White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY RX2540 M6

This document provides an overview of benchmarks executed on the FUJITSU Server PRIMERGY RX2540 M6.

Explains PRIMERGY RX2540 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
1.1
2021/07/21

Contents

Document history

Version 1.0 (2021/06/17)

New:

- Technical data
- SPECcpu2017, OLTP-2, STREAM, LINPACK
  Measured and calculated with 3rd Generation Intel Xeon Processor Scalable Family
- SPECpower_ssj2008, SAP SD, VMmark V3
  Measured with Intel Xeon Platinum 8380
- Disk I/O
  Measured with 2.5", 3.5", and EDSSF storage media

Version 1.1 (2021/07/21)

Updated:

- SPECcpu2017, OLTP-2, STREAM, LINPACK
  Measured and calculated with 3rd Generation Intel Xeon Processor Scalable Family

Technical data

PRIMERGY RX2540 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 RX2540 M6
Form factor	Rack server
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	32 (16 per processor)
Maximum memory configuration	12,288 GB
Onboard HDD controller	Controller with RAID (0, 1, 10) function (supports up to 8 SATA HDD/SSD)
PCI slots	PCI-Express 4.0 (x16 lane): 4 (Low Profile) PCI-Express 4.0 (x8 lane) : 2 (Low Profile)
Max. number of internal storage	30 x 2.5 inches or 12 x 3.5 inches or 64 x EDSFF

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 Platinum 8380	40	80	60	11.2	2.3	3.4	3200	270
Xeon Platinum 8368Q	38	76	57	11.2	2.6	3.7	3200	270
Xeon Platinum 8368	38	76	57	11.2	2.4	3.4	3200	270
Xeon Platinum 8362*1	32	64	48	11.2	2.8	3.6	3200	265
Xeon Platinum 8360Y	36	72	54	11.2	2.4	3.5	3200	250
Xeon Platinum 8358P	32	64	48	11.2	2.6	3.4	3200	240
Xeon Platinum 8358	32	64	48	11.2	2.6	3.4	3200	250
Xeon Platinum 8352Y	32	64	48	11.2	2.2	3.4	3200	205
Xeon Platinum 8352V	36	72	54	11.2	2.1	3.5	2933	195
Xeon Platinum 8352M*1	32	64	48	11.2	2.3	3.5	3200	185
Xeon Gold 6354	18	36	39	11.2	3.0	3.6	3200	205
Xeon Gold 6348	28	56	42	11.2	2.6	3.5	3200	235
Xeon Gold 6346	16	32	36	11.2	3.1	3.6	3200	205
Xeon Gold 6342*1	24	48	36	11.2	2.8	3.5	3200	230
Xeon Gold 6338T*1	24	48	36	11.2	2.1	3.4	2933	165
Xeon Gold 6338	32	64	48	11.2	2.0	3.2	3200	205
Xeon Gold 6336Y*1	24	48	36	11.2	2.4	3.6	3200	185
Xeon Gold 6334*1	8	16	18	11.2	3.6	3.7	3200	165
Xeon Gold 6330N	28	56	42	11.2	2.2	3.4	2667	165
Xeon Gold 6330	28	56	42	11.2	2.0	3.1	2933	205
Xeon Gold 6326*1	16	32	24	11.2	2.9	3.5	3200	185
Xeon Gold 6314U	32	64	48	11.2	2.3	3.4	3200	205
Xeon Gold 6312U*1	24	48	36	11.2	2.4	3.6	2933	185
Xeon Gold 5320*1	26	52	39	11.2	2.2	3.4	2933	185
Xeon Gold 5318Y*1	24	48	36	11.2	2.1	3.4	2933	165
Xeon Gold 5318S*1	24	48	36	11.2	2.1	3.4	2933	165
Xeon Gold 5317*1	12	24	18	11.2	3.0	3.6	2933	150
Xeon Gold 5315Y*1	8	16	12	11.2	3.2	3.6	2933	140
Xeon Silver 4316*1	20	40	30	10.4	2.3	3.4	2666	150
Xeon Silver 4314*1	16	32	24	10.4	2.4	3.4	2666	135
Xeon Silver 4310*1	12	24	18	10.4	2.1	3.3	2666	120
Xeon Silver 4309Y*1	8	16	12	10.4	2.8	3.6	2666	105

*1: To be supported in July 2021 or later

All processors that can be ordered with PRIMERGY RX2540 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 optmized for the following feature.

Suffix	Additional feature
M	Processor which is optimized for media, AI, and HPC
N	Processor which is optimized for NFV usage
P	Processor which is optimized for IaaS (orchestration efficiency targeting higher frequency for VM Market)
Q	Processor which requires liquid cooling
S	Processor which supports max SGX enclave size of 512GB
T	Processor with higher Tcase, which is designed for Network Environment-Building system environments and supports for up to 10-year reliability
U	Processor which is optimized for single socket
V	Processor which is optimized for SaaS (orchestration efficiency targeting high density, lower power VM environment)
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	3200			✓		✓
16 GB (1x 16 GB) 2Rx8 DDR4-3200 R ECC	16	2	8	3200			✓		✓
16 GB (1x 16 GB) 1Rx4 DDR4-3200 R ECC	16	1	4	3200			✓		✓
32 GB (1x 32 GB) 2Rx4 DDR4-3200 R ECC	32	2	4	3200			✓		✓
64 GB (1x 64 GB) 2Rx4 DDR4-3200 R ECC	64	2	4	3200			✓		✓
64 GB (1x 64 GB) 2Rx4 DDR4-3200 LR ECC	64	2	4	3200		✓	✓		✓
128 GB (1x128 GB) 4Rx4 DDR4-3200 3DS ECC	128	4	4	3200	✓		✓		✓
128 GB (1x128 GB) 4Rx4 DDR4-3200 LR ECC	128	4	4	3200		✓	✓		✓
256 GB (1x256 GB) 8Rx4 DDR4-3200 3DS ECC	256	8	4	3200	✓		✓		✓
128 GB (1x128GB) Optane PMem-3200	128			3200				✓	✓
256 GB (1x256GB) Optane PMem-3200	256			3200				✓	✓
512 GB (1x512GB) Optane PMem-3200	512			3200				✓	✓

Power supplies	Maximum number
Modular PSU 500W platinum	2
Modular PSU 900W platinum	2
Modular PSU 1600W platinum	2
Modular PSU 2200W platinum	2
Modular PSU 900W Titanum	2

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

Detailed technical information is available in the data sheet of PRIMERGY RX2540 M6.

SPECcpu2017

Benchmark description

SPECcpu2017 is a benchmark which measures the system efficiency with integer and floating-point operations. It consists of an integer test suite (SPECrate 2017 Integer, 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.

SPECcpu2017 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.

SPECcpu2017 contains two different performance measurement methods. The first method (SPECspeed 2017 Integer or SPECspeed 2017 Floating Point) determines the time which is required to process a single task. The second method (SPECrate 2017 Integer or SPECrate 2017 Floating Point) determines the throughput, i.e. the number of tasks that can be handled in parallel. 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 SPECcpu2017 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 RX2540 M6
Processor	2 x 3rd Generation Intel® Xeon® Scalable Processors Family
Memory	32 x 32 GB 2Rx4 PC4-3200AA-R
Software	
BIOS settings	SPECspeed2017_int: CPU C1E Support = Disabled UPI Link Frequency Select = 10.4GT/s SPECSpeed2017_fp: Hyper Threading = Disabled DCU Streamer Prefetcher = Disabled Override OS energy Performance = Enabled Energy Performance = Performance FAN Control = Full SPECrate2017_int: DCU Streamer Prefetcher = Disabled CPU C1E Support = Disabled Package C State Limit = C2 UPI Link Frequency Select = 10.4GT/s XPT Prefetch = Enable LLC Prefetch = Enabled UPI Prefetch = Disabled Sub NUMA (SNC) =Enable SNC2 Fan Control = Ful SPECrate2017_fp: Hyper Threading = Disabled Adjacent Cache Line Prefetch = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Override OS energy Performance = Enabled Energy Performance = Performance CPU C1E Support = Disabled Patrol Scrub = Enabled Sub NUMA (SNC) = Enable SNC2
Operating system	SPECrate2017_fp: SUSE Linux Enterprise Server 15 SP2 5.3.18-22-default SPECrate2017_int , SPECspeed2017_int, SPECspeed2017_fp: Red Hat Enterprise Linux Server release 8.2 4.18.0-193.el8.x86_64
Operating system settings	Stack size set to unlimited using "ulimit -s unlimited" SPECspeed2017_int Sched_migration_cost_ns = 1000 SPECspeed2017_fp Sched_latency_ns = 16000000 SPECrate2017_int: Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1)
Compiler	SPECspeed2017, SPECrate2017_int C/C++: Version 19.1.2.275 of Intel C/C++ Compiler for Linux Fortran: Version 19.1.2.275 of Intel Fortran Compiler for Linux SPECrate2017_fp 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.

SPECrate2017				
Processor	**Number of cores**	**Number of processors**	**SPECrate2017 int_base**	**SPECrate2017 fp_base**
Xeon Platinum 8380	40	2	**533**	**463**
Xeon Platinum 8368Q	38	2	**518**	**459**
Xeon Platinum 8368	38	2	**507**	**449**
Xeon Platinum 8362*1	32	2	**465**	**429**
Xeon Platinum 8360Y	36	2	**484**	**438**
Xeon Platinum 8358P	32	2	**439**	**413**
Xeon Platinum 8358	32	2	**455**	**425**
Xeon Platinum 8352Y	32	2	**414**	**396**
Xeon Platinum 8352V	36	2	**425**	**391**
Xeon Platinum 8352M*1	32	2	**389**	**378**
Xeon Gold 6354	18	2	**297**	**314**
Xeon Gold 6348	28	2	**409**	**393**
Xeon Gold 6346	16	2	**275**	**294**
Xeon Gold 6342*1	24	2	**373**	**368**
Xeon Gold 6338T*1	24	2	**310**	**323**
Xeon Gold 6338	32	2	**398**	**386**
Xeon Gold 6336Y*1	24	2	**338**	**343**
Xeon Gold 6334*1	8	2	**145**	**172**
Xeon Gold 6330N	28	2	**332**	**325**
Xeon Gold 6330	28	2	**360**	**349**
Xeon Gold 6326*1	16	2	**260**	**277**
Xeon Gold 6314U	32	2	**214**	**205**
Xeon Gold 6312U*1	24	2	**175**	**179**
Xeon Gold 5320*1	26	2	**350**	**345**
Xeon Gold 5318Y*1	24	2	**312**	**317**
Xeon Gold 5318S*1	24	2	**311**	**313**
Xeon Gold 5317*1	12	2	**195**	**212**
Xeon Gold 5315Y*1	8	2	**138**	**155**
Xeon Silver 4316*1	20	2	**270**	**278**
Xeon Silver 4314*1	16	2	**223**	**238**
Xeon Silver 4310*1	12	2	**167**	**189**
Xeon Silver 4309Y*1	8	2	**125**	**139**

*1: To be supported in July 2021 or later

SPECspeed2017				
Processor	**Number of cores**	**Number of processors**	**SPECspeed2017 int_base**	**SPECspeed2017 fp_base**
Xeon Platinum 8380	40	2	**12.0**	**233**
Xeon Platinum 8368Q	38	2	**12.5**	

The following graph compares the throughput of PRIMERGY RX2540 M6 and its older model, PRIMERGY RX2540 M5, with maximum performance configurations.

