White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY RX2540 M4

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY RX2540 M4.

The PRIMERGY RX2540 M4 performance data are compared with the data of other PRIMERGY models and discussed. In addition to the benchmark results, an explanation has been included for each benchmark and for the benchmark environment.

Version
1.3
2018/11/27

Contents

Document history ... 2
Technical data ... 4
SPECcpu2006 ... 8
SPECpower_ssj2008 ... 15
SPECjbb2015 ... 19
SAP SD ... 22
OLTP-2 ... 25
TPC-E ... 30
vServCon ... 34
VMmark V3 ... 43
STREAM ... 48
LINPACK ... 53
Literature ... 57
Contact ... 59

Document history

Version 1.0 (2017/10/31)

New:

- Technical data
- SPECcpu2006
  Measurements with Intel® Xeon® Processor Scalable Family
- SAP SD
  Certification number 2017014
- OLTP-2
  Measurements with Intel® Xeon® Processor Scalable Family
- vServCon
  Measurements with Intel® Xeon® Processor Scalable Family
- STREAM
  Measurements with Intel® Xeon® Processor Scalable Family
- LINPACK
  Measurements with Intel® Xeon® Processor Scalable Family

Version 1.1 (2018/02/08)

New:

- SPECpower_ssj2008
  Measurement with Intel® Xeon® Platinum 8176M
- VMmark V3
  "Performance Only" measurement with Intel® Xeon® Platinum 8180
  "Performance with Server Power" measurement Intel® Xeon® Platinum 8180
  "Performance with Server and Storage Power" measurement Intel® Xeon® Platinum 8180

Updated:

- SPECcpu2006
  Additional measurements with Intel® Xeon® Processor Scalable Family
- vServCon
  Additional measurements with Intel® Xeon® Processor Scalable Family

Version 1.2 (2018/04/10)

New:

- TPC-E
  Measurement with Intel® Xeon® Platinum 8180

Updated:

- vServCon
  Additional measurements with Intel® Xeon® Processor Scalable Family

Version 1.3 (2018/11/27)

New:

- SPECjbb2015
  Measurement with Intel® Xeon® Platinum 8180

Technical data

PRIMERGY RX2540 M4

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

Model	PRIMERGY RX2540 M4
Model versions	PY RX2540 M4 4x 3.5' expandable PY RX2540 M4 12x 3.5' PY RX2540 M4 8x 2.5' expandable PY RX2540 M4 24x 2.5'
Form factor	Rack server
Chipset	Intel® C620
Number of sockets	2
Number of processors orderable	1 or 2
Processor type	Intel® Xeon® Processor Scalable Family
Number of memory slots	24 (12 per processor)
Maximum memory configuration	3,072 GB
Onboard HDD controller	Controller with RAID 0, RAID 1 or RAID 10 for up to 8 SATA HDDs
PCI slots	PCI-Express 3.0 x8 × 3 PCI-Express 3.0 x16 × 3
Max. number of internal hard disks	PY RX2540 M4 4x 3.5' expandable: 3.5" × 8 + 2.5" × 4 PY RX2540 M4 12x 3.5': 3.5" × 12 + 2.5" × 4 PY RX2540 M4 8x 2.5' expandable: 2.5" × 28 PY RX2540 M4 24x 2.5': 2.5" × 28

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]
Xeon Bronze 3104	6	6	8.3	9.6	1.7	n/a	2133	85
Xeon Bronze 3106	8	8	11.0	9.6	1.7	n/a	2133	85
Xeon Silver 4108	8	16	11.0	9.6	1.8	3.0	2400	85
Xeon Silver 4110	8	16	11.0	9.6	2.1	3.0	2400	85
Xeon Silver 4114	10	20	13.8	9.6	2.2	3.0	2400	85
Xeon Silver 4116	12	24	16.5	9.6	2.1	3.0	2400	85
Xeon Gold 5115	10	20	13.8	10.4	2.4	3.2	2400	85
Xeon Gold 5118	12	24	16.5	10.4	2.3	3.2	2400	105
Xeon Gold 5120	14	28	19.3	10.4	2.2	3.2	2400	105
Xeon Gold 6130	16	32	22.0	10.4	2.1	3.7	2666	125
Xeon Gold 6140	18	36	24.8	10.4	2.3	3.7	2666	140
Xeon Gold 6138	20	40	27.5	10.4	2.0	3.7	2666	125
Xeon Gold 6148	20	40	27.5	10.4	2.4	3.7	2666	150
Xeon Gold 6152	22	44	30.3	10.4	2.1	3.7	2666	140
Xeon Platinum 8153	16	32	22.0	10.4	2.0	2.8	2666	125
Xeon Platinum 8160	24	48	33.0	10.4	2.1	3.7	2666	150
Xeon Platinum 8164	26	52	35.8	10.4	2.0	3.7	2666	150
Xeon Platinum 8170	26	52	35.8	10.4	2.1	3.7	2666	165
Xeon Platinum 8176	28	56	38.5	10.4	2.1	3.8	2666	165
Xeon Platinum 8180	28	56	38.5	10.4	2.5	3.8	2666	205
Xeon Silver 4112	4	8	8.3	9.6	2.6	3.0	2400	85
Xeon Gold 5122	4	8	16.5	10.4	3.6	3.7	2666	105
Xeon Gold 6128	6	12	19.3	10.4	3.4	3.7	2666	115
Xeon Gold 6134	8	16	24.8	10.4	3.2	3.7	2666	130
Xeon Gold 6144	8	16	24.8	10.4	3.5	4.2	2666	150
Xeon Gold 6126	12	24	19.3	10.4	2.6	3.7	2666	125
Xeon Gold 6136	12	24	24.8	10.4	3.0	3.7	2666	150
Xeon Gold 6146	12	24	24.8	10.4	3.2	4.2	2666	165
Xeon Gold 6132	14	28	19.3	10.4	2.6	3.7	2666	140
Xeon Gold 6142	16	32	22.0	10.4	2.6	3.7	2666	150
Xeon Gold 6150	18	36	24.8	10.4	2.7	3.7	2666	165
Xeon Gold 6154	18	36	24.8	10.4	3.0	3.7	2666	200
Xeon Platinum 8168	24	48	33.0	10.4	2.7	3.7	2666	205
Xeon Silver 4114T	10	20	13.8	9.6	2.2	3.0	2400	85
Xeon Gold 5119T	14	28	19.3	10.4	1.9	3.2	2400	85

Xeon Gold 6134M	8	16	24.8	10.4	3.2	3.7	2666	130
Xeon Gold 6140M	18	36	24.8	10.4	2.3	3.7	2666	140
Xeon Gold 6142M	16	32	22.0	10.4	2.6	3.7	2666	150
Xeon Platinum 8160M	24	48	33.0	10.4	2.1	3.7	2666	150
Xeon Platinum 8170M	26	52	35.8	10.4	2.1	3.7	2666	165
Xeon Platinum 8176M	28	56	38.5	10.4	2.1	3.8	2666	165
Xeon Platinum 8180M	28	56	38.5	10.4	2.5	3.8	2666	205

All the processors that can be ordered with the PRIMERGY RX2540 M4, apart from Xeon Bronze 3104 and Xeon Bronze 3106, 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.

Memory modules (since system release)								
Memory module	Capacity [GB]	Ranks	Bit width of the memory chips	Frequency [MHz]	Low voltage	Load reduced	Registered	ECC
8 GB (1x8 GB) 2Rx8 DDR4-2666 R ECC	8	2	8	2666			✓	✓
16 GB (1x16 GB) 2Rx8 DDR4-2666 R ECC	16	2	8	2666			✓	✓
8 GB (1x8 GB) 1Rx4 DDR4-2666 R ECC	8	1	4	2666			✓	✓
16 GB (1x16 GB) 1Rx4 DDR4-2666 R ECC	16	1	4	2666			✓	✓
16 GB (1x16 GB) 2Rx4 DDR4-2666 R ECC	16	2	4	2666			✓	✓
32 GB (1x32 GB) 2Rx4 DDR4-2666 R ECC	32	2	4	2666			✓	✓
64 GB (1x64 GB) 4Rx4 DDR4-2666 3DS ECC	64	4	4	2666			✓	✓
128 GB (1x128 GB) 8Rx4 DDR4-2666 3DS ECC	128	8	4	2666			✓	✓
64 GB (1x64 GB) 4Rx4 DDR4-2666 LR ECC	64	4	4	2666		✓	✓	✓

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

Some components may not be available in all countries or sales regions.
Detailed technical information is available in the data sheet PRIMERGY RX2540 M4.

SPECcpu2006

Benchmark description

SPECcpu2006 is a benchmark which measures the system efficiency with integer and floating-point operations. It consists of an integer test suite (SPECint2006) containing 12 applications and a floating-point test suite (SPECfp2006) containing 17 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.

SPECcpu2006 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.

SPECcpu2006 contains two different performance measurement methods: The first method (SPECint2006 or SPECfp2006) determines the time which is required to process a single task. The second method (SPECint_rate2006 or SPECfp_rate2006) determines the throughput, i.e. the number of tasks that can be handled in parallel. Both methods are also divided into two measurement runs, “base” and “peak”, which differ in the use of compiler optimization. When publishing the results, the base values are always used and the peak values are optional.

Benchmark	Arithmetic	Type	Compiler optimization	Measurement result	Application
SPECint2006	integer	peak	aggressive	Speed	single-threaded
SPECint_base2006	integer	base	conservative		
SPECint_rate2006	integer	peak	aggressive	Throughput	multi-threaded
SPECint_rate_base2006	integer	base	conservative		
SPECfp2006	floating point	peak	aggressive	Speed	single-threaded
SPECfp_base2006	floating point	base	conservative		
SPECfp_rate2006	floating point	peak	aggressive	Throughput	multi-threaded
SPECfp_rate_base2006	floating point	base	conservative		

The measurement results are the geometric average from normalized ratio values which have been determined for individual benchmarks. The geometric average - in contrast to the arithmetic average - means that there is a weighting in favor of the lower individual results. Normalized means that the measurement is how fast is the test system compared to a reference system. Value “1” was defined for the SPECint_base2006, SPECint_rate_base2006, SPECfp_base2006, and SPECfp_rate_base2006 results of the reference system. For example, a SPECint_base2006 value of 2 means that the measuring system has handled this benchmark twice as fast as the reference system. A SPECfp_rate_base2006 value of 4 means that the measuring system has handled this benchmark some 4/[# base copies] times faster than the reference system. “# base copies” specifies how many parallel instances of the benchmark have been executed.

Not every SPECcpu2006 measurement is submitted by us for publication at SPEC. This is why the SPEC web pages do not have every result. As we archive the log files for all measurements, we can prove the correct implementation of the measurements at any time.

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M4
Processor	Intel® Xeon® Processor Scalable Family × 2
Memory	16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC × 24
Software	
BIOS settings	Xeon Platinum 8180, Gold 61XX: HWPM Support = Disabled Intel Virtualization Technology = Disabled Sub NUMA Clustering = Enabled IMC Interleaving = 1-way LLC Dead Line Alloc = Disabled Stale AtoS = Enabled Link Frequency Select = 10.4 GT/s Xeon Silver 4110 HWPM Support = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Sub NUMA Clustering = Disabled IMC Interleaving = 2-way LLC Dead Line Alloc = Disabled Stale AtoS = Enabled Xeon Silver 4116 HWPM Support = Disabled Intel Virtualization Technology = Disabled Sub NUMA Clustering = Enabled IMC Interleaving = 1-way LLC Dead Line Alloc = Disabled Stale AtoS = Enabled Link Frequency Select = 10.4 GT/s
Operating system	SUSE Linux Enterprise Server 12 SP2 (x86_64)
Operating system settings	Stack size set to unlimited using "ulimit -s unlimited" Kernel Boot Parameter set with : nohz_full=1-xx cpupower -c all frequency-set -g performance Tmpfs filesystem can be set with: mkdir /home/memory mount -t tmpfs -o size=752g,rw tmpfs /home/memory Process tunning setting: echo 10000000 > /proc/sys/kernel/sched_min_granularity_ns echo 15000000 > /proc/sys/kernel/sched_wakeup_granularity_ns echo 0 > /proc/sys/kernel/numa_balancing cpupower idle-set -d 1 cpupower idle-set -d 2
Compiler	C/C++: Version 17.0.3.191 of Intel C/C++ Compiler for Linux Version 18.0.0.128 of Intel C++ Compiler Fortran: Version 17.0.3.191 of Intel Fortran Compiler for Linux

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

Benchmark results

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

This results in italic are estimated values from the results of RX2530 M4.

Processor	Number of processors	SPECint_rate_base2006	SPECint_rate2006	SPECint_rate_base2006 Version 18.0.0.128 of Intel C++ Compiler
Xeon Bronze 3104	2	*330*		
Xeon Bronze 3106	2	*440*		
Xeon Silver 4108	2	*641*		
Xeon Silver 4110	2	711	748	
Xeon Silver 4114	2	*911*		
Xeon Silver 4116	2	1060		
Xeon Gold 5115	2	*980*		
Xeon Gold 5118	2	*1160*		
Xeon Gold 5120	2	*1310*		
Xeon Gold 6130	2	1550	1630	
Xeon Gold 6140	2	1760		
Xeon Gold 6138	2	*1770*		
Xeon Gold 6148	2	1960		
Xeon Gold 6152	2	*1990*	*2090*	
Xeon Platinum 8153	2	*1370*		
Xeon Platinum 8160	2	*2170*		
Xeon Platinum 8164	2	*2220*		
Xeon Platinum 8170	2	*2310*		
Xeon Platinum 8176	2	*2440*		
Xeon Platinum 8180	2	2710	2820	2870
Xeon Silver 4112	2	*426*		
Xeon Gold 5122	2	*547*		
Xeon Gold 6128	2	*822*		
Xeon Gold 6134	2	*1060*		
Xeon Gold 6144	2	*1120*		
Xeon Gold 6126	2	*1310*		
Xeon Gold 6136	2	*1480*		
Xeon Gold 6146	2	*1540*		

Xeon Gold 6132	2	*1540*		
Xeon Gold 6142	2	*1710*		
Xeon Gold 6150	2	*1900*		
Xeon Gold 6154	2	*2090*		
Xeon Platinum 8168	2	*2460*		
Xeon Silver 4114T	2	*910*		
Xeon Gold 5119T	2	*1190*		
Xeon Gold 6134M	2	*1060*		
Xeon Gold 6140M	2	*1540*		
Xeon Gold 6142M	2	*1710*		
Xeon Platinum 8160M	2	*2170*		
Xeon Platinum 8170M	2	*2310*		
Xeon Platinum 8176M	2	*2440*		
Xeon Platinum 8180M	2	*2710*	*2820*	

Processor	Number of processors	SPECfp_rate_base2006	SPECfp_rate2006	SPECfp_rate_base2006 Version 18.0.0.128 of Intel C++ Compiler
Xeon Bronze 3104	2	*364*		
Xeon Bronze 3106	2	*481*		
Xeon Silver 4108	2	640		
Xeon Silver 4110	2	690		
Xeon Silver 4114	2	*838*		
Xeon Silver 4116	2	938		
Xeon Gold 5115	2	*874*		
Xeon Gold 5118	2	*993*		
Xeon Gold 5120	2	*1080*		
Xeon Gold 6130	2	*1270*		
Xeon Gold 6140	2	*1380*		
Xeon Gold 6138	2	*1390*		
Xeon Gold 6148	2	1480		
Xeon Gold 6152	2	*1490*	*1520*	
Xeon Platinum 8153	2	*1170*		
Xeon Platinum 8160	2	*1560*		
Xeon Platinum 8164	2	*1590*		
Xeon Platinum 8170	2	*1630*		
Xeon Platinum 8176	2	*1680*		
Xeon Platinum 8180	2	1790	1820	
Xeon Silver 4112	2	*430*		
Xeon Gold 5122	2	*534*		
Xeon Gold 6128	2	*780*		
Xeon Gold 6134	2	*968*		
Xeon Gold 6144	2	*996*		
Xeon Gold 6126	2	*1130*		
Xeon Gold 6136	2	*1230*		
Xeon Gold 6146	2	*1260*		
Xeon Gold 6132	2	*1250*		
Xeon Gold 6142	2	*1350*		
Xeon Gold 6150	2	*1430*		
Xeon Gold 6154	2	*1520*		
Xeon Platinum 8168	2	*1690*		
Xeon Silver 4114T	2	*839*		
Xeon Gold 5119T	2	*1020*		

Xeon Gold 6134M	2	*968*		
Xeon Gold 6140M	2	*1380*		
Xeon Gold 6142M	2	*1350*		
Xeon Platinum 8160M	2	*1560*		
Xeon Platinum 8170M	2	*1630*		
Xeon Platinum 8176M	2	*1680*		
Xeon Platinum 8180M	2	*1790*	*1820*	

The following two diagrams illustrate the throughput of the PRIMERGY RX2540 M4 in comparison to its predecessor PRIMERGY RX2540 M2, in their respective most performant configuration.

