Fujitsu Server PRIMERGY
Performance Report
PRIMERGY RX4770 M5

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMERGY RX4770 M5.

Explaines PRIMERGY RX4770 M5 performance data in comparison to other PRIMERGY models. In addition to the benchmark results, the explanation for each benchmark and benchmark environment are also included.

Version
1.5
2023-10-03

Contents

Technical data

PRIMERGY RX4770 M5

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 RX4770 M5
Model versions	PY RX4770 M5 PY RX4770 M5 LC PY RX4770 M5 Performance
Form factor	Rack server
Chipset	Intel C624
Number of sockets	4
Number of processors orderable	PY RX4770 M5 : 2 or 4 PY RX4770 M5 LC : 4 PY RX4770 M5 Performance : 4
Processor type	2nd Generation Intel Xeon Scalable Processors Family
Number of memory slots	48 (12 per processor)
Maximum memory configuration	6,144 GB
Onboard HDD controller	Controller with RAID 0, RAID 1 or RAID 10 for up to 8 SATA HDDs
PCI slots	1 × PCI-Express 3.0 x8 3 × PCI-Express 3.0 x16
Max. number of internal hard disks	PY RX4770 M5 : 16 × 2.5" PY RX4770 M5 LC : 16 × 2.5" PY RX4770 M5 Performance : 8 × 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	2,933	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	2,933	205
Xeon Platinum 8276L	28	56	38.5	10.4	2.2	4.0	2,933	165
Xeon Platinum 8276M	28	56	38.5	10.4	2.2	4.0	2,933	165
Xeon Platinum 8276	28	56	38.5	10.4	2.2	4.0	2,933	165
Xeon Platinum 8270	26	52	35.8	10.4	2.7	4.0	2,933	205
Xeon Platinum 8268	24	48	35.8	10.4	2.9	3.9	2,933	205
Xeon Platinum 8260L	24	48	35.8	10.4	2.4	3.9	2,933	165
Xeon Platinum 8260M	24	48	35.8	10.4	2.4	3.9	2,933	165
Xeon Platinum 8260Y	24	48	35.8	10.4	2.4	3.9	2,933	165
	20	40						
	16	32						
Xeon Platinum 8260	24	48	35.8	10.4	2.4	3.9	2,933	165
Xeon Platinum 8256	4	8	16.5	10.4	3.8	3.9	2,933	105
Xeon Platinum 8253	16	32	22.0	10.4	2.2	3.0	2,933	125
Xeon Gold 6262V	24	48	33.0	10.4	1.9	3.6	2,933	135
Xeon Gold 6254	18	36	24.8	10.4	3.1	4.0	2,933	200
Xeon Gold 6252	24	48	35.8	10.4	2.1	3.7	2,933	150
Xeon Gold 6248	20	40	27.5	10.4	2.5	3.9	2,933	150
Xeon Gold 6246	12	24	24.8	10.4	3.3	4.2	2,933	165
Xeon Gold 6244	8	16	24.8	10.4	3.6	4.4	2,933	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	2,933	150
Xeon Gold 6240M	18	36	24.8	10.4	2.6	3.9	2,933	150
Xeon Gold 6240Y	18	36	24.8	10.4	2.6	3.9	2,933	150
	14	28						
	8	16						
Xeon Gold 6240	18	36	24.8	10.4	2.6	3.9	2,933	150
Xeon Gold 6238M	22	44	30.3	10.4	2.1	3.7	2,933	140
Xeon Gold 6238L	22	44	30.3	10.4	2.1	3.7	2,933	140
Xeon Gold 6238	22	44	30.3	10.4	2.1	3.7	2,933	140
Xeon Gold 6234	8	16	24.8	10.4	3.3	4.0	2,933	130
Xeon Gold 6230	20	40	27.5	10.4	2.1	3.9	2,933	125
Xeon Gold 6226	12	24	19.3	10.4	2.7	3.7	2,933	125
Xeon Gold 6222V	20	40	27.5	10.4	1.8	3.6	2,400	115
Xeon Gold 5222	4	8	16.5	10.4	3.8	3.9	2,933	105
Xeon Gold 5220S	18	36	24.8	10.4	2.7	3.9	2,666	125
Xeon Gold 5220	18	36	24.8	10.4	2.2	3.9	2,666	125
Xeon Gold 5218B	16	32	22.0	10.4	2.3	3.9	2,666	125
Xeon Gold 5218	16	32	22.0	10.4	2.3	3.9	2,666	125
Xeon Gold 5217	8	16	11.0	10.4	3.0	3.7	2,666	115

Xeon Gold 5215L	10	20	13.8	10.4	2.5	3.4	2,666	85
Xeon Gold 5215M	10	20	13.8	10.4	2.5	3.4	2,666	85
Xeon Gold 5215	10	20	13.8	10.4	2.5	3.4	2,666	85

Model	PY RX4770 M5	PY RX4770 M5 LC	PY RX4770 M5 Performance
Supported CPU	Xeon Platinum 8276L Xeon Platinum 8276M Xeon Platinum 8276 Xeon Platinum 8260L Xeon Platinum 8260M Xeon Platinum 8260Y Xeon Platinum 8260 Xeon Platinum 8256 Xeon Platinum 8253 Xeon Gold 6262V 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 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 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 Platinum 8256 Xeon Platinum 8253 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 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 Platinum 8280L Xeon Platinum 8280M Xeon Platinum 8280 Xeon Platinum 8270 Xeon Platinum 8268 Xeon Gold 6254

All the processors that can be ordered with the PRIMERGY RX4770 M5 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 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
V	VM Density Optimized
Y	Speed Select

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	2,933		✓		✓
16 GB (1x16 GB) 2Rx8 DDR4-2933 R ECC	16	2	8	2,933		✓		✓
16 GB (1x16 GB) 1Rx4 DDR4-2933 R ECC	16	1	4	2,933		✓		✓
32 GB (1x32 GB) 2Rx4 DDR4-2933 R ECC	32	2	4	2,933		✓		✓
64 GB (1x64 GB) 2Rx4 DDR4-2933 R ECC	64	2	4	2,933		✓		✓
64 GB (1x64 GB) 4Rx4 DDR4-2933 LR ECC	64	4	4	2,933	✓	✓		✓
128 GB (1x128 GB) 4Rx4 DDR4-2933 LR ECC	128	4	4	2,933	✓	✓		✓
128GB (1x128GB) DCPMM-2666	128	1	4	2,666			✓	✓
256GB (1x256GB) DCPMM-2666	256	2	4	2,666			✓	✓
512GB (1x512GB) DCPMM-2666	512	4	4	2,666			✓	✓

Power supplies		Maximum number
Modular redundant PSU	1,600W platinum PSU	2

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

Detailed technical information is available in the data sheet PRIMERGY RX4770 M5.

SPEC CPU2017

Benchmark description

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

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

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

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

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

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

Benchmark environment

System Under Test (SUT)

Hardware

• Model	PRIMERGY RX4770 M5
• Processor	2nd Generation Intel Xeon Scalable Processors Family
• Memory	48 × 32 GB (1x32 GB) 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: • Adjacent Cache Line Prefetch = Disabled • Energy Performance = Performance • Fan Control = Full • Hyper-Threading = Disabled • Intel Virtualization Technology = Disabled • Override OS Energy Performance = Enabled • Patrol Scrub = Disabled • Stale AtoS = Disabled • Sub NUMA Clustering = Disabled • WR CRC feature Control = Disabled SPECrate2017_int: • DCU Streamer Prefetcher = Disabled • Energy Performance = Performance • Hardware Prefetcher = Disabled*1 • HWPM = Native Mode with no legacy • LLC Prefetcher = Disabled*2 • Override OS Energy Performance = Enable • Patrol Scrub = Disabled • Sub Numa Clustering = Disabled*3 • WR CRC feature Control = Disabled • XPT Prefetch = Enable • Fan Control = Ful SPECrate2017_fp: • HWPM = Native Mode with no legacy • Hyper Threading = Disabled*4 • LLC Dead Line Alloc = Disabled*2 • Local/Remote Threshold = Disable • Patrol Scrub = Disabled • Sub Numa Clustering = Disabled*3 • Uncore Frequency Override = Power Balanced*1

	• WR CRC feature Control = Disabled • Fan Control = Full *1: Xeon Platinum 8280L/8280M/8280, Xeon Platinum 8276L/8276M/8276, Xeon Platinum 8270, Xeon Platinum 8268, Xeon Platinum 8260L/8260M/8260 *2: except *1 *3: Xeon Bold 5217, Xeon Gold 5215 *4: Xeon Platinum 8280L/8280M/8280, Xeon Platinum 8276L/8276M/8276, Xeon Platinum 8270, Xeon Platinum 8268, Xeon Platinum 8260L/8260M/8260
• Operating system	SPECspeed2017_int, SPECrate2017_int , SPECrate2017_fp: SUSE Linux Enterprise Server 15 4.12.14-25.28-default SPECspeed2017_fp: Red Hat Enterprise Linux Server release 7.6 3.10.0-957.el7.x86_64
• Operating system settings	Stack size set to unlimited using "ulimit -s unlimited" SPECspeed2017_fp, SPECrate2017 Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1) SPECrate2017_fp echo 10000000 > /proc/sys/kernel/sched_min_granularity_ns
• Compiler	SPECspeed2017_int, SPECspeed2017_fp, 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 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

Benchmark results

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

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

Processor	Cores	Number of Processors	SPECrate2017 int_base	SPECrate2017 fp_base
April 2019 released				
Xeon Platinum 8280L	28	4	**664**	**557**
Xeon Platinum 8280M	28	4	664 est.	557 est.
Xeon Platinum 8280	28	4	664 est.	557 est.
Xeon Platinum 8276L	28	4	**584**	**516**
Xeon Platinum 8276M	28	4	584 est.	516 est.
Xeon Platinum 8276	28	4	584 est.	516 est.
Xeon Platinum 8270	26	4	**610**	**530**
Xeon Platinum 8268	24	4	**586**	**519**
Xeon Platinum 8260L	24	4	**536**	488 est.
Xeon Platinum 8260M	24	4	536 est.	488 est.
Xeon Platinum 8260Y	24	4	**542**	**489**
	20	4	476 est.	437 est.
	16	4	409 est.	399 est.
Xeon Platinum 8260	24	4	536 est.	488 est.
Xeon Platinum 8256	4	4	**127**	**152**
Xeon Platinum 8253	16	4	**334**	**344**
Xeon Gold 6262V	24	4	**475**	**417**
Xeon Gold 6254	18	4	**500**	**452**
Xeon Gold 6252	24	4	**528**	**476**
Xeon Gold 6248	20	4	**479**	**444**
Xeon Gold 6246	12	4	**360**	**370**
Xeon Gold 6244	8	4	**264**	**293**
Xeon Gold 6242	16	4	**418**	**402**
Xeon Gold 6240L	18	4	442 est.	418 est.
Xeon Gold 6240M	18	4	442 est.	418 est.
Xeon Gold 6240Y	18	4	**438**	**416**
	14	4	373 est.	376 est.
	8	4	236 est.	273 est.
Xeon Gold 6240	18	4	**442**	**418**
Xeon Gold 6238L	22	4	484 est.	448 est.
Xeon Gold 6238M	22	4	484 est.	448 est.
Xeon Gold 6238	18	4	**484**	**448**
Xeon Gold 6234	22	4	**247**	**275**
Xeon Gold 6230	20	4	**433**	**412**
Xeon Gold 6226	12	4	**324**	**337**

Xeon Gold 6222V	20	4	**395**	**369**
Xeon Gold 5222	4	4	**127**	**151**
Xeon Gold 5220S	18	4	**397**	**383**
Xeon Gold 5220	18	4	**398**	**383**
Xeon Gold 5218B	16	4	364 est.	358 est.
Xeon Gold 5218	16	4	**364**	**358**
Xeon Gold 5217	8	4	**213**	**232**
Xeon Gold 5215L	10	4	232 est.	251 est.
Xeon Gold 5215M	10	4	232 est.	251 est.
Xeon Gold 5215	10	4	**232**	**251**

Processor	Cores	Number of Processors	SPECspeed2017 int_base	SPECspeed2017 fp_base
April 2019 released				
Xeon Platinum 8280	28	4	-	157 est.
Xeon Gold 6244	8	4	10.8 est.	-

The following two diagrams illustrate the throughput of the PRIMERGY RX4770 M5 in comparison to its predecessor PRIMERGY RX2540 M4, in their respective most performant configuration.
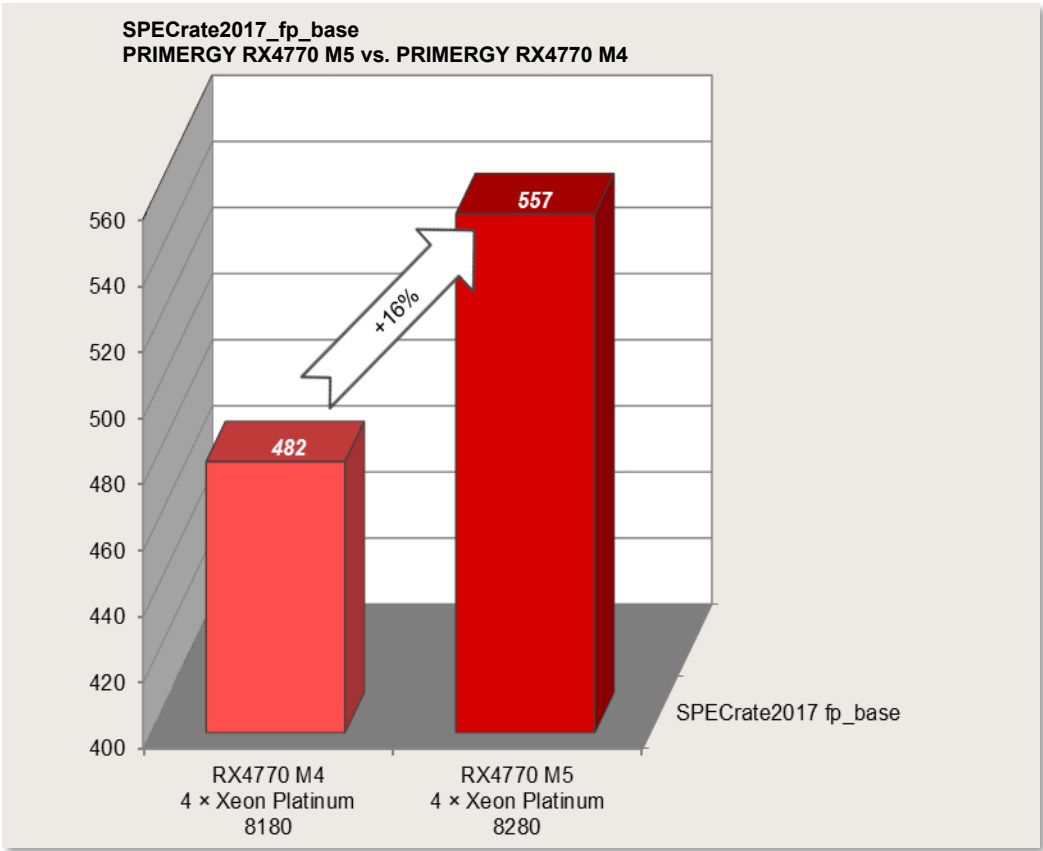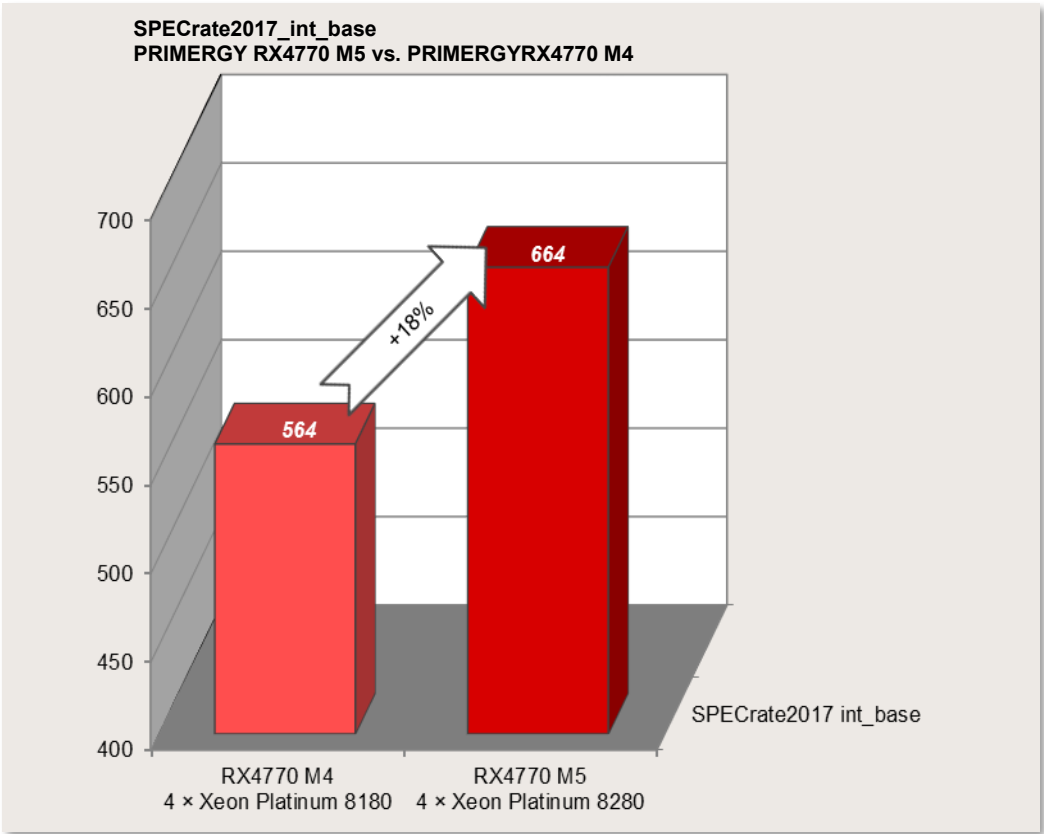
STREAM

