FUJITSU

White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY CX2550 M4/ CX2560 M4/ CX2570 M4

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY CX2550 M4/ CX2560 M4/ CX2570 M4.

The PRIMERGY CX2550 M4/ CX2560 M4/ CX2570 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
2018/10/23

Contents

Document history

Version 1.0 (2018/10/23)

New:

- Technical data
- SPECcpu2006
  Measurements with and results for Intel® Xeon® Processor Scalable Family
- OLTP-2
  Results for Intel® Xeon® Processor Scalable Family
- vServCon
  Results for Intel® Xeon® Processor Scalable Family
- STREAM
  Measurements with and results for Intel® Xeon® Processor Scalable Family
- LINPACK
  Measurements with and results for Intel® Xeon® Processor Scalable Family

Version 1.1 (2018/10/23)

New:

- VMmark V3
  "Performance Only" measurement with Intel® Xeon® Platinum 8164

Updated:

- SPECcpu2006
  Additional measurements with and results for Intel® Xeon® Processor Scalable Family

Technical data

PRIMERGY CX2550 M4/ CX2560 M4

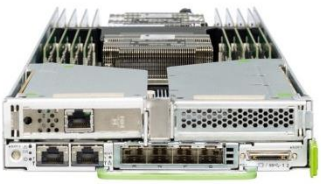

PRIMERGY CX2570 M4

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.

<table>
<tr><th>Model</th><th colspan="2">PRIMERGY CX2550 M4</th><th>PRIMERGY CX2560M4</th><th colspan="2">PRIMERGY CX2570 M4</th></tr>
<tr><td>cooling method</td><td>air cooling</td><td>liquid cooling</td><td>air cooling</td><td>air cooling</td><td>liquid cooling</td></tr>
<tr><td>Form factor</td><td colspan="5">Server node</td></tr>
<tr><td>Chipset</td><td colspan="2">Intel® C621</td><td colspan="3">Intel® C624</td></tr>
<tr><td>Number of sockets</td><td colspan="5">2</td></tr>
<tr><td>Number of processors orderable</td><td colspan="5">1 - 2</td></tr>
<tr><td>Processor type</td><td colspan="5">Intel® Xeon® Processor Scalable Family</td></tr>
<tr><td>Number of memory slots</td><td colspan="5">16 (8 per processor)</td></tr>
<tr><td>Maximum memory configuration</td><td colspan="5">2048 GB</td></tr>
<tr><td>PCI slots</td><td>PCI-Express 3.0 x16 × 2</td><td>PCI-Express 3.0 x16 × 1</td><td>PCI-Express 3.0 x16 × 3</td><td colspan="2">PCI-Express 3.0 x16 × 2</td></tr>
</table>

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 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
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

This CPUs in italic are only supported by liquid cooling model.

Model	PRIMERGY CX2550 M4	PRIMERGY CX2560M4	PRIMERGY CX2570 M4
Supported Processors	Xeon Gold 6130 Xeon Gold 6140 Xeon Gold 6138 Xeon Gold 6148 Xeon Gold 6152 Xeon Platinum 8153 Xeon Platinum 8160 Xeon Platinum 8164 Xeon Platinum 8170 Xeon Platinum 8176 *Xeon Platinum 8180* Xeon Gold 6128 Xeon Gold 6134 *Xeon Gold 6144* Xeon Gold 6126 Xeon Gold 6136 Xeon Gold 6146 Xeon Gold 6132 Xeon Gold 6142 Xeon Gold 6150 *Xeon Gold 6154* *Xeon Platinum 8168* Xeon Gold 6134M Xeon Gold 6140M Xeon Gold 6142M Xeon Platinum 8160M Xeon Platinum 8170M Xeon Platinum 8176M *Xeon Platinum 8180M*	Xeon Bronze 3104 Xeon Bronze 3106 Xeon Silver 4108 Xeon Silver 4110 Xeon Silver 4114 Xeon Silver 4116 Xeon Gold 5115 Xeon Gold 5118 Xeon Gold 5120 Xeon Gold 6130 Xeon Gold 6140 Xeon Gold 6138 Xeon Gold 6148 Xeon Gold 6152 Xeon Platinum 8153 Xeon Platinum 8160 Xeon Platinum 8164 Xeon Silver 4112 Xeon Gold 5122 Xeon Gold 6128 Xeon Gold 6134 Xeon Gold 6126 Xeon Gold 6136 Xeon Gold 6132 Xeon Gold 6142 Xeon Gold 6134M Xeon Gold 6140M Xeon Gold 6142M Xeon Platinum 8160M Xeon Silver 4114T Xeon Gold 5119T	Xeon Bronze 3104 Xeon Bronze 3106 Xeon Silver 4108 Xeon Silver 4110 Xeon Silver 4114 Xeon Silver 4116 Xeon Gold 5115 Xeon Gold 5118 Xeon Gold 5120 Xeon Gold 6130 Xeon Gold 6140 Xeon Gold 6138 Xeon Gold 6148 Xeon Gold 6152 Xeon Platinum 8153 Xeon Platinum 8160 Xeon Platinum 8164 *Xeon Platinum 8170* *Xeon Platinum 8176* *Xeon Platinum 8180* Xeon Silver 4112 Xeon Gold 5122 Xeon Gold 6128 Xeon Gold 6134 *Xeon Gold 6144* Xeon Gold 6126 Xeon Gold 6136 *Xeon Gold 6146* Xeon Gold 6132 Xeon Gold 6142 *Xeon Gold 6150* *Xeon Gold 6154* *Xeon Platinum 8168* Xeon Gold 6134M Xeon Gold 6140M Xeon Gold 6142M Xeon Platinum 8160M *Xeon Platinum 8170M* *Xeon Platinum 8176M* *Xeon Platinum 8180M* Xeon Silver 4114T Xeon Gold 5119T

All the processors that can be ordered with the PRIMERGY CX2550 M4/ CX2560 M4/ CX2570 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
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 (1x16GB) 1Rx4 DDR4-2666 R ECC	16	1	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			✓	✓

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

Detailed technical information is available in the data sheet PRIMERGY CX2550 M4/ CX2560 M4/ CX2570 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 CX2550 M4/ CX2560 M4/ CX2570 M4
Processor	Intel® Xeon® Processor Scalable Family × 2
Memory	16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC × 12
Software	
BIOS settings	Xeon Platinum 81xx, Gold 61xx: DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Power Technology = Custom HWPM Support = Disabled UPI Link Frequency Select = 10.4 GT/s Sub NUMA Clustering = Enabled Stale Atos = Enabled LLC dead line alloc = Disabled
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	Int_rate: C/C++: Version 18.0.0.128 of Intel C/C++ Compiler for Linux fp_rate: C/C++: Version 17.0.3.191 of Intel C/C++ Compiler for Linux 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 on 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.

