Fujitsu Server PRIMERGY
Performance Report
PRIMERGY RX2520 M5

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMERGY RX2520 M5.

Explaines PRIMERGY RX2520 M5 performance data in comparison to other PRIMERGY models. In addition to the benchmark results, the explanation for each benchmark and benchmark environment are also included.

Version
1.5 2023-10-03

Contents

Technical data

PRIMERGY RX2520 M5

PY PRIMERGY RX2520 M5 12 x 3.5'

PRIMERGY RX2520 M5

PY PRIMERGY RX2520 M5 16 x 2.5' expandable

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 RX2520 M5
Model versions	PRIMERGY RX2520 M5 / 4x3.5” PRIMERGY RX2520 M5 / 12x3.5” PRIMERGY RX2520 M5 / 8x2.5” PRIMERGY RX2520 M5 / 16x2.5” PRIMERGY RX2520 M5 / 24x2.5” PRIMERGY RX2520 M5 / PCIe-SSD
Form factor	Rack server
Chipset	Intel C624
Number of sockets	2
Number of processors orderable	1 or 2
Processor type	2nd Generation Intel Xeon Scalable Processors Family
Number of memory slots	12 (6 per processor)
Maximum memory configuration	768 GB
Onboard HDD controller	Controller with RAID 0, RAID 1 or RAID 10 for up to 8 SATA HDDs
PCI slots	1 × PCI-Express 3.0 x8 3 × PCI-Express 3.0 x16
Max. number of internal hard disks	PRIMERGY RX2520 M5 / 4x 3.5” : 8 PRIMERGY RX2520 M5 / 12x 3.5” : 12 PRIMERGY RX2520 M5 / 8x 2.5” : 24 PRIMERGY RX2520 M5 / 16x2.5” : 16 PRIMERGY RX2520 M5 / 24x2.5” : 24 PRIMERGY RX2520 M5 / PCIe-SSD : 8x2.5” SAS/SATA + 4xPCIe-SSD (2.5”)

Processors (since system release)								
Processor	Cores	Threads	Cache [MB]	UPI Speed [GT/s]	Rated Frequency [GHz]	Max. Turbo Frequency [GHz]	Max. Memory Frequency [MHz]	TDP [W]
July 2019 released								
Xeon Gold 5222	4	8	16.5	10.4	3.8	3.9	2,933	105
Xeon Gold 5220	18	36	24.8	10.4	2.2	3.9	2,666	125
Xeon Gold 5218	16	32	22.0	10.4	2.3	3.9	2,666	125
Xeon Gold 5218B	16	32	22.0	10.4	2.3	3.9	2,666	125
Xeon Gold 5217	8	16	11.0	10.4	3.0	3.7	2,666	115
Xeon Gold 5215	10	20	13.8	10.4	2.5	3.4	2,666	85
Xeon Silver 4216	16	32	22.0	9.6	2.1	3.2	2,400	100
Xeon Silver 4215	8	16	11.0	9.6	2.5	3.5	2,400	85
Xeon Silver 4214Y	12	24	16.5	9.6	2.2	3.2	2,400	85
	10	20						
	8	16						
Xeon Silver 4214	12	24	16.5	9.6	2.2	3.2	2,400	85
Xeon Silver 4210	10	20	13.8	9.6	2.2	3.2	2,400	85
Xeon Silver 4208	8	16	11.0	9.6	2.1	3.2	2,400	85
Xeon Bronze 3204	6	6	8.3	9.6	1.9		2,133	85
March 2020 released								
Xeon Gold 5218R	20	40	27.5	10.4	2.1	4.0	2,666	125
Xeon Silver 4215R	8	16	11.0	9.6	3.2	4.0	2,400	130
Xeon Silver 4214R	12	24	16.5	9.6	2.4	3.5	2,400	100
Xeon Silver 4210R	10	20	13.8	9.6	2.4	3.2	2,400	100
Xeon Bronze 3206R	8	8	11.0	9.6	1.9		2,133	85

All the processors that can be ordered with the PRIMERGY RX2520 M5, apart from Xeon Bronze 3204 and Xeon Bronze 3206R, 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.

Suffix of Processor number shows additional feature of Xeon Processor.

The processors with M/L suffix support larger memory capacity of 2TB/socket(M-suffix) or 4.5TB/socket(L-suffix) whereas normal processors support 1TB/socket memory capacity.

The processors with S suffix are specifically designed to offer consistent performance for search workloads.

The processors with U suffix are only capable of single socket but the prices are lower than comparable normal processors with the same core count and frequency.

The processors with V suffix are specifically designed to help maximize $/VM

The processors with Y suffix support Intel Speed Select Technology. It enables to provide 3 distinct configurations(number of active cores and frequencies) which customer can choose in BIOS option.

Specifications of Xeon Gold 5218B and Xeon Gold 5218 including core count and frequencies are the same. The difference is minor electrical specifications only.

Suffix	Additional feature
M	Support up to 2TB/socket memory
L	Support up to 4.5TB/socket memory
S	Search Optimized
U	Single Socket
V	VM Density Optimized
Y	Speed Select

Memory modules								
Type	Capacity	Number of ranks	Bit width of the memory chips	Frequency	Load Reduced	Registered	NVDIMM	ECC
8 GB (1x8 GB) 1Rx8 DDR4-2933 R ECC	8	1	8	2,933		✓		✓
16 GB (1x16 GB) 2Rx8 DDR4-2933 R ECC	16	2	8	2,933		✓		✓
16 GB (1x16 GB) 1Rx4 DDR4-2933 R ECC	16	1	4	2933		✓		✓
32 GB (1x32 GB) 2Rx4 DDR4-2933 R ECC	32	2	4	2,933		✓		✓
64 GB (1x64 GB) 2Rx4 DDR4-2933 R ECC	64	2	4	2,933		✓		✓
8 GB (1x8 GB) 1Rx8 DDR4-2933 R ECC	8	1	8	2,933		✓		✓

Power supplies		Maximum number
Modular redundant PSU	450W platinum PSU	2
	800W platinum PSU	2
	800W titanium PSU	2

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

Detailed technical information is available in the data sheet PRIMERGY RX2520 M5.

SPEC CPU2017

Benchmark description

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

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

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

Benchmark	Number of single benchmarks	Arithmetics	Type	Compiler optimization	Measurement result
SPECspeed2017_int_peak	10	integer	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 RX2520 M5
• Processor	2nd Generation Intel Xeon Scalable Processors Family
• Memory	12 × 32 GB (1 x 32 GB) 2Rx4 PC4-2933Y-R
Software	
• BIOS settings	SPECrate2017_int_base: • Patrol Scrub = Disabled • DCU Ip Prefetcher = Disabled*1 • DCU Streamer Prefetcher = Disabled*1 • Fan Control = Full • Stale AtoS = Enable • WR CRC feature Control = Disabled • Sub NUMA Clustering = Disabled*2 • Hyper-Threading = Disabled*3 SPECrate2017_fp_base: • Patrol Scrub = Disabled • WR CRC feature Control = Disabled • Sub NUMA Clustering = Disabled *2 • Hyper-Threading = Disabled*3 • Fan Control = Full *1: Xeon Bronze 3206R、Xeon Silver 4210R、Xeon Silver 4214R、Xeon Silver 4215 Res *2: Xeon Gold 5217, Xeon Gold 5215, Xeon Silver 4215, Xeon Silver 4210, Xeon Silver 4208, Xeon Bronze 3204, Xeon Bronze 3206R, Xeon Silver 4210R, Xeon Silver 4215R *3: Xeon Bronze 3204, Xeon Bronze 3206R
• Operating system	SUSE Linux Enterprise Server 15 4.12.14-25.28-default
• Operating system settings	Stack size set to unlimited using "ulimit -s unlimited" Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1) SPECrate2017_fp: echo 10000000 > /proc/sys/kernel/sched_min_granularity_ns
• Compiler	SPECrate2017_int: CPU released in July 2019 C/C++: Version 19.0.1.144 of Intel C/C++ Compiler for Linux Fortran: Version 19.0.1.144 of Intel Fortran Compiler for Linux CPU released in March 2020 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 SPECrate2017_fp:

	CPU released in July 2019 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 CPU released in March 2020 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

