte/s).

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY TX1320 M4
Processor	Intel® Celeron® G4900 Intel® Celeron® G4930 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 [MHz]	Max. Memory Bandwidth [GB/s]	Cores	Processor Frequency [GHz]	TRIAD [GB/s]
November 2018 released					
Celeron G4900	2400	37.5	2	3.1	**33.6**
Pentium Gold G5400	2400	37.5	2	3.7	*33.8 est.*
Core i3-8100	2400	37.5	4	3.6	*33.0 est.*
Xeon E-2124	2666	41.6	4	3.3	*36.7 est.*
Xeon E-2124G	2666	41.6	4	3.4	*37.0 est.*
Xeon E-2126G	2666	41.6	6	3.3	*36.2 est.*
Xeon E-2134	2666	41.6	4	3.5	*36.7 est.*
Xeon E-2136	2666	41.6	4	3.3	*35.8 est.*
Xeon E-2144G	2666	41.6	4	3.6	*36.6 est.*
Xeon E-2146G	2666	41.6	6	3.5	*35.9 est.*
Xeon E-2174G	2666	41.6	4	3.8	*36.9 est.*
Xeon E-2176G	2666	41.6	6	3.7	*36.2 est.*
Xeon E-2186G	2666	41.6	6	3.8	*35.9 est.*
November 2019 released					
Celeron G4930	2400	37.5	2	3.2	**33.8**
Pentium Gold G5420	2400	37.5	2	3.8	*34.2 est.*
Core i3-9100	2400	37.5	4	3.6	*33.2 est.*
Xeon E-2224	2666	41.6	4	3.4	*36.8 est.*
Xeon E-2224G	2666	41.6	4	3.5	*36.6 est.*
Xeon E-2226G	2666	41.6	6	3.4	*35.4 est.*
Xeon E-2234	2666	41.6	4	3.6	*36.8 est.*
Xeon E-2236	2666	41.6	6	3.4	*35.4 est.*
Xeon E-2244G	2666	41.6	4	3.8	*36.8 est.*
Xeon E-2246G	2666	41.6	6	3.6	*35.4 est.*
Xeon E-2274G	2666	41.6	4	4.0	*36.8 est.*
Xeon E-2276G	2666	41.6	6	3.8	*35.5 est.*
Xeon E-2278G	2666	41.6	8	3.4	*34.8 est.*
Xeon E-2286G	2666	41.6	6	4.0	*35.4 est.*
Xeon E-2288G	2666	41.6	8	3.7	*34.9 est.*

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 TX1320 M4.

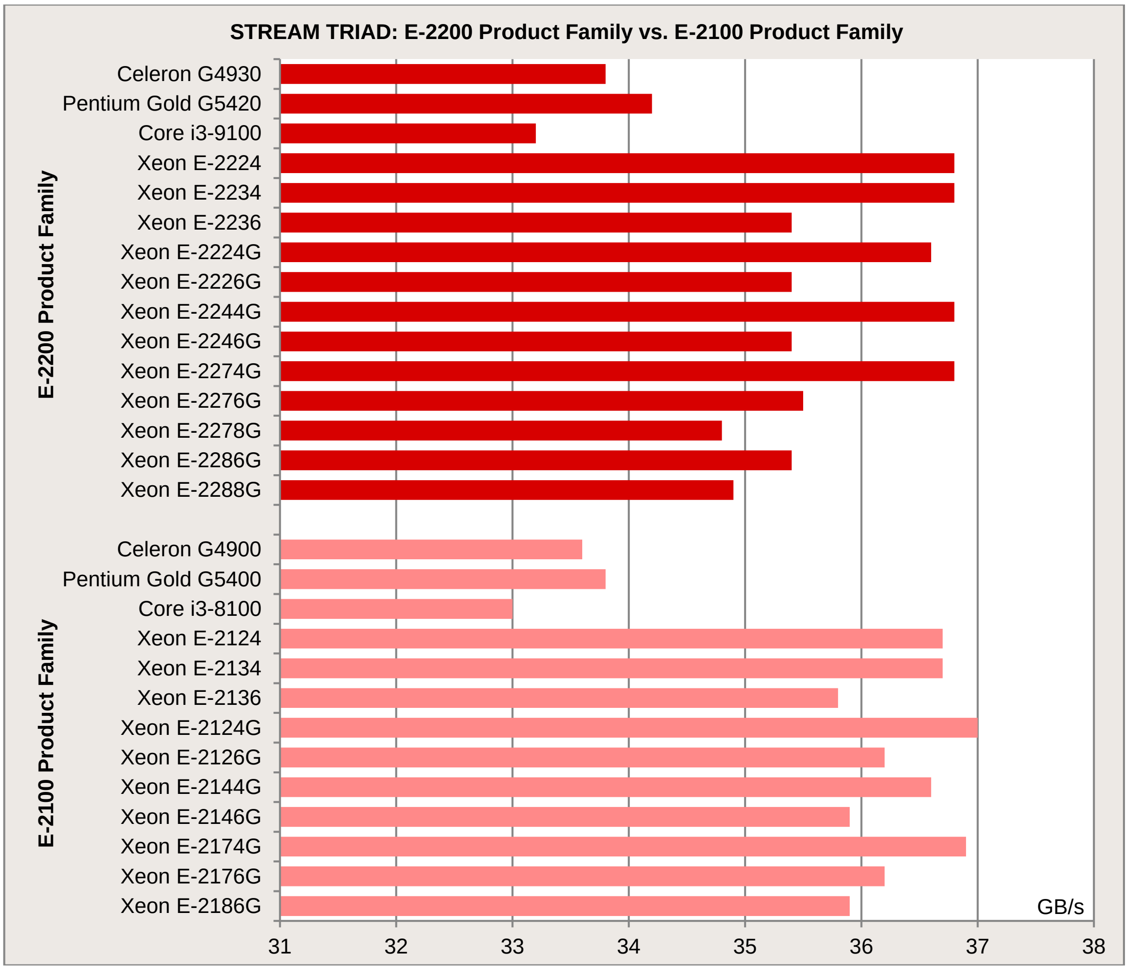

Literature

PRIMERGY Servers

http://primergy.com/

PRIMERGY TX1320 M4

This White Paper:

http://docs.ts.fujitsu.com/dl.aspx?id=54b0597c-ad5d-4ee1-bbb2-39b8a7afebf6

http://docs.ts.fujitsu.com/dl.aspx?id=4e7e9010-a606-417d-9e0d-5ac25ba1cf59

Data sheet

http://docs.ts.fujitsu.com/dl.aspx?id=bc629a48-cc8a-499a-927f-4884a3def54f

PRIMERGY Performance

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

SPECcpu2017

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

Benchmark Overview SPECcpu2017

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

SPECjbb2015

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

STREAM

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

Contact

FUJITSU

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

PRIMERGY Product Marketing

mailto:Primergy-PM@ts.fujitsu.com

PRIMERGY Performance and Benchmarks

mailto:primergy.benchmark@ts.fujitsu.com

2020/04/17 WW EN