White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY RX2530 M4

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY RX2530 M4

The PRIMERGY RX2530 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
1.1 2018/04/10

Contents

Document history

Version 1.0 (2017/10/31)

New:

- Technical data
- SPECcpu2006
  Measurements with Intel® Xeon® Processor Scalable Family
- 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 (2018/04/10)

Updated:

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

Technical data

PRIMERGY RX2530 M4
PY RX2530 M4 4x 3.5'

PRIMERGY RX2530 M4
PY RX2530 M4 10x 2.5'

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

Model	PRIMERGY RX2530 M4
Model versions	PY RX2530 M4 4x 3.5' expandable PY RX2530 M4 2x 2.5' expandable PY RX2530 M4 10x 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 × 1 PCI-Express 3.0 x16 × 3
Max. number of internal hard disks	PY RX2530 M4 4x 3.5' : 4 PY RX2530 M4 8x 2.5' expandable: 8 PY RX2530 M4 10 x 2.5': 10

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	3.7	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	3.7	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 RX2530 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 RX2530 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 RX2530 M4
Processor	Intel® Xeon® Processor Scalable Family × 2
Memory	16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC × 24
Software	
BIOS settings	Xeon Platinum 81xx, Gold 61xx, Gold 5120: 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 Gold 51xx (5120 を除く), HWPM Support = Disabled Intel Virtualization Technology = Disabled Sub NUMA Clustering = Disabled IMC Interleaving = 2-way LLC Dead Line Alloc = Disabled Stale AtoS = Enabled Link Frequency Select = 10.4 GT/s Xeon Silver 41xx HWPM Support = Disabled Intel Virtualization Technology = Disabled Sub NUMA Clustering = Disabled IMC Interleaving = 2-way LLC Dead Line Alloc = Disabled Stale AtoS = Enabled Xeon Bronze 31xx 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	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.

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		
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	1560		
Xeon Gold 6140	2	1760		
Xeon Gold 6138	2	1770		
Xeon Gold 6148	2	1940		
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	2670	2780	*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	*2670*	*2780*	

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	1470		
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	1770	1800	
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	*1770*	*1800*	

The following two diagrams illustrate the throughput of the PRIMERGY RX2530 M4 in comparison to its predecessor PRIMERGY RX2530 M2, in their respective most performant configuration.