Benchmark description

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

The benchmark is designed in such a way that it can be used both on PCs and on server systems. The unit of measurement of the benchmark is GB/s, i.e. the number of gigabytes that can be read and written per second.

STREAM measures the memory throughput for sequential accesses. These can generally be performed more efficiently than accesses that are randomly distributed on the memory, because the processor caches are used for sequential access.

Before execution the source code is adapted to the environment to be measured. Therefore, the size of the data area must be at least 12 times larger than the total of all last-level processor caches so that these have as little influence as possible on the result. The OpenMP program library is used to enable selected parts of the program to be executed in parallel during the runtime of the benchmark. This provides optimal load distribution for the available processor cores.

In the STREAM benchmark, a data area consisting of 8-byte elements is continuously copied to four operation types. Arithmetic operations are also performed on operation types other than COPY.

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

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

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

In this chapter, throughputs are indicated as a power of 10. (1 GB/s = 10^9 Byte/s)

Benchmark environment

System Under Test (SUT)

Hardware

• Model	PRIMERGY RX4770 M5
• Processor	2nd Generation Intel Xeon Scalable Processors Family
• Memory	48 × 32 GB (1x32 GB) 2Rx4 PC4-2933Y-R

Software

• BIOS settings	• MC Interleaving = 1-way • Override OS Energy Performance = Enabled • Energy Performance = Performance • HWPM Support = Disabled • Intel Virtualization Technology = Disabled • LLC Dead Line Alloc = Disabled • Stale AtoS = Enabled
• Operating system	SUSE Linux Enterprise Server 15
• Operating system settings	Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1) Transparent Huge Pages inactivated sched_cfs_bandwidth_slice_us = 50000 sched_latency_ns = 240000000 sched_migration_cost_ns = 5000000 sched_min_granularity_ns = 100000000 sched_wakeup_granularity_ns = 150000000 cpupower -c all frequency-set -g performance cpupower idle-set -d 1 cpupower idle-set -d 2 cpupower idle-set -d 3 echo 0 > /proc/sys/kernel/numa_balancing echo 1 > /proc/sys/vm/drop_caches ulimit -s unlimited
• Compiler	CPU released in April 2019 C/C++: Version 2019.3.0.591499 of Intel C/C++ Compiler for Linux
• Benchmark	STREAM Version 5.10

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

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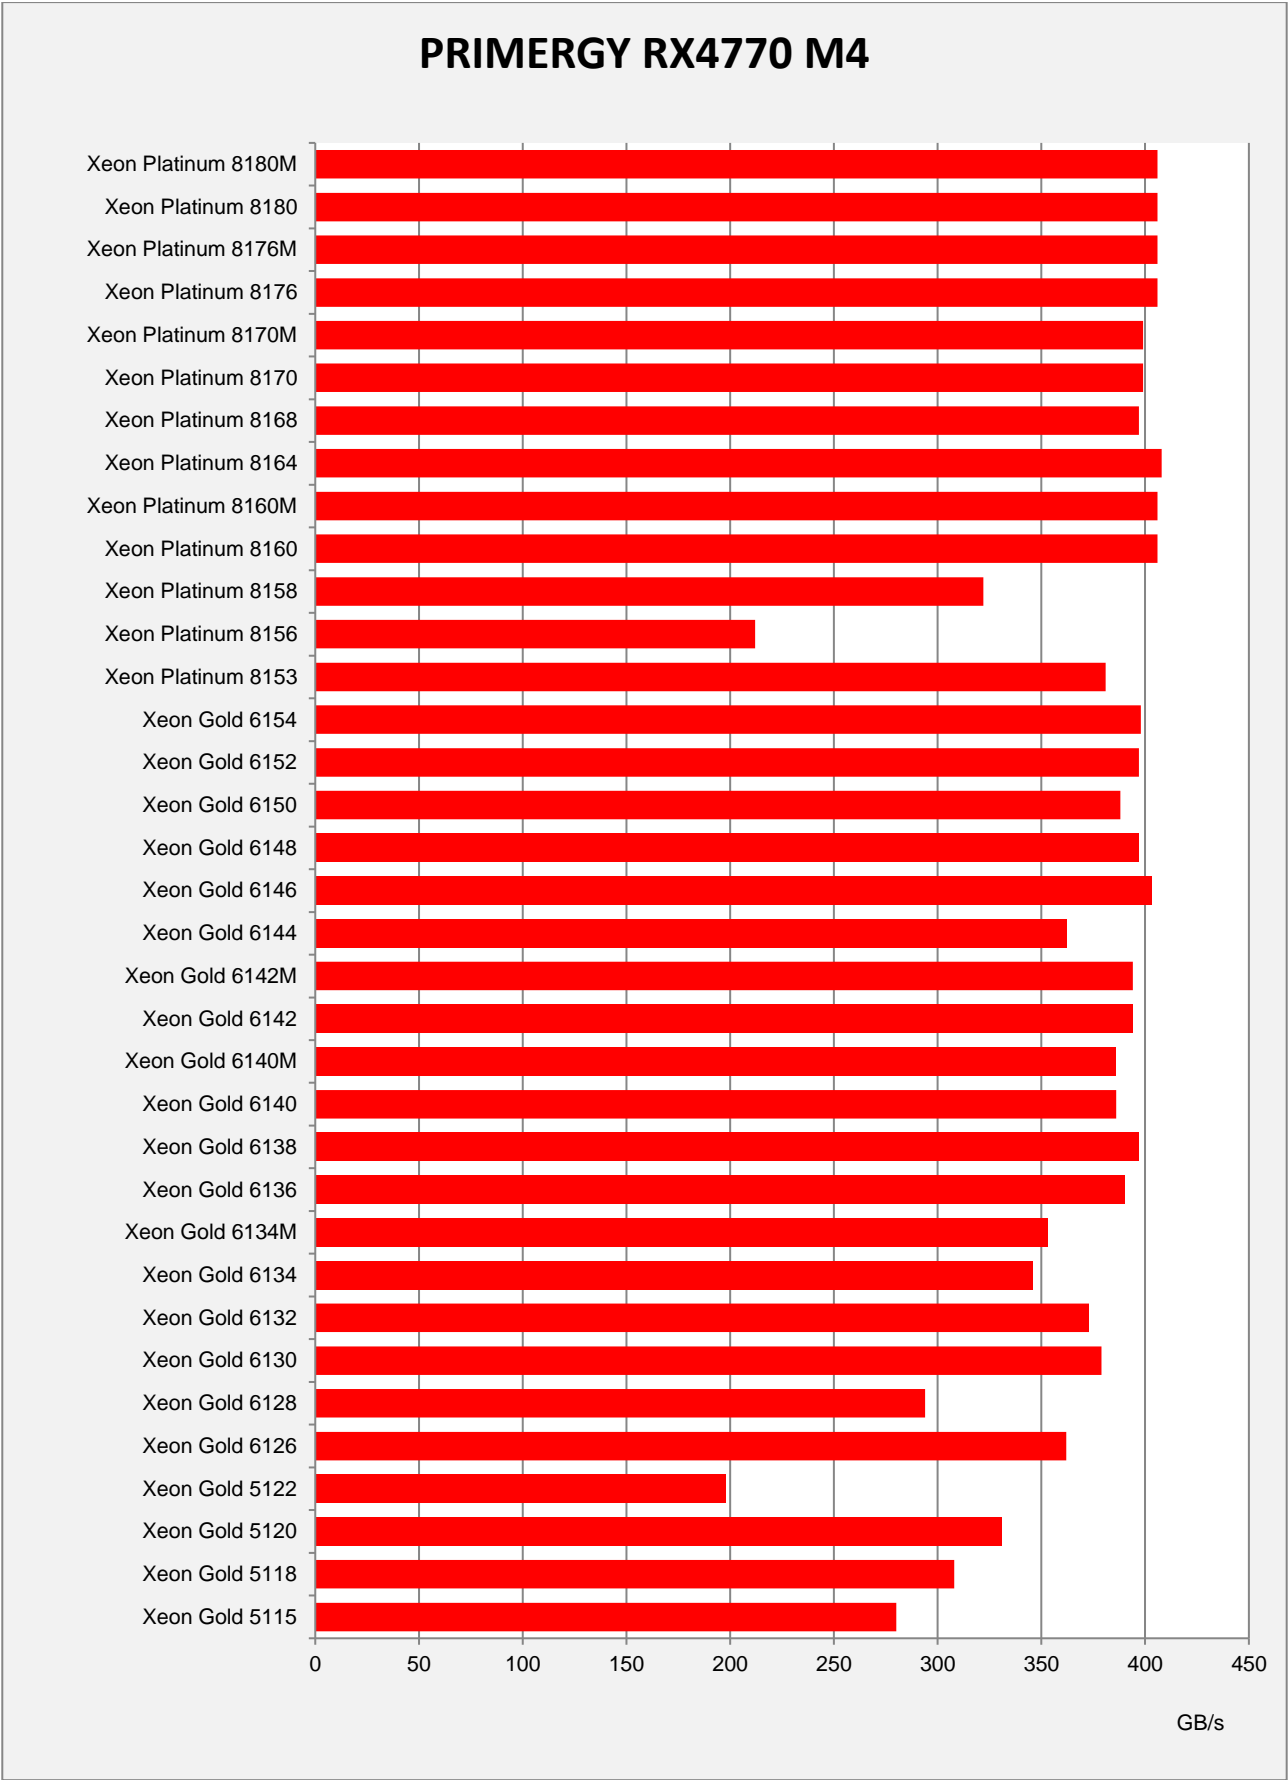
The results with "est." are the estimated values.

Processor	Memory Frequency [MHz]	Max. Memory Bandwidth*[1] [GB/s]	Cores	Processor Frequency [GHz]	Number of Processors	TRIAD [GB/s]	
April 2019 released							
Xeon Platinum 8280L	2,933	140.8	28	2.7	4	**422**	
Xeon Platinum 8280M	2,933	140.8	28	2.7	4	422	est.
Xeon Platinum 8280	2,933	140.8	28	2.7	4	422	est.
Xeon Platinum 8276L	2,933	140.8	28	2.2	4	**422**	
Xeon Platinum 8276M	2,933	140.8	28	2.2	4	422	est.
Xeon Platinum 8276	2,933	140.8	28	2.2	4	422	est.
Xeon Platinum 8270	2,933	140.8	26	2.7	4	**417**	
Xeon Platinum 8268	2,933	140.8	24	2.9	4	**414**	
Xeon Platinum 8260L	2,933	140.8	24	2.4	4	**414**	
Xeon Platinum 8260M	2,933	140.8	24	2.4	4	414	est.
Xeon Platinum 8260Y	2,933	140.8	24	2.4	4	**414**	
	2,933	140.8	20	2.4	4	**397**	
	2,933	140.8	16	2.4	4	**392**	
Xeon Platinum 8260	2,933	140.8	24	2.4	4	414	est.
Xeon Platinum 8256	2,933	140.8	4	3.8	4	**213**	
Xeon Platinum 8253	2,933	140.8	16	2.2	4	**388**	
Xeon Gold 6262V	2,933	140.8	24	1.9	4	**364**	
Xeon Gold 6254	2,933	140.8	18	3.1	4	**396**	
Xeon Gold 6252	2,933	140.8	24	2.1	4	**415**	
Xeon Gold 6248	2,933	140.8	20	2.5	4	**405**	
Xeon Gold 6246	2,933	140.8	12	3.3	4	**380**	
Xeon Gold 6244	2,933	140.8	8	3.6	4	**337**	
Xeon Gold 6242	2,933	140.8	16	2.8	4	**389**	
Xeon Gold 6240L	2,933	140.8	18	2.6	4	372	est.
Xeon Gold 6240M	2,933	140.8	18	2.6	4	372	est.
Xeon Gold 6240Y	2,933	140.8	18	2.6	4	**397**	
	2,933	140.8	14	2.6	4	**382**	
	2,933	140.8	8	2.6	4	**339**	
Xeon Gold 6240	2,933	140.8	18	2.6	4	**372**	
Xeon Gold 6238M	2,933	140.8	22	2.1	4	386	est.
Xeon Gold 6238L	2,933	140.8	22	2.1	4	386	est.
Xeon Gold 6238	2,933	140.8	22	2.1	4	**386**	
Xeon Gold 6234	2,933	140.8	8	3.3	4	**337**	
Xeon Gold 6230	2,933	140.8	20	2.1	4	**404**	
Xeon Gold 6226	2,933	140.8	12	2.7	4	**367**	
Xeon Gold 6222V	2,400	140.8	20	1.8	4	**351**	
Xeon Gold 5222	2,933	140.8	4	3.8	4	**183**	
Xeon Gold 5220S	2,666	128.0	18	2.7	4	**338**	
Xeon Gold 5220	2,666	128.0	18	2.2	4	**353**	

Xeon Gold 5218B	2,666	128.0	16	2.3	4	338	est.
Xeon Gold 5218	2,666	128.0	16	2.3	4	**338**	
Xeon Gold 5217	2,666	128.0	8	3	4	**211**	
Xeon Gold 5215L	2,666	128.0	10	2.5	4	228	est.
Xeon Gold 5215M	2,666	128.0	10	2.5	4	228	est.
Xeon Gold 5215	2,666	128.0	10	2.5	4	**228**	

*1: Value per Processor

The following diagram illustrates the throughput of the PRIMERGY RX4770 M5 in comparison to its predecessor, the RX4770 M4

PRIMERGY RX4770 M5

Processor	Throughput (GB/s)
Xeon Platinum 8280L	420
Xeon Platinum 8280M	420
Xeon Platinum 8280	420
Xeon Platinum 8276L	420
Xeon Platinum 8276M	420
Xeon Platinum 8276	420
Xeon Platinum 8270	415
Xeon Platinum 8268	412
Xeon Platinum 8260L	412
Xeon Platinum 8260M	412
Xeon Platinum 8260Y	412
Xeon Platinum 8260	412
Xeon Platinum 8256	212
Xeon Platinum 8253	388
Xeon Gold 6262V	363
Xeon Gold 6254	395
Xeon Gold 6252	415
Xeon Gold 6248	405
Xeon Gold 6246	380
Xeon Gold 6244	336
Xeon Gold 6242	388
Xeon Gold 6240L	371
Xeon Gold 6240M	371
Xeon Gold 6240Y	397
Xeon Gold 6240	371
Xeon Gold 6238M	385
Xeon Gold 6238L	385
Xeon Gold 6238	385
Xeon Gold 6234	337
Xeon Gold 6230	405
Xeon Gold 6226	367
Xeon Gold 6222V	350
Xeon Gold 5222	182
Xeon Gold 5220S	337
Xeon Gold 5220	355
Xeon Gold 5218B	337
Xeon Gold 5218	337
Xeon Gold 5217	210
Xeon Gold 5215L	229
Xeon Gold 5215M	229
Xeon Gold 5215	229
Xeon Gold 6256	393
Xeon Gold 6250	327

GB/s

.