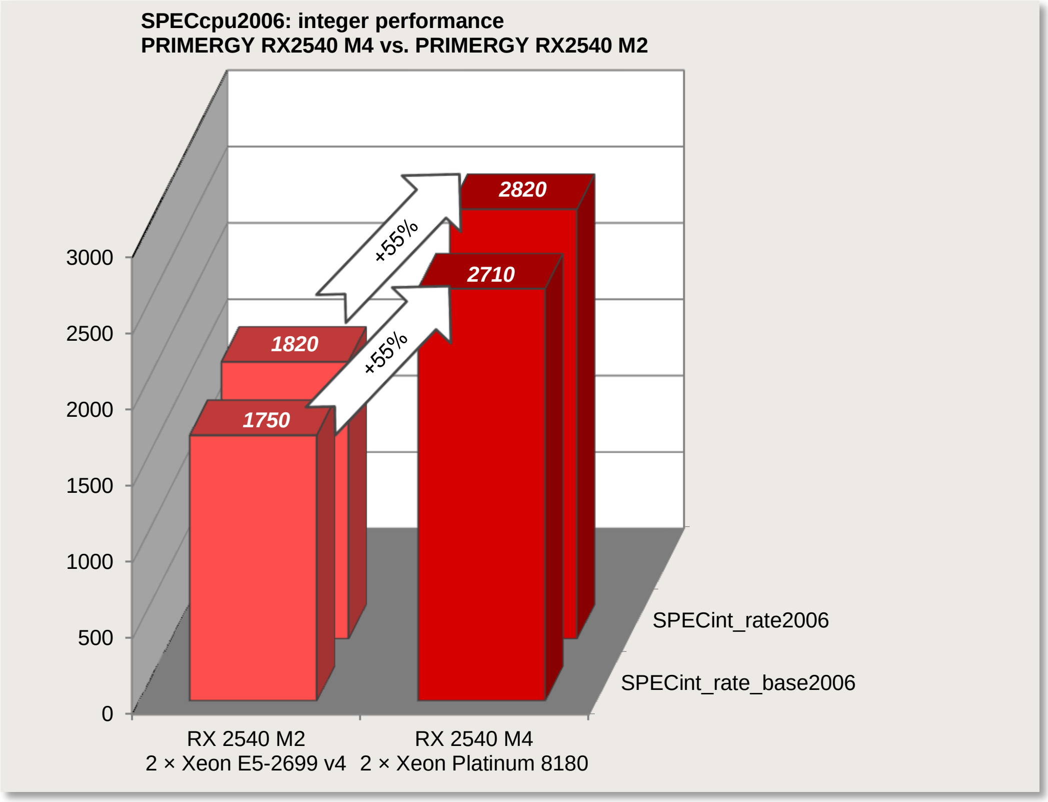

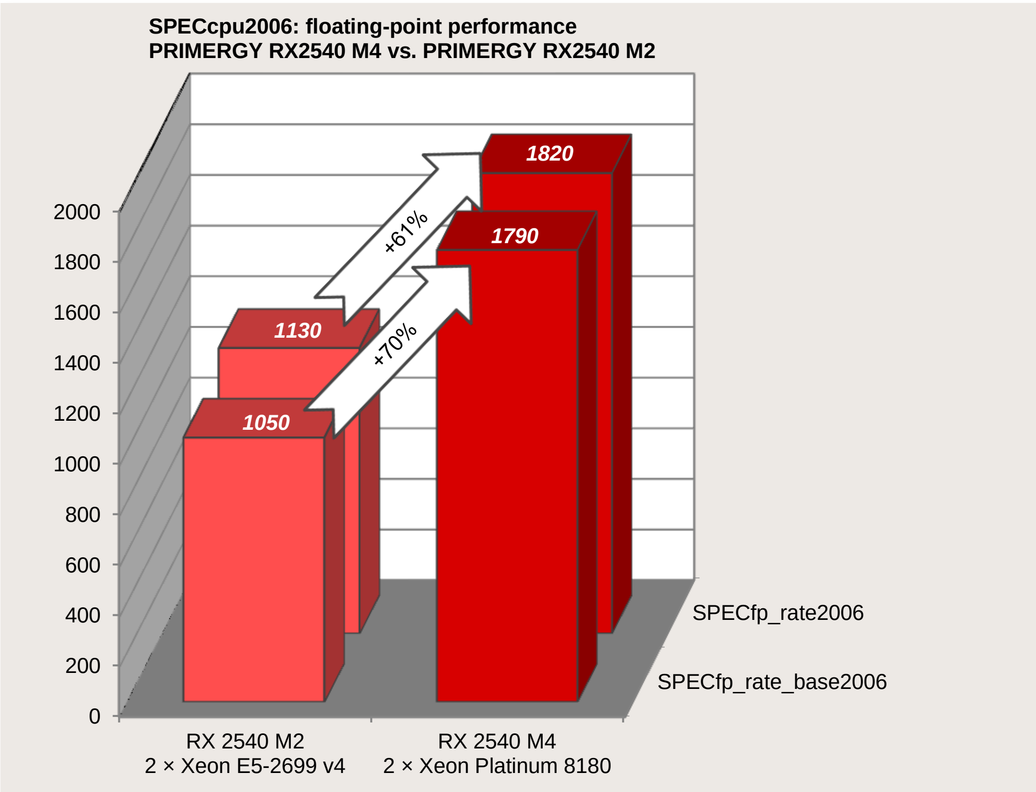

SPECpower_ssj2008

Benchmark description

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

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

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

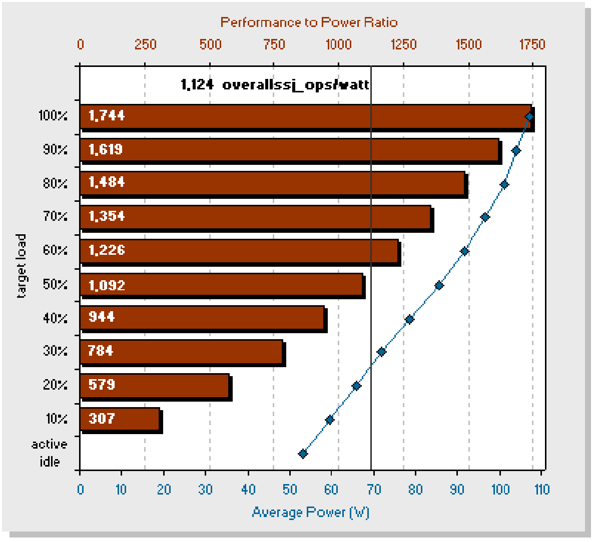

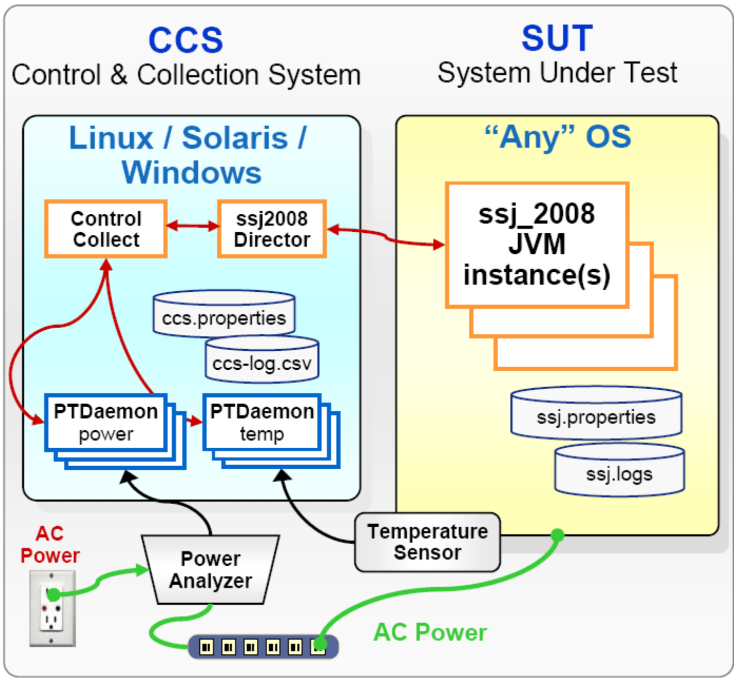

The benchmark runs on a wide variety of operating systems and hardware architectures, and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System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

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M4
Processor	Intel® Xeon® Platinum 8176M
Memory	12 ×16 GB (1x16 GB) 2Rx8 PC4-2666 R ECC
Network interface	1 ×Intel(R) I350 Gigabit Network Connection
Disk subsystem	Onboard SATA. controller 1 × SSD SATA 6G 150 GB DOM N H-P
Power Supply Unit	1 × Modular PSU 800 W titanium hp
Software	
BIOS	R1.16.0
BIOS settings	LAN Controller = LAN1 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 = Balanced Energy DDR Performance = Power balanced Autonomous C-state Support = Enabled Package C State limit = No Limit Link Frequency Select = 9.6 GT/s Uncore Frequency Override = Balanced Power IMC Interleaving = 1-way
Firmware	1.10P
Operating system	Microsoft Windows Server 2012 R2 Standard
Operating system settings	Set "Turn off hard disk after = 1 Minute" in OS. Using the local security settings console, "lock pages in memory" was enabled for the user running the benchmark. Benchmark was started via Windows Remote Desktop Connection.
JVM	Oracle Java HotSpot(TM) 64-Bit Server VM (build 24.80-b11, mixed mode), version 1.7.0_80
JVM settings	-server -Xmn1300m -Xms1550m -Xmx1550m -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:LoopUnrollLimit=45 -XX:InitialTenuringThreshold=12 -XX:MaxTenuringThreshold=15 -XX:ParallelGCThreads=2 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:FreqInlineSize=2500 -XX:+AggressiveOpts -XX:+UseLargePages -XX:+UseParallelOldGC

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

Benchmark results

The PRIMERGY RX2540 M4 achieved the following result:

SPECpower_ssj2008 = 12,842 overall ssj_ops/watt

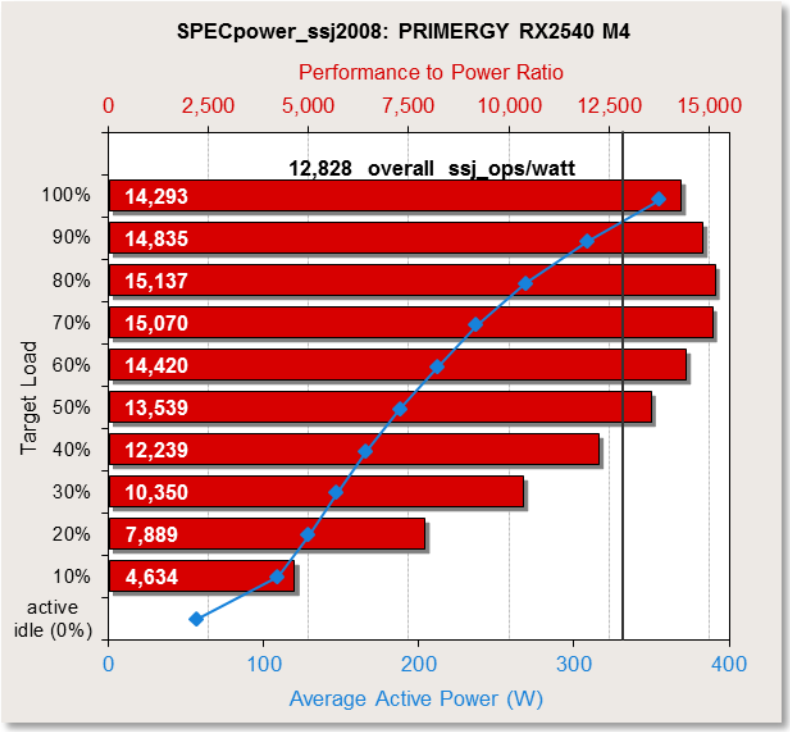

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 12,842 overall ssj_ops/watt for the PRIMERGY RX2540 M2. This is the quotient of the sum of the transaction throughputs for each load level and the sum of the average power consumed for each measurement interval.

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

Performance		**Power**	**Energy Efficiency**
Target Load	ssj_ops	Average Power (W)	ssj_ops/watt
100 %	5,079,736	355	14,293
90 %	4,580,629	309	14,835
80 %	4,076,351	269	15,137
70 %	3,564,103	237	15,070
60 %	3,049,103	212	14,420
50 %	2,543,083	188	13,539
40 %	2,035,379	166	12,239
30 %	1,526,277	147	10,350
20 %	1,019,110	129	7,889
10 %	506,845	109	4,634
Active Idle	0	57.3	0
			∑ssj_ops / ∑power = 12,842

The following diagram shows for each load level the power consumption (on the right y-axis) and the throughput (on the left y-axis) of the PRIMERGY RX2540 M4 compared to the predecessor PRIMERGY RX2540 M2.