SPECrate2017: Comparison of PRIMERGY RX2540 M6 and PRIMERGYRX2540 M5

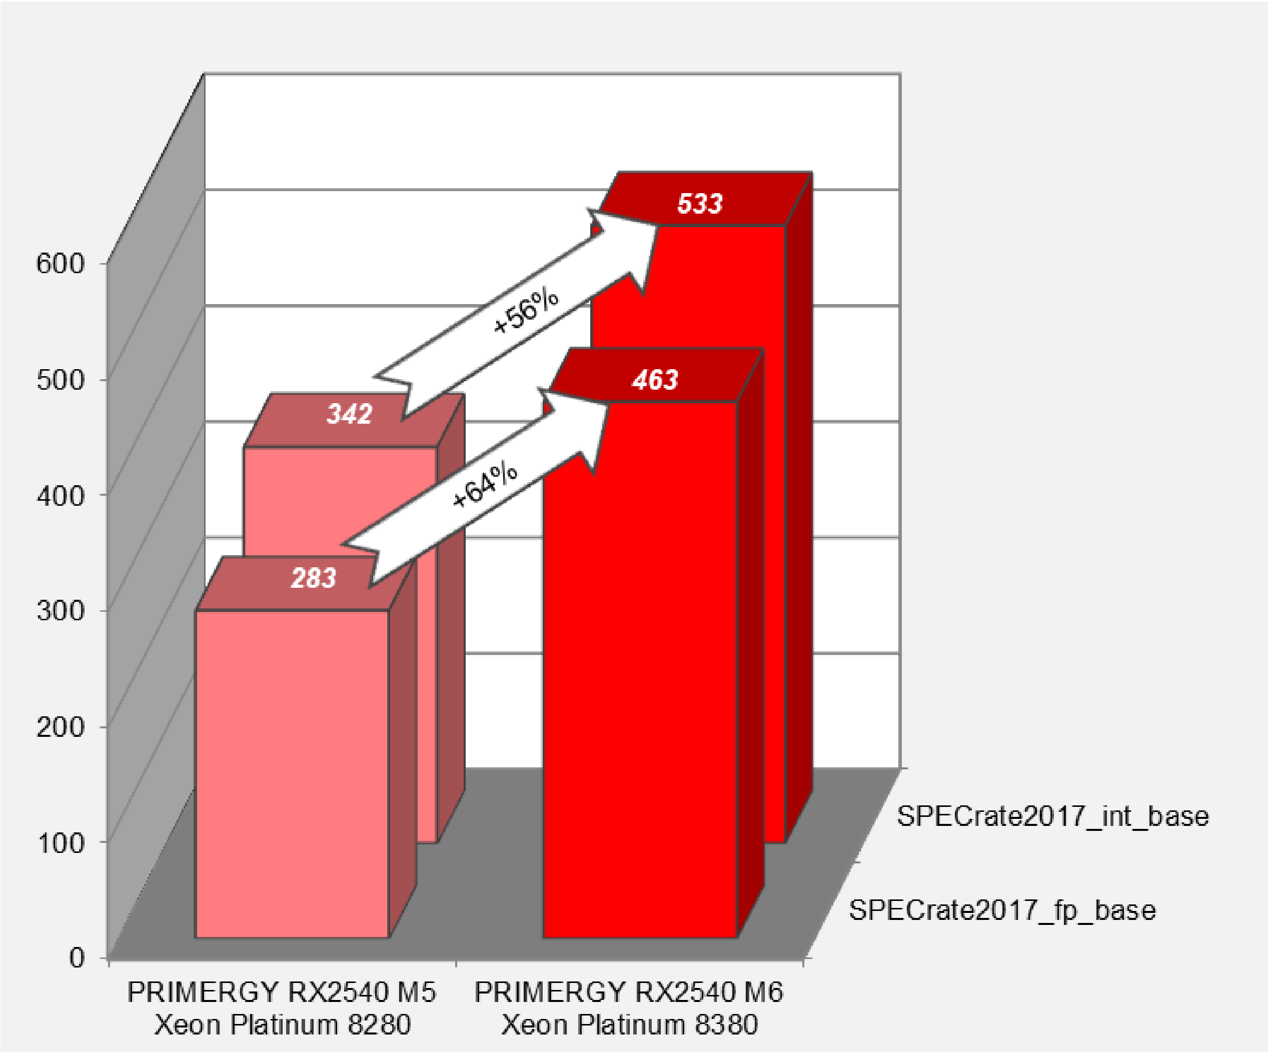

On April 28, 2021, PRIMERGY RX2540 M6 with the Xeon Platinum 8380 processor won first place in the 2-socket Xeon category of the SPECspeed2017_fp_base benchmark.

SPECpower_ssj2008

Benchmark description

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

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

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

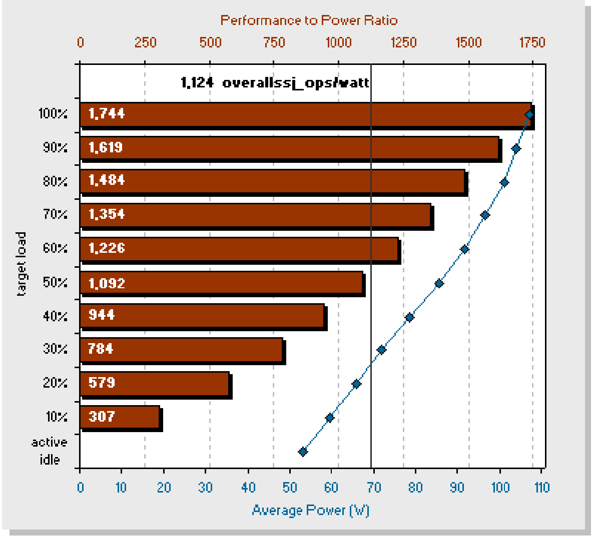

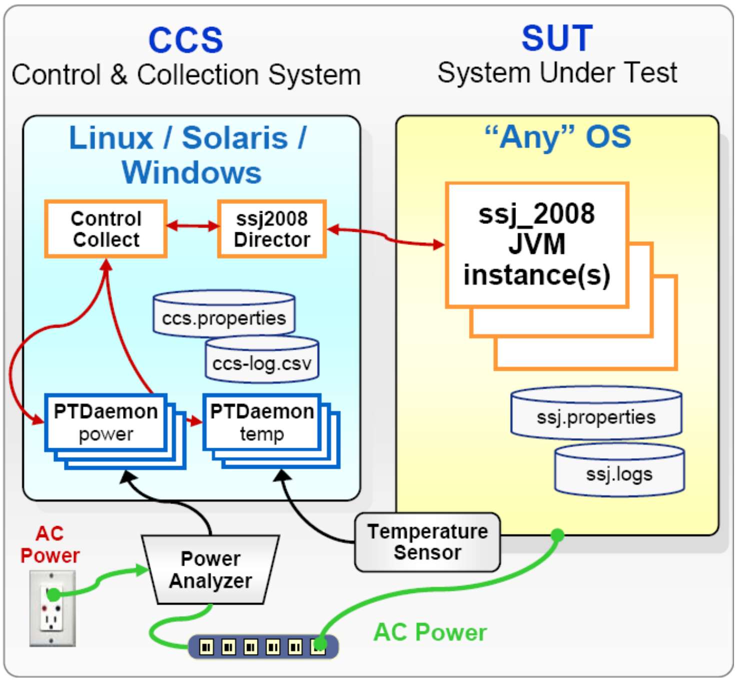

The benchmark runs on a wide variety of operating systems and hardware architectures, and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System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

Benchmark environment

System Under Test (SUT)	
For Linux OS measurement	
Hardware	
Model	PRIMERGY RX2540 M6
Processor	2 x Intel® Xeon® Platinum 8380
Memory	16 x 16 GB 2Rx8 PC4-3200Y-R
Network interface	1 x PLAN CP I350-T4 4X 1000BASE-T OCPv3, PYBLA274U
Disk subsystem	1 x SSD M.2 240GB, S26361-F5787-E240
Power Supply Unit	1 x 900 W, S26113-E629-V50-1
Software	
BIOS	R1.4.0
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 = Power balanced SNC(Sub NUMA) = Enable SNC2 ASPM Support = L1 only SATA Controller = Disabled Serial Port = Disabled USB Port Control = Disable all ports Network Stack = Disabled
iRMC Firmware	3.22P
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 94 > /sys/devices/system/cpu/cpufreq/conservative/up_threshold echo 1 > /sys/devices/system/cpu/cpufreq/conservative/freq_step echo 93 > /sys/devices/system/cpu/cpufreq/conservative/down_threshold echo always > /sys/kernel/mm/transparent_hugepage/enabled cpupower frequency-set -u 2400MHz sysctl -w kernel.nmi_watchdog=0
JVM	Oracle Java HotSpot(TM) 64-Bit Server VM 18.9 (build 11.0.9+7-LTS, mixed mode)
JVM settings	-server -Xmn16000m -Xms18000m -Xmx18000m -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	
Model	PRIMERGY RX2540 M6
Processor	2 x Intel® Xeon® Platinum 8380
Memory	16 x 16 GB 2Rx8 PC4-3200Y-R
Network interface	1 x PLAN CP I350-T4 4X 1000BASE-T OCPv3, PYBLA274U
Disk subsystem	1 x SSD M.2 240GB, S26361-F5787-E240
Power Supply Unit	1 x 900 W, S26113-E629-V50-1
Software	
BIOS	R1.4.0
BIOS settings	Hardware Prefetcher = Disabled Adjacent Cache Line Prefetch = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Turbo Mode = Disabled Energy Performance = Energy Efficient Package C State limit = C6 UPI Link Frequency Select = 9.6GT/s Uncore Frequency Scaling = Power Balanced DDR Performance = Power balanced SNC(Sub NUMA) = Enable SNC2 ASPM Support = L1 only SATA Controller = Disabled Serial Port = Disabled USB Port Control = Disable all ports Network Stack = Disabled
iRMC Firmware	3.22P
Operating system	Microsoft Windows Server 2019 Standard
Operating system settings	Turn off hard disk after = 1 Minute Turn off display after = 1 Minute Minimum processor state = 0% Maximum processor state = 100% Using the local security settings console, "lock pages in memory" was enabled for the user running the benchmark. Benchmark was started via Windows Remote Desktop Connection.
JVM	Oracle Java HotSpot(TM) 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

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

Benchmark results (Linux)

The PRIMERGY RX2540 M6 in SUSE Linux Enterprise Server 12 SP4 achieved the following result:

SPECpower_ssj2008 = 11,652 overall ssj_ops/watt

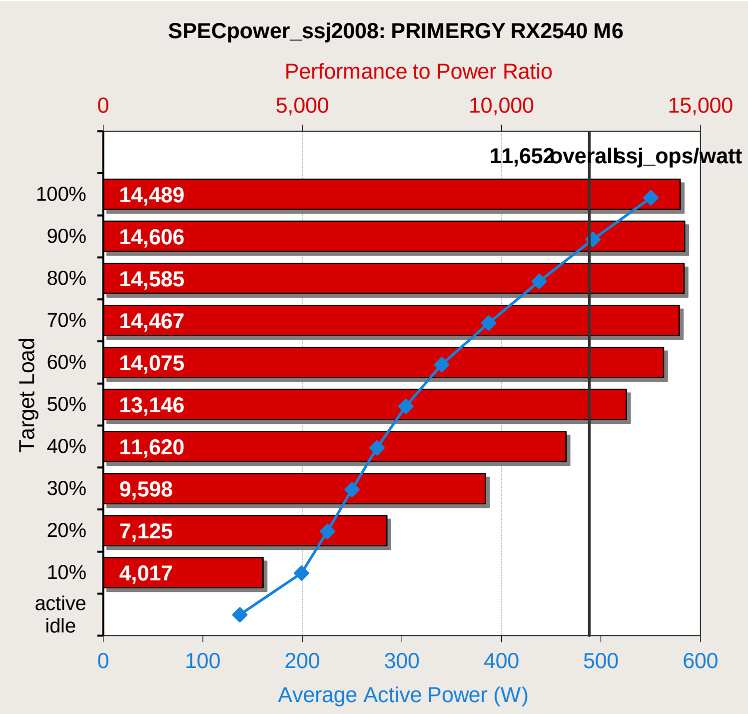

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 11,652 overall ssj_ops/watt for the PRIMERGYRX2540 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%	7,594,068	564	13,469
90%	6,823,877	489	13,968
80%	6,068,791	437	13,880
70%	5,320,946	378	14,077
60%	4,554,273	335	13,590
50%	3,801,003	302	12,602
40%	3,041,484	274	11,109
30%	2,273,651	249	9,135
20%	1,517,770	224	6,771
10%	759,382	198	3,842
Active Idle	0	135	0
			∑ssj_ops / ∑power = 11,652

Benchmark results (Windows)

The PRIMERGY RX2540 M6 in Microsoft Windows Server 2016 Standard achieved the following result:

SPECpower_ssj2008 = 12,211 overall ssj_ops/watt

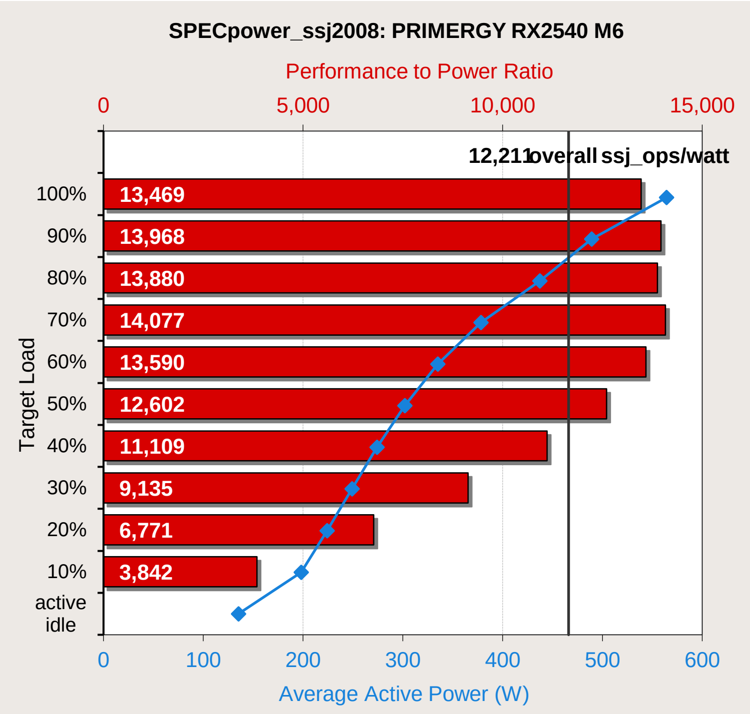

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 12,211 overall ssj_ops/watt for the PRIMERGYRX2540 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%	7,972,718	550	14,489
90%	7,186,063	492	14,606
80%	6,393,133	438	14,585
70%	5,593,093	387	14,467
60%	4,791,610	340	14,075
50%	3,997,496	304	13,146
40%	3,195,994	275	11,620
30%	2,399,228	250	9,598
20%	1,605,244	225	7,125
10%	798,865	199	4,017
Active Idle	0	137	0
			∑ssj_ops / ∑power = 12,211

Comparison with the predecessor

The following diagram shows for each load level (on the x-axis) the throughput (on the left y-axis) and the power consumption (on the right y-axis) of the PRIMERGY RX2540 M6 compared to the predecessor PRIMERGY RX2540 M5 with the JVM versions that affect the SPECpower_ssj2008 benchmark.