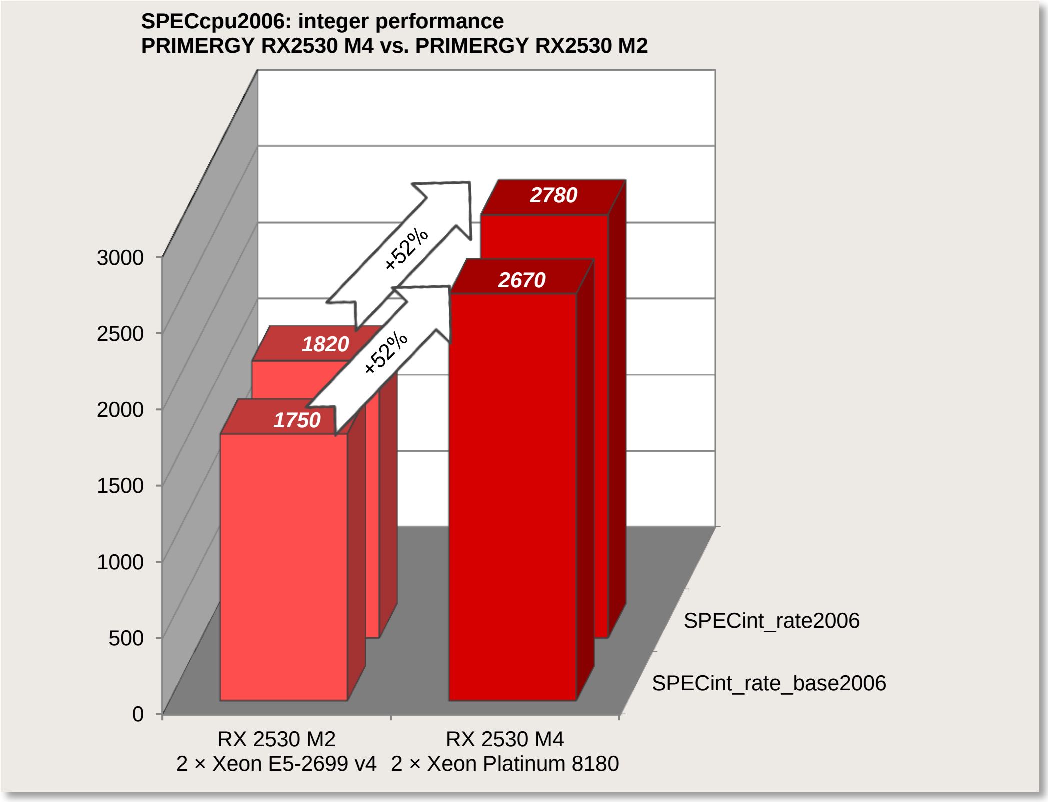

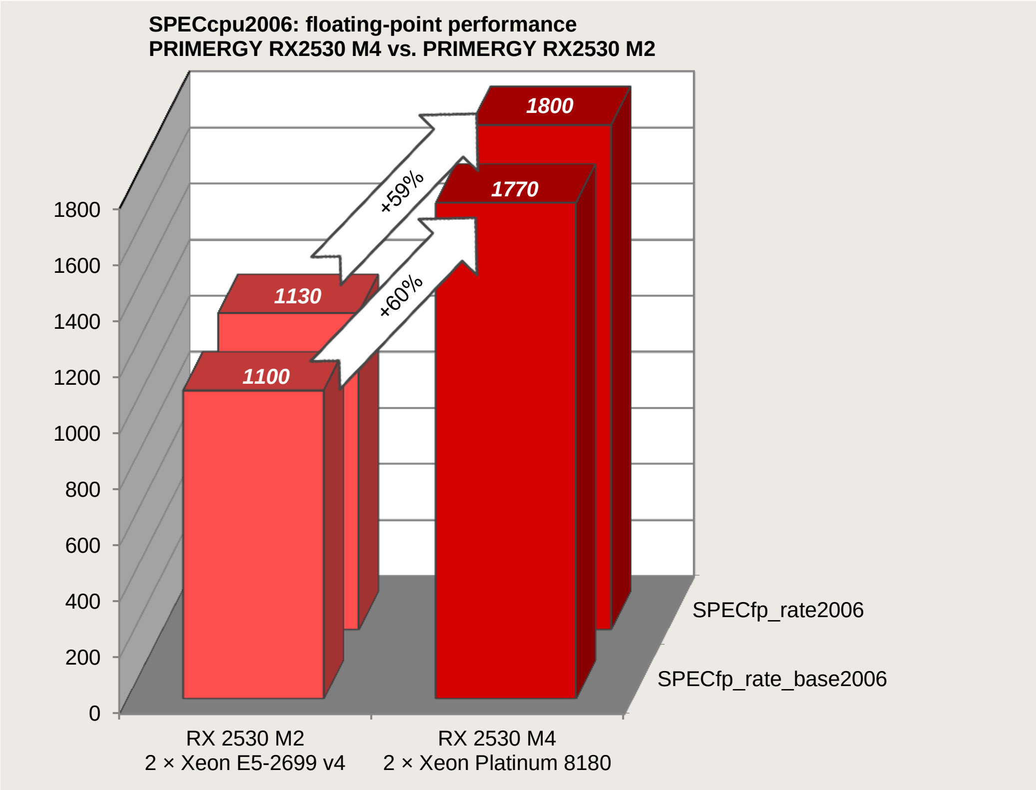

OLTP-2

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:

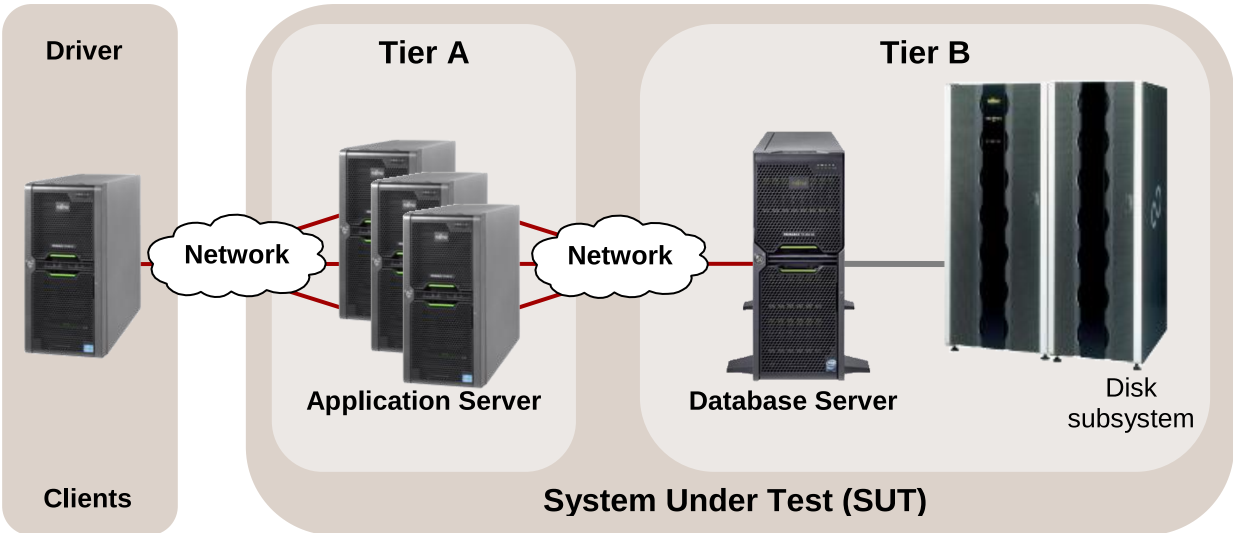

All results were determined by way of example on a PRIMERGY RX2540 M4.

Database Server (Tier B)	
Hardware	
Model	PRIMERGY RX2530 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	RX2530 M4: Onboard RAID controller PRAID EP420i 2 × 300 GB 10k rpm SAS Drive, RAID1 (OS), 4 × 600 GB 10k rpm SAS Drive, RAID10 (LOG) 2 × 1.2 TB 10k rpm SAS Drive, RAID1 (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..

