FUJITSU

White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY RX2540 M5

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY RX2540 M5.

The PRIMERGY RX2540 M5 performance data are compared with the data of other PRIMERGY models and discussed. In addition to the benchmark results, an explanation has been included for each benchmark and for the benchmark environment.

Version
1.3 2020/05/29

Contents

Document history

Version 1.0 (2019/04/30)

New:

- Technical data
- SPECcpu2017
  Measurements with 2nd Generation Intel® Xeon® Processor Scalable Family
- SPECpower_ssj2008
  Measurement with Intel® Xeon® Platinum 8276L and Intel® Xeon® Platinum 8280L
- SPECjbb2015
  Measurement with Intel® Xeon® Platinum 8280M
- SAP SD
- Certification number 2019010
- OLTP-2
  Calculated with 2nd Generation Intel® Xeon® Processor Scalable Family
- VMmark V3
  “Performance Only” measurement with Intel® Xeon® Platinum 8280
  “Performance with Server Power” measurement Intel® Xeon® Platinum 8280
  “Performance with Server and Storage Power” measurement Intel® Xeon® Platinum 8280
- vServCon
  Measurements with 2nd Generation Intel® Xeon® Processor Scalable Family

Version 1.1 (2019/10/04)

New:

- Disk I/O: Performance of storage media
  Results for 2.5" and 3.5" storage media
- STREAM, LINPACK
  Measured with 2nd Generation Intel® Xeon® Processor Scalable Family

Updated:

- SPECcpu2017
  Measured additionally with 2nd Generation Intel® Xeon® Processor Scalable Family

Version 1.2 (2020/04/24)

New:

- TPC-E
  Measurement with Intel® Xeon® Platinum 8280

Updated:

- Technical data
  Added 2nd Generation Intel® Xeon® Processor Scalable Family
- SPECcpu2017, OLTP-2 , vServCon, STREAM, LINPACK
  Measured or calculated additionally with 2nd Generation Intel® Xeon® Processor Scalable Family

Version 1.3 (2020/05/29)

Updated:

- Technical data, LINPAC
  Fixed typo in processor specifications
- STREAM
  Fixed typo in processor specifications and updated measurements

Technical data

PRIMERGY RX2540 M5 2.5'

PRIMERGY RX2540 M5 3.5'

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 M5
Model versions	PY RX2540 M5 4 x 3.5' PY RX2540 M5 12 x 3.5' PY RX2540 M5 8 x 2.5' PY RX2540 M5 16 x 2.5' PY RX2540 M5 24 x 2.5'
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	24 (12 per processor)
Maximum memory configuration	3,072 GB
Onboard HDD controller	Controller with RAID 0, RAID 1 or RAID 10 for up to 8 SATA HDDs
PCI slots	3 × PCI-Express 3.0 x8 3 × PCI-Express 3.0 x16
Max. number of internal hard disks	PY RX2540 M5 4x 3.5': 8 × 3.5" PY RX2540 M5 12x 3.5': 12 × 3.5" + 4 × 2.5" PY RX2540 M5 8x 2.5' : 16 × 2.5" PY RX2540 M5 16x 2.5' :16 × 2.5" PY RX2540 M5 24x 2.5': 24 × 2.5" + 4 × 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 [Watt]
April 2019 released								
Xeon Platinum 8280L	28	56	38.5	10.4	2.7	4.0	2933	205
Xeon Platinum 8280M	28	56	38.5	10.4	2.7	4.0	2933	205
Xeon Platinum 8280	28	56	38.5	10.4	2.7	4.0	2933	205
Xeon Platinum 8276L	28	56	38.5	10.4	2.2	4.0	2933	165
Xeon Platinum 8276M	28	56	38.5	10.4	2.2	4.0	2933	165
Xeon Platinum 8276	28	56	38.5	10.4	2.2	4.0	2933	165
Xeon Platinum 8270	26	52	35.8	10.4	2.7	4.0	2933	205
Xeon Platinum 8268	24	48	35.8	10.4	2.9	3.9	2933	205
Xeon Platinum 8260L	24	48	35.8	10.4	2.4	3.9	2933	165
Xeon Platinum 8260M	24	48	35.8	10.4	2.4	3.9	2933	165
Xeon Platinum 8260Y	24	48	35.8	10.4	2.4	3.9	2933	165
	20	40						
	16	32						
Xeon Platinum 8260	24	48	35.8	10.4	2.4	3.9	2933	165
Xeon Gold 6262V	24	48	33.0	10.4	1.9	3.6	2933	135
Xeon Gold 6254	18	36	24.8	10.4	3.1	4.0	2933	200
Xeon Gold 6252	24	48	35.8	10.4	2.1	3.7	2933	150
Xeon Gold 6248	20	40	27.5	10.4	2.5	3.9	2933	150
Xeon Gold 6246	12	24	24.8	10.4	3.3	4.2	2933	165
Xeon Gold 6244	8	16	24.8	10.4	3.6	4.4	2933	150
Xeon Gold 6242	16	32	22.0	10.4	2.8	3.9	2933	150
Xeon Gold 6240L	18	36	24.8	10.4	2.6	3.9	2933	150
Xeon Gold 6240M	18	36	24.8	10.4	2.6	3.9	2933	150
Xeon Gold 6240Y	18	36	24.8	10.4	2.6	3.9	2933	150
	14	28						
	8	16						
Xeon Gold 6240	18	36	24.8	10.4	2.6	3.9	2933	150
Xeon Gold 6238M	22	44	30.3	10.4	2.1	3.7	2933	140
Xeon Gold 6238L	22	44	30.3	10.4	2.1	3.7	2933	140
Xeon Gold 6238	22	44	30.3	10.4	2.1	3.7	2933	140
Xeon Gold 6234	8	16	24.8	10.4	3.3	4.0	2933	130
Xeon Gold 6230	20	40	27.5	10.4	2.1	3.9	2933	125
Xeon Gold 6226	12	24	19.3	10.4	2.7	3.7	2933	125
Xeon Gold 6222V	20	40	27.5	10.4	1.8	3.6	2400	115
Xeon Gold 6212U	24	48	33.0	10.4	2.4	3.9	2933	165
Xeon Gold 6210U	20	40	27.5	10.4	2.5	3.9	2933	150
Xeon Gold 6209U	20	40	27.5	10.4	2.1	3.9	2933	125
Xeon Gold 5222	4	8	16.5	10.4	3.8	3.9	2933	105
Xeon Gold 5220S	18	36	24.8	10.4	2.7	3.9	2666	125

Xeon Gold 5220	18	36	24.8	10.4	2.2	3.9	2666	125
Xeon Gold 5218B	16	32	22.0	10.4	2.3	3.9	2666	125
Xeon Gold 5218	16	32	22.0	10.4	2.3	3.9	2666	125
Xeon Gold 5217	8	16	11.0	10.4	3.0	3.7	2666	115
Xeon Gold 5215L	10	20	13.8	10.4	2.5	3.4	2666	85
Xeon Gold 5215M	10	20	13.8	10.4	2.5	3.4	2666	85
Xeon Gold 5215	10	20	13.8	10.4	2.5	3.4	2666	85
Xeon Silver 4216	16	32	22.0	9.6	2.1	3.2	2400	100
Xeon Silver 4215	8	16	11.0	9.6	2.5	3.5	2400	85
Xeon Silver 4214Y	12	24	16.5	9.6	2.2	3.2	2400	85
	10	20						
	8	16						
Xeon Silver 4214	12	24	16.5	9.6	2.2	3.2	2400	85
Xeon Silver 4210	10	20	13.8	9.6	2.2	3.2	2400	85
Xeon Silver 4208	8	16	11.0	9.6	2.1	3.2	2400	85
Xeon Bronze 3204	6	6	8.3	9.6	1.9		2133	85
March 2020 released								
Xeon Gold 6258R	28	56	38.5	10.4	2.7	4.0	2933	205
Xeon Gold 6256	12	24	33.0	10.4	3.6	4.5	2933	205
Xeon Gold 6250	8	16	35.8	10.4	3.9	4.5	2933	185
Xeon Gold 6248R	24	48	35.8	10.4	3.0	4.0	2933	205
Xeon Gold 6246R	16	32	35.8	10.4	3.4	4.1	2933	205
Xeon Gold 6242R	20	40	35.8	10.4	3.1	4.1	2933	205
Xeon Gold 6240R	24	48	35.8	10.4	2.4	4.0	2933	165
Xeon Gold 6238R	28	56	38.5	10.4	2.2	4.0	2933	165
Xeon Gold 6230R	26	52	35.8	10.4	2.1	4.0	2933	150
Xeon Gold 6226R	16	32	22.0	10.4	2.9	3.9	2933	150
Xeon Gold 6208U	16	32	22.0	10.4	2.9	3.9	2933	150
Xeon Gold 5220R	24	48	35.8	10.4	2.2	4.0	2666	150
Xeon Gold 5218R	20	40	27.5	10.4	2.1	4.0	2666	125
Xeon Silver 4215R	8	16	11.0	9.6	3.2	4.0	2400	130
Xeon Silver 4214R	12	24	16.5	9.6	2.4	3.5	2400	100
Xeon Silver 4210R	10	20	13.8	9.6	2.4	3.2	2400	100
Xeon Bronze 3206R	8	8	11.0	9.6	1.9		2133	85

All the processors that can be ordered with the PRIMERGY RX2540 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. Listed in the processor table is "Max. Turbo Frequency" for 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 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 (since system release)								
Memory module	Capacity [GB]	Ranks	Bit width of the memory chips	Frequency [MHz]	Load Reduced	Registered	NVDIMM	ECC
8 GB (1x8 GB) 1Rx8 DDR4-2933 R ECC	8	1	8	2933		✓		✓
16 GB (1x16 GB) 2Rx8 DDR4-2933 R ECC	16	2	8	2933		✓		✓
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	2933		✓		✓
64 GB (1x64 GB) 2Rx4 DDR4-2933 R ECC	64	2	4	2933		✓		✓
64 GB (1x64 GB) 4Rx4 DDR4-2933 LR ECC	64	4	4	2933	✓	✓		✓
128 GB (1x128 GB) 4Rx4 DDR4-2933 LR ECC	128	4	4	2933	✓	✓		✓
128GB (1x128GB) DCPMM-2666	128			2666			✓	✓
256GB (1x256GB) DCPMM-2666	256			2666			✓	✓
512GB (1x512GB) DCPMM-2666	512			2666			✓	✓

Power supplies (since system release)	Max. number
Modular PSU 450 W platinum hp	2
Modular PSU 800 W platinum hp	2
Modular PSU 800 W titanium hp	2
Modular PSU 1200 W platinum hp	2

Some components may not be available in all countries or sales regions.
Detailed technical information is available in the data sheet PRIMERGY RX2540 M5.

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”, which differ in the use of compiler optimization. When publishing the results, the base values are always used and the peak values are optional.

Benchmark	Number of single benchmarks	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. Value “1” was defined for the SPECspeed2017_int_base, SPECrate2017_int_base, SPECspeed2017_fp_base, and SPECrate2017_fp_base results of the reference system. For example, a SPECspeed2017_int_base value of 2 means that the measuring system has handled this benchmark twice as fast as the reference system. A SPECrate2017_fp_base value of 4 means that the measuring system has handled this benchmark some 4/[# base copies] times faster than the reference system. “# base copies” specifies how many parallel instances of the benchmark have been executed.

Not every SPECcpu2017 measurement is submitted by us for publication at SPEC. This is why the SPEC web pages do not have every result. As we archive the log files for all measurements, we can prove the correct implementation of the measurements at any time.

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M5
Processor	2nd Generation Intel® Xeon® Scalable Processors Family
Memory	24 × 32GB (1x32GB) 2Rx4 PC4-2933Y-R
Software	
BIOS settings	SPECspeed2017_int: Patrol Scrub = Disabled Override OS Energy Performance = Enabled Energy Performance = Performance Fan Control = Full Sub NUMA Clustering = Disabled WR CRC feature Control = Disabled Hyper-Threading = Disabled SPECspeed2017_fp Hyper-Threading = Disabled Adjacent Cache Line Prefetch = Disabled Override OS Energy Performance = Enabled Energy Performance = Performance Patrol Scrub = Disabled Sub NUMA Clustering = Disabled WR CRC feature Control = Disabled Fan Control = Full UPI Link L0p Enable = Disable UPI Link L1 Enable = Disable Max Page Table Size Select = 2M IO Directory Cashe (IODC) = Disable SPECrate2017_int: 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 Patrol Scrub = Disabled WR CRC feature Control = Disabled Fan Control = Full Sub NUMA Clustering = Disabled *2 Hyper-Threading = Disabled*3 *1: Except Xeon Bronze 3206R, Xeon Silver 4210R, Xeon Silver 4214R, Xeon Silver 4215R, Xeon Gold 6226R, Xeon Gold 6246R, Xeon Gold 6250, Xeon Gold 6256, Xeon Gold 6208U *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" SPECrate2017: Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1) echo 10000000 > /proc/sys/kernel/sched_min_granularity_ns

Compiler	SPECspeed2017_int, SPECrate2017_int: CPU released in April 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 SPECspeed2017_fp C/C++: Version 19.0.2.187 of Intel C/C++ Compiler Build 20190117 for Linux Fortran: Version 19.0.2.187 of Intel Fortran Compiler Build 20190117 for Linux SPECrate2017_fp: CPU released in April 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

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 result with "est." are the estimated values.

SPECrate2017				
Processor	**Cores**	**Number of Processors**	**SPECrate2017 int_base**	**SPECrate2017 fp_base**
April 2019 released				
Xeon Platinum 8280L	28	2	342(est.)	283(est.)
Xeon Platinum 8280M	28	2	342(est.)	283(est.)
Xeon Platinum 8280	28	2	**342**	**283**
Xeon Platinum 8276L	28	2	**304**	**262**
Xeon Platinum 8276M	28	2	304(est.)	262(est.)
Xeon Platinum 8276	28	2	304(est.)	262(est.)
Xeon Platinum 8270	26	2	**319**	**270**
Xeon Platinum 8268	24	2	**304**	**265**
Xeon Platinum 8260L	24	2	276(est.)	249(est.)
Xeon Platinum 8260M	24	2	276(est.)	249(est.)
Xeon Platinum 8260Y	24	2	**285**	**251**
	20	2	248(est.)	232(est.)
	16	2	215(est.)	214(est.)
Xeon Platinum 8260	24	2	276(est.)	249(est.)
Xeon Gold 6262V	24	2	**237**	**208**
Xeon Gold 6254	18	2	**251**	**230**
Xeon Gold 6252	24	2	**268**	**245**
Xeon Gold 6248	20	2	**249**	**229**
Xeon Gold 6246	12	2	**182**	**192**
Xeon Gold 6244	8	2	**131**	**150**
Xeon Gold 6242	16	2	**214**	**208**
Xeon Gold 6240L	18	2	224(est.)	212(est.)
Xeon Gold 6240M	18	2	224(est.)	212(est.)
Xeon Gold 6240Y	18	2	**225**	**214**
	14	2	184(est.)	190(est.)
	8	2	115(est.)	137(est.)
Xeon Gold 6240	18	2	**224**	**212**
Xeon Gold 6238L	22	2	248(est.)	230(est.)
Xeon Gold 6238M	22	2	248(est.)	230(est.)
Xeon Gold 6238	18	2	**248**	**230**
Xeon Gold 6234	22	2	**125**	**140**
Xeon Gold 6230	20	2	**220**	**211**
Xeon Gold 6226	12	2	**164**	**174**

Xeon Gold 6222V	20	2	**199**	**188**
Xeon Gold 6212U	24	1	**143**	**127**
Xeon Gold 6210U	20	1	124(est.)	116(est.)
Xeon Gold 6209U	20	1	**113**	**109**
Xeon Gold 5222	4	2	**62.8**	**77.5**
Xeon Gold 5220S	18	2	**199**	**195**
Xeon Gold 5220	18	2	**197**	**193**
Xeon Gold 5218B	16	2	180(est.)	181(est.)
Xeon Gold 5218	16	2	**180**	**181**
Xeon Gold 5217	8	2	**106**	**118**
Xeon Gold 5215L	10	2	119(est.)	128(est.)
Xeon Gold 5215M	10	2	119(est.)	128(est.)
Xeon Gold 5215	10	2	**119**	**128**
Xeon Silver 4216	16	2	**174**	**171**
Xeon Silver 4215	8	2	**95.6**	**108**
Xeon Silver 4214Y	12	2	132(est.)	140(est.)
	10	2	110(est.)	124(est.)
	8	2	94.9(est.)	113(est.)
Xeon Silver 4214	12	2	**132**	**139**
Xeon Silver 4210	10	2	**108**	**119**
Xeon Silver 4208	8	2	**81.5**	**93.3**
Xeon Bronze 3204	6	2	**38.9**	**55.0**
March 2020 released				
Xeon Gold 6258R	28	2	**330**	**274**
Xeon Gold 6256	12	2	**193**	**200**
Xeon Gold 6250	8	2	**136**	**155**
Xeon Gold 6248R	24	2	**303**	**261**
Xeon Gold 6246R	16	2	**236**	**229**
Xeon Gold 6242R	20	2	**273**	**247**
Xeon Gold 6240R	24	2	**274**	**242**
Xeon Gold 6238R	28	2	**294**	**253**
Xeon Gold 6230R	26	2	**273**	**239**
Xeon Gold 6226R	16	2	**206**	**201**
Xeon Gold 6208U	16	1	**108**	**105**
Xeon Gold 5220R	24	2	**257**	**227**
Xeon Gold 5218R	20	2	**217**	**200**
Xeon Silver 4215R	8	2	**100**	**109**
Xeon Silver 4214R	12	2	**133**	**145**
Xeon Silver 4210R	10	2	**108**	**121**
Xeon Bronze 3206R	8	2	**50.4**	**72.8**

SPECspeed2017				
Processor	**Cores**	**Number of Processors**	**SPECspeed2017 int_base**	**SPECspeed2017 fp_base**
April 2019 released				
Xeon Platinum 8280	28	2		157(est.)
Xeon Gold 6244	28	2	10.8(est.)	

The following two diagrams illustrate the throughput of the PRIMERGY RX2540 M5 in comparison to its predecessor PRIMERGY RX2540 M4, in their respective most performant configuration.

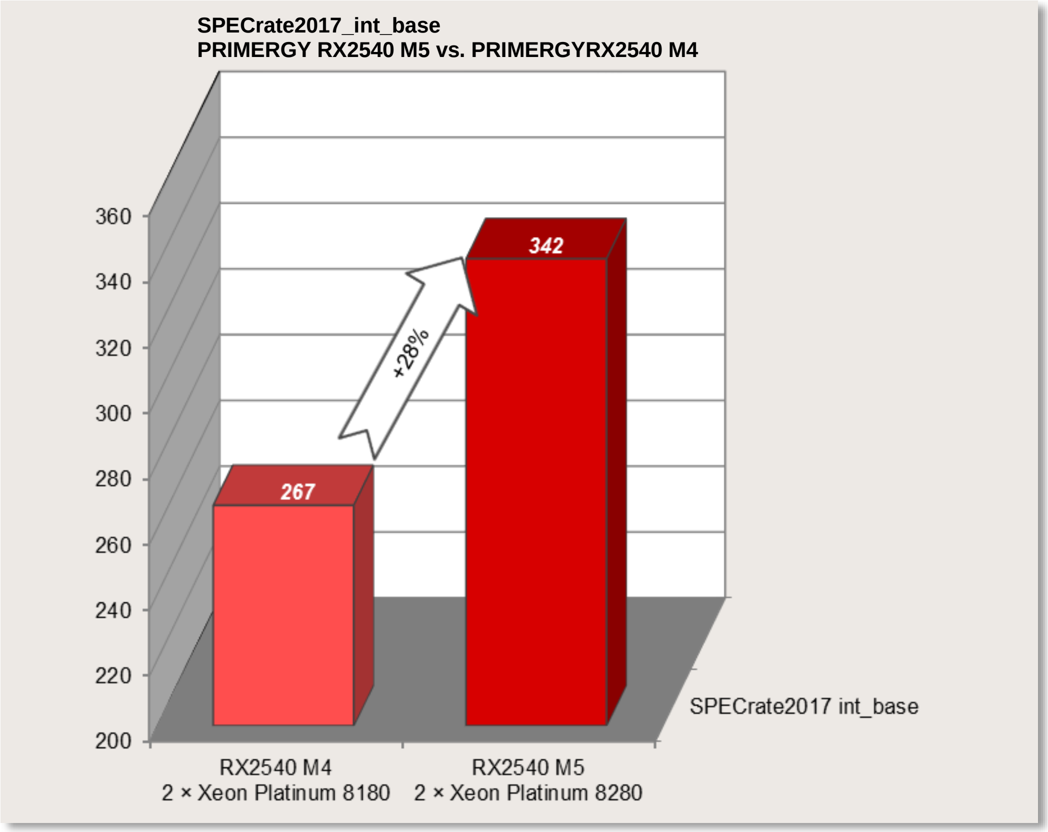

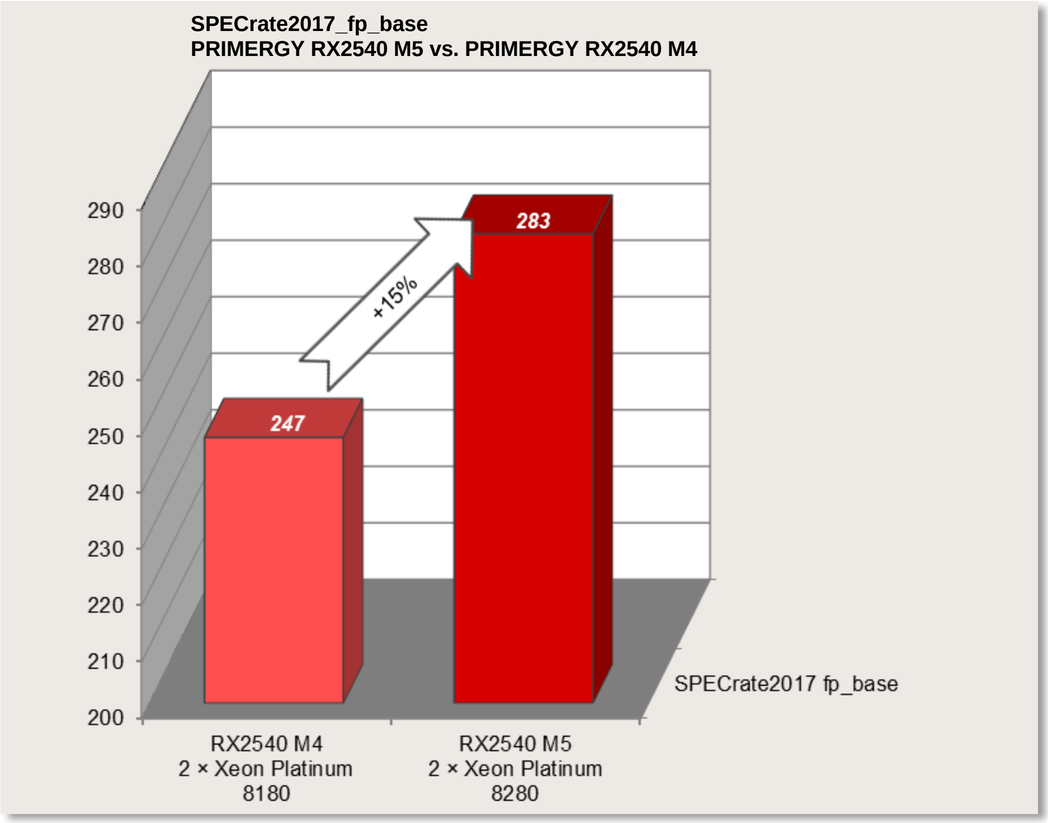

SPECpower_ssj2008

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.