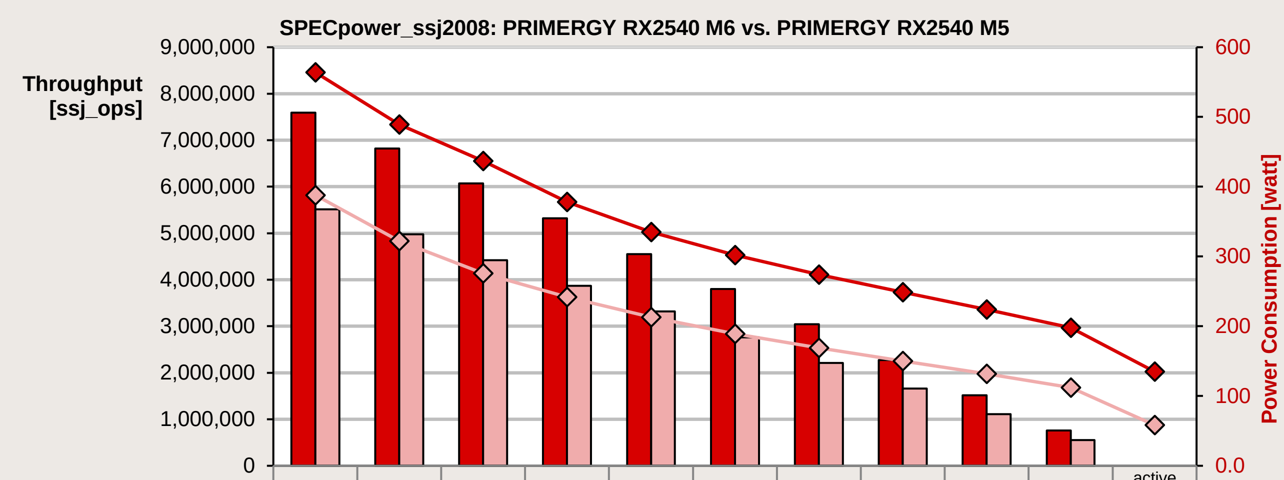

Load Level	100%	90%	80%	70%	60%	50%	40%	30%	20%	10%	active idle
RX2540 M6 ssj_ops	7,594,068	6,823,877	6,068,791	5,320,946	4,554,273	3,801,003	3,041,484	2,273,651	1,517,770	759,382	0
RX2540 M5 ssj_ops	5,514,468	4,979,547	4,421,965	3,873,008	3,321,077	2,759,717	2,211,380	1,660,248	1,109,117	555,330	0
RX2540 M6 watt	564	489	437	378	335	302	274	249	224	198	135
RX2540 M5 watt	388	322	276	242	213	189	169	150	132	112	58.3

The total average throughput of the PRIMERGY RX2540 M6 is 43,933,444 ssj_ops, an improvement of 71.3% over the 25,641,201 ssj_ops of the PRIMERGY RX2540 M5.

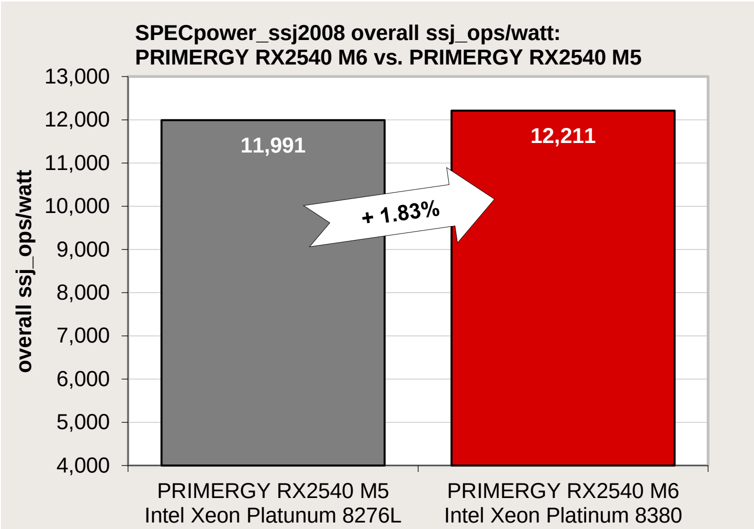

On the other hand, the total average power consumption of the PRIMERGY RX2540 M6 is 3,597 W, which is 68.2% higher than the 2,138 W of the PRIMERGY RX2540 M5. One of the reasons for the increased power consumption is the increased power consumption of the CPU used for the measurement. The TDP of the Intel Xeon Platinum 8276L used in the PRIMERGY RX2540 M5 was 165W, while the TDP of the Intel Xeon Platinum 8380 used in the PRIMERGY RX2540 M6 was 270W, an increase of 63.6%. Another reason is the increased power consumption when idle. It was 58.2W on the PRIMERGY RX2540 M5, but it increased 2.35 times to 137W on the PRIMERGY RX2540 M6.

The overall energy efficiency of the PRIMERGY RX2540 M6 has improved by 1.83% due to a 68.2% increase in power consumption but a 71.3% improvement in performance.

Measurement results of SPECpower_ssj2008 (June 9, 2021)

11,652 SPECpower_ssj2008

On June 9, 2021, PRIMERGY RX2540 M6 with two Xeon Platinum 8380 processors achieved a performance value of 11,652 on the SUSE Linux Enterprise Server 15 SP2 in the SPECpower_ssj2008 benchmark, in the Linux division of the 3rd Generation Intel® Xeon® Processor Scalable Family category and won first place in 2-socket server SPECpower_ssj2008 performance.

12,221 SPECpower_ssj2008

On June 9, 2021, PRIMERGY RX2540 M6 with two Xeon Platinum 8380 processors achieved a performance value of 12,221 on the Windows Server 2019 Standard in the SPECpower_ssj2008 benchmark, in the Windows division of the 3rd Generation Intel® Xeon® Processor Scalable Family category and won first place in 2-socket server SPECpower_ssj2008 performance.

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

SAP Sales and Distribution (SD) Standard Application Benchmark

Benchmark description

Since 1993 the SAP Standard Application Benchmarks have been developed by SAP in order to verify the performance, stability and scaling of a SAP application system and to provide information for configuring, sizing and for platform comparison. By far the most popular benchmarks from the many available are the SAP SD benchmark and the BW Edition for SAP HANA benchmark (see corresponding section).

The Sales and Distribution benchmark is one of the most CPU consuming benchmarks available and has become a de-facto standard for SAP's platform partners and in the ERP (Enterprise Resource Planning) environment.

During the benchmark a defined sequence of business transactions are run through as shown in the table below. The Sales and Distribution (SD) benchmark covers a sell-from-stock scenario (including a customer order creation, the corresponding delivery with subsequent goods movement and creation of the invoice) and consists of the following SAP transactions:

Create an order with five line items (SAP transaction VA01)
Create a delivery for this order (SAP transaction VL01N)
Display the customer order (SAP transaction VA03)
Change the delivery (SAP transaction VL02N) and post goods issue
List 40 orders for one sold-to party (SAP transaction VA05)
Create an invoice (SAP transaction VF01)

Each of the simulated users repeats this series of transactions from the start to the end of a benchmark run. The think time between two user actions is 10 seconds. During the so-called ramp-up phase the number of concurrently working users is increased until the expected limit is reached. When all users are active, the test interval starts. This performance level must be maintained for at least 15 minutes (benchmark rule). After at least 5 minutes of the high load phase one or more database checkpoints must be enforced (i.e. all log file data is flushed back to the database within the high load phase) or the amount of created dirty blocks must be written to disk for at least 5 minutes to stress the I/O subsystem in a realistic way (benchmark rule). At the end of the high load phase users are gradually taken off the system until none is active. When the test concludes, all relevant data (some are gathered with a SAP developed Operating System monitor) are then transferred to the presentation server for further evaluation.

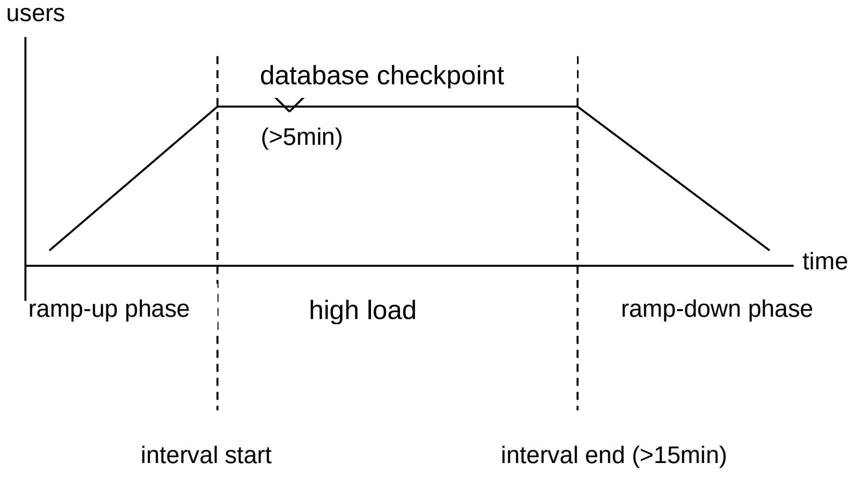

A benchmark can only be certified if the average dialog response time is less than 1 second.

Certified and published SAP SD Benchmarks are published on SAP’s benchmark site [here](#) .

Benchmark environment

The benchmark differentiates between a two-tier and a three-tier configuration. The two-tier configuration has the SAP application and database installed on one server. With a three-tier configuration the individual components of the SAP application can be distributed via several servers and an additional server handles the database.

The SD benchmark users are simulated by the presentation server aka benchmark driver.

The SAP SD Benchmark for PRIMERGY RX2540 M6 was performed on a two-tier configuration.

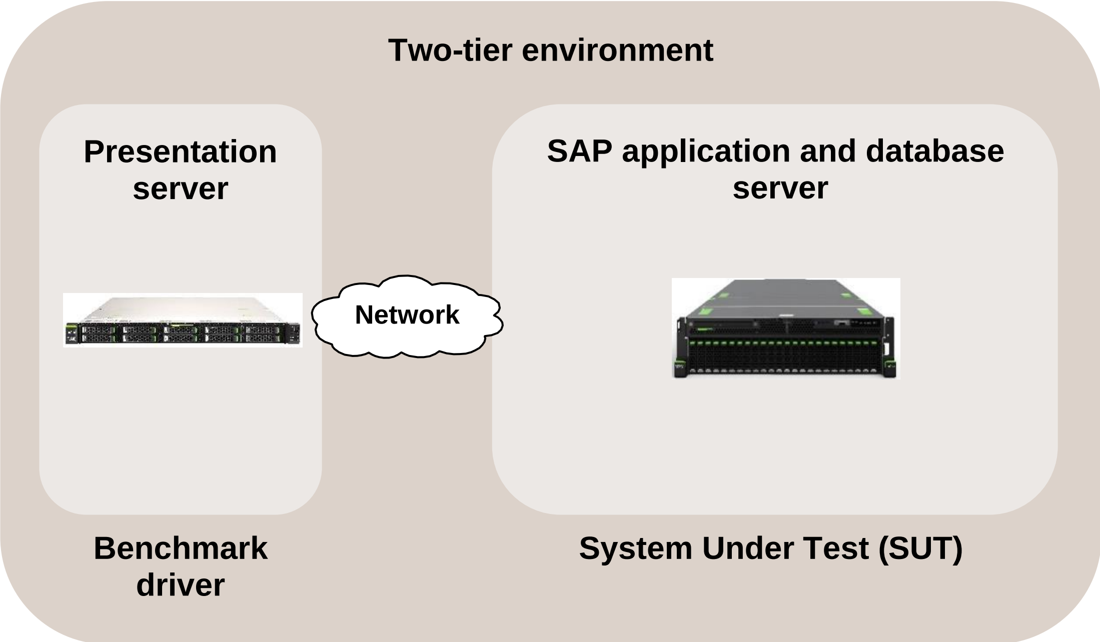

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M6
Processor	2 × Intel® Xeon® Platinum 8380 processor
Memory	32 × 32 GB 2Rx4 DDR4-3200 R ECC
Network interface	1 Gbit LAN
Disk subsystem	PRIMERGY RX2540 M6: 1 × PRAID EP580i RAID Controller 1 × Internal SSD 480GB 1 × NVMe P4610-3.2TB 2 × PRAID EP540e RAID Controller JX40 S2 × 3: 25 × JX40 S2 MLC SSD 960GB 3DWPD
oftware	
Operating system	Windows Server 2016
Database	Microsoft SQL Server 2012
SAP Business Suite Software	SAP enhancement package 5 for SAP ERP 6.0

Benchmark driver	
Hardware	
Model	PRIMERGY RX2530 M1
Processor	2 × Intel® Xeon® E5-2699 v3 processor
Memory	236 GB
Network interface	1 Gbit LAN
Software	
Operating system	SUSE Linux Enterprise Server 12 SP2

Benchmark results

Certification number 2021024	
Number of SAP SD benchmark users	42,750
Average dialog response time	0.99 seconds
Throughput Fully processed order line items/hour Dialog steps/hour SAPS	 4,668,330 14,005,000 233,420
Average database request time (dialog/update)	0.009 sec / 0.022 sec
CPU utilization of central server	98%
Operating system, central server	Windows Server 2016
RDBMS	Microsoft SQL Server 2012
SAP Business Suite software	SAP enhancement package 5 for SAP ERP 6.0
Configuration Central Server	Fujitsu Server PRIMERGY RX2540 M6, 2 processors / 80 cores / 160 threads, Intel Xeon Platinum 8380 processor, 2.30 GHz, 80 KB L1 cache and 1,280 KB L2 cache per core, 60 MB L3 cache per processor, 1,024 GB main memory

The following chart compares the two-tier SAP SD Standard Application Benchmarks for PRIMERGY RX2540 M6 and its predecessor RX2540 M5, shown are the number of benchmark users.

