Fujitsu Server PRIMERGY
Performance Report
PRIMERGY CX2550 M6/ CX2560 M6

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMERGY CX2550 M6/ CX2560 M6.

Explaines PRIMERGY CX2550 M6/ CX2560 M6 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
1.2
2023-10-03

Contents

Technical data

PRIMERGY CX400 M6 シャーシ

PRIMERGY CX2550 M6

PRIMERGY CX2560 M6

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

Model	PRIMERGY CX2550 M6	PRIMERGY CX2560 M6
Cooling method	Air cooling/ Liquid cooling	Air cooling
Form factor	Server node	
Chipset	Intel C621A	
Number of sockets	2	
Number of processors orderable	1 or 2	
Processor type	3rd Generation Intel Xeon Scalable Processors Family	
Number of memory slots	16 (8 per processor)	24 (12 per processor)
Maximum memory configuration	2,048 GB	3,072 GB
Storage controller	2 x onboard SATA controller	
SATA interface (onboard)	2 x SATA port 2 x M.2 SSD	6 x SATA port 2 x M.2 SSD
PCI slots	PCI-Express 4.0 (x16 lane): 2 (Low Profile)	

Processor								
Processor model	Number of cores	Number of threads	Cache [MB]	UPI speed [GT/s]	Rated frequency [GHz]	Maximum turbo frequency [GHz]	Maximum memory frequency [MHz]	TDP [W]
Xeon Gold 5318Y	24	48	36	11.2	2.1	3.4	2,933	165
Xeon Gold 5318S	24	48	36	11.2	2.1	3.4	2,933	165
Xeon Gold 5317	12	24	18	11.2	3.0	3.6	2,933	150
Xeon Gold 5315Y	8	16	12	11.2	3.2	3.6	2,933	140
Xeon Silver 4316	20	40	30	10.4	2.3	3.4	2,666	150
Xeon Silver 4314	16	32	24	10.4	2.4	3.4	2,666	135
Xeon Silver 4310	12	24	18	10.4	2.1	3.3	2,666	120
Xeon Silver 4309Y	8	16	12	10.4	2.8	3.6	2,666	105

All processors that can be ordered with PRIMERGY CX2550 M6/ CX2560 M6 support Intel Turbo Boost Technology 2.0. This technology allows you to operate the processor with higher frequencies than the rated frequency. The "maximum turbo frequency" listed in the processor list above is the theoretical maximum frequency when there is only one active core per processor. The maximum frequency that can actually be achieved depends on the number of active cores, current consumption, power consumption, and processor temperature.

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

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

The processor with the suffix means it is optimized for the following feature.

Suffix	Additional feature
S	Processor which supports max SGX enclave size of 512GB
Y	Processor which supports Speed Select Technology

Memory modules									
Type	Capacity [GB]	Number of ranks	Bit width of the memory chips	Frequency [MHz]	3DS	Load Reduced	Registered	NVDIMM	ECC
8 GB (1x 8 GB) 1Rx8 DDR4-3200 R ECC	8	1	8	3,200			✓		✓
16 GB (1x 16 GB) 2Rx8 DDR4-3200 R ECC	16	2	8	3,200			✓		✓
16 GB (1x 16 GB) 1Rx4 DDR4-3200 R ECC	16	1	4	3,200			✓		✓
32 GB (1x 32 GB) 2Rx4 DDR4-3200 R ECC	32	2	4	3,200			✓		✓
64 GB (1x 64 GB) 2Rx4 DDR4-3200 R ECC	64	2	4	3,200			✓		✓
64 GB (1x 64 GB) 2Rx4 DDR4-3200 LR ECC	64	2	4	3,200		✓	✓		✓
128 GB (1x128 GB) 4Rx4 DDR4-3200 LR ECC	128	4	4	3,200		✓	✓		✓
256 GB (1x256 GB) 8Rx4 DDR4-3200 3DS ECC	256	8	4	3,200	✓		✓		✓
128 GB (1x128GB) Optane PMem-3200	128			3,200				✓	✓
256 GB (1x256GB) Optane PMem-3200	256			3,200				✓	✓

SPEC CPU2017

Benchmark description

SPEC CPU2017 is a benchmark which measures the system efficiency with integer and floating-point operations. It consists of an integer test suite (SPECrate 2017 Integer, SPECspeed 2017 Integer) containing 10 applications and a floating-point test suite (SPECrate 2017 Floating Point, SPECspeed 2017 Floating Point) containing 14 applications. Both test suites are extremely computing-intensive and concentrate on the CPU and the memory. Other components, such as Disk I/O and network, are not measured by this benchmark.

SPEC CPU2017 is not tied to a special operating system. The benchmark is available as source code and is compiled before the actual measurement. The used compiler version and their optimization settings also affect the measurement result.

SPEC CPU2017 contains two different performance measurement methods. The first method (SPECspeed 2017 Integer or SPECspeed 2017 Floating Point) determines the time which is required to process a single task. The second method (SPECrate 2017 Integer or SPECrate 2017 Floating Point) determines the throughput, i.e. the number of tasks that can be handled in parallel. Both methods are also divided into two measurement runs, "base" and "peak." They differ in the use of compiler optimization. When publishing the results, the base values are always used and the peak values are optional.

Benchmark	Number of single benchmarks	Arithmetics	Type	Compiler optimization	Measurement result
SPECspeed2017_int_peak	10	Integer	peak	aggressive	Speed
SPECspeed2017_int_base	10	Integer	base	conservative	
SPECrate2017_int_peak	10	Integer	peak	aggressive	Throughput
SPECrate2017_int_base	10	Integer	base	conservative	
SPECspeed2017_fp_peak	10	floating point	peak	aggressive	Speed
SPECspeed2017_fp_base	10	floating point	base	conservative	
SPECrate2017_fp_peak	13	floating point	peak	aggressive	Throughput
SPECrate2017_fp_base	13	floating point	base	conservative	

The measurement results are the geometric average from normalized ratio values which have been determined for individual benchmarks. The geometric average - in contrast to the arithmetic average - means that there is a weighting in favor of the lower individual results. "Normalized" means that the measurement is how fast is the test system compared to a reference system. For example, value "1" was defined for the SPECspeed2017_int_base, SPECrate2017_int_base, SPECspeed2017_fp_base, and SPECrate2017_fp_base results of the reference system. A SPECspeed2017_int_base value of 2 means that the measuring system has handled this benchmark twice as fast as the reference system. A SPECrate2017_fp_base value of 4 means that the measuring system has handled this benchmark about 4/[# base copies] times faster than the reference system. "# base copies" specifies how many parallel instances of the benchmark have been executed.