PRIMERGY RX4770 M4

Processor	GB/s
Xeon Platinum 8180M	405
Xeon Platinum 8180	405
Xeon Platinum 8176M	405
Xeon Platinum 8176	405
Xeon Platinum 8170M	398
Xeon Platinum 8170	398
Xeon Platinum 8168	396
Xeon Platinum 8164	408
Xeon Platinum 8160M	405
Xeon Platinum 8160	405
Xeon Platinum 8158	322
Xeon Platinum 8156	212
Xeon Platinum 8153	380
Xeon Gold 6154	398
Xeon Gold 6152	397
Xeon Gold 6150	388
Xeon Gold 6148	396
Xeon Gold 6146	403
Xeon Gold 6144	362
Xeon Gold 6142M	394
Xeon Gold 6142	394
Xeon Gold 6140M	385
Xeon Gold 6140	385
Xeon Gold 6138	397
Xeon Gold 6136	390
Xeon Gold 6134M	352
Xeon Gold 6134	345
Xeon Gold 6132	372
Xeon Gold 6130	378
Xeon Gold 6128	294
Xeon Gold 6126	362
Xeon Gold 5122	197
Xeon Gold 5120	331
Xeon Gold 5118	308
Xeon Gold 5115	280

GB/s

LINPACK

Benchmark description

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

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

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

A memory of $8n^2$ bytes is required for the matrix. In case of an n x n matrix the number of arithmetic operations required for the solution is 2/3n3 + 2n2. Thus, the choice of n determines the duration of the measurement. In other words, if n is doubled, the measurement time will be approximately eight times longer. The size of n also has an influence on the measurement result itself. As n increases, the measured value asymptotically approaches its limit. The size of the matrix is therefore usually adapted to the amount of memory available. Furthermore, the memory bandwidth of the system only plays a minor role for the measurement result, but a role that cannot be fully ignored. he processor performance is the decisive factor for the measurement result. Since the algorithm used permits parallel processing, in particular the number of processors used and their processor cores are - in addition to the clock rate - of outstanding significance.

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

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

Rpeak =   Maximum number of floating point operations per clock cycle

            x Number of processor cores of the computer

            x Rated processor frequency [GHz]

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

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

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

http://software.intel.com/en-us/articles/intel-math-kernel-library-linpack-download/.

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

Benchmark environment

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY RX4770 M5
• Processor	2nd Generation Intel Xeon Scalable Processors Family
• Memory	48 × 32 GB (1x32 GB) 2Rx4 PC4-2933Y-R
Software	
• BIOS settings	• HyperThreading = Disabled • Link Frequency Select = 10.4 GT/s • HWPM Support = Disabled • Intel Virtualization Technology = Disabled • Sub NUMA Clustering = Disabled • LLC Dead Line Alloc = Disabled • Stale AtoS = Enabled • Write CRC = 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 sched_min_granularity_ns = 100000000 sched_wakeup_granularity_ns = 150000000 run with avx512
• Compiler	CPU released in April 2019 C/C++: Version 2019.3.0.591499 of Intel C/C++ Compiler for Linux
• Benchmark	Intel Optimized MP LINPACK Benchmark for Clusters

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

Benchmark results

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

Processor	Cores	Processor Frequency [GHz]	Number of Processors	Rpeak [GFlops]	Rmax [GFlops]	Efficiency
April 2019 released						
Xeon Platinum 8280L	28	2.7	4	9,677	**7,028**	73%
Xeon Platinum 8280M	28	2.7	4	9,677	7,028 est.	73%
Xeon Platinum 8280	28	2.7	4	9,677	7,028 est.	73%
Xeon Platinum 8276L	28	2.2	4	7,885	**5,403**	69%
Xeon Platinum 8276M	28	2.2	4	7,885	5,403 est.	69%
Xeon Platinum 8276	28	2.2	4	7,885	5,403 est.	69%
Xeon Platinum 8270	26	2.7	4	8,986	**6,166**	69%
Xeon Platinum 8268	24	2.9	4	8,909	**6,070**	68%
Xeon Platinum 8260L	24	2.4	4	7,373	5,230 est.	71%
Xeon Platinum 8260M	24	2.4	4	7,373	5,230 est.	71%
Xeon Platinum 8260Y	24	2.4	4	7,373	**5,304**	72%
	20	2.4	4	6,144	**4,694**	76%
	16	2.4	4	4,915	**4,149**	84%
Xeon Platinum 8260	24	2.4	4	7,373	5,230 est.	71%
Xeon Platinum 8256	4	3.8	4	1,946	**1,542**	79%
Xeon Platinum 8253	16	2.2	4	4,506	**3,145**	70%
Xeon Gold 6262V	24	1.9	4	5,837	**4,212**	72%
Xeon Gold 6254	18	3.1	4	7,142	**5,206**	73%
Xeon Gold 6252	24	2.1	4	6,451	**4,930**	76%
Xeon Gold 6248	20	2.5	4	6,400	**4,551**	71%
Xeon Gold 6246	12	3.3	4	5,069	**3,673**	72%
Xeon Gold 6244	8	3.6	4	3,686	**2,861**	78%
Xeon Gold 6242	16	2.8	4	5,734	**4,172**	73%
Xeon Gold 6240L	18	2.6	4	5,990	4,286 est.	72%
Xeon Gold 6240M	18	2.6	4	5,990	4,286 est.	72%
Xeon Gold 6240Y	18	2.6	4	5,990	**4,251**	71%
	14	2.6	4	4,659	**3,704**	80%
	8	2.6	4	2,662	**2,744**	103%
Xeon Gold 6240	18	2.6	4	5,990	**4,286**	72%
Xeon Gold 6238M	22	2.1	4	5,914	4,413 est.	75%
Xeon Gold 6238L	22	2.1	4	5,914	4,413 est.	75%
Xeon Gold 6238	22	2.1	4	5,914	**4,413**	75%
Xeon Gold 6234	8	3.3	4	3,379	**2,607**	77%
Xeon Gold 6230	20	2.1	4	5,376	**3,773**	70%
Xeon Gold 6226	12	2.7	4	4,147	**3,295**	79%
Xeon Gold 6222V	20	1.8	4	4,608	**3,623**	79%
Xeon Gold 5222	4	3.8	4	1,946	**1,539**	79%
Xeon Gold 5220S	18	2.7	4	3,110	**2,443**	79%
Xeon Gold 5220	18	2.2	4	2,534	**2,448**	97%
Xeon Gold 5218B	16	2.3	4	2,355	2,236 est.	95%

Xeon Gold 5218	16	2.3	4	2,355	**2,236**		95%
Xeon Gold 5217	8	3	4	1,536	**1,434**		93%
Xeon Gold 5215L	10	2.5	4	1,600	1,468	est.	92%
Xeon Gold 5215M	10	2.5	4	1,600	1,468	est.	92%
Xeon Gold 5215	10	2.5	4	1,600	**1,468**		92%

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

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

SPECpower_ssj2008

Benchmark description

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

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

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

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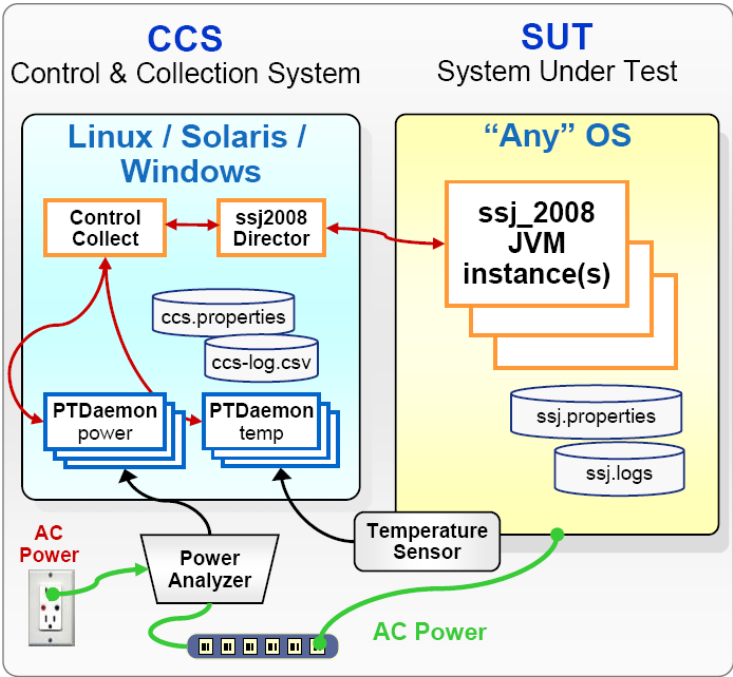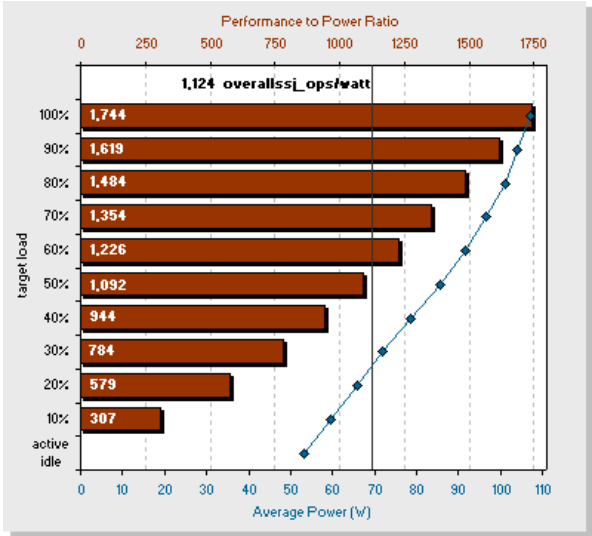
The benchmark runs on a wide variety of operating systems and hardware architectures, and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System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

Benchmark environment

System Under Test (SUT)	
For Windows OS measurement	
Hardware	
• Model	PRIMERGY RX4770 M5
• Processor	Intel Xeon Platinum 8276L
• Memory	24 ×16 GB (1x16 GB) 2Rx8 PC4-2933 R ECC
• Network interface	1 × PLAN EM 4 x 1Gb T OCP(Dynamic LOM), S26361-F3953-E401
• Disk subsystem	1 × SSD M.2 240GB, S26361-F5706-E240
• Power Supply Unit	1 × Modular PSU 1600 W platinum hp
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 = 2666 MHz)
	Autonomous C-state Support = Enabled.
	ASPM Support = Force L0s.
	UPI Link Frequency Select = 9.6GT/s.
	Uncore Frequency Override = Power balanced.
	IMC Interleaving = 1-way.
	HWPM = Disabled.
	USB Port Control = Enable internal ports only.
	Network Stack = Disabled.
	Link L0p Enable = Enabled.
	Link L1 Enable = Enabled.
• Firmware	2.20c
• 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%.
	Disabled 3 NIC ports.
	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 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: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 RX4770 M5
• Processor	Intel Xeon Platinum 8280L
• Memory	24 ×16 GB (1x16 GB) 2Rx8 PC4-2933 R ECC
• Network interface	1 × PLAN EM 4 x 1Gb T OCP(Dynamic LOM), S26361-F3953-E401
• Disk subsystem	1 × SSD M.2 240GB, S26361-F5706-E240
• Power Supply Unit	1 × Modular PSU 1600 W platinum hp

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 = 2666 MHz) Autonomous C-state Support = Enabled. ASPM Support = Force L0s. UPI Link Frequency Select = 9.6GT/s. Uncore Frequency Override = Power balanced. IMC Interleaving = 1-way. HWPM = Disabled. USB Port Control = Enable internal ports only. Network Stack = Disabled. Link L0p Enable = Enabled. Link L1 Enable = Enabled.
• Firmware	2.20c
• 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=1-223 nohz_full=1-223 isolcpus=1-223 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 2400MHz
	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
	-XX:UseAVX=0

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

Benchmark results(Windows)

The PRIMERGY RX4770 M5 in Microsoft Windows Server 2016 Standard achieved the following result:

SPECpower_ssj2008 = 11,444 overall ssj_ops/watt
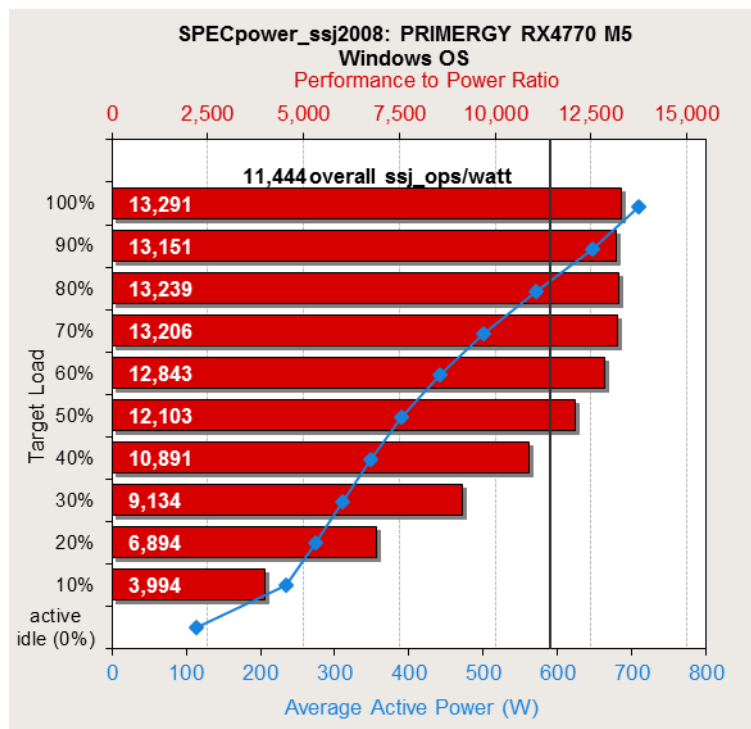
The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 11,444 overall ssj_ops/watt for the PRIMERGY RX4770 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%	9,432,104	710	13,291
90%	8,512,403	647	13,151
80%	7,566,232	572	13,239
70%	6,622,170	501	13,206
60%	5,675,935	442	12,843
50%	4,725,788	390	12,103
40%	3,787,516	348	10,891
30%	2,836,687	311	9,134
20%	1,888,531	274	6,894
10%	940,508	235	3,994
Active Idle	0	113	0
			∑ssj_ops / ∑power = 11,444