Model	CX2550 M4		CX2560 M4			CX2570 M4	
Processor	Number of processors	SPECint_rate_base2006	Number of processors	SPECint_rate_base2006	SPECint_rate2006	Number of processors	SPECint_rate_base2006
Xeon Bronze 3104			2	327		2	*329*
Xeon Bronze 3106			2	434		2	*438*
Xeon Silver 4108			2	633		2	*638*
Xeon Silver 4110			2	718		2	*708*
Xeon Silver 4114			2	927		2	*907*
Xeon Silver 4116			2	1,040		2	*1056*
Xeon Gold 5115			2	998		2	*976*
Xeon Gold 5118			2	1,200		2	*1,155*
Xeon Gold 5120			2	1,330		2	*1,305*
Xeon Gold 6130	2	*1,554*	2	1,510		2	*1,554*
Xeon Gold 6140	2	*1,753*	2	1,750		2	*1,753*
Xeon Gold 6138	2	*1,763*	2	1,720		2	*1,763*
Xeon Gold 6148	2	*1,932*	2	2,010		2	*1,932*
Xeon Gold 6152	2	*1,982*	2	2,000		2	*1,982*
Xeon Platinum 8153	2	*1,364*	2	1,420		2	*1,364*
Xeon Platinum 8160	2	*2,161*	2	2,160		2	*2,161*
Xeon Platinum 8164	2	*2,211*	2	2,180		2	*2,211*
Xeon Platinum 8170	2	2,320				2	*2,320*
Xeon Platinum 8176	2	2,430				2	*2,430*
Xeon Platinum 8180	2	*2,659*				2	*2,659*
Xeon Silver 4112			2	430		2	*424*
Xeon Gold 5122			2	571		2	*545*
Xeon Gold 6128	2	*819*	2	838		2	*819*
Xeon Gold 6134	2	*1,056*	2	1,080		2	*1,056*
Xeon Gold 6144	2	1,115				2	*1,150*
Xeon Gold 6126	2	*1,305*	2	1,340	1,410	2	*1,305*
Xeon Gold 6136	2	*1,474*	2	1,470		2	*1,474*
Xeon Gold 6146	2	1,570				2	*1,570*

Xeon Gold 6132	2	*1,534*	2	1,510		2	*1,534*
Xeon Gold 6142	2	*1,703*	2	1,730		2	*1,703*
Xeon Gold 6150	2	1,940				2	*1,940*
Xeon Gold 6154	2	*2,081*				2	*2,081*
Xeon Platinum 8168	2	*2450*				2	*2,450*
Xeon Gold 6134M	2	*1,056*	2	*1,041*		2	*1,056*
Xeon Gold 6140M	2	*1,753*	2	*1,728*		2	*1,753*
Xeon Gold 6142M	2	*1,703*	2	*1,679*		2	*1,703*
Xeon Platinum 8160M	2	*2,161*	2	*2,131*		2	*2,161*
Xeon Platinum 8170M	2	*2,301*				2	*2,320*
Xeon Platinum 8176M	2	*2,430*				2	*2,430*
Xeon Platinum 8180M	2	*2,659*				2	*2,659*
Xeon Silver 4114T			2	901		2	*907*
Xeon Gold 5119T			2	1,200		2	*1,265*

Model	CX2550 M4		CX2560 M4			CX2570 M4	
Processor	Number of processors	SPECfp_rate_base2006	Number of processors	SPECfp_rate_base2006	SPECfp_rate2006	Number of processors	SPECfp_rate_base2006
Xeon Bronze 3104			2	353		2	*353*
Xeon Bronze 3106			2	464		2	*467*
Xeon Silver 4108			2	614		2	*621*
Xeon Silver 4110			2			2	*669*
Xeon Silver 4114			2	808		2	*813*
Xeon Silver 4116			2	907		2	*910*
Xeon Gold 5115			2	829		2	*848*
Xeon Gold 5118			2	963		2	*963*
Xeon Gold 5120			2	1,050		2	*1,048*
Xeon Gold 6130	2	*1,232*	2	1,190		2	*1,232*
Xeon Gold 6140	2	*1,339*	2	1,310		2	*1,339*
Xeon Gold 6138	2	*1,349*	2	1,290		2	*1,349*
Xeon Gold 6148	2	*1,426*	2	1,440		2	*1,426*
Xeon Gold 6152	2	*1,446*	2	1,440		2	*1,446*
Xeon Platinum 8153	2	*1,135*	2	1,140		2	*1,135*
Xeon Platinum 8160	2	*1,514*	2	1,510		2	*1,514*
Xeon Platinum 8164	2	*1,543*	2	1,520		2	*1,543*
Xeon Platinum 8170	2	1,580				2	*1,580*
Xeon Platinum 8176	2	1,630				2	*1,630*
Xeon Platinum 8180	2	*1,717*				2	*1,717*
Xeon Silver 4112			2	415		2	*417*
Xeon Gold 5122			2	538		2	*518*
Xeon Gold 6128	2	*757*	2	755		2	*757*
Xeon Gold 6134	2	*939*	2	929		2	*939*
Xeon Gold 6144	2	966				2	*966*
Xeon Gold 6126	2	*1,096*	2	1,080	1,100	2	*1,096*
Xeon Gold 6136	2	*1,193*	2	1,170		2	*1,193*
Xeon Gold 6146	2	1,220				2	*1,220*
Xeon Gold 6132	2	*1,213*	2	1,170		2	*1,213*
Xeon Gold 6142	2	*1,310*	2	1,290		2	*1,310*
Xeon Gold 6150	2	1,400				2	*1,400*
Xeon Gold 6154	2	*1,475*				2	*1,475*
Xeon Platinum 8168	2	*1,600*				2	*1,600*
Xeon Gold 6134M	2	*939*	2	*925*		2	*939*

Xeon Gold 6140M	2	*1,339*	2	*1,319*		2	*1,339*
Xeon Gold 6142M	2	*1,310*	2	*1,291*		2	*1,310*
Xeon Platinum 8160M	2	*1,514*	2	*1,491*		2	*1,514*
Xeon Platinum 8170M	2	*1,581*				2	*1,581*
Xeon Platinum 8176M	2	*1,630*				2	*1,630*
Xeon Platinum 8180M	2	*1,717*				2	*1,717*
Xeon Silver 4114T			2	810		2	*813*
Xeon Gold 5119T			2	975		2	*1,007*

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:

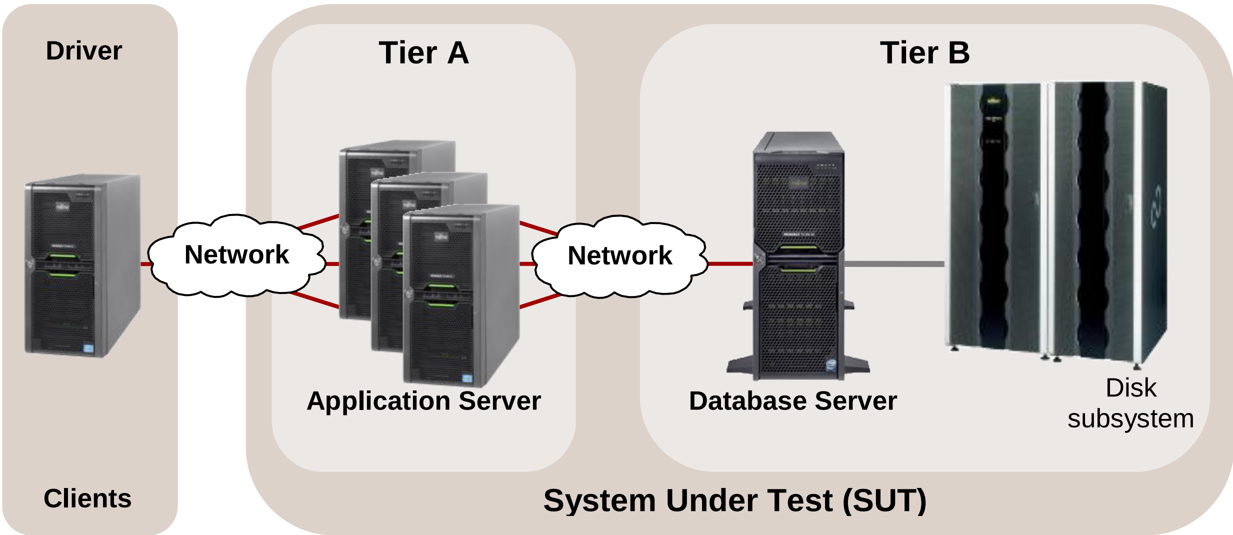

All results were determined by way of example on a PRIMERGY RX2540 M4

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, 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..