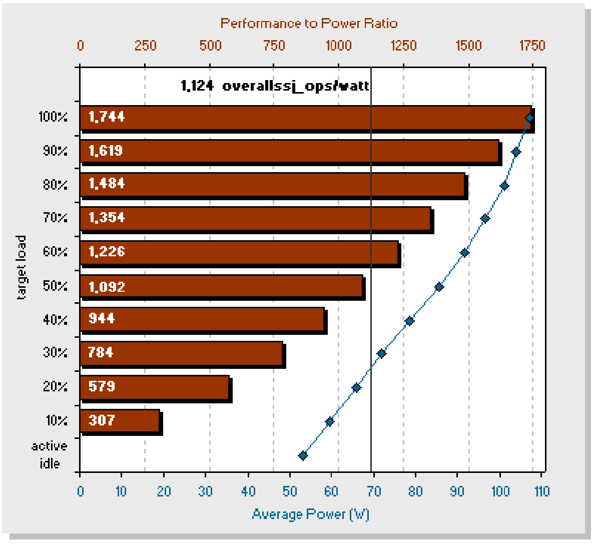

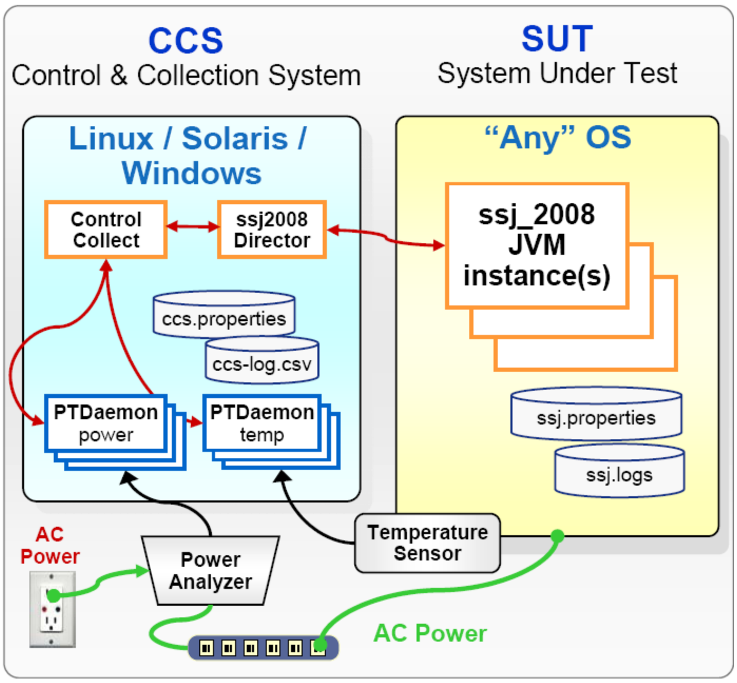

The benchmark runs on a wide variety of operating systems and hardware architectures, and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System 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 Windows OS measurement	
Hardware	
Model	PRIMERGY RX2540 M5
Processor	Intel® Xeon® Platinum 8276L
Memory	12 ×16 GB (1x16 GB) 2Rx8 PC4-2933Y-R
Network interface	2 x Intel® I350 Gigabit Network Connection (onboard)
Disk subsystem	1 x SSD M.2 240GB, S26361-F5706-E240
Power Supply Unit	1 × Fujitsu Technology Solutions S26113-F615-E10
Software	
BIOS	R1.8.0
BIOS settings	SATA Controller = Disabled. Serial Port = Disabled. Hardware Prefetcher = Disabled. Adjacent Cache Line Prefetch = Disabled. DCU Streamer Prefetcher = Disabled. Intel Virtualization Technology = Disabled. Turbo Mode = 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. Package C State limit = C6(Retention). HWPM = Disabled. USB Port Control = Enable internal ports only. Network Stack = Disabled. LAN Controller = LAN1.
Firmware	2.00c
Operating system	Microsoft Windows Server 2016 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. <N/A>: The test sponsor attests, as of date of publication, that CVE-2017-5754 (Meltdown) is mitigated in the system as tested and documented. <Yes>: The test sponsor attests, as of date of publication, that CVE-2017-5753 (Spectre variant 1) is mitigated in the system as tested and documented. <Yes>: The test sponsor attests, as of date of publication, that CVE-2017-5715 (Spectre variant 2) is mitigated in the system as tested and documented.
JVM	Oracle Java HotSpot(TM) 64-Bit Server VM 18.9(build 11+28, mixed mode), version 11
JVM settings	-server -Xmn1700m -Xms1950m -Xmx1950m -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:ParallelGCThreads=2 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:FreqInlineSize=2500 -XX:+UseLargePages -XX:+UseParallelOldGC -XX:AllocatePrefetchInstr=0 -XX:MinJumpTableSize=18 -XX:UseAVX=0

For Linux OS measurement	
Hardware	
Model	PRIMERGY RX2540 M5
Processor	Intel® Xeon® Platinum 8280L
Memory	12 ×16 GB (1x16 GB) 2Rx8 PC4-2933Y-R
Network interface	2 x Intel® I350 Gigabit Network Connection (onboard)
Disk subsystem	1 x SSD M.2 240GB, S26361-F5706-E240
Power Supply Unit	1 × Fujitsu Technology Solutions S26113-F615-E10
Software	
BIOS	R1.8.0
BIOS settings	SATA Controller = Disabled. Serial Port = Disabled. Hardware Prefetcher = Disabled. Adjacent Cache Line Prefetch = Disabled. DCU Streamer Prefetcher = Disabled. Intel Virtualization Technology = Disabled. Turbo Mode = 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. Package C State limit = C6(Retention). HWPM = Disabled. USB Port Control = Enable internal ports only. Network Stack = Disabled. LAN Controller = LAN1.
Firmware	2.00c
Operating system	SUSE Linux Enterprise Server 12 SP4 4.12.14-94.41-default

Operating system settings	kernal parameter:pcie_aspm=force pcie_aspm.policy=powersave intel_pstate=disable rcu_nocbs=1-111 nohz_full=1-111 isolcpus=1-111 modprobe cpufreq_conservative cpupower frequency-set --governor conservative echo -n 98 > /sys/devices/system/cpu/cpufreq/conservative/up_threshold echo -n 1 > /sys/devices/system/cpu/cpufreq/conservative/freq_step echo -n 1000000 > /sys/devices/system/cpu/cpufreq/conservative/sampling_rate echo -n 0 > /sys/devices/system/cpu/cpufreq/conservative/ignore_nice_load sysctl -w kernel.sched_migration_cost_ns=6000 echo -n 97 > /sys/devices/system/cpu/cpufreq/conservative/down_threshold echo -n 1 > /sys/devices/system/cpu/cpufreq/conservative/sampling_down_factor sysctl -w kernel.sched_min_granularity_ns=10000000 echo always > /sys/kernel/mm/transparent_hugepage/enabled powertop --auto-tune echo 0 > /proc/sys/kernel/nmi_watchdog cpupower frequency-set -u 2500MHz sysctl -w vm.swappiness=50 sysctl -w vm.laptop_mode=5 <Yes>: The test sponsor attests, as of date of publication, that CVE-2017-5754 (Meltdown) is mitigated in the system as tested and documented. <Yes>: The test sponsor attests, as of date of publication, that CVE-2017-5753 (Spectre variant 1) is mitigated in the system as tested and documented. <Yes>: The test sponsor attests, as of date of publication, that CVE-2017-5715 (Spectre variant 2) is mitigated in the system as tested and documented.
JVM	Oracle Java HotSpot(TM) 64-Bit Server VM (build 24.80-b11, mixed mode), version 1.7.0_80
JVM settings	-server -Xmn1700m -Xms1950m -Xmx1950m -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:LoopUnrollLimit=45 -XX:InitialTenuringThreshold=12 -XX:MaxTenuringThreshold=15 -XX:ParallelGCThreads=8 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:FreqInlineSize=2500 -XX:+AggressiveOpts -XX:+UseLargePages -XX:+UseParallelOldGC -XX:+UseHugeTLBFS -XX:+UseTransparentHugePages

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

Benchmark results(Windows)

The PRIMERGY RX2540 M5 in Microsoft Windows Server 2016 Standard achieved the following result:

SPECpower_ssj2008 = 11,991 overall ssj_ops/watt

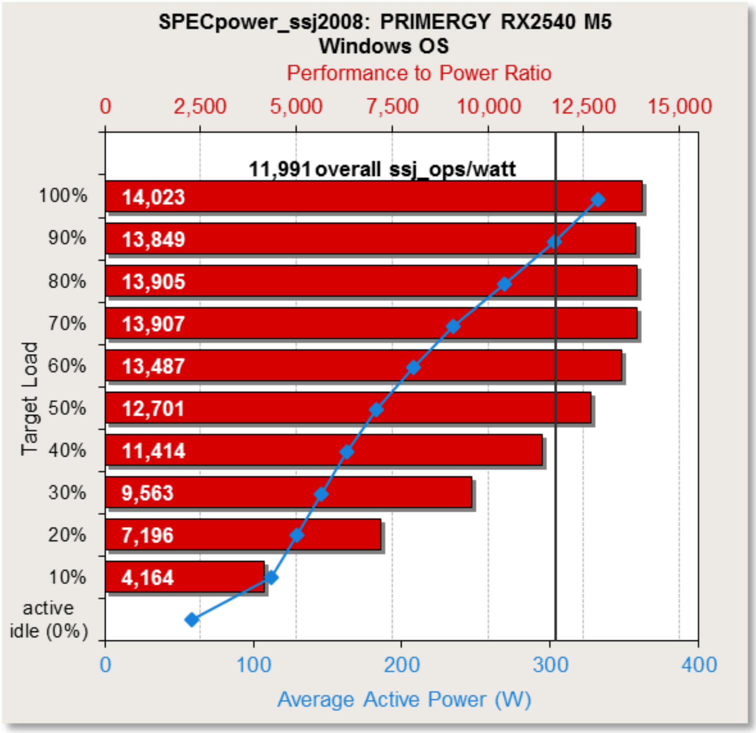

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 11,991 overall ssj_ops/watt for the PRIMERGYRX2540 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%	4,657,126	332	14,023
90%	4,195,308	303	13,849
80%	3,736,211	269	13,905
70%	3,265,708	235	13,907
60%	2,799,049	208	13,487
50%	2,327,926	183	12,701
40%	1,864,928	163	11,414
30%	1,398,828	146	9,563
20%	930,994	129	7,196
10%	465,123	112	4,164
Active Idle	0	58.2	0
			∑ssj_ops / ∑power = 11,991

Benchmark results(Linux)

The PRIMERGY RX2540 M5 in SUSE Linux Enterprise Server 12 SP4 achieved the following result:

SPECpower_ssj2008 = 13,502 overall ssj_ops/watt

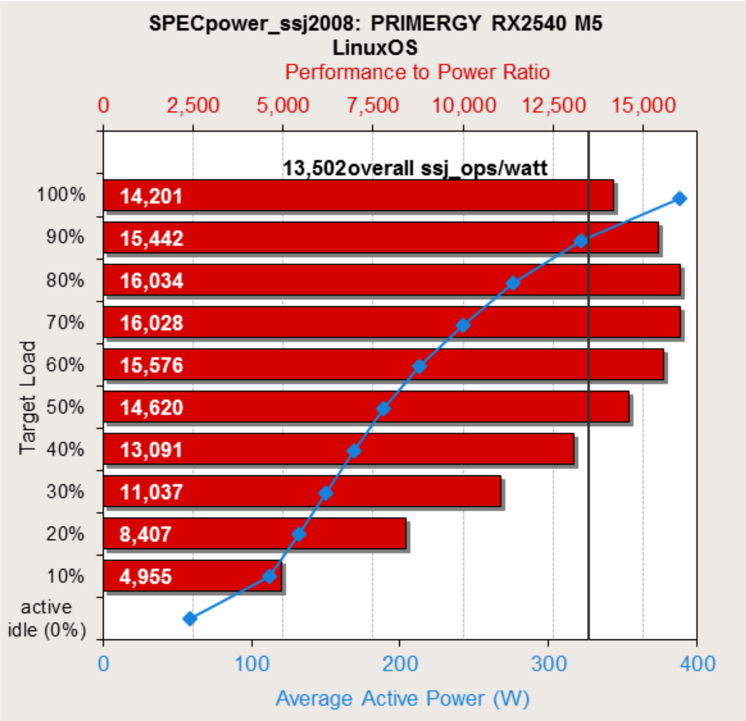

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 13,502 overall ssj_ops/watt for the PRIMERGYRX2540 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%	5,514,468	388	14,201
90%	4,979,547	322	15,442
80%	4,421,965	276	16,034
70%	3,873,008	242	16,028
60%	3,321,077	213	15,576
50%	2,759,717	189	14,620
40%	2,211,380	169	13,091
30%	1,660,248	150	11,037
20%	1,109,117	132	8,407
10%	555,330	112	4,955
Active Idle	0	58.3	0
			∑ssj_ops / ∑power = 13,502

The following diagram shows for each load level the power consumption (on the right y-axis) and the throughput (on the left y-axis) of the PRIMERGY RX2540 M5 compared to the predecessor PRIMERGY RX2540 M4.

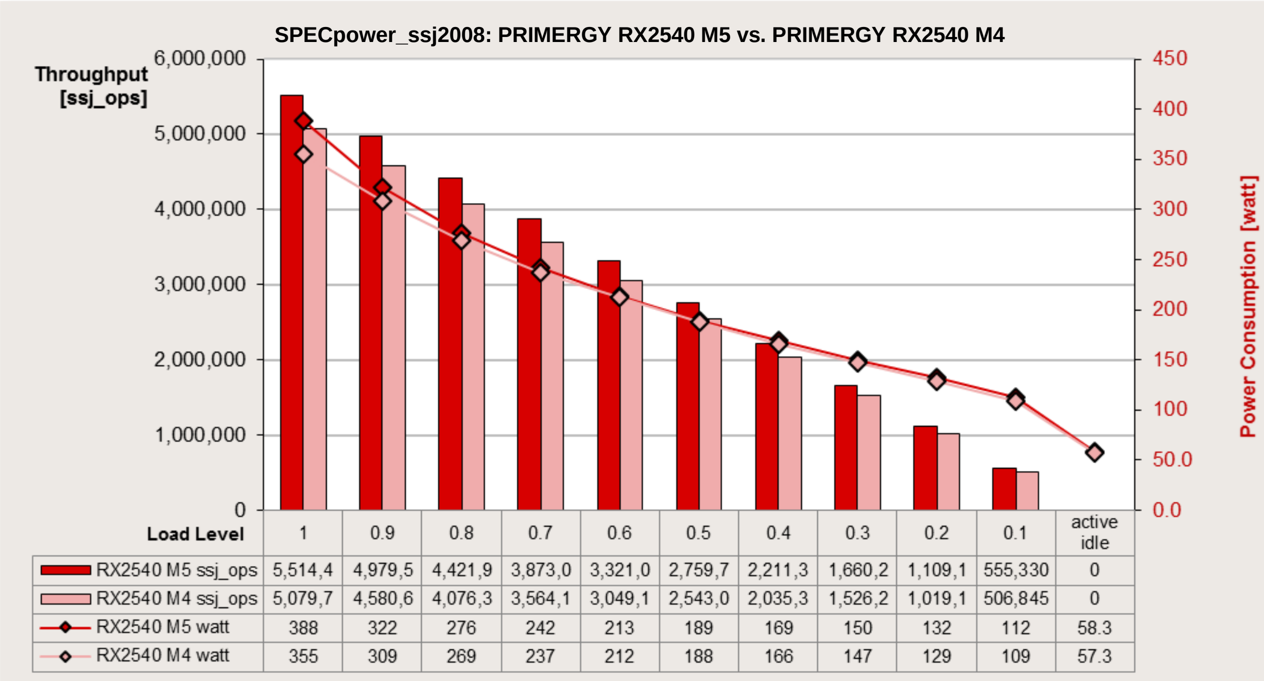

Load Level	1	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	active idle
RX2540 M5 ssj_ops	5,514,4	4,979,5	4,421,9	3,873,0	3,321,0	2,759,7	2,211,3	1,660,2	1,109,1	555,330	0
RX2540 M4 ssj_ops	5,079,7	4,580,6	4,076,3	3,564,1	3,049,1	2,543,0	2,035,3	1,526,2	1,019,1	506,845	0
RX2540 M5 watt	388	322	276	242	213	189	169	150	132	112	58.3
RX2540 M4 watt	355	309	269	237	212	188	166	147	129	109	57.3

Thanks to the latest Scalable Family processors, the PRIMERGY RX2540 M5 has a higher throughput. This results in an overall 5.3% increase in energy efficiency in the PRIMERGY RX2540 M5.

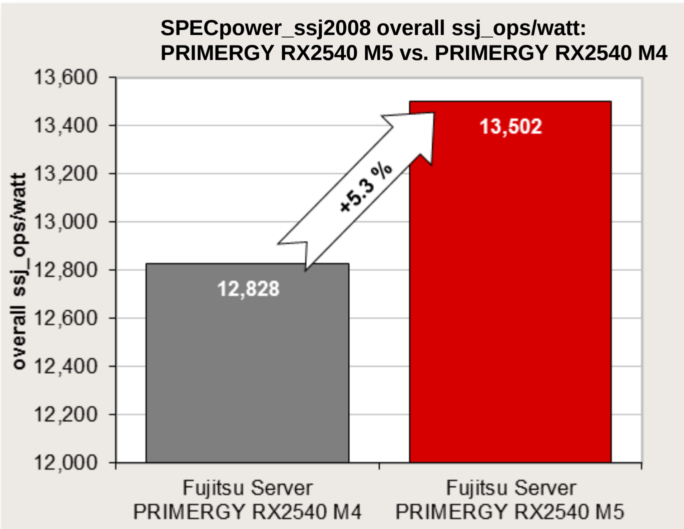

Difference of score by OS&JVM version

A score of SPECpower_ssj2008 differs about 10% in maximum caused by OS used in the system. OS has performance influence in itself. Thus depending on OS type, usable JVM versin is different. Currently combinations of Windows Server2012 R2&JVM7, Windows Server2016&JVM11 and Linux&JVM11 are used in Fujitsu and other vendor's submission results.
Under appropriate OS settings and JVM options, the score becomes high in order as Linux&JVM7≧ Windows Server2012 R2&JVM7>Windows Server2016&JVM11.
There is so few difference between Linux&JVM7 and Windows Server2012. On the other hand, a combination of Windows Server2016&JVM11's score is about 10% lower than the other two combination's score.
Under the rule of SPECpower_ssj2008, Windows Server2016, relatively new OS, is not allowed to measure with JVM7. Therefore it needs to use later JVM version. Alt-rt.jar , a module including in JVM7, is related to accelerate collection type HashMap. However, the module is deleted in JVM11. This is the main reason of SPECpower_ssj2008 score measured with JVM11 gets lower.

SPECjbb2015

Benchmark description

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

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

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

The SPECjbb2015 benchmark has a two performance metrics:

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

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

- SPECjbb2015 Composite

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

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

  Back-ends are existed on hosts separated from hosts on which the other components are running. Back-ends and the other components are connected by networks.

Results are not comparable to those in other categories.

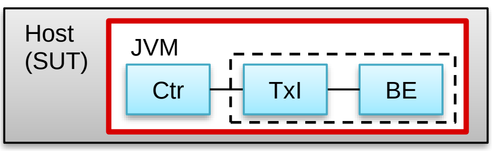

(a) example of SPECjbb2015 Composite configuration

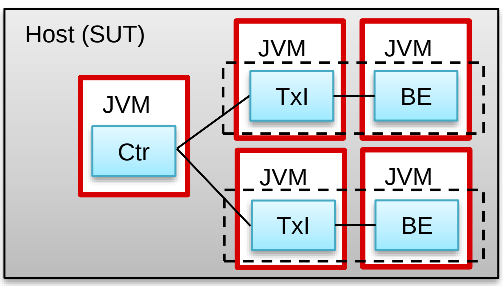

(b) example of SPECjbb2015 MultiJVM configuration

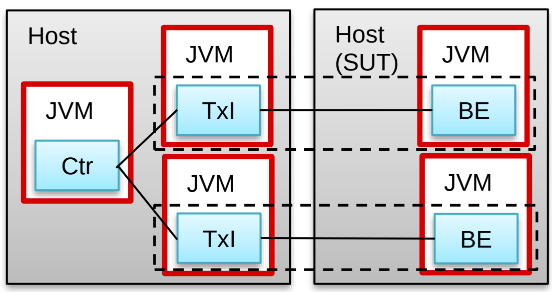

(c) example of SPECjbb2015 Distributed configuration

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

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

Benchmark environment

PRIMERGY RX2540 M5 was configured for the SPECjbb2015 Composite & MultiJVM benchmark measurement.