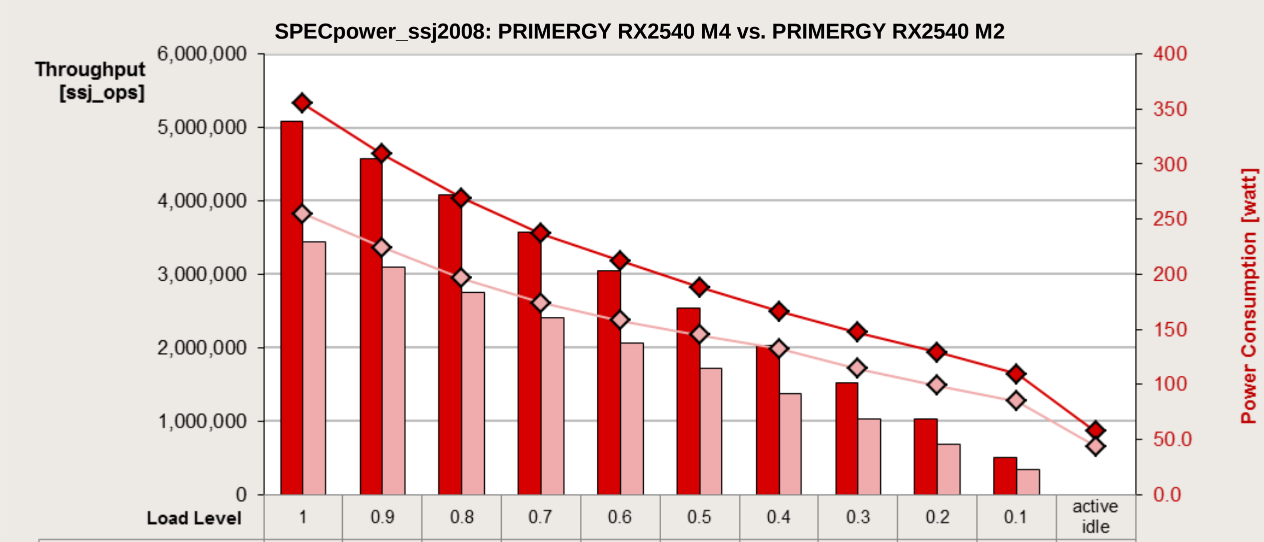

Load Level	1	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	active idle
RX2540 M4 ssj_ops	5,079,73	4,580,62	4,076,35	3,564,10	3,049,10	2,543,08	2,035,37	1,526,27	1,019,11	506,845	0
RX2540 M2 ssj_ops	3,445,02	3,096,29	2,748,12	2,410,31	2,065,07	1,720,55	1,374,69	1,030,25	688,822	342,368	0
RX2540 M4 watt	355	309	269	237	212	188	166	147	129	109	57.3
RX2540 M2 watt	255	224	196	174	158	145	132	114	98.9	84.7	43.6

Thanks to the new Scalable Family processors, the PRIMERGY RX2540 M4 has a higher throughput at substantially lower power consumption compared to the PRIMERGY RX2540 M2. Both result in an overall increase in energy efficiency in the PRIMERGY RX2540 M4 of 10.2%.

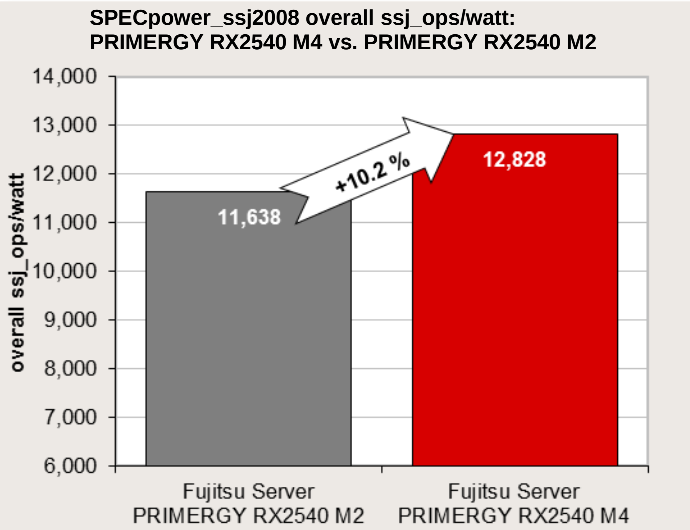

SPECjbb2015

Benchmark description

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

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

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

The SPECjbb2015 benchmark has a two performance metrics:

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

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

- SPECjbb2015 Composite

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

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

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

Results are not comparable to those in other categories.

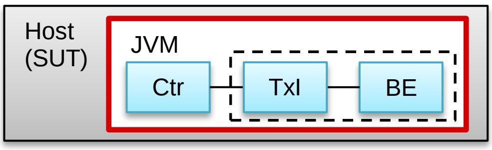

(a) example of SPECjbb2015 Composite configuration

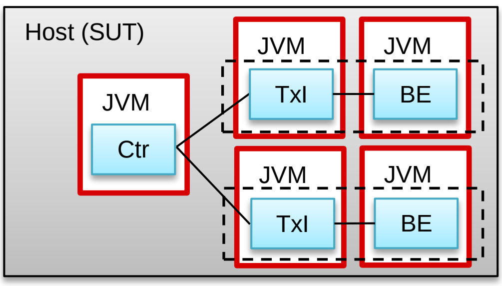

(b) example of SPECjbb2015 MultiJVM configuration

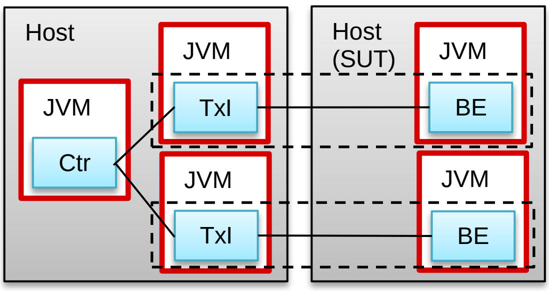

(c) example of SPECjbb2015 Distributed configuration

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

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

Benchmark environment

PRIMERGY RX2540 M4 was configured for the SPECjbb2015 Composite benchmark measurement.

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M4
Processor	2 × Intel® Xeon® Platinum 8180
Memory	24 × 32 GB (1x32 GB) 2Rx4 DDR4-2666 R ECC
Network interface	1 Gbit/s LAN
Disk subsystem	RAID : 1 × PRAID EP420i Disk : 1 × SSD SAS 12 Gb/s 2.5" 400 GB
Software	
BIOS settings	Patrol Scrub set to Disable SNC set to Enabled IMC Interleaving set to 1way
Operating system	Red Hat Enterprise Linux 7.4 (kernel 3.10.0-693.11.6.el7.x86_64 was applied)
Operating system settings	ulimit -l 800000000 ulimit -m 800000000 ulimit -v 800000000 echo 0 > /proc/sys/kernel/numa_balancing tuned-adm profile latency-performance echo 16000000 > /proc/sys/kernel/sched_latency_ns 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
JVM	Oracle Java SE 10.0.2
JVM settings	-server -Xms690g -Xmx690g -Xmn660g -XX:SurvivorRatio=69 --add-modules=java.xml.bind -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=112 -verbose:gc -XX:+PrintGCDetails -XX:+UseHugeTLBFS -XX:+UseTransparentHugePages -XX:+AggressiveHeap
SPECjbb2015 settings	specjbb.comm.connect.client.pool.size = 300 specjbb.comm.connect.worker.pool.max = 300 specjbb.comm.connect.worker.pool.min = 64 specjbb.forkjoin.workers = {Tier1=180, Tier2=1, Tier3=25}

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

Benchmark results

"SPECjbb2015 Composite" measurement result (September 13, 2018)

146,473 SPECjbb2015-Composite max-jOPS

129,399 SPECjbb2015-Composite critical-jOPS

On September 13, 2018 PRIMERGY RX2540 M4 with two Xeon Platinum 8180 processors achieved the scores of 129,399 SPECjbb2015-Composite critical-jOPS. With the result, it ranked first in the 2-socket Xeon server category for SPECjbb2015-Composite critical-jOPS.

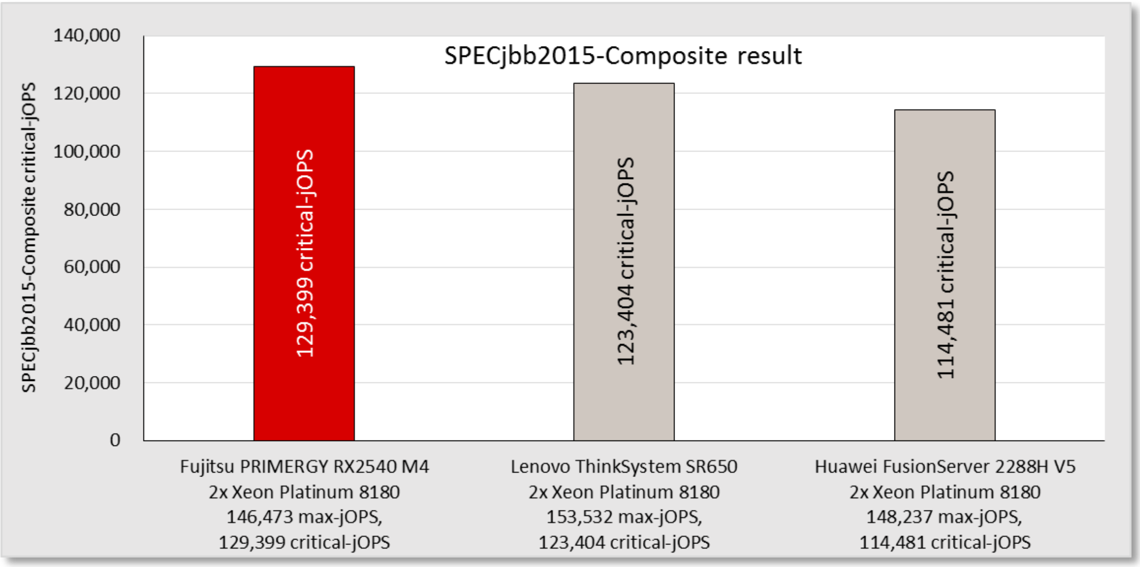

The latest results of the SPECjbb2015 benchmark can be found at https://www.spec.org/jbb2015/results/ .

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: http://www.sap.com/benchmark.

Benchmark environment

The typical measurement set-up is illustrated below:

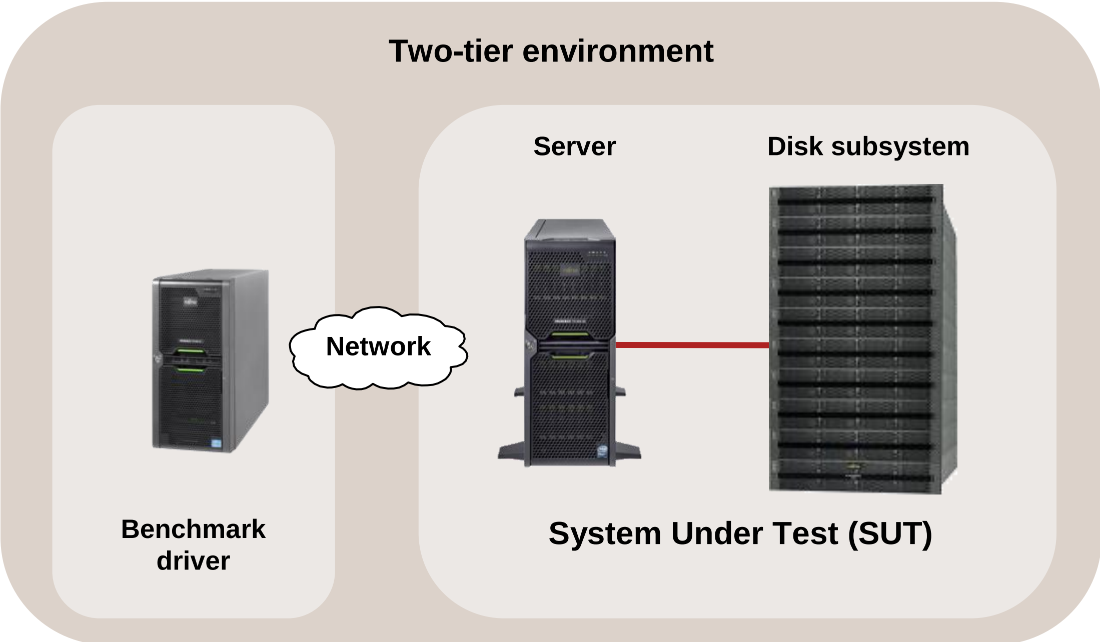

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M4
Processor	2 × Xeon Platinum 8180
Memory	24 × 16 GB (1x16 GB) 2Rx4 DDR4-2666 R ECC
Network interface	1 Gbit/s LAN
Disk subsystem	PRIMERGY RX2540 M4: 1 × SSD SATA 6 Gb/s 2.5" 480 GB 1 × SSD SATA 6 Gb/s 2.5" 1.2 TB
Software	
BIOS settings	Enable SNC
Operating system	Microsoft Windows Server 2012 R2 Standard Edition
Database	Microsoft SQL Server 2012 (64-bit)
SAP Business Suite Software	SAP enhancement package 5 for SAP ERP 6.0

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

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

Benchmark results

Certification number 2017014	
Number of SAP SD benchmark users	29,600
Average dialog response time	0.95 sec
Throughput Fully processed order line items/hour Dialog steps/hour SAPS	 3,244,000 9,732,000 162,200
Average database request time (dialog/update)	0.012 sec / 0.029 sec
CPU utilization of central server	98%
Operating system, central server	Windows Server 2012 R2 Standard Edition
RDBMS	SQL Server 2012
SAP Business Suite software	SAP enhancement package 5 for SAP ERP 6.0
Configuration Central Server	Fujitsu PRIMERGY RX2540 M4 2 processors / 44 cores / 88 threads Intel Xeon Platinum, 2.50 GHz, 64 KB L1 cache and 1024KB L2 cache per core, 38.5 MB L3 cache per processor 384 GB main memory

The following chart shows a comparison of two-tier SAP SD Standard Application Benchmark results for 2-way Xeon Processor Scalable Family based servers with Windows OS and SQL Server database (as of July 11, 2017). The PRIMERGY RX2540 M4 outperforms the comparably configured servers from HPE. The latest SAP SD 2-tier results can be found at http://global.sap.com/solutions/benchmark/sd2tier.epx

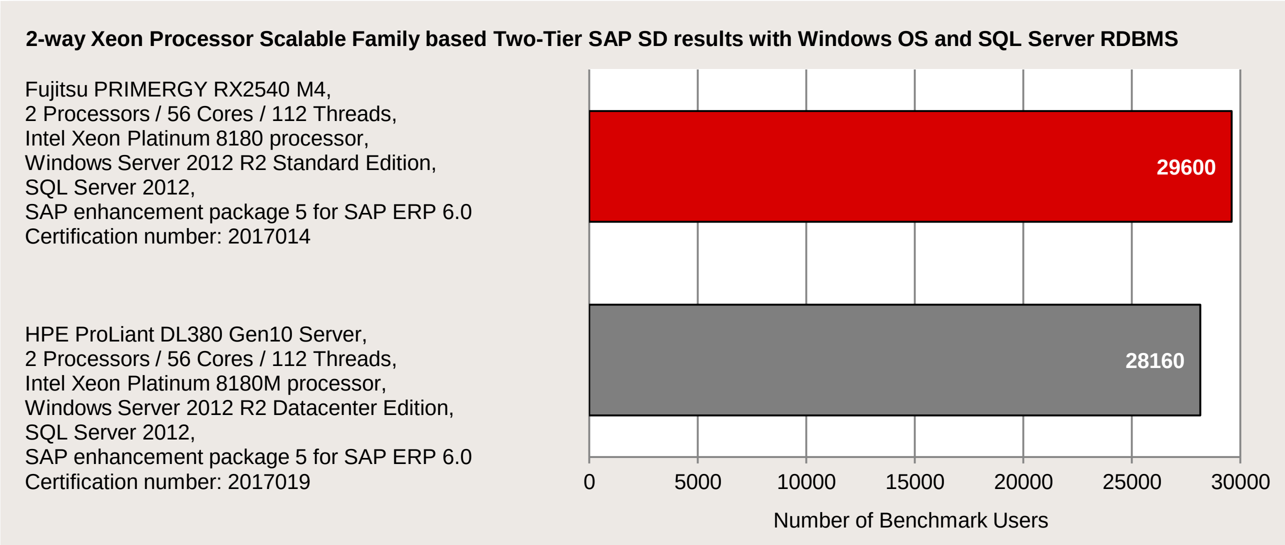

OLTP-2

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:

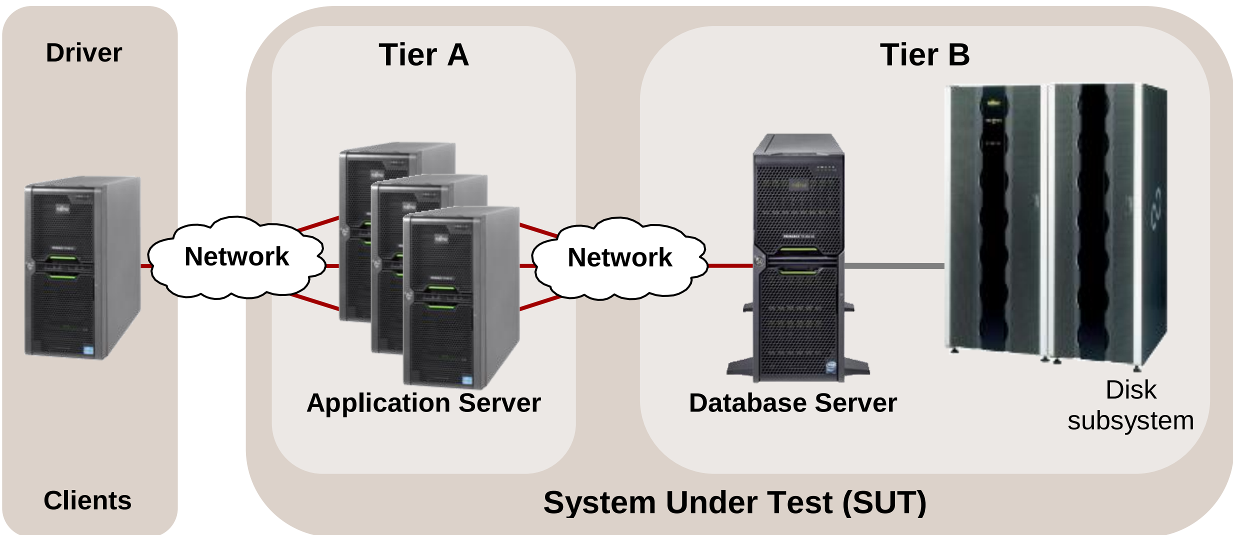

All results were determined by way of example on a PRIMERGY RX2540 M4.