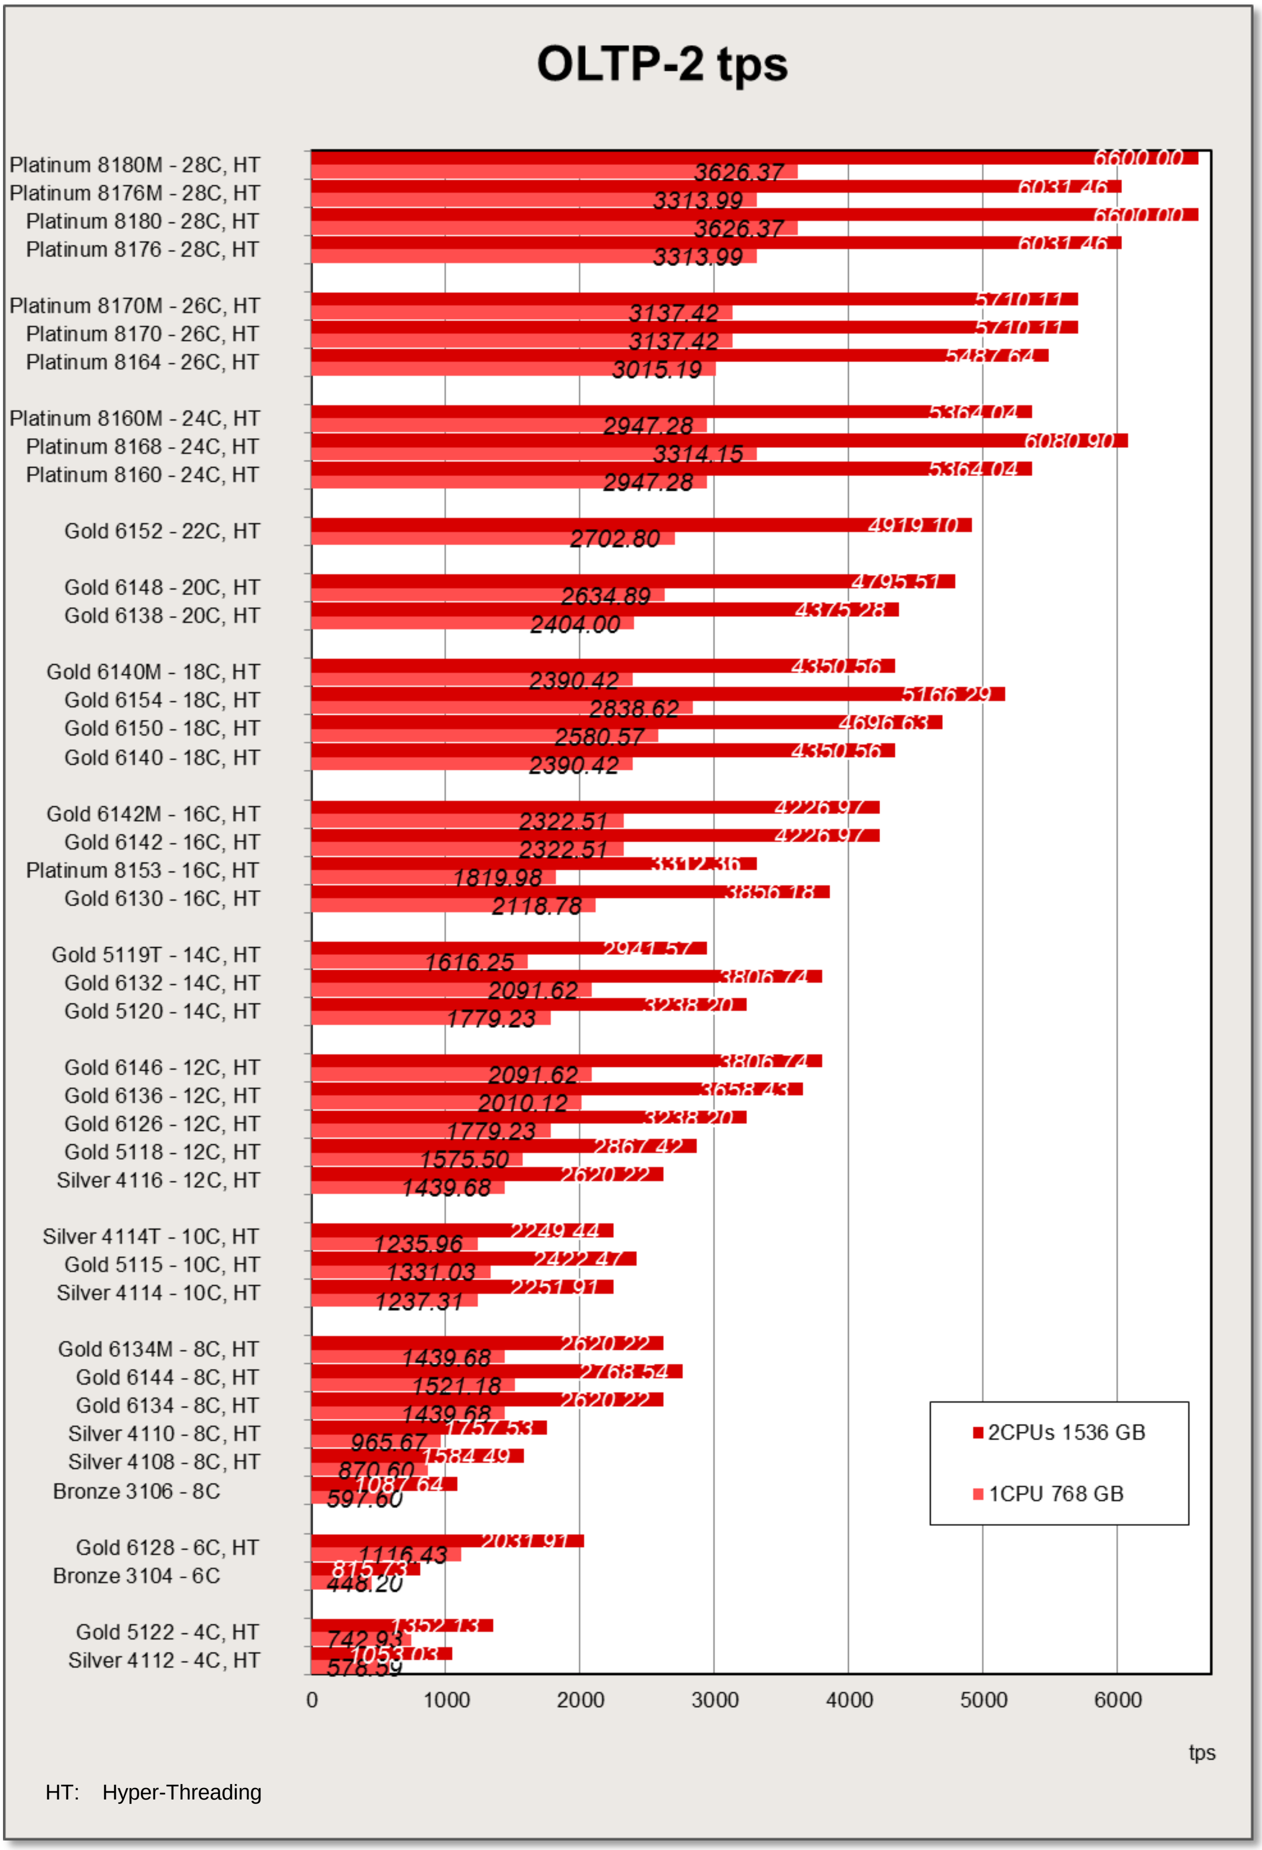
OLTP-2 tps
Platinum 8180M - 28C, HT 6600.00 3626.37
Platinum 8176M - 28C, HT 6031.46 3313.99
Platinum 8180 - 28C, HT 6600.00 3626.37
Platinum 8176 - 28C, HT 6031.46 3313.99
Platinum 8170M - 26C, HT 5710.11 3137.42
Platinum 8170 - 26C, HT 5710.11 3137.42
Platinum 8164 - 26C, HT 5487.64 3015.19
Platinum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