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

Benchmark environment

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY CX2550 M6/ CX2560 M6
• Processor	2 x 3rd Generation Intel Xeon Scalable Processors Family
• Memory	16 x 32 GB 2Rx4 PC4-3200AA-R
Software	
• BIOS settings	SPECspeed2017_int_base: • CPU C1E Support = Disabled SPECSpeed2017_fp_base: • Hyper Threading = Disabled SPECrate2017_int_base: • DCU Streamer Prefetcher = Disabled • UPI Link Frequency Select = 10.4GT/s • LLC Prefetch = Enabled • Sub NUMA (SNC) =Enable SNC2 SPECrate2017_fp_base: • Hyper Threading = Disabled • DCU Streamer Prefetcher = Disabled • Sub NUMA (SNC) = Enable SNC2
• Operating system	Red Hat Enterprise Linux Server release 8.2 4.18.0-193.el8.x86_64
• Operating system settings	Default
• Compiler	C/C++: Version 2021.1 of Intel C/C++ Compiler for Linux Fortran: Version 2021.1 of Intel Fortran Compiler for Linux

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

Benchmark results

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

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

Processor	Number of cores	Number of processors	SPECrate2017_int_base		SPECrate2017_fp_base	
			CX2550 M6	CX2560 M6	CX2550 M6	CX2560 M6
Xeon Gold 5318Y	24	2	**298**	**299**	**306**	**304**
Xeon Gold 5318S	24	2	304 est.	305 est.	309 est.	307 est.
Xeon Gold 5317	12	2	191 est.	192 est.	207 est.	206 est.
Xeon Gold 5315Y	8	2	136 est.	136 est.	152 est.	151 est.
Xeon Silver 4316	20	2	250 est.	251 est.	261 est.	259 est.
Xeon Silver 4314	16	2	216 est.	217 est.	232 est.	231 est.
Xeon Silver 4310	12	2	165 est.	166 est.	186 est.	185 est.
Xeon Silver 4309Y	8	2	124 est.	125 est.	139 est.	138 est.

Processor	Number of cores	Number of processors	SPECspeed2017_int_base		SPECspeed2017_fp_base	
			CX2550 M6	CX2560 M6	CX2550 M6	CX2560 M6
Xeon Gold 5318Y	24	2	**11.3**	**11.3**	**171**	**170**

The following graph compares the throughput of PRIMERGY CX2550 M6/ CX2560 M6 with the air-cooled maximum performance confiugration (Xeon Gold 5318Y) and its older model, PRIMERGY CX2550 M5/ CX2560 M5 with the equivalent conifugration (Xeon Gold 6252).