Database Server (Tier B)	
Hardware	
Model	PRIMERGY RX2540 M4
Processor	Intel® Xeon® Processor Scalable Family
Memory	1 processor: 12 × 64 GB (1x64 GB) 4Rx4 DDR4-2666 3DS ECC 2 processors: 24 × 64 GB (1x64 GB) 4Rx4 DDR4-2666 3DS ECC
Network interface	2 × onboard LAN 10 Gb/s
Disk subsystem	RX2540 M4: Onboard RAID controller PRAID EP420i 2 × 300 GB 10k rpm SAS Drive, RAID 1 (OS), 4 × 600 GB 10k rpm SAS Drive, RAID 10 (LOG) 2 × 1.2 TB 10k rpm SAS Drive, RAID 1 (temp) 5 × PRAID EP420e 5 × JX40: 12 × 960 GB SSD Drive each, RAID5 (data)
Software	
BIOS	Version R1.4.1
Operating system	Microsoft Windows Server 2016 Standard
Database	Microsoft SQL Server 2017 Enterprise

Application Server (Tier A)	
Hardware	
Model	1 × PRIMERGY RX2530 M2
Processor	2 × Xeon E5-2690 v4
Memory	128 GB, 2400 MHz registered ECC DDR4
Network interface	2 × 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 2012 R2 Standard

Client	
Hardware	
Model	1 × PRIMERGY RX2530 M2
Processor	2 × Xeon E5-2667 v4
Memory	128 GB, 2400 MHz registered ECC DDR3
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 2666 MHz..

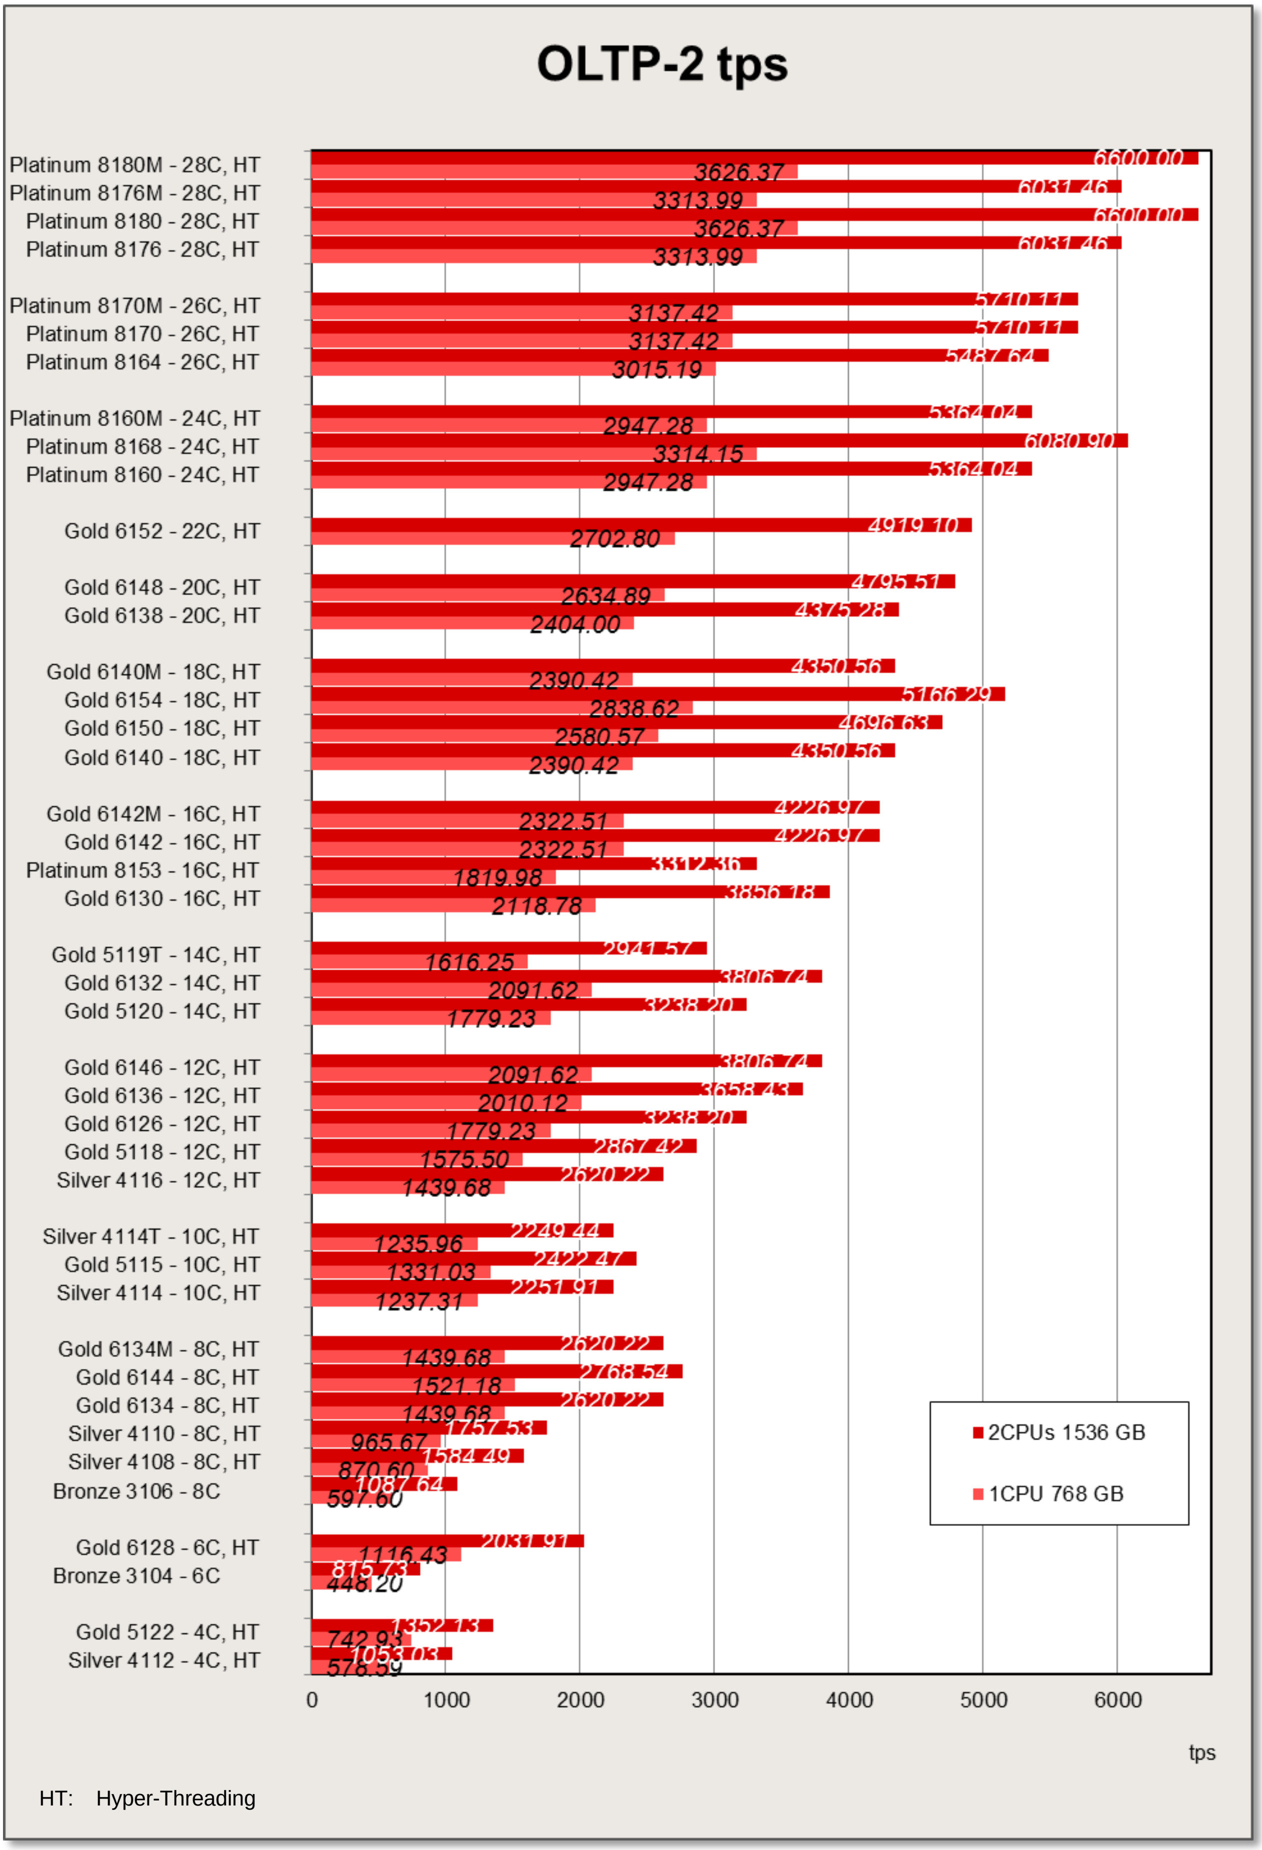
OLTP-2 tps
Platinum 8180M - 28C, HT
Platinum 8176M - 28C, HT
Platinum 8180 - 28C, HT
Platinum 8176 - 28C, HT
Platinum 8170M - 26C, HT
Platinum 8170 - 26C, HT
Platinum 8164 - 26C, HT
Platinum 8160M - 24C, HT
Platinum 8168 - 24C, HT
Platinum 8160 - 24C, HT
Gold 6152 - 22C, HT
Gold 6148 - 20C, HT
Gold 6138 - 20C, HT
Gold 6140M - 18C, HT
Gold 6154 - 18C, HT
Gold 6150 - 18C, HT
Gold 6140 - 18C, HT
Gold 6142M - 16C, HT
Gold 6142 - 16C, HT
Platinum 8153 - 16C, HT
Gold 6130 - 16C, HT
Gold 5119T - 14C, HT
Gold 6132 - 14C, HT
Gold 5120 - 14C, HT
Gold 6146 - 12C, HT
Gold 6136 - 12C, HT
Gold 6126 - 12C, HT
Gold 5118 - 12C, HT
Silver 4116 - 12C, HT
Silver 4114T - 10C, HT
Gold 5115 - 10C, HT
Silver 4114 - 10C, HT
Gold 6134M - 8C, HT
Gold 6144 - 8C, HT
Gold 6134 - 8C, HT
Silver 4110 - 8C, HT
Silver 4108 - 8C, HT
Bronze 3106 - 8C
Gold 6128 - 6C, HT
Bronze 3104 - 6C
Gold 5122 - 4C, HT
Silver 4112 - 4C, HT
6600.00
3626.37
6031.46
3313.99
6600.00
3626.37
6031.46
3313.99
5710.11
3137.42
5710.11
3137.42
5487.64
3015.19
5364.04
2947.28
6080.90
3314.15
5364.04
2947.28
4919.10
2702.80
4795.51
2634.89
4375.28
2404.00
4350.56
2390.42
5166.29
2838.62
4696.63
2580.57
4350.56
2390.42
4226.97
2322.51
4226.97
2322.51
3312.36
1819.98
3856.18
2118.78
2941.57
1616.25
3806.74
2091.62
3238.20
1779.23
3806.74
2091.62
3658.43
2010.12
3238.20
1779.23
2867.42
1575.50
2620.22
1439.68
2249.44
1235.96
2422.47
1331.03
2251.91
1237.31
2620.22
1439.68
2768.54
1521.18
2620.22
1439.68
1757.53
965.67
1584.49
870.60
1087.64
597.60
2031.91
1116.43
815.73
448.20
1352.13
742.93
1053.03
578.59
0
1000
2000
3000
4000
5000
6000
tps
2CPUs 1536 GB
1CPU 768 GB
HT: Hyper-Threading

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

It is evident that a wide performance range is covered by the variety of released processors. If you compare the OLTP-2 value of the processor with the lowest performance (Xeon Bronze 3104) with the value of the processor with the highest performance (Xeon Platinum 8180), 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.

A low performance can be seen in the Xeon Bronze 3104 and Bronze 3106 processors, as they have to manage without Hyper-Threading (HT) and turbo mode (TM).

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 40%.

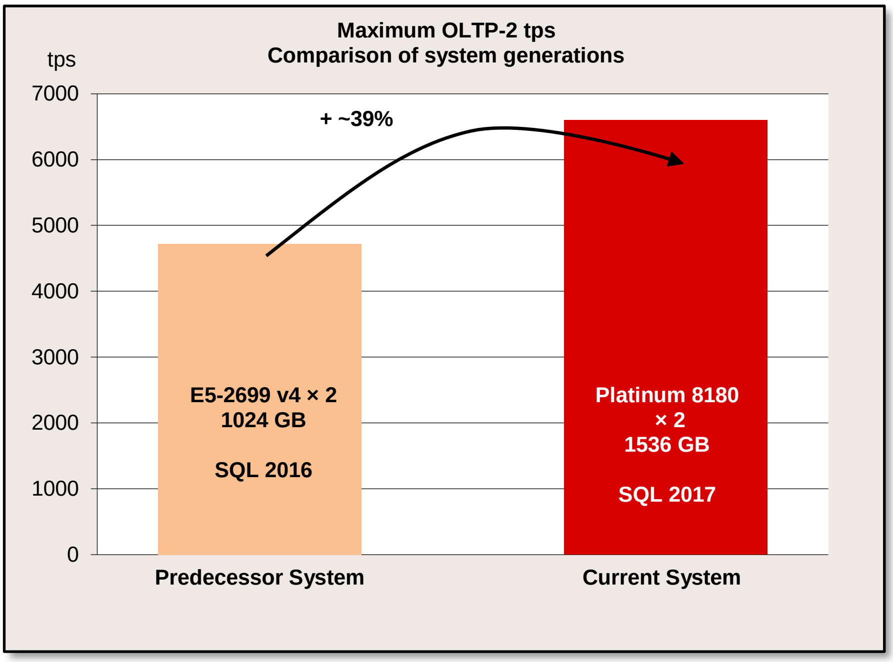

Current System	RX2530 M4	RX2540 M4
Predecessor System	RX2530 M2	RX2540 M2

TPC-E

Benchmark description

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

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

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

Benchmark results

In March 2016 Fujitsu submitted a TPC-E benchmark result for the PRIMERGY RX2540 M4 with the 28-core processor Intel Xeon Platinum8180 and 1536 GB memory.