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M5
Processor	2 × Intel® Xeon® Platinum 8280M
Memory	24 × 32 GB (1x32 GB) 2Rx4 PC4-2933Y-R
Network interface	1 Gbit/s LAN
Disk subsystem	Disk : 1 × SSD SAS 12 Gb/s 2.5" 1.6TB
Software	
For measurement result (1)	
BIOS settings	Hyper-Threading set to Disabled Hardware Prefetcher set to Disabled Adjacent Cache Line Prefetch set to Disabled DCU Streamer Prefetcher set to Disabled Intel Virtualization Technology set to Disabled VT-d set to Disabled Override OS Energy Performance set to Enabled Energy Performance set to Performance Link Frequency Select set to 10.4GT/s Patrol Scrub set to Disable SNC set to Disabled Write CRC set to Disabled
Operating system	Windows Server 2016 Standard
Operating system settings	Power Options is set to High performance Processor scheduling is set to Programs Lock Pages in Memory is Enabled Performance Options is set to Adjust for best performance Total paging file size for all drives is set to 28,672 MB
JVM	Oracle Java SE 11.0.2
JVM settings	-server -Xms670g -Xmx670g -Xmn625g -XX:SurvivorRatio=100 -XX:MaxTenuringThreshold=15 -XX:+UseLargePages -XX:LargePageSizeInBytes=2m -XX:+UseParallelOldGC -Xnoclassgc -XX:+UseNUMA -XX:-UseBiasedLocking -XX:+AlwaysPreTouch -XX:-UseAdaptiveSizePolicy -XX:-UsePerfData -XX:TargetSurvivorRatio=95 -XX:ParallelGCThreads=56 -verbose:gc -XX:+AggressiveHeap
SPECjbb2015 settings	specjbb.comm.connect.timeouts.connect = 600000 specjbb.comm.connect.timeouts.read = 600000 specjbb.comm.connect.timeouts.write = 600000 specjbb.comm.connect.worker.pool.max = 64 specjbb.comm.connect.worker.pool.min = 64 specjbb.customerDriver.threads = {saturate=96} specjbb.forkjoin.workers = {Tier1=180, Tier2=28, Tier3=20} specjbb.mapreducer.pool.size = 4

For measurement result (2)	
BIOS settings	Hyper-Threading set to Disabled Hardware Prefetcher set to Disabled Adjacent Cache Line Prefetch set to Disabled DCU Streamer Prefetcher set to Disabled Intel Virtualization Technology set to Disabled VT-d set to Disabled Override OS Energy Performance set to Enabled Energy Performance set to Performance Link Frequency Select set to 10.4GT/s Patrol Scrub set to Disable SNC set to Disabled Write CRC set to Disabled
Operating system	Windows Server 2016 Standard
Operating system settings	Power Options is set to High performance Processor scheduling is set to Programs Lock Pages in Memory is Enabled Performance Options is set to Adjust for best performance Total paging file size for all drives is set to 28,672 MB
JVM	Oracle Java SE 11.0.2
JVM settings	-server -Xms670g -Xmx670g -Xmn625g -XX:SurvivorRatio=100 -XX:MaxTenuringThreshold=15 -XX:+UseLargePages -XX:LargePageSizeInBytes=2m -XX:+UseParallelOldGC -Xnoclassgc -XX:+UseNUMA -XX:-UseBiasedLocking -XX:+AlwaysPreTouch -XX:-UseAdaptiveSizePolicy -XX:-UsePerfData -XX:TargetSurvivorRatio=95 -XX:ParallelGCThreads=56 -verbose:gc -XX:+AggressiveHeap
SPECjbb2015 settings	specjbb.comm.connect.timeouts.connect = 600000 specjbb.comm.connect.timeouts.read = 600000 specjbb.comm.connect.timeouts.write = 600000 specjbb.comm.connect.worker.pool.max = 64 specjbb.comm.connect.worker.pool.min = 64 specjbb.customerDriver.threads = {saturate=96} specjbb.forkjoin.workers = {Tier1=180, Tier2=28, Tier3=20} specjbb.mapreducer.pool.size = 4
For measurement result (3)	
BIOS settings	Hardware Prefetcher set to Disabled Adjacent Cache Line Prefetch set to Disabled DCU Streamer Prefetcher set to Disabled Intel Virtualization Technology set to Disabled VT-d set to Disabled Override OS Energy Performance set to Enabled Energy Performance set to Performance Link Frequency Select set to 10.4GT/s Patrol Scrub set to Disable SNC set to Disabled Write CRC set to Disabled
Operating system	Windows Server 2016 Standard
Operating system settings	Power Options is set to High performance Processor scheduling is set to Programs Lock Pages in Memory is Enabled Performance Options is set to Adjust for best performance Total paging file size for all drives is set to 28,672 MB

JVM	Oracle Java SE 11.0.2
JVM settings Controller (Ctr)	-server -Xms2g -Xmx2g -Xmn1536m -XX:+UseParallelOldGC
JVM settings Backends (BE)	-server -Xms350g -Xmx350g -Xmn340g -XX:SurvivorRatio=50 -XX:+UseLargePages -XX:LargePageSizeInBytes=2m -XX:+UseParallelOldGC -XX:+AggressiveOpts -XX:+AlwaysPreTouch -XX:-UseAdaptiveSizePolicy -XX:-UsePerfData -XX:TargetSurvivorRatio=98 -XX:ParallelGCThreads=56 -XX:MaxTenuringThreshold=15
JVM settings Transaction Injector (TxI)	-server -Xms2g -Xmx2g -Xmn1536m -XX:+UseParallelOldGC -XX:+AlwaysPreTouch -XX:ParallelGCThreads=56
SPECjbb2015 settings	specjbb. specjbb.comm.connect.selector.runner.count = 4 specjbb.comm.connect.timeouts.connect = 600000 specjbb.comm.connect.timeouts.read = 600000 specjbb.comm.connect.timeouts.write = 600000 specjbb.comm.connect.worker.pool.max = 64 specjbb.comm.connect.worker.pool.min = 64 specjbb.forkjoin.workers = {Tier1=180, Tier2=1, Tier3=25} specjbb.group.count = 2 specjbb.txi.pergroup.count = 1

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

Benchmark results

Measurement result (1) : "SPECjbb2015 Composite" (April 12, 2019)

101,742 SPECjbb2015-Composite max-jOPS

67,948 SPECjbb2015-Composite critical-jOPS

On April 12, 2019 PRIMERGY RX2540 M5 with two Intel® Xeon® Platinum 8280M processor achieved the scores of 101,742 SPECjbb2015-Composite max-jOPS in Windows Server 2016 Standard. With the result, it ranked first in the 2-socket Windows server category for SPECjbb2015-Composite max-jOPS.

Measurement result (2) : "SPECjbb2015 Composite" (April 12, 2019)

98,065 SPECjbb2015-Composite max-jOPS

71,031 SPECjbb2015-Composite critical-jOPS

On April 12, 2019 PRIMERGY RX2540 M5 with two Intel® Xeon® Platinum 8280M processor achieved the scores of 71,031 SPECjbb2015-Composite critical-jOPS in Windows Server 2016 Standard. With the result, it ranked first in the 2-socket Windows server category for SPECjbb2015-Composite critical-jOPS.

Measurement result (3) : "SPECjbb2015 MultiJVM" (April 2, 2019)

155,295 SPECjbb2015-MultiJVM max-jOPS

81,233 SPECjbb2015-MultiJVM critical-jOPS

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

SAP SD

Benchmark description

The SAP application software consists of modules used to manage all standard business processes. These include modules for ERP (Enterprise Resource Planning), such as Assemble-to-Order (ATO), Financial Accounting (FI), Human Resources (HR), Materials Management (MM), Production Planning (PP), and Sales and Distribution (SD), as well as modules for SCM (Supply Chain Management), Retail, Banking, Utilities, BI (Business Intelligence), CRM (Customer Relation Management) or PLM (Product Lifecycle Management).

The application software is always based on a database so that a SAP configuration consists of the hardware, the software components operating system, the database, and the SAP software itself.

SAP AG has developed SAP Standard Application Benchmarks in order to verify the performance, stability and scaling of a SAP application system. The benchmarks, of which SD Benchmark is the most commonly used and most important, analyze the performance of the entire system and thus measure the quality of the integrated individual components.

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 entire specification of the benchmark developed by SAP AG, Walldorf, Germany, can be found at: http://www.sap.com/benchmark.

Benchmark environment

The typical measurement set-up is illustrated below:

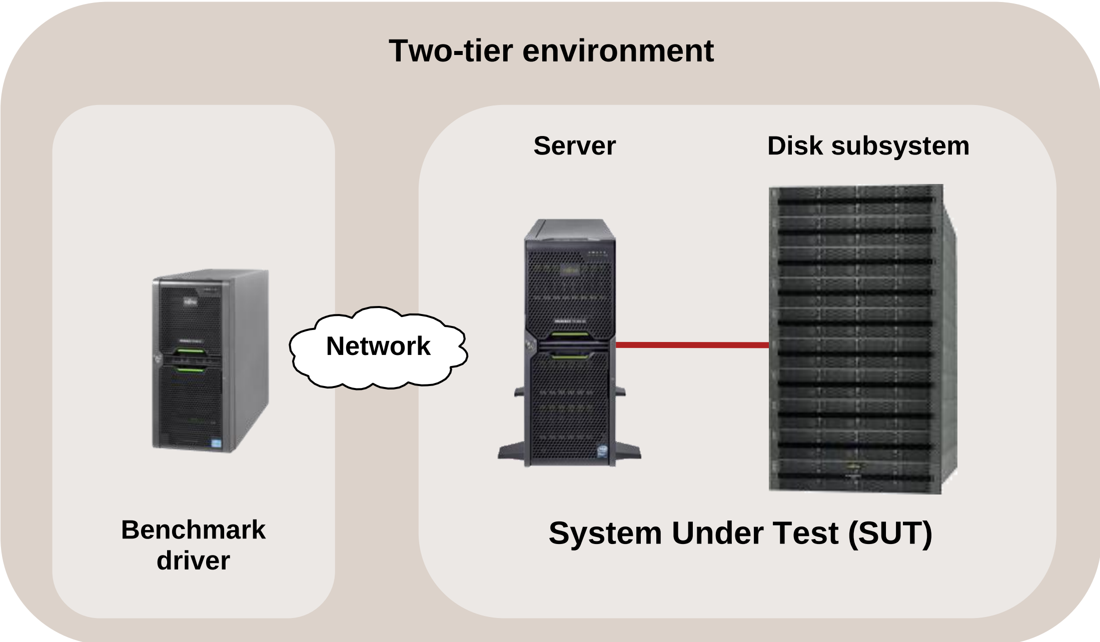

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M5
Processor	2 × Xeon Platinum 8280
Memory	24 × 32 GB (1x32 GB) 2Rx4 DDR4-2933 R ECC
Network interface	1 Gbit/s LAN
Disk subsystem	PRIMERGY RX2540 M5: 1 × HDD SAS 12 Gb/s 2.5" 15K 600 GB 1 × PCIe-SSD 750 GB
Software	
BIOS settings	Enable SNC
Operating system	Windows Server 2016 Standard
Database	Microsoft SQL Server 2012 (64-bit)
SAP Business Suite Software	SAP enhancement package 5 for SAP ERP 6.0

Benchmark driver	
Hardware	
Model	PRIMERGY RX2540 M2
Processor	2 × Xeon E5-2637 v4
Memory	256 GB
Network interface	1 Gbit/s LAN
Software	
Operating system	SUSE Linux Enterprise Server 12 SP2

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

Benchmark results

Certification number 2019010	
Number of SAP SD benchmark users	30,500
Average dialog response time	0.99 seconds
Throughput Fully processed order line items/hour Dialog steps/hour SAPS	 3,330,330 9,991,000 166,520
Average database request time (dialog/update)	0.012 sec / 0.027 sec
CPU utilization of central server	99%
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 M5, 2 processors / 56 cores / 112 threads, Intel Xeon Platinum 8280 Processor, 2.70 GHz, 64 KB L1 cache per core and 1,024 KB L2 cache per core, 38.5 MB L3 cache per processor, 768 GB main memory

The following chart shows a comparison of two-tier SAP SD Standard Application Benchmark results for 2-way Xeon Processor Scalable Family based servers with Windows OS and SQL Server database (as of April 2, 2019). The PRIMERGY RX2540 M5 outperforms the comparably configured servers from HPE. The latest SAP SD 2-tier results can be found at https://www.sap.com/dmc/exp/2018-benchmark-directory/#/sd

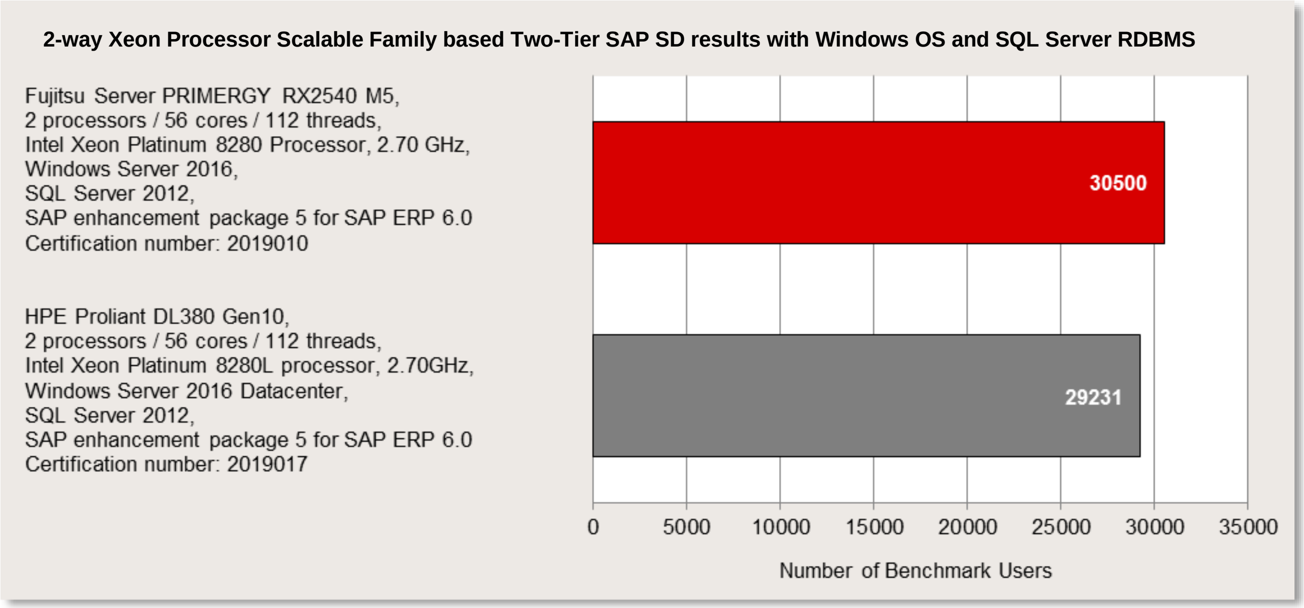

Disk I/O: Performance of storage media

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. 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:

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

A given 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		
File copy	random	50%	50%	64	Copying of files
File server	random	67%	33%	64	File server
Database	random	67%	33%	8	Database (data transfer) Mail server
Streaming	sequential	100%	0%	64	Database (log file), Data backup; Video streaming (partial)
Restore	sequential	0%	100%	64	Restoring of files

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

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

The main results of a measurement are:

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

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

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

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

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

Benchmark environment

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

System Under Test (SUT)			
Hardware			
3.5 inch Model:			
Controller: 1x PRAID CP400i			
	Storage media	Category	Drive Name
	HDD	SAS HDD(SAS 12Gbps, 10krpm)[512e]	AL15SEB18EQ *2 *3
		SAS HDD(SAS 12Gbps, 10krpm)[512n]	AL15SEB030N *2 *3
		SAS HDD(SAS 12Gbps, 15krpm)[512n]	ST300MP0006 *1 *3
		NL-SAS HDD(SAS 12Gbps, 7.2krpm)[512e]	HUH721212AL5204 *2 *3
		NL-SAS HDD(SAS 12Gbps, 7.2krpm)[512n]	ST2000NM0045 *1 *3
		BC-SATA HDD(SATA 6Gbps, 7.2krpm)[512e]	ST6000NM0115 *1 *3 HUH721212ALE604 *2 *3
		BC-SATA HDD(SATA 6Gbps, 7.2krpm)[512n]	HUS722T1TALA604 *2 *3 ST2000NM0055 *1 *3
	SSD	SAS SSD(SAS 12Gbps, Write Intensive)	KPM51MUG400G *2 *3 KPM51MUG800G *2 *3 KPM51MUG1T60 *2 *3
		SAS SSD(SAS 12Gbps, Mixed Use)	WUSTR6440ASS204 *2 *3 WUSTR6480ASS204 *2 *3 WUSTR6416ASS204 *2 *3 WUSTR6432ASS204 *2 *3
		SAS SSD(SAS 12Gbps, Read Intensive)	WUSTR1548ASS204 *2 *3 WUSTR1596ASS204 *2 *3 WUSTR1519ASS204 *2 *3 WUSTR1538ASS204 *2 *3 WUSTR1576ASS204 *2 *3
		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(R) Xeon(R) Gold 5222 (3.80GHz)			
	Storage media	Category	Drive Name
	SSD	PCIe SSD AIC(Write Intensive)	SSDPED1K375GA *2 *4 SSDPED1K750GA *2 *4

Controller: Intel(R) C620 Standard SATA AHCI controller			
	Storage media	Category	Drive Name
	SSD	M.2 Flash Module	MTFDDAV240TCB *2 *4 MTFDDAV480TCB *2 *4

2.5 inch Model:			
Controller: 1x PRAID CP400i			
	Storage media	Category	Drive Name
	HDD	SAS HDD(SAS 12Gbps, 10krpm)[512e]	AL15SEB06EQ *2 *3
		SAS HDD(SAS 12Gbps, 10krpm)[512n]	AL15SEB030N *2 *3
		SAS HDD(SAS 12Gbps, 15krpm)[512n]	ST300MP0006 *1 *3
		NL-SAS HDD (SAS 12Gbps, 7.2krpm)[512n]	ST1000NX0453 *1 *3
		BC-SATA HDD(SATA 6Gbps, 7.2krpm)[512e]	ST1000NX0313 *1 *3
		BC-SATA HDD(SATA 6Gbps, 7.2krpm)[512n]	ST2000NX0403 *1 *3
	SSD	SAS SSD(SAS 12Gbps, Write Intensive)	KPM51MUG400G *2 *3 KPM51MUG800G *2 *3 KPM51MUG1T60 *2 *3
		SAS SSD(SAS 12Gbps, Mixed Use)	WUSTR6440ASS204 *2 *3 WUSTR6480ASS204 *2 *3 WUSTR6416ASS204 *2 *3 WUSTR6432ASS204 *2 *3 WUSTR6464ASS204 *2 *3
		SAS SSD(SAS 12Gbps, Read Intensive)	WUSTR1548ASS204 *2 *3 WUSTR1596ASS204 *2 *3 WUSTR1519ASS204 *2 *3 WUSTR1538ASS204 *2 *3 WUSTR1576ASS204 *2 *3 WUSTR1515ASS204 *2 *3
		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(R) Xeon(R) 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(R) C620 Standard SATA AHCI controller			
	Storage media	Category	Drive Name
	SSD	M.2 Flash Module	MTFDDAV240TCB *2 *4 MTFDDAV480TCB *2 *4

*1) The operating system uses Microsoft Windows Server 2012 Standard R2.
*2) The operating system uses Microsoft Windows Server 2016 Standard.
*3) Measurement area is type 1.
*4) Measurement area is type 2.

Software		
Operating system		Microsoft Windows Server 2012 Standard R2 Microsoft Windows Server 2016 Standard
Benchmark version		3.0
RAID type		Logical drive of type RAID 0 consisting of 1 hard disk
Stripe size		Controller default (here 64 kB)
Measuring tool		Iometer 1.1.0
Measurement area	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 32GB 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 / 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	-	PCIe 3.0 x8	SATA 6G SAS 12G	0, 1, 1E, 10, 5, 50
PCIe SSD	Integrated PCI Express controller	-	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 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 inch, or Add in card

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 modi “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 (# 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.