SPECrate2017: Comparison of PRIMERGY CX2550 M6 and PRIMERGY CX2550 M5

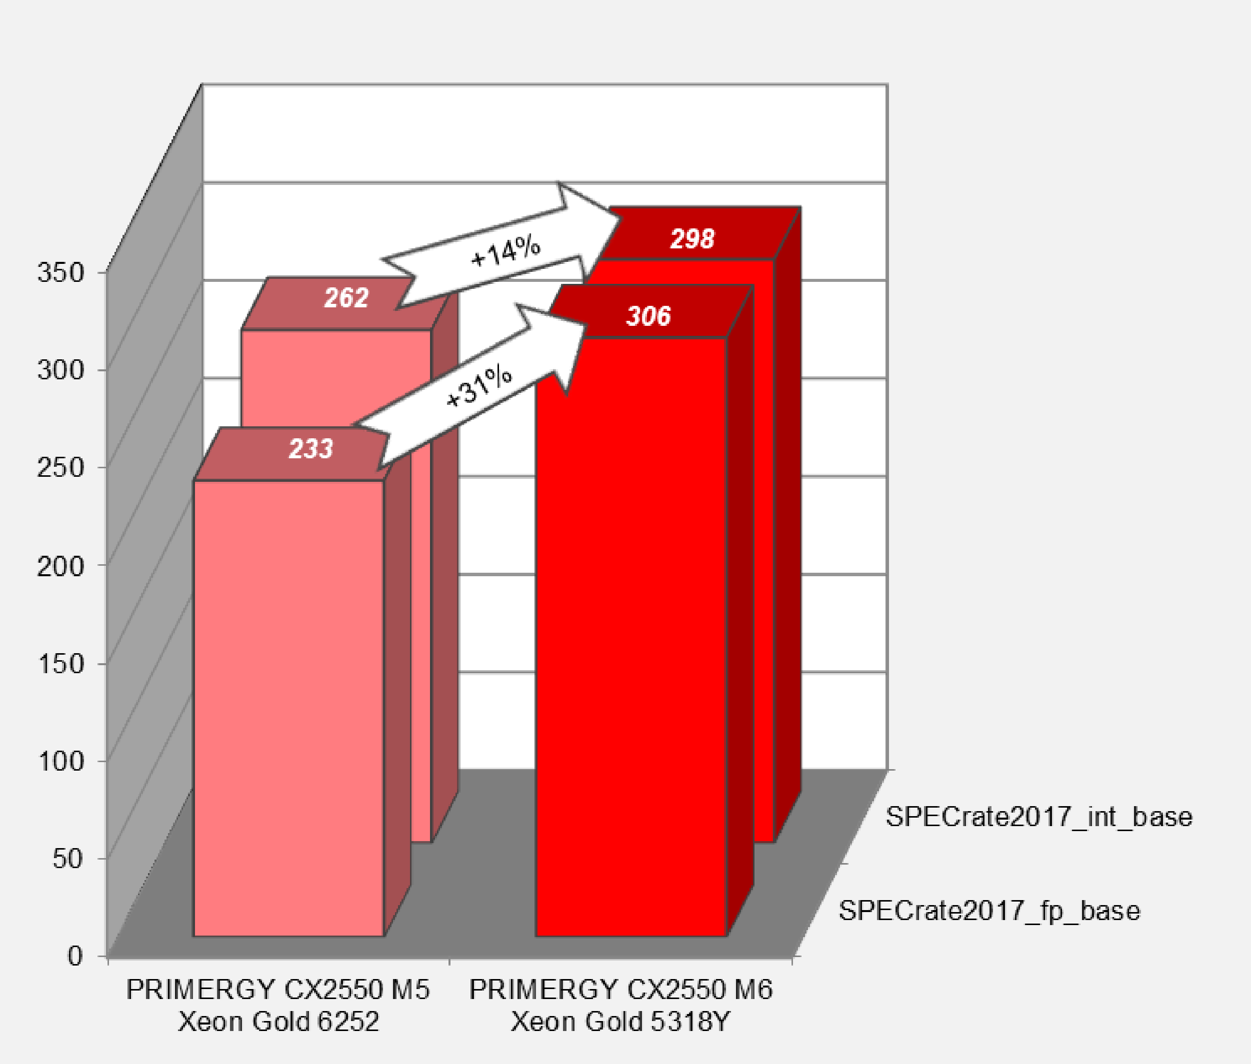

SPECrate2017: Comparison of PRIMERGY CX2560 M6 and PRIMERGY CX2560 M5

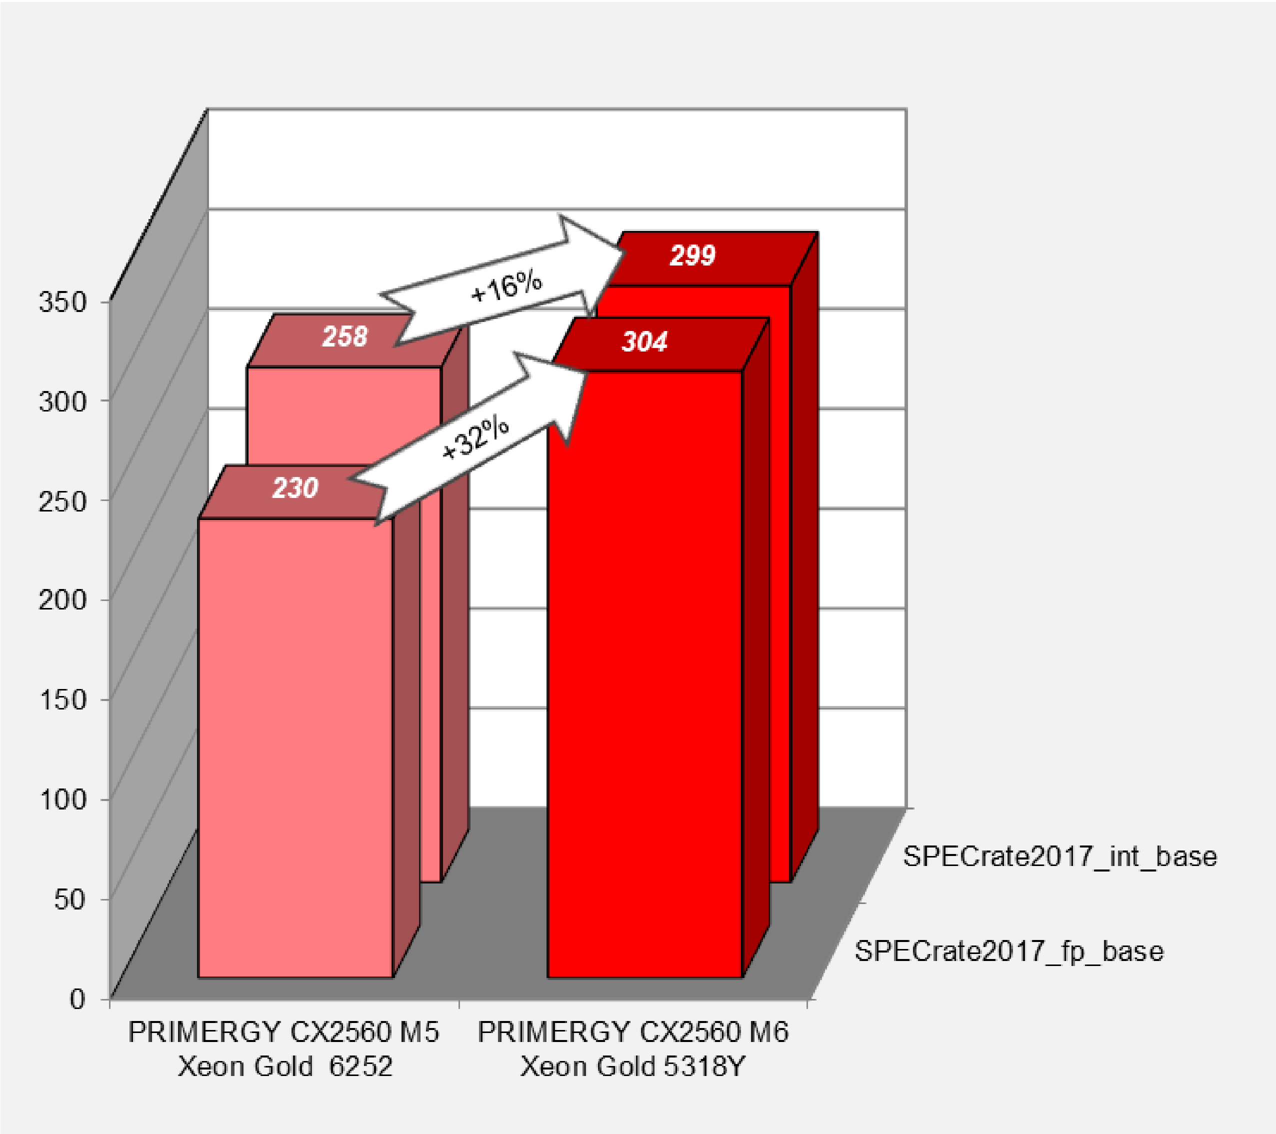

STREAM

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 CX2550 M6/ CX2560 M6
• Processor	2 x 3rd Generation Intel Xeon Scalable Processors Family
• Memory	16 x 32 GB 2Rx4 PC4-3200AA-R
Software	
• BIOS settings	• Override OS Energy Performance = Enabled • HWPM Support = Disabled • Intel Virtualization Technology = Disabled • LLC Dead Line Alloc = Disabled • Stale AtoS = Enabled • Sub NUMA (SNC) = Enable SNC2 • XPT Prefetch = Enabled
• Operating system	Red Hat Enterprise Linux Server release 8.2 4.18.0-193.el8.x86_64
• Operating system settings	Transparent Huge Pages inactivated echo 50000 > /proc/sys/kernel/sched_cfs_bandwidth_slice_us echo 240000000 > /proc/sys/kernel/sched_latency_ns echo 5000000 > /proc/sys/kernel/sched_migration_cost_ns echo 100000000 > /proc/sys/kernel/sched_min_granularity_ns echo 150000000 > /proc/sys/kernel/sched_wakeup_granularity_ns cpupower -c all frequency-set -g performance cpupower idle-set -d 1 cpupower idle-set -d 2 cpupower idle-set -d 3 echo 0 > /proc/sys/kernel/numa_balancing ulimit -s unlimited
• Compiler	C/C++: Version 19.1.2.254 of Intel C/C++ Compiler for Linux
• Benchmark	STREAM Version 5.10

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

Benchmark results

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

Processor	Memory frequency [MHz]	Maximum memory bandwidth*1 [GB/s]	Number of cores	Rated frequency [GHz]	Number of processors	TRIAD [GB/s] CX2550M6	TRIAD [GB/s] CX2560M6
Xeon Gold 5318Y	2,933	188	24	2.1	2	**295**	**293**
Xeon Gold 5318S	2,933	188	24	2.1	2	**295**	293 est.
Xeon Gold 5317	2,933	188	12	3.0	2	**229**	228 est.
Xeon Gold 5315Y	2,933	188	8	3.2	2	**175**	175 est.
Xeon Silver 4316	2,666	171	20	2.3	2	**261**	260 est.
Xeon Silver 4314	2,666	171	16	2.4	2	**249**	248 est.
Xeon Silver 4310	2,666	171	12	2.1	2	**224**	223 est.
Xeon Silver 4309Y	2,666	171	8	2.8	2	**161**	160 est.

LINPACK

Benchmark description

LINPACK was developed in the 1970s by Jack Dongarra and some other people to show the performance of supercomputers. The benchmark consists of a collection of library functions for the analysis and solution of linear system of equations. The description can be found in the following document.

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

LINPACK can be used to measure the speed of computers when solving a linear equation system. For this purpose, an n x n matrix is set up and filled with random numbers between -2 and +2. The calculation is then performed via LU decomposition with partial pivoting.

A memory of $8n^2$ bytes is required for the matrix. In case of an n x n matrix the number of arithmetic operations required for the solution is $2/3n^3 + 2n^2$. Thus, the choice of n determines the duration of the measurement. In other words, if n is doubled, the measurement time will be approximately eight times longer. The size of n also has an influence on the measurement result itself. As n increases, the measured value asymptotically approaches its limit. The size of the matrix is therefore usually adapted to the amount of memory available. Furthermore, the memory bandwidth of the system only plays a minor role for the measurement result, but a role that cannot be fully ignored. he processor performance is the decisive factor for the measurement result. Since the algorithm used permits parallel processing, in particular the number of processors used and their processor cores are - in addition to the clock rate - of outstanding significance.