Benchmark results(Linux)

The PRIMERGY RX4770 M5 in SUSE Linux Enterprise Server 12 SP4 achieved the following result:

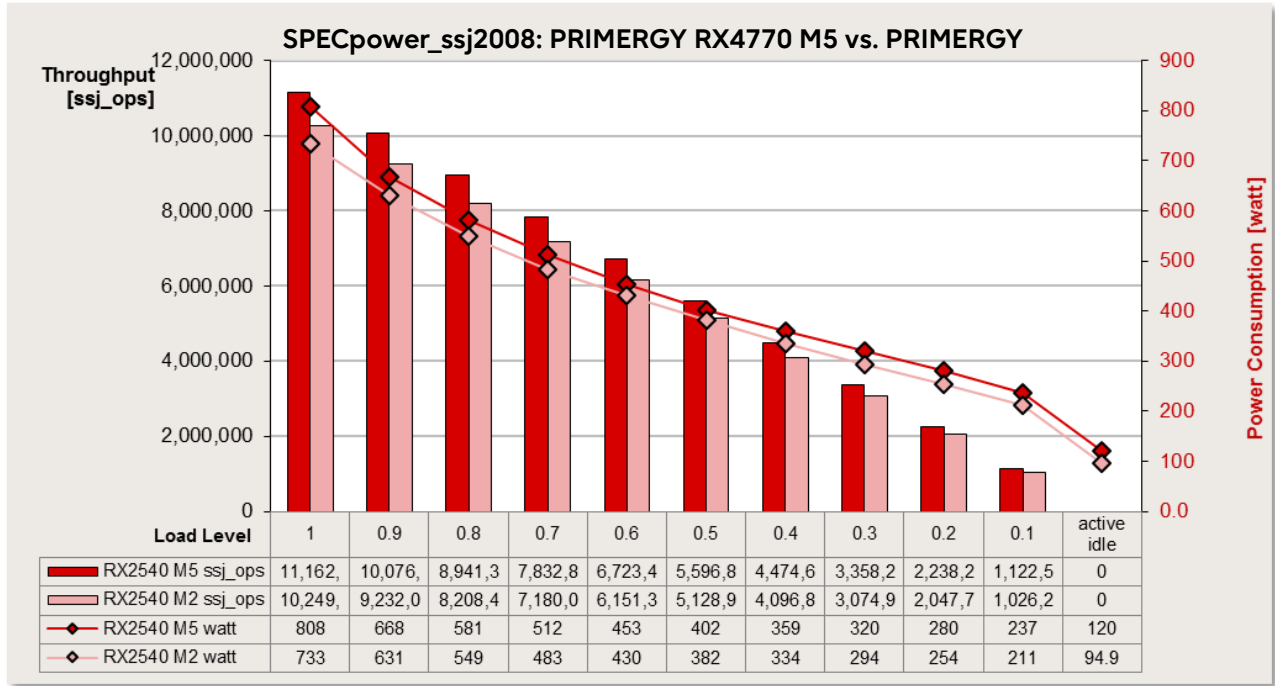### SPECpower_ssj2008 = 12,979 overall ssj_ops/watt
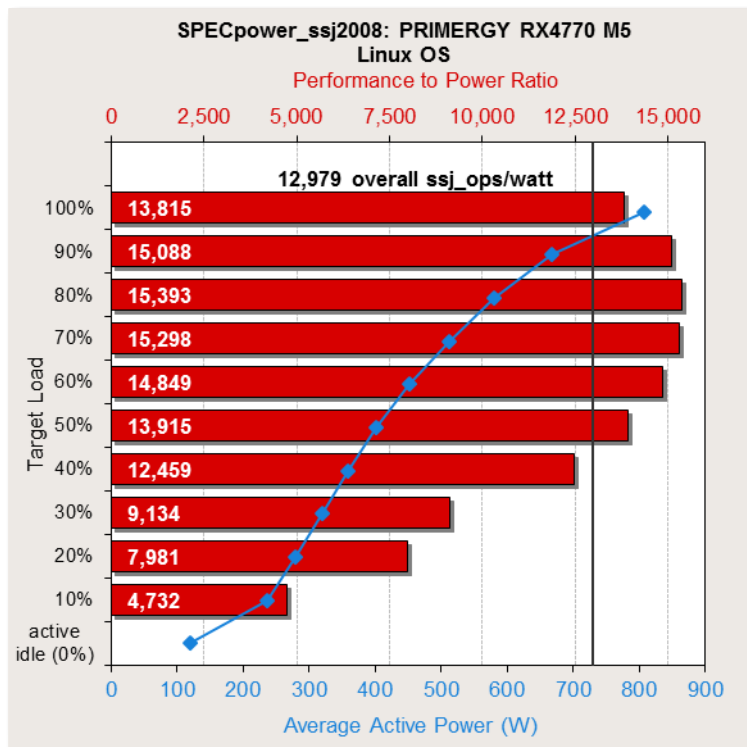
The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 12,979 overall ssj_ops/watt for the PRIMERGY RX4770 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.
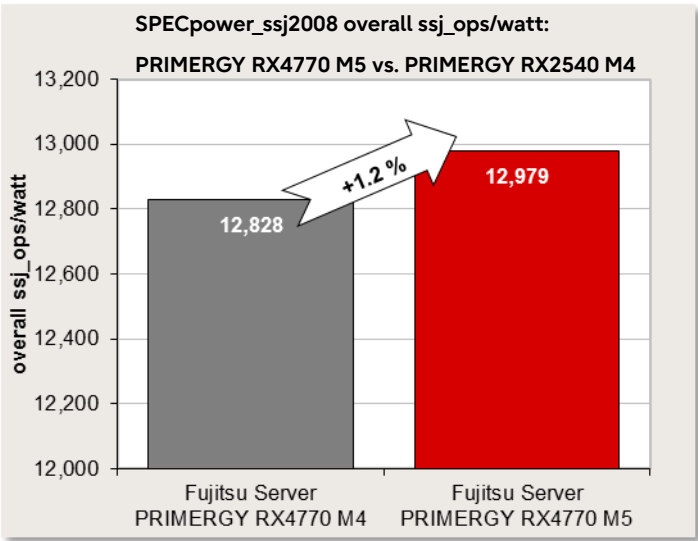
The following table shows the benchmark results for the throughput in ssj_ops, the power consumption in watts and the resulting energy efficiency for each load level.

Performance		Power	Energy Efficiency
Target Load	ssj_ops	Average Power (W)	ssj_ops/watt
100%	11,162,949	808	13,815
90%	10,076,651	668	15,088
80%	8,941,321	581	15,393
70%	7,832,822	512	15,298
60%	6,723,475	453	14,849
50%	5,596,891	402	13,915
40%	4,474,650	359	12,459
30%	3,358,285	320	9,134
20%	2,238,257	280	7,981
10%	1,122,533	237	4,732
Active Idle	0	120	0
			∑ssj_ops / ∑power = 12,979

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 RX4770 M5 compared to the predecessor PRIMERGY RX2540 M4.

SPECpower_ssj2008: PRIMERGY RX4770 M5 vs. PRIMERGY

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	11,162,	10,076,	8,941,3	7,832,8	6,723,4	5,596,8	4,474,6	3,358,2	2,238,2	1,122,5	0
RX2540 M2 ssj_ops	10,249,	9,232,0	8,208,4	7,180,0	6,151,3	5,128,9	4,096,8	3,074,9	2,047,7	1,026,2	0
RX2540 M5 watt	808	668	581	512	453	402	359	320	280	237	120
RX2540 M2 watt	733	631	549	483	430	382	334	294	254	211	94.9

Thanks to the latest Scalable Family processors, the PRIMERGY RX4770 M5 has a higher throughput. This results in an overall 1.2% increase in energy efficiency in the PRIMERGY RX4770 M5.

SPECpower_ssj2008 overall ssj_ops/watt:
PRIMERGY RX4770 M5 vs. PRIMERGY RX2540 M4

+1.2 %

Fujitsu Server PRIMERGY RX4770 M4: 12,828
Fujitsu Server PRIMERGY RX4770 M5: 12,979

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
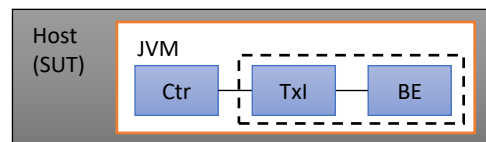

The SPECjbb2015 benchmark has a two performance metrics:

- 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
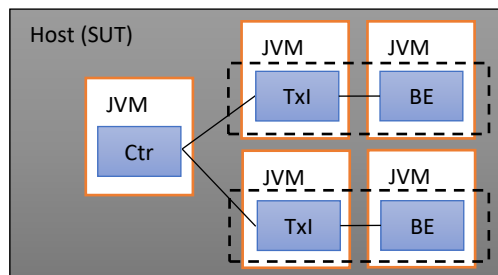
(b) example of SPECjbb2015 MultiJVM configuration
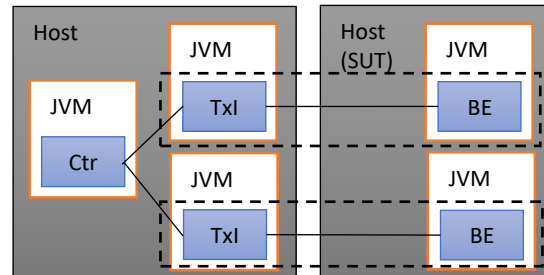
(c) example of SPECjbb2015 Distributed configuration

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 RX4770 M5 was configured for the SPECjbb2015 Composite benchmark measurement.

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY RX4770 M5
• Processor	4 × Xeon Platinum 8280L
• Memory	48 × 32 GB (1x32 GB) 2Rx4 DDR4-2933 R ECC
• Network interface	1 Gbit/s LAN
• Disk subsystem	Disk : 1 × M.2 SSD 240 GB
Software	
• 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	Red Hat Enterprise Linux 7.6 (kernel 3.10.0-957.el7.x86_64-default was applied)
• Operating system settings	ulimit -n 1048576 ulimit -i unlimited ulimit -s unlimited ulimit -u unlimited ulimit -l unlimited ulimit -v unlimited ulimit -m unlimited cpupower -c all frequency-set -g performance echo 6000000 > /proc/sys/kernel/sched_min_granularity_ns echo 1000 > /proc/sys/kernel/sched_migration_cost_ns echo 990000 > /proc/sys/kernel/sched_rt_runtime_us echo 0 > /proc/sys/kernel/numa_balancing tuned-adm profile latency-performance Disable SELinux
• JVM	Oracle Java SE 11.0.2
• JVM settings	-server -Xms1380g -Xmx1380g -Xmn1350g -XX:SurvivorRatio=69 -XX:MaxTenuringThreshold=15 -XX:+UseLargePages -XX:LargePageSizeInBytes=2m -XX:+UseParallelOldGC -Xnoclassgc -XX:+AggressiveOpts -XX:+UseNUMA -XX:-UseBiasedLocking -XX:+AlwaysPreTouch -XX:-UseAdaptiveSizePolicy -XX:-UsePerfData -XX:TargetSurvivorRatio=90 -XX:ParallelGCThreads=224 -verbose:gc -XX:+PrintGCDetails -XX:+UseHugeTLBFS -XX:+UseTransparentHugePages -XX:+AggressiveHeap

• SPECjbb2015 settings	specjbb.comm.connect.client.pool.size = 64;
	specjbb.comm.connect.timeouts.connect = 700000;
	specjbb.comm.connect.timeouts.read = 700000;
	specjbb.comm.connect.timeouts.write = 700000;
	specjbb.controller.rtcurve.warmup.step = 0.9;
	specjbb.customerDriver.threads = {probe=96, saturare=96};
	specjbb.forkjoin.workers = {Tier1=222, Tier2=1, Tier3=50};
	specjbb.mapreducer.pool.size = 256;

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

Benchmark results

"SPECjbb2015 Composite" measurement result (April 2, 2019)

259,141 SPECjbb2015-Composite max-jOPS

207,097 SPECjbb2015-Composite critical-jOPS

On April 2, 2019 PRIMERGY RX4770 M5 with two Intel Xeon Platinum 8280L processor achieved the scores of 207,097 SPECjbb2015-Composite critical-jOPS in Red Hat Enterprise Linux 7.6. With the result, it ranked first in the 4-socket server category for SPECjbb2015-Composite 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: https://www.sap.com/about/benchmark.html.

Benchmark environment

The typical measurement set-up is illustrated below:
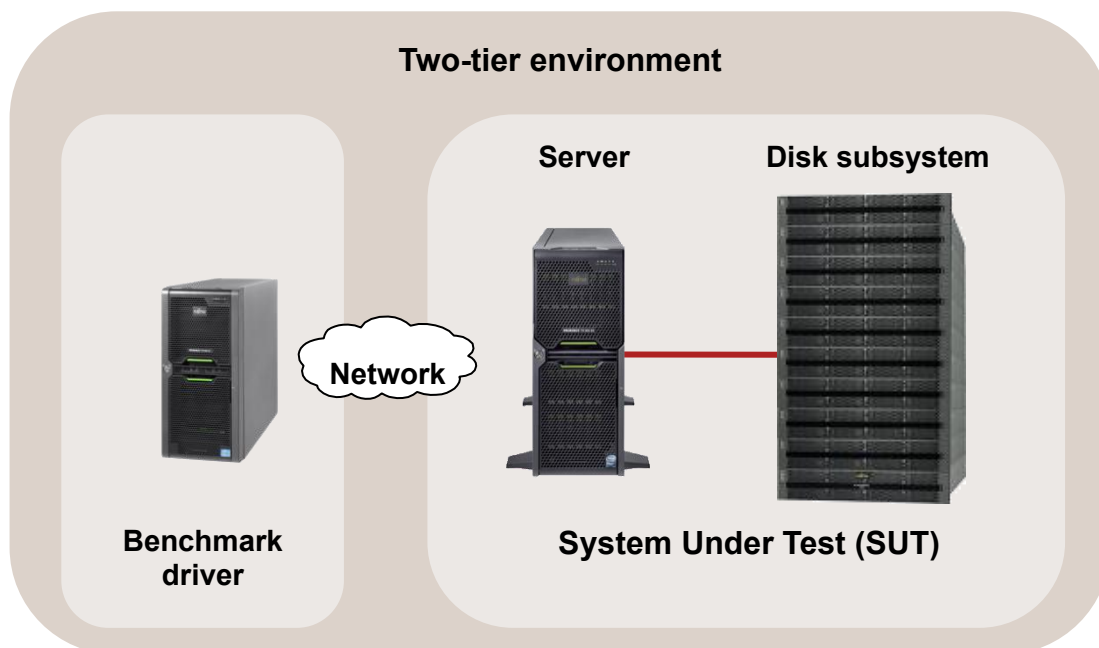
System Under Test (SUT)

Hardware

• Model	PRIMERGY RX4770 M5
• Processor	4 × Xeon Platinum 8280
• Memory	48 × 32 GB (1x32 GB) 2Rx4 DDR4-2933 R ECC
• Network interface	1 Gbit/s LAN
• Disk subsystem	PRIMERGY RX4770 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
• 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