3.5 inch model storage media

HDDs

Random accesses (units: IO/s):

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]		
			Database	Fileserver	filecopy
1,800	AL15SEB18EQ	SAS 12G	600	512	547
300	AL15SEB030N	SAS 12G	645	546	568
300	ST300MP0006	SAS 12G	768	662	472
12,000	HUH721212AL5204	SAS 12G	396	339	364
2,000	ST2000NM0045	SAS 12G	376	336	343
6,000	ST6000NM0115	SATA 6G	392	362	371
12,000	HUH721212ALE604	SATA 6G	350	313	341
1,000	HUS722T1TALA604	SATA 6G	287	264	269
2,000	ST2000NM0055	SATA 6G	339	301	314

Sequential accesses (units: MB/s):

Capacity [GB]	Storage device	Inter face	Throughput [MB/s]	
			Streaming	Restore
1,800	AL15SEB18EQ	SAS 12G	258	255
300	AL15SEB030N	SAS 12G	231	230
300	ST300MP0006	SAS 12G	304	304
12,000	HUH721212AL5204	SAS 12G	245	244
2,000	ST2000NM0045	SAS 12G	206	206
6,000	ST6000NM0115	SATA 6G	213	208
12,000	HUH721212ALE604	SATA 6G	246	246
1,000	HUS722T1TALA604	SATA 6G	201	201
2,000	ST2000NM0055	SATA 6G	196	195

SSDs

Random accesses (units: IO/s):

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]		
			Database	Fileserver	filecopy
400	KPM51MUG400G	SAS 12G	84,469	13,329	13,677
800	KPM51MUG800G	SAS 12G	99,728	14,549	18,049
1,600	KPM51MUG1T60	SAS 12G	108,428	17,243	19,634
400	WUSTR6440ASS204	SAS 12G	83,427	14,459	13,924
800	WUSTR6480ASS204	SAS 12G	94,899	22,414	21,187
1,600	WUSTR6416ASS204	SAS 12G	97,107	24,053	22,802
3,200	WUSTR6432ASS204	SAS 12G	106,745	23,975	22,793
480	WUSTR1548ASS204	SAS 12G	77,846	11,663	9,904
960	WUSTR1596ASS204	SAS 12G	88,384	18,834	16,636
1,920	WUSTR1519ASS204	SAS 12G	89,397	21,635	21,597
3,840	WUSTR1538ASS204	SAS 12G	99,644	23,727	22,831
7,680	WUSTR1576ASS204	SAS 12G	106,933	23,688	22,644
240	MZ7KH240HAHQ	SATA 6G	49,159	7,313	7,431
480	MZ7KH480HAHQ	SATA 6G	50,558	7,774	7,810
960	MZ7KH960HAJR	SATA 6G	50,647	7,793	7,916
1,920	MZ7KH1T9HAJR	SATA 6G	50,702	8,040	7,960
3,840	MZ7KH3T8HALS	SATA 6G	50,766	8,039	7,936
240	MTFDDAK240TCB	SATA 6G	18,959	3,367	4,516
480	MTFDDAK480TDC	SATA 6G	24,710	3,799	5,006
960	MTFDDAK960TDC	SATA 6G	30,152	4,625	5,553
1,920	MTFDDAK1T9TDC	SATA 6G	37,234	5,606	5,566
3,840	MTFDDAK3T8TDC	SATA 6G	41,711	6,429	6,133
7,680	MTFDDAK7T6TDC	SATA 6G	40,683	6,874	6,672
375	SSDPED1K375GA	PCIe3 x4	212,118	37,121	36,123
750	SSDPED1K750GA	PCIe3 x4	209,628	37,592	36,941
240	MTFDDAV240TCB	SATA 6G	19,773	3,844	4,968
480	MTFDDAV480TCB	SATA 6G	22,258	4,935	6,294

Sequential accesses (units: MB/s):

Capacity [GB]	Storage device	Inter face	Throughput [MB/s]	
			Streaming	Restore
400	KPM51MUG400G	SAS 12G	1,056	1,041
800	KPM51MUG800G	SAS 12G	1,056	1,042
1,600	KPM51MUG1T60	SAS 12G	1,057	1,042
400	WUSTR6440ASS204	SAS 12G	1,073	626
800	WUSTR6480ASS204	SAS 12G	1,073	1,008
1,600	WUSTR6416ASS204	SAS 12G	1,073	1,029
3,200	WUSTR6432ASS204	SAS 12G	1,073	1,030
480	WUSTR1548ASS204	SAS 12G	1,055	554
960	WUSTR1596ASS204	SAS 12G	1,067	965
1,920	WUSTR1519ASS204	SAS 12G	1,073	1,030
3,840	WUSTR1538ASS204	SAS 12G	1,073	1,030
7,680	WUSTR1576ASS204	SAS 12G	1,073	1,030
240	MZ7KH240HAHQ	SATA 6G	526	486
480	MZ7KH480HAHQ	SATA 6G	526	485
960	MZ7KH960HAJR	SATA 6G	525	485
1,920	MZ7KH1T9HAJR	SATA 6G	526	485
3,840	MZ7KH3T8HALS	SATA 6G	526	485
240	MTFDDAK240TCB	SATA 6G	487	258
480	MTFDDAK480TDC	SATA 6G	507	362
960	MTFDDAK960TDC	SATA 6G	507	440
1,920	MTFDDAK1T9TDC	SATA 6G	507	483
3,840	MTFDDAK3T8TDC	SATA 6G	504	481
7,680	MTFDDAK7T6TDC	SATA 6G	469	482
375	SSDPED1K375GA	PCIe3 x4	2,460	2,197
750	SSDPED1K750GA	PCIe3 x4	2,546	2,296
240	MTFDDAV240TCB	SATA 6G	487	258
480	MTFDDAV480TCB	SATA 6G	509	403

2.5 inch model storage media

HDDs

Random accesses (units: IO/s):

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]		
			Database	Fileserver	filecopy
600	AL15SEB06EQ	SAS 12G	592	516	544
300	AL15SEB030N	SAS 12G	645	546	568
300	ST300MP0006	SAS 12G	768	662	472
1,000	ST1000NX0453	SAS 12G	371	321	306
1,000	ST1000NX0313	SATA 6G	324	281	288
2,000	ST2000NX0403	SATA 6G	326	286	294

Sequential accesses (units: MB/s):

Capacity [GB]	Storage device	Inter face	Throughput [MB/s]	
			Streaming	Restore
600	AL15SEB06EQ	SAS 12G	260	260
300	AL15SEB030N	SAS 12G	231	230
300	ST300MP0006	SAS 12G	304	304
1,000	ST1000NX0453	SAS 12G	137	137
1,000	ST1000NX0313	SATA 6G	131	131
2,000	ST2000NX0403	SATA 6G	133	133

SSDs

Random accesses (units: IO/s):

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]		
			Database	Fileserver	filecopy
400	KPM51MUG400G	SAS 12G	84,469	13,329	13,677
800	KPM51MUG800G	SAS 12G	99,728	14,549	18,049
1,600	KPM51MUG1T60	SAS 12G	108,428	17,243	19,634
400	WUSTR6440ASS204	SAS 12G	83,427	14,459	13,924
800	WUSTR6480ASS204	SAS 12G	94,899	22,414	21,187
1,600	WUSTR6416ASS204	SAS 12G	97,107	24,053	22,802
3,200	WUSTR6432ASS204	SAS 12G	106,745	23,975	22,793
6,400	WUSTR6464ASS204	SAS 12G	111,695	23,911	22,639
480	WUSTR1548ASS204	SAS 12G	77,846	11,663	9,904
960	WUSTR1596ASS204	SAS 12G	88,384	18,834	16,636
1,920	WUSTR1519ASS204	SAS 12G	89,397	21,635	21,597
3,840	WUSTR1538ASS204	SAS 12G	99,644	23,727	22,831
7,680	WUSTR1576ASS204	SAS 12G	106,933	23,688	22,644
15,360	WUSTR1515ASS204	SAS 12G	107,687	23,590	22,686
240	MZ7KH240HAHQ	SATA 6G	49,159	7,313	7,431
480	MZ7KH480HAHQ	SATA 6G	50,558	7,774	7,810
960	MZ7KH960HAJR	SATA 6G	50,647	7,793	7,916
1,920	MZ7KH1T9HAJR	SATA 6G	50,702	8,040	7,960
3,840	MZ7KH3T8HALS	SATA 6G	50,766	8,039	7,936
240	MTFDDAK240TCB	SATA 6G	18,959	3,367	4,516
480	MTFDDAK480TDC	SATA 6G	24,710	3,799	5,006
960	MTFDDAK960TDC	SATA 6G	30,152	4,625	5,553
1,920	MTFDDAK1T9TDC	SATA 6G	37,234	5,606	5,566
3,840	MTFDDAK3T8TDC	SATA 6G	41,711	6,429	6,133
7,680	MTFDDAK7T6TDC	SATA 6G	40,683	6,874	6,672
750	SSDPE21K750GA	PCIe3 x4	214,231	37,611	36,957
1,600	SSDPE2KE016T8	PCIe3 x4	135,500	41,066	37,080
3,200	SSDPE2KE032T8	PCIe3 x4	136,782	48,210	45,348
6,400	SSDPE2KE064T8	PCIe3 x4	192,245	51,767	51,438
375	SSDPED1K375GA	PCIe3 x4	212,118	37,121	36,123
750	SSDPED1K750GA	PCIe3 x4	209,628	37,592	36,941
240	MTFDDAV240TCB	SATA 6G	19,773	3,844	4,968
480	MTFDDAV480TCB	SATA 6G	22,258	4,935	6,294

Sequential accesses (units: MB/s):

Capacity [GB]	Storage device	Inter face	Throughput [MB/s]	
			Streaming	Restore
400	KPM51MUG400G	SAS 12G	1,056	1,041
800	KPM51MUG800G	SAS 12G	1,056	1,042
1,600	KPM51MUG1T60	SAS 12G	1,057	1,042
400	WUSTR6440ASS204	SAS 12G	1,073	626
800	WUSTR6480ASS204	SAS 12G	1,073	1,008
1,600	WUSTR6416ASS204	SAS 12G	1,073	1,029
3,200	WUSTR6432ASS204	SAS 12G	1,073	1,030
6,400	WUSTR6464ASS204	SAS 12G	1,073	1,030
480	WUSTR1548ASS204	SAS 12G	1,055	554
960	WUSTR1596ASS204	SAS 12G	1,067	965
1,920	WUSTR1519ASS204	SAS 12G	1,073	1,030
3,840	WUSTR1538ASS204	SAS 12G	1,073	1,030
7,680	WUSTR1576ASS204	SAS 12G	1,073	1,030
15,360	WUSTR1515ASS204	SAS 12G	1,073	1,029
240	MZ7KH240HAHQ	SATA 6G	526	486
480	MZ7KH480HAHQ	SATA 6G	526	485
960	MZ7KH960HAJR	SATA 6G	525	485
1,920	MZ7KH1T9HAJR	SATA 6G	526	485
3,840	MZ7KH3T8HALS	SATA 6G	526	485
240	MTFDDAK240TCB	SATA 6G	487	258
480	MTFDDAK480TDC	SATA 6G	507	362
960	MTFDDAK960TDC	SATA 6G	507	440
1,920	MTFDDAK1T9TDC	SATA 6G	507	483
3,840	MTFDDAK3T8TDC	SATA 6G	504	481
7,680	MTFDDAK7T6TDC	SATA 6G	469	482
750	SSDPE21K750GA	PCIe3 x4	2,546	2,295
1,600	SSDPE2KE016T8	PCIe3 x4	3,213	1,917
3,200	SSDPE2KE032T8	PCIe3 x4	3,209	2,800
6,400	SSDPE2KE064T8	PCIe3 x4	3,205	3,048
375	SSDPED1K375GA	PCIe3 x4	2,460	2,197
750	SSDPED1K750GA	PCIe3 x4	2,546	2,296
240	MTFDDAV240TCB	SATA 6G	487	258
480	MTFDDAV480TCB	SATA 6G	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) 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 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:

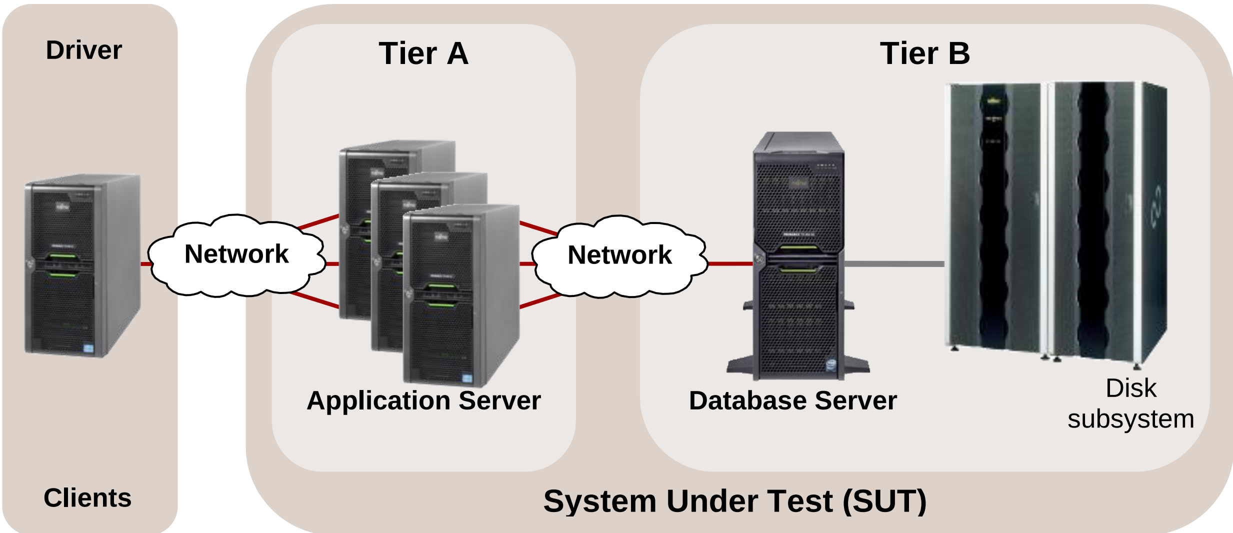

All OLTP-2 results were Calculated based on the configuration of the next following pages of PRIMERGY RX2540 M5

Database Server (Tier B)	
Hardware	
Model	PRIMERGY RX2540 M5
Processor	2nd Generation Intel® Xeon® Scalable Processors Family
Memory	1 processor: 12 ×64 GB (1x64 GB) 2Rx4 DDR4-2933 ECC 2 processors:24 ×64 GB (1x64 GB) 2Rx4 DDR4-2933 ECC
Network interface	1 × Dual Port onboard LAN 10 Gb/s
Disk subsystem	RX2540 M5: Onboard RAID controller PRAID EP420i 2 × 300 GB 10k rpm 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 10k rpm 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 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 1536 GB was considered for the measurements with two processors and a configuration with a total memory of 768 GB for the measurements with one processor. Both memory configurations have memory access of 2933 MHz.

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

Processor	Cores	Threads	2CPU Score	1CPU Score
April 2019 released				
Xeon Platinum 8280L	28	56	6,732 (est.)	3,699 (est.)
Xeon Platinum 8280M	28	56	6,732 (est.)	3,699 (est.)
Xeon Platinum 8280	28	56	6,732 (est.)	3,699 (est.)
Xeon Platinum 8276L	28	56	5,848 (est.)	3,213 (est.)
Xeon Platinum 8276M	28	56	5,848 (est.)	3,213 (est.)
Xeon Platinum 8276	28	56	5,848 (est.)	3,213 (est.)
Xeon Platinum 8270	26	52	6,233 (est.)	3,425 (est.)
Xeon Platinum 8268	24	48	5,984 (est.)	3,288 (est.)
Xeon Platinum 8260L	24	48	5,440 (est.)	2,989 (est.)
Xeon Platinum 8260M	24	48	5,440 (est.)	2,989 (est.)
Xeon Platinum 8260Y	24	48	5,440 (est.)	2,989 (est.)
	20	40	4,734 (est.)	2,601 (est.)
	16	32	4,104 (est.)	2,255 (est.)
Xeon Platinum 8260	24	48	5,440 (est.)	2,989 (est.)
Xeon Gold 6262V	24	48	4,805 (est.)	2,640 (est.)
Xeon Gold 6254	18	36	4,987 (est.)	2,740 (est.)
Xeon Gold 6252	24	48	5,191 (est.)	2,852 (est.)
Xeon Gold 6248	20	40	4,760 (est.)	2,615 (est.)
Xeon Gold 6246	12	24	3,547 (est.)	1,949 (est.)
Xeon Gold 6244	8	16	2,697 (est.)	1,482 (est.)
Xeon Gold 6242	16	32	4,148 (est.)	2,279 (est.)
Xeon Gold 6240L	18	36	4,352 (est.)	2,391 (est.)
Xeon Gold 6240M	18	36	4,352 (est.)	2,391 (est.)
Xeon Gold 6240Y	18	36	4,352 (est.)	2,391 (est.)
	14	28	3,559 (est.)	1,955 (est.)
	8	16	2,224 (est.)	1,222 (est.)
Xeon Gold 6240	18	36	4,352 (est.)	2,391 (est.)
Xeon Gold 6238L	22	44	4,533 (est.)	2,491 (est.)
Xeon Gold 6238M	22	44	4,533 (est.)	2,491 (est.)
Xeon Gold 6238	22	44	4,533 (est.)	2,491 (est.)
Xeon Gold 6234	8	16	2,516 (est.)	1,382 (est.)
Xeon Gold 6230	20	40	4,284 (est.)	2,354 (est.)
Xeon Gold 6226	12	24	3,196 (est.)	1,756 (est.)
Xeon Gold 6222V	20	40	3,876 (est.)	2,130 (est.)
Xeon Gold 6212U	24	48		2,765 (est.)

Xeon Gold 6210U	20	40		2,448 (est.)
Xeon Gold 6209U	20	40		2,132 (est.)
Xeon Gold 5222	4	8	1,321 (est.)	726 (est.)
Xeon Gold 5220S	18	36	4,035 (est.)	2,217 (est.)
Xeon Gold 5220	18	36	4,035 (est.)	2,217 (est.)
Xeon Gold 5218B	16	32	3,672 (est.)	2,018 (est.)
Xeon Gold 5218	16	32	3,672 (est.)	2,018 (est.)
Xeon Gold 5217	8	16	2,171 (est.)	1,193 (est.)
Xeon Gold 5215L	10	20	2,403 (est.)	1,320 (est.)
Xeon Gold 5215M	10	20	2,403 (est.)	1,320 (est.)
Xeon Gold 5215	10	20	2,403 (est.)	1,320 (est.)
Xeon Silver 4216	16	32	3,332 (est.)	1,831 (est.)
Xeon Silver 4215	8	16	1,952 (est.)	1,072 (est.)
Xeon Silver 4214Y	12	24	2,607 (est.)	1,432 (est.)
	10	20	2,172 (est.)	1,194 (est.)
	8	16	1,876 (est.)	1,031 (est.)
Xeon Silver 4214	12	24	2,607 (est.)	1,432 (est.)
Xeon Silver 4210	10	20	2,221 (est.)	1,221 (est.)
Xeon Silver 4208	8	16	1,675 (est.)	920 (est.)
Xeon Bronze 3204	6	6	811 (est.)	446 (est.)
March 2020 released				
Xeon Gold 6258R	28	56	6,499 (est.)	3,571 (est.)
Xeon Gold 6256	12	24	3,767 (est.)	2,070 (est.)
Xeon Gold 6250	8	16	2,804 (est.)	1,541 (est.)
Xeon Gold 6248R	24	48	5,957 (est.)	3,273 (est.)
Xeon Gold 6246R	16	32	4,572 (est.)	2,512 (est.)
Xeon Gold 6242R	20	40	5,217 (est.)	2,867 (est.)
Xeon Gold 6240R	24	48	5,239 (est.)	2,878 (est.)
Xeon Gold 6238R	28	56	5,656 (est.)	3,108 (est.)
Xeon Gold 6230R	26	52	5,336 (est.)	2,932 (est.)
Xeon Gold 6226R	16	32	4,001 (est.)	2,198 (est.)
Xeon Gold 6208U	16	32		1,153 (est.)
Xeon Gold 5220R	24	48	4,987 (est.)	2,740 (est.)
Xeon Gold 5218R	20	40	4,203 (est.)	2,309 (est.)
Xeon Silver 4215R	8	16	2,043 (est.)	1,122 (est.)
Xeon Silver 4214R	12	24	2,622 (est.)	1,441 (est.)
Xeon Silver 4210R	10	20	2,225 (est.)	1,223 (est.)
Xeon Bronze 3206R	8	8	1,052 (est.)	578 (est.)

The following diagram shows the OLTP-2 transaction rates that can be achieved with one and two processors of the 2nd Generation Intel® Xeon® Processor Scalable Family.