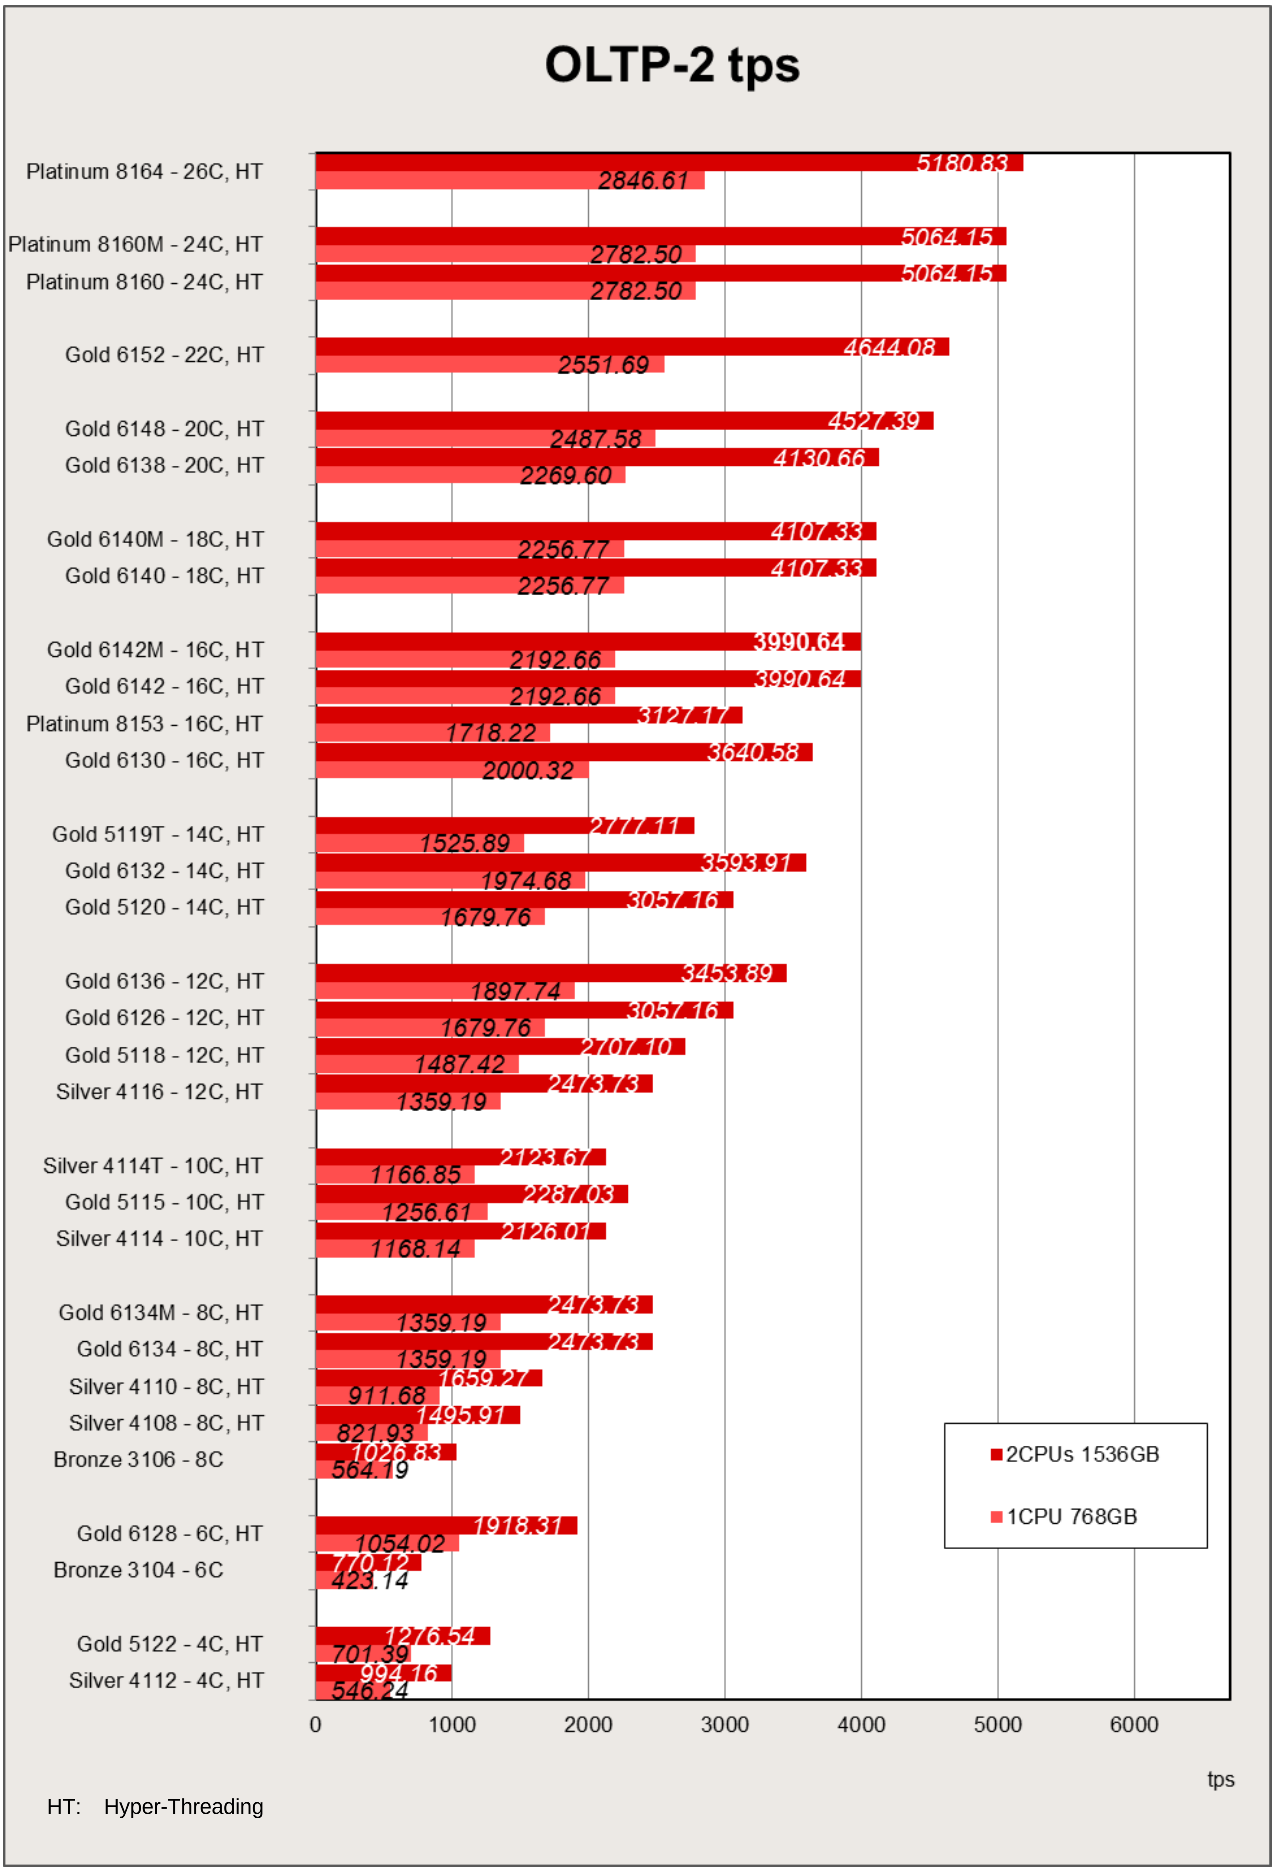
OLTP-2 tps
Platinum 8164 - 26C, HT
5180.83
2846.61
Platinum 8160M - 24C, HT
5064.15
2782.50
Platinum 8160 - 24C, HT
5064.15
2782.50
Gold 6152 - 22C, HT
4644.08
2551.69
Gold 6148 - 20C, HT
4527.39
2487.58
Gold 6138 - 20C, HT
4130.66
2269.60
Gold 6140M - 18C, HT
4107.33
2256.77
Gold 6140 - 18C, HT
4107.33
2256.77
Gold 6142M - 16C, HT
3990.64
2192.66
Gold 6142 - 16C, HT
3990.64
2192.66
Platinum 8153 - 16C, HT
3127.17
1718.22
Gold 6130 - 16C, HT
3640.58
2000.32
Gold 5119T - 14C, HT
2777.11
1525.89
Gold 6132 - 14C, HT
3593.91
1974.68
Gold 5120 - 14C, HT
3057.16
1679.76
Gold 6136 - 12C, HT
3453.89
1897.74
Gold 6126 - 12C, HT
3057.16
1679.76
Gold 5118 - 12C, HT
2707.10
1487.42
Silver 4116 - 12C, HT
2473.73
1359.19
Silver 4114T - 10C, HT
2123.67
1166.85
Gold 5115 - 10C, HT
2287.03
1256.61
Silver 4114 - 10C, HT
2126.01
1168.14
Gold 6134M - 8C, HT
2473.73
1359.19
Gold 6134 - 8C, HT
2473.73
1359.19
Silver 4110 - 8C, HT
1659.27
911.68
Silver 4108 - 8C, HT
1495.91
821.93
Bronze 3106 - 8C
1026.83
564.19
Gold 6128 - 6C, HT
1918.31
1054.02
Bronze 3104 - 6C
770.12
423.14
Gold 5122 - 4C, HT
1276.54
701.39
Silver 4112 - 4C, HT
994.16
546.24
0
1000
2000
3000
4000
5000
6000
tps
2CPUs 1536GB
1CPU 768GB
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.

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.

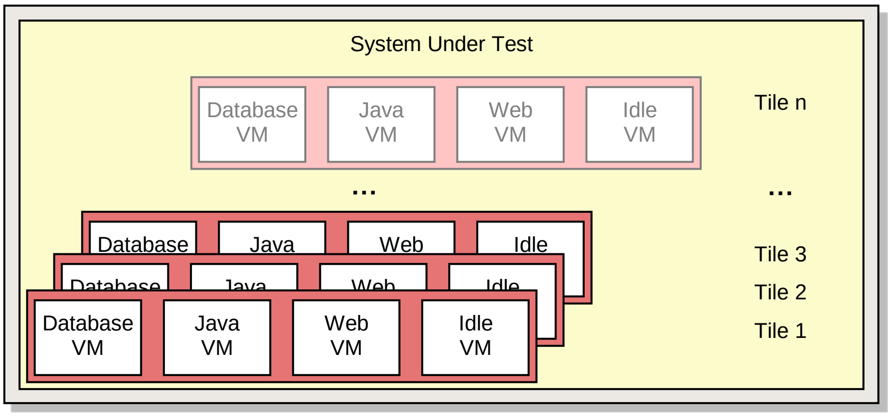

Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score, the individual benchmark result for one tile is put in relation to the respective result of a reference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