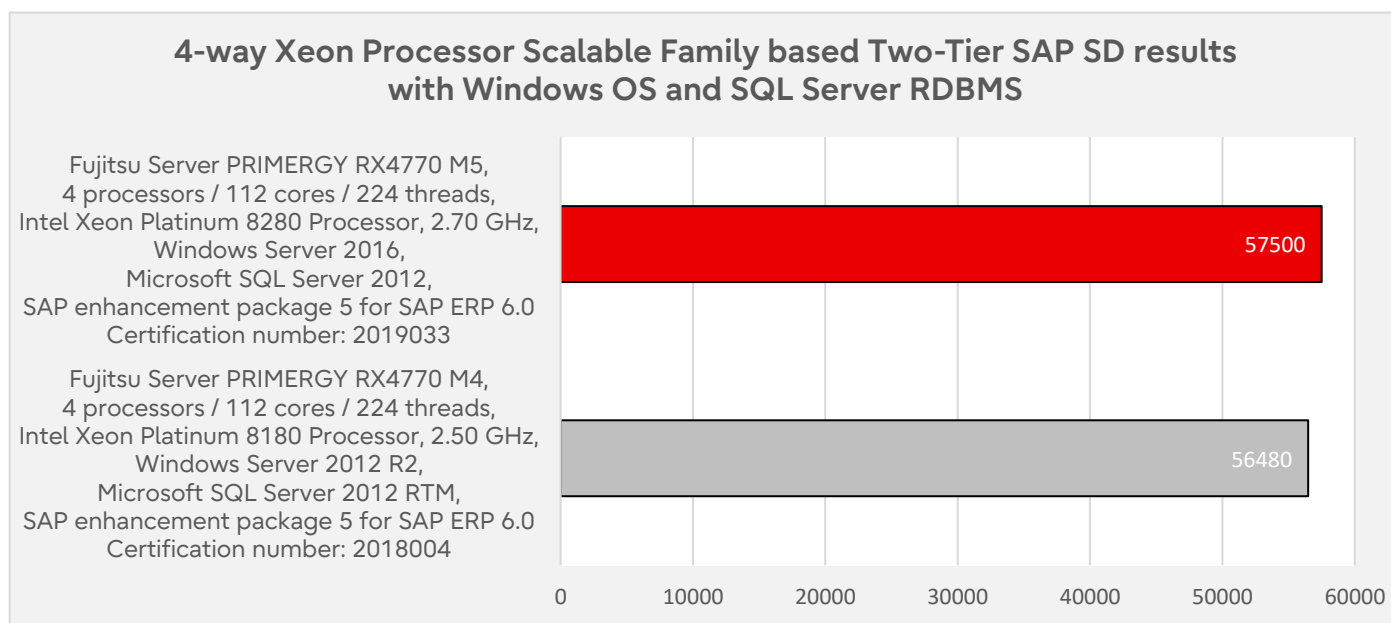|	
• Operating system	SUSE Linux Enterprise Server 12 SP2

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

Benchmark results

Certification number 2019010	
Number of SAP SD benchmark users	57,500
Average dialog response time	0.97 seconds

Throughput
    Fully processed order line items/hour
    Dialog steps/hour
    SAPS
	6,287,670
	18,863,000
	314,380

Average database request time (dialog/update)	0.010 sec / 0.021 sec
CPU utilization of central server	98%
Operating system, central server	Windows Server 2016
RDBMS	Microsoft SQL Server 2012
SAP Business Suite software	SAP enhancement package 5 for SAP ERP 6.0
Configuration Central Server	Fujitsu Server PRIMERGY RX4770 M5, 4 processors / 112 cores / 224 threads, Intel Xeon Platinum 8280 Processor, 2.70 GHz, 64 KB L1 cache per core and 1,024KB L2 cache per core, 38.5 MB L3 cache per socket, 1,536 GB main memory

The following chart shows a comparison of two-tier SAP SD Standard Application Benchmark results for 4 way Xeon Processor Scalable Family based servers with Windows OS and SQL Server database (as of June 18, 2019). The PRIMERGY RX4770 M5 outperforms its predecessor PRIMERGY RX4770 M4. The latest SAP SD 2-tier results can be found at https://www.sap.com/dmc/exp/2018-benchmark-directory/#/sd.

4-way Xeon Processor Scalable Family based Two-Tier SAP SD results with Windows OS and SQL Server RDBMS

Fujitsu Server PRIMERGY RX4770 M5, 4 processors / 112 cores / 224 threads, Intel Xeon Platinum 8280 Processor, 2.70 GHz, Windows Server 2016, Microsoft SQL Server 2012, SAP enhancement package 5 for SAP ERP 6.0 Certification number: 2019033	57500
Fujitsu Server PRIMERGY RX4770 M4, 4 processors / 112 cores / 224 threads, Intel Xeon Platinum 8180 Processor, 2.50 GHz, Windows Server 2012 R2, Microsoft SQL Server 2012 RTM, SAP enhancement package 5 for SAP ERP 6.0 Certification number: 2018004	56480

0   10000   20000   30000   40000   50000   60000

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

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

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

In order to model applications that access in parallel with a different load intensity the "# 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 [MiB/s]   Throughput in megabytes per second
- Transactions [IO/s]   Transaction rate in I/O operations per second
- Latency [ms]   Average response time in ms

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

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

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 MiB/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			
Controller:	1x PRAID CP400i		
Storage media	Category		Drive Name
HDD	SAS HDD (SAS 12Gbps, 10k rpm) [512e]		AL15SEB06EQ *2 *3
	SAS HDD (SAS 12Gbps, 10k rpm) [512n]		AL15SEB030N *2 *3
	SAS HDD (SAS 12Gbps, 15k rpm) [512n]		ST300MP0006 *1 *3
	NL-SAS HDD (SAS 12Gbps, 7.2k rpm) [512n]		ST1000NX0453 *1 *3
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512e]		ST1000NX0313 *1 *3
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512n]		ST2000NX0403 *1 *3
SSD	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 Xeon Gold 5222 (3.80GHz)		
	Storage media	Category	Drive Name
	SSD	2.5 inch PCIe SSD (Write Intensive)	SSDPE21K750GA *2 *4
		2.5 inch PCIe SSD (Mixed Use)	SSDPE2KE016T8 *2 *4
			SSDPE2KE032T8 *2 *4
			SSDPE2KE064T8 *2 *4
	SSD	PCIe SSD (Write Intensive)	SSDPED1K375GA *2 *4
			SSDPED1K750GA *2 *4

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

*1) 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 kiB)	
Measuring tool	Iometer 1.1.0	
Measurement area	Type1	RAW file system is used. The first 10% of the usable LBA area is used for sequential accesses; the next 25% for random accesses.
	Type2	NTFS file system is used. The 32GiB area is secured for the first of the target drive and is used for sequential access and random access.
Total number of Iometer workers	1	
Alignment of Iometer accesses	Aligned to whole multiples of 4096 bytes	

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

Benchmark results

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

Controller

The measurements were made using controllers in the table below.

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

Storage media

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

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

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.

2.5 inch Storage media
HDDs

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

SSDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
400	KPM51MUG400G	SAS 12G	84,469	13,329	13,677	1,056	1,041
800	KPM51MUG800G	SAS 12G	99,728	14,549	18,049	1,056	1,042
1,600	KPM51MUG1T60	SAS 12G	108,428	17,243	19,634	1,057	1,042
400	WUSTR6440ASS204	SAS 12G	83,427	14,459	13,924	1,073	626
800	WUSTR6480ASS204	SAS 12G	94,899	22,414	21,187	1,073	1,008
1,600	WUSTR6416ASS204	SAS 12G	97,107	24,053	22,802	1,073	1,029
3,200	WUSTR6432ASS204	SAS 12G	106,745	23,975	22,793	1,073	1,030
6,400	WUSTR6464ASS204	SAS 12G	111,695	23,911	22,639	1,073	1,030
480	WUSTR1548ASS204	SAS 12G	77,846	11,663	9,904	1,055	554
960	WUSTR1596ASS204	SAS 12G	88,384	18,834	16,636	1,067	965
1,920	WUSTR1519ASS204	SAS 12G	89,397	21,635	21,597	1,073	1,030
3,840	WUSTR1538ASS204	SAS 12G	99,644	23,727	22,831	1,073	1,030
7,680	WUSTR1576ASS204	SAS 12G	106,933	23,688	22,644	1,073	1,030
15,360	WUSTR1515ASS204	SAS 12G	107,687	23,590	22,686	1,073	1,029
240	MZ7KH240HAHQ	SATA 6G	49,159	7,313	7,431	526	486
480	MZ7KH480HAHQ	SATA 6G	50,558	7,774	7,810	526	485
960	MZ7KH960HAJR	SATA 6G	50,647	7,793	7,916	525	485
1,920	MZ7KH1T9HAJR	SATA 6G	50,702	8,040	7,960	526	485
3,840	MZ7KH3T8HALS	SATA 6G	50,766	8,039	7,936	526	485
240	MTFDDAK240TCB	SATA 6G	18,959	3,367	4,516	487	258
480	MTFDDAK480TDC	SATA 6G	24,710	3,799	5,006	507	362
960	MTFDDAK960TDC	SATA 6G	30,152	4,625	5,553	507	440
1,920	MTFDDAK1T9TDC	SATA 6G	37,234	5,606	5,566	507	483
3,840	MTFDDAK3T8TDC	SATA 6G	41,711	6,429	6,133	504	481
7,680	MTFDDAK7T6TDC	SATA 6G	40,683	6,874	6,672	469	482
750	SSDPE21K750GA	PCIe3 x4	214,231	37,611	36,957	2,546	2,295
1,600	SSDPE2KE016T8	PCIe3 x4	135,500	41,066	37,080	3,213	1,917
3,200	SSDPE2KE032T8	PCIe3 x4	136,782	48,210	45,348	3,209	2,800
6,400	SSDPE2KE064T8	PCIe3 x4	192,245	51,767	51,438	3,205	3,048
375	SSDPED1K375GA	PCIe3 x4	212,118	37,121	36,123	2,460	2,197
750	SSDPED1K750GA	PCIe3 x4	209,628	37,592	36,941	2,546	2,296
240	MTFDDAV240TCB	SATA 6G	19,773	3,844	4,968	487	258
480	MTFDDAV480TCB	SATA 6G	22,258	4,935	6,294	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:
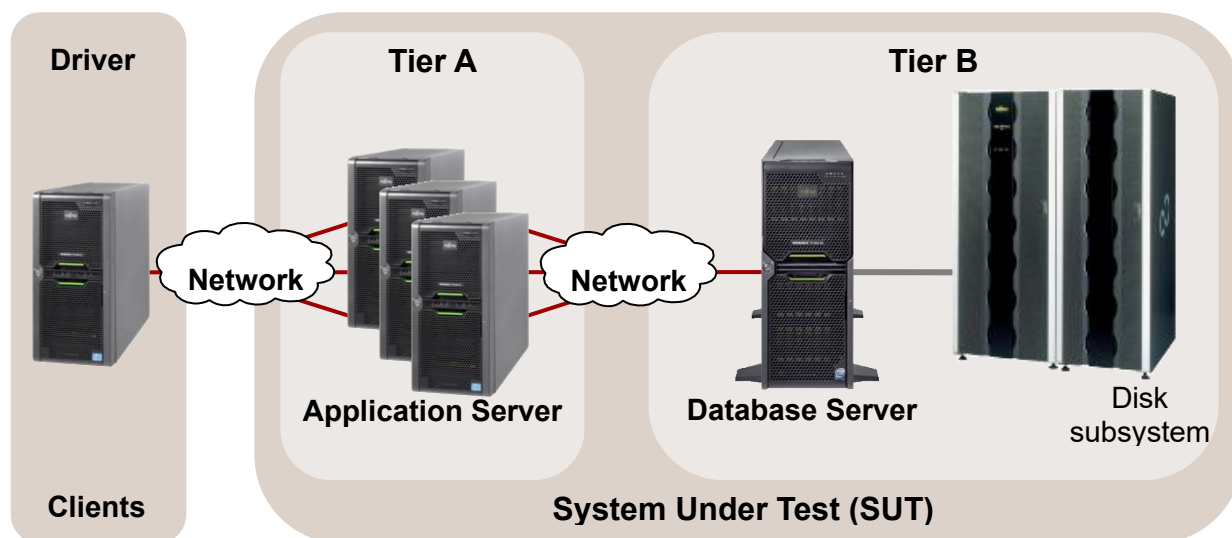
All OLTP-2 results were Calculated based on the configuration of the next following pages of PRIMERGY RX2540 M5

Database Server (Tier B)