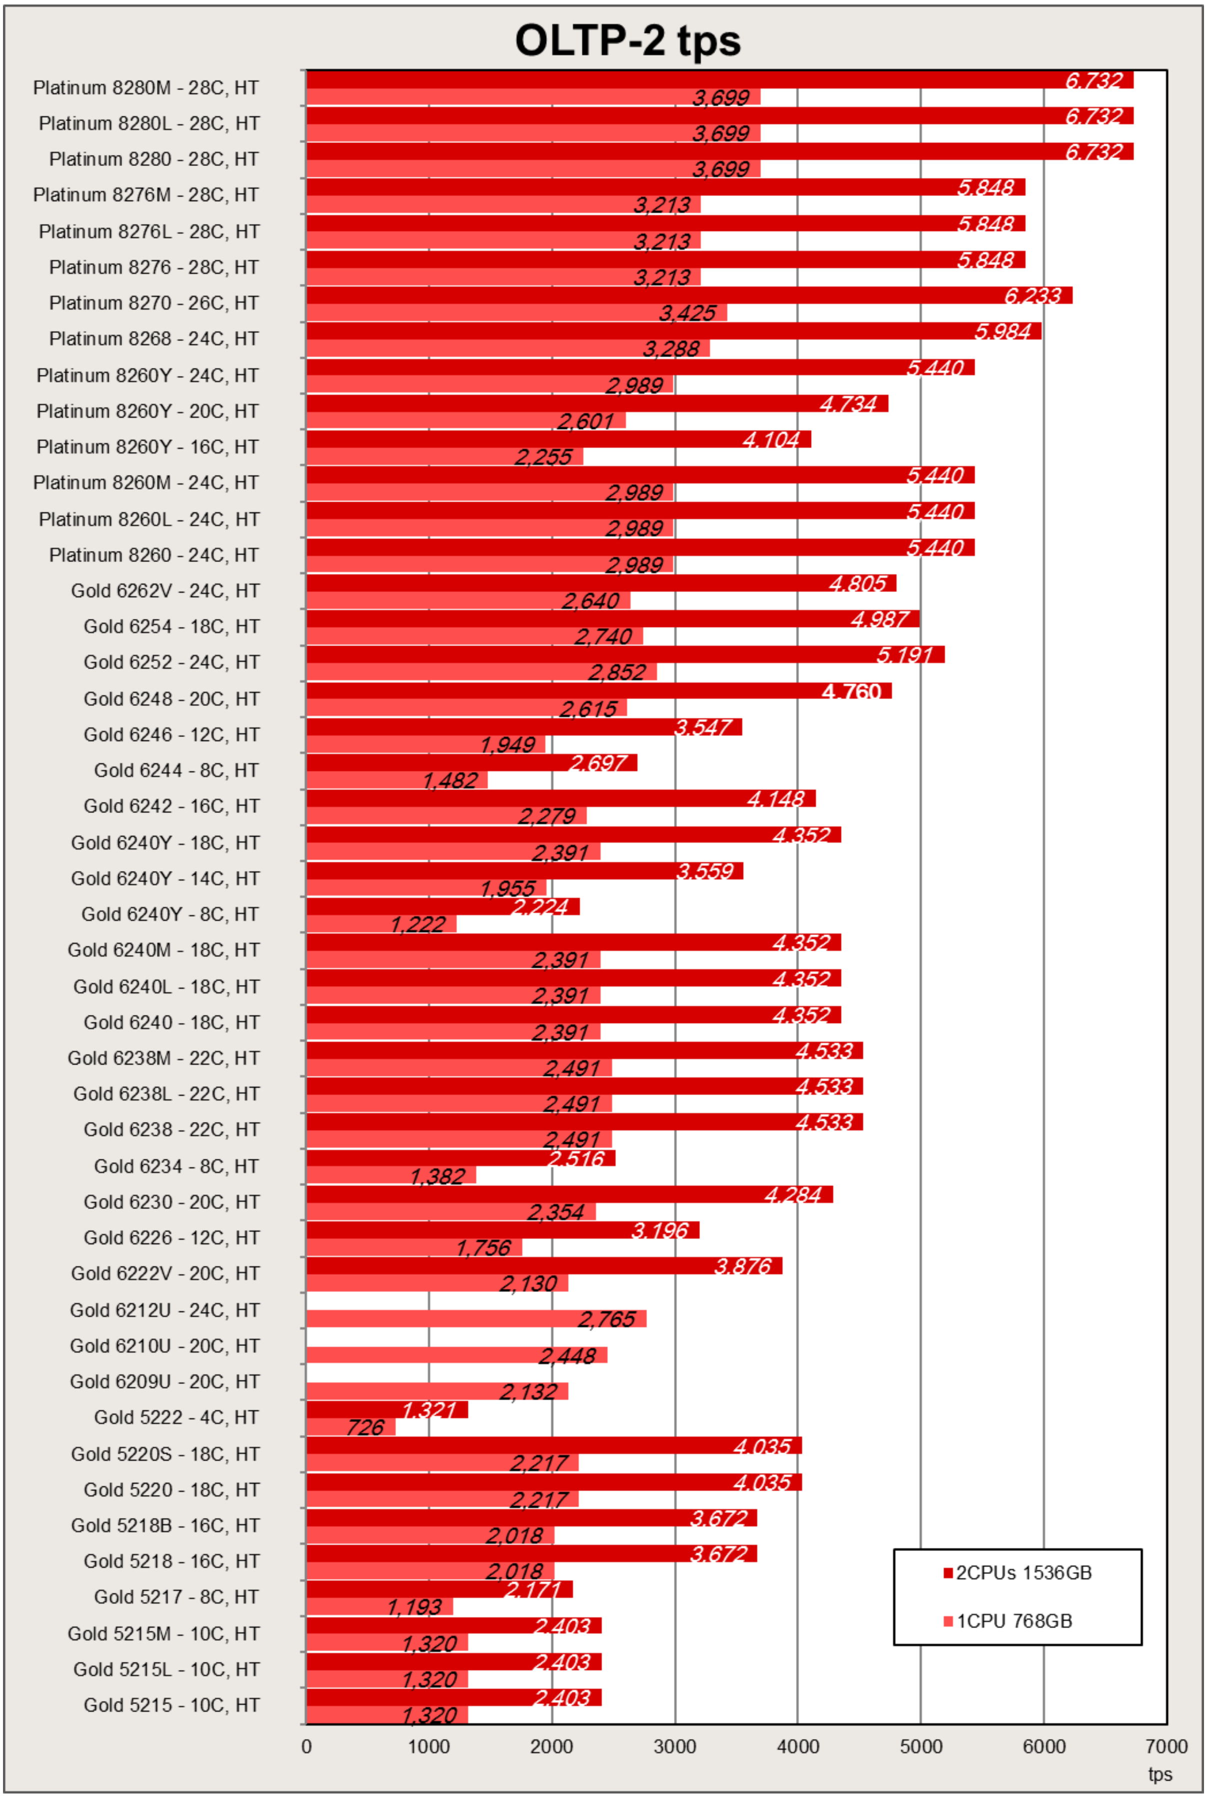

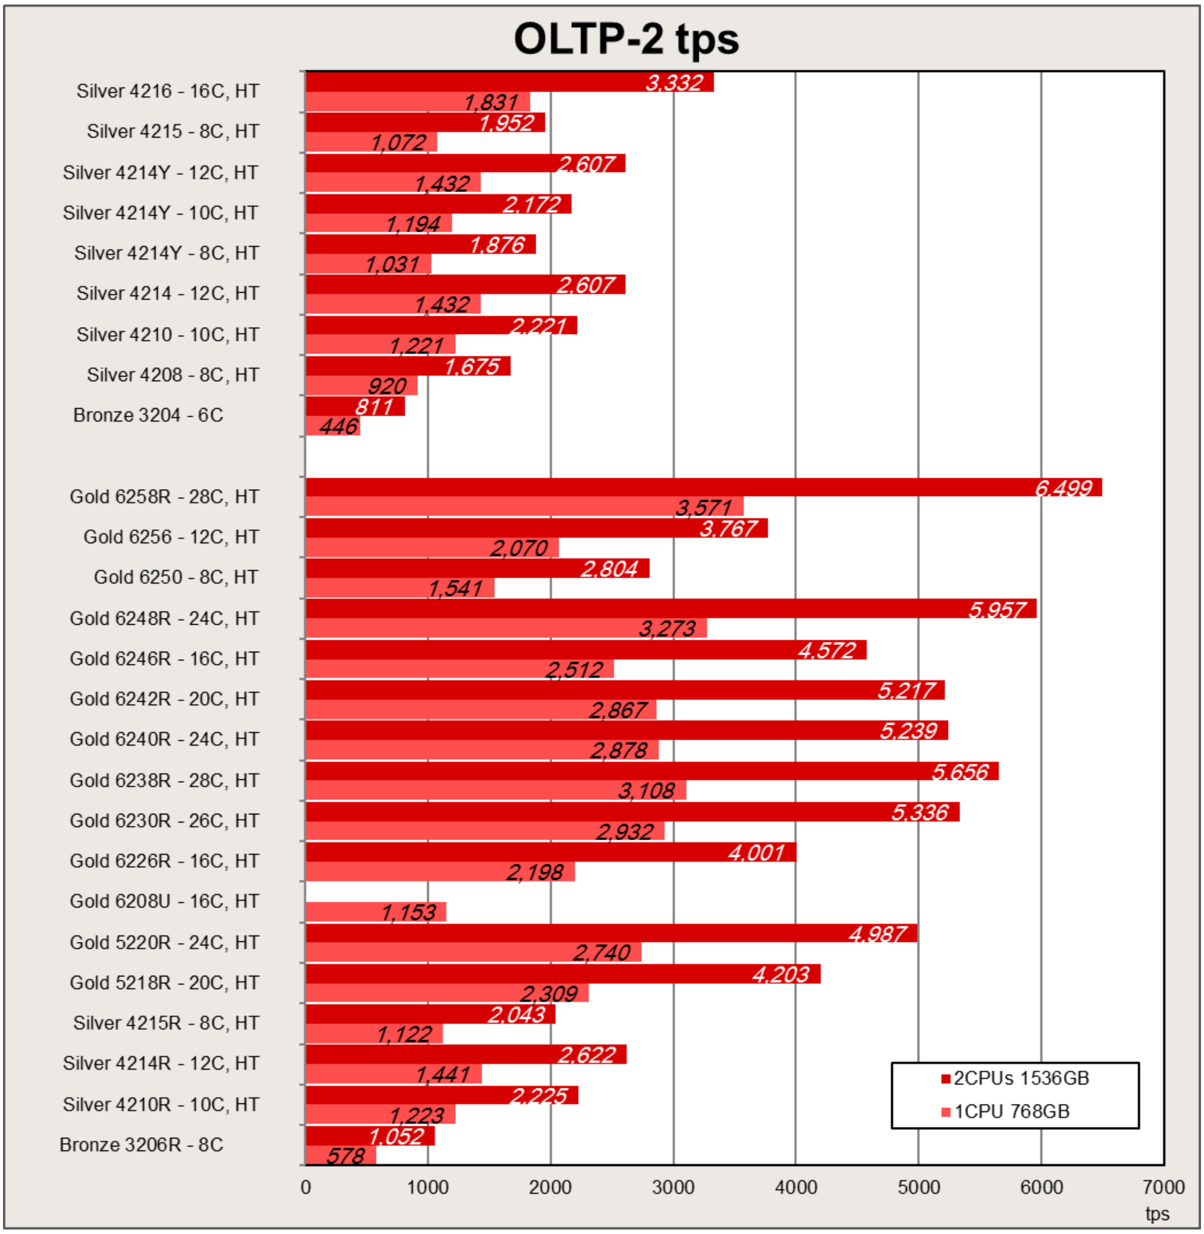

It is evident that a wide performance range is covered by the variety of released processors. If you compare the OLTP-2 value of the processor with the lowest performance (Xeon Bronze 3204) with the value of the processor with the highest performance (Xeon Platinum 8280), the result is an 8-fold increase 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.

If you compare the maximum achievable OLTP-2 values of the current system generation with the values that were achieved on the predecessor systems, the result is an increase of about 2%.

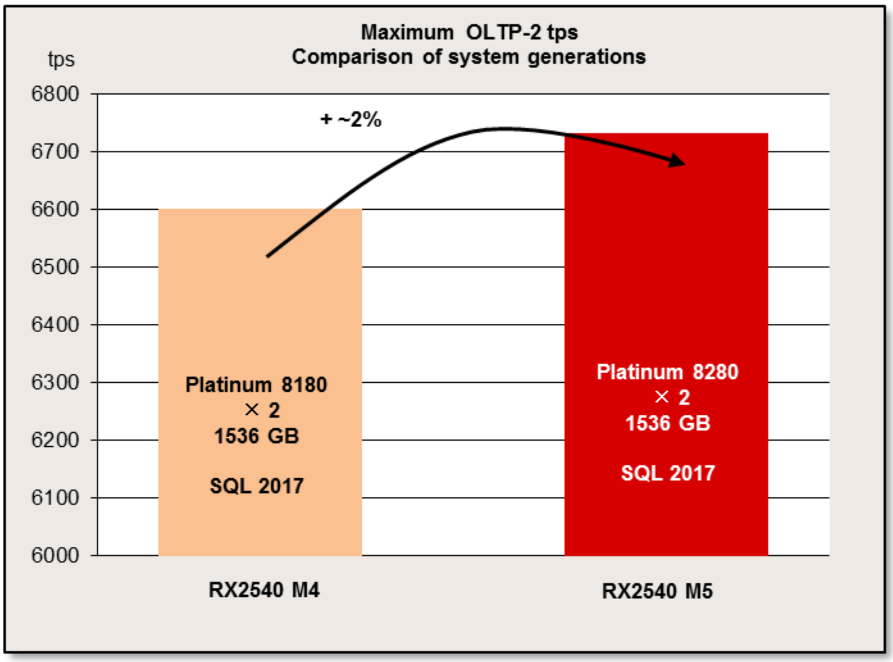

TPC-E

Benchmark description

The TPC-E benchmark measures the performance of online transaction processing systems (OLTP) and is based on a complex database and a number of different transaction types that are carried out on it. TPC-E is not only a hardware-independent but also a software-independent benchmark and can thus be run on every test platform, i.e. proprietary or open. In addition to the results of the measurement, all the details of the systems measured and the measuring method must also be explained in a measurement report (Full Disclosure Report or FDR). Consequently, this ensures that the measurement meets all benchmark requirements and is reproducible. TPC-E does not just measure an individual server, but a rather extensive system configuration. Keys to performance in this respect are the database server, disk I/O and network communication.

The performance metric is tpsE, where tps means transactions per second. tpsE is the average number of Trade-Result-Transactions that are performed within a second. The TPC-E standard defines a result as the tpsE rate, the price per performance value (e.g. $/tpsE) and the availability date of the measured configuration.

Further information about TPC-E can be found in the overview document Benchmark Overview TPC-E.

Benchmark results

In October 2019 Fujitsu submitted a TPC-E benchmark result for the PRIMERGY RX2540 M5 with the 28-core processor Intel Xeon Platinum 8280 and 1536 GB memory.

The results show an increase in performance compared with the PRIMERGY RX2540 M4 with a simultaneous reduction in price per performance ratio.

FUJITSU	**FUJITSU Server PRIMERGY RX2540 M5**	**TPC-E™ 1.14.0** **TPC Pricing 2.4.0** Report Date: Oct 24, 2019 Revision Date: Oct 24, 2019

TPC-E Throughput	Price/Performance	Availability Date	Total System Cost
6,844.20 tpsE	**$ 85.13 USD per tpsE**	**Oct 24, 2019**	**$ 582,623 USD**

Database Server Configuration

Operating System	Database Manager	Processors/Cores/Threads	Memory
Microsoft Windows Server 2016 Standard Edition	**Microsoft SQL Server 2017 Enterprise Edition**	**2/56/112**	**1,536 GB**

SUT

Tier A
PRIMERGY RX2530 M5
2x Intel Xeon Platinum 8280 2.70 GHz
192 GB Memory
2x 300 GB 10k rpm SAS Drive
1x onboard dual port LAN 10 Gb/s
1x onboard dual port LAN 1 Gb/s
1x SAS RAID controller

Tier B
PRIMERGY RX2540 M5
2x Intel Xeon Platinum 8280 2.70 GHz
1,536 GB Memory
2x 300 GB 15k rpm SAS Drives
6x 1.6 TB SAS SSD
1x onboard dual port LAN 10 Gb/s
1x onboard dual port LAN 1 Gb/s
6x SAS RAID Controller

Storage
1x PRIMECENTER Rack
5x ETERNUS JX40 S2
49x 1.6 TB SSD Drives

Initial Database Size	Redundancy Level 1	Storage
33,336 GB	**RAID-5 for data** **RAID-10 for tempDB and log**	**55 x 1.6 TB SSD**

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

More details about this TPC-E result, in particular the Full Disclosure Report, can be found via the TPC web page http://www.tpc.org/tpce/results/tpce_result_detail5.asp?id=119102301.

In October 2019, Fujitsu is represented with two results in the TPC-E list (without historical results).

System and Processors	Throughput	Price / Performance	Availability Date
PRIMERGY RX2540 M4 with 2 × Xeon Platinum 8180	6606.75 tpsE	$92.85 per tpsE	March 31, 2018
PRIMERGY RX2540 M5 with 2 × Xeon Platinum 8280	6844.20 tpsE	$85.13 per tpsE	October 24, 2019

See the TPC web site for more information and all the TPC-E results (including historical results) (http://www.tpc.org/tpce).

The following diagram for two-socket PRIMERGY systems with different processor types shows the good performance of the two-socket system PRIMERGY RX2540 M5 .

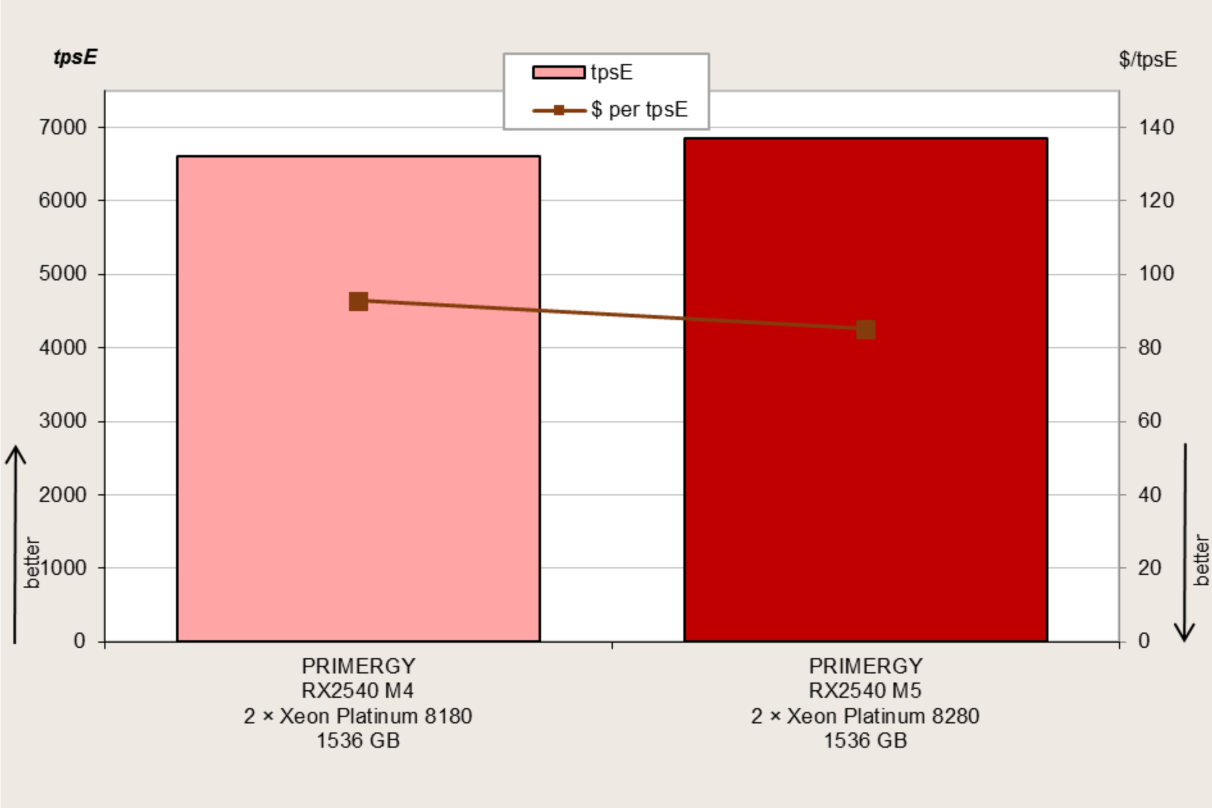

Performance of the PRIMERGY RX2540 M5 is 6844.20 tpsE and improves by 3.6% in comparison with the PRIMERGY RX2540 M4. The price per performance is $85.13 per tpsE and is reduced to 92% compared with the PRIMERGY RX2540 M4, making it a more cost-effective system.

The following overview, sorted according to price/performance, shows the best TPC-E price per performance ratios (as of January 31, 2020, without historical results) and the corresponding TPC-E throughputs. PRIMERGY RX2540 M5 with a price per performance ratio of $85.13 per tpsE achieved the best cost-effectiveness.

System		Processor Type /Number of processors	tpsE (higher is better)	$/tpsE (lower is better)	Availability date
Fujitsu	**PRIMERGY RX2540 M5**	2 x Intel Xeon Platinum 8280	6,844.20	**85.13**	2019-10-24
Lenovo	Thnk System SR650	2 x Intel Xeon Platinum 8280	7,012.53	90.99	2019-04-17
Lenovo	Thnk System SR650	2 x Intel Xeon Platinum 8180	6,779.53	92.49	2018-09-10
Fujitsu	PRIMERGY RX2540 M4	2 x Intel Xeon Platinum 8180	6,606.75	92.85	2018-03-31
Lenovo	Thnk System SR650	2 x Intel Xeon Platinum 8180	6,598.36	93.48	2017-10-19
Lenovo	Thnk System SR950	4 x Intel Xeon Platinum 8180	11,357.28	98.83	2017-11-01
Lenovo	Thnk System SR655	1 x AMD EPYC 7742	6,716.88	99.99	2019-12-31

See the TPC web site for more information and all the TPC-E results (including historical results) (http://www.tpc.org/tpce).

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.