The results show an enormous increase in performance compared with the PRIMERGY RX2540 M2 with a simultaneous reduction in price per performance ratio.

FUJITSU	**FUJITSU Server PRIMERGY RX2540 M4**	**TPC-E 1.14.0** **TPC Pricing 2.2.0** Report Date March 31, 2018

TPC-E Throughput	Price/Performance	Availability Date	Total System Cost
6,606.75 tpsE	**$ 92.85 USD per tpsE**	**March 31, 2018**	**$ 613,391 USD**

Database Server Configuration

Operating System	Database Manager	Processors/Cores/Threads	Memory
Microsoft Windows Server 2016 Standard Edition	**Microsoft SQL Server 2017 Enterprise Edition**	**2/56/112**	**1536 GB**

SUT

Tier A
PRIMERGY RX2530 M4
2x Intel Xeon Platinum 8180 2.50 GHz
192 GB Memory
2x 300 GB 10k rpm SAS Drive
1x onboard dual port LAN 10 Gb/s
1x onboard dual port LAN 1 Gb/s
1x SAS RAID controller

Tier B
PRIMERGY RX2540 M4
2x Intel Xeon Platinum 8180 2.50 GHz
1,536 GB Memory
2x 300 GB 15k rpm SAS Drives
6x 960 GB SAS SSD
1x onboard dual port LAN 10 Gb/s
1x onboard dual port LAN 1 Gb/s
6x SAS RAID Controller

Storage
1x PRIMECENTER Rack
5x ETERNUS JX40 S2
80x 400 GB SSD Drives

Initial Database Size	Redundancy Level 1	Storage
33,388 GB	**RAID-5 data and RAID-10 log**	**80 x 960 GB SSD**

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

More details about this TPC-E result, in particular the Full Disclosure Report, can be found via the TPC web page http://www.tpc.org/tpce/results/tpce_result_detail.asp?id=118033101.

In March 2018, Fujitsu is represented with five results in the TPC-E list (without historical results).

System and Processors	Throughput	Price / Performance	Availability Date
PRIMERGY RX4770 M2 with 4 × Xeon E7-8890 v3	6904.53 tpsE	$126.49 per tpsE	June 1, 2015
PRIMEQUEST 2800E2 with 8 × Xeon E7-8890 v3	10058.28 tpsE	$187.53 per tpsE	November 11, 2015
PRIMERGY RX2540 M2 with 2 × Xeon E5-2699 v4	4734.87 tpsE	$111.65 per tpsE	July 31, 2016
PRIMERGY RX4770 M3 with 4 × Xeon E7-8890 v4	6904.53 tpsE	$116.62 per tpsE	July 31, 2016
PRIMERGY RX2540 M4 with 2 × Xeon Platinum 8180	6606.75 tpsE	$92.85 per tpsE	March 31, 2018

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

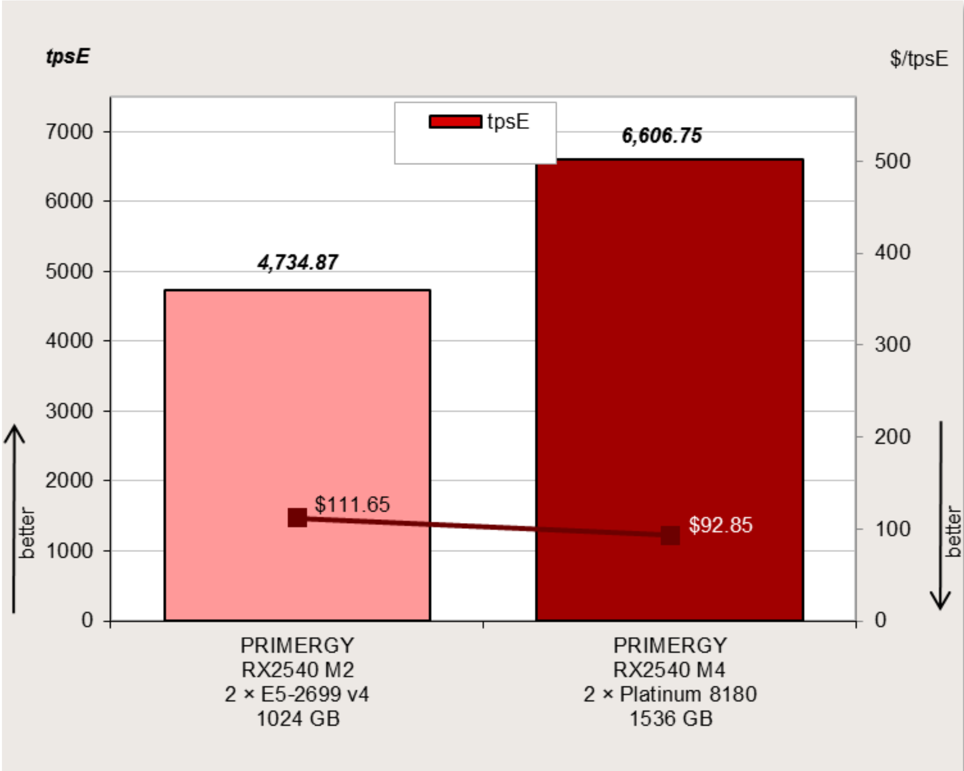

The following diagram for two-socket PRIMERGY systems with different processor types shows the good performance of the two-socket system PRIMERGY RX2540 M4.

System and Processors	Throughput	Price / Performance	Availability Date
PRIMERGY RX2540 M2 with 2 × Xeon E5-2699 v4	4734.87 tpsE	$111.65 per tpsE	July 31, 2016
PRIMERGY RX2540 M4 with 4 × Xeon Platinum 8180	6606.75 tpsE	$92.85 per tpsE	March 31, 2018

In comparison with the PRIMERGY RX2540 M2 the increase in performance is +40%. The price per performance is $92.85 per tpsE. Compared with the PRIMERGY RX2540 M2 are reduced to 83%.

The following overview, sorted according to price/performance, shows the best TPC-E price per performance ratios (as of March 31, 2018, without historical results) and the corresponding TPC-E throughputs. PRIMERGY RX2540 M4 with a price per performance ratio of $92.85 per tpsE achieved the best cost-effectiveness. In addition, PRIMERGY RX2540 M4 with TPC-E throughputs of 6,606.75 tpsE has the best performance value of all two-socket systems.

System		Processor type processors/ cores/threads	tpsE (higher is better)	$/tpsE (lower is better)	availability date
Fujitsu	**PRIMERGY RX2540 M4**	2 × Intel Xeon Platinum 8180	**6,606.75**	**92.85**	**2018-03-31**
Lenovo	ThinkSystem SR650	2 × Intel Xeon Platinum 8180	6,598.36	93.48	2017-10-19
Lenovo	ThinkSystem SR950	4 × Intel Xeon Platinum 8180	11,357.28	98.83	2017-11-16
Fujitsu	PRIMERGY RX2540 M2	2 × Intel Xeon E5-2699 v4	4,734.87	111.65	2016-07-31
Fujitsu	PRIMERGY RX4770 M3	2 ×Intel Xeon E7-8890 v4	8,796.47	116.62	2016-07-31
Lenovo	System x3650 M5	2 × Intel Xeon E5-2699 v4	4,938.14	117.91	2016-07-31
Fujitsu	PRIMERGY RX4770 M2	4 × Intel Xeon E7-8890 v3	6,904.53	126.94	2015-06-01
Lenovo	System x3850 X6	4 × Intel Xeon E7-8890 v3	9,068.00	139.85	2016-07-31
Lenovo	System x3950 X6	8 × Intel Xeon E7-8890 v3	11,058.99	143.91	2015-12-17
Fujitsu	PRIMEQUEST 2800E2	8 × Intel Xeon E7-8890 v3	10,058.28	187.53	2015-11-11

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

vServCon

Benchmark description

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

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

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

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

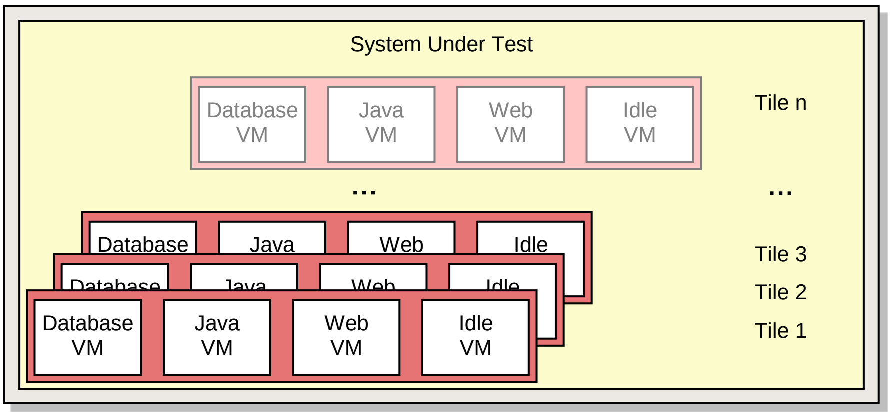

Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score, the individual benchmark result for one tile is put in relation to the respective result of a reference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

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

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

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

Benchmark environment

The typical measurement set-up is illustrated below:

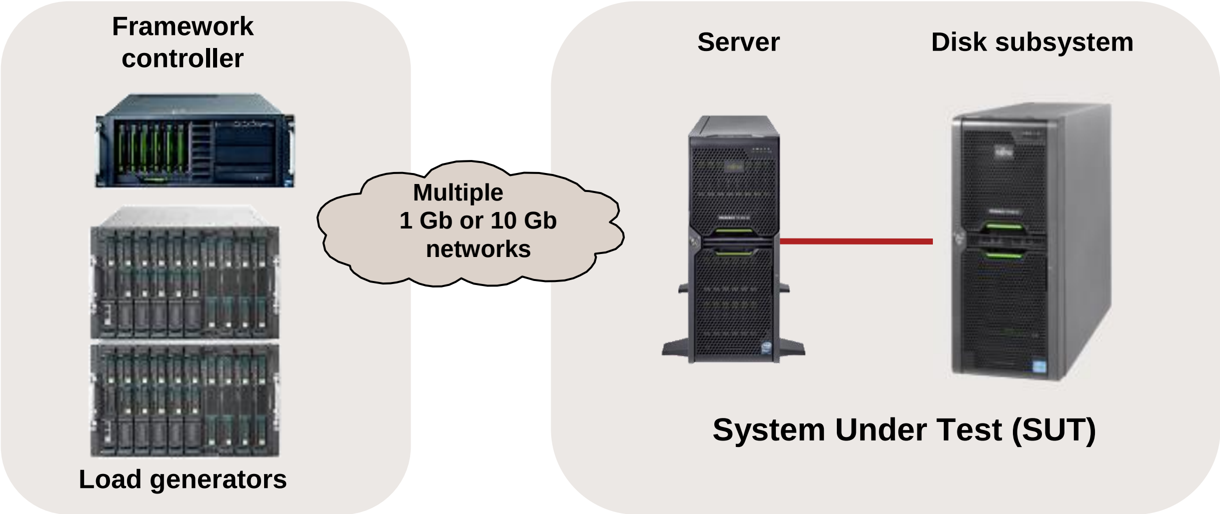

All results were determined by way of example on a PRIMERGY RX2530 M4.

System Under Test (SUT)	
Hardware	
Processor	2 × Intel® Xeon® Processor Scalable Family
Memory	24 × 32 GB (1x32 GB) 2Rx4 DDR4-2666 R ECC
Network interface	1 × Emulex OneConnect OCe14000 Dual Port Adapter with 10Gb SFP+ DynamicLoM interface module
Disk subsystem	1 ×dual-channel FC controller Emulex LPe160021 LINUX/LIO based flash storage system
Software	
Operating system	VMware ESXi 6.5.0b Build 5146846

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

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

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 results in italic are estimated values.

Processor			Score	#Tiles
Intel® Xeon® Processor Scalable Family	4 Cores Hyper-Threading, Turbo-Modus	Silver 4112	*9.09*	*8*
		Gold 5122	11.6	10
	6 Cores	Bronze 3104	*7.05*	*6*
	6 Cores Hyper-Threading, Turbo-Modus	Gold 6128	16.8	12
	8 Cores	Bronze 3106	8.87	7
	8 Cores Hyper-Threading, Turbo-Modus	Silver 4108	*13.7*	*11*
		Silver 4110	*16.5*	*12*
		Gold 6134	*22.6*	*14*
		Gold 6144	24.3	14
		Gold 6134M	*22.6*	*14*
	10 Cores Hyper-Threading, Turbo-Modus	Silver 4114	*19.5*	*16*
		Gold 5115	21.0	16
		Silver 4114T	*19.4*	*16*
	12 Cores Hyper-Threading, Turbo-Modus	Silver 4116	*22.6*	*18*
		Gold 5118	*24.8*	*19*
		Gold 6126	28.3	19
		Gold 6136	*31.6*	*20*
		Gold 6146	33.6	20
	14 Cores Hyper-Threading, Turbo-Modus	Gold 5120	28.2	19
		Gold 6132	32.4	21
		Gold 5119T	*25.4*	*20*
	16 Cores Hyper-Threading, Turbo-Modus	Gold 6130	34.1	21
		Platinum 8153	*28.6*	*21*
		Gold 6142	35.7	22
		Gold 6142M	*35.7*	*22*

	18 Cores Hyper-Threading, Turbo-Modus	Gold 6140	***37.6***	***25***
		Gold 6150	***40.6***	***26***
		Gold 6154	**42.6**	**26**
		Gold 6140M	***37.6***	***26***
	20 Cores Hyper-Threading, Turbo-Modus	Gold 6138	***37.8***	***20***
		Gold 6148	**41.4**	**28**
	22 Cores Hyper-Threading, Turbo-Modus	Gold 6152	**43.4**	**25**
	24 Cores Hyper-Threading, Turbo-Modus	Platinum 8160	***46.3***	***29***
		Platinum 8168	**52.7**	**30**
		Platinum 8160M	***46.3***	***29***
	26 Cores Hyper-Threading, Turbo-Modus	Platinum 8164	**50.9**	**32**
		Platinum 8170	**51.2**	**32**
		Platinum 8170M	***51.2***	***32***
	28 Cores Hyper-Threading, Turbo-Modus	Platinum 8176	***52.1***	***34***
		Platinum 8180	**59.4**	**34**
		Platinum 8176M	***52.1***	***34***
		Platinum 8180M	***59.4***	***34***

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 53% 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.