Hardware

• Model	PRIMERGY RX4770 M5
• Processor	2nd Generation Intel® Xeon® Scalable Processors Family
• Memory	2 processors:24 × 64 GB (1x64 GB) 2Rx4 DDR4-2933 ECC
	4 processors:48 × 64 GB (1x64 GB) 2Rx4 DDR4-2933 ECC
• Network interface	1 × Dual Port onboard LAN 10 Gb/s
• Disk subsystem	PRIMERGY RX4770 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 S2:14 × 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.
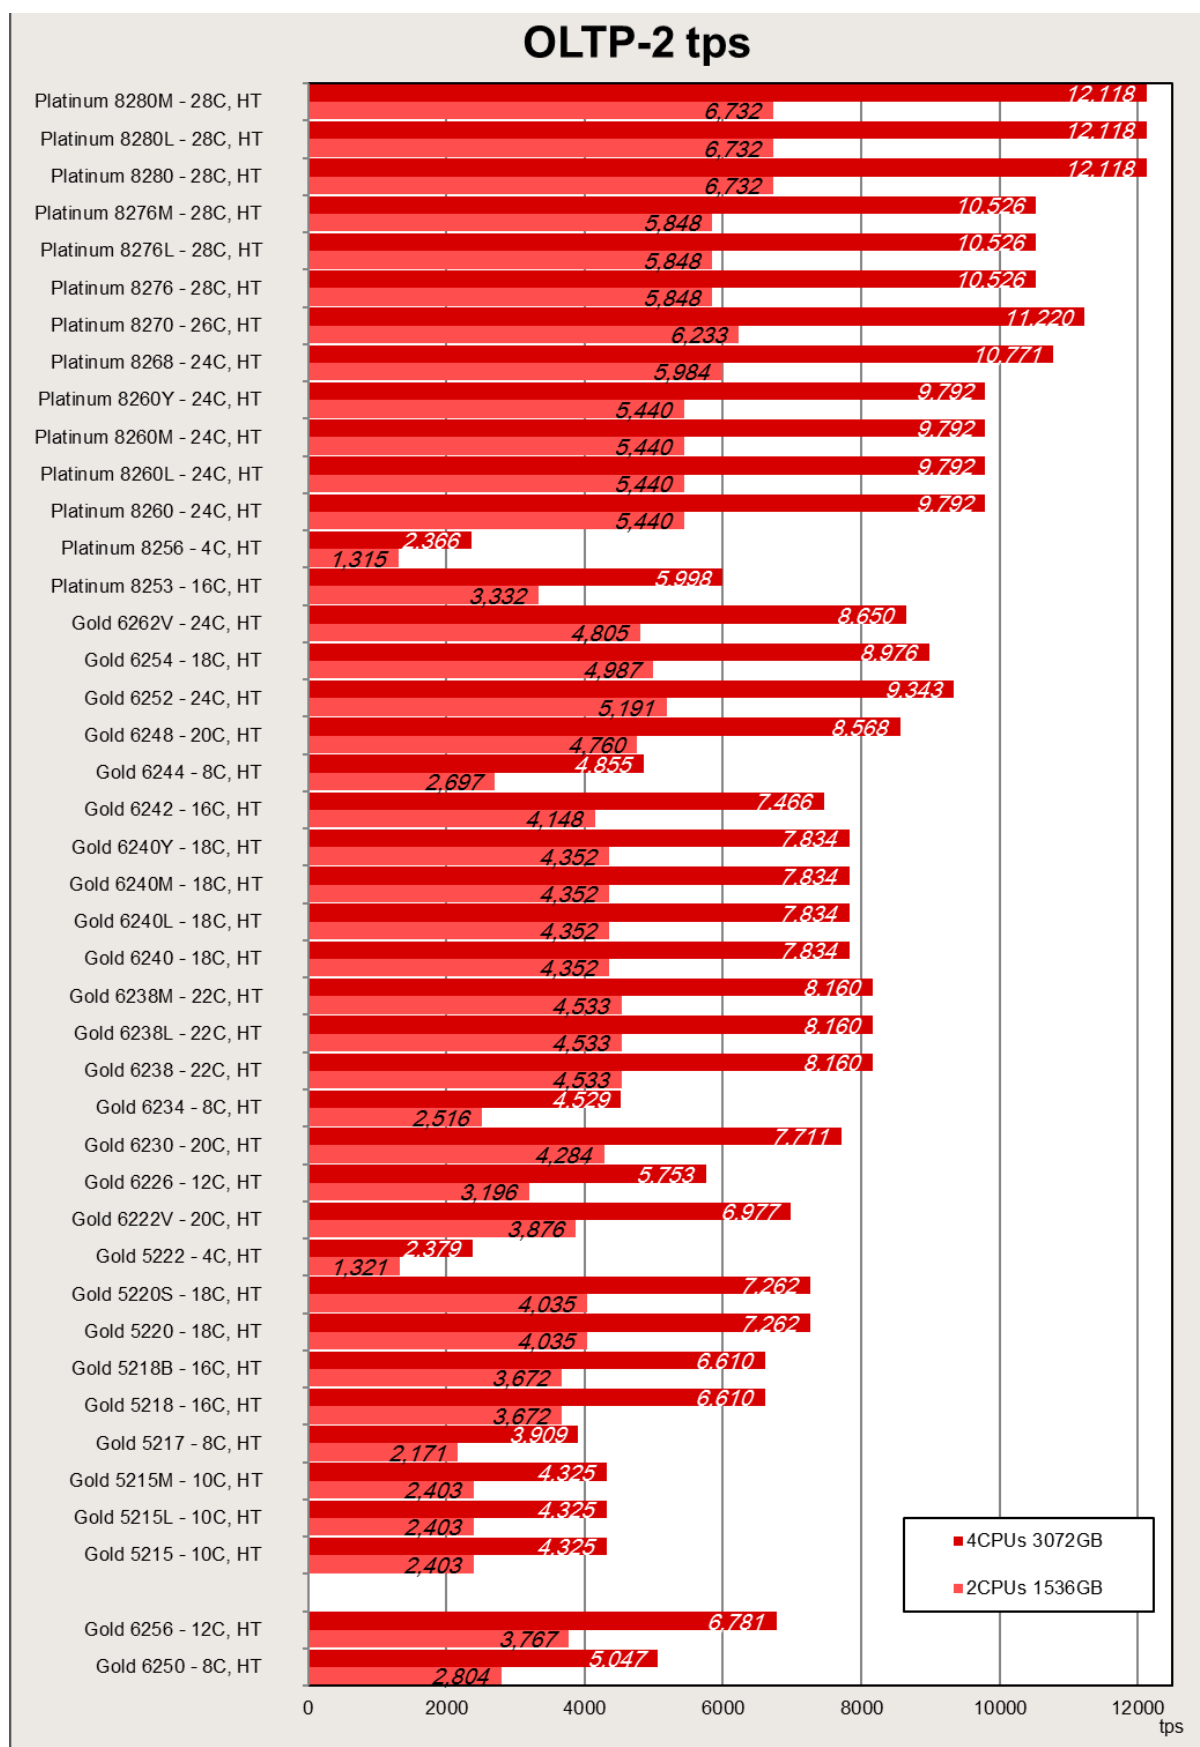
A guideline in the database environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses. This why a configuration with a total memory of 1536 GB was considered for the measurements with two processors and a configuration with a total memory of 768 GB for the measurements with one processor. Both memory configurations have memory access of 2933 MHz.
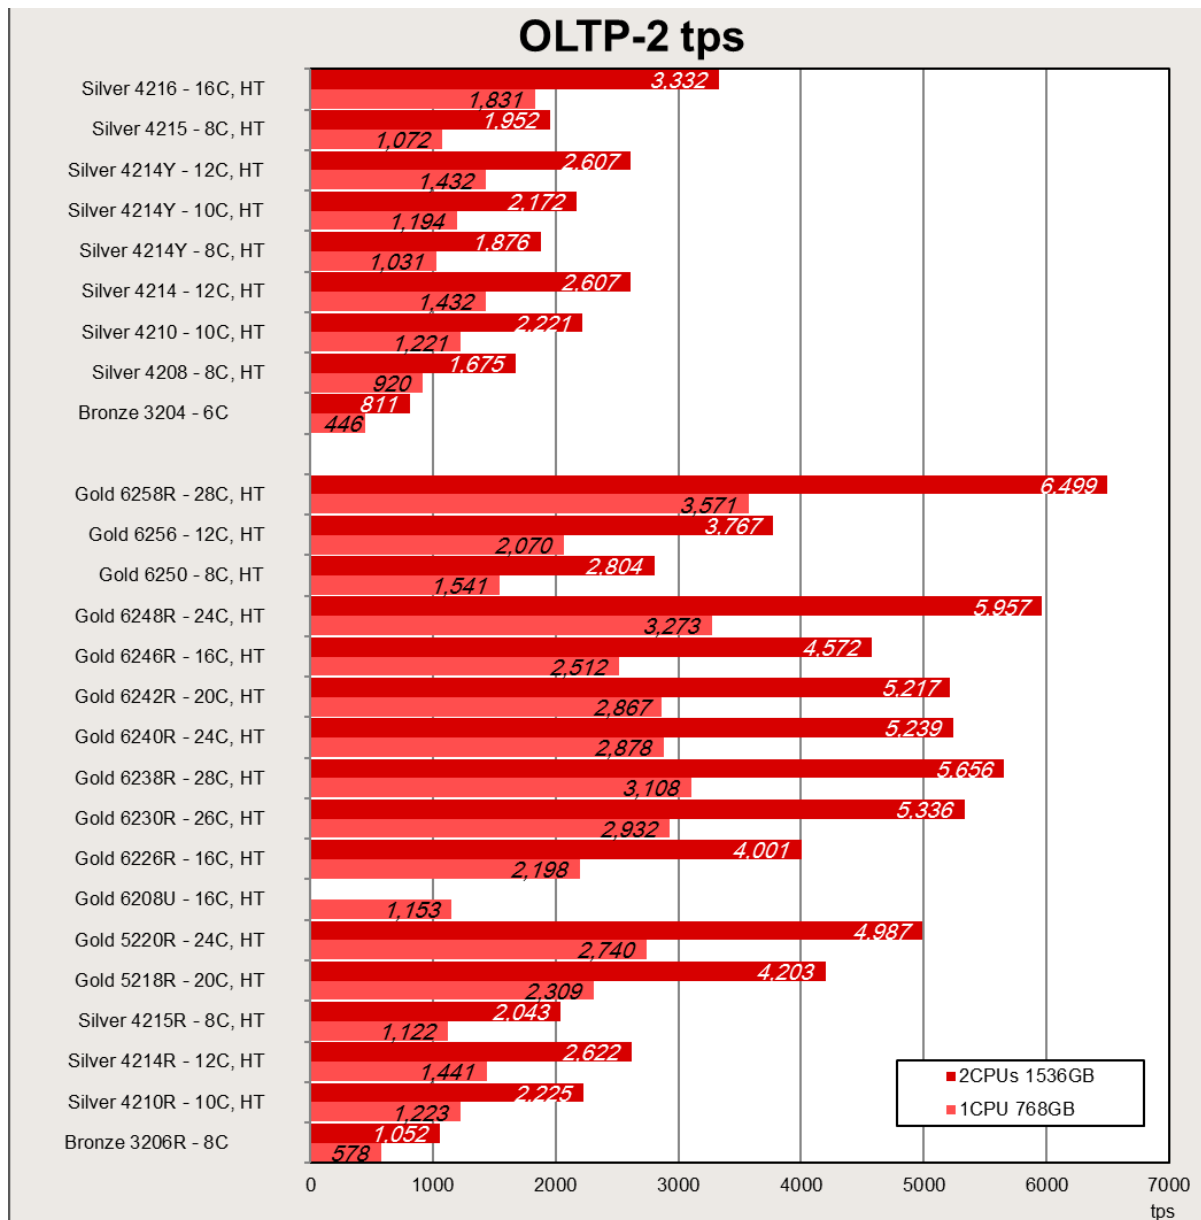
The result with "est." are the estimated values.

Processor	Cores	Threads	4CPU Score		2CPU Score	
April 2019 released						
Xeon Platinum 8280L	28	56	12,118	est.	6,732	est.
Xeon Platinum 8280M	28	56	12,118	est.	6,732	est.
Xeon Platinum 8280	28	56	12,118	est.	6,732	est.
Xeon Platinum 8276L	28	56	10,526	est.	5,848	est.
Xeon Platinum 8276M	28	56	10,526	est.	5,848	est.
Xeon Platinum 8276	28	56	10,526	est.	5,848	est.
Xeon Platinum 8270	26	52	11,220	est.	6,233	est.
Xeon Platinum 8268	24	48	10,771	est.	5,984	est.
Xeon Platinum 8260L	24	48	9,792	est.	5,440	est.
Xeon Platinum 8260M	24	48	9,792	est.	5,440	est.
Xeon Platinum 8260Y	24	48	9,792	est.	5,440	est.
	20	40	8,521	est.	4,734	est.
	16	32	7,387	est.	4,104	est.
Xeon Platinum 8260	24	48	9,792	est.	5,440	est.
Xeon Platinum 8256	4	8	2,366	est.	1,315	est.
Xeon Platinum 8253	16	32	5,998	est.	3,332	est.
Xeon Gold 6262V	24	48	8,650	est.	4,805	est.
Xeon Gold 6254	18	36	8,976	est.	4,987	est.
Xeon Gold 6252	24	48	9,343	est.	5,191	est.
Xeon Gold 6248	20	40	8,568	est.	4,760	est.
Xeon Gold 6246	12	24	5,753	est.	3,196	est.
Xeon Gold 6244	8	16	4,855	est.	2,697	est.
Xeon Gold 6242	16	32	7,466	est.	4,148	est.
Xeon Gold 6240L	18	36	7,834	est.	4,352	est.
Xeon Gold 6240M	18	36	7,834	est.	4,352	est.
	18	36	7,834	est.	4,352	est.
Xeon Gold 6240Y	14	28	6,406	est.	3,559	est.
	8	16	4,004	est.	2,224	est.
Xeon Gold 6240	18	36	7,834	est.	4,352	est.
Xeon Gold 6238L	22	44	8,160	est.	4,533	est.
Xeon Gold 6238M	22	44	8,160	est.	4,533	est.
Xeon Gold 6238	22	44	8,160	est.	4,533	est.
Xeon Gold 6234	8	16	4,529	est.	2,516	est.

Xeon Gold 6230	20	40	7,711	est.	4,284	est.
Xeon Gold 6226	12	24	5,753	est.	3,196	est.
Xeon Gold 6222V	20	40	6,977	est.	3,876	est.
Xeon Gold 5222	4	8	2,379	est.	1,321	est.
Xeon Gold 5220S	18	36	7,262	est.	4,035	est.
Xeon Gold 5220	18	36	7,262	est.	4,035	est.
Xeon Gold 5218B	16	32	6,610	est.	3,672	est.
Xeon Gold 5218	16	32	6,610	est.	3,672	est.
Xeon Gold 5217	8	16	3,909	est.	2,171	est.
Xeon Gold 5215L	10	20	4,325	est.	2,403	est.
Xeon Gold 5215M	10	20	4,325	est.	2,403	est.
Xeon Gold 5215	10	20	4,325	est.	2,403	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.

OLTP-2 tps

Processor	2CPUs 1536GB	4CPUs 3072GB
Platinum 8280M - 28C, HT	6,732	12,118
Platinum 8280L - 28C, HT	6,732	12,118
Platinum 8280 - 28C, HT	6,732	12,118
Platinum 8276M - 28C, HT	5,848	10,526
Platinum 8276L - 28C, HT	5,848	10,526
Platinum 8276 - 28C, HT	5,848	10,526
Platinum 8270 - 26C, HT	6,233	11,220
Platinum 8268 - 24C, HT	5,984	10,771
Platinum 8260Y - 24C, HT	5,440	9,792
Platinum 8260M - 24C, HT	5,440	9,792
Platinum 8260L - 24C, HT	5,440	9,792
Platinum 8260 - 24C, HT	5,440	9,792
Platinum 8256 - 4C, HT	1,315	2,366
Platinum 8253 - 16C, HT	3,332	5,998
Gold 6262V - 24C, HT	4,805	8,650
Gold 6254 - 18C, HT	4,987	8,976
Gold 6252 - 24C, HT	5,191	9,343
Gold 6248 - 20C, HT	4,760	8,568
Gold 6244 - 8C, HT	2,697	4,855
Gold 6242 - 16C, HT	4,148	7,466
Gold 6240Y - 18C, HT	4,352	7,834
Gold 6240M - 18C, HT	4,352	7,834
Gold 6240L - 18C, HT	4,352	7,834
Gold 6240 - 18C, HT	4,352	7,834
Gold 6238M - 22C, HT	4,533	8,160
Gold 6238L - 22C, HT	4,533	8,160
Gold 6238 - 22C, HT	4,533	8,160
Gold 6234 - 8C, HT	2,516	4,529
Gold 6230 - 20C, HT	4,284	7,711
Gold 6226 - 12C, HT	3,196	5,753
Gold 6222V - 20C, HT	3,876	6,977
Gold 5222 - 4C, HT	1,321	2,379
Gold 5220S - 18C, HT	4,035	7,262
Gold 5220 - 18C, HT	4,035	7,262
Gold 5218B - 16C, HT	3,672	6,610
Gold 5218 - 16C, HT	3,672	6,610
Gold 5217 - 8C, HT	2,171	3,909
Gold 5215M - 10C, HT	2,403	4,325
Gold 5215L - 10C, HT	2,403	4,325
Gold 5215 - 10C, HT	2,403	4,325
Gold 6256 - 12C, HT	3,767	6,781
Gold 6250 - 8C, HT	2,804	5,047

Legend:
■ 4CPUs 3072GB
■ 2CPUs 1536GB

OLTP-2 tps

CPU	2CPUs 1536GB	1CPU 768GB
Silver 4216 - 16C, HT	3,332	1,831
Silver 4215 - 8C, HT	1,952	1,072
Silver 4214Y - 12C, HT	2,607	1,432
Silver 4214Y - 10C, HT	2,172	1,194
Silver 4214Y - 8C, HT	1,876	1,031
Silver 4214 - 12C, HT	2,607	1,432
Silver 4210 - 10C, HT	2,221	1,221
Silver 4208 - 8C, HT	1,675	920
Bronze 3204 - 6C	811	446
Gold 6258R - 28C, HT	6,499	3,571
Gold 6256 - 12C, HT	3,767	2,070
Gold 6250 - 8C, HT	2,804	1,541
Gold 6248R - 24C, HT	5,957	3,273
Gold 6246R - 16C, HT	4,572	2,512
Gold 6242R - 20C, HT	5,217	2,867
Gold 6240R - 24C, HT	5,239	2,878
Gold 6238R - 28C, HT	5,656	3,108
Gold 6230R - 26C, HT	5,336	2,932
Gold 6226R - 16C, HT	4,001	2,198
Gold 6208U - 16C, HT	1,153	
Gold 5220R - 24C, HT	4,987	2,740
Gold 5218R - 20C, HT	4,203	2,309
Silver 4215R - 8C, HT	2,043	1,122
Silver 4214R - 12C, HT	2,622	1,441
Silver 4210R - 10C, HT	2,225	1,223
Bronze 3206R - 8C	1,052	578