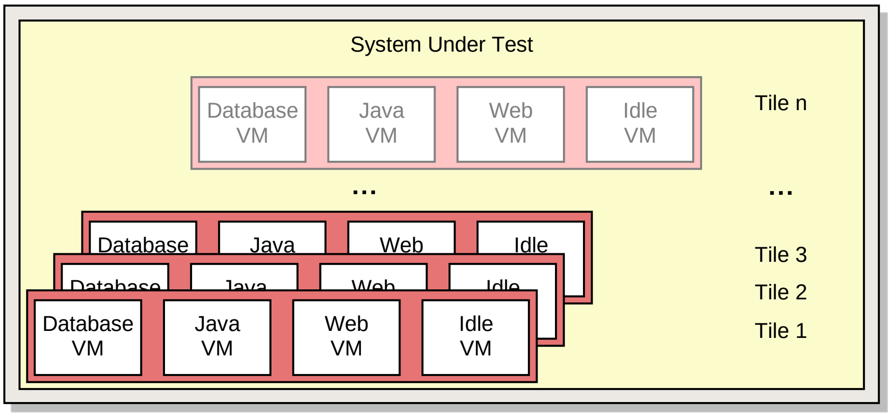

Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score, the individual benchmark result for one tile is put in relation to the respective result of a reference 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:

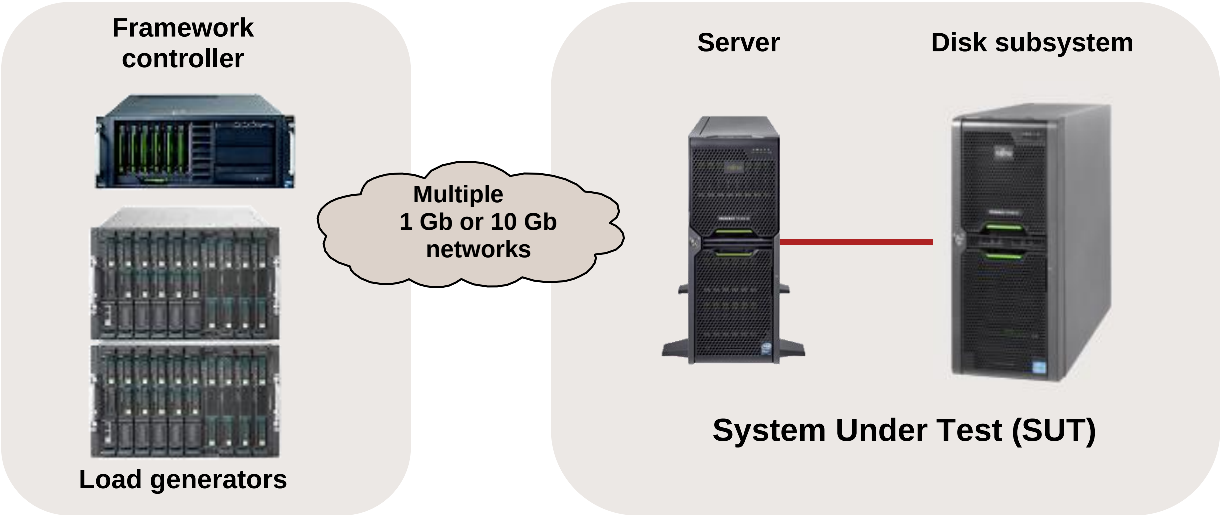

System Under Test (SUT)	
Hardware	
Processor	2nd Generation Intel® Xeon® Scalable Processors Family
Memory	24 × 32 GB (1x32 GB) 2Rx4 PC4-2933Y-R
Network interface	2 × Intel® Ethernet Controller X710 for 10GbE SFP+
Disk subsystem	1 × dual-channel FC controller Emulex LPe160021 LINUX/LIO based flash storage 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 U2 Build 3620759

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

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

Benchmark results

The PRIMERGY dual-socket rack and tower systems dealt with here are based on processors of the Intel® Xeon® Processor Scalable Family. 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 Platinum 8280L	28	56	2	35 (est.)	61.7 (est.)
Xeon Platinum 8280M	28	56	2	35 (est.)	61.7 (est.)
Xeon Platinum 8280	28	56	2	35	61.7
Xeon Platinum 8276L	28	56	2	34 (est.)	53.1 (est.)
Xeon Platinum 8276M	28	56	2	34 (est.)	53.1 (est.)
Xeon Platinum 8276	28	56	2	34 (est.)	53.1 (est.)
Xeon Platinum 8270	26	52	2	33	57.3
Xeon Platinum 8268	24	48	2	33	55.6
Xeon Platinum 8260L	24	48	2	31 (est.)	49.9 (est.)
Xeon Platinum 8260M	24	48	2	31 (est.)	49.9 (est.)
Xeon Platinum 8260Y	24	48	2	31 (est.)	49.9 (est.)
	20	40	2	28 (est.)	43.4 (est.)
	16	32	2	26 (est.)	37.6 (est.)
Xeon Platinum 8260	24	48	2	31 (est.)	49.9 (est.)
Xeon Gold 6262V	24	48	2	28 (est.)	44.0 (est.)
Xeon Gold 6254	18	36	2	26 (est.)	44.0 (est.)
Xeon Gold 6252	24	48	2	29 (est.)	49.6 (est.)
Xeon Gold 6248	20	40	2	28	43.7
Xeon Gold 6246	12	24	2	20	32.4
Xeon Gold 6244	8	16	2	15 (est.)	24.4 (est.)
Xeon Gold 6242	16	32	2	26 (est.)	38.0 (est.)
Xeon Gold 6240L	18	36	2	27 (est.)	39.9 (est.)
Xeon Gold 6240M	18	36	2	27 (est.)	39.9 (est.)
Xeon Gold 6240Y	18	36	2	27 (est.)	39.9 (est.)
	14	28	2	23 (est.)	29.5 (est.)
	8	16	2	13 (est.)	19.9 (est.)

Xeon Gold 6240	18	36	2	27 (est.)	39.9 (est.)
Xeon Gold 6238L	22	44	2	29 (est.)	45.8 (est.)
Xeon Gold 6238M	22	44	2	29 (est.)	45.8 (est.)
Xeon Gold 6238	22	44	2	29	45.8
Xeon Gold 6234	8	16	2	15 (est.)	23.1 (est.)
Xeon Gold 6230	20	40	2	26 (est.)	39.3 (est.)
Xeon Gold 6226	12	24	2	21 (est.)	29.3 (est.)
Xeon Gold 6222V	20	40	2	25 (est.)	35.5 (est.)
Xeon Gold 6212U	24	48	1	29 (est.)	25.1 (est.)
Xeon Gold 6210U	20	40	1	20 (est.)	22.2 (est.)
Xeon Gold 6209U	20	40	1	17 (est.)	19.5 (est.)
Xeon Gold 5222	4	8	2	8 (est.)	11.3 (est.)
Xeon Gold 5220S	18	36	2	26 (est.)	37.0 (est.)
Xeon Gold 5220	18	36	2	26 (est.)	37.0 (est.)
Xeon Gold 5218B	16	32	2	23 (est.)	33.6 (est.)
Xeon Gold 5218	16	32	2	23 (est.)	33.6 (est.)
Xeon Gold 5217	8	16	2	14 (est.)	19.9 (est.)
Xeon Gold 5215L	10	20	2	18 (est.)	20.9 (est.)
Xeon Gold 5215M	10	20	2	18 (est.)	20.9 (est.)
Xeon Gold 5215	10	20	2	18 (est.)	20.9 (est.)
Xeon Silver 4216	16	32	2	25 (est.)	30.5 (est.)
Xeon Silver 4215	8	16	2	14 (est.)	17.9 (est.)
Xeon Silver 4214Y	12	24	2	19 (est.)	24.3 (est.)
	10	20	2	16 (est.)	20.9 (est.)
	8	16	2	12 (est.)	18.3 (est.)
Xeon Silver 4214	12	24	2	20 (est.)	23.9 (est.)
Xeon Silver 4210	10	20	2	18 (est.)	20.4 (est.)
Xeon Silver 4208	8	16	2	13 (est.)	15.4 (est.)
Xeon Bronze 3204	6	6	2	10 (est.)	7.4 (est.)
March 2020 released					
Xeon Gold 6258R	28	56	2	34 (est.)	59.6 (est.)
Xeon Gold 6256	12	24	2	25 (est.)	34.5 (est.)
Xeon Gold 6250	8	16	2	16 (est.)	25.7 (est.)
Xeon Gold 6248R	24	48	2	32 (est.)	54.6 (est.)
Xeon Gold 6246R	16	32	2	29 (est.)	41.9 (est.)
Xeon Gold 6242R	20	40	2	31 (est.)	47.8 (est.)
Xeon Gold 6240R	24	48	2	30 (est.)	48.0 (est.)
Xeon Gold 6238R	28	56	2	33 (est.)	51.8 (est.)
Xeon Gold 6230R	26	52	2	28 (est.)	48.9 (est.)
Xeon Gold 6226R	16	32	2	25 (est.)	36.7 (est.)
Xeon Gold 6208U	16	32	1	18 (est.)	19.6 (est.)
Xeon Gold 5220R	24	48	2	27 (est.)	45.7 (est.)
Xeon Gold 5218R	20	40	2	25 (est.)	38.5 (est.)
Xeon Silver 4215R	8	16	2	15 (est.)	18.7 (est.)
Xeon Silver 4214R	12	24	2	20 (est.)	24.0 (est.)
Xeon Silver 4210R	10	20	2	18 (est.)	20.4 (est.)
Xeon Bronze 3206R	8	8	2	13 (est.)	9.6 (est.)

These PRIMERGY dual-socket rack and tower systems are very suitable for application virtualization owing to the progress made in processor technology. Compared with a system based on the previous processor generation, approximately 2.9% higher virtualization performance can be achieved (measured in vServCon score in their maximum configuration).

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

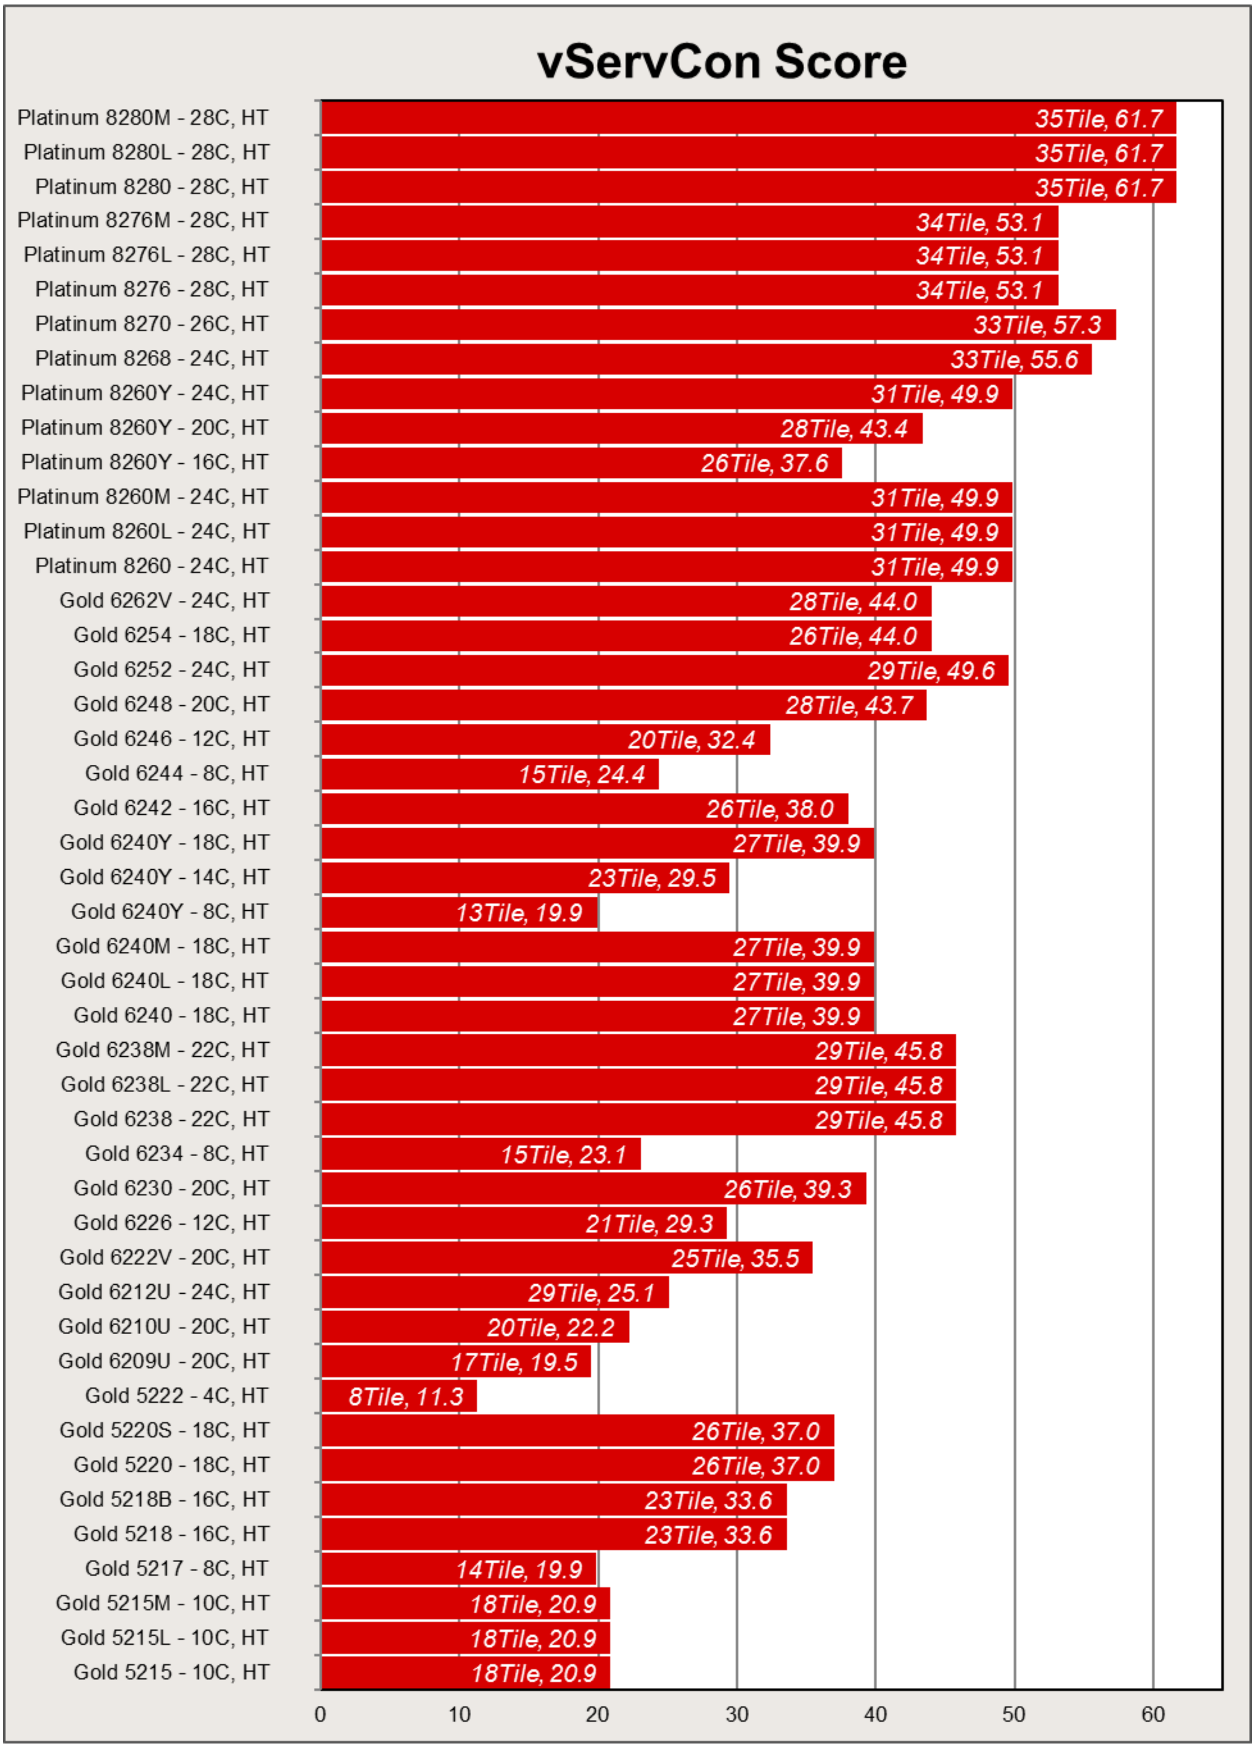

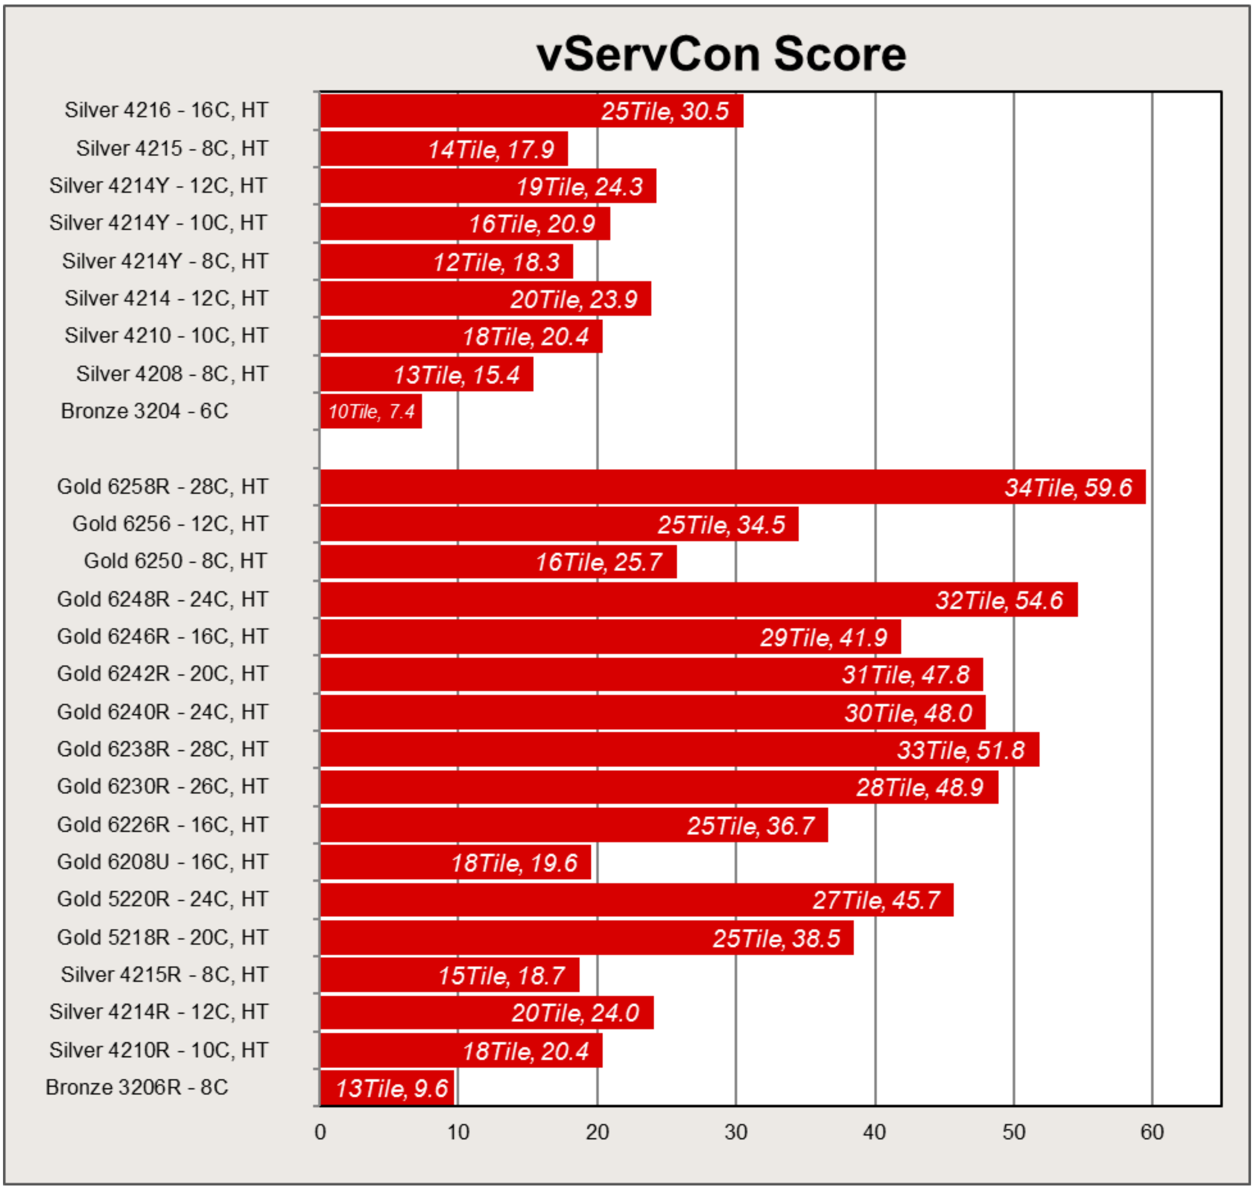
vServCon Score
Silver 4216 - 16C, HT
25Tile, 30.5
Silver 4215 - 8C, HT
14Tile, 17.9
Silver 4214Y - 12C, HT
19Tile, 24.3
Silver 4214Y - 10C, HT
16Tile, 20.9
Silver 4214Y - 8C, HT
12Tile, 18.3
Silver 4214 - 12C, HT
20Tile, 23.9
Silver 4210 - 10C, HT
18Tile, 20.4
Silver 4208 - 8C, HT
13Tile, 15.4
Bronze 3204 - 6C
10Tile, 7.4
Gold 6258R - 28C, HT
34Tile, 59.6
Gold 6256 - 12C, HT
25Tile, 34.5
Gold 6250 - 8C, HT
16Tile, 25.7
Gold 6248R - 24C, HT
32Tile, 54.6
Gold 6246R - 16C, HT
29Tile, 41.9
Gold 6242R - 20C, HT
31Tile, 47.8
Gold 6240R - 24C, HT
30Tile, 48.0
Gold 6238R - 28C, HT
33Tile, 51.8
Gold 6230R - 26C, HT
28Tile, 48.9
Gold 6226R - 16C, HT
25Tile, 36.7
Gold 6208U - 16C, HT
18Tile, 19.6
Gold 5220R - 24C, HT
27Tile, 45.7
Gold 5218R - 20C, HT
25Tile, 38.5
Silver 4215R - 8C, HT
15Tile, 18.7
Silver 4214R - 12C, HT
20Tile, 24.0
Silver 4210R - 10C, HT
18Tile, 20.4
Bronze 3206R - 8C
13Tile, 9.6
0
10
20
30
40
50
60

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.

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.

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

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.

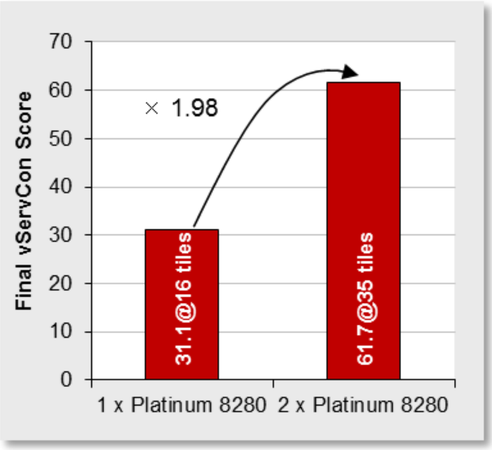

Until now, we have looked at the virtualization performance of a fully configured system. However, with a server with four sockets, the question also arises as to how good performance scaling is from one to two processors. The better the scaling, the lower the overhead usually caused by the shared use of resources within a server. The scaling factor also depends on the application. If the server is used as a virtualization platform for server consolidation, the system scales with a factor of 1.98. When operated with two processors, the system thus achieves twice the performance as with one processor, as is illustrated in this diagram using the processor version Xeon Platinum 8280 as an example.

The next diagram illustrates the virtualization performance for increasing numbers of VMs based on the Xeon Gold 6244 (20-Core) processors.