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	Cores	Processor	SPECrate2017	
			SPECrate2017_int_base	SPECrate2017_fp_base
July 2019 released				
Xeon Gold 5222	4	2	63.0 est.	76.9 est.
Xeon Gold 5220	18	2	**198**	**191**
Xeon Gold 5218	16	2	181 est.	179 est.
Xeon Gold 5217	8	2	106 est.	117 est.
Xeon Gold 5215	10	2	120 est.	127 est.
Xeon Silver 4216	16	2	175 est.	170 est.
Xeon Silver 4215	8	2	95.9 est.	107 est.
Xeon Silver 4214Y	12	2	132 est.	139 est.
	10	2	111 est.	123 est.
	8	2	95 est.	112 est.
Xeon Silver 4214	12	2	132 est.	138 est.
Xeon Silver 4210	10	2	109 est.	118 est.
Xeon Silver 4208	8	2	81.8 est.	92.5 est.
Xeon Bronze 3204	6	2	39.0 est.	54.5 est.
April 2019 released				
Xeon Gold 5218R	20	2	**198**	**197**
Xeon Silver 4215R	8	2	91.3 est.	108 est.
Xeon Silver 4214R	12	2	121 est.	142 est.
Xeon Silver 4210R	10	2	98.7 est.	120 est.
Xeon Bronze 3206R	8	2	46.0 est.	71.6 est.

The following two diagrams illustrate the throughput of the PRIMERGY RX2520 M5 in comparison to its predecessor PRIMERGY RX2520 M4, in their respective most performant configurations.

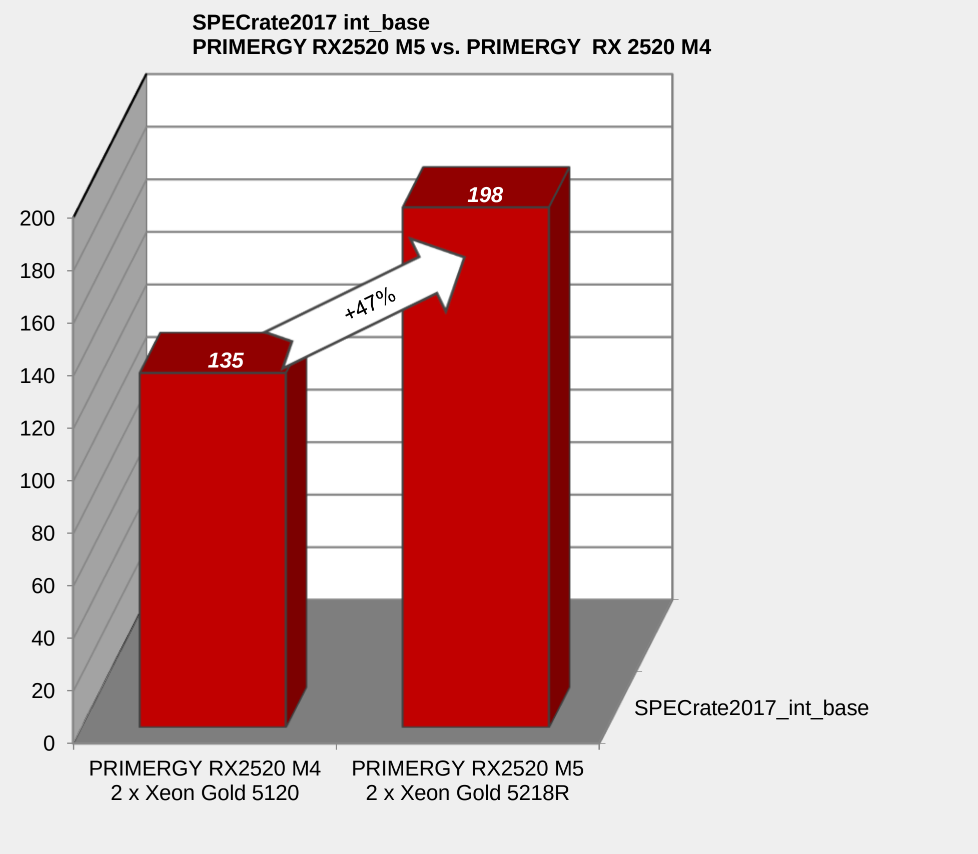

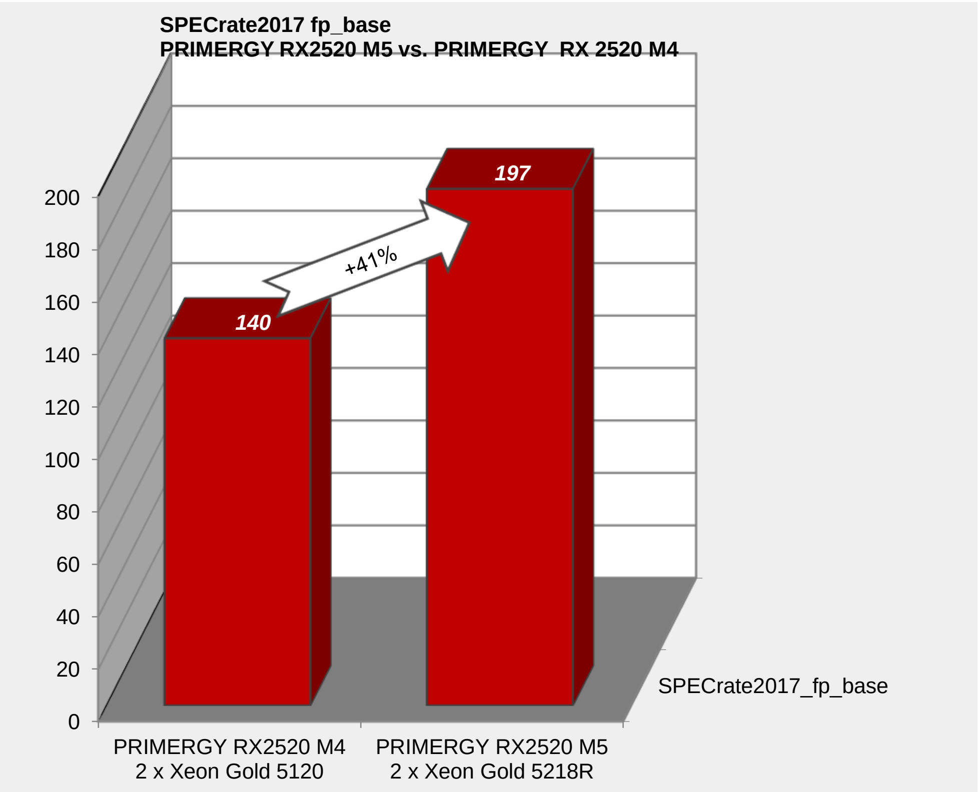

STREAM

Benchmark description

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

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 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 RX2520 M5
• Processor	2nd Generation Intel Xeon Scalable Processors Family
• Memory	24 x 32GB (1x32GB) 2Rx4 PC4-2933Y-R
Software	
• BIOS settings	• IMC Interleaving = 1-way • Override OS Energy Performance = Enabled • HWPM Support = Disable • Intel Virtualization Technology = Disabled • Energy Performance = Performance • LLC Dead Line Alloc = Disabled • Stale AtoS = Enabled • WR CRC feature Control = Disabled
• Operating system	SUSE Linux Enterprise Server 15
• Operating system settings	nohz_full=1-111 echo never > /sys/kernel/mm/transparent_hugepage/enabled run with avx512
• Compiler	C/C++: Version 2019.3.0.591499 of Intel C/C++ Compiler for Linux
• Benchmark	STREAM Version 5.10

Benchmark results

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

Processor	Memory Frequency [MHz]	Maximum memory bandwidth [GB/s]	Number of cores	Rated frequency [GHz]	Number of processors	TRIAD [GB/s]
July 2019 released						
Xeon Gold 5222	2,933	140.8	4	3.8	2	103 est.
Xeon Gold 5220	2,666	128.0	18	2.2	2	**203**
Xeon Gold 5218	2,666	128.0	16	2.3	2	202 est.
Xeon Gold 5217	2,666	128.0	8	3	2	134 est.
Xeon Gold 5215	2,666	128.0	10	2.5	2	151 est.
Xeon Silver 4216	2,400	115.2	16	2.1	2	190 est.
Xeon Silver 4215	2,400	115.2	8	2.5	2	94.6 est.
Xeon Silver 4214Y	2,400	115.2	12	2.2	2	161 est.
	2,400	115.2	10	2.2	2	tbd.
	2,400	115.2	8	2.2	2	tbd.
Xeon Silver 4214	2,400	115.2	12	2.2	2	161 est.
Xeon Silver 4210	2,400	115.2	10	2.2	2	92.7 est.
Xeon Silver 4208	2,400	115.2	8	2.1	2	92.5 est.
Xeon Bronze 3204	2,133	102.4	6	1.9	2	79.1 est.
March 2020 released						
Xeon Gold 5218R	2,666	128.0	20	2.1	2	**198**
Xeon Silver 4215R	2,400	115.2	8	3.2	2	109 est.
Xeon Silver 4214R	2,400	115.2	12	2.4	2	153 est.
Xeon Silver 4210R	2,400	115.2	10	2.4	2	87.0 est.
Xeon Bronze 3206R	2,133	102.4	8	1.9	2	75.7 est.