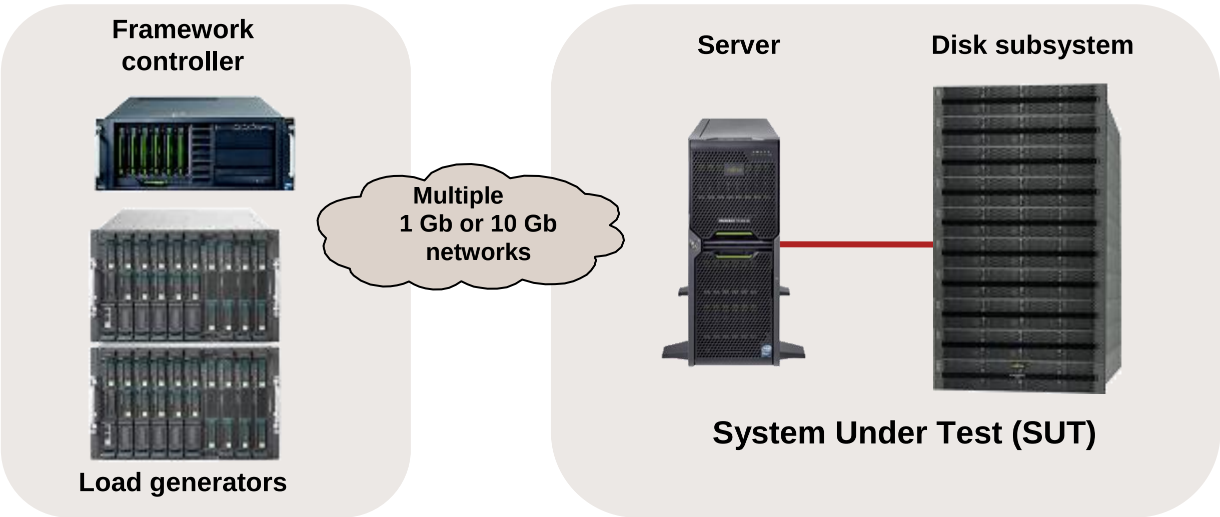

All results were determined by way of example on a PRIMERGY RX2530 M4.

System Under Test (SUT)	
Hardware	
Processor	2 × Intel® Xeon® Processor Scalable Family
Memory	24 × 16 GB (1x16 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 x 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 (per tile 3 load generator VMs on various server blades)	
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.

Processor			Score	#Tiles
Intel® Xeon® Processor Scalable Family	4 Cores Hyper-Threading, Turbo-Modus	Silver 4112	8.91	8
		Gold 5122	11.3	10
	6 Cores	Bronze 3104	6.91	6
	6 Cores Hyper-Threading, Turbo-Modus	Gold 6128	16.3	12
	8 Cores	Bronze 3106	8.71	7
	8 Cores Hyper-Threading, Turbo-Modus	Silver 4108	13.3	11
		Silver 4110	16.0	12
		Gold 6134	22.0	14
		Gold 6144	23.6	14
		Gold 6134M	22.0	14
	10 Cores Hyper-Threading, Turbo-Modus	Silver 4114	18.9	16
		Gold 5115	20.4	16
		Silver 4114T	18.9	16
	12 Cores Hyper-Threading, Turbo-Modus	Silver 4116	22.0	18
		Gold 5118	24.1	19
		Gold 6126	27.2	19
		Gold 6136	30.7	20
		Gold 6146	32.7	20
	14 Cores Hyper-Threading, Turbo-Modus	Gold 5120	27.2	21
		Gold 6132	31.5	21
		Gold 5119T	24.7	20
	16 Cores Hyper-Threading, Turbo-Modus	Gold 6130	33.2	21
		Platinum 8153	27.8	21
		Gold 6142	34.7	22
		Gold 6142M	34.7	22

	18 Cores Hyper-Threading, Turbo-Modus	Gold 6140	**36.6**	**25**
		Gold 6150	**39.5**	**26**
		Gold 6154	**41.5**	**26**
		Gold 6140M	**36.6**	**26**
	20 Cores Hyper-Threading, Turbo-Modus	Gold 6138	**36.8**	**20**
		Gold 6148	**40.3**	**28**
	22 Cores Hyper-Threading, Turbo-Modus	Gold 6152	**41.4**	**23**
	24 Cores Hyper-Threading, Turbo-Modus	Platinum 8160	**45.1**	**29**
		Platinum 8168	**51.3**	**30**
		Platinum 8160M	**45.1**	**29**
	26 Cores Hyper-Threading, Turbo-Modus	Platinum 8164	**46.1**	**32**
		Platinum 8170	**49.8**	**32**
		Platinum 8170M	**49.8**	**32**
	28 Cores Hyper-Threading, Turbo-Modus	Platinum 8176	**50.7**	**34**
		Platinum 8180	**57.8**	**34**
		Platinum 8176M	**50.7**	**34**
		Platinum 8180M	**57.8**	**34**

These PRIMERGY dual-socket rack and tower systems are very suitable for application virtualization thanks to the progress made in processor technology.