In addition to the increased number of physical cores, Hyper-Threading, which is supported by almost all processors of the 2nd Generation Intel® Xeon® Processor Scalable Product Family, is an additional reason for the high number of VMs that can be operated. As is known, a physical processor core is consequently divided into two logical cores so that the number of cores available for the hypervisor is doubled. This standard feature thus generally increases the virtualization performance of a system.

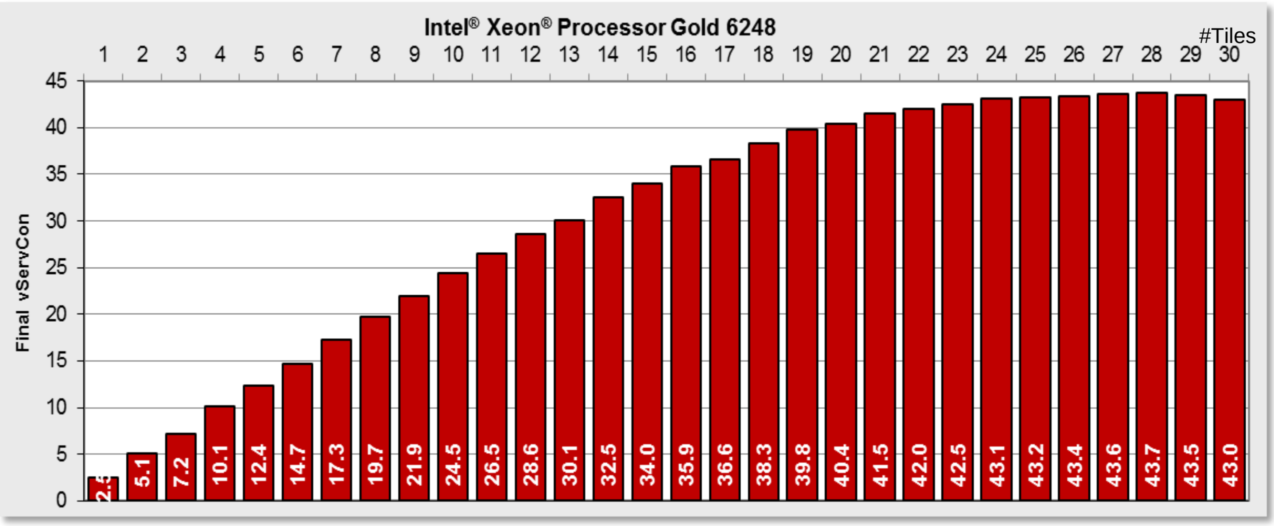

The previous diagram examined the total performance of all application VMs of a host. However, studying the performance from an individual application VM viewpoint is also interesting. This information is in the previous diagram. For example, the total optimum is reached in the above Xeon Gold 6248 situation with 84 application VMs (28 tiles, not including the idle VMs) The low load case is represented by three application VMs (one tile, not including the idle VM). Remember, the vServCon score for one tile is an average value across the three application scenarios in vServCon. This average performance of one tile drops when changing from the low load case to the total optimum of the vServCon score – from 2.5 to 43.7/28=1.5, i.e. to 67%. The individual types of application VMs can react very differently in the high load situation. It is thus clear that in a specific situation the performance requirements of an individual application must be balanced against the overall requirements regarding the numbers of VMs on a virtualization host.

The performance for an individual VM in low-load situations has only slightly increased for the processors compared here with the highest clock frequency per core. We must explicitly point out that the increased virtualization performance as seen in the score cannot be completely deemed as an improvement for one individual VM.

Performance increases in the virtualization environment since 2010 are mainly achieved by increases in the maximum number of VMs that can be operated.

	Best Maximum Performance CPU	vServCon Score max.
2008	X5460	2.94@2 Tile
2009	X5570	6.08@ 6 Tile
2011	X5690	9.61@ 9 Tile
2012	E5-2690	13.5@ 8 Tile
2013	E5-2697 v2	17.1@11 Tile
2014	E5-2699 v3	30.3@18 Tile
2016	E5-2699 v4	38.7@22 Tile
2017	Platinum 8180	59.4@34 Tile
2019	Platinum 8280	61.7@35 Tile

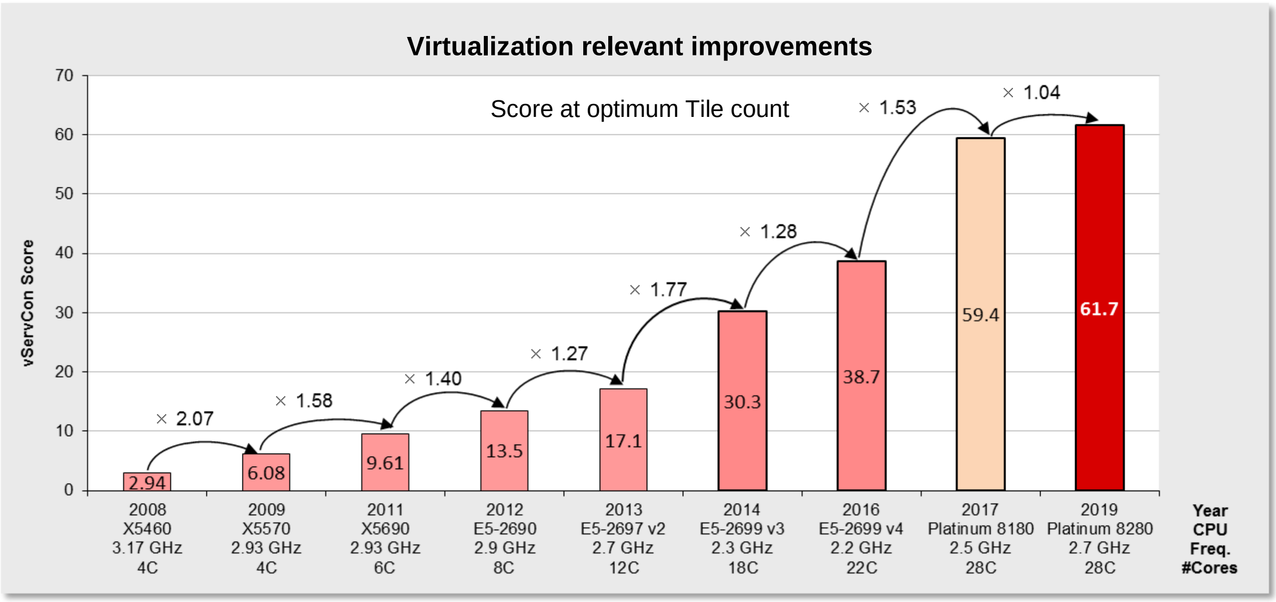

VMmark V3

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 and 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.

A new feature of VMmark V3 is 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 reflects the maximum total consolidation benefit of all VMs for a server configuration with hypervisor and is used as a comparison criterion of various hardware platforms.

This score is determined from the individual results of the VMs and an infrastructure 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 workload components.

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, which 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:

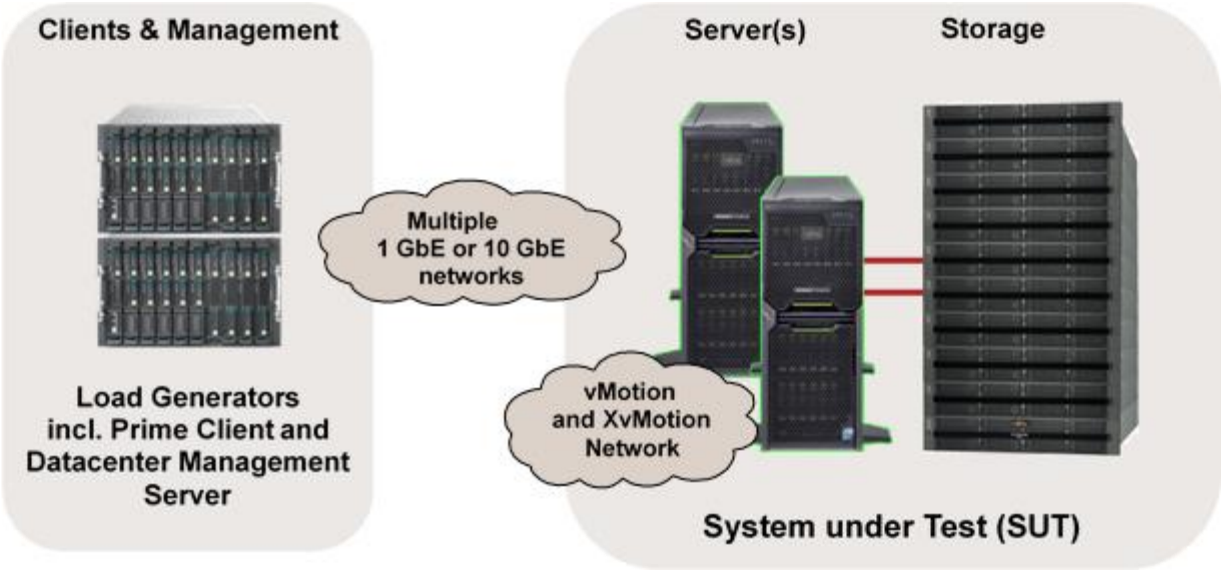

System Under Test (SUT)	
Hardware	
Number of servers	2
Model	PRIMERGY RX2540 M5
Processor	2 × Intel® Xeon® Platinum 8280
Memory	768 GB: 24 × 32 GB (1x32 GB) 2Rx4 DDR4-2933 R ECC
Network interface	2 × Intel ® Ethernet Controller X710 for 10GbE SFP+ 1 × Intel I350 Dual Port 1 GbE Adapter
Disk subsystem	2 × Dual port PFC EP LPe31002 4 × PRIMERGY RX2540 M4 configured as Fibre Channel target: 8 × Micron MTFDDAK480 TDC SATA-SSD (480 GB) 8 × Intel P4800X 750GB PCIe SSD (750 GB) 4 × Intel P4600 2TB PCIe SSD (2TB) 4 × Intel P4600 4TB PCIe SSD (4TB) RAID 0 with several LUNs
Software	
BIOS	R1.2.0
BIOS settings	See details
Operating system	VMware ESXi 6.7 EP 06, Build 11675023
Operating system settings	ESX settings: see details

Details	
See disclosure	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2019-04-02-Fujitsu-RX2540M5.pdf https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2019-04-02-Fujitsu-RX2540M5-serverPPKW.pdf https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2019-04-02-Fujitsu-RX2540M5-serverstoragePPKW.pdf

Datacenter Management Server (DMS)	
Hardware	
Model	1 × PRIMERGY RX2530 M2
Processor	2 × Intel Xeon E5-2698 v4
Memory	64 GB
Network interface	1 × Emulex One Connect Oce14000 1 GbE Dual Port Adapter
Software	
Operating system	VMware ESXi 6.7 EP 02a Build 9214924
Datacenter Management Server (DMS) VM	
Hardware	
Processor	4 × logical CPU
Memory	16 GB
Network interface	1 × 1 Gbit/s LAN
Software	
Operating system	VMware vCenter Server Appliance 6.7.0d Build 9451876

Load generator	
Hardware	
Model	3 × PRIMERGY RX2530 M2
Processor	2 × Xeon E5-2699 v4
Memory	258 GB
Network interface	1 × Emulex One Connect Oce14000 1GbE Dual Port Adapter 1 × Emulex One Connect Oce14000 10GbE Dual Port Adapter
Software	
Operating system	VMware ESXi 6.7 U1 Build 10302608

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

Benchmark results

"Performance Only" measurement result (April 2 2019)

On April 2, 2019 Fujitsu achieved with a PRIMERGY RX2540 M5 with Xeon Platinum 8280 processors and VMware VMware ESXi 6.7 EP 06 a VMmark V3 score of "9.02@9 tiles" in a system configuration with a total of 2 × 56 processor cores and when using two identical servers in the "System under Test" (SUT). With this result the PRIMERGY RX2540 M5 is in the official VMmark V3 "Performance Only" ranking the most powerful two-socket server in a "matched pair" configuration consisting of two identical hosts (valid as of benchmark results publication date).

All comparisons for the competitor products reflect the status of April 2, 2019. The current VMmark V3 "Performance Only" results as well as the detailed results and configuration data are available at https://www.vmware.com/products/vmmark/results3x.html .

The processors used, which with a good hypervisor setting could make optimal use of their processor features, were the essential prerequisites for achieving the PRIMERGY RX2540 M5 result. These features include Hyper-Threading. All this has a particularly positive effect during virtualization.

All VMs, their application data, the host operating system as well as additionally required data were on a powerful Fibre Channel disk subsystem. As far as possible, the configuration of the disk subsystem takes the specific requirements of the benchmark into account. The use of flash technology in the form of SAS SSDs and PCIe-SSDs in the powerful Fibre Channel disk subsystem resulted in further advantages in response times of the storage medium used.

The network connection to the load generators and the infrastructure-workload connection between the hosts were implemented via 10GbE LAN ports.

All the components used were optimally attuned to each other.

"Performance with Server Power" measurement result (April 2 2019)

On April 2, 2019 Fujitsu achieved with a PRIMERGY RX2540 M5 with Xeon Platinum 8280 processors and VMware ESXi 6.7 EP 06 a VMmark V3 "Server PPKW Score" of "6.3290@9 tiles" in a system configuration with a total of 2 × 56 processor cores and when using two identical servers in the "System under Test" (SUT). With this result the PRIMERGY RX2540 M5 is in the official VMmark V3 "Performance with Server Power" ranking the most energy-efficient virtualization server worldwide (valid as of benchmark results publication date).

The current VMmark V3 "Performance with Server Power" results as well as the detailed results and configuration data are available at https://www.vmware.com/products/vmmark/results3x.html .

"Performance with Server and Storage Power" measurement result (April 2 2019)

On April 2, 2019 Fujitsu achieved with a PRIMERGY RX2540 M5 with Xeon Platinum 8280 processors and VMware ESXi 6.7 EP 06 a VMmark V3 "Server and Storage PPKW Score" of "3.5013 @9 tiles" in a system configuration with a total of 2 × 56 processor cores and when using two identical servers in the "System under Test" (SUT).

The current VMmark V3 "Performance with Server and Storage Power" results as well as the detailed results and configuration data are available at https://www.vmware.com/products/vmmark/results3x.html .

VMmark® is a product of VMware, Inc.

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, consequently achieving optimal load distribution to the available processor cores.

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

Type	Execution	Bytes per step	Floating-point calculation per step
COPY	a(i) = b(i)	16	0
SCALE	a(i) = q × b(i)	16	1
SUM	a(i) = b(i) + c(i)	24	1
TRIAD	a(i) = b(i) + q × c(i)	24	2

The throughput is output in GB/s for each type of calculation. The differences between the various values are usually only minor on modern systems. In general, only the determined TRIAD value is used as a comparison.

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

This chapter specifies throughputs on a basis of 10 (1 GB/s = 10^9 Byte/s).

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M5
Processor	2nd Generation Intel® Xeon® Scalable Processors Family
Memory	24 × 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 Sub NUMA Clustering = Disabled*1 WR CRC feature Control = Disabled *1: 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
Operating system	SUSE Linux Enterprise Server 15
Operating system settings	Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1) echo never > /sys/kernel/mm/transparent_hugepage/enabled run with avx512 or avx2*1 *1: Xeon Gold 5220R, Xeon Gold 5218R, Xeon Silver 4215R, Xeon Silver 4214R, Xeon Silver 4210R, Xeon Bronze 3206R
Compiler	CPU released in April 2019 C/C++: Version 2019.3.0.591499 of Intel C/C++ Compiler for Linux CPU released in March 2020 C/C++: Version 19.0.4.227 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 result with "est." are the estimated values.

Processor	Memory Frequency [MHz]	Max. Memory Bandwidth *1 [GB/s]	Cores	Processor Frequency [GHz]	Number of Processors	TRIAD [GB/s]
April 2019 released						
Xeon Platinum 8280L	2933	140.8	28	2.7	2	242(est.)
Xeon Platinum 8280M	2933	140.8	28	2.7	2	242(est.)
Xeon Platinum 8280	2933	140.8	28	2.7	2	**242**
Xeon Platinum 8276L	2933	140.8	28	2.2	2	**242**
Xeon Platinum 8276M	2933	140.8	28	2.2	2	242(est.)
Xeon Platinum 8276	2933	140.8	28	2.2	2	242(est.)
Xeon Platinum 8270	2933	140.8	26	2.7	2	**241**
Xeon Platinum 8268	2933	140.8	24	2.9	2	**243**
Xeon Platinum 8260L	2933	140.8	24	2.4	2	243(est.)
Xeon Platinum 8260M	2933	140.8	24	2.4	2	243(est.)
Xeon Platinum 8260Y	2933	140.8	24	2.4	2	**243**
	2933	140.8	20	2.4	2	**245**
	2933	140.8	16	2.4	2	**245**
Xeon Platinum 8260	2933	140.8	24	2.4	2	243(est.)
Xeon Gold 6262V	2933	140.8	24	1.9	2	**201**
Xeon Gold 6254	2933	140.8	18	3.1	2	**227**
Xeon Gold 6252	2933	140.8	24	2.1	2	**242**
Xeon Gold 6248	2933	140.8	20	2.5	2	**235**
Xeon Gold 6246	2933	140.8	12	3.3	2	**228**
Xeon Gold 6244	2933	140.8	8	3.6	2	**203**
Xeon Gold 6242	2933	140.8	16	2.8	2	**223**
Xeon Gold 6240L	2933	140.8	18	2.6	2	228(est.)
Xeon Gold 6240M	2933	140.8	18	2.6	2	228(est.)
Xeon Gold 6240Y	2933	140.8	18	2.6	2	**227**
	2933	140.8	14	2.6	2	**229**
	2933	140.8	8	2.6	2	**192**
Xeon Gold 6240	2933	140.8	18	2.6	2	**228**
Xeon Gold 6238M	2933	140.8	22	2.1	2	237(est.)
Xeon Gold 6238L	2933	140.8	22	2.1	2	237(est.)
Xeon Gold 6238	2933	140.8	22	2.1	2	**237**
Xeon Gold 6234	2933	140.8	8	3.3	2	**161**
Xeon Gold 6230	2933	140.8	20	2.1	2	**235**
Xeon Gold 6226	2933	140.8	12	2.7	2	**213**
Xeon Gold 6222V	2400	140.8	20	1.8	2	**198**
Xeon Gold 6212U	2933	140.8	24	2.4	1	**124**
Xeon Gold 6210U	2933	140.8	20	2.5	1	119(est.)
Xeon Gold 6209U	2933	140.8	20	2.1	1	**122**

Xeon Gold 5222	2933	140.8	4	3.8	2	**103**
Xeon Gold 5220S	2666	128.0	18	2.7	2	**209**
Xeon Gold 5220	2666	128.0	18	2.2	2	**212**
Xeon Gold 5218B	2666	128.0	16	2.3	2	209(est.)
Xeon Gold 5218	2666	128.0	16	2.3	2	**209**
Xeon Gold 5217	2666	128.0	8	3.0	2	**139**
Xeon Gold 5215L	2666	128.0	10	2.5	2	160(est.)
Xeon Gold 5215M	2666	128.0	10	2.5	2	160(est.)
Xeon Gold 5215	2666	128.0	10	2.5	2	**160**
Xeon Silver 4216	2400	115.2	16	2.1	2	**194**
Xeon Silver 4215	2400	115.2	8	2.5	2	**97.5**
Xeon Silver 4214Y	2400	115.2	12	2.2	2	**174**
	2400	115.2	10	2.2	2	**175**
	2400	115.2	8	2.2	2	**165**
Xeon Silver 4214	2400	115.2	12	2.2	2	**174**
Xeon Silver 4210	2400	115.2	10	2.2	2	**98.5**
Xeon Silver 4208	2400	115.2	8	2.1	2	**95.3**
Xeon Bronze 3204	2133	102.4	6	1.9	2	**76.9**
March 2020 released						
Xeon Gold 6258R	2933	140.8	28	2.7	2	**243**
Xeon Gold 6256	2933	140.8	12	3.6	2	**232**
Xeon Gold 6250	2933	140.8	8	3.9	2	**185**
Xeon Gold 6248R	2933	140.8	24	3.0	2	**243**
Xeon Gold 6246R	2933	140.8	16	3.4	2	**247**
Xeon Gold 6242R	2933	140.8	20	3.1	2	**247**
Xeon Gold 6240R	2933	140.8	24	2.4	2	**244**
Xeon Gold 6238R	2933	140.8	28	2.2	2	**242**
Xeon Gold 6230R	2933	140.8	26	2.1	2	**241**
Xeon Gold 6226R	2933	140.8	16	2.9	2	**222**
Xeon Gold 6208U	2933	140.8	16	2.9	1	**121**
Xeon Gold 5220R	2666	128.0	24	2.2	2	**223**
Xeon Gold 5218R	2666	128.0	20	2.1	2	**217**
Xeon Silver 4215R	2400	115.2	8	3.2	2	**119**
Xeon Silver 4214R	2400	115.2	12	2.4	2	**167**
Xeon Silver 4210R	2400	115.2	10	2.4	2	**95.1**
Xeon Bronze 3206R	2133	102.4	8	1.9	2	**82.8**

*1: Value per Processor

The following diagram illustrates the throughput of the PRIMERGY RX2540 M5 in comparison to its predecessor, the PRIMERGY RX2540 M4.