*1: Value per Processor

The following diagram shows the throughput of the PRIMERGY RX2520 M5 in comparison to its predecessor, the PRIMERGY RX2520 M4.

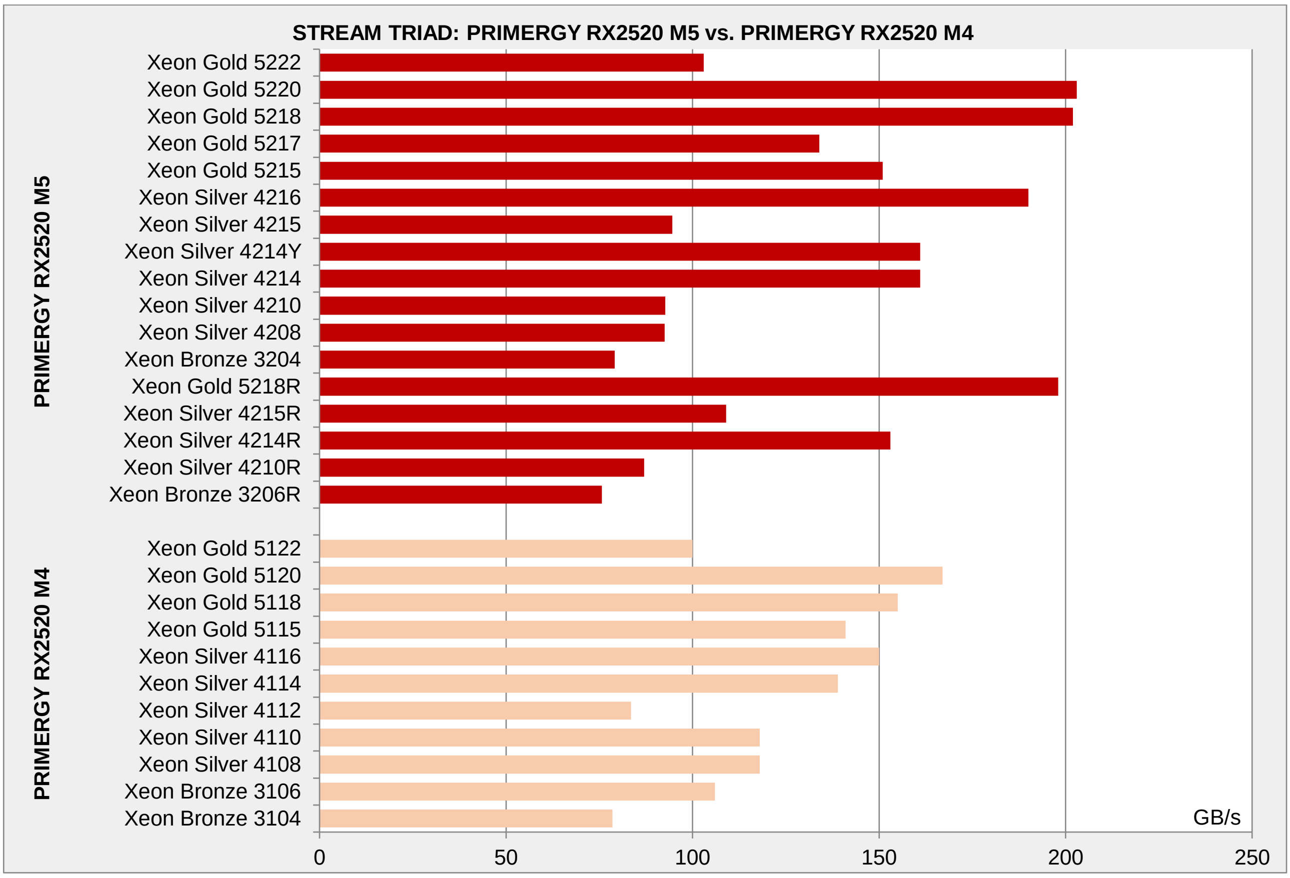

SPECpower_ssj2008

Benchmark description

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

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

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

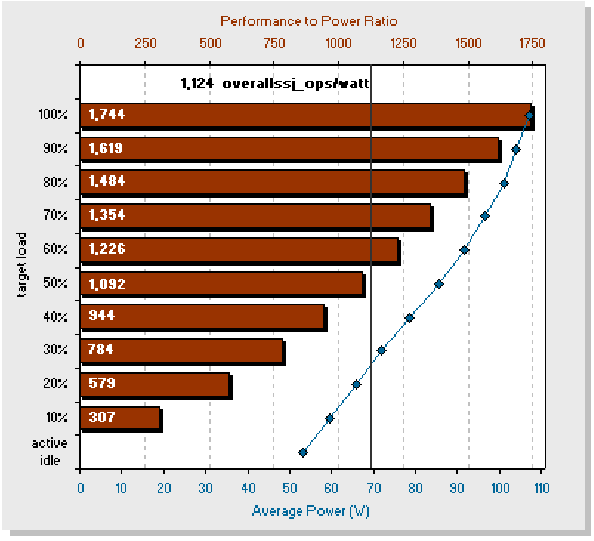

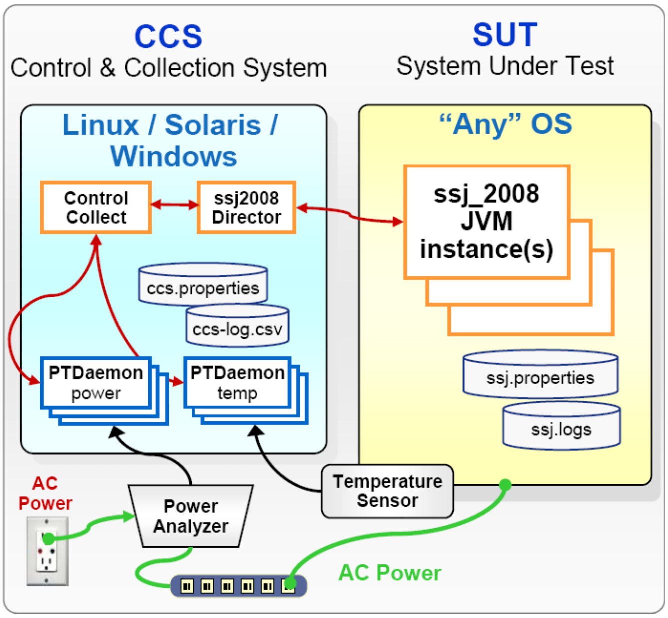

The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System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

Benchmark environment

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY RX2520 M5
• Processor	Intel Xeon Gold 5220
• Memory	12 ×16 GB (1 x 16 GB) 2Rx8 PC4-2933Y-R
• Network interface	1 x Intel I350 Gigabit Network Connection (onboard)
• Disk subsystem	1 x SSD M.2 240 GB, S26361-F5706-E240
• Power Supply Unit	1 × Fujitsu Technology Solutions S26113-F574-E13
Software	
• BIOS	R1.11.0
• BIOS settings	SATA Controller = Disabled. Serial Port = Disabled. Hardware Prefetcher = Disabled. Adjacent Cache Line Prefetch = Disabled. DCU Streamer Prefetcher = Disabled. Intel Virtualization Technology = Disabled. Override OS Energy Performance = Enabled. Energy Performance = Energy Efficient. DDR Performance = Power balanced.(effective memory frequency = 2400 MHz) Autonomous C-state Support = Enabled. ASPM Support = Auto. UPI Link Frequency Select = 9.6GT/s. Uncore Frequency Override = Power balanced. IMC Interleaving = 1-way. USB Port Control = Enable internal ports only. Network Stack = Disabled.
• Firmware	2.43P
• 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=1-71 nohz_full=1-71 isolcpus=1-71 Benchmark started via ssh modprobe cpufreq_conservative cpupower frequency-set --governor conservative echo -n 95 > /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 94 > /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 sysctl -w vm.swappiness=50 sysctl -w vm.laptop_mode=5
• JVM	Oracle Java HotSpot 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 -XX:UseAVX=0