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

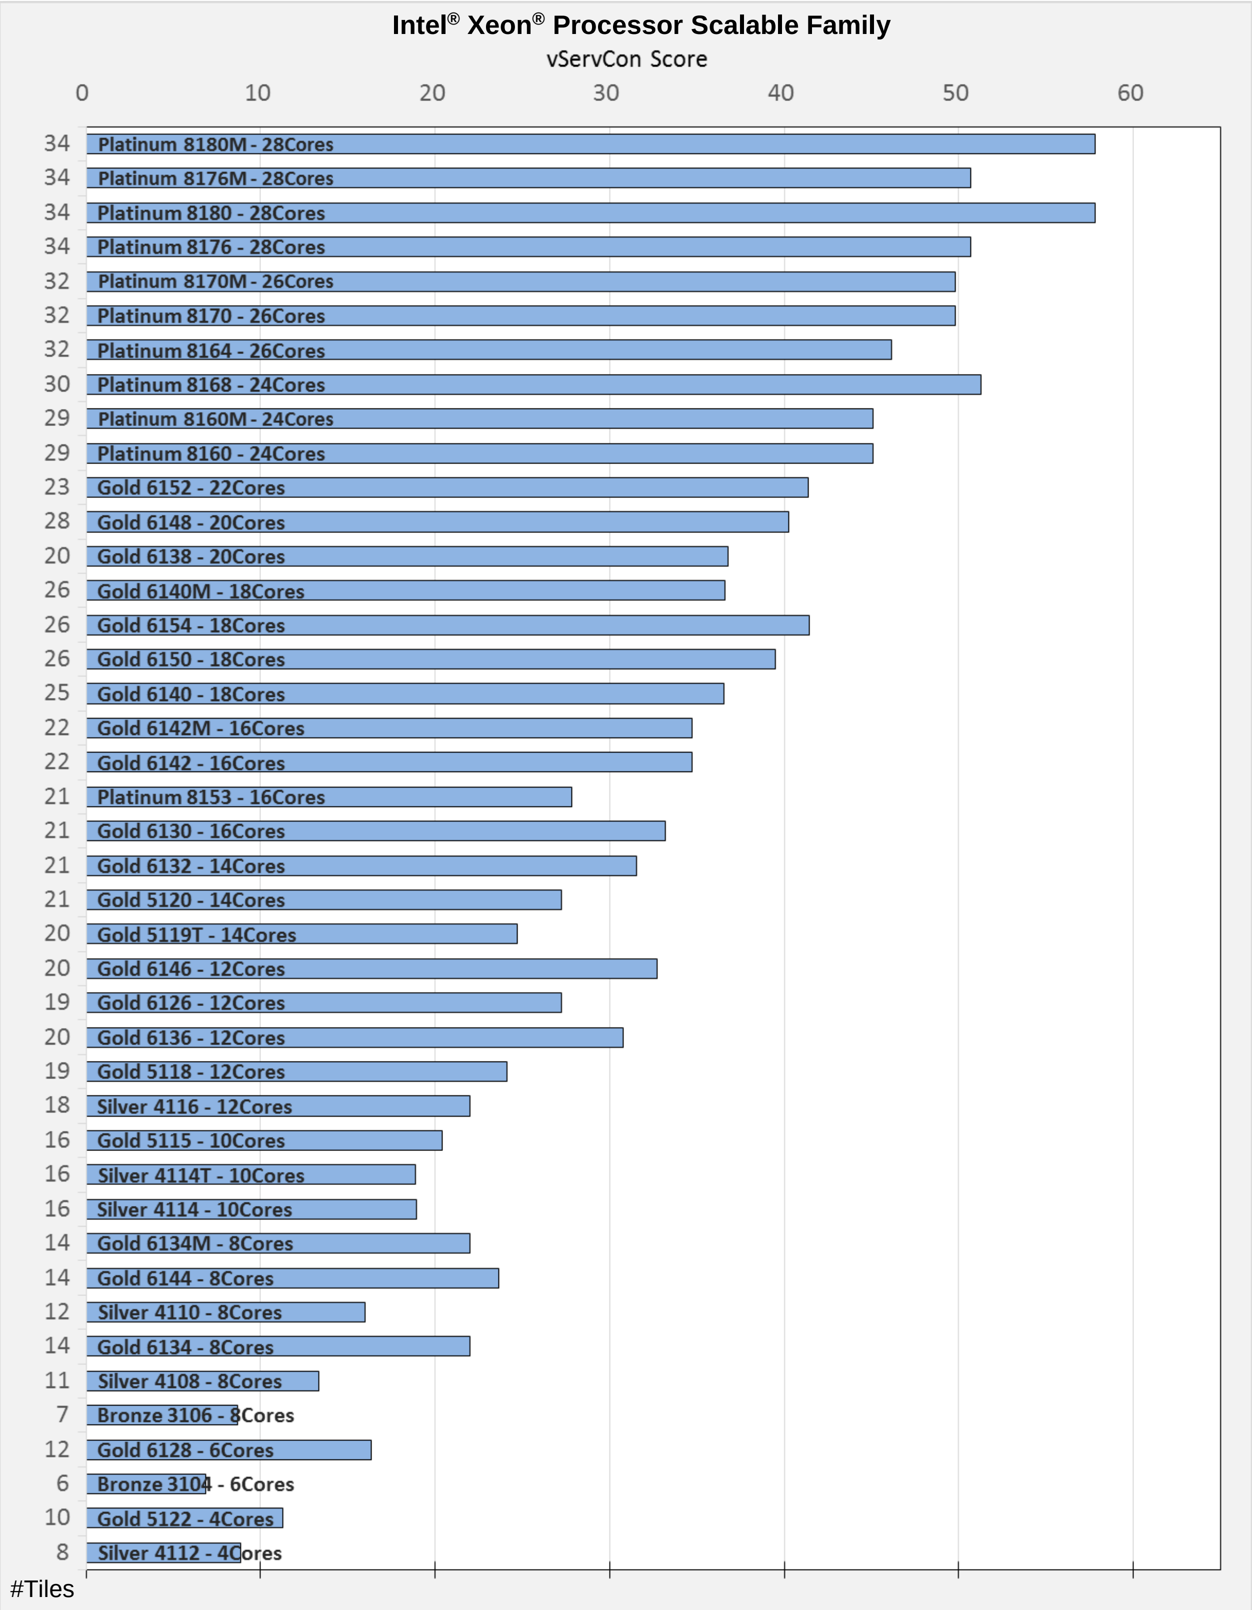

The relatively large performance differences between the processors can be explained by their features. The values scale on the basis of the number of cores, the size of the L3 cache and the CPU clock frequency and as a result of the features of Hyper-Threading and turbo mode, which are available in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

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 measurement set-up is symbolically illustrated below:

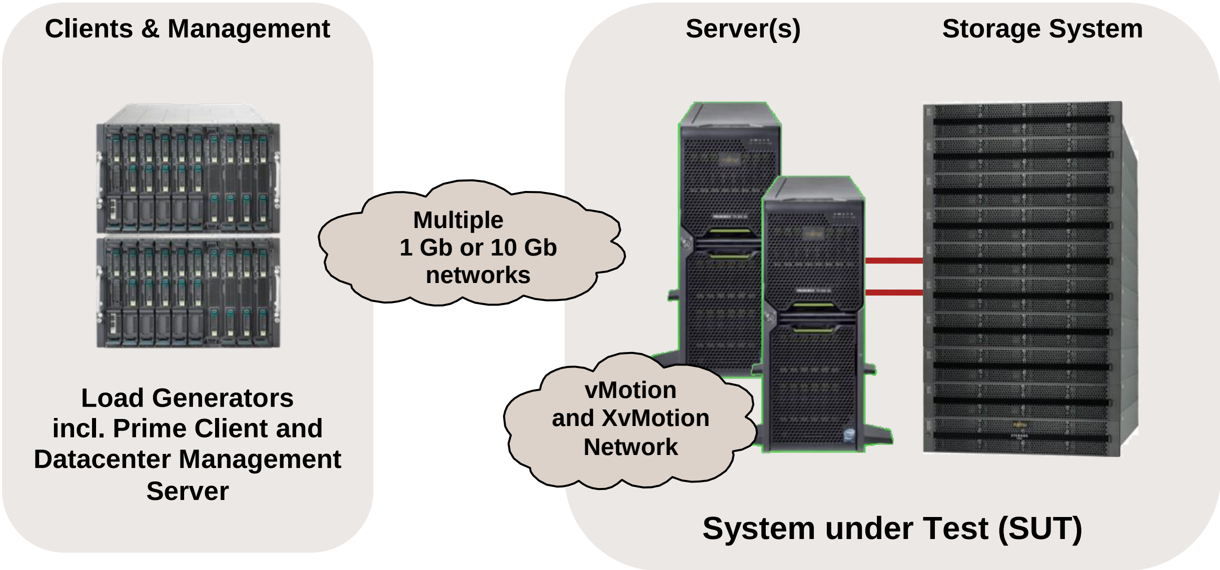

System Under Test (SUT)	
Hardware(Shared)	
Enclosure	PRIMERGY CX400 M4
Enclosure version	PRIMERGY CX400 M4 chassis for CX2560 M4 2U Chassis
Hardware(per Node)	
Number of servers	4
Model	PRIMERGY CX2560 M4
Cooling method	Air cooling
Processor	2 × Xeon Platinum 8164
Memory	768 GB: 12 × 64 GB (1x64 GB) 4Rx4 DDR4-2666 R 3DS
Network interface	1 × Emulex OneConnect OCe14000 Dual Port 10 GbE Adapter 1 × Intel I210 Gigabit Network Connection
Disk subsystem	1 × 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 1 × PRIMEQUEST 2800E3 configured as Fibre Channel target: 3 × SAS-SSD (400 GB) 3 × Intel P3700 (800 GB) RAID 0 with several LUNs Total: 3.6 TB
Software	
BIOS	R1.32.0
BIOS settings	See details
Operating system	VMware ESXi 6.7.0 EP 02a Build 9214924
Operating system settings	ESX settings: see details

Details	
See disclosure	http://www.vmware.com/a/assets/vmmark/pdf/2018-09-25-Fujitsu-PRIMERGY-CX2560M4.pdf

Datacenter Management Server (DMS)	
Hardware	
Model	1 × PRIMERGY RX2530 M2
Processor	2 × Intel Xeon E5-2698 v4
Memory	64 GB
Network interface	1 × Emulex One Connect Oce14000 1GbE 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 2012 R2 Standard (64-bit)

Load generator	
Hardware	
Model	4 × PRIMERGY RX2530 M2
Processor	Client Host 1-2 : 2 × Xeon E5-2699 v4 Client Host 3-4 : 2 × Xeon E5-2699A v4
Memory	256 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 6.5.0 U1 Build 5969303

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

Benchmark results

“Performance Only” measurement result (September 25 2018)

On September 25, 2018 Fujitsu achieved with a PRIMERGY CX400 M4 with 4 PRIMERGY CX2560 M4 server nodes with Xeon Platinum 8164 processors and VMware ESXi 6.7.0 EP 02a a VMmark V3 score of “11.45@11 tiles” in a system configuration with a total of 2 × 52 processor cores and when using four identical server nodes in the “System under Test” (SUT). With this result the PRIMERGY CX400 M4 with 4 PRIMERGY CX2560 M4 server nodes is in the official VMmark V3 “Performance Only” ranking the most powerful multi server nodes (valid as of benchmark results publication date).

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

The processors used, which with a good hypervisor setting could make optimal use of their processor features, were the essential prerequisites for achieving the PRIMERGY CX400 M4 with 4 PRIMERGY CX2560 M4 server nodes result. These features include Hyper-Threading. All this has a particularly positive effect during virtualization.

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

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

All the components used were optimally attuned to each other.

VMmark® is a product of VMware, Inc.

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 CX2550 M4/ CX2560 M4/ CX2570 M4
Processor	2 × Intel® Xeon® Processor Scalable Family
Memory	12 × 16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC
Software	
BIOS settings	Xeon Platinum 81xx, Gold 61xx: DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Power Technology = Custom HWPM Support = Disabled UPI Link Frequency Select = 10.4 GT/s Sub NUMA Clustering = Enabled Stale Atos = Enabled LLC dead line alloc = Disabled
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 aio-max-nr = 1048576 cpupower -c all frequency-set -g performance cpupower idle-set -d 1 cpupower idle-set -d 2 cpupower idle-set -d 3 echo 0 > /proc/sys/kernel/numa_balancing echo 1 > /proc/sys/vm/drop_caches ulimit -s unlimited
Compiler	C/C++: Version 17.0.0 20160721 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	CX2550M4 TRIAD [GB/s]	CX2560M4 TRIAD [GB/s]	CX2570M4 TRIAD [GB/s]
Xeon Bronze 3104	2133	102.4	6	1.7	2		*83*	*84*
Xeon Bronze 3106	2133	102.4	8	1.7	2		*113*	*114*
Xeon Silver 4108	2400	115.2	8	1.8	2		*126*	*127*
Xeon Silver 4110	2400	115.2	8	2.1	2		*126*	*127*
Xeon Silver 4114	2400	115.2	10	2.2	2		*147*	*149*
Xeon Silver 4116	2400	115.2	12	2.1	2		*159*	*161*
Xeon Gold 5115	2400	115.2	10	2.4	2		*150*	*152*
Xeon Gold 5118	2400	115.2	12	2.3	2		*164*	*167*
Xeon Gold 5120	2400	115.2	14	2.2	2		*177*	*180*
Xeon Gold 6130	2666	128.0	16	2.1	2	*205*	*202*	*205*
Xeon Gold 6140	2666	128.0	18	2.3	2	*209*	*206*	*209*
Xeon Gold 6138	2666	128.0	20	2.0	2	*215*	*212*	*215*
Xeon Gold 6148	2666	128.0	20	2.4	2	*215*	208	*215*
Xeon Gold 6152	2666	128.0	22	2.1	2	*217*	209	*217*
Xeon Platinum 8153	2666	128.0	16	2.0	2	*206*	*203*	*206*
Xeon Platinum 8160	2666	128.0	24	2.1	2	*220*	215	*220*
Xeon Platinum 8164	2666	128.0	26	2.0	2	*221*	218	*221*
Xeon Platinum 8170	2666	128.0	26	2.1	2	212		*212*
Xeon Platinum 8176	2666	128.0	28	2.1	2	220		*220*
Xeon Platinum 8180	2666	128.0	28	2.5	2	*220*		*220*
Xeon Silver 4112	2400	115.2	4	2.6	2		*89*	*90*
Xeon Gold 5122	2666	128.0	4	3.6	2		*106*	*108*
Xeon Gold 6128	2666	128.0	6	3.4	2	*163*	*161*	*163*
Xeon Gold 6134	2666	128.0	8	3.2	2	*191*	*189*	*191*
Xeon Gold 6144	2666	128.0	8	3.5	2	195		*195*
Xeon Gold 6126	2666	128.0	12	2.6	2	*196*	*194*	*196*
Xeon Gold 6136	2666	128.0	12	3.0	2	*211*	*208*	*211*
Xeon Gold 6146	2666	128.0	12	3.2	2	206		*206*
Xeon Gold 6132	2666	128.0	14	2.6	2	*194*	*192*	*194*
Xeon Gold 6142	2666	128.0	16	2.6	2	*200*	198	*200*
Xeon Gold 6150	2666	128.0	18	2.7	2	201		*201*
Xeon Gold 6154	2666	128.0	18	3.0	2	*209*		*209*
Xeon Platinum 8168	2666	128.0	24	2.7	2	*218*		*218*
Xeon Gold 6134M	2666	128.0	8	3.2	2	*191*	*189*	*191*
Xeon Gold 6140M	2666	128.0	18	2.3	2	*209*	*206*	*209*
Xeon Gold 6142M	2666	128.0	16	2.6	2	*200*	*198*	*200*