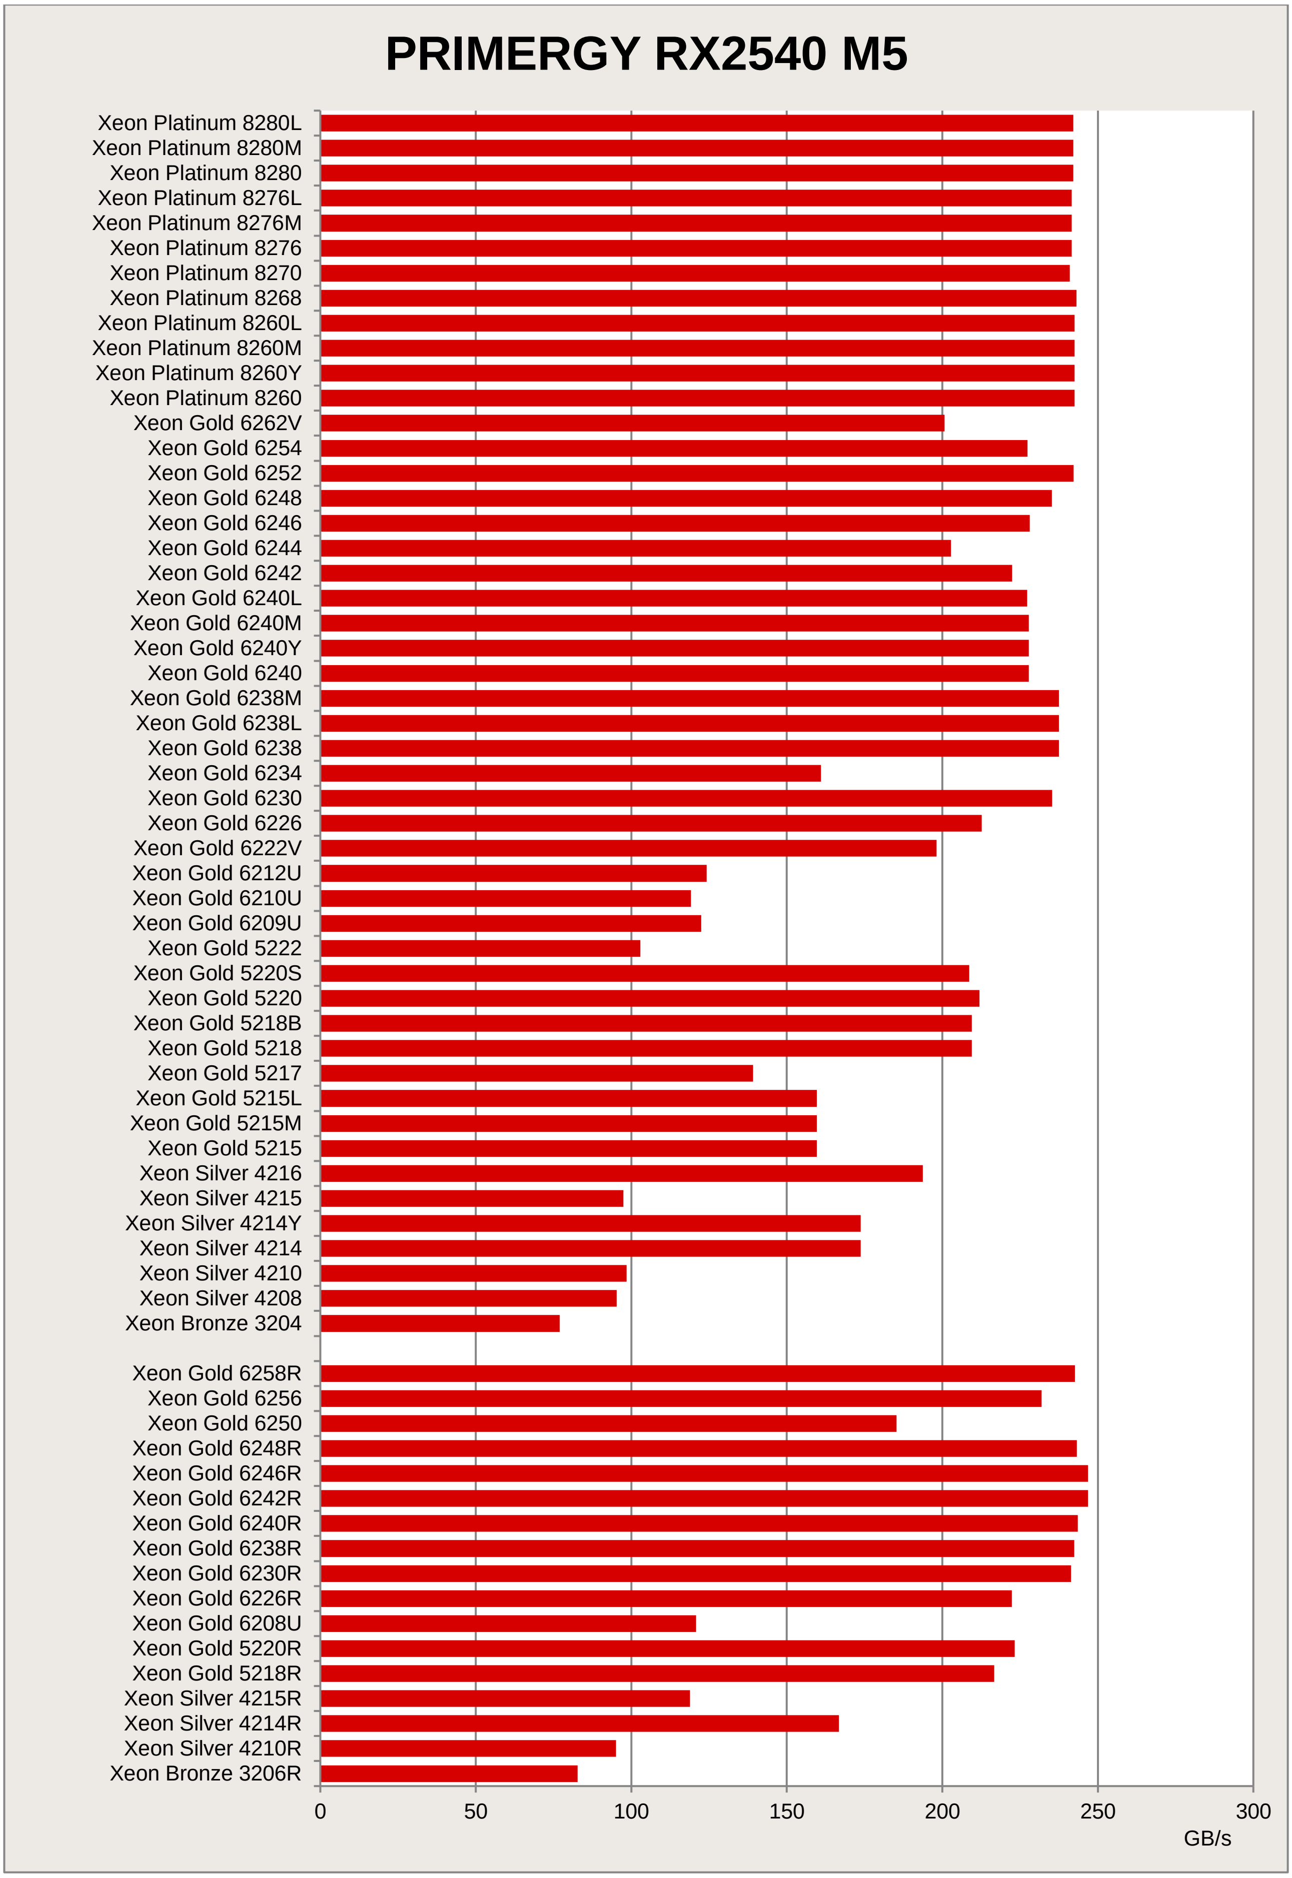
PRIMERGY RX2540 M5
Xeon Platinum 8280L
Xeon Platinum 8280M
Xeon Platinum 8280
Xeon Platinum 8276L
Xeon Platinum 8276M
Xeon Platinum 8276
Xeon Platinum 8270
Xeon Platinum 8268
Xeon Platinum 8260L
Xeon Platinum 8260M
Xeon Platinum 8260Y
Xeon Platinum 8260
Xeon Gold 6262V
Xeon Gold 6254
Xeon Gold 6252
Xeon 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
GB/s

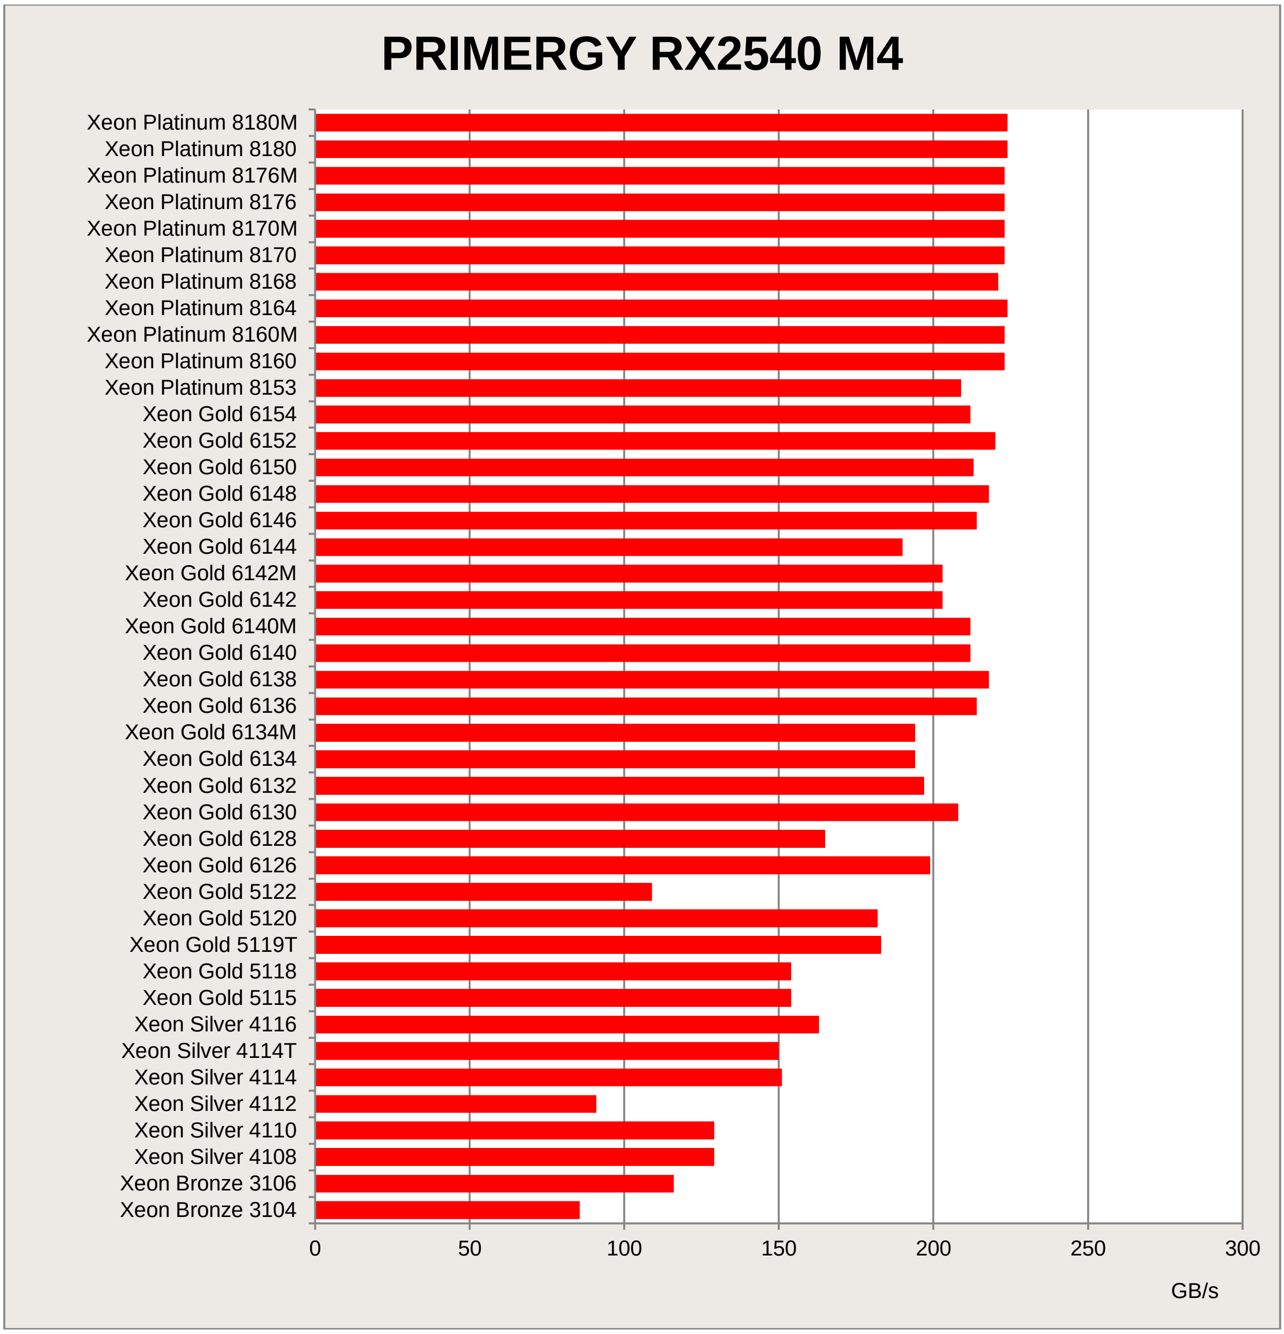
PRIMERGY RX2540 M4
Xeon Platinum 8180M
Xeon Platinum 8180
Xeon Platinum 8176M
Xeon Platinum 8176
Xeon Platinum 8170M
Xeon Platinum 8170
Xeon Platinum 8168
Xeon Platinum 8164
Xeon Platinum 8160M
Xeon Platinum 8160
Xeon Platinum 8153
Xeon Gold 6154
Xeon Gold 6152
Xeon Gold 6150
Xeon Gold 6148
Xeon Gold 6146
Xeon Gold 6144
Xeon Gold 6142M
Xeon Gold 6142
Xeon Gold 6140M
Xeon Gold 6140
Xeon Gold 6138
Xeon Gold 6136
Xeon Gold 6134M
Xeon Gold 6134
Xeon Gold 6132
Xeon Gold 6130
Xeon Gold 6128
Xeon Gold 6126
Xeon Gold 5122
Xeon Gold 5120
Xeon Gold 5119T
Xeon Gold 5118
Xeon Gold 5115
Xeon Silver 4116
Xeon Silver 4114T
Xeon Silver 4114
Xeon Silver 4112
Xeon Silver 4110
Xeon Silver 4108
Xeon Bronze 3106
Xeon Bronze 3104
0
50
100
150
200
250
300
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. A description can be found in the 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 × 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 × 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: a doubling of n results in an approximately eight-fold increase in the duration of the measurement. The size of n also has an influence on the measurement result itself. As n increases, the measured value asymptotically approaches a 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. The 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).

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
× Number of processor cores of the computer
× 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/. The use of a LINPACK version based on HPL is prerequisite for this (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 M5
Processor	2nd Generation Intel® Xeon® Scalable Processors Family
Memory	24 × 32GB (1x32GB) 2Rx4 PC4-2933Y-R
Software	
BIOS settings	HyperThreading = Disabled HWPM Support = Disabled Link Frequency Select = 10.4 GT/s Intel Virtualization Technology = Disabled Sub NUMA Clustering = Disabled LLC Dead Line Alloc = Disabled Stale AtoS = Enabled WR CRC feature Control = Disabled Fan Control = Full
Operating system	SUSE Linux Enterprise Server 15
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 or avx2*1 *1: Xeon Gold 5220R, Xeon Gold 5218R, Xeon Silver 4215R, Xeon Silver 4214R, Xeon Silver 4210R, Xeon Bronze 3206R
Compiler	CPU released in April 2019 C/C++: Version 2019.3.0.591499 of Intel C/C++ Compiler for Linux CPU released in March 2020 C/C++: Version 19.0.4.227 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 result with "est." are the estimated values.

Processor	Cores	Processor Frequency [GHz]	Number of Processors	Rpeak [GFlops]	Rmax [GFlops]	Efficiency
April 2019 released						
Xeon Platinum 8280L	28	2.7	2	4,838	**3,522**	73%
Xeon Platinum 8280M	28	2.7	2	4,838	3,522(est.)	73%
Xeon Platinum 8280	28	2.7	2	4,838	3,522(est.)	73%
Xeon Platinum 8276L	28	2.2	2	3,942	**2,768**	70%
Xeon Platinum 8276M	28	2.2	2	3,942	2,768(est.)	70%
Xeon Platinum 8276	28	2.2	2	3,942	2,768(est.)	70%
Xeon Platinum 8270	26	2.7	2	4,493	**3,200**	71%
Xeon Platinum 8268	24	2.9	2	4,454	**3,096**	70%
Xeon Platinum 8260L	24	2.4	2	3,686	2,735(est.)	74%
Xeon Platinum 8260M	24	2.4	2	3,686	2,735(est.)	74%
Xeon Platinum 8260Y	24	2.4	2	3,686	**2,735**	74%
	20	2.4	2	3,072	**2,423**	79%
	16	2.4	2	2,458	**2,149**	87%
Xeon Platinum 8260	24	2.4	2	3,686	2,735(est.)	74%
Xeon Gold 6262V	24	1.9	2	2,918	**2,061**	71%
Xeon Gold 6254	18	3.1	2	3,571	**2,705**	76%
Xeon Gold 6252	24	2.1	2	3,226	**2,674**	83%
Xeon Gold 6248	20	2.5	2	3,200	**2,375**	74%
Xeon Gold 6246	12	3.3	2	2,534	**1,915**	76%
Xeon Gold 6244	8	3.6	2	1,843	**1,460**	79%
Xeon Gold 6242	16	2.8	2	2,867	**2,253**	79%
Xeon Gold 6240L	18	2.6	2	2,995	2,169(est.)	72%
Xeon Gold 6240M	18	2.6	2	2,995	2,169(est.)	72%
Xeon Gold 6240Y	18	2.6	2	2,995	**2,210**	74%
	14	2.6	2	2,330	**1,894**	81%
	8	2.6	2	1,331	**1,401**	105%
Xeon Gold 6240	18	2.6	2	2,995	**2,169**	72%
Xeon Gold 6238M	22	2.1	2	2,957	2,334(est.)	79%
Xeon Gold 6238L	22	2.1	2	2,957	2,334(est.)	79%
Xeon Gold 6238	22	2.1	2	2,957	**2,334**	79%
Xeon Gold 6234	8	3.3	2	1,690	1,325	78%
Xeon Gold 6230	20	2.1	2	2,688	**1,976**	74%
Xeon Gold 6226	12	2.7	2	2,074	**1,732**	84%
Xeon Gold 6222V	20	1.8	2	2,304	**1,885**	82%
Xeon Gold 6212U	24	2.4	1	1,843	**1,387**	76%
Xeon Gold 6210U	20	2.5	1	1,600	tbd.	
Xeon Gold 6209U	20	2.1	1	1,344	tbd.	
Xeon Gold 5222	4	3.8	2	973	**775**	80%
Xeon Gold 5220S	18	2.7	2	1,555	**1,259**	81%

Xeon Gold 5220	18	2.2	2	1,267	**1,234**	97%
Xeon Gold 5218B	16	2.3	2	1,178	1,113(est.)	94%
Xeon Gold 5218	16	2.3	2	1,178	**1,113**	94%
Xeon Gold 5217	8	3	2	768	**714**	93%
Xeon Gold 5215L	10	2.5	2	800	766(est.)	96%
Xeon Gold 5215M	10	2.5	2	800	766(est.)	96%
Xeon Gold 5215	10	2.5	2	800	**766**	96%
Xeon Silver 4216	16	2.1	2	1,075	**1,066**	99%
Xeon Silver 4215	8	2.5	2	640	**626**	98%
Xeon Silver 4214Y	12	2.2	2	845	**798**	95%
	10	2.2	2	704	**724**	103%
	8	2.2	2	563	**654**	116%
Xeon Silver 4214	12	2.2	2	845	**805**	95%
Xeon Silver 4210	10	2.2	2	704	**698**	99%
Xeon Silver 4208	8	2.1	2	538	**488**	91%
Xeon Bronze 3204	6	1.9	2	365	**275**	75%
March 2020 released						
Xeon Gold 6258R	28	2.7	2	4,838	**3,333**	69%
Xeon Gold 6256	12	3.6	2	2,765	**2,136**	77%
Xeon Gold 6250	8	3.9	2	1,997	**1,559**	78%
Xeon Gold 6248R	24	3.0	2	4,608	**3,124**	68%
Xeon Gold 6246R	16	3.4	2	3,482	**2,538**	73%
Xeon Gold 6242R	20	3.1	2	3,968	**2,876**	72%
Xeon Gold 6240R	24	2.4	2	3,686	**2,574**	70%
Xeon Gold 6238R	28	2.2	2	3,942	**2,696**	68%
Xeon Gold 6230R	26	2.1	2	3,494	**2,465**	71%
Xeon Gold 6226R	16	2.9	2	2,970	**2,120**	71%
Xeon Gold 6208U	16	2.9	1	1,485	**1,129**	76%
Xeon Gold 5220R	24	2.2	2	1,690	**1,494**	88%
Xeon Gold 5218R	20	2.1	2	1,344	**1,210**	90%
Xeon Silver 4215R	8	3.2	2	819	**621**	76%
Xeon Silver 4214R	12	2.4	2	922	**877**	95%
Xeon Silver 4210R	10	2.4	2	768	**748**	97%
Xeon Bronze 3206R	8	1.9	2	486	**439**	90%

Rpeak values in the table above were calculated by the base frequency of each processor. Since we enabled Turbo mode in measurements of *Rmax*, the average Turbo frequency exceeded the base frequency for some processors. That is the reason why *Efficiency* of some processors exceeds 100%.

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 cases, you should disable the turbo functionality via BIOS option.

Literature

PRIMERGY Servers

http://primergy.com/

PRIMERGY RX2540 M5

This White Paper:

http://docs.ts.fujitsu.com/dl.aspx?id=9667bde2-ed29-42be-b718-1273d8422b22

http://docs.ts.fujitsu.com/dl.aspx?id=62e48ebf-b2e4-435b-863e-abfe978a79f2

Data sheet

http://docs.ts.fujitsu.com/dl.aspx?id=3ac53e9b-a567-4c9b-8bc1-be7f5e186e0b

PRIMERGY Performance

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

SPECcpu2017

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

Benchmark Overview SPECcpu2017

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

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

SPECjbb2015

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

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

vServCon

Benchmark Overview vServCon

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

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 Product Marketing
mailto:Primergy-PM@ts.fujitsu.com
PRIMERGY Performance and Benchmarks
mailto:primergy.benchmark@ts.fujitsu.com