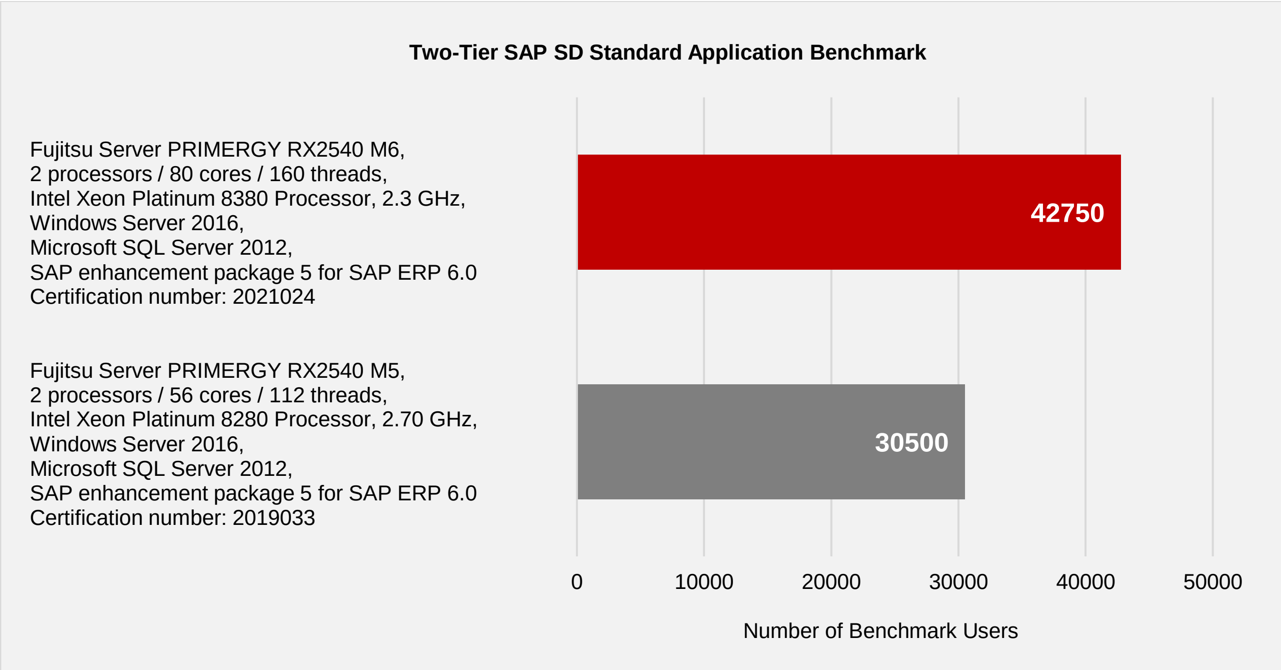

The SAP SD Benchmark certificates can be found here: Certificate 2021024, Certificate 2019010

Disk I/O: Performance of storage media

Benchmark description

Performance measurements of disk subsystems for PRIMERGY servers are used for performance evaluation. It is possible to compare different storage connections as well. As standard, these performance measurements are carried out with a defined measurement method, which models the accesses of real application scenarios on the basis of specifications.

The essential specifications are as follows.

- Random access / sequential access ratio
- Read / write access ratio
- Block size (kB)
- Number of parallel accesses (number 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 [kB]	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 "number of Outstanding I/Os" 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 [MB/s] Throughput in megabytes per second
- Transactions [IO/s] Transaction rate in I/O operations per second
- Latency [ms] Average response time in ms

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

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

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 MB / s = 2^{20} bytes) is used to indicate the capacity of other media, file size, block size, and throughput.

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

Benchmark environment

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

System Under Test (SUT)			
Hardware			
3.5 inch model			
Controller: PRAID EP540i			
	Storage media	Category	Drive name
	HDD	SAS HDD(SAS 12Gbps, 15krpm)[512n]	ST300MP0006 ST600MP0006
		SAS HDD(SAS 12Gbps, 10krpm)[512e]	AL15SEB18EQ AL15SEB24EQ
		SAS HDD(SAS 12Gbps, 10krpm)[512n]	AL15SEB060N AL15SEB120N
		NL-SAS HDD (SAS 12Gbps, 7.2krpm)[512e]	ST6000NM029A ST8000NM001A
		NL-SAS HDD (SAS 12Gbps, 7.2krpm)[512n]	ST2000NM003A ST4000NM003A
		BC-SATA HDD(SATA 6Gbps, 7.2krpm)[512e]	ST6000NM021A ST8000NM000A
		BC-SATA HDD(SATA 6Gbps, 7.2krpm)[512n]	ST2000NM000A ST4000NM000A
	SSD	SAS SSD(SAS 12Gbps, Write Intensive)	XS400ME70084 XS800ME70084 XS1600ME70084
		SAS SSD(SAS 12Gbps, Mixed Use)	XS800LE70084 XS1600LE70084 XS3200LE70084
		SAS SSD(SAS 12Gbps, Read Intensive)	XS960SE70084 XS1920SE70084 XS3840SE70084 XS7680SE70084
		SATA SSD(SATA 6Gbps, Mixed Use)	MTFDDAK240TDT MTFDDAK480TDT MTFDDAK960TDT MTFDDAK1T9TDT MTFDDAK3T8TDT
		SATA SSD(SATA 6Gbps, Read Intensive)	MTFDDAK240TDS MTFDDAK480TDS MTFDDAK960TDS MTFDDAK1T9TDS MTFDDAK3T8TDS MTFDDAK7T6TDS

Controller: Intel C620 Standard SATA AHCI controller			
	Storage media	Category	Drive name
	SSD	M.2 Flash module	MTFDDAV240TDS MTFDDAV480TDS

2.5 inch model			
Controller: PRAID EP540i			
	Storage media	Category	Drive name
	HDD	SAS HDD(SAS 12Gbps, 15krpm)[512n]	ST300MP0006 ST600MP0006
		SAS HDD(SAS 12Gbps, 10krpm)[512e]	AL15SEB18EQ AL15SEB24EQ
		SAS HDD(SAS 12Gbps, 10krpm)[512n]	AL15SEB060N AL15SEB120N
	SSD	SAS SSD(SAS 12Gbps, Write Intensive)	XS400ME70084 XS800ME70084 XS1600ME70084
		SAS SSD(SAS 12Gbps, Mixed Use)	XS800LE70084 XS1600LE70084 XS3200LE70084
		SAS SSD(SAS 12Gbps, Read Intensive)	XS960SE70084 XS1920SE70084 XS3840SE70084 XS7680SE70084
		SATA SSD(SATA 6Gbps, Mixed Use)	MTFDDAK240TDT MTFDDAK480TDT MTFDDAK960TDT MTFDDAK1T9TDT MTFDDAK3T8TDT
		SATA SSD(SATA 6Gbps, Read Intensive)	MTFDDAK240TDS MTFDDAK480TDS MTFDDAK960TDS MTFDDAK1T9TDS MTFDDAK3T8TDS MTFDDAK7T6TDS
		PCIe SSD(Write intensive)	SSDPE21K750GA
		PCIe SSD(Mixed Use)	SSDPE2KE016T8 SSDPE2KE032T8 SSDPE2KE064T8
		PCIe SSD(Read intensive)	SSDPE2KX010T8 SSDPE2KX020T8 SSDPE2KX040T8

Controller:	Intel C620 Standard SATA AHCI controller		
	Storage media	Category	Drive name
	SSD	M.2 Flash module	MTFDDAV240TDS MTFDDAV480TDS

EDSFF NVMe model			
Controller:	Integrated PCI Express controller CPU: 2x Xeon Gold 6330 processor (2GHz, 28core, 42MB)		
	Storage media	Category	Drive name
	SSD	E1.S form factor PCIe SSD(Read Intensive)	SSDPEYKX040T8C
Controller:	Intel C620 Standard SATA AHCI controller		
	Storage media	Category	Drive name
	SSD	M.2 Flash module	MTFDDAV240TDS MTFDDAV480TDS

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

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

Benchmark results

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

Controller

The measurements were made using controllers in the table below.

Storage media	Controller name	Cache	Supported interfaces		RAID levels
			host	drive	
SSD/HDD PCIe SSD 2.5"	PRAID EP540i	-	PCIe 3.0 x8	SATA 6G SAS 12G PCIe x16	0, 1, 1E, 10, 5, 50
PCIe SSD E1.S	Integrated PCI Express controller	-	PCIe 4.0 x4	PCIe 3.0 x4	-
M.2 Flash	C620 Standard SATA AHCI controller	-	DMI 3.0 x4	SATA 6G	-

Storage media

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

Storage media type	Interface	Form factor
HDD	SAS 12G	2.5 inch
	SATA 6G	2.5 inch
SSD	SAS 12G	2.5 inch
	SATA 6G	2.5 inch or M.2
	PCIe 3.0	2.5 inch

model	Storage media type	Interface	Form factor
3.5 inch model	HDD	SAS 12G	3.5 inch or 2.5 inch [1)]
		SATA 6G	3.5 inch
	SSD	SAS 12G	2.5 inch [1)]
		SATA 6G	2.5 inch [1)] or M.2
2.5 inch model	HDD	SAS 12G	2.5 inch
		SATA 6G	2.5 inch
	SSD	SAS 12G	2.5 inch
		SATA 6G	2.5 inch or M.2
		PCIe 3.0	2.5 inch
EDSFF NVMe model	SSD	PCIe 3.0	E1.S
		SATA 6G	M.2

1) It is available with a 3.5 inch cage.

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

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. To avoid any confusion among the measurement units the tables have been separated for the two access types.

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

3.5 inch model

HDDs

Random accesses (units: IO/s)

Capacity [GB]	Storage device	interface	Transactions [IO/s]		
			Database	Fileserver	Filecopy
300	ST300MP0006	SAS 12G	790	696	666
600	ST600MP0006	SAS 12G	736	651	601
1,800	AL15SEB18EQ	SAS 12G	767	631	624
2,400	AL15SEB24EQ	SAS 12G	754	620	617
600	AL15SEB060N	SAS 12G	698	586	600
1,200	AL15SEB120N	SAS 12G	732	604	615
6,000	ST6000NM029A	SAS 12G	369	333	323
8,000	ST8000NM001A	SAS 12G	354	310	310
2,000	ST2000NM003A	SAS 12G	378	343	336
4,000	ST4000NM003A	SAS 12G	369	333	330
6,000	ST6000NM021A	SATA 6G	326	293	302
8,000	ST8000NM000A	SATA 6G	325	290	301
2,000	ST2000NM000A	SATA 6G	331	304	313
4,000	ST4000NM000A	SATA 6G	313	290	297

Sequential accesses (units: MB/s)

Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
300	ST300MP0006	SAS 12G	304	304
600	ST600MP0006	SAS 12G	301	300
1,800	AL15SEB18EQ	SAS 12G	255	249
2,400	AL15SEB24EQ	SAS 12G	264	260
600	AL15SEB060N	SAS 12G	232	232
1,200	AL15SEB120N	SAS 12G	230	226
6,000	ST6000NM029A	SAS 12G	252	252
8,000	ST8000NM001A	SAS 12G	255	255
2,000	ST2000NM003A	SAS 12G	237	237
4,000	ST4000NM003A	SAS 12G	214	215
6,000	ST6000NM021A	SATA 6G	253	253
8,000	ST8000NM000A	SATA 6G	249	248
2,000	ST2000NM000A	SATA 6G	230	207
4,000	ST4000NM000A	SATA 6G	211	210

SSDs

Random accesses (units: IO/s)