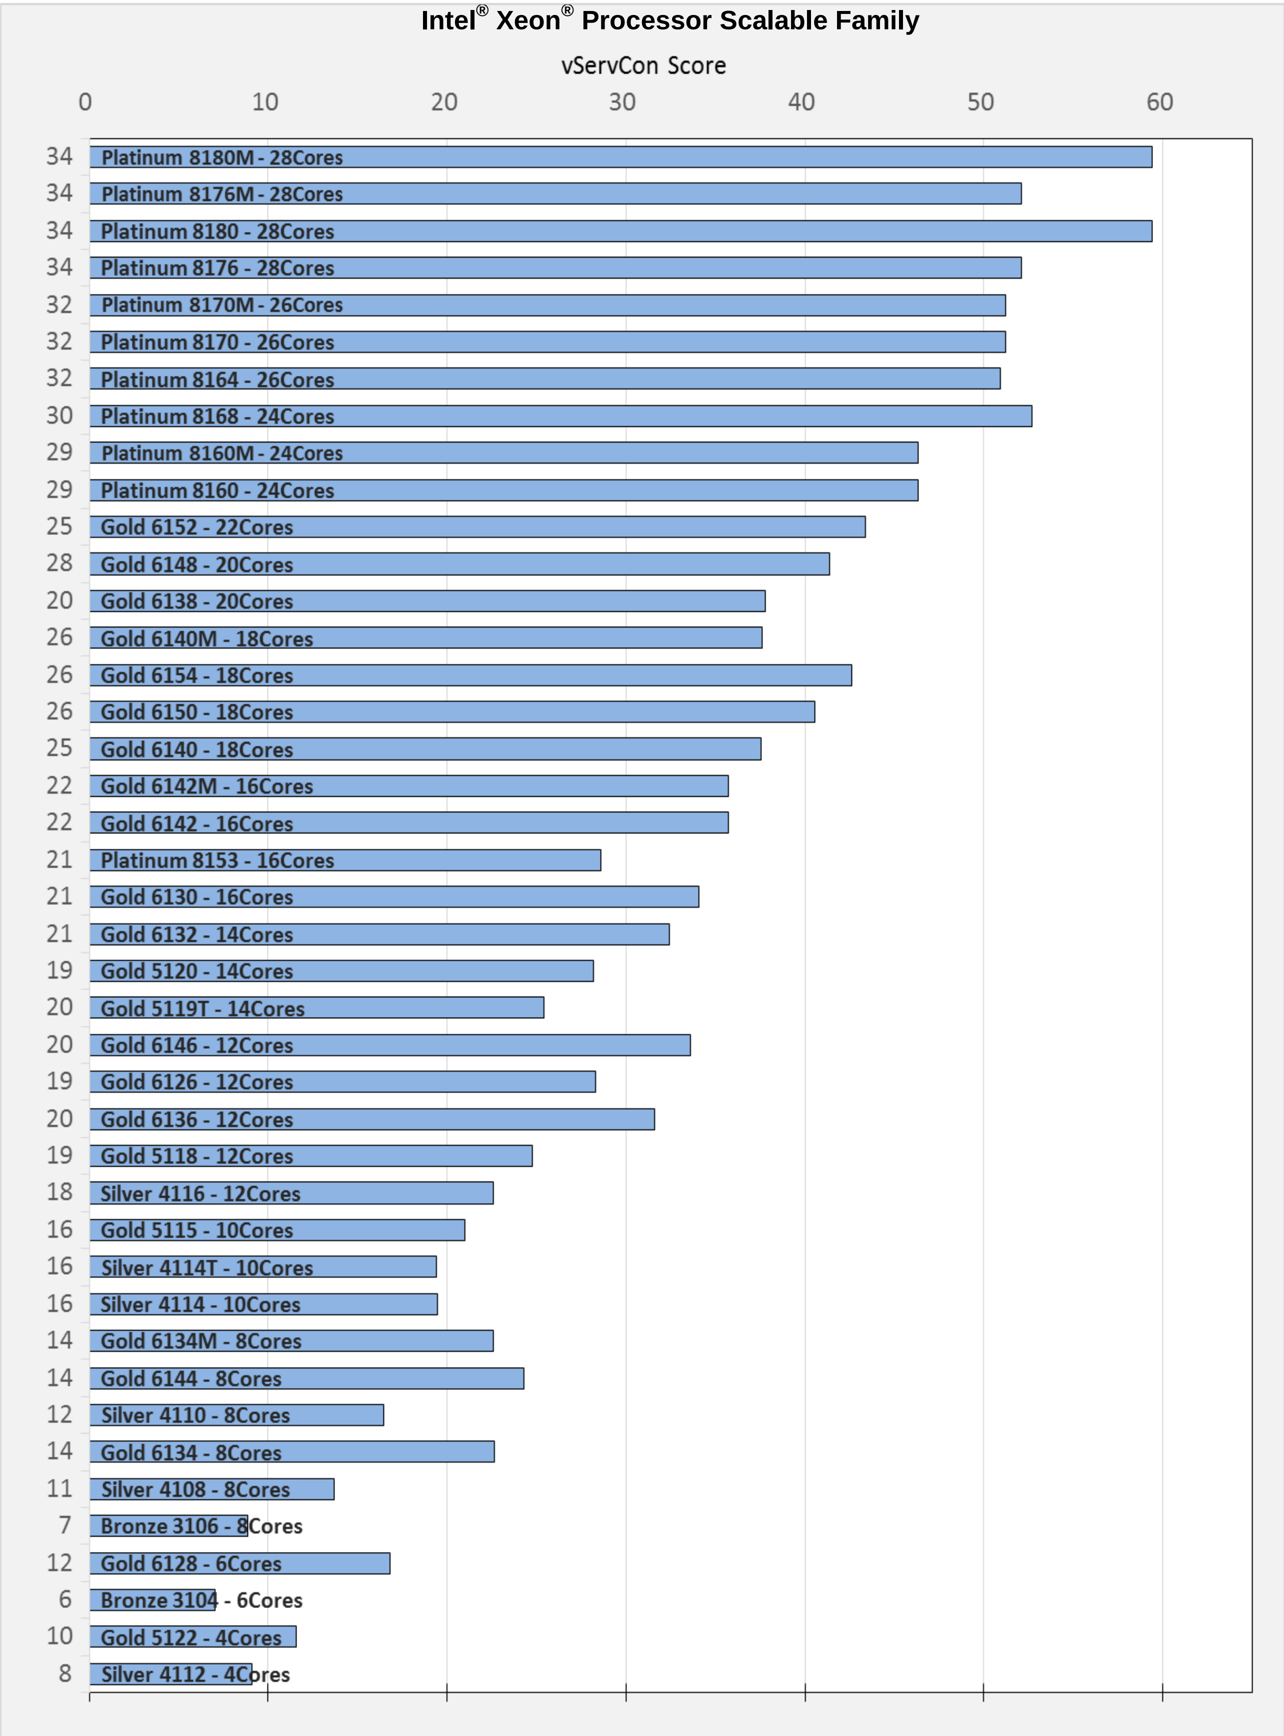

The relatively large performance differences between the processors can be explained by their features. The values scale on the basis of the number of cores, the size of the L3 cache and the CPU clock frequency and as a result of the features of Hyper-Threading and turbo mode, which are available in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

A low performance can be seen in the Xeon Bronze 3104 and Bronze 3106 processors, as they have to manage without Hyper-Threading (HT) and turbo mode (TM). In principle, these weakest processors are only to a limited extent suitable for the virtualization environment.

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

As a matter of principle, the memory access speed also influences performance. A guideline in the virtualization environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses. The vServCon scaling measurements presented here were all performed with a memory access speed – depending on the processor type – of at most 2666 MHz.

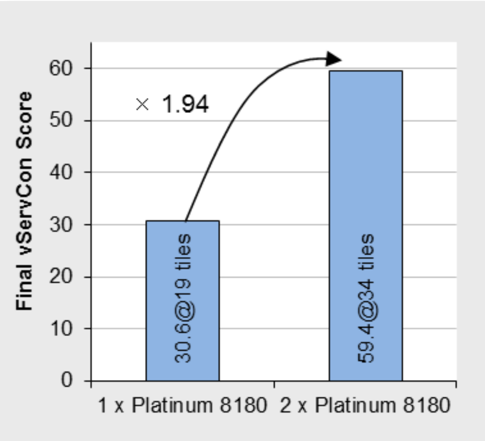

Until now, we have looked at the virtualization performance of a fully configured system. However, with a server with two sockets, the question also arises as to how good performance scaling is from one to two processors. The better the scaling, the lower the overhead usually caused by the shared use of resources within a server. The scaling factor also depends on the application. If the server is used as a virtualization platform for server consolidation, the system scales with a factor of 1.94. When operated with two processors, the system thus achieves a significantly better performance than with one processor, as is illustrated in this diagram using the processor version Xeon Platinum 8180 as an example.

The next diagram illustrates the virtualization performance for increasing numbers of VMs based on the Xeon Gold6130 (16 core) processors.

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

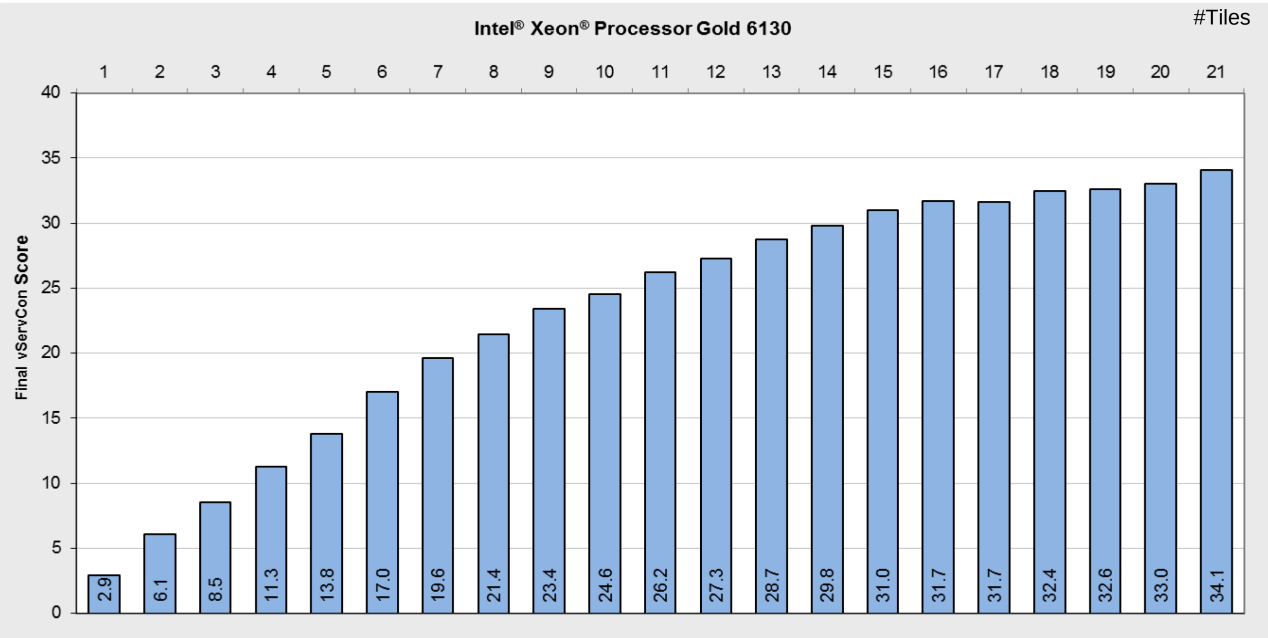

The previous diagram examined the total performance of all application VMs of a host. However, studying the performance from an individual application VM viewpoint is also interesting. This information is in the previous diagram. For example, the total optimum is reached in the above Xeon Gold 6130 situation with 63 application VMs (21 tiles, not including the idle VMs). The low load case is represented by three application VMs (one tile, not including the idle VM). Remember, the vServCon score for one tile is an average value across the three application scenarios in vServCon. This average performance of one tile drops when changing from the low load case to the total optimum of the vServCon score - from 2.93 to 34.1/21=1.62, i.e. to 55%. 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 virtualization-relevant progress in processor technology since 2008 has an effect on the one hand on an individual VM and, on the other hand, on the possible maximum number of VMs up to CPU full utilization. The following comparison shows the proportions for both types of improvements.

Seven systems with similar housing construction are compared with the best processors each (see table below) for few VMs and for highest maximum performance.

Year	2008	2009	2011	2012	2013	2014/2015	2016	2017
Comparison Server	RX200 S4	RX200 S5	RX200 S6	RX200 S7	RX200 S8	RX2530 M1	RX2530 M2	RX2530 M4
	RX300 S4	RX300 S5	RX300 S6	RX300 S7	RX300 S8	RX2540 M1	RX2540 M2	RX2540 M4
	-	-	TX300 S6	RX350 S7	RX350 S8	RX2560 M1	RX2560 M2	-
	TX300 S4	TX300 S5	TX300 S6	TX300 S7	TX300 S8	TX2560 M1	TX2560 M2	-

	Best Performance Few VMs	vServCon Score 1 Tile	Best Maximum Performance	vServCon Score max.
2008	X5460	1.91	X5460	2.94@2 tiles
2009	X5570	2.45	X5570	6.08@ 6 tiles
2011	X5690	2.63	X5690	9.61@ 9 tiles
2012	E5-2643	2.73	E5-2690	13.5@ 8 tiles
2013	E5-2667 v2	2.85	E5-2697 v2	17.1@11 tiles
2014	E5-2643 v3	3.22	E5-2699 v3	30.3@18 tiles
2016	E5-2637 v4	3.29	E5-2699 v4	38.7@22 tiles
2017	Gold 6144	3.56	Platinum 8180	59.4@34 tiles

The clearest performance improvements arose from 2008 to 2009 with the introduction of the Xeon 5500 processor generation (e. g. via the feature “Extended Page Tables” (EPT)[1]). One sees an increase of the vServCon score by a factor of 1.28 with a few VMs (one tile).

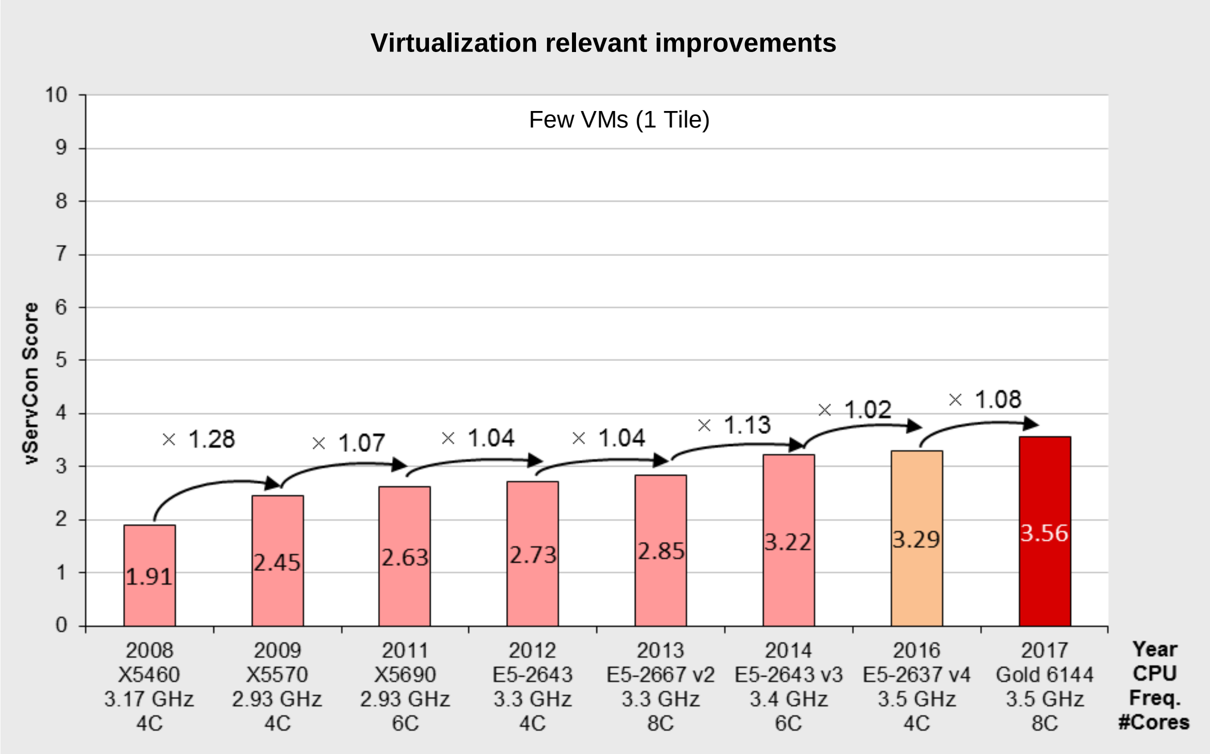

With full utilization of the systems with VMs there was an increase by a factor of 2.07. The one reason was the performance increase that could be achieved for an individual VM (see score for a few VMs). The other reason was that more VMs were possible with total optimum (via Hyper-Threading). However, it can be seen that the optimum was “bought” with a triple number of VMs with a reduced performance of the individual VM.