LINPACK is used to measure how many floating point operations were carried out per second. The result is referred to as **Rmax** and specified in GFlops (Giga Floating Point Operations per Second: 1 billion floating point operations/second).

An upper limit, referred to as **Rpeak**, for the speed of a computer can be calculated from the maximum number of floating point operations that its processor cores could theoretically carry out in one clock cycle.

Rpeak = Maximum number of floating point operations per clock cycle

x Number of processor cores of the computer

x Rated processor frequency [GHz]

LINPACK is classed as one of the leading benchmarks in the field of high performance computing (HPC). LINPACK is one of the seven benchmarks currently included in the HPC Challenge benchmark suite, which takes other performance aspects in the HPC environment into account.

Manufacturer-independent publication of LINPACK results is possible at http://www.top500.org/. This requires using an HPL-based LINPACK version (see http://www.netlib.org/benchmark/hpl/).

Intel offers a highly optimized LINPACK version (shared memory version) for individual systems with Intel processors. Parallel processes communicate here via "shared memory," i.e. jointly used memory. Another version provided by Intel is based on HPL (High Performance Linpack). Intercommunication of the LINPACK processes here takes place via OpenMP and MPI (Message Passing Interface). This enables communication between the parallel processes - also from one computer to another. Both versions can be downloaded from http://software.intel.com/en-us/articles/intel-math-kernel-library-linpack-download/.

Manufacturer-specific LINPACK versions also come into play when graphics cards for General Purpose Computation on Graphics Processing Unit (GPGPU) are used. These are based on HPL and include extensions which are needed for communication with the graphics cards.Benchmark environment

Benchmark environment

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY CX2550 M6/ CX2560 M6
• Processor	2 x 3rd Generation Intel Xeon Scalable Processors Family
• Memory	16 x 32 GB 2Rx4 PC4-3200AA-R
Software	
• BIOS settings	• HyperThreading = Disabled • Link Frequency Select = 10.4 GT/s • HWPM Support = Disabled • Intel Virtualization Technology = Disabled • LLC Dead Line Alloc = Disabled • Stale AtoS = Enabled • Fan Control = Full
• Operating system	Red Hat Enterprise Linux Server release 8.2 4.18.0-193.el8.x86_64
• Operating system settings	Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1) cpupower -c all frequency-set -g performance echo 50000 > /proc/sys/kernel/sched_cfs_bandwidth_slice_us echo 240000000 > /proc/sys/kernel/sched_latency_ns echo 5000000 > /proc/sys/kernel/sched_migration_cost_ns echo 100000000 > /proc/sys/kernel/sched_min_granularity_ns echo 150000000 > /proc/sys/kernel/sched_wakeup_granularity_ns echo always > /sys/kernel/mm/transparent_hugepage/enabled echo 1048576 > /proc/sys/fs/aio-max-nr run with avx512
• Compiler	C/C++: Version 19.1.2.254 of Intel C/C++ Compiler for Linux
• Benchmark	Intel Optimized MP LINPACK Benchmark for Clusters

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

Benchmark results

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

Processor	Number of cores	Rated frequency [GHz]	Number of processors	Rpeak [GFlops]	CX2550 M6		CX2560 M6	
					Rmax [GFlops]	Efficiency	Rmax [GFlops]	Efficiency
Xeon Gold 5318Y	24	2.1	2	3,226	2,356	73%	2,385	74%
Xeon Gold 5318S	24	2.1	2	3,226	2,468	77%	2,499	77%
Xeon Gold 5317	12	3.0	2	2,304	1,645	71%	1,665	72%
Xeon Gold 5315Y	8	3.2	2	1,638	1,243	76%	1,258	77%
Xeon Silver 4316	20	2.3	2	2,944	2,002	68%	2,027	69%
Xeon Silver 4314	16	2.4	2	2,458	1,694	69%	1,714	70%
Xeon Silver 4310	12	2.1	2	1,613	1,473	91%	1,491	92%
Xeon Silver 4309Y	8	2.8	2	1,434	1,051	73%	1,064	74%

Rpeak values in the table above were calculated by the base frequency of each processor. Since we enabled Turbo mode in the measurements, the average Turbo frequency exceeded the base frequency for some processors.

As explained in the section "Technical Data," Intel generally does not guarantee that the maximum turbo frequency can be reached in the processor models due to manufacturing tolerances. A further restriction applies for workloads, such as those generated by LINPACK, with intensive use of AVX instructions and a high number of instructions per clock unit. Here the frequency of a core can also be limited if the upper limits of the processor for power consumption and temperature are reached before the upper limit for the current consumption. This can result in the achievement of a lower performance with turbo mode than without turbo mode. In such a case, disable the turbo function in the BIOS option.

SPECpower_ssj2008

Benchmark description

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

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

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

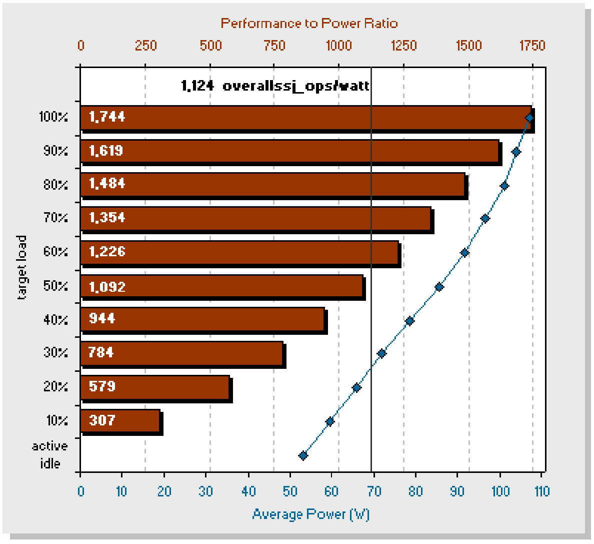

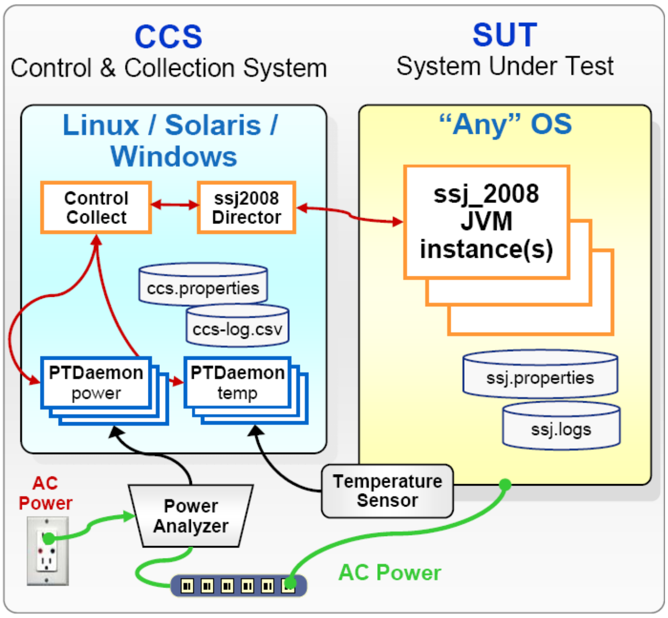

The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System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