(x-axis: tps, 0 to 7000)

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%.
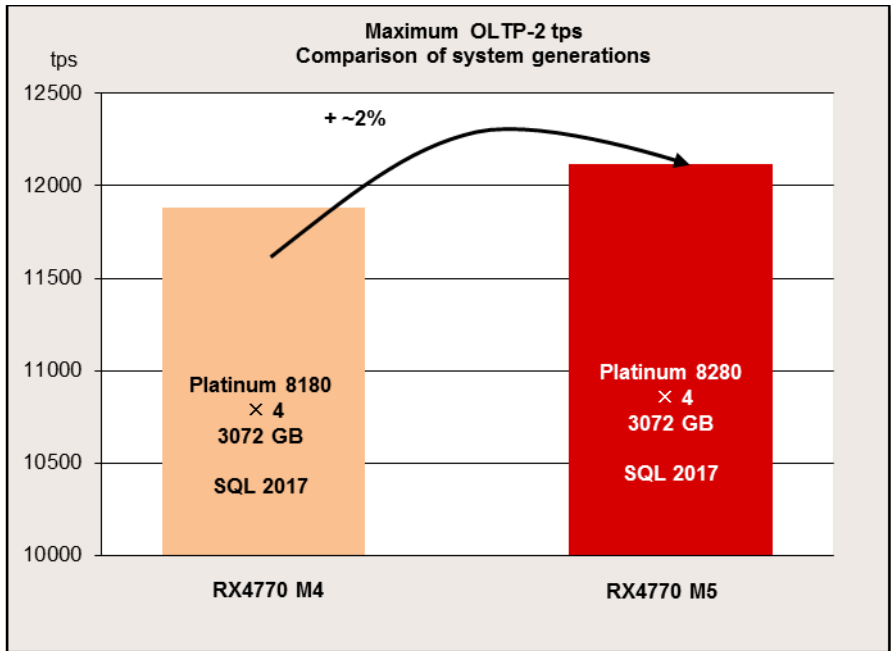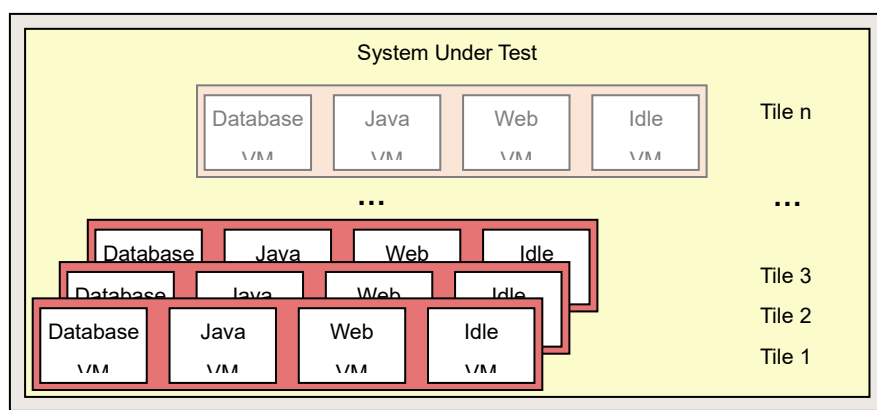
vServCon

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:
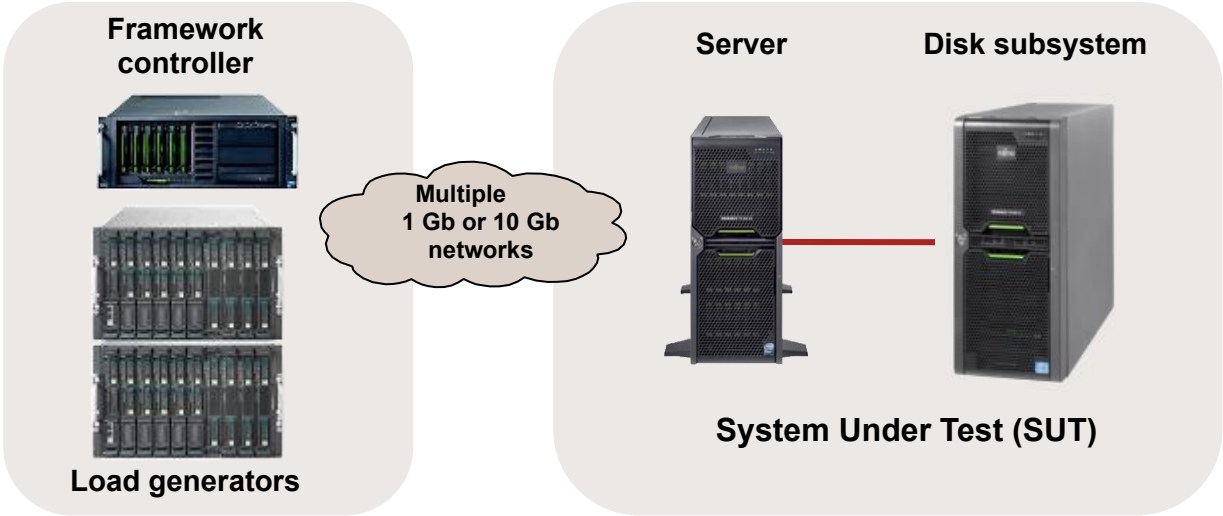
System Under Test (SUT)	
Hardware	
• Processor	4 × 2nd Generation Intel Xeon Scalable Processors Family
• Memory	48 × 32 GB (1x32 GB) 2Rx4 DDR4-2933 R ECC
• Network interface	2 × Intel Ethernet Controller X710 for 10GbE SFP+
• Disk subsystem	1 × dual-channel FC controller Emulex LPe160021 LINUX/LIO based flash storage system
Software	
• Operating system	VMware ESXi 6.7 EP06 Build 11675023

Load generator (incl. Framework controller)	
Hardware (Shared)	
• Enclosure	4 × 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".

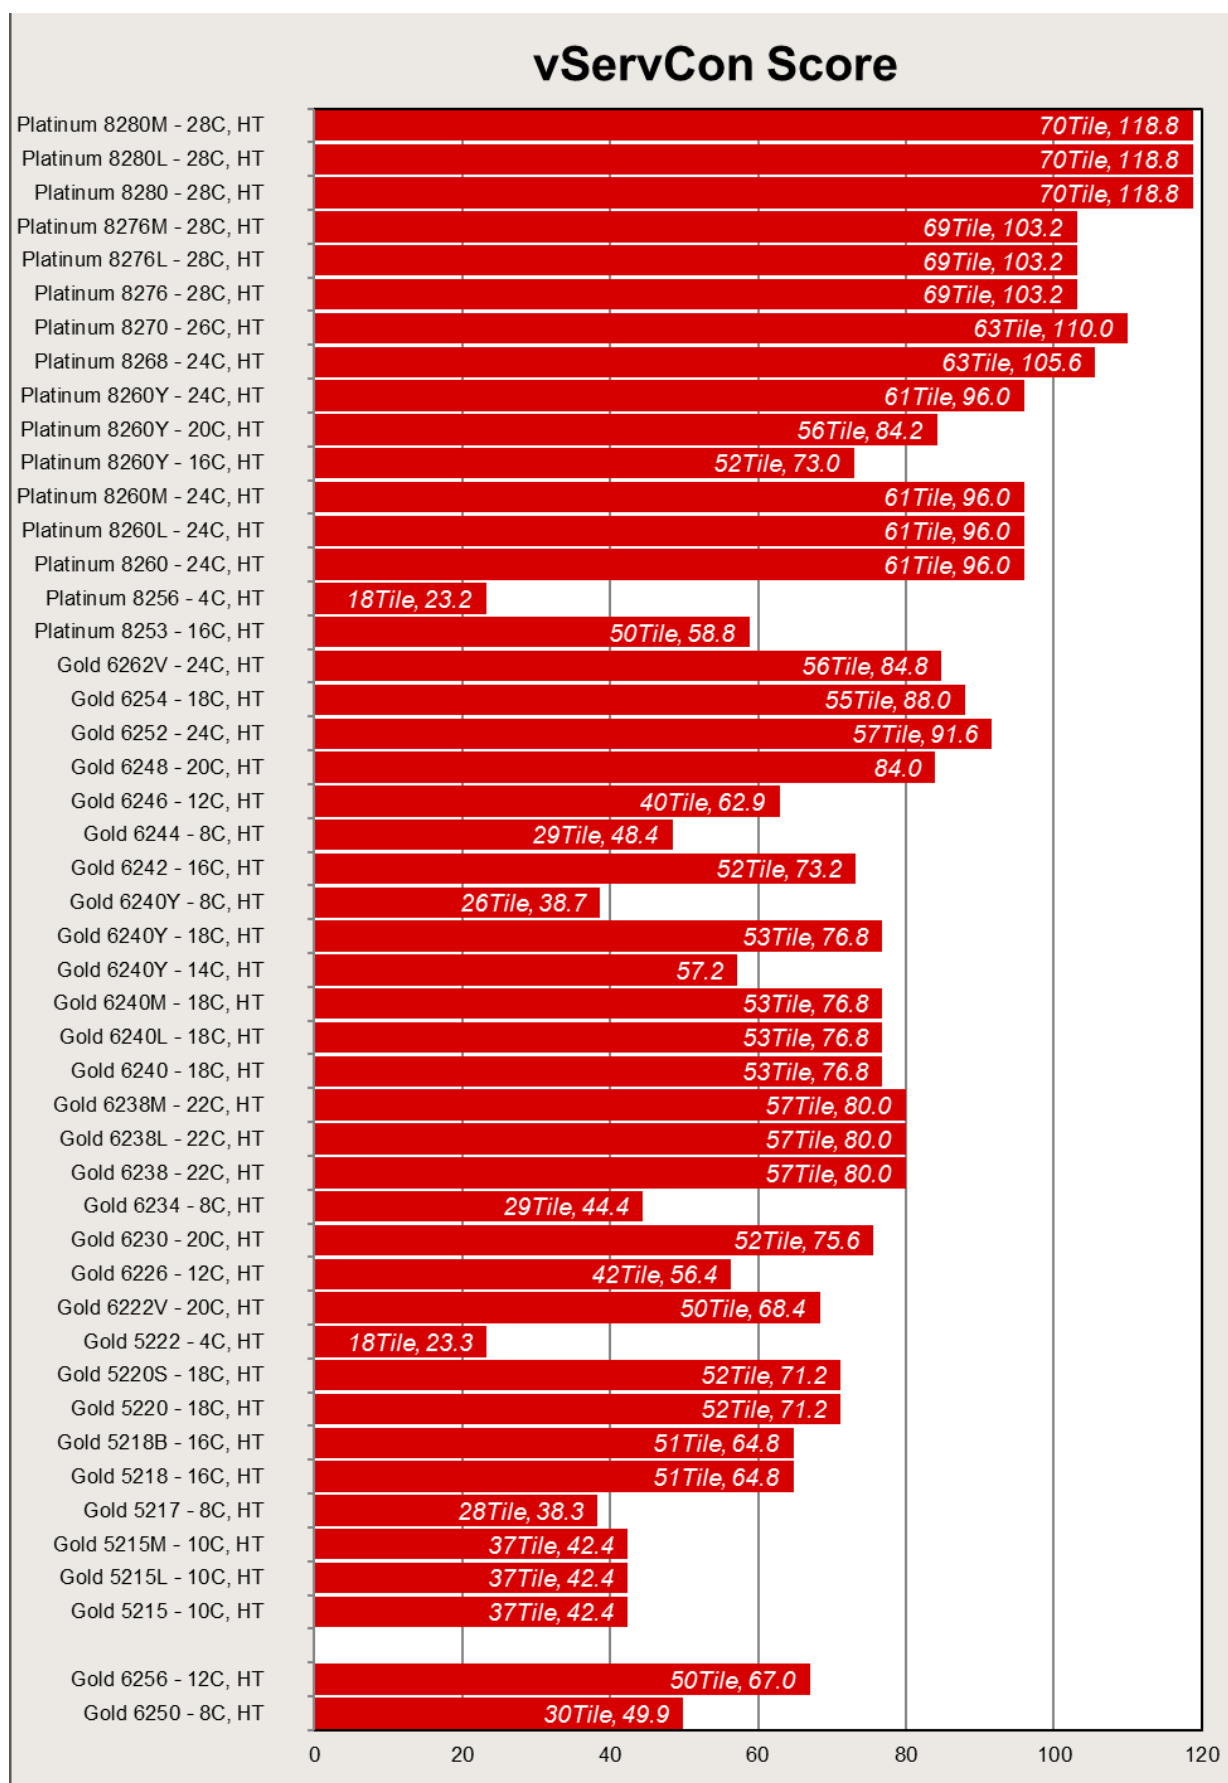The available processors of these systems with their results can be seen in the following table.

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

Processor	Cores	Threads	#Tiles		Score	
April 2019 released						
Xeon Platinum 8280L	28	56	70	est.	118.8	est.
Xeon Platinum 8280M	28	56	70	est.	118.8	est.
Xeon Platinum 8280	28	56	**70**		**118.8**	
Xeon Platinum 8276L	28	56	69	est.	103.2	est.
Xeon Platinum 8276M	28	56	69	est.	103.2	est.
Xeon Platinum 8276	28	56	69	est.	103.2	est.
Xeon Platinum 8270	26	52	63	est.	110	est.
Xeon Platinum 8268	24	48	63	est.	105.6	est.
Xeon Platinum 8260L	24	48	61	est.	96	est.
Xeon Platinum 8260M	24	48	61	est.	96	est.
Xeon Platinum 8260Y	24	48	61	est.	96	est.
	20	40	56	est.	84.2	est.
	16	32	52	est.	73	est.
Xeon Platinum 8260	24	48	61	est.	96	est.
Xeon Platinum 8256	4	8	18	est.	23.2	est.
Xeon Platinum 8253	16	32	50	est.	58.8	est.
Xeon Gold 6262V	24	48	56	est.	84.8	est.
Xeon Gold 6254	18	36	55	est.	88	est.
Xeon Gold 6252	24	48	57	est.	91.6	est.
Xeon Gold 6248	20	40	53	est.	84	est.
Xeon Gold 6246	12	24	40	est.	**62.9**	
Xeon Gold 6244	8	16	**29**		**48.4**	
Xeon Gold 6242	16	32	52	est.	73.2	est.
Xeon Gold 6240L	18	36	53	est.	76.8	est.
Xeon Gold 6240M	18	36	53	est.	76.8	est.
Xeon Gold 6240Y	18	36	53	est.	76.8	est.
	14	28	46	est.	57.2	est.