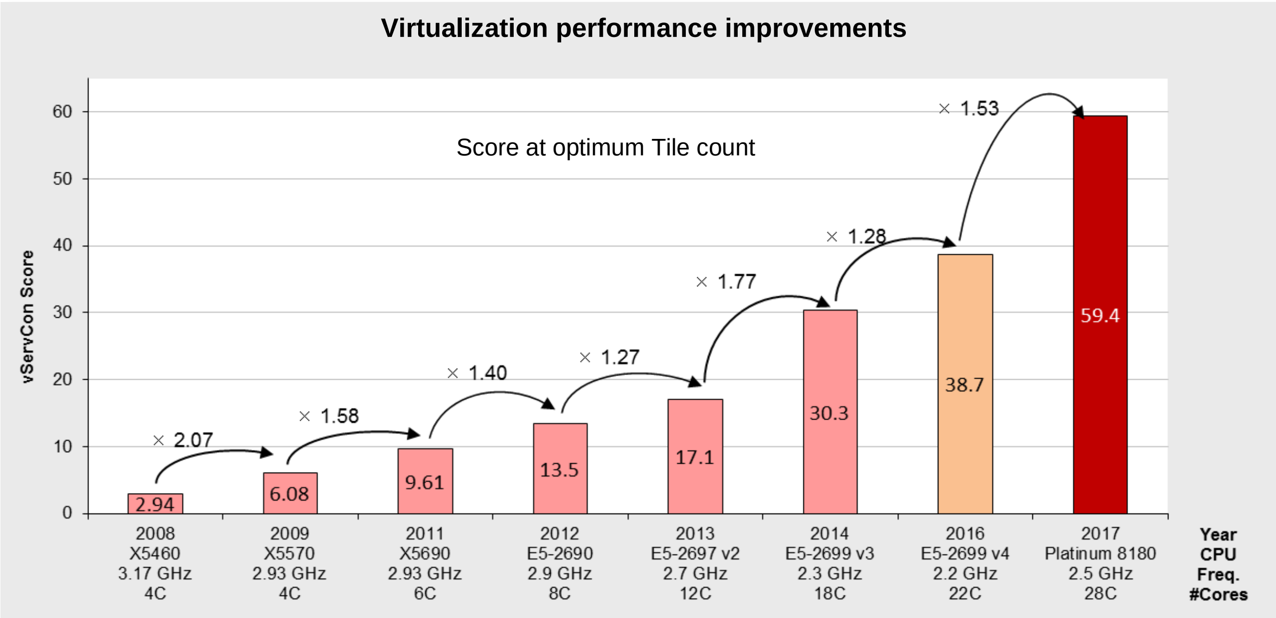

[1] EPT accelerates memory virtualization via hardware support for the mapping between host and guest memory addresses.

Where exactly is the technology progress between 2009 and 2017?

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.

The decisive progress is in the higher number of physical cores and – associated with it – in the increased values of maximum performance (factor 1.58, 1.40, 1.27, 1.77, 1.28 and 1.53 in the diagram).

Up to and including 2011 the best processor type of a processor generation had both the highest clock frequency and the highest number of cores. From 2012 there have been differently optimized processors on offer: Versions with a high clock frequency per core for few cores and versions with a high number of cores, but with a lower clock frequency per core. The features of the processors are summarized in the section "Technical data".

Performance increases in the virtualization environment since 2009 are mainly achieved by increased VM numbers due to the increased number of available logical or physical cores. However, since 2012 it has been possible - depending on the application scenario in the virtualization environment – to also select a CPU with an optimized clock frequency if a few or individual VMs require maximum computing power.

VMmark V3

Benchmark description

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

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

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

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

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

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

The result of VMmark V3 for test type “Performance Only” is a number, known as a “score”, which provides information about the performance of the measured virtualization solution. The score reflects the maximum total consolidation benefit of all VMs for a server configuration with hypervisor and is used as a comparison criterion of various hardware platforms.

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

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

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

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

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

Benchmark environment

The typical measurement set-up is illustrated below:

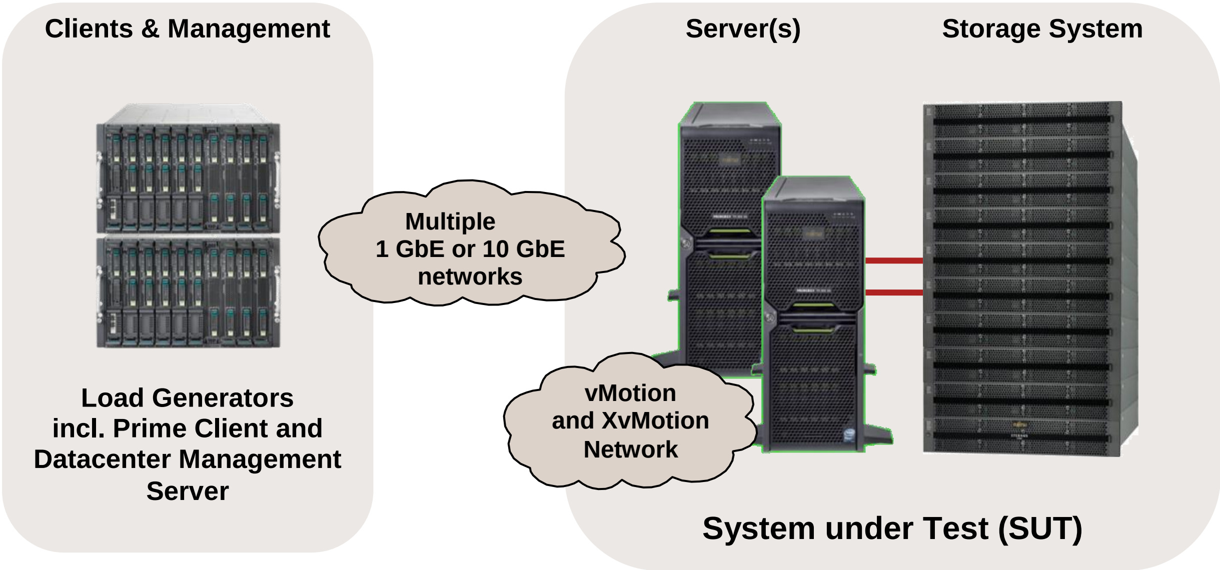

System Under Test (SUT)	
Hardware	
Number of servers	2
Model	PRIMERGY RX2540 M4
Processor	2 × Intel® Xeon® Platinum 8180
Memory	768 GB: 24 × 32 GB (1x32 GB) 2Rx4 DDR4-2666 R ECC
Network interface	2 × Emulex OneConnect OCe14000 Dual Port 10 GbE Adapter 1 × Intel I350-T2 Dual Port 1 GbE Adapter
Disk subsystem	2 × Dual port PFC EP LPe31002 3 × PRIMERGY RX2540 M2 configured as Fibre Channel target: 1 × SAS-SSD (400 GB) 1 × Fusion-io ioMemory PX600(1.3 TB) 3 × Fusion-io ioMemory PX600(2.6 TB) RAID 0 with several LUNs Total: 28.5 TB
Software	
BIOS	R1.7.0
BIOS settings	See details
Operating system	VMware ESXi 6.5.0b Build 5146846
Operating system settings	ESX settings: see details

Details	
See disclosure	http://www.vmware.com/a/assets/vmmark/pdf/2018-01-02-Fujitsu-RX2540M4.pdf http://www.vmware.com/a/assets/vmmark/pdf/2018-01-02-Fujitsu-RX2540M4-serverPPKW.pdf http://www.vmware.com/a/assets/vmmark/pdf/2018-01-16-Fujitsu-RX2540M4.pdf

Datacenter Management Server (DMS)	
Hardware	
Model	1 × PRIMERGY RX2530 M2
Processor	2 × Intel Xeon E5-2698
Memory	64 GB
Network interface	1 × Emulex One Connect Oce14000 1 GbE Dual Port Adapter
Software	
Operating system	VMware ESXi 6.0.0 U2 Build 3620759
Datacenter Management Server (DMS) VM	
Hardware	
Processor	8 × logical CPU
Memory	32 GB
Network interface	1 × 1 Gbit/s LAN
Software	
Operating system	Microsoft Windows Server 2008 R2 Enterprise x64 Edition

Load generator	
Hardware	
Model	2 × PRIMERGY RX2530 M2
Processor	2 × Xeon E5-2699 v4
Memory	128 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 6.0.0 U2 Build 3620759

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

Benchmark results

"Performance Only" measurement result (January 2 2018)

On January 2, 2018 Fujitsu achieved with a PRIMERGY RX2540 M4 with Xeon Platinum 8180 processors and VMware ESXi 6.5.0b a VMmark V3 score of "8.11@8 tiles" in a system configuration with a total of 2 × 56 processor cores and when using two identical servers in the "System under Test" (SUT). With this result the PRIMERGY RX2540 M4 is in the official VMmark V3 "Performance Only" ranking the most powerful two-socket server in a "matched pair" configuration consisting of two identical hosts (valid as of benchmark results publication date).

All comparisons for the competitor products reflect the status of January 2, 2018. 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 diagram shows the "Performance Only" result of the PRIMERGY RX2540 M4 in comparison with the best two-socket systems in a "matched pair" configuration.

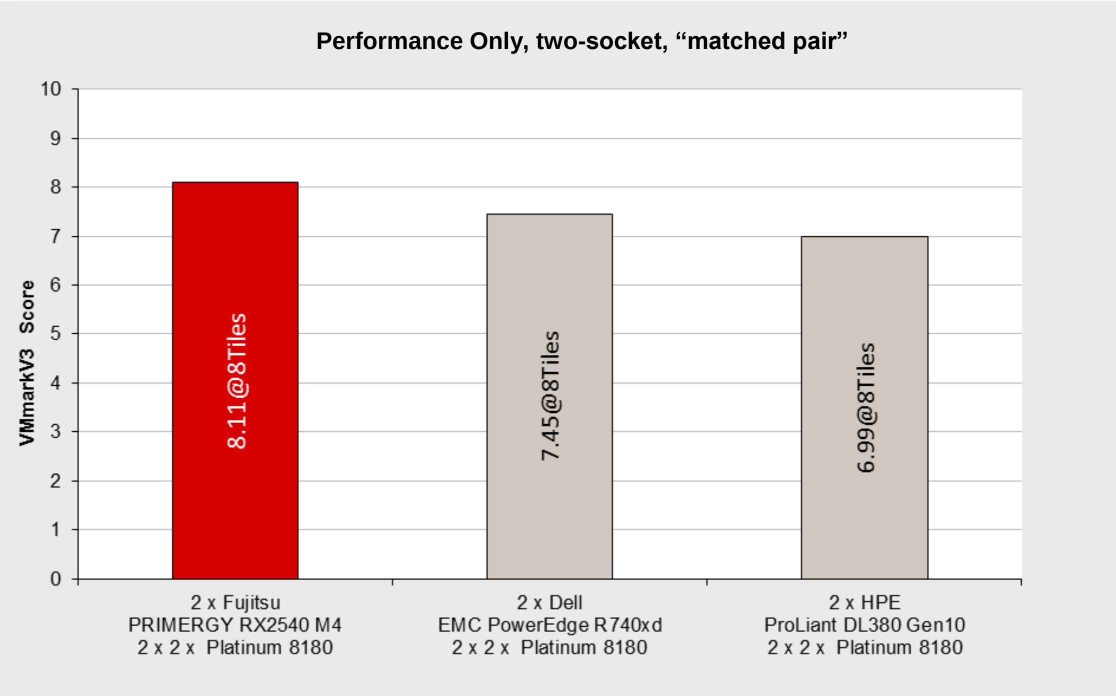

two-socket systems, "matched pair"	VMmark V3 Score	Difference
Fujitsu PRIMERGY RX2540 M4	8.11@8 tiles	
EMC PowerEdge R740xd	7.45@8 tiles	8.86 %
HPE ProLiant ML350 Gen10	6.99@8 tiles	16.02 %

The processors used, which with a good hypervisor setting could make optimal use of their processor features, were the essential prerequisites for achieving the PRIMERGY RX2540 M4 result. These features include Hyper-Threading. All this has a particularly positive effect during virtualization.

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

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

All the components used were optimally attuned to each other.

“Performance with Server Power” measurement result (January 2 2018)

On January 2, 2018 Fujitsu achieved with a PRIMERGY RX2540 M4 with Xeon Platinum 8180 processors and VMware ESXi 6.5.0b a VMmark V3 “Server PPKW Score” of “6.0863@8 tiles” in a system configuration with a total of 2 × 56 processor cores and when using two identical servers in the “System under Test” (SUT). With this result the PRIMERGY RX2540 M4 is in the official VMmark V3 “Performance with Server Power” ranking the most energy-efficient virtualization server worldwide (valid as of benchmark results publication date).

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

“Performance with Server and Storage Power” measurement result (January 16 2018)

On January 16, 2018 Fujitsu achieved with a PRIMERGY RX2540 M4 with Xeon Platinum 8180 processors and VMware ESXi 6.5.0b a VMmark V3 “Server and Storage PPKW Score” of “3.6750@8 tiles” in a system configuration with a total of 2 × 56 processor cores and when using two identical servers in the “System under Test” (SUT). With this result the PRIMERGY RX2540 M4 is in the official VMmark V3 “Performance with Server and Storage Power” ranking the most energy-efficient virtualization platform worldwide (valid as of benchmark results publication date).

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

STREAM

Benchmark description

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

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

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

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

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

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

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

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

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

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M4
Processor	2 × Intel® Xeon® Processor Scalable Family
Memory	24 × 16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC
Software	
BIOS settings	Link Frequency Select = 10.4 GT/s HWPM Support = Disabled Intel Virtualization Technology = Disabled Sub NUMA Clustering = Disabled IMC Interleaving = 2-way LLC Dead Line Alloc = Disabled Stale AtoS = Enabled
Operating system	SUSE Linux Enterprise Server 12 SP2 (x86_64)
Operating system settings	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 nohz_full = 1-xx Xeon Platinum 8180 : run with avx512 Xeon Silver 4116 : run with avx2
Compiler	Version 17.0.0.098 of Intel C++ Compiler for Linux
Benchmark	Stream.c Version 5.10

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

Benchmark results

This results in italic are estimated values from the results of RX2530 M4.