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

Benchmark results

The PRIMERGY RX2520 M5 in SUSE Linux Enterprise Server 12 SP4 achieved the following result:

SPECpower_ssj2008 = 10,765 overall ssj_ops/watt

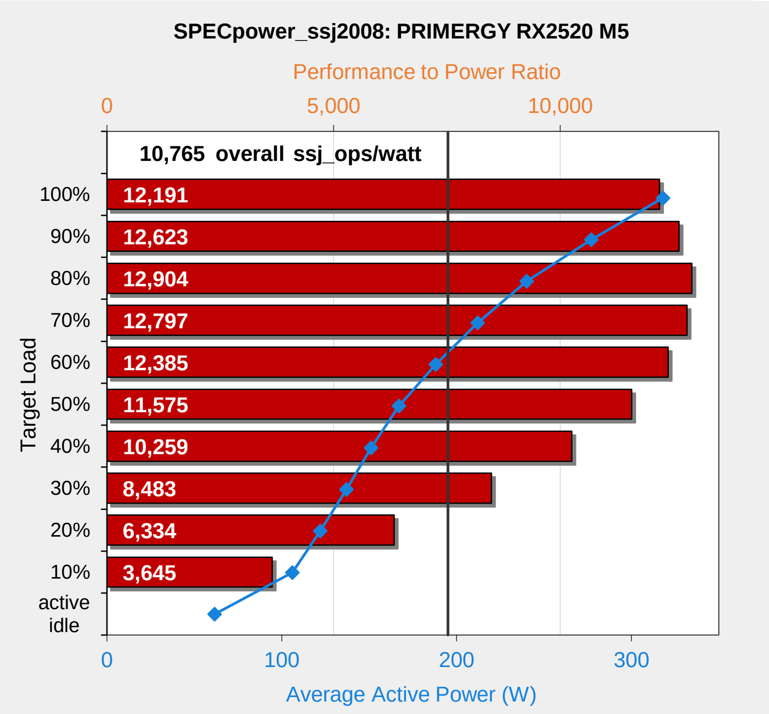

The adjoining diagram shows the result for the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (top x-axis) for each target load level on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhombus. The black vertical line shows the benchmark result of 10,765 overall ssj_ops/watt for the PRIMERGY RX2520 M5. This is the quotient of the sum of the transaction throughputs for each load level and the sum of the average power consumed for each measurement interval.

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

Performance		Power	Energy Efficiency
Target Load	ssj_ops	Average Power (W)	ssj_ops/watt
100%	3,872,265	318	12,191
90%	3,493,729	277	12,623
80%	3,093,414	240	12,904
70%	2,712,055	212	12,797
60%	2,325,712	188	12,385
50%	1,938,275	167	11,575
40%	1,553,224	151	10,259
30%	1,161,157	137	8,483
20%	774,176	122	6,334
10%	387,384	106	3,645
Active Idle	0	61.5	0
			Σssj_ops / Σpower = 10,765

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 (kiB)
- Number of parallel accesses (# of outstanding I/Os)

A given value combination of these specifications is known as "load profile." The following five standard load profiles can be allocated to typical application scenarios.

Standard load profile	Access	Type of access		Block size [kiB]	Application
		read	write		
Filecopy	Random	50%	50%	64	Copying files
Fileserver	Random	67%	33%	64	Fileserver
Database	Random	67%	33%	8	Database (data transfer) Mail server
Streaming	Sequential	100%	0%	64	Database (log file), Data backup, Video streaming (partial)
Restore	Sequential	0%	100%	64	Restoring files

In order to model applications that access in parallel with a different load intensity the Queue Depth is increased from 1 to 512 (in steps to the power of two).

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

The main measurement items are as follows.

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

The data throughput has established itself as the normal measurement variable for sequential load profiles, whereas the measurement variable "transaction rate" is mostly used for random load profiles with their small block sizes. Data throughput and transaction rate are directly proportional to each other and can be transferred to each other according to the following formula.

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

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

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

Benchmark environment

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

System Under Test (SUT)

Hardware

3.5 inch Model:

Controller:

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

Controller: Integrated PCI Express controller
CPU: 2x Intel Xeon Gold 5222 (3.80GHz)

Storage media	Category	Drive Name
SSD	PCIe SSD AIC (Write Intensive)	SSDPED1K375GA *2 *4 SSDPED1K750GA *2 *4

Controller: Intel C620 Standard SATA AHCI controller

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

2.5 inch Model:

Controller: 1x PRAID CP400i

Storage media	Category	Drive Name
HDD	SAS HDD (SAS 12Gbps, 10k rpm) [512e]	AL15SEB06EQ *2 *3
	SAS HDD (SAS 12Gbps, 10k rpm) [512n]	AL15SEB030N *2 *3
	SAS HDD (SAS 12Gbps, 15k rpm) [512n]	ST300MP0006 *1 *3
	NL-SAS HDD (SAS 12Gbps, 7.2k rpm) [512n]	ST1000NX0453 *1 *3
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512e]	ST1000NX0313 *1 *3
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512n]	ST2000NX0403 *1 *3
SSD	SATA SSD (SATA 6Gbps, Mixed Use)	MZ7KH240HAHQ *2 *3 MZ7KH480HAHQ *2 *3 MZ7KH960HAJR *2 *3 MZ7KH1T9HAJR *2 *3 MZ7KH3T8HALS *2 *3
	SATA SSD (SATA 6Gbps, Read Intensive)	MTFDDAK240TCB *2 *3 MTFDDAK480TDC *2 *3 MTFDDAK960TDC *2 *3 MTFDDAK1T9TDC *2 *3 MTFDDAK3T8TDC *2 *3 MTFDDAK7T6TDC *2 *3

Controller: Integrated PCI Express controller
CPU: 2x Intel Xeon Gold 5222 (3.80GHz)

Storage media	Category	Drive Name
SSD	2.5 inch PCIe SSD (Write Intensive)	SSDPE21K750GA *2 *4
	2.5 inch PCIe SSD (Mixed Use)	SSDPE2KE016T8 *2 *4 SSDPE2KE032T8 *2 *4 SSDPE2KE064T8 *2 *4
SSD	PCIe SSD (Write Intensive)	SSDPED1K375GA *2 *4 SSDPED1K750GA *2 *4

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

Software

Operating system		Microsoft Windows Server 2012 Standard R2 Microsoft Windows Server 2016 Standard
Stripe size		Controller default (here 64 KiB)
Measuring tool		Iometer 1.1.0
Measurement area	Type1	RAW file system is used. The first 10% of the usable LBA area is used for sequential accesses; the next 25% for random accesses.
	Type2	NTFS file system is used. The 32GiB area is secured for the first of the target drive and is used for sequential access and random access.
Total number of Iometer workers		1
Alignment of Iometer accesses		Aligned to whole multiples of 4096 bytes

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

Benchmark results

The results shown here are intended to help you select the appropriate storage media under the aspect of disk-I/O performance. For this purpose, a single storage medium was measured in the configuration specified in the subsection "Benchmark environment".

Controller

The measurements were made using controllers in the table below.