Capacity [GB]	Storage device	interface	Transactions [IO/s]		
			Database	Fileserver	Filecopy
400	XS400ME70084	SAS 12G	122,956	22,969	19,438
800	XS800ME70084	SAS 12G	123,848	23,784	19,435
1,600	XS1600ME70084	SAS 12G	123,277	23,725	19,270
800	XS800LE70084	SAS 12G	121,914	23,707	19,257
1,600	XS1600LE70084	SAS 12G	122,949	23,771	19,455
3,200	XS3200LE70084	SAS 12G	123,090	22,816	19,418
960	XS960SE70084	SAS 12G	123,014	23,678	19,424
1,920	XS1920SE70084	SAS 12G	123,093	23,760	19,423
3,840	XS3840SE70084	SAS 12G	122,810	22,949	19,406
7,680	XS7680SE70084	SAS 12G	123,461	22,899	19,516
240	MTFDDAK240TDT	SATA 6G	46,406	5,989	6,121
480	MTFDDAK480TDT	SATA 6G	49,138	6,383	6,600
960	MTFDDAK960TDT	SATA 6G	50,488	6,970	7,136
1,920	MTFDDAK1T9TDT	SATA 6G	50,669	7,183	7,336
3,840	MTFDDAK3T8TDT	SATA 6G	49,490	7,115	7,208
240	MTFDDAK240TDS	SATA 6G	42,594	5,435	5,510
480	MTFDDAK480TDS	SATA 6G	47,577	6,109	6,310
960	MTFDDAK960TDS	SATA 6G	50,134	6,633	6,852
1,920	MTFDDAK1T9TDS	SATA 6G	50,638	7,078	7,286
3,840	MTFDDAK3T8TDS	SATA 6G	49,542	7,097	7,196
7,680	MTFDDAK7T6TDS	SATA 6G	47,200	7,134	7,563
240	MTFDDAV240TDS	SATA 6G	32,805	5,482	5,518
480	MTFDDAV480TDS	SATA 6G	39,927	6,384	6,575

Sequential accesses (units: MB/s)

Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
400	XS400ME70084	SAS 12G	1,052	872
800	XS800ME70084	SAS 12G	1,052	874
1,600	XS1600ME70084	SAS 12G	1,051	884
800	XS800LE70084	SAS 12G	1,052	871
1,600	XS1600LE70084	SAS 12G	1,052	874
3,200	XS3200LE70084	SAS 12G	1,051	872
960	XS960SE70084	SAS 12G	1,052	870
1,920	XS1920SE70084	SAS 12G	1,052	874
3,840	XS3840SE70084	SAS 12G	1,051	871
7,680	XS7680SE70084	SAS 12G	1,051	880
240	MTFDDAK240TDT	SATA 6G	508	370
480	MTFDDAK480TDT	SATA 6G	508	437
960	MTFDDAK960TDT	SATA 6G	508	486
1,920	MTFDDAK1T9TDT	SATA 6G	508	487
3,840	MTFDDAK3T8TDT	SATA 6G	493	474
240	MTFDDAK240TDS	SATA 6G	508	301
480	MTFDDAK480TDS	SATA 6G	508	401
960	MTFDDAK960TDS	SATA 6G	506	480
1,920	MTFDDAK1T9TDS	SATA 6G	508	488
3,840	MTFDDAK3T8TDS	SATA 6G	495	477
7,680	MTFDDAK7T6TDS	SATA 6G	508	487
240	MTFDDAV240TDS	SATA 6G	504	300
480	MTFDDAV480TDS	SATA 6G	497	397

2.5 inch model

HDDs

Random accesses (units: IO/s)

Capacity [GB]	Storage device	interface	Transactions [IO/s]		
			Database	Fileserver	Filecopy
300	ST300MP0006	SAS 12G	790	696	666
600	ST600MP0006	SAS 12G	736	651	601
1,800	AL15SEB18EQ	SAS 12G	767	631	624
2,400	AL15SEB24EQ	SAS 12G	754	620	617
600	AL15SEB060N	SAS 12G	698	586	600
1,200	AL15SEB120N	SAS 12G	732	604	615

Sequential accesses (units: MB/s)

Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
300	ST300MP0006	SAS 12G	304	304
600	ST600MP0006	SAS 12G	301	300
1,800	AL15SEB18EQ	SAS 12G	255	249
2,400	AL15SEB24EQ	SAS 12G	264	260
600	AL15SEB060N	SAS 12G	232	232
1,200	AL15SEB120N	SAS 12G	230	226

SSDs

Random accesses (units: IO/s)

Capacity [GB]	Storage device	interface	Transactions [IO/s]		
			Database	Fileserver	Filecopy
400	XS400ME70084	SAS 12G	122,956	22,969	19,438
800	XS800ME70084	SAS 12G	123,848	23,784	19,435
1,600	XS1600ME70084	SAS 12G	123,277	23,725	19,270
800	XS800LE70084	SAS 12G	121,914	23,707	19,257
1,600	XS1600LE70084	SAS 12G	122,949	23,771	19,455
3,200	XS3200LE70084	SAS 12G	123,090	22,816	19,418
960	XS960SE70084	SAS 12G	123,014	23,678	19,424
1,920	XS1920SE70084	SAS 12G	123,093	23,760	19,423
3,840	XS3840SE70084	SAS 12G	122,810	22,949	19,406
7,680	XS7680SE70084	SAS 12G	123,461	22,899	19,516
240	MTFDDAK240TDT	SATA 6G	46,406	5,989	6,121
480	MTFDDAK480TDT	SATA 6G	49,138	6,383	6,600
960	MTFDDAK960TDT	SATA 6G	50,488	6,970	7,136
1,920	MTFDDAK1T9TDT	SATA 6G	50,669	7,183	7,336
3,840	MTFDDAK3T8TDT	SATA 6G	49,490	7,115	7,208
240	MTFDDAK240TDS	SATA 6G	42,594	5,435	5,510
480	MTFDDAK480TDS	SATA 6G	47,577	6,109	6,310
960	MTFDDAK960TDS	SATA 6G	50,134	6,633	6,852
1,920	MTFDDAK1T9TDS	SATA 6G	50,638	7,078	7,286
3,840	MTFDDAK3T8TDS	SATA 6G	49,542	7,097	7,196
7,680	MTFDDAK7T6TDS	SATA 6G	47,200	7,134	7,563
750	SSDPE21K750GA	PCIe3 x4	194,085	37,392	36,626
1,600	SSDPE2KE016T8	PCIe3 x4	276,785	45,739	40,923
3,200	SSDPE2KE032T8	PCIe3 x4	306,446	53,059	50,093
6,400	SSDPE2KE064T8	PCIe3 x4	297,505	56,338	56,632
1,000	SSDPE2KX01	PCIe3 x4	153,263	25,891	21,942
2,000	SSDPE2KX02	PCIe3 x4	237,530	38,336	34,740
4,000	SSDPE2KX04	PCIe3 x4	242,546	39,242	38,151
240	MTFDDAV240TDS	SATA 6G	32,805	5,482	5,518
480	MTFDDAV480TDS	SATA 6G	39,927	6,384	6,575

Sequential accesses (units: MB/s)

Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
400	XS400ME70084	SAS 12G	1,052	872
800	XS800ME70084	SAS 12G	1,052	874
1,600	XS1600ME70084	SAS 12G	1,051	884
800	XS800LE70084	SAS 12G	1,052	871
1,600	XS1600LE70084	SAS 12G	1,052	874
3,200	XS3200LE70084	SAS 12G	1,051	872
960	XS960SE70084	SAS 12G	1,052	870
1,920	XS1920SE70084	SAS 12G	1,052	874
3,840	XS3840SE70084	SAS 12G	1,051	871
7,680	XS7680SE70084	SAS 12G	1,051	880
240	MTFDDAK240TDT	SATA 6G	508	370
480	MTFDDAK480TDT	SATA 6G	508	437
960	MTFDDAK960TDT	SATA 6G	508	486
1,920	MTFDDAK1T9TDT	SATA 6G	508	487
3,840	MTFDDAK3T8TDT	SATA 6G	493	474
240	MTFDDAK240TDS	SATA 6G	508	301
480	MTFDDAK480TDS	SATA 6G	508	401
960	MTFDDAK960TDS	SATA 6G	506	480
1,920	MTFDDAK1T9TDS	SATA 6G	508	488
3,840	MTFDDAK3T8TDS	SATA 6G	495	477
7,680	MTFDDAK7T6TDS	SATA 6G	508	487
750	SSDPE21K750GA	PCIe3 x4	2,561	2,334
1,600	SSDPE2KE016T8	PCIe3 x4	3,214	1,972
3,200	SSDPE2KE032T8	PCIe3 x4	3,220	2,461
6,400	SSDPE2KE064T8	PCIe3 x4	3,219	2,499
1,000	SSDPE2KX01	PCIe3 x4	2,799	1,109
2,000	SSDPE2KX02	PCIe3 x4	3,181	1,979
4,000	SSDPE2KX04	PCIe3 x4	2,905	2,417
240	MTFDDAV240TDS	SATA 6G	504	300
480	MTFDDAV480TDS	SATA 6G	497	397

EDSFF NVMe model

SSDs

Random accesses (units: IO/s)

Capacity [GB]	Storage device	interface	Transactions [IO/s]		
			Database	Fileserver	Filecopy
4,000	SSDPEYKX040T8C	PCIe3 x4	170,015	33,691	32,143
240	MTFDDAV240TDS	SATA 6G	32,805	5,482	5,518
480	MTFDDAV480TDS	SATA 6G	39,927	6,384	6,575

Sequential accesses (units: MB/s)

Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
4,000	SSDPEYKX040T8C	PCIe3 x4	2,589	2,202
240	MTFDDAV240TDS	SATA 6G	504	300
480	MTFDDAV480TDS	SATA 6G	497	397

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:

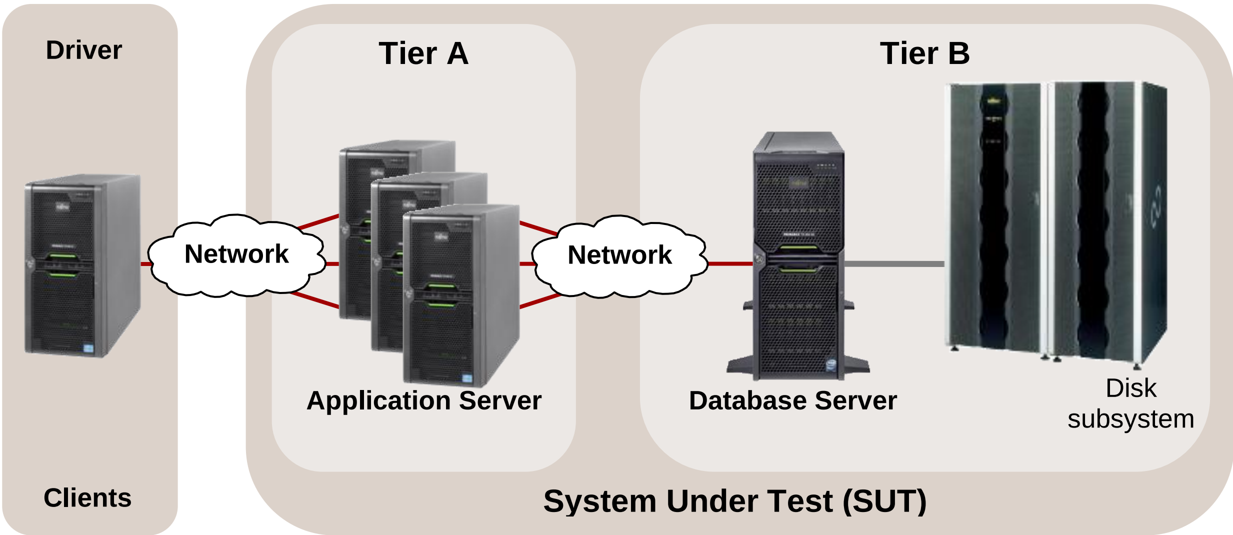

All OLTP-2 results were Calculated based on the configuration of the next following pages of PRIMERGY RX2540 M6

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 why a configuration with a total memory of 2048 GB was considered for the measurements with two processors and a configuration with a total memory of 1024 GB for the measurements with one processor. Both memory configurations have memory access of 3200 MHz.