Benchmark environment

System Under Test (SUT)	
For Linux OS measurement	
Hardware (chassis)	
• Enclosure	PRIMERGY CX400 M6
• Power Supply Unit	1 x S26113-E649-V90-1 2600W
• Number of servers	4
• Model	PRIMERGY CX2560 M6
Hardware (per node)	
• Processor	2 x Intel Xeon Gold 5318Y
• Memory	16 x 16 GB 2Rx8 PC4-3200Y-R
• Network interface	1 x Intel 10Gb X550T Gigabit Network Connection (onboard)
• Disk subsystem	1 x SSD M.2 240GB, S26361-F5787-E240
Software	
• BIOS	R1.23.0
• BIOS settings	Hardware Prefetcher = Disabled Adjacent Cache Line Prefetch = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Energy Performance = Energy Efficient Package C State limit = C6 UPI Link Frequency Select = 9.6GT/s Uncore Frequency Scaling = Power Balanced DDR Performance = Energy optimized SNC(Sub NUMA) = Enable SNC2 ASPM Support = L1 only SATA Controller = Disabled USB Port Control = Disable all ports Network Stack = Disabled
• iRMC Firmware	3.26P
• Operating system	SUSE Linux Enterprise Server 15 SP2, 5.3.18-24.43-default
• Operating system settings	kernel parameter: pcie_aspm=force pcie_aspm.policy=powersave intel_pstate=passive Benchmark started via ssh. modprobe cpufreq_conservative cpupower frequency-set -g conservative echo 3000000 > /sys/devices/system/cpu/cpufreq/conservative/sampling_rate echo 93 > /sys/devices/system/cpu/cpufreq/conservative/up_threshold echo 1 > /sys/devices/system/cpu/cpufreq/conservative/freq_step echo 92 > /sys/devices/system/cpu/cpufreq/conservative/down_threshold echo always > /sys/kernel/mm/transparent_hugepage/enabled cpupower frequency-set -u 2600MHz sysctl -w kernel.nmi_watchdog=0
• JVM	Oracle Java HotSpot 64-Bit Server VM 18.9 (build 11.0.9+7-LTS, mixed mode)
• JVM settings	-server -Xmn20000m -Xms22000m -Xmx22000m -XX:+UseHugeTLBFS -XX:+UseLargePages -XX:+UseTransparentHugePages -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchInstr=0 -XX:AllocatePrefetchLines=4 -XX:InlineSmallCode=3900

-XX:MaxInlineSize=270 -XX:ParallelGCThreads=8 -XX:SurvivorRatio=1
-XX:TargetSurvivorRatio=99 -XX:+UseParallelOldGC -XX:FreqInlineSize=2500
-XX:MinJumpTableSize=18 -XX:UseAVX=0 -XX:+UseBiasedLocking

For Windows OS measurement	
Hardware (chassis)	
• Enclosure	PRIMERGY CX400 M6
• Power Supply Unit	1 x S26113-E649-V90-1 2600W
• Number of servers	4
• Model	PRIMERGY CX2560 M6
Hardware (per node)	
• Processor	2 x Intel Xeon Gold 5318Y
• Memory	16 x 16 GB 2Rx8 PC4-3200Y-R
• Network interface	1 x Intel 10Gb X550T Gigabit Network Connection (onboard)
• Disk subsystem	1 x SSD M.2 240GB, S26361-F5787-E240
Software	
• BIOS	R1.23.0
• BIOS settings	Hardware Prefetcher = Disabled Adjacent Cache Line Prefetch = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Energy Performance = Energy Efficient Package C State limit = C6 UPI Link Frequency Select = 9.6GT/s Uncore Frequency Scaling = Power Balanced DDR Performance = Energy optimized SNC(Sub NUMA) = Enable SNC2 ASPM Support = L1 only SATA Controller = Disabled USB Port Control = Disable all ports Network Stack = Disabled
• iRMC Firmware	3.26P
• Operating system	Windows Server 2019 Standard
• Operating system settings	Turn off hard disk after = 1 Minute Turn off display after = 1 Minute Minimum processor state = 0% Maximum processor state = 100% Using the local security settings console, "lock pages in memory" was enabled for the user running the benchmark. Benchmark was started via Windows Remote Desktop Connection.
• JVM	Oracle Java HotSpot 64-Bit Server VM 18.9 (build 11.0.9+7-LTS, mixed mode)
• JVM settings	-server -Xmn1700m -Xms1950m -Xmx1950m -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:ParallelGCThreads=2 - XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:LoopUnrollLimit=45 - XX:InitialTenuringThreshold=12 -XX:MaxTenuringThreshold=15 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:FreqInlineSize=2500 -XX:+UseLargePages -XX:+UseParallelOldGC - XX:UseAVX=0 -XX:-UseAdaptiveSizePolicy -XX:-ThreadLocalHandshakes

Benchmark results (Linux)

The PRIMERGY CX2560 M6 in SUSE Linux Enterprise Server 15 SP2 achieved the following result:

SPECpower_ssj2008 = 8,903 overall ssj_ops/watt

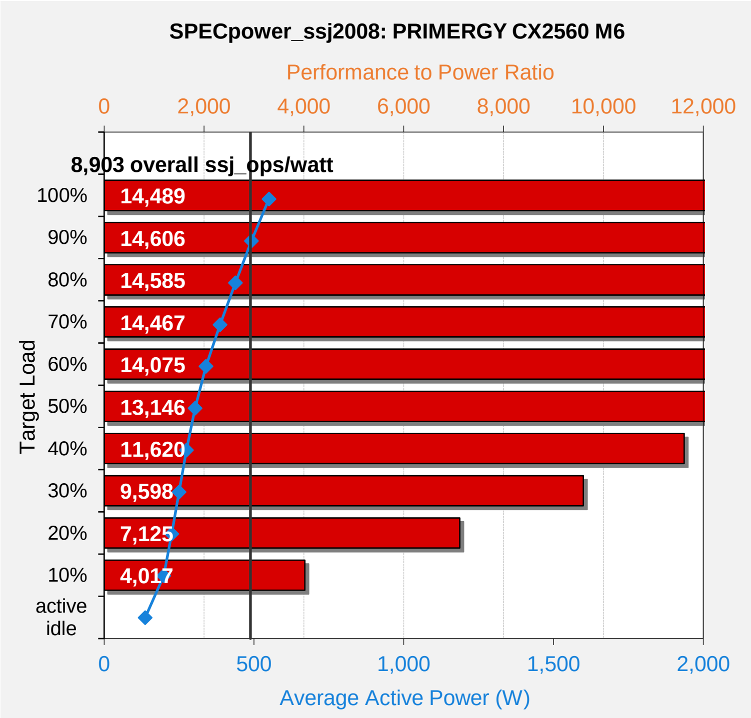

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 8,903 overall ssj_ops/watt for the PRIMERGY CX2560 M6. 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%	19,762,291	1,906	10,367
90%	17,822,302	1,715	10,394
80%	15,860,766	1,474	10,757
70%	13,904,810	1,292	10,759
60%	11,895,706	1,109	10,725
50%	9,899,337	1,009	9,812
40%	7,920,473	929	8,529
30%	5,940,102	856	6,942
20%	3,961,149	778	5,089
10%	1,980,785	696	2,845
Active Idle	0	472	0
			Σ ssj_ops / Σ power = 8,903