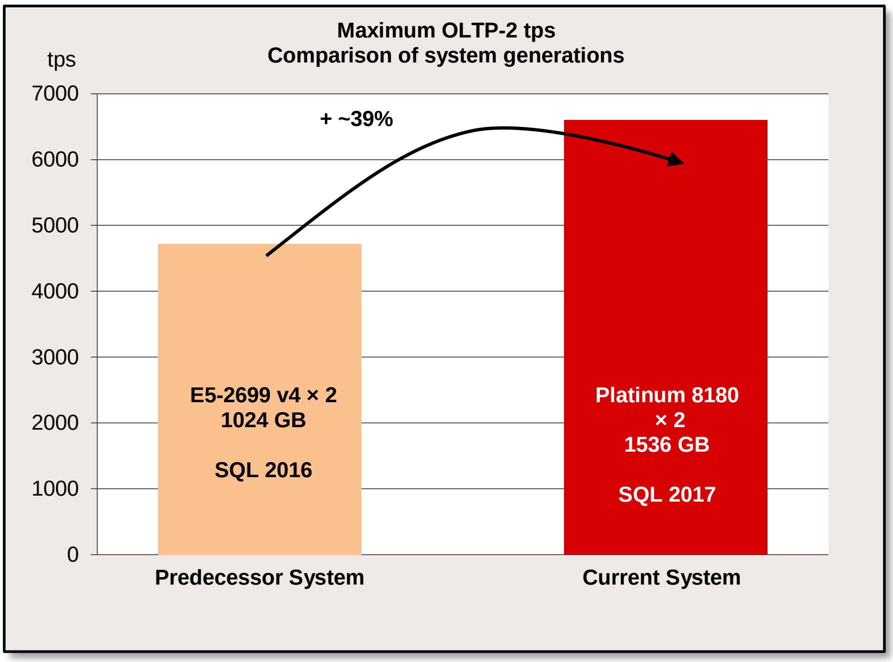

Current System	RX2530 M4	RX2540 M4
Predecessor System	RX2530 M2	RX2540 M2

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.

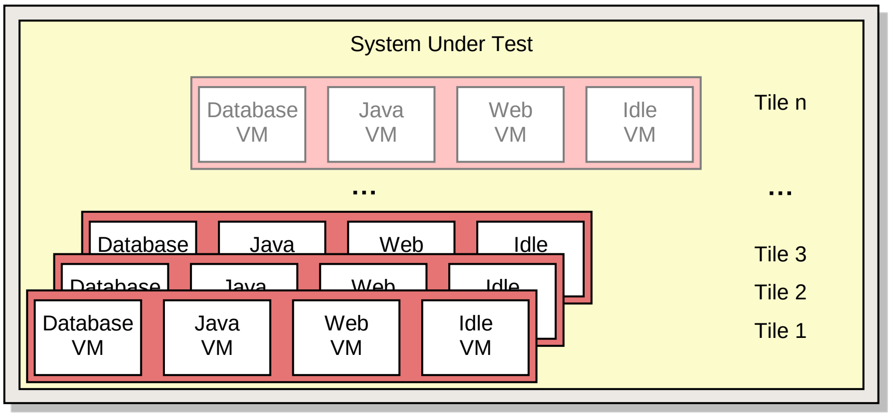

Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score, the individual benchmark result for one tile is put in relation to the respective result of a reference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

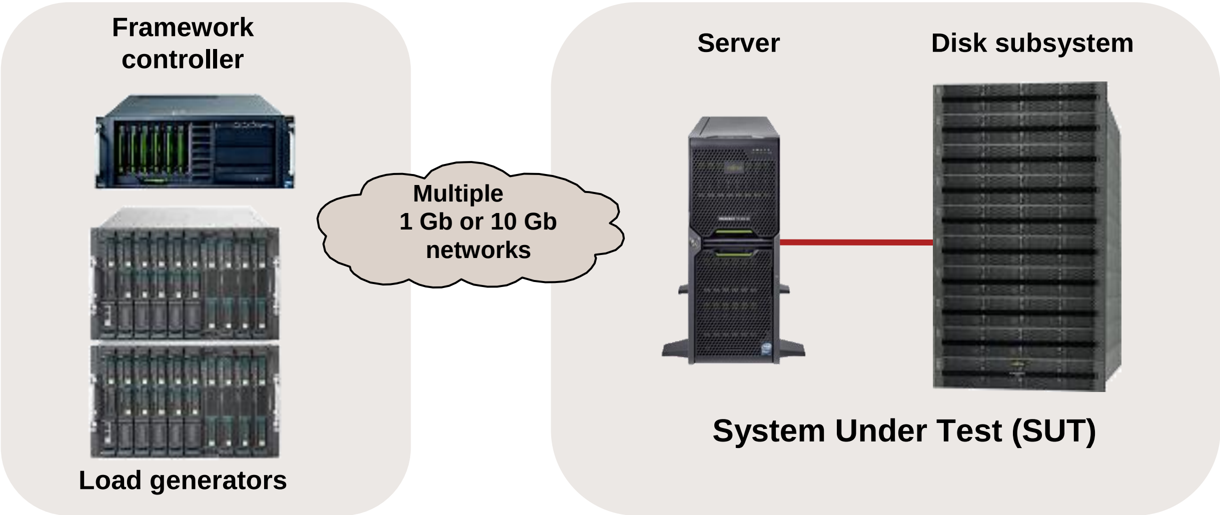

All results were determined by way of example on a PRIMERGY RX2530 M4.