Storage medium	Storage medium	Cache	Supported interfaces		RAID levels
			host	drive	
SSD/HDD	PRAID CP400i	-	8x PCIe 3.0	SATA 6G SAS 12G	0, 1, 1E, 10, 5, 50
PCIe SSD	Integrated PCI Express controller	-	4x PCIe 3.0		-
M.2 Flash	C620 Standard SATA AHCI controller	-	4x DMI 3.0	SATA 6G	-

Storage media

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

Model type	Storage medium type	Interface	Form factor
3.5 inch Model	HDD	SAS 12G	3.5 inch, or 2.5 inch *1
		SATA 6G	3.5 inch
	SSD	SAS 12G	2.5 inch *1
		SATA 6G	2.5 inch *1, or M.2
		PCIe 3.0	Add in card
2.5 inch Model	HDD	SAS 12G	2.5 inch
		SATA 6G	2.5 inch
	SSD	SAS 12G	2.5 inch
		SATA 6G	2.5 inch, or M.2
		PCIe 3.0	2.5

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

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

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

Cache settings

In most cases, the cache of HDDs has a great influence on disk I/O performance. It is frequently regarded as a security problem in case of power failure and is thus switched off. On the other hand, it was integrated by hard disk manufacturers for the good reason of increasing the write performance. For performance reasons it is therefore advisable to enable the hard disk cache. To prevent data loss in case of power failure you are recommended to equip the system with a UPS.

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

Performance values

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

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

Storage media performance

3.5 inch model

HDDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
1,800	AL15SEB18EQ	SAS 12G	600	512	547	258	255
300	AL15SEB030N	SAS 12G	645	546	568	231	230
300	ST300MP0006	SAS 12G	768	662	472	304	304
12,000	HUH721212AL5204	SAS 12G	396	339	364	245	244
2,000	ST2000NM0045	SAS 12G	376	336	343	206	206
6,000	ST6000NM0115	SATA 6G	392	362	371	213	208
12,000	HUH721212ALE604	SATA 6G	350	313	341	246	246
1,000	HUS722T1TALA604	SATA 6G	287	264	269	201	201
2,000	ST2000NM0055	SATA 6G	339	301	314	196	195

SSDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
240	MZ7KH240HAHQ	SATA 6G	49,159	7,313	7,431	526	486
480	MZ7KH480HAHQ	SATA 6G	50,558	7,774	7,810	526	485
960	MZ7KH960HAJR	SATA 6G	50,647	7,793	7,916	525	485
1,920	MZ7KH1T9HAJR	SATA 6G	50,702	8,040	7,960	526	485
3,840	MZ7KH3T8HALS	SATA 6G	50,766	8,039	7,936	526	485
240	MTFDDAK240TCB	SATA 6G	18,959	3,367	4,516	487	258
480	MTFDDAK480TDC	SATA 6G	24,710	3,799	5,006	507	362
960	MTFDDAK960TDC	SATA 6G	30,152	4,625	5,553	507	440
1,920	MTFDDAK1T9TDC	SATA 6G	37,234	5,606	5,566	507	483
3,840	MTFDDAK3T8TDC	SATA 6G	41,711	6,429	6,133	504	481
7,680	MTFDDAK7T6TDC	SATA 6G	40,683	6,874	6,672	469	482
375	SSDPED1K375GA	PCIe3 x4	212,118	37,121	36,123	2,460	2,197
750	SSDPED1K750GA	PCIe3 x4	209,628	37,592	36,941	2,546	2,296
240	MTFDDAV240TCB	SATA 6G	20,113	3,936	5,021	510	271
480	MTFDDAV480TCB	SATA 6G	22,596	4,993	6,331	509	403

2.5 inch model

HDDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
600	AL15SEB06EQ	SAS 12G	592	516	544	260	260
300	AL15SEB030N	SAS 12G	645	546	568	231	230
300	ST300MP0006	SAS 12G	768	662	472	304	304
1,000	ST1000NX0453	SAS 12G	371	321	306	137	137
1,000	ST1000NX0313	SATA 6G	324	281	288	131	131
2,000	ST2000NX0403	SATA 6G	326	286	294	133	133

SSDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
240	MZ7KH240HAHQ	SATA 6G	49,159	7,313	7,431	526	486
480	MZ7KH480HAHQ	SATA 6G	50,558	7,774	7,810	526	485
960	MZ7KH960HAJR	SATA 6G	50,647	7,793	7,916	525	485
1,920	MZ7KH1T9HAJR	SATA 6G	50,702	8,040	7,960	526	485
3,840	MZ7KH3T8HALS	SATA 6G	50,766	8,039	7,936	526	485
240	MTFDDAK240TCB	SATA 6G	18,959	3,367	4,516	487	258
480	MTFDDAK480TDC	SATA 6G	24,710	3,799	5,006	507	362
960	MTFDDAK960TDC	SATA 6G	30,152	4,625	5,553	507	440
1,920	MTFDDAK1T9TDC	SATA 6G	37,234	5,606	5,566	507	483
3,840	MTFDDAK3T8TDC	SATA 6G	41,711	6,429	6,133	504	481
7,680	MTFDDAK7T6TDC	SATA 6G	40,683	6,874	6,672	469	482
750	SSDPE21K750GA	PCIe3 x4	214,231	37,611	36,957	2,546	2,295
1,600	SSDPE2KE016T8	PCIe3 x4	135,500	41,066	37,080	3,213	1,917
3,200	SSDPE2KE032T8	PCIe3 x4	136,782	48,210	45,348	3,209	2,800
6,400	SSDPE2KE064T8	PCIe3 x4	192,245	51,767	51,438	3,205	3,048
375	SSDPED1K375GA	PCIe3 x4	212,118	37,121	36,123	2,460	2,197
750	SSDPED1K750GA	PCIe3 x4	209,628	37,592	36,941	2,546	2,296
240	MTFDDAV240TCB	SATA 6G	20,113	3,936	5,021	510	271
480	MTFDDAV480TCB	SATA 6G	22,596	4,993	6,331	509	403

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) is determined as the unit of measurement for the system.

In contrast to benchmarks such as SPECint and TPC-E, which were standardized by independent bodies and for which adherence to the respective rules and regulations is monitored, OLTP-2 is an internal benchmark developed by Fujitsu. OLTP-2 is based on the well-known database benchmark TPC-E. It has been designed in such a way that a wide range of configurations can be measured to present the scaling performance 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. However, direct comparison, or even referring to the OLTP-2 result as TPC-E, is not permitted, in particular 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:

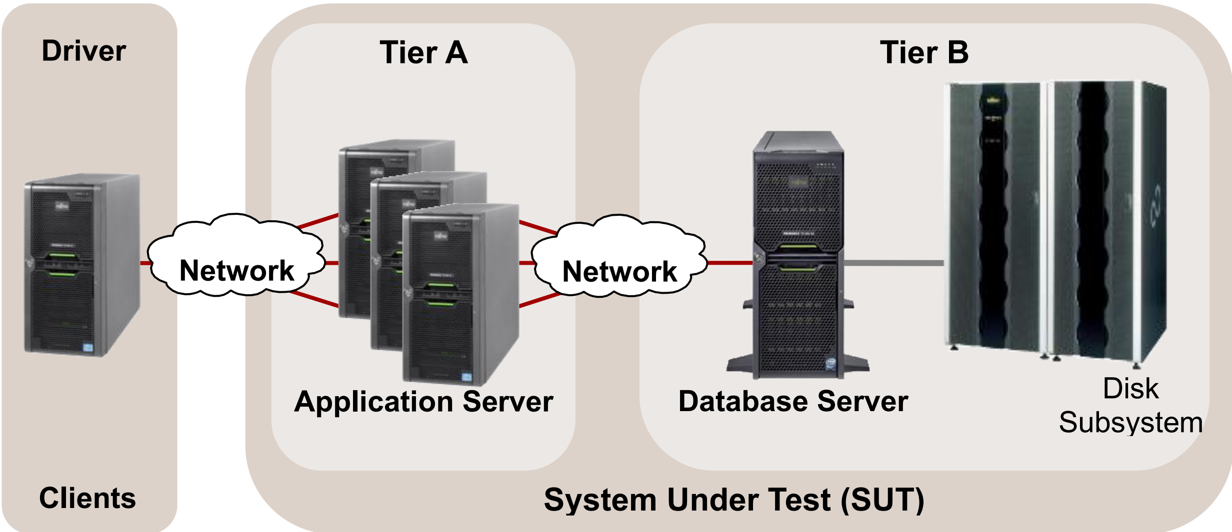

All OLTP-2 results were calculated based on the configuration of the next following pages of PRIMERGY RX2540 M5.