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

Processor	Cores	Threads	2CPU Score	1CPU Score
Xeon Platinum 8380	40	80	9,515	5,228 (est.)
Xeon Platinum 8368Q	38	76	9,146 (est.)	5,025 (est.)
Xeon Platinum 8368	38	76	9,111	5,006 (est.)
Xeon Platinum 8362	32	64	8,481 (est.)	4,660 (est.)
Xeon Platinum 8360Y	36	72	8,686	4,772 (est.)
Xeon Platinum 8358P	32	64	8,006 (est.)	4,399 (est.)
Xeon Platinum 8358	32	64	8,320	4,571 (est.)
Xeon Platinum 8352Y	32	64	7,654 (est.)	4,205 (est.)
Xeon Platinum 8352V	36	72	7,697 (est.)	4,229 (est.)
Xeon Platinum 8352M	32	64	7,094 (est.)	3,898 (est.)
Xeon Gold 6354	18	36	5,491 (est.)	3,017 (est.)
Xeon Gold 6348	28	56	7,318	4,021 (est.)
Xeon Gold 6346	16	32	5,005	2,750 (est.)
Xeon Gold 6342	24	48	6,952	3,820 (est.)
Xeon Gold 6338T	24	48	6,170 (est.)	3,390 (est.)
Xeon Gold 6338	32	64	7,454	4,096 (est.)
Xeon Gold 6336Y	24	48	6,253 (est.)	3,436 (est.)
Xeon Gold 6334	8	16	2,737	1,504 (est.)
Xeon Gold 6330N	28	56	6,024	3,310 (est.)
Xeon Gold 6330	28	56	6,798	3,735 (est.)
Xeon Gold 6326	16	32	5,003	2,749 (est.)
Xeon Gold 6314U	32	64	-	4,178 (est.)
Xeon Gold 6312U	24	48	-	3,380 (est.)
Xeon Gold 5320	26	52	6,526	3,586 (est.)
Xeon Gold 5318Y	24	48	6,109	3,357 (est.)
Xeon Gold 5318S	24	48	6,109 (est.)	3,357 (est.)
Xeon Gold 5317	12	24	3,561	1,957 (est.)
Xeon Gold 5315Y	8	16	2,548	1,400 (est.)
Xeon Silver 4316	20	40	4,867	2,674 (est.)
Xeon Silver 4314	16	32	4,026 (est.)	2,212 (est.)
Xeon Silver 4310	12	24	3,092 (est.)	1,699 (est.)
Xeon Silver 4309Y	8	16	2,306 (est.)	1,267 (est.)

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

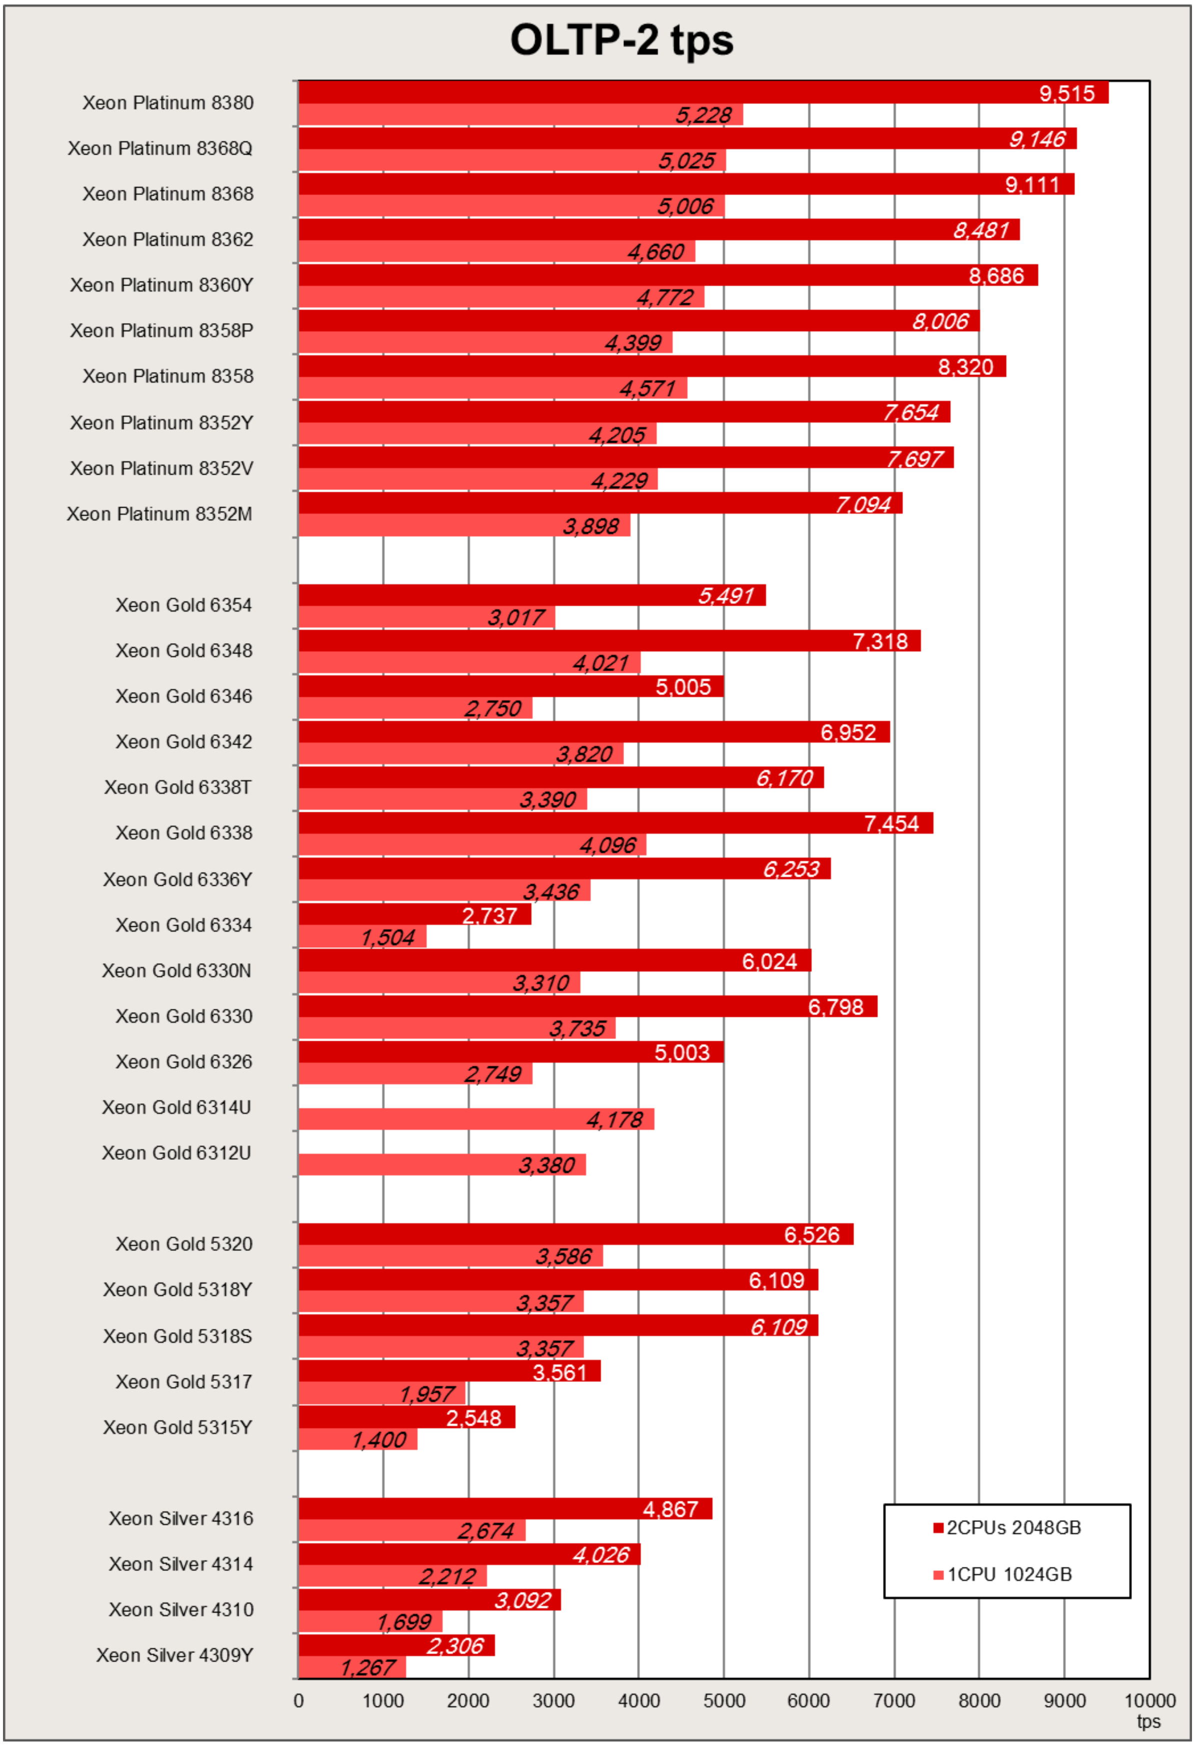

It is evident that a wide performance range is covered by the variety of released processors. If you compare the OLTP-2 value of the processor with the lowest performance (Xeon Platinum 8380) with the value of the processor with the highest performance (Xeon Sliver 4309Y) the OLTP-2 value increased by a factor of 4.

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.

The highest value for OLTP-2 on the current PRIMERGY model is about 40% higher than the highest value on the previous model.

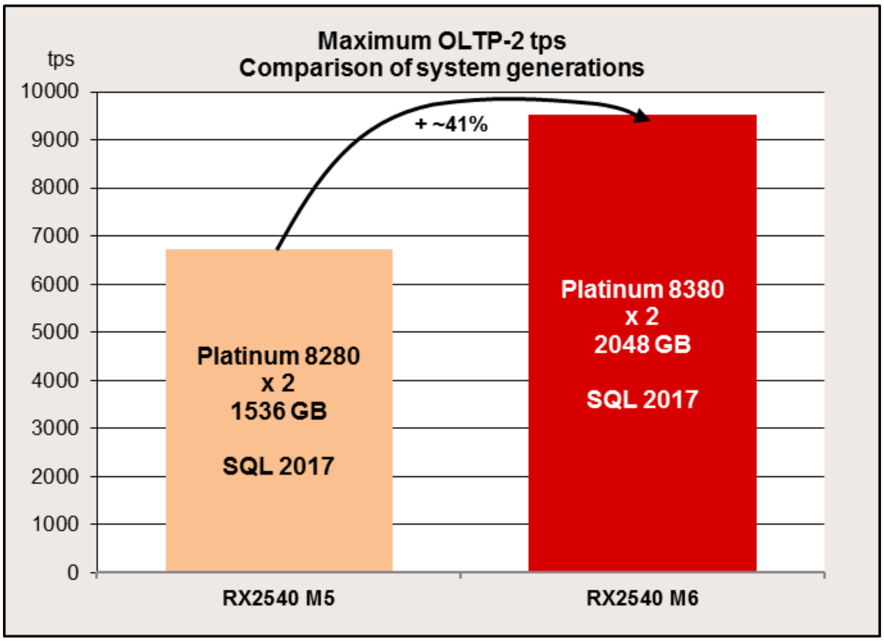

VMmark V3

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).

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.

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

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

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

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

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

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

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

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

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

Benchmark environment

The typical measurement set-up is illustrated below:

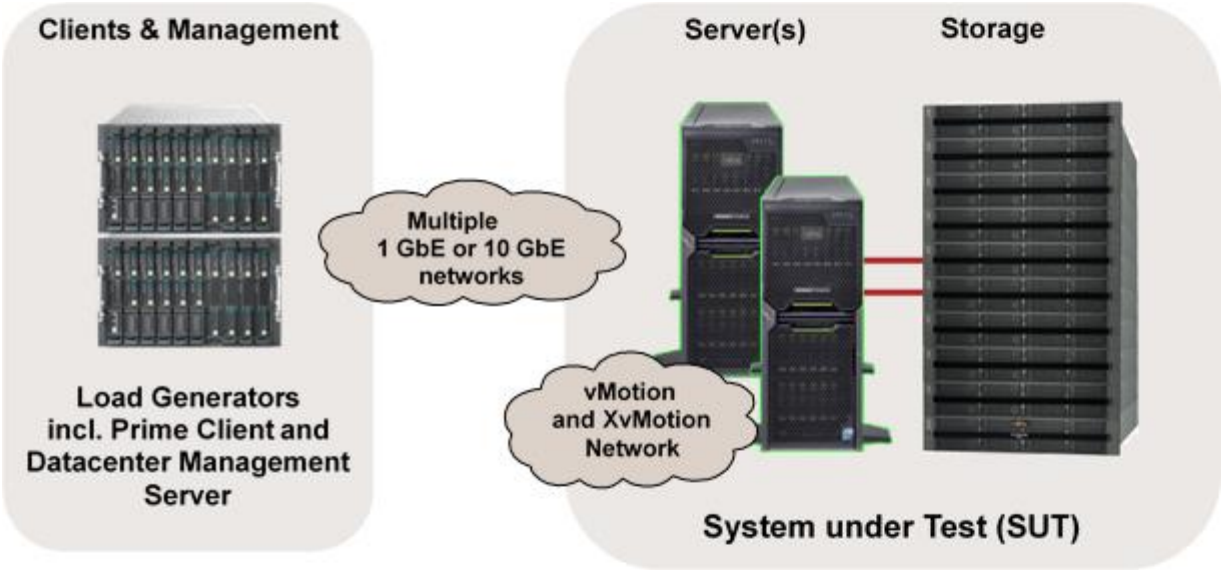

System Under Test (SUT)	
Hardware	
Number of servers	2
Model	PRIMERGY RX2540 M6
Processor	2 x Intel® Xeon® Platinum 8380
Memory	2048 GB：32 x 64 GB (1x64 GB) 2Rx4 DDR4-3200 R ECC
Network interface	2 x Mellanox MCX4121A-ACAT dual port 25Gb SFP28 PCIe adapter 1 x Intel® I350-T2 1Gb quad port OCPv3
Disk subsystem	2 x Qlogic QLE2772 dual port 32Gb PCIe adapter 5 x PRIMERGY RX2540 M4 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)
Software	
BIOS	V1.0.0.0 R1.2.0 for D3891-A1x
BIOS settings	See "Details"
Operating system	VMware ESXi 7.0 U2, Build 17630552
Operating system settings	ESX settings: see "Details"

Details	
See disclosure	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2021-04-20-Fujitsu-PRIMERGY-RX2540M6.pdf https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2021-04-20-Fujitsu-PRIMERGY-RX2540M6-serverPPKW.pdf

	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2021-04-20-Fujitsu-PRIMERGY-RX2540M6-serverstoragePPKW.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 U1 Build 16860138

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	256 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 (April 21, 2021)

On April 21, 2021, Fujitsu achieved a VMmark V3.1.1 score of "14.19@14 tiles" using PRIMERGY RX2540 M6 with Xeon Platinum 8380 processors and VMware ESXi 7.0 U2. At this time, the system configuration had a total of 2 x 80 processor cores, and two identical servers were used for the "System Under Test" (SUT). Based on the above results, PRIMERGY RX2540 M6 is rated as the most powerful 2-socket Intel processor based rack server in a "matched pair" configuration with two identical hosts in the official VMmark V3 "Performance Only" ranking (as of the date the benchmark results were published).