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 32bit

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	**42.5**	**23**
	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	*47.4*	*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.

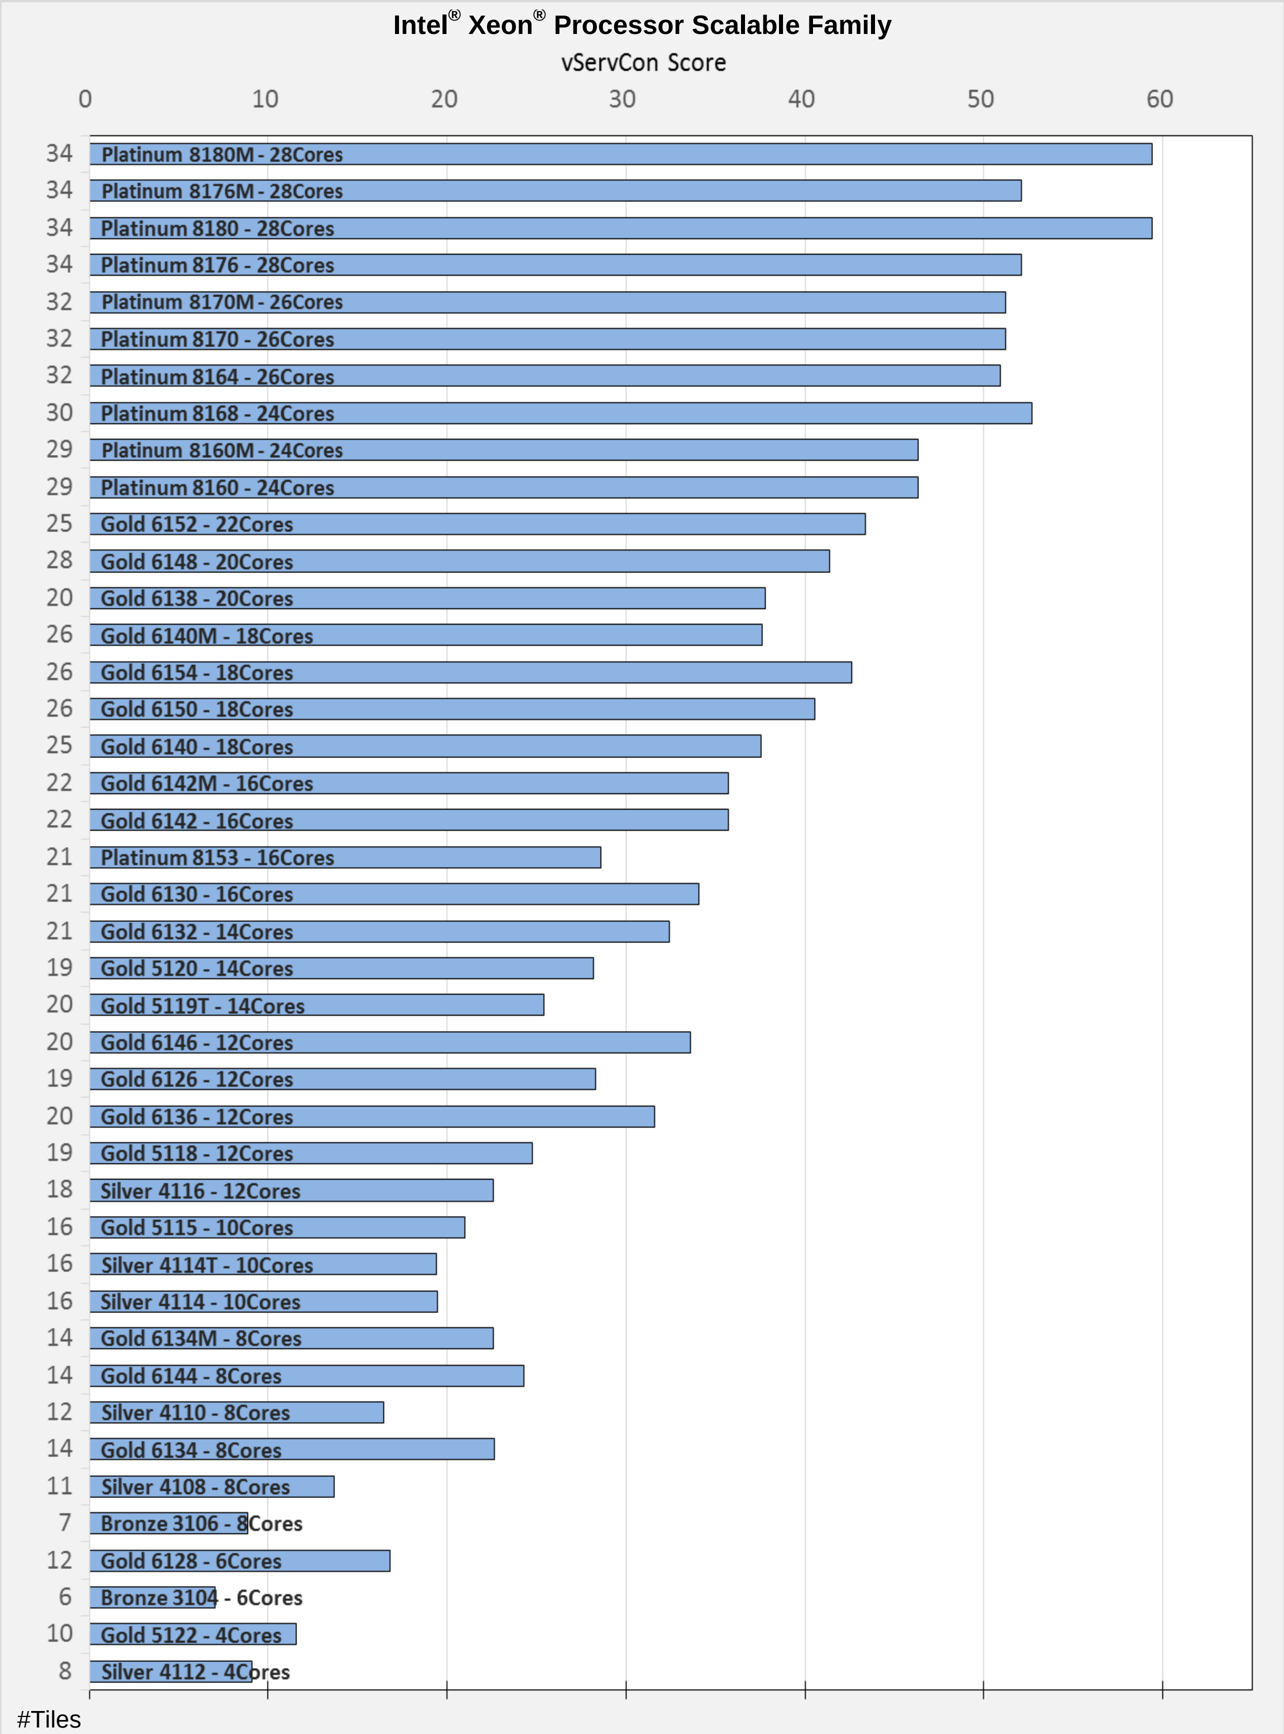

The relatively large performance differences between the processors can be explained by their features. The values scale on the basis of the number of cores, the size of the L3 cache and the CPU clock frequency and as a result of the features of Hyper-Threading and turbo mode, which are available in 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.

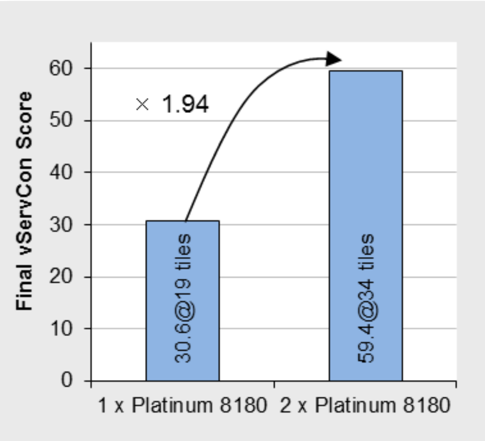