Benchmark results (Windows)

The PRIMERGY CX2560 M6 in Microsoft Windows Server 2016 Standard achieved the following result:

SPECpower_ssj2008 = 9,077 overall ssj_ops/watt

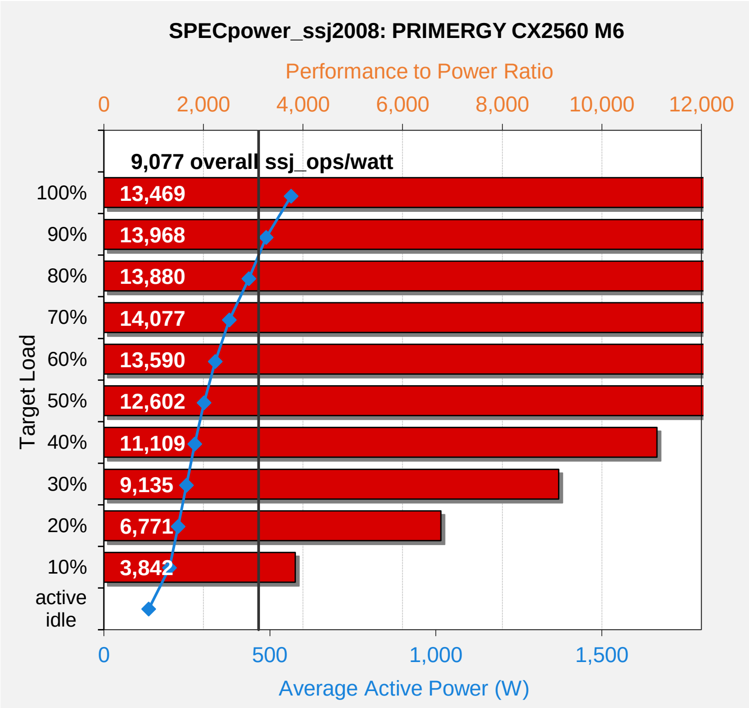

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 9,077 overall ssj_ops/watt for the PRIMERGY CX2560 M6. 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%	18,526,911	1,600	11,580
90%	16,684,528	1,464	11,399
80%	14,844,421	1,300	11,423
70%	12,987,497	1,159	11,210
60%	11,123,811	1,056	10,532
50%	9,267,551	972	9,532
40%	7,428,723	905	8,211
30%	5,561,324	837	6,641
20%	3,713,805	769	4,832
10%	1,852,036	694	2,669
Active Idle	0	482	0
			Σ ssj_ops / Σ power = 9,077

OLTP-2

Benchmark description

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

In contrast to benchmarks such as SPEC CPU and TPC-E, which were standardized by independent bodies and for which adherence to the respective rules and regulations are monitored, OLTP-2 is an internal benchmark of Fujitsu. OLTP-2 is based on the well-known database benchmark TPC-E. OLTP-2 was designed in such a way that a wide range of configurations can be measured to present the scaling of a system with regard to the CPU and memory configuration.

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

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

Benchmark environment

The typical measurement set-up is illustrated below:

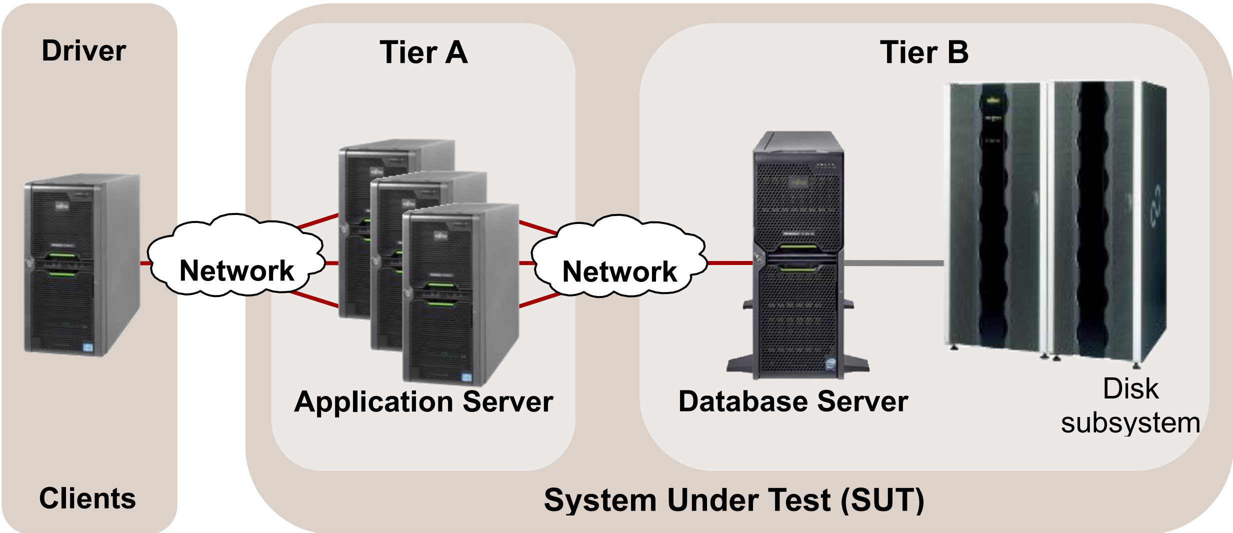

All OLTP-2 results were Calculated based on the configuration of the next following pages of PRIMERGY RX2540 M6

Database Server (Tier B)	
Hardware	
• Model	PRIMERGY RX2540 M6
• Processor	3rd Generation Intel Xeon Processor Scalable Family
• Memory	1 processors:16 x 64 GB (1x64 GB) 2Rx4 DDR4-3200 ECC 2 processors:32 x 64 GB (1x64 GB) 2Rx4 DDR4-3200 ECC
• Network interface	1 x Dual port LAN 10 Gbps 1 x Quad port OCPv3 LAN 1 Gbps
• Disk subsystem	RX2540 M6:RAID controller PRAID EP540i 6 x 1.6 TB SSD drive, RAID10 (LOG), 5 x RAID controller PRAID EP540e 10 x JX40 S2:4 x 1.6 TB SSD drive, RAID10 (temp), 68 x 1.6 TB SSD drive, RAID5 (data)
Software	
• BIOS	Version R1.6.0
• Operating system	Microsoft Windows Server 2016 Standard
• Database	Microsoft SQL Server 2017 Enterprise + KB4341265

Application Server (Tier A)	
Hardware	
• Model	1 x PRIMERGY RX2530 M4
• Processor	2 x Xeon Platinum 8180
• Memory	192 GB, 2666 MHz Registered ECC DDR4
• Network interface	1 x Dual port LAN 10 Gbps 1 x Dual port onboard LAN 1 Gbps
• Disk subsystem	2 x 300 GB 10k rpm SAS drive
Software	
• Operating system	Microsoft Windows Server 2016 Standard

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

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

Benchmark results

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

A guideline in the database environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses. This is why a configuration with a total memory of 1024 GB was considered for the calculation of two processors configuration and a configuration with a total memory of 512 GB for one processor configuration. Both memory configurations have memory access of 3200 MHz.