Processor	Memory Frequency [MHz]	Max. Memory Bandwidth [GB/s]	Cores	Processor Frequency [GHz]	Number of Processors	TRIAD [GB/s]
Xeon Bronze 3104	2133	102.4	6	1.7	2	*85.6*
Xeon Bronze 3106	2133	102.4	8	1.7	2	*116*
Xeon Silver 4108	2400	115.2	8	1.8	2	*129*
Xeon Silver 4110	2400	115.2	8	2.1	2	*129*
Xeon Silver 4114	2400	115.2	10	2.2	2	*151*
Xeon Silver 4116	2400	115.2	12	2.1	2	163
Xeon Gold 5115	2400	115.2	10	2.4	2	*154*
Xeon Gold 5118	2400	115.2	12	2.3	2	*154*
Xeon Gold 5120	2400	115.2	14	2.2	2	*182*
Xeon Gold 6130	2666	128.0	16	2.1	2	*208*
Xeon Gold 6140	2666	128.0	18	2.3	2	*212*
Xeon Gold 6138	2666	128.0	20	2.0	2	*218*
Xeon Gold 6148	2666	128.0	20	2.4	2	*218*
Xeon Gold 6152	2666	128.0	22	2.1	2	*220*
Xeon Platinum 8153	2666	128.0	16	2.0	2	*209*
Xeon Platinum 8160	2666	128.0	24	2.1	2	*223*
Xeon Platinum 8164	2666	128.0	26	2.0	2	*224*
Xeon Platinum 8170	2666	128.0	26	2.1	2	*223*
Xeon Platinum 8176	2666	128.0	28	2.1	2	*223*
Xeon Platinum 8180	2666	128.0	28	2.5	2	224
Xeon Silver 4112	2400	115.2	4	2.6	2	*91.0*
Xeon Gold 5122	2666	128.0	4	3.6	2	*109*
Xeon Gold 6128	2666	128.0	6	3.4	2	*165*
Xeon Gold 6134	2666	128.0	8	3.2	2	*194*
Xeon Gold 6144	2666	128.0	8	3.5	2	*190*
Xeon Gold 6126	2666	128.0	12	2.6	2	*199*
Xeon Gold 6136	2666	128.0	12	3.0	2	*214*
Xeon Gold 6146	2666	128.0	12	3.2	2	*214*
Xeon Gold 6132	2666	128.0	14	2.6	2	*197*
Xeon Gold 6142	2666	128.0	16	2.6	2	*203*
Xeon Gold 6150	2666	128.0	18	2.7	2	*213*
Xeon Gold 6154	2666	128.0	18	3.0	2	*212*
Xeon Platinum 8168	2666	128.0	24	2.7	2	*221*
Xeon Silver 4114T	2400	115.2	10	2.2	2	*150*
Xeon Gold 5119T	2400	115.2	14	1.9	2	*183*
Xeon Gold 6134M	2666	128.0	8	3.2	2	*194*

Xeon Gold 6140M	2666	128.0	18	2.3	2	*212*
Xeon Gold 6142M	2666	128.0	16	2.6	2	*203*
Xeon Platinum 8160M	2666	128.0	24	2.1	2	*223*
Xeon Platinum 8170M	2666	128.0	26	2.1	2	*223*
Xeon Platinum 8176M	2666	128.0	28	2.1	2	*223*
Xeon Platinum 8180M	2666	128.0	28	2.5	2	224

The following diagram illustrates the throughput of the PRIMERGY RX2540 M4 in comparison to its predecessor, the PRIMERGY RX2540 M2.

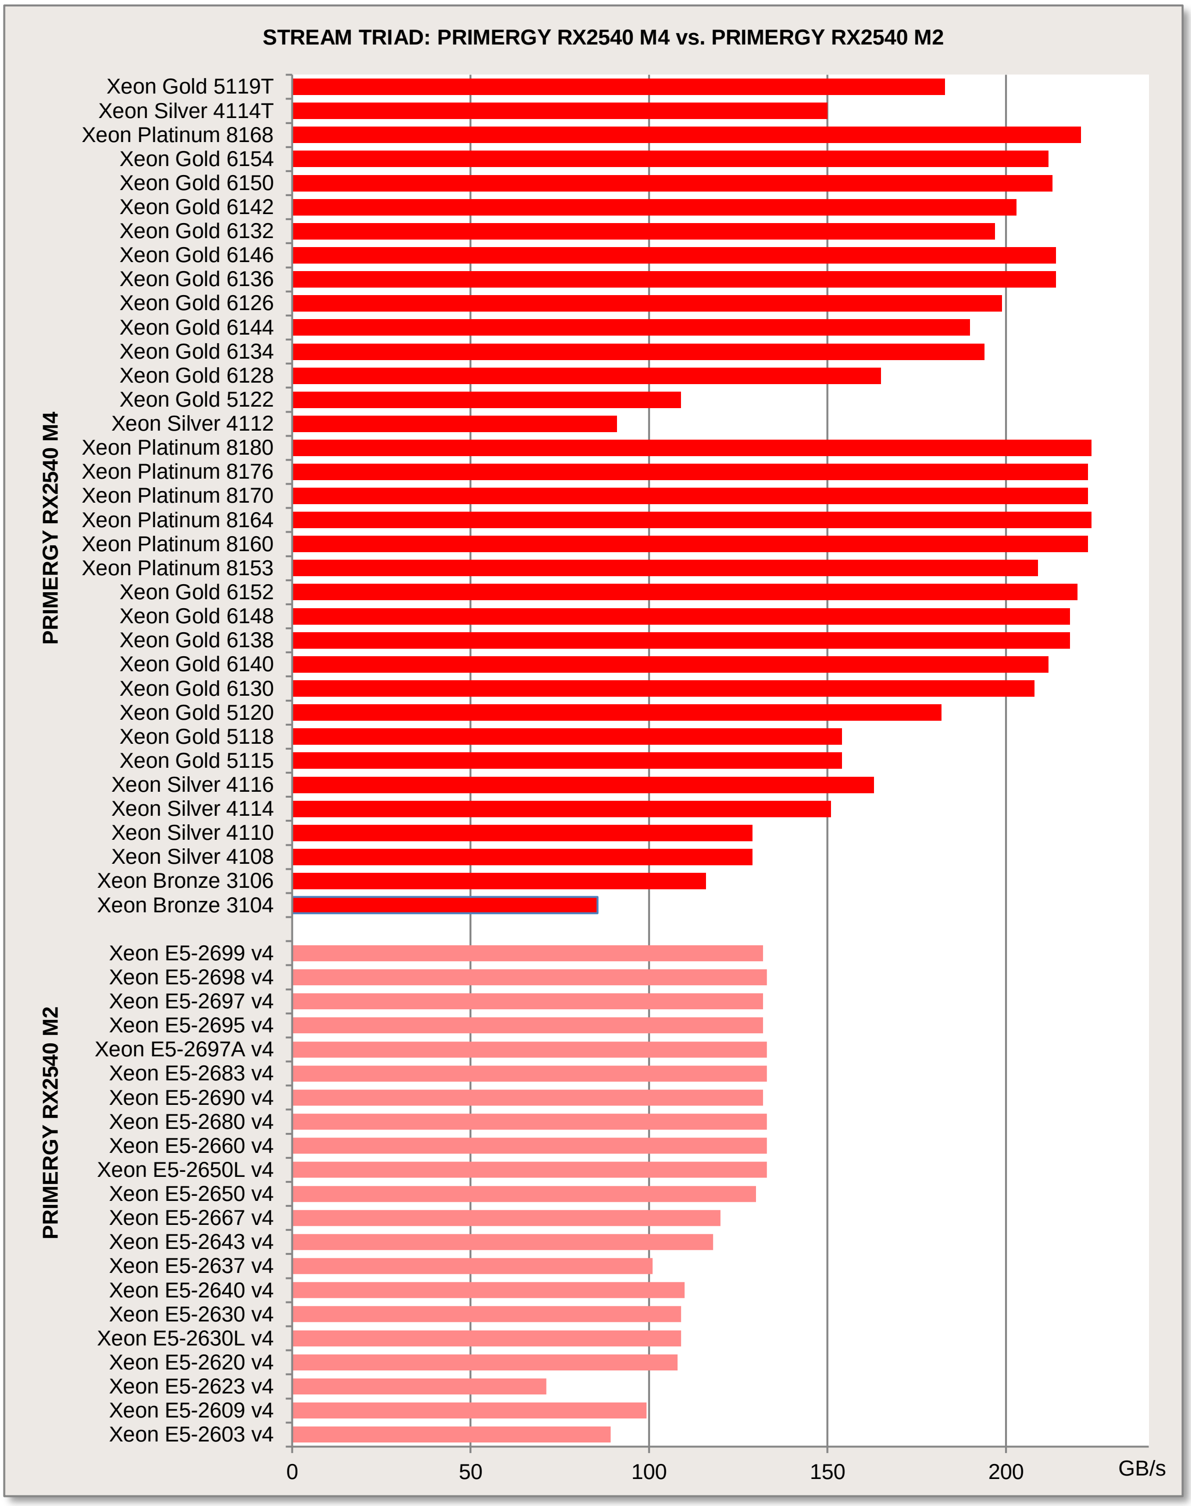

LINPACK

Benchmark description

LINPACK was developed in the 1970s by Jack Dongarra and some other people to show the performance of supercomputers. The benchmark consists of a collection of library functions for the analysis and solution of linear system of equations. A description can be found in the document http://www.netlib.org/utk/people/JackDongarra/PAPERS/hplpaper.pdf.

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

A memory of $8n^2$ bytes is required for the matrix. In case of an n × n matrix the number of arithmetic operations required for the solution is $^2/_3n^3 + 2n^2$. Thus, the choice of n determines the duration of the measurement: a doubling of n results in an approximately eight-fold increase in the duration of the measurement. The size of n also has an influence on the measurement result itself. As n increases, the measured value asymptotically approaches a limit. The size of the matrix is therefore usually adapted to the amount of memory available. Furthermore, the memory bandwidth of the system only plays a minor role for the measurement result, but a role that cannot be fully ignored. The processor performance is the decisive factor for the measurement result. Since the algorithm used permits parallel processing, in particular the number of processors used and their processor cores are - in addition to the clock rate - of outstanding significance.

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

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

Rpeak = Maximum number of floating point operations per clock cycle
× Number of processor cores of the computer
× Rated processor frequency [GHz]

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

Manufacturer-independent publication of LINPACK results is possible at http://www.top500.org/. The use of a LINPACK version based on HPL is prerequisite for this (see http://www.netlib.org/benchmark/hpl/).

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

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

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX2540 M4
Processor	Intel® Xeon® Processor Scalable Family × 2
Memory	16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC × 24
Software	
BIOS settings	HyperThreading = Disabled Link Frequency Select = 10.4 GT/s HWPM Support = Disabled Intel Virtualization Technology = Disabled Sub NUMA Clustering = Disabled IMC Interleaving = 1-way LLC Dead Line Alloc = Disabled Stale AtoS = Enabled
Operating system	SUSE Linux Enterprise Server 12 SP2 (x86_64)
Operating system settings	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 aio-max-nr = 1048576 ulimit -s unlimited nohz_full = 1-xx Xeon Platinum 8180 : run with avx512 Xeon Silver 4116 : run with avx2
Benchmark	MPI version: Intel® Math Kernel Library Benchmarks for Linux OS (l_mklb_p_2017.3.017)

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

Benchmark results

This results in italic are estimated values from the results of RX2530 M4.

Processor	Cores	Processor Frequency [GHz]	Number of Processors	Rpeak [GFlops]	Rmax [GFlops]	Efficiency
Xeon Bronze 3104	6	1.7	2	326	***238***	73%
Xeon Bronze 3106	8	1.7	2	435	***318***	73%
Xeon Silver 4108	8	1.8	2	461	***298***	65%
Xeon Silver 4110	8	2.1	2	538	***512***	95%
Xeon Silver 4114	10	2.2	2	704	***670***	95%
Xeon Silver 4116	12	2.1	2	806	**765**	95%
Xeon Gold 5115	10	2.4	2	768	***683***	89%
Xeon Gold 5118	12	2.3	2	883	***838***	95%
Xeon Gold 5120	14	2.2	2	986	***702***	71%
Xeon Gold 6130	16	2.1	2	2150	***1810***	84%
Xeon Gold 6140	18	2.3	2	2650	***2020***	76%
Xeon Gold 6138	20	2.0	2	2560	***1930***	75%
Xeon Gold 6148	20	2.4	2	3072	***2210***	72%
Xeon Gold 6152	22	2.1	2	2957	***2180***	74%
Xeon Platinum 8153	16	2.0	2	2048	***1546***	75%
Xeon Platinum 8160	24	2.1	2	3226	***2370***	73%
Xeon Platinum 8164	26	2.0	2	3328	***2474***	74%
Xeon Platinum 8170	26	2.1	2	3494	***2722***	78%
Xeon Platinum 8176	28	2.1	2	3763	***2779***	74%
Xeon Platinum 8180	28	2.5	2	4480	**3409**	74%
Xeon Silver 4112	4	2.6	2	333	***315***	95%
Xeon Gold 5122	4	3.6	2	922	***736***	80%
Xeon Gold 6128	6	3.4	2	1306	***990***	76%
Xeon Gold 6134	8	3.2	2	1638	***1270***	78%
Xeon Gold 6144	8	3.5	2	1792	***1300***	73%
Xeon Gold 6126	8	2.6	2	1997	***1560***	78%
Xeon Gold 6136	12	3.0	2	2304	***1780***	77%
Xeon Gold 6146	12	3.2	2	2458	***1880***	76%
Xeon Gold 6132	14	2.6	2	2330	***1890***	81%
Xeon Gold 6142	16	2.6	2	2662	***2090***	79%
Xeon Gold 6150	18	2.7	2	3110	***2240***	72%
Xeon Gold 6154	18	3.0	2	3456	***2700***	78%
Xeon Platinum 8168	24	2.7	2	4147	***2747***	66%
Xeon Silver 4114T	10	2.2	2	704	***670***	95%
Xeon Gold 5119T	14	1.9	2	851	***806***	95%
Xeon Gold 6134M	8	3.2	2	1638	***1270***	78%

Xeon Gold 6140M	18	2.3	2	2650	***2020***	76%
Xeon Gold 6142M	16	2.6	2	2662	***2090***	79%
Xeon Platinum 8160M	24	2.1	2	3226	***2370***	73%
Xeon Platinum 8170M	26	2.1	2	3494	***2722***	78%
Xeon Platinum 8176M	28	2.1	2	3763	***2779***	74%
Xeon Platinum 8180M	28	2.5	2	4480	***3409***	74%

Rmax = Measurement result

Rpeak = Maximum number of floating point operations per clock cycle
× Number of processor cores of the computer
× Rated frequency [GHz]

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

Literature

PRIMERGY Servers

http://primergy.com/

PRIMERGY RX2540 M4

This White Paper:

http://docs.ts.fujitsu.com/dl.aspx?id=2b079d6b-a1de-47d5-88e9-d4124a99dbff

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

Data sheet
http://docs.ts.fujitsu.com/dl.aspx?id=e6102f2f-76da-4673-909c-c1d191ce2b31

PRIMERGY Performance

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

OLTP-2

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

SAP SD

http://www.sap.com/benchmark

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

SPECcpu2006

http://www.spec.org/osg/cpu2006

Benchmark overview SPECcpu2006
http://docs.ts.fujitsu.com/dl.aspx?id=1a427c16-12bf-41b0-9ca3-4cc360ef14ce

SPECpower_ssj2008

http://www.spec.org/power_ssj2008

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

SPECjbb2015

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

STREAM

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

vServCon

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

VMmark V3

VMmark 3
http://www.vmmark.com

LINPACK

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

TOP500
http://www.top500.org/

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

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

Contact

FUJITSU
Website: http://www.fujitsu.com/
PRIMERGY Product Marketing
mailto:Primergy-PM@ts.fujitsu.com
PRIMERGY Performance and Benchmarks
mailto:primergy.benchmark@ts.fujitsu.com