All comparisons for the competitor products reflect the status of April 21, 2021. For the latest VMmark V3 "Performance Only" results, as well as detailed results and configuration data, see https://www.vmware.com/products/vmmark/results3x.html.

All VMs, their application data, the host operating system, and any additional data needed are stored in a powerful Fiber Channel disk subsystem. This disk subsystem uses fast PCIe SSDs such as Intel® Optane™ to improve storage media response time. Network connectivity with host-side load generators and infrastructure load connectivity between hosts are implemented using 25GbE LAN ports.

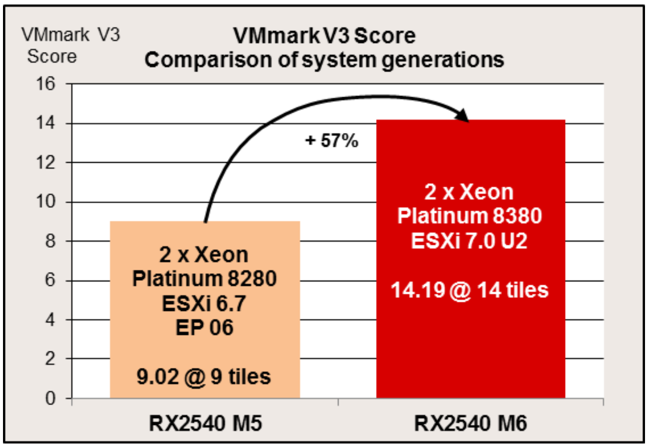

The graph on the left compares the VMmark V3 scores of the PRIMEGY RX2540 M6 and the previous generation PRIMEGY RX2540 M5.
The PRIMERGY RX2540 M6 achieved a 57% improvement in score compared to the previous generation PRIMERGY RX2540 M5. This is due to the improved performance of the 3rd generation Intel Xeon scalable processor and the effective use of the capabilities of the VMware ESXi hypervisor.

"Performance with Server Power" measurement results (April 21, 2021)
"Performance with Server and Storage Power" measurement results (April 21, 2021)

On April 21, 2021, Fujitsu achieved a VMmark V3.1.1 "Server PPKW" score of "7.1922@14 tiles" using PRIMERGY RX2540 M6 with Xeon Platinum 8380 processors and VMware ESXi 7.0 U2. At the same time, it also achieved a VMmark V3.1.1 "Server and Storage PPKW" score of "4.1138@14 tiles." These were system configurations with a total of 2 x 80 processor cores, and two identical servers were used for the "System Under Test" (SUT). Based on the above results, PRIMERGY RX2540 M6 is rated as the most energy efficient virtual server in the world in the official VMmark V3 "Performance with Server Power" ranking and "Performance with Server and Storage Power" ranking (as of the date the benchmark results were published).

For the latest VMmark V3 "Performance with Server Power" results, detailed results, and configuration data, see https://www.vmware.com/products/vmmark/results3x.1.html.
For the latest VMmark V3 "Performance Server and Storage Power" results, detailed results, and configuration data, see https://www.vmware.com/products/vmmark/results3x.2.html.

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 x b(i)	16	1
SUM	a(i) = b(i) + c(i)	24	1
TRIAD	a(i) = b(i) + q x 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 RX2540 M6
Processor	2 x 3rd Generation Intel® Xeon® Scalable Processors Family
Memory	32 x 32 GB 2Rx4 PC4-3200AA-R
Software	
BIOS settings	Override OS Energy Performance = Enabled Energy Performance = Performance HWPM Support = Disabled Intel Virtualization Technology = Disabled LLC Dead Line Alloc = Disabled Stale AtoS = Enabled 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	Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1) Transparent Huge Pages inactivated 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]
Xeon Platinum 8380	3200	205	40	2.3	2	**328**
Xeon Platinum 8368Q	3200	205	38	2.6	2	**326**
Xeon Platinum 8368	3200	205	38	2.4	2	**319**
Xeon Platinum 8362*2	3200	205	32	2.8	2	**325**
Xeon Platinum 8360Y	3200	205	36	2.4	2	**326**
Xeon Platinum 8358P	3200	205	32	2.6	2	**327**
Xeon Platinum 8358	3200	205	32	2.6	2	**330**
Xeon Platinum 8352Y	3200	205	32	2.2	2	**327**
Xeon Platinum 8352V	2933	188	36	2.1	2	**299**
Xeon Platinum 8352M*2	3200	205	32	2.3	2	**319**
Xeon Gold 6354	3200	205	18	3.0	2	**312**
Xeon Gold 6348	3200	205	28	2.6	2	**327**
Xeon Gold 6346	3200	205	16	3.1	2	**310**
Xeon Gold 6342*2	3200	205	24	2.8	2	**322**
Xeon Gold 6338T*2	2933	188	24	2.1	2	**318**
Xeon Gold 6338	3200	205	32	2.0	2	**329**
Xeon Gold 6336Y*2	3200	205	24	2.4	2	**320**
Xeon Gold 6334*2	3200	205	8	3.6	2	**242**
Xeon Gold 6330N	2666	171	28	2.2	2	**278**
Xeon Gold 6330	2933	188	28	2.0	2	**299**
Xeon Gold 6326*2	3200	205	16	2.9	2	**294**
Xeon Gold 6314U	3200	205	32	2.3	2	**167**
Xeon Gold 6312U*2	2933	205	24	2.4	2	**165**
Xeon Gold 5320*2	2933	188	26	2.2	2	**298**
Xeon Gold 5318Y*2	2933	188	24	2.1	2	**296**
Xeon Gold 5318S*2	2933	188	24	2.1	2	**293**
Xeon Gold 5317*2	2933	188	12	3.0	2	**240**
Xeon Gold 5315Y*2	2933	188	8	3.2	2	**183**
Xeon Silver 4316*2	2666	171	20	2.3	2	**264**
Xeon Silver 4314*2	2666	171	16	2.4	2	**256**
Xeon Silver 4310*2	2666	171	12	2.1	2	**230**
Xeon Silver 4309Y*2	2666	171	8	2.8	2	**166**

*1: Value per Processor
*2: To be supported in July 2021 or later

The following diagram illustrates the throughput of the RX2540 M6in comparison to its predecessor, the RX2540 M5.