Until now, we have looked at the virtualization performance of a fully configured system. However, with a server with two sockets, the question also arises as to how good performance scaling is from one to two processors. The better the scaling, the lower the overhead usually caused by the shared use 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 Gold 6130 (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.

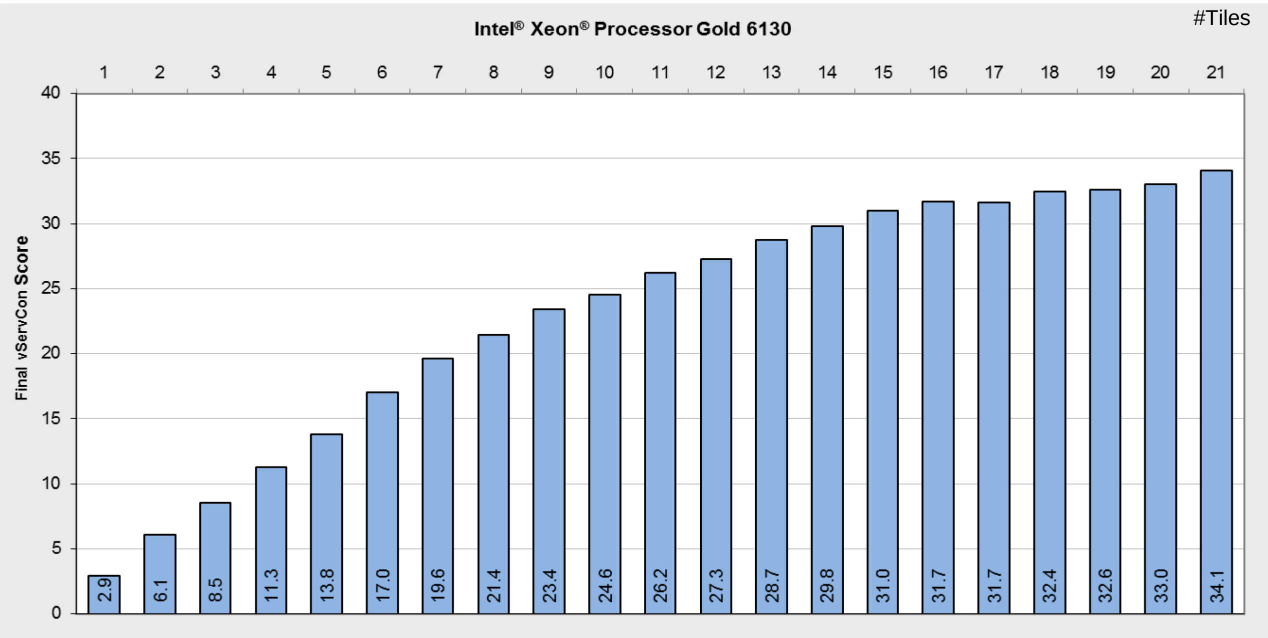

The previous diagram examined the total performance of all application VMs of a host. However, studying the performance from an individual application VM viewpoint is also interesting. This information is in the previous diagram. For example, the total optimum is reached in the above Xeon Gold 6130 situation 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).