Database Server (Tier B)	
Hardware	
• Model	PRIMERGY RX2540 M5
• Processor	2nd Generation Intel Xeon Processor Scalable Family
• Memory	1 processor: 12 ×64 GB (1 x 64 GB) 2Rx4 DDR4-2933 ECC 2 processors: 24 ×64 GB (1 x 64 GB) 2Rx4 DDR4-2933 ECC
• Network interface	1 × Dual port onboard LAN 10 Gb/s
• Disk subsystem	PRIMERGY RX2520 M5RX2540 M5: Onboard RAID controller PRAID EP420i 2 × 300 GB 10 krpm SAS Drive, RAID 1 (OS), 6 × 1.6 TB SSD, RAID 10 (LOG) 4 × 1.6 TB SSD, RAID 10 (temp) 5 × PRAID EP540e 5 × JX40 S2 : 9 × 1.6 TB SSD Drive each, RAID5 (data)
Software	
• BIOS	Version R1.2.0
• Operating system	Microsoft Windows Server 2016 Standard + KB4462928
• Database	Microsoft SQL Server 2017 Enterprise + KB4341265

Application Server (Tier A)	
Hardware	
• Model	1 × PRIMERGY RX2530 M4
• Processor	2 × Xeon Platinum 8180
• Memory	192 GB, 2666 MHz Registered ECC DDR4
• Network interface	1 × Dual Port onboard LAN 10 Gb/s 1 × Dual Port LAN 1 Gb/s
• Disk subsystem	2 × 300 GB 10 krpm SAS Drive
Software	
• Operating system	Microsoft Windows Server 2016 Standard

Client	
Hardware	
• Model	1 × PRIMERGY RX2530 M2
• Processor	2 × Xeon E5-2667 v4
• Memory	128 GB, 2400 MHz registered ECC DDR4
• Network interface	1 × onboard Quad Port LAN 1 Gb/s
• Disk subsystem	1 × 300 GB 10 krpm 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 or sales regions.

Benchmark results

Database performance greatly depends on the configuration options with the CPU and 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 have been appropriately chosen and are not bottlenecks.

A guideline for the database environment for selecting the main memory is that sufficient quantity is more important than the speed of memory access. This why a configuration with a total memory of 768 GB was considered for the measurements with two processors and a configuration with a total memory of 384 GB for the measurements with one processor. Both the memory configurations had a memory access of 2933 MHz.

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

Processor	Cores	Threads	2CPU Score		1CPU Score	
April 2019 released						
Xeon Gold 5222	4	8	1,212	est.	666	est.
Xeon Gold 5220	18	36	3,702	est.	2,034	est.
Xeon Gold 5218	16	32	3,369	est.	1,851	est.
Xeon Gold 5217	8	16	1,992	est.	1,095	est.
Xeon Gold 5215	10	20	2,204	est.	1,211	est.
Xeon Silver 4216	16	32	3,057	est.	1,680	est.
Xeon Silver 4215	8	16	1,790	est.	984	est.
Xeon Silver 4214Y	12	24	2,391	est.	1,314	est.
	10	20	1,993	est.	1,095	est.
	8	16	1,721	est.	946	est.
Xeon Silver 4214	12	24	2,391	est.	1,314	est.
Xeon Silver 4210	10	20	2,038	est.	1,120	est.
Xeon Silver 4208	8	16	1,537	est.	844	est.
Xeon Bronze 3204	6	6	744	est.	409	est.
March 2020 released						
Xeon Gold 5218R	20	40	3,856	est.	2,119	est.
Xeon Silver 4215R	8	16	1,874	est.	1,030	est.
Xeon Silver 4214R	12	24	2,406	est.	1,322	est.
Xeon Silver 4210R	10	20	2,041	est.	1,122	est.
Xeon Bronze 3206R	8	8	965	est.	530	est.

The following diagram shows the OLTP-2 transaction rates that can be achieved with the 2nd Generation Intel Xeon Processor Scalable Family.

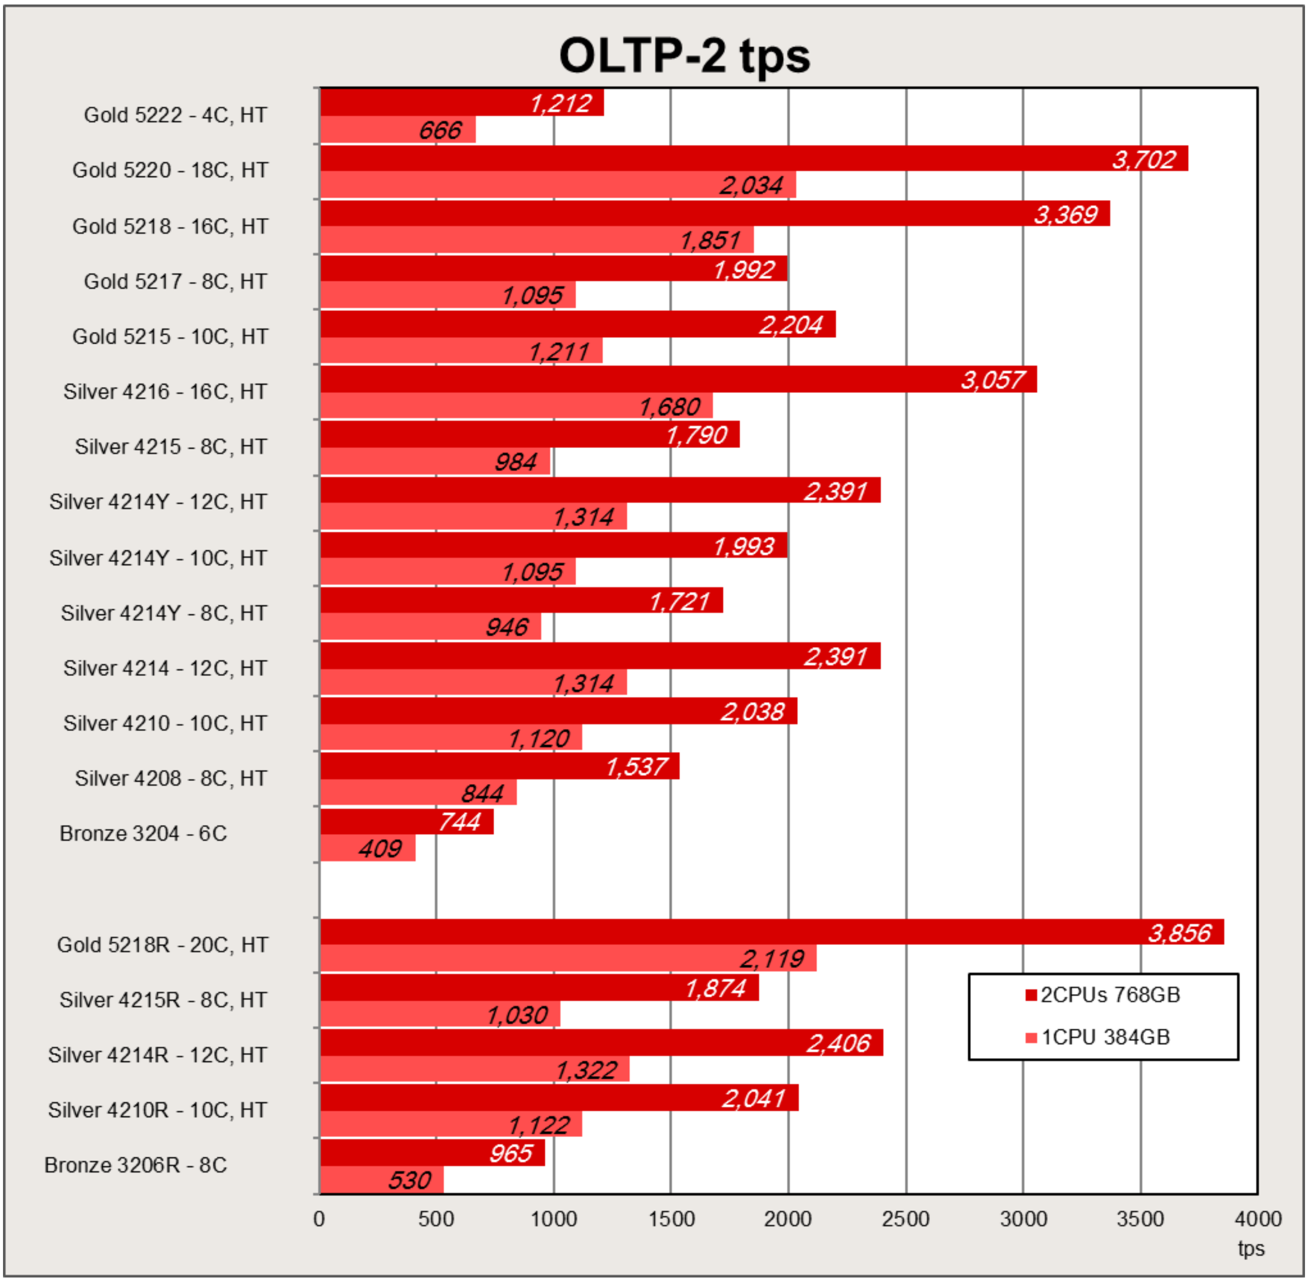

It is evident that a wide performance range is covered by a variety of released processors. If you compare the OLTP-2 value for the processor with the lowest performance (Xeon Bronze 3204) with the value for the processor with the highest performance (Xeon Platinum 8280), the result is a 5-fold increase in performance.