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

Processor	Cores	Threads	2CPU		1CPU	
			Score		Score	
Xeon Gold 5318Y	24	48	5,865	est.	3,222	est.
Xeon Gold 5318S	24	48	5,865	est.	3,222	est.
Xeon Gold 5317	12	24	3,419	est.	1,878	est.
Xeon Gold 5315Y	8	16	2,447	est.	1,344	est.
Xeon Silver 4316	20	40	4,673	est.	2,567	est.
Xeon Silver 4314	16	32	3,865	est.	2,123	est.
Xeon Silver 4310	12	24	2,969	est.	1,631	est.
Xeon Silver 4309Y	8	16	2,213	est.	1,216	est.

The following graph shows the OLTP-2 transaction rates obtained with the 3rd Generation Intel Xeon Processor Scalable Family (one or two processors).

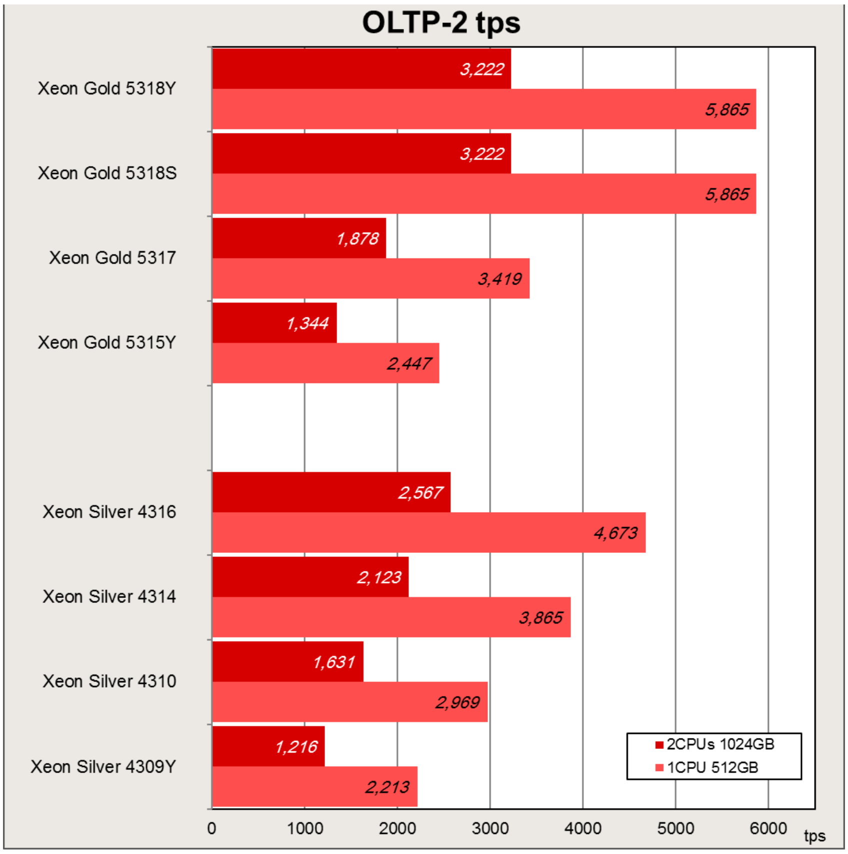

If you compare the OLTP-2 value of the processor with the lowest performance (Xeon Silver 4309Y) with the value of the processor with the highest performance (Xeon Gold 5318Y) the OLTP-2 value increased by a factor of 2.7.

The features of the processors are summarized in the section "Technical data."

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

The highest value for OLTP-2 on the CX2560 M6 is about 20% higher than the highest value on the previous model CX2560 M5.

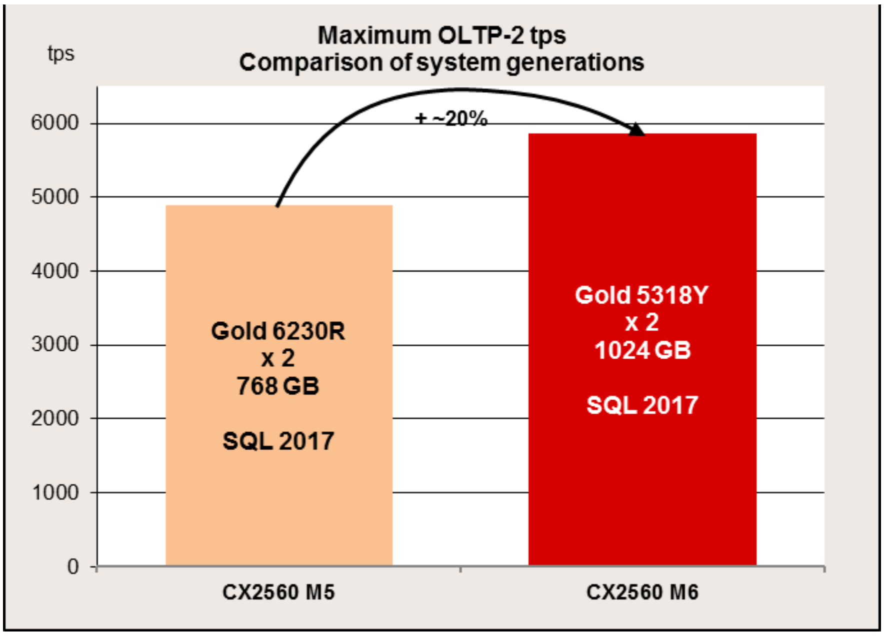

VMmark V3

Benchmark description

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

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

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

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

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

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

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

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

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

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

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

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

Benchmark environment

The typical measurement set-up is illustrated below:

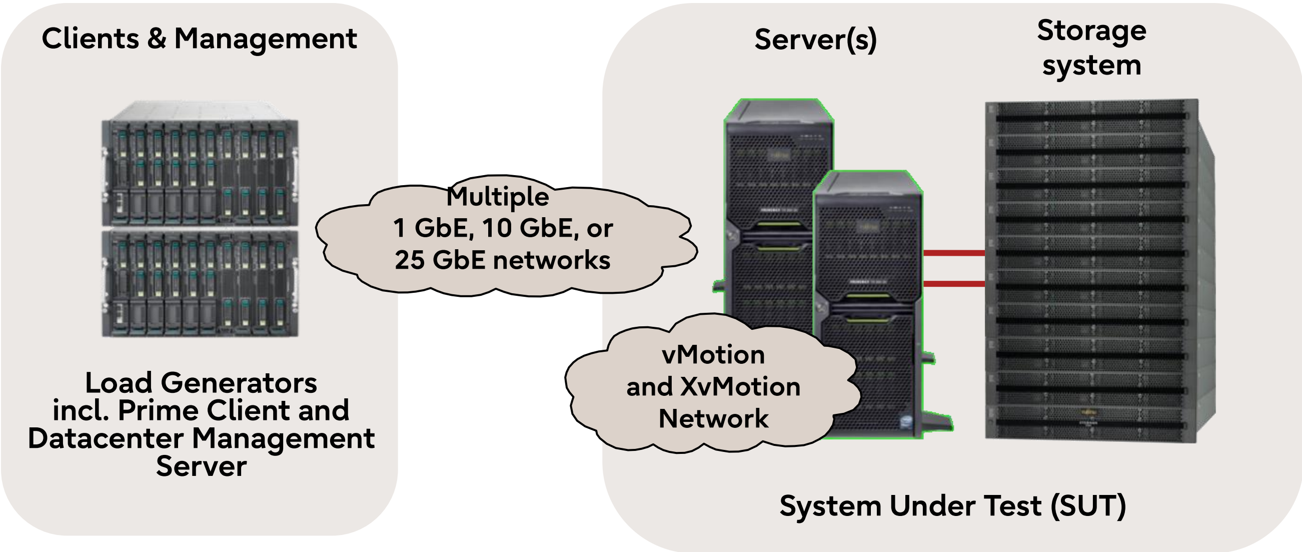

All the benchmark results were measured with the following environment:

System Under Test (SUT)	
Hardware	
• Number of servers	4
• Model	PRIMERGY CX2560 M6
• Processor	2 x Intel Xeon Gold 5318Y
• Memory	1,024 GB : 16 x 64 GB (1x64 GB) 2Rx4 DDR4-3200 R ECC
• Network interface	1 x Mellanox MCX4121A-ACAT dual port 25Gb SFP28 PCIe adapter 1 x Intel X550 1Gb onboard controller
• Disk subsystem	1 x Emulex LPe35002 dual port 32Gb PCIe adapter 6 x PRIMERGY RX2540 M4 & M5 configured as Fiber Channel targets 4 x PRIMERGY RX2540 M4: 2 x Micron MTFDDAK480TDC SATA SSD (480 GB, RAID1) 3 x Intel P4800X PCIe SSD (750 GB) 1 x Intel P4600 PCIe SSD (4 TB) 1 x PRIMERGY RX2540 M4: 1 x Micron MTFDDAK480TDC SATA SSD (480 GB) 3 x Intel P4800X PCIe SSD (750 GB) 1 x Intel P4600 PCIe SSD (2 TB) 1 x PRIMERGY RX2540 M5: 1 x Micron MTFDDAK480TDC SATA SSD (480 GB) 2 x Intel P4610 PCIe SSD (3.2 TB)
Software	
• BIOS	V1.0.0.0 R1.23.0 for D3894-A1x
• BIOS settings	See "Details"
• Operating system	VMware ESXi 7.0 U2a, Build 17867351
• Operating system settings	ESX settings: see "Details"

Detail	
See disclosure	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2021-11-16-Fujitsu-PRIMERGY-CX2560M6.pdf

Datacenter Management Server (DMS)	
Hardware	
• Model	1 x PRIMERGY RX2540 M2
• Processor	1 x Intel Xeon E5-2698 v4
• Memory	64 GB
• Network interface	1 x Emulex One Connect Oce14000 1GbE dual port PCIe adapter
Software	
• Operating system	VMware ESXi 6.7 EP 02a Build 9214924
Datacenter Management Server (DMS) VM	
Hardware	
• Processor	4 x Logical CPU
• Memory	19 GB
• Network interface	1 x 1 Gbit/s LAN
Software	
• Operating system	VMware vCenter Server Appliance 7.0 Update 2 Build 17694817

Load generator	
Hardware	
• Model	4 x PRIMERGY RX2530 M2
• Processor	3 x PRIMERGY RX2530 M2 2 x Intel Xeon E5-2699 v4 1 x PRIMERGY RX2530 M2 2 x Intel Xeon E5-2699A v4
• Memory	3 x 256 GB, 1 x 240 GB
• Network interface	1 x Emulex One Connect Oce14000 1GbE dual port PCIe adapter 1 x Emulex One Connect Oce14000 10GbE dual port PCIe adapter
Software	
• Operating system	VMware ESXi 6.7 EP 08 Build 13473784

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

Benchmark results

"Performance Only" measurement results (November 16, 2021)

On November 16, 2021, Fujitsu achieved a VMmark V3.1.1 score of "15.48@15 tiles" using PRIMERGY CX2560 M6 with Xeon Gold 5318Y processors and VMware ESXi 7.0 U2a. At this time, the system configuration had a total of 4 x 48 processor cores, and four identical server nodes were used for the "System Under Test" (SUT).

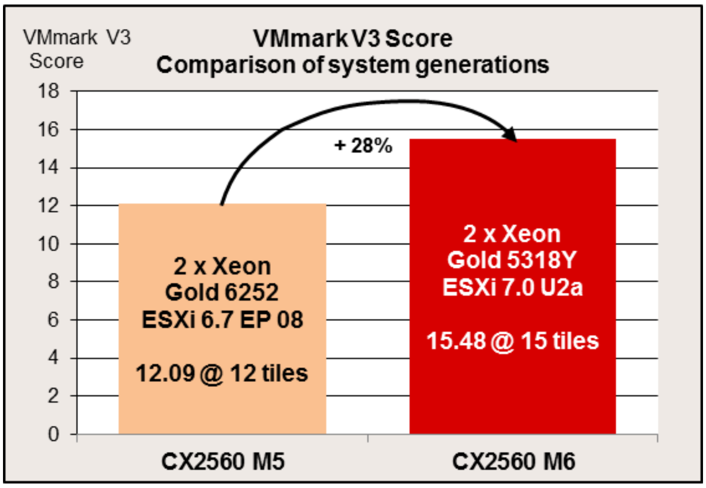

The graph on the left compares the VMmark V3 scores of the PRIMEGY CX2560 M6 and the previous generation PRIMEGY CX2560 M5.

The PRIMERGY CX2560 M6 achieved a 28% improvement in score compared to the previous generation PRIMERGY CX2560 M5. This is due to the improved performance of the 3rd generation Intel Xeon scalable processor and the effective use of the capabilities of the VMware ESXi hypervisor.

Literature

PRIMERGY Servers

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

PRIMERGY CX2550 M6/ CX2560 M6

This Whitepaper

https://docs.ts.fujitsu.com/dl.aspx?id=994cfea6-22e9-4232-9dfb-0e5e2abfcdea

https://docs.ts.fujitsu.com/dl.aspx?id=167c9238-027a-47a9-aacf-ff5d3eaf0a82

Data sheet

https://docs.ts.fujitsu.com/dl.aspx?id=42897e4f-bc18-4ca9-89e1-bc2d53e49d9f

PRIMERGY Performance

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

SPEC CPU2017

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

Benchmark Overview SPECcpu2017

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

STREAM

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

LINPACK

The LINPACK Benchmark: Past, Present, and Future

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

TOP500

https://www.top500.org/

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

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

Intel Math Kernel Library – LINPACK Download

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

SPECpower_ssj2008

https://www.spec.org/power_ssj2008

Benchmark Overview SPECpower_ssj2008

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

OLTP-2

Benchmark Overview OLTP-2

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

VMmark V3

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

Benchmark Overview VMmark V3

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

Document change history

Version	Date	Description
1.2	2023-10-03	Update: • New Visual Identity format
1.1	2022-07-05	Update: • Fixed contact information
1.0	2021-12-10	New: • Technical data • SPECcpu2017, OLTP-2, STREAM, LINPACK Measured and calculated with 3rd Generation Intel Xeon Processor Scalable Family • SPECpower_ssj2008, VMmark V3 Measured with Intel Xeon Gold 5318Y

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