Processor					
	8	16	26	est.	38.7 est.
Xeon Gold 6240	18	36	53	est.	76.8 est.
Xeon Gold 6238L	22	44	57	est.	80 est.
Xeon Gold 6238M	22	44	57	est.	80 est.
Xeon Gold 6238	22	44	57	est.	80 est.
Xeon Gold 6234	8	16	29	est.	44.4 est.
Xeon Gold 6230	20	40	52	est.	75.6 est.
Xeon Gold 6226	12	24	42	est.	56.4 est.
Xeon Gold 6222V	20	40	50	est.	68.4 est.
Xeon Gold 5222	4	8	18	est.	23.3 est.
Xeon Gold 5220S	18	36	52	est.	71.2 est.
Xeon Gold 5220	18	36	52	est.	71.2 est.
Xeon Gold 5218B	16	32	51	est.	64.8 est.
Xeon Gold 5218	16	32	51	est.	64.8 est.
Xeon Gold 5217	8	16	28	est.	38.3 est.
Xeon Gold 5215L	10	20	37	est.	42.4 est.
Xeon Gold 5215M	10	20	37	est.	42.4 est.
Xeon Gold 5215	10	20	37	est.	42.4 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

Processor	Score
Platinum 8280M - 28C, HT	70Tile, 118.8
Platinum 8280L - 28C, HT	70Tile, 118.8
Platinum 8280 - 28C, HT	70Tile, 118.8
Platinum 8276M - 28C, HT	69Tile, 103.2
Platinum 8276L - 28C, HT	69Tile, 103.2
Platinum 8276 - 28C, HT	69Tile, 103.2
Platinum 8270 - 26C, HT	63Tile, 110.0
Platinum 8268 - 24C, HT	63Tile, 105.6
Platinum 8260Y - 24C, HT	61Tile, 96.0
Platinum 8260Y - 20C, HT	56Tile, 84.2
Platinum 8260Y - 16C, HT	52Tile, 73.0
Platinum 8260M - 24C, HT	61Tile, 96.0
Platinum 8260L - 24C, HT	61Tile, 96.0
Platinum 8260 - 24C, HT	61Tile, 96.0
Platinum 8256 - 4C, HT	18Tile, 23.2
Platinum 8253 - 16C, HT	50Tile, 58.8
Gold 6262V - 24C, HT	56Tile, 84.8
Gold 6254 - 18C, HT	55Tile, 88.0
Gold 6252 - 24C, HT	57Tile, 91.6
Gold 6248 - 20C, HT	84.0
Gold 6246 - 12C, HT	40Tile, 62.9
Gold 6244 - 8C, HT	29Tile, 48.4
Gold 6242 - 16C, HT	52Tile, 73.2
Gold 6240Y - 8C, HT	26Tile, 38.7
Gold 6240Y - 18C, HT	53Tile, 76.8
Gold 6240Y - 14C, HT	57.2
Gold 6240M - 18C, HT	53Tile, 76.8
Gold 6240L - 18C, HT	53Tile, 76.8
Gold 6240 - 18C, HT	53Tile, 76.8
Gold 6238M - 22C, HT	57Tile, 80.0
Gold 6238L - 22C, HT	57Tile, 80.0
Gold 6238 - 22C, HT	57Tile, 80.0
Gold 6234 - 8C, HT	29Tile, 44.4
Gold 6230 - 20C, HT	52Tile, 75.6
Gold 6226 - 12C, HT	42Tile, 56.4
Gold 6222V - 20C, HT	50Tile, 68.4
Gold 5222 - 4C, HT	18Tile, 23.3
Gold 5220S - 18C, HT	52Tile, 71.2
Gold 5220 - 18C, HT	52Tile, 71.2
Gold 5218B - 16C, HT	51Tile, 64.8
Gold 5218 - 16C, HT	51Tile, 64.8
Gold 5217 - 8C, HT	28Tile, 38.3
Gold 5215M - 10C, HT	37Tile, 42.4
Gold 5215L - 10C, HT	37Tile, 42.4
Gold 5215 - 10C, HT	37Tile, 42.4
Gold 6256 - 12C, HT	50Tile, 67.0
Gold 6250 - 8C, HT	30Tile, 49.9

(x-axis: 0, 20, 40, 60, 80, 100, 120)

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

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

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.
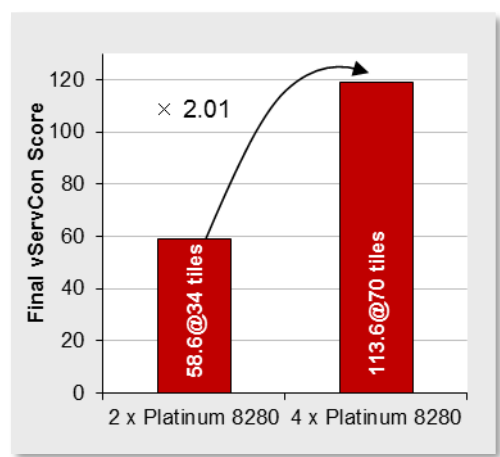
Until now, we have looked at the virtualization performance of a fully configured system. However, with a server with four sockets, the question also arises as to how good performance scaling is from two to four processors. The better the scaling, the lower the overhead usually caused by the shared use 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 2.01. When operated with four processors, the system thus achieves twice the performance as with two processors, 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 (8-Core) processors.
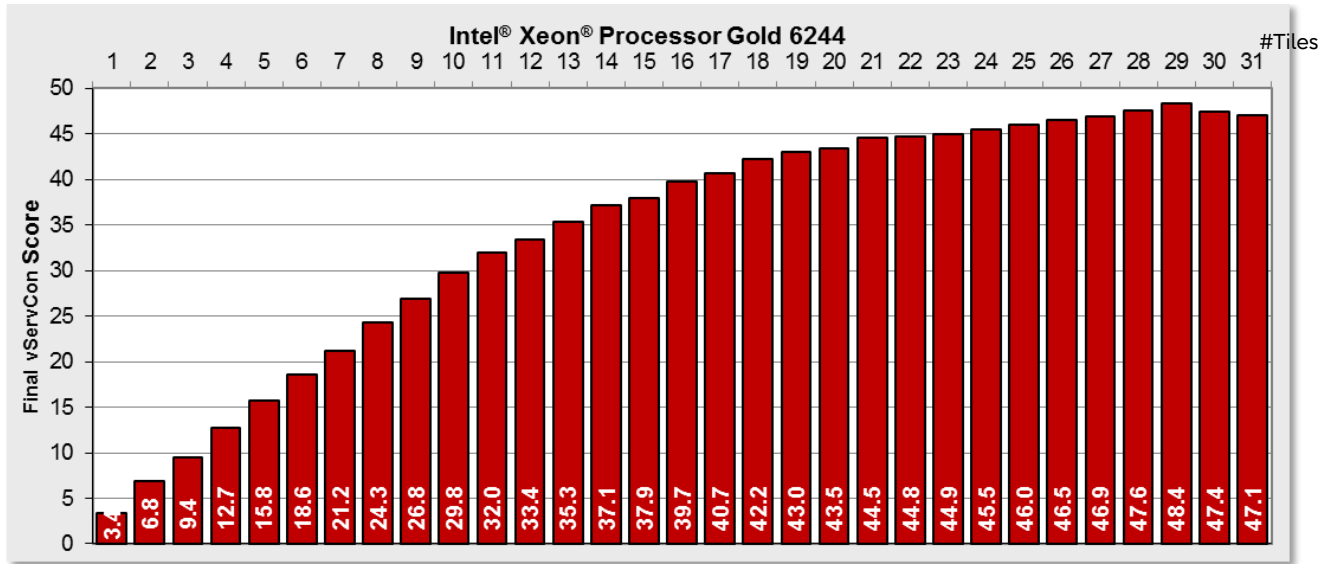
The previous diagram examined the total performance of all application VMs of a host. However, studying the performance from an individual application VM viewpoint is also interesting. This information is in the previous diagram. For example, the total optimum is reached in the above Xeon Gold 6244 situation with 87 application VMs (29 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 3.4 to 48.4/29=1.7, i.e. to 50%. 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

Virtualization relevant improvements

Score at optimum Tile count

Year CPU Freq. #Cores	Score	Multiplier

vServCon Score

- 2008 X7460 2.67 GHz 6C — 6.66
- 2010 X7560 2.27 GHz 8C — 17.80 (× 2.67)
- 2011 E5-2690 2.4 GHz 10C — 23.9 (× 1.34)
- 2014 E7-4890 v2 2.8 GHz 15C — 44.9 (× 1.88)
- 2015 E5-8890 v3 2.5 GHz 18C — 58.2 (× 1.30)
- 2016 E7-8890 v4 2.2 GHz 24C — 75.6 (× 1.30)
- 2017 Platinum 8180 2.5 GHz 28C — 113.6 (× 1.50)
- 2019 Platinum 8280 2.7 GHz 28C — 118.8 (× 1.05)

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".
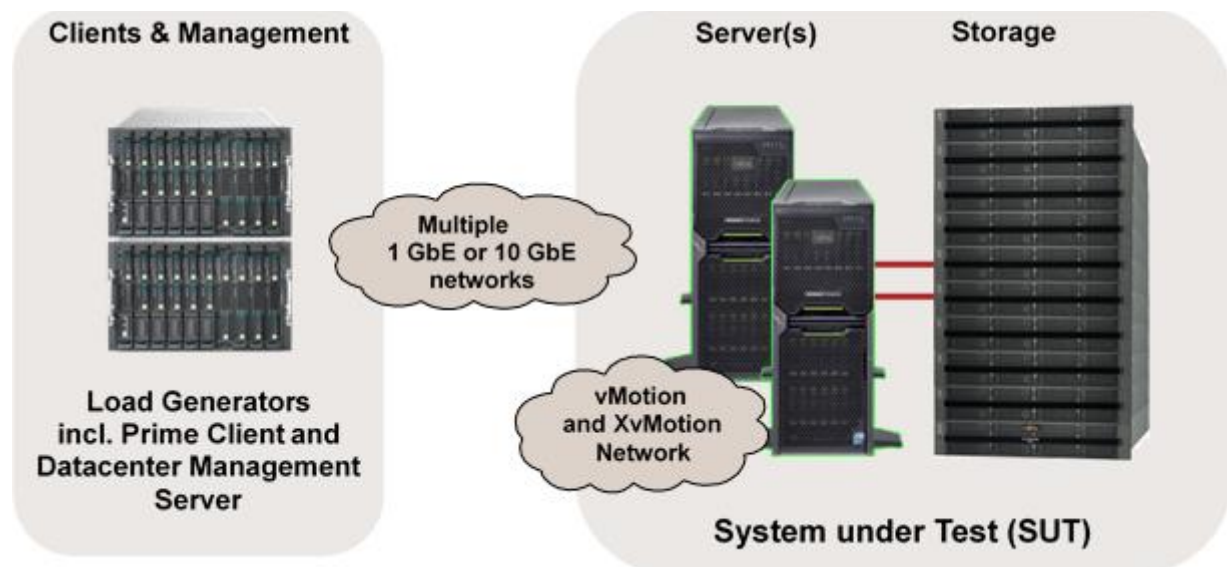
In the case of the two test types "Performance with Server Power" and "Performance with Server and Storage Power", a so-called "Server PPKW Score" and "Server and Storage PPKW Score" are determined, which are the performance scores divided by the average power consumption in kilowatts (PPKW = performance 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 RX4770 M5
• Processor	4 × Intel Xeon Platinum 8280
• Memory	768 GB: 48 × 32 GB (1x32 GB) 2Rx4 DDR4-2933 R ECC
• Network interface	2 × Intel Ethernet Controller X710 for 10GbE SFP+
	1 × Intel Ethernet Connection X722 for 1GbE
• Disk subsystem	2 × Dual port PFC EP LPe31002
	6 × PRIMERGY RX2540 M5 configured as Fibre Channel target
	4 × PRIMERGY RX2540 M5:
	2 × Micron MTFDDAK480 TDC SATA-SSD (480 GB)
	3 × Intel P4800X 750GB PCIe SSD (750 GB)
	1 × Intel P4600 4TB PCIe SSD (4 TB)
	SATA-SSD1, 2(RAID 1)
	1 × PRIMERGY RX2540 M5:
	1 × Micron MTFDDAK960 TDC SATA-SSD (960 GB)
	2 × Intel P4800X 750GB PCIe SSD (750 GB)
	2 × Intel P4600 2TB PCIe SSD (2 TB)
	1 × PRIMERGY RX2540 M5:
	1 × Micron MTFDDAK960 TDC SATA-SSD (960 GB)
	4 × Intel P4800X 750GB PCIe SSD (750 GB)
Software	
• BIOS	V5.0.0.14 R1.8.0 for D3753-C1x

• BIOS settings	See details
• Operating system	VMware ESXi 6.7 EP 08, Build 13473784
• Operating system settings	ESX settings: see details

Details

See disclosure	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2019-08-27-Fujitsu-RX4770M5.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	5 × PRIMERGY RX2530 M2
• Processor	3 × Fujitsu PRIMERGY RX2530M2 2 × Intel Xeon E5-2699 v4 2 × Fujitsu PRIMERGY RX2530M2 2 × Intel Xeon E5-2699A 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 (August 27   2019)

On August 27, 2019 Fujitsu achieved with a PRIMERGY RX4770 M5 with Xeon Platinum 8280 processors and VMware ESXi 6.7 EP 08 a VMmark V3 score of "17.23@18 tiles" in a system configuration with a total of 4 × 56 processor cores and when using two identical servers in the "System under Test" (SUT). With this result the PRIMERGY RX4770 M5 is in the official VMmark V3 "Performance Only" ranking the most powerful four-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 August 27, 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 RX4770 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.

VMmark is a product of VMware, Inc.

Literature

PRIMERGY Servers

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

PRIMERGY RX4770 M5

This White Paper:

https://docs.ts.fujitsu.com/dl.aspx?id=71064ca4-408d-4f40-b08d-49741dbc086e

https://docs.ts.fujitsu.com/dl.aspx?id=fa5b3124-e575-406c-b4ba-a8ecd10ab6a2

Data sheet

https://docs.ts.fujitsu.com/dl.aspx?id=c73aa330-ea96-4e4a-bfcc-5ac79a09ed7d

PRIMERGY Performance

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

SPEC CPU2017

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

Benchmark Overview SPECcpu2017

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

STREAM

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

LINPACK

The LINPACK Benchmark: Past, Present, and Future

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

TOP500

https://www.top500.org/

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

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

Intel Math Kernel Library – LINPACK Download

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

SPECpower_ssj2008

https://www.spec.org/power_ssj2008

Benchmark Overview SPECpower_ssj2008

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

SPECjbb2015

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

SAP SD

https://www.sap.com/about/benchmark.html

Benchmark overview SAP SD

https://docs.ts.fujitsu.com/dl.aspx?id=0a1e69a6-e366-4fd1-a1a6-0dd93148ea10

OLTP-2

Benchmark Overview OLTP-2

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

vServCon

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

VMmark V3

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

Document change history

Version	Date	Description
1.5	2023-10-03	Updated: • New Visual Identity format.
1.4	2021-07-28	Updated: • Contact information and URLs Updated to the latest one • Minor correction
1.3	2020-05-29	Updated: • Technical data, STREAM, LINPACK Fixed typo in processor specifications
1.2	2020-04-24	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
1.1	2019-10-04	New: • SAP SD Certification number 2019033 • Disk I/O: Performance of storage media Results for 2.5" storage media • VMmark V3 "Performance Only" measurement with Intel Xeon Platinum 8280 • STREAM, LINPACK Measured with 2nd Generation Intel Xeon Processor Scalable Family Updated: • SPECcpu2017 Measured additionally with 2nd Generation Intel Xeon Processor Scalable Family

Document change history

Version	Date	Description
1.0	2019-04-30	New: • Technical data • SPECcpu2017 Measurements with 2nd Generation Intel Xeon Processor Scalable Family • SPECpower_ssj2008 Measurements with Intel Xeon Platinum 8276L and Intel Xeon Platinum 8280L • SPECjbb2015 Measurement with Intel Xeon Platinum 8280L • OLTP-2 Calculated with 2nd Generation Intel Xeon Processor Scalable Family • vServCon Measurements with 2nd Generation Intel Xeon Processor Scalable Family

Contact

Fujitsu

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

PRIMERGY Performance and Benchmarks

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