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.

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

vServCon

Benchmark description

vServCon is a benchmark used by Fujitsu to compare server configurations with hypervisor with regard to their suitability for server consolidation. This allows both the comparison of systems, processors and I/O technologies as well as the comparison of hypervisors, virtualization forms, and additional drivers for virtual machines.

vServCon is not a new benchmark in the true sense of the word. It is more a framework that combines already established benchmarks (or in modified form) as workloads in order to reproduce the load of a consolidated and virtualized server environment. Three proven benchmarks are used which cover the application scenarios database, application server, and web server.

Application scenario	Benchmark	No. of logical CPU cores	Memory
Database	Sysbench (adapted)	2	1.5 GB
Java application server	SPECjbb (adapted, with 50% - 60% load)	2	2 GB
Web server	WebBench	1	1.5 GB

Each of the three application scenarios is allocated to a dedicated virtual machine (VM). A fourth machine, the so-called idle VM, is added to these. These four VMs make up a "tile". Depending on the performance capability of the underlying server hardware, you may as part of a measurement also have to start several identical tiles in parallel in order to achieve a maximum performance score.

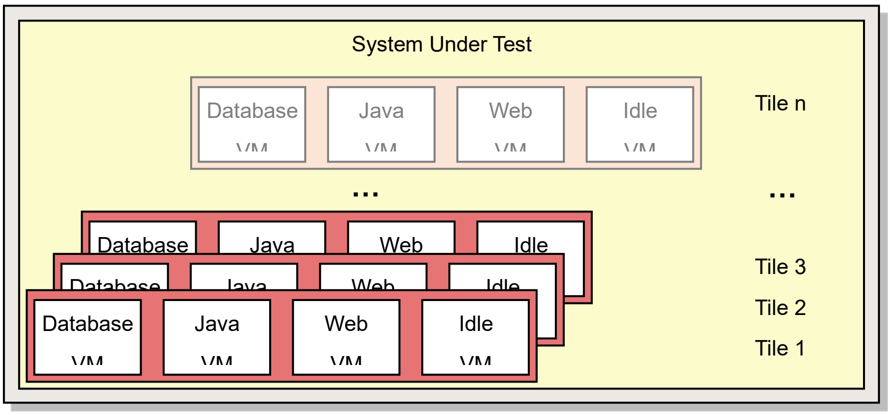

Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score, the individual benchmark result for one tile is put in relation to the respective result of a reference system. The resulting relative performance value is then suitably weighted and finally added up for all VMs and tiles. The outcome is a score for this tile number.

As a general rule, start with one tile, and this procedure is performed for an increasing number of tiles until no further significant increase in this vServCon score occurs. The final vServCon score is then the maximum of the vServCon scores for all tile numbers. This score thus reflects the maximum total throughput that can be achieved by running the mix defined in vServCon that consists of numerous VMs up to the possible full utilization of CPU resources. This is why the measurement environment for vServCon measurements is designed in such a way that only the CPU is the limiting factor and that no limitations occur as a result of other resources.

The progression of the vServCon scores for the tile numbers provides useful information about the scaling behavior of the "System under Test".

A detailed description of vServCon is in the document: Benchmark Overview vServCon.

Benchmark environment

The typical measurement set-up is illustrated below:

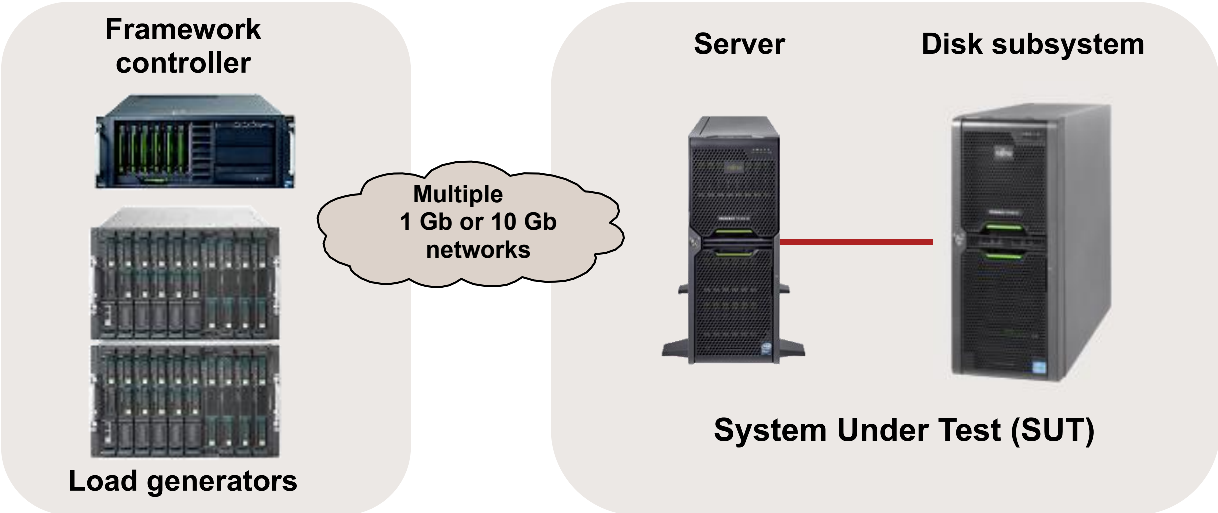

All vServCon results were Calculated based on the configuration of the next following pages of PRIMERGY RX2540 M5.

System Under Test (SUT)	
Hardware	
• Processor	2nd Generation Intel Xeon Scalable Processors Family
• Memory	24 × 32 GB (1x32 GB) 2Rx4 DDR4-2933 R ECC
• Network interface	2 × Intel Ethnet Controller X710 for 10GbE SFP+
• Disk subsystem	1 ×dual-channel FC controller Emulex LPe160021 LINUX/LIO based flash storage system
Software	
• Operating system	VMware ESXi 6.7 EP06 Build 11675023

Load generator (incl. Framework controller)	
Hardware (Shared)	
• Enclosure	5 × PRIMERGY RX2530 M2
Hardware	
• Processor	2 × XeonE5-2683 v4
• Memory	128 GB
• Network interface	3 × 1 Gbit LAN
Software	
• Operating system	VMware ESXi 6.0.0 U1b Build 3380124

Load generator VM (on various servers)	
Hardware	
• Processor	1 × logical CPU
• Memory	4048 MB
• Network interface	2 × 1 Gbit/s LAN
Software	
• Operating system	Microsoft Windows Server 2008 Standard Edition 32 bit

Benchmark results

The PRIMERGY dual-socket rack and tower systems dealt with here are based on processors of the 2nd Generation Intel Xeon Processor. The features of the processors are summarized in the section “Technical data”.

The available processors of these systems with their results can be seen in the following table.