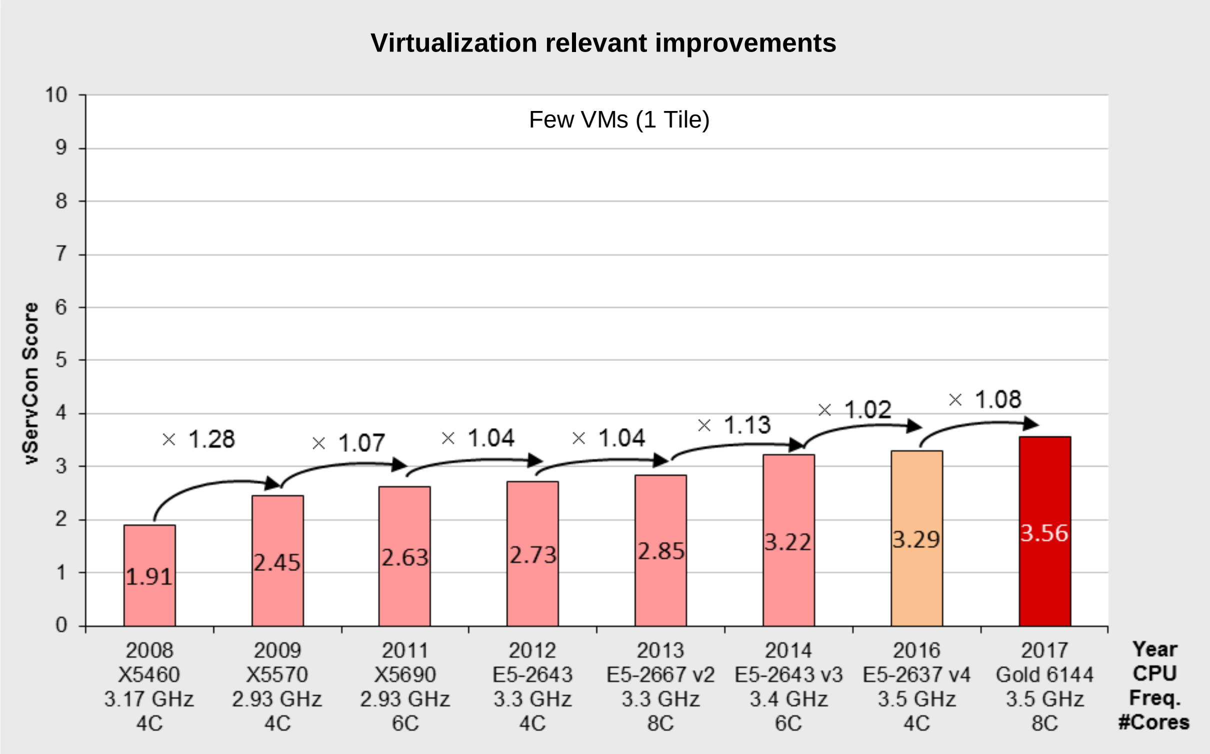

With full utilization of the systems with VMs there was an increase by a factor of 2.07. The one reason was the performance increase that could be achieved for an individual VM (see score for a few VMs). The other reason was that more VMs were possible with total optimum (via Hyper-Threading). However,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

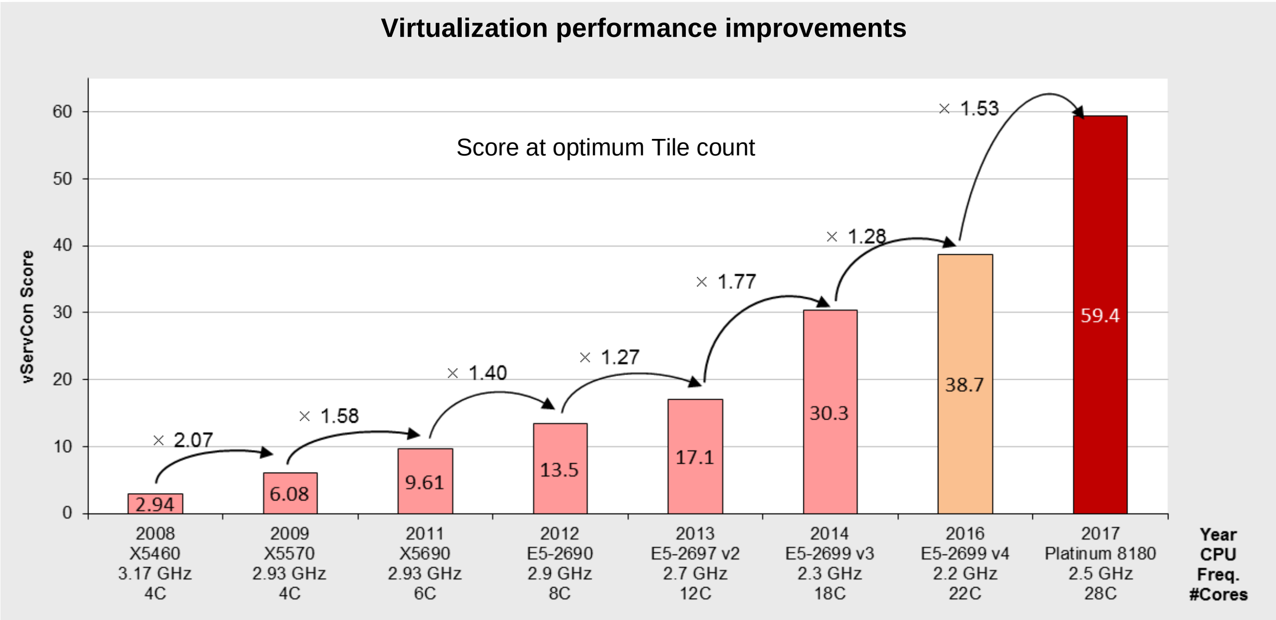

[1] EPT accelerates memory virtualization via hardware support for the mapping between host and guest memory addresses.

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.

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 RX2530 M4
Processor	Intel® Xeon® Processor Scalable Family × 2
Memory	16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC × 24
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 81xx, Gold 61xx, Gold 5122 : run with avx512 All others : 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.

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	223
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	*223*

The following diagram illustrates the throughput of the PRIMERGY RX2530 M4 in comparison to its predecessor, the PRIMERGY RX2530 M2.