STREAM TRIAD: Comparison of PRIMERGY RX2540 M6 and PRIMERGY RX2540 M5

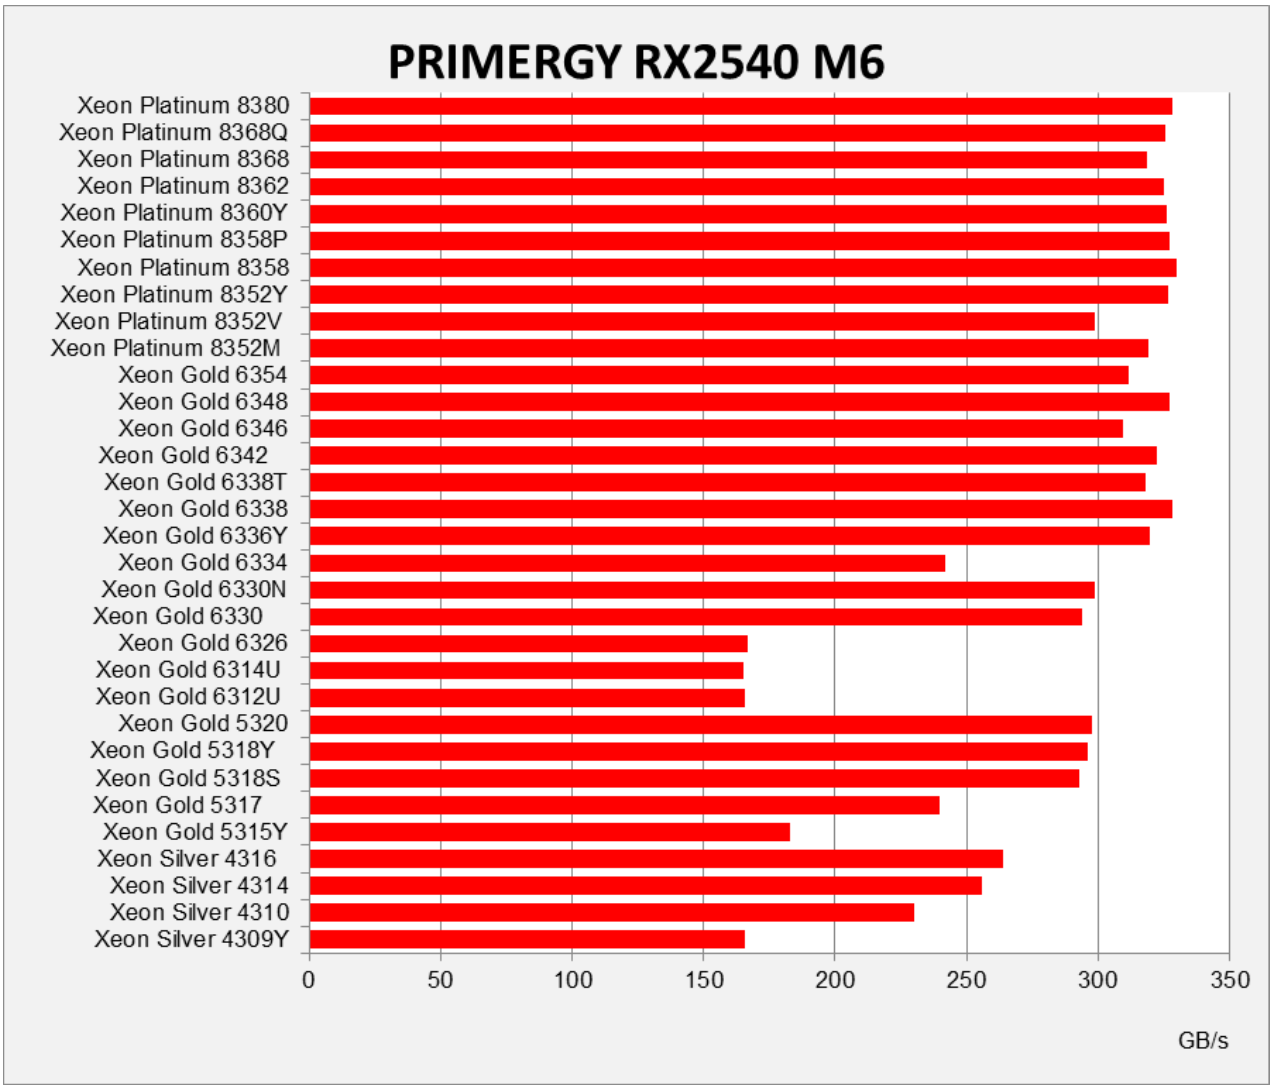

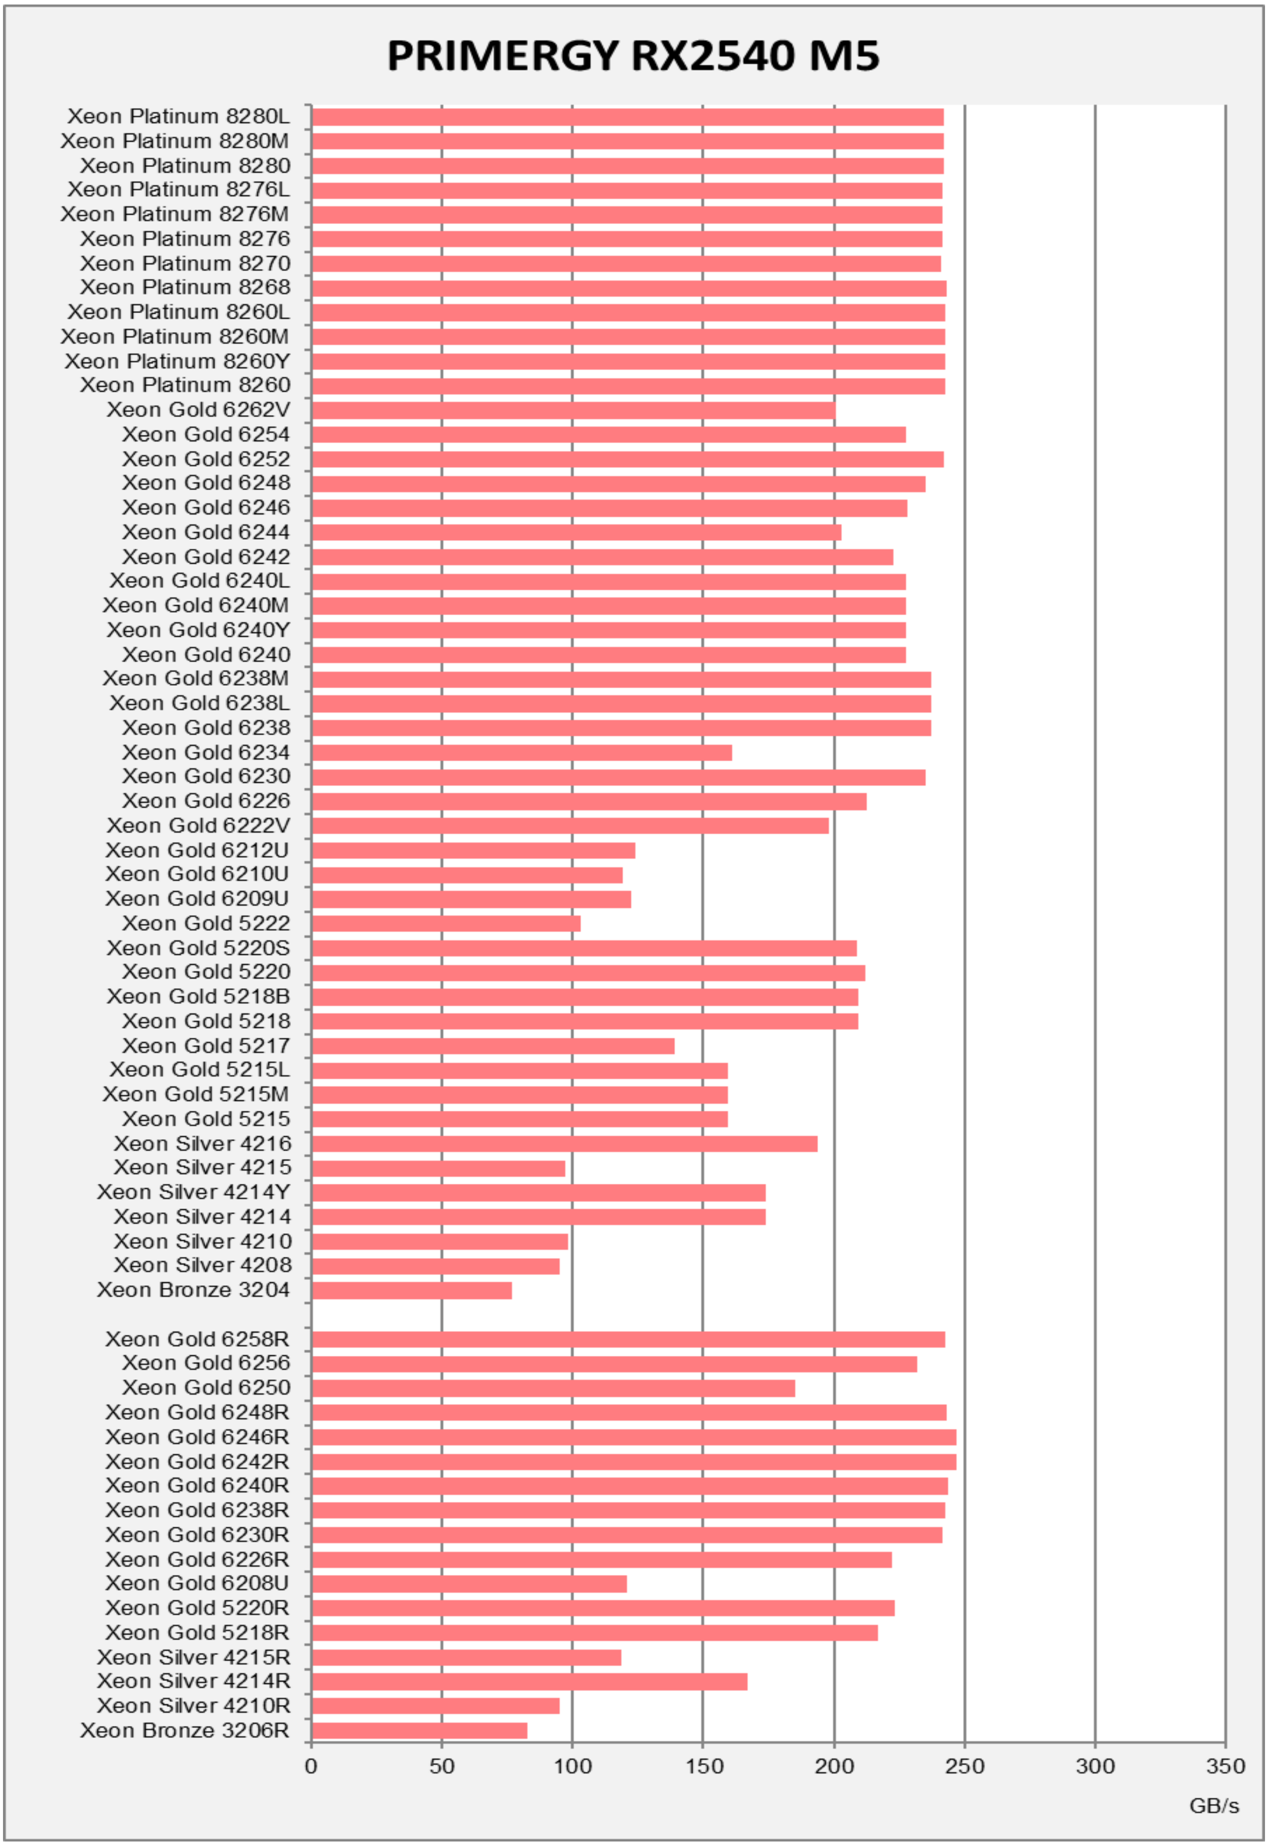
PRIMERGY RX2540 M5
Xeon Platinum 8280L
Xeon Platinum 8280M
Xeon Platinum 8280
Xeon Platinum 8276L
Xeon Platinum 8276M
Xeon Platinum 8276
Xeon Platinum 8270
Xeon Platinum 8268
Xeon Platinum 8260L
Xeon Platinum 8260M
Xeon Platinum 8260Y
Xeon Platinum 8260
Xeon Gold 6262V
Xeon Gold 6254
Xeon Gold 6252
Xeon Gold 6248
Xeon Gold 6246
Xeon Gold 6244
Xeon Gold 6242
Xeon Gold 6240L
Xeon Gold 6240M
Xeon Gold 6240Y
Xeon Gold 6240
Xeon Gold 6238M
Xeon Gold 6238L
Xeon Gold 6238
Xeon Gold 6234
Xeon Gold 6230
Xeon Gold 6226
Xeon Gold 6222V
Xeon Gold 6212U
Xeon Gold 6210U
Xeon Gold 6209U
Xeon Gold 5222
Xeon Gold 5220S
Xeon Gold 5220
Xeon Gold 5218B
Xeon Gold 5218
Xeon Gold 5217
Xeon Gold 5215L
Xeon Gold 5215M
Xeon Gold 5215
Xeon Silver 4216
Xeon Silver 4215
Xeon Silver 4214Y
Xeon Silver 4214
Xeon Silver 4210
Xeon Silver 4208
Xeon Bronze 3204
Xeon Gold 6258R
Xeon Gold 6256
Xeon Gold 6250
Xeon Gold 6248R
Xeon Gold 6246R
Xeon Gold 6242R
Xeon Gold 6240R
Xeon Gold 6238R
Xeon Gold 6230R
Xeon Gold 6226R
Xeon Gold 6208U
Xeon Gold 5220R
Xeon Gold 5218R
Xeon Silver 4215R
Xeon Silver 4214R
Xeon Silver 4210R
Xeon Bronze 3206R
0
50
100
150
200
250
300
350
GB/s

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

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M6
Processor	2 x 3rd Generation Intel® Xeon® Scalable Processors Family
Memory	32 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]	Rmax [GFlops]	Efficiency
Xeon Platinum 8380	40	2.3	2	5,888	**4,431**	75%
Xeon Platinum 8368Q	38	2.6	2	6,323	**4,406**	70%
Xeon Platinum 8368	38	2.4	2	5,837	**4,249**	73%
Xeon Platinum 8362*1	32	2.8	2	5,734	**4,040**	70%
Xeon Platinum 8360Y	36	2.4	2	5,530	**3,938**	71%
Xeon Platinum 8358P	32	2.6	2	5,325	**3,560**	67%
Xeon Platinum 8358	32	2.6	2	5,325	**3,786**	71%
Xeon Platinum 8352Y	32	2.2	2	4,506	**3,166**	70%
Xeon Platinum 8352V	36	2.1	2	4,838	**3,327**	69%
Xeon Platinum 8352M*1	32	2.3	2	4,710	**2,939**	62%
Xeon Gold 6354	18	3.0	2	3,456	**2,446**	71%
Xeon Gold 6348	28	2.6	2	4,659	**3,350**	72%
Xeon Gold 6346	16	3.1	2	3,174	**2,326**	73%
Xeon Gold 6342*1	24	2.8	2	4,301	**3,173**	74%
Xeon Gold 6338T*1	24	2.1	2	3,226	**2,363**	73%
Xeon Gold 6338	32	2.0	2	4,096	**3,156**	77%
Xeon Gold 6336Y*1	24	2.4	2	3,686	**2,657**	72%
Xeon Gold 6334*1	8	3.6	2	1,843	**1,316**	71%
Xeon Gold 6330N	28	2.2	2	3,942	**2,506**	64%
Xeon Gold 6330	28	2.0	2	3,584	**2,994**	84%
Xeon Gold 6326*1	16	2.9	2	2,970	**2,204**	74%
Xeon Gold 6314U	32	2.3	2	2,355	**1,592**	68%
Xeon Gold 6312U*1	24	2.4	2	1,843	**1,341**	73%
Xeon Gold 5320*1	26	2.2	2	3,661	**2,816**	77%
Xeon Gold 5318Y*1	24	2.1	2	3,226	**2,456**	76%
Xeon Gold 5318S*1	24	2.1	2	3,226	**2,457**	76%
Xeon Gold 5317*1	12	3.0	2	2,304	**1,614**	70%
Xeon Gold 5315Y*1	8	3.2	2	1,638	**1,211**	74%
Xeon Silver 4316*1	20	2.3	2	2,944	**2,085**	71%
Xeon Silver 4314*1	16	2.4	2	2,458	**1,702**	69%
Xeon Silver 4310*1	12	2.1	2	1,613	**1,480**	92%
Xeon Silver 4309Y*1	8	2.8	2	1,434	**1,026**	72%

*1: To be supported in July 2021 or later

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.

Literature

PRIMERGY Servers

http://primergy.com/

PRIMERGY RX2540 M6

This White Paper:

https://docs.ts.fujitsu.com/dl.aspx?id=f4ca266a-f3e0-4832-82d5-c5b5b0a3b6d2

https://docs.ts.fujitsu.com/dl.aspx?id=a4cdb252-a45a-47a5-b32d-49cdd428beca

Data sheet
https://docs.ts.fujitsu.com/dl.aspx?id=81d5b31d-e617-4c4d-bc0c-1c14c01140d7

PRIMERGY Performance

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

SPECcpu2017

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

Benchmark Overview SPECcpu2017
http://docs.ts.fujitsu.com/dl.aspx?id=20f1f4e2-5b3c-454a-947f-c169fca51eb1

SPECpower_ssj2008

http://www.spec.org/power_ssj2008

Benchmark Overview SPECpower_ssj2008
http://docs.ts.fujitsu.com/dl.aspx?id=166f8497-4bf0-4190-91a1-884b90850ee0

SAP SD

http://www.sap.com/benchmark

Benchmark overview SAP SD
http://docs.ts.fujitsu.com/dl.aspx?id=0a1e69a6-e366-4fd1-a1a6-0dd93148ea10

OLTP-2

Benchmark Overview OLTP-2
http://docs.ts.fujitsu.com/dl.aspx?id=e6f7a4c9-aff6-4598-b199-836053214d3f

VMmark V3

VMmark 3
http://www.vmmark.com

STREAM

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

LINPACK

The LINPACK Benchmark: Past, Present, and Future
http://www.netlib.org/utk/people/JackDongarra/PAPERS/hplpaper.pdf

TOP500
http://www.top500.org/

HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers
http://www.netlib.org/benchmark/hpl/

Intel Math Kernel Library – LINPACK Download
http://software.intel.com/en-us/articles/intel-math-kernel-library-linpack-download/

Contact

FUJITSU
Website: http://www.fujitsu.com/
PRIMERGY Performance and Benchmarks
mailto:fj-benchmark@dl.jp.fujitsu.com