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

Processor	Cores	Threads	Number of Processors	#Tiles		Score	
April 2019 released							
Xeon Gold 5222	4	8	2	7	est.	11.2	est.
Xeon Gold 5220	18	36	2	26	est.	36.3	est.
Xeon Gold 5218	16	32	2	23	est.	32.9	est.
Xeon Gold 5217	8	16	2	14	est.	19.5	est.
Xeon Gold 5215	10	20	2	18	est.	20.5	est.
Xeon Silver 4216	16	32	2	25	est.	29.9	est.
Xeon Silver 4215	8	16	2	14	est.	17.5	est.
Xeon Silver 4214Y	12	24	2	19	est.	23.8	est.
	10	20	2	16	est.	20.5	est.
	8	16	2	12	est.	17.9	est.
Xeon Silver 4214	12	24	2	19	est.	23.8	est.
Xeon Silver 4210	10	20	2	18	est.	20	est.
Xeon Silver 4208	8	16	2	13	est.	15.1	est.
Xeon Bronze 3204	6	6	2	7	est.	7.2	est.
March 2020 released							
Xeon Gold 5218R	20	40	2	25	est.	37.7	est.
Xeon Silver 4215R	8	16	2	15	est.	18.3	est.
Xeon Silver 4214R	12	24	2	19	est.	23.6	est.
Xeon Silver 4210R	10	20	2	18	est.	20	est.
Xeon Bronze 3206R	8	8	2	9	est.	9.5	est.

The following diagram compares the virtualization performance values that can be achieved with the processors reviewed here.

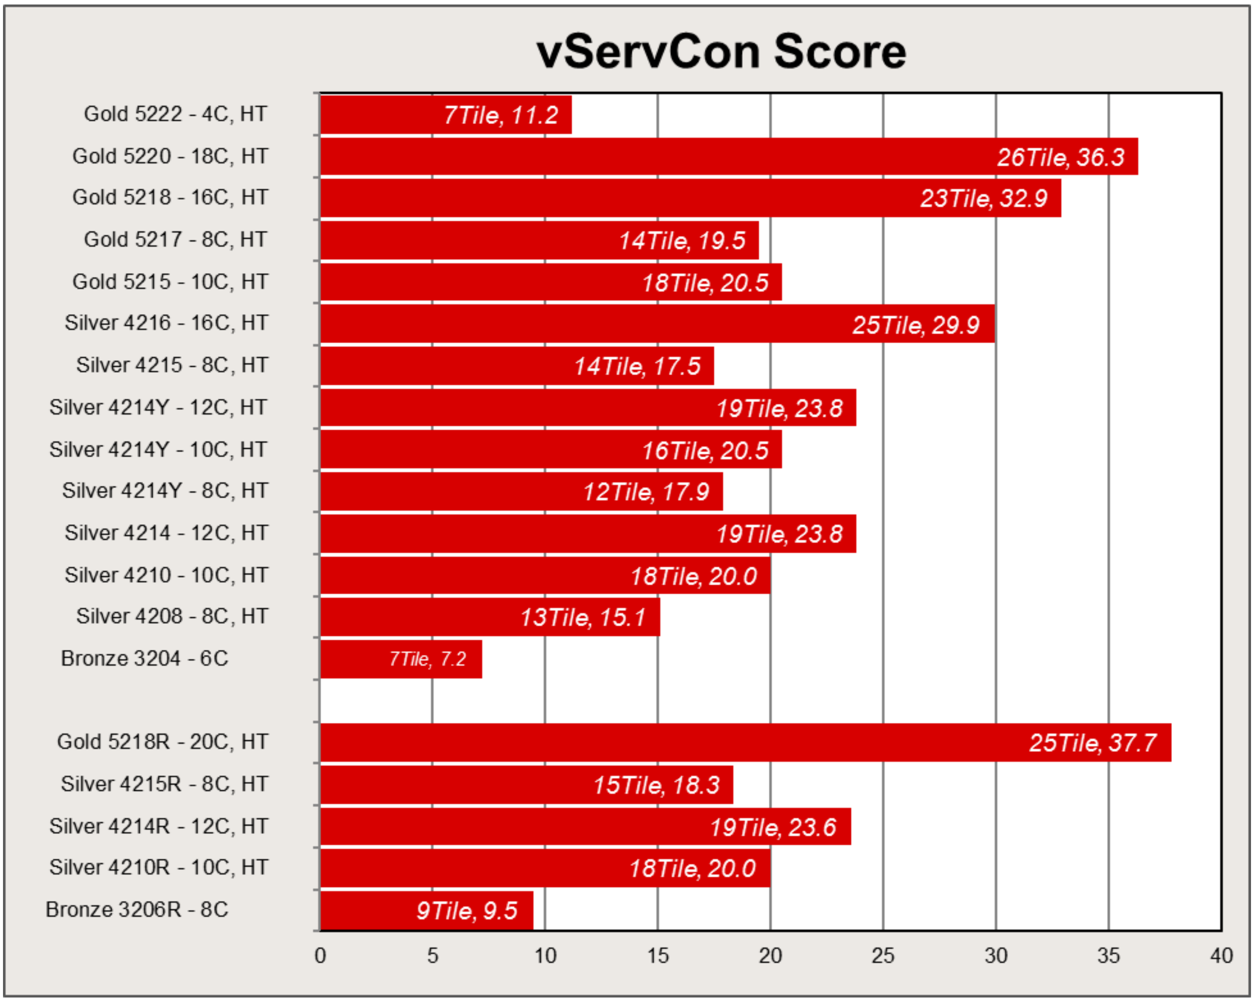

The relatively large performance differences between the processors can be explained by their features. The values scale on the basis of the number of cores, the size of the L3 cache and the CPU clock frequency and as a result of the features of Hyper-Threading and turbo mode, which are available in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

Within a group of processors with the same number of cores scaling can be seen via the CPU clock frequency.
A low performance can be seen in the Xeon Bronze 3204 processor, as they have to manage without Hyper-Threading (HT) and turbo mode (TM). In principle, these weakest processors are only to a limited extent suitable for the virtualization environment.

As a matter of principle, the memory access speed also influences performance. A guideline in the virtualization environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses. The vServCon scaling measurements presented here were all performed with a memory access speed – depending on the processor type – of at most 2933 MHz.

Literature

PRIMERGY Servers

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

PRIMERGY RX2520 M5

This Whitepaper

https://docs.ts.fujitsu.com/dl.aspx?id=188e42f9-9a54-4a1d-a8ee-b71c94d5cbeb

https://docs.ts.fujitsu.com/dl.aspx?id=95c8a62b-e41b-406e-8dc3-c1c226532ea2

Data sheet

https://docs.ts.fujitsu.com/dl.aspx?id=e65225a5-06be-4e16-a4c1-5e7dac6e9f6a

PRIMERGY Performance

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

SPEC CPU2017

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

Benchmark Overview SPECcpu2017

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

STREAM

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

SPECpower_ssj2008

https://www.spec.org/power_ssj2008

Benchmark Overview SPECpower_ssj2008

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

OLTP-2

Benchmark Overview OLTP-2

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

vServCon

Benchmark Overview vServCon

https://docs.ts.fujitsu.com/dl.aspx?id=b953d1f3-6f98-4b93-95f5-8c8ba3db4e59

Document change history

Version	Date	Description
1.5	2023-10-03	Updated: • New Visual Identity format
1.4	2021-07-28	Updated: • Contact information and URLs Updated to the latest one • Minor correction
1.3	2020-05-29	Updated: • Technical data, STREAM Fixed typo in processor specifications
1.2	2020-04-24	New: • LINPACK Measured and calculated with 2nd Generation Intel Xeon Processor Scalable Family Updated: • Technical data Added 2nd Generation Intel Xeon Processor Scalable Family • SPECcpu2017 Measured and calculated with 2nd Generation Intel Xeon Processor Scalable Family • OLTP-2 Calculated with 2nd Generation Intel Xeon Processor Scalable Family • vServCon Calculated with 2nd Generation Intel Xeon Processor Scalable Family • STREAM Measured and calculated with 2nd Generation Intel Xeon Processor Scalable Family
1.1	2019-12-27	New • SPECpower_ssj2008 Measured with Intel Xeon Gold 5220 • Disk I/O: Performance of storage media Results for 2.5" and 3.5" storage media • STREAM Measured with 2nd Generation Intel Xeon Processor Scalable Family

Version	Date	Description
1.0	2019-07-26	New • Technical data • SPECcpu2017 Measured with 2nd Generation Intel Xeon Processor Scalable Family • OLTP-2 Calculated with 2nd Generation Intel Xeon Processor Scalable Family • vServCon Calculated with 2nd Generation Intel Xeon Processor Scalable Family

Contact

Fujitsu

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

PRIMERGY Performance and Benchmarks

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