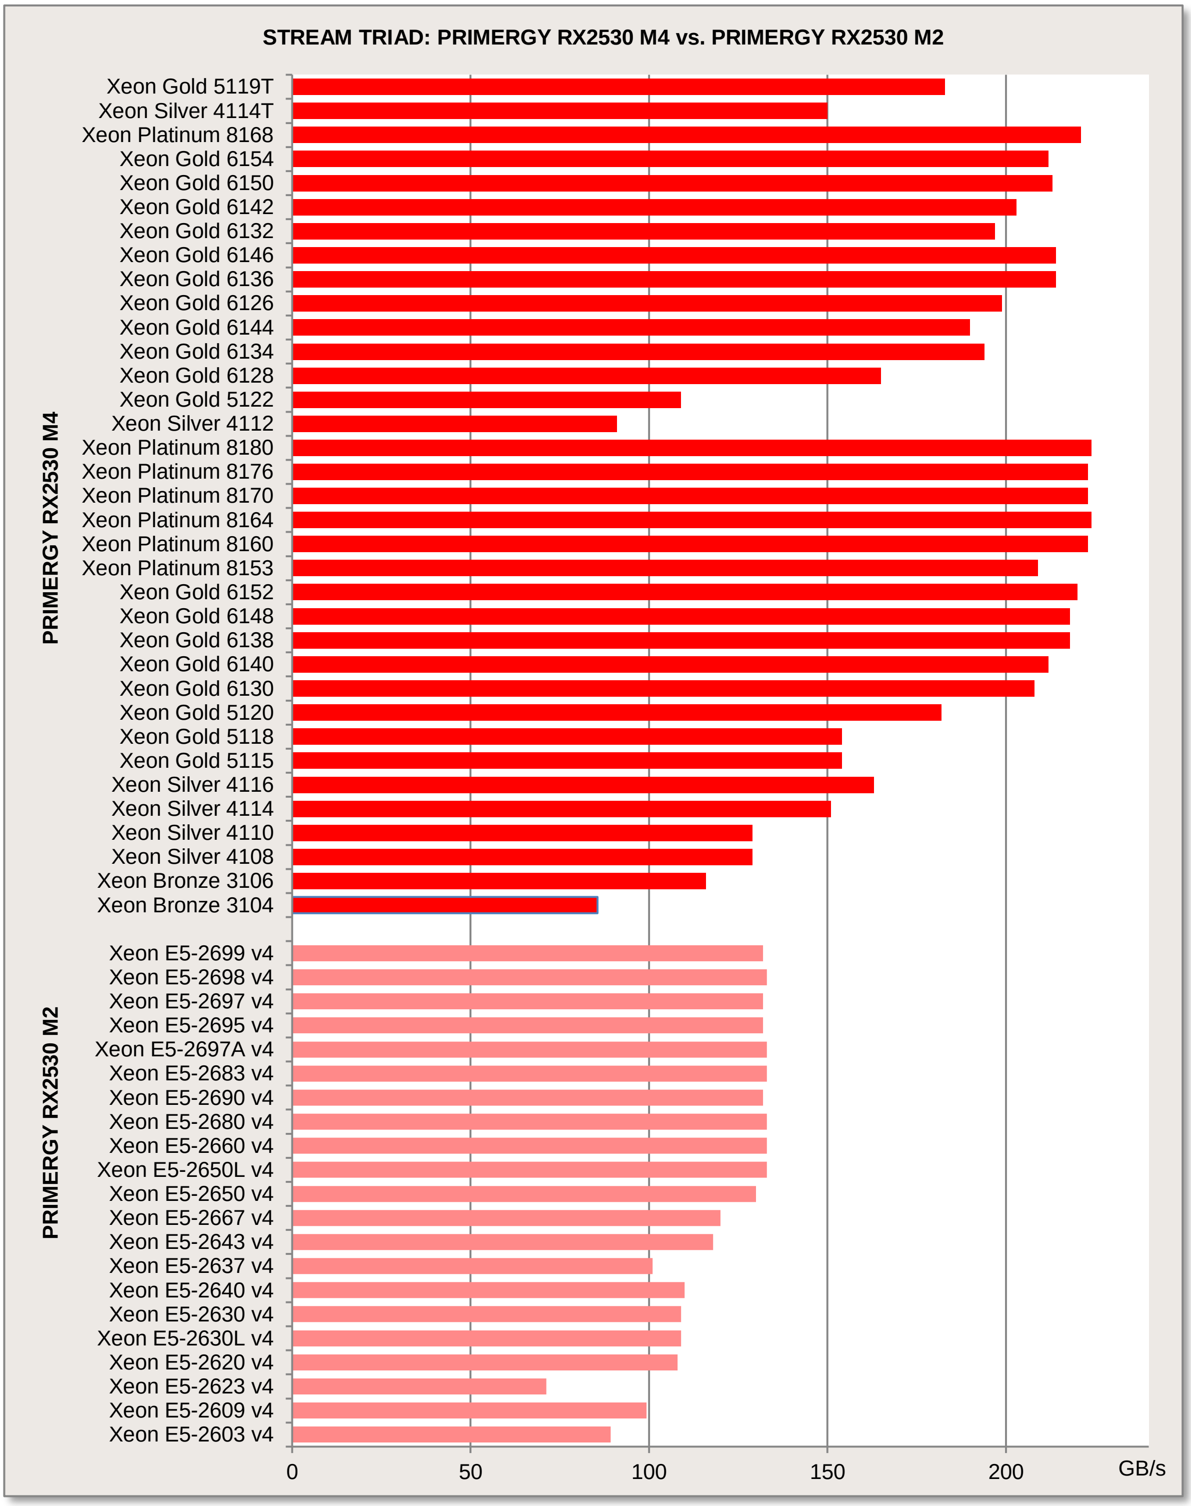

LINPACK

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 RX2530 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 81xx, Gold 61xx, Gold 5122 : run with avx512 All others : 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.

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	**3335**	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	12	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	***3335***	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 RX2530 M4

This White Paper:

- http://docs.ts.fujitsu.com/dl.aspx?id=89a24c07-02d3-44f7-8539-e719ac42cb5f
- http://docs.ts.fujitsu.com/dl.aspx?id=7491b025-dbf5-485c-bfa2-2ddd6fa08c28

Data sheet
http://docs.ts.fujitsu.com/dl.aspx?id=3dc20cf0-3801-4464-b5ea-570510e5f105

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

STREAM

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

vServCon

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

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