Xeon Platinum 8160M	2666	128.0	24	2.1	2	*220*	*217*	*220*
Xeon Platinum 8170M	2666	128.0	26	2.1	2	*220*		*220*
Xeon Platinum 8176M	2666	128.0	28	2.1	2	*220*		*220*
Xeon Platinum 8180M	2666	128.0	28	2.5	2	*224*		*224*
Xeon Silver 4114T	2400	115.2	10	2.2	2		*146*	*149*
Xeon Gold 5119T	2400	115.2	14	1.9	2		*178*	*176*

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 CX2550 M4/ CX2560 M4/ CX2570 M4
Processor	Intel® Xeon® Processor Scalable Family × 2
Memory	16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC × 12
Software	
BIOS settings	Xeon Platinum 81xx,Gold 61xx: HyperThreading = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Power Technology = Custom HWPM Support = Disabled UPI Link Frequency Select = 10.4 GT/s Sub NUMA Clustering = Disabled Stale Atos = Enabled LLC dead line alloc = Disabled
Operating system	SUSE Linux Enterprise Server 12 SP2 (x86_64)
Operating system settings	run with avx512 cpupower -c all frequency-set -g performance 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 aio-max-nr = 1048576
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 result of cursive is an estimated value.

Processor	Cores	Processor Frequency [GHz]	Number of Processors	Rpeak [GFlops]	CX2550 M4		CX2560 M4		CX2570 M4	
					Rmax [GFlops]	Efficiency	Rmax [GFlops]	Efficiency	Rmax [GFlops]	Efficiency
Xeon Bronze 3104	6	1.7	2	326			*214*	66%	*217*	67%
Xeon Bronze 3106	8	1.7	2	435			*286*	66%	*290*	67%
Xeon Silver 4108	8	1.8	2	461			*268*	58%	*272*	59%
Xeon Silver 4110	8	2.1	2	538			*460*	86%	*467*	87%
Xeon Silver 4114	10	2.2	2	704			*603*	86%	*612*	87%
Xeon Silver 4116	12	2.1	2	806			*688*	85%	*698*	87%
Xeon Gold 5115	10	2.4	2	768			*614*	80%	*624*	81%
Xeon Gold 5118	12	2.3	2	883			*754*	85%	*765*	87%
Xeon Gold 5120	14	2.2	2	986			*631*	64%	*641*	65%
Xeon Gold 6130	16	2.1	2	2,150	*1,653*	77%	*1,628*	76%	*1,653*	77%
Xeon Gold 6140	18	2.3	2	2,650	*1,844*	70%	*1,817*	69%	*1,844*	70%
Xeon Gold 6138	20	2.0	2	2,560	*1,762*	69%	*1,736*	68%	*1,762*	69%
Xeon Gold 6148	20	2.4	2	3,072	*2,018*	66%	2,072	67%	*2,018*	66%
Xeon Gold 6152	22	2.1	2	2,957	*1,990*	67%	2,050	69%	*1,990*	67%
Xeon Platinum 8153	16	2.0	2	2,048	*1,412*	69%	*1,390*	68%	*1,412*	69%
Xeon Platinum 8160	24	2.1	2	3,226	*2,164*	67%	2,269	70%	*2,164*	67%
Xeon Platinum 8164	26	2.0	2	3,328	*2,259*	68%	2,225	67%	*2,259*	68%
Xeon Platinum 8170	26	2.1	2	3,494	2,519	72%			*2,519*	72%
Xeon Platinum 8176	28	2.1	2	3,763	2,537	67%			*2,537*	67%
Xeon Platinum 8180	28	2.5	2	4,480	*3,045*	68%			*3,045*	68%
Xeon Silver 4112	4	2.6	2	333			*283*	85%	*288*	95%
Xeon Gold 5122	4	3.6	2	922			*662*	72%	*672*	73%
Xeon Gold 6128	6	3.4	2	1,306	*904*	69%	*890*	68%	*904*	69%
Xeon Gold 6134	8	3.2	2	1,638	*1,160*	71%	*1,142*	70%	*1,160*	71%
Xeon Gold 6144	8	3.5	2	1,792	*1,187*	66%			*1,187*	66%
Xeon Gold 6126	12	2.6	2	1,997	*1,424*	71%	*1,403*	70%	*1,424*	71%
Xeon Gold 6136	12	3.0	2	2,304	*1,625*	71%	*1,601*	69%	*1,625*	71%
Xeon Gold 6146	12	3.2	2	2,458	1,806	73%			*1,806*	73%
Xeon Gold 6132	14	2.6	2	2,330	*1,726*	74%	*1,700*	73%	*1,726*	74%
Xeon Gold 6142	16	2.6	2	2,662	*1,908*	72%	1,880	71%	*1,908*	72%
Xeon Gold 6150	18	2.7	2	3,110	*2,508*	60%			*2,508*	60%
Xeon Gold 6154	18	3.0	2	3,456	*2,508*	60%			*2,508*	60%
Xeon Platinum 8168	24	2.7	2	4,147	*2,508*	60%			*2,508*	60%
Xeon Gold 6134M	8	3.2	2	1,638	*1,160*	71%	*1,142*	70%	*1,160*	71%
Xeon Gold 6140M	18	2.3	2	2,650	*1,844*	70%	*1,817*	69%	*1,844*	70%
Xeon Gold 6142M	16	2.6	2	2,662	*1,908*	72%	*1,880*	71%	*1,908*	72%

Xeon Platinum 8160M	24	2.1	2	3,226	*2,164*	67%	*2,132*	66%	*2,164*	67%
Xeon Platinum 8170M	26	2.1	2	3,494	*2,485*	71%			*2,485*	71%
Xeon Platinum 8176M	28	2.1	2	3,763	*2,537*	67%			*2,537*	67%
Xeon Platinum 8180M	28	2.5	2	4,480	*3,045*	68%			*3,045*	68%
Xeon Silver 4114T	10	2.2	2	704			*603*	86%	*612*	87%
Xeon Gold 5119T	14	1.9	2	851			*725*	85%	*838*	98%

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 CX2550 M4/ CX2560 M4/ CX2570 M4

This White Paper:
http://docs.ts.fujitsu.com/dl.aspx?id=691cd43e-d5b6-42a3-9499-2d75cba4eb48
http://docs.ts.fujitsu.com/dl.aspx?id=c35ec51b-f227-4930-9546-4fb608252f9e

Data sheet
CX2550 M4: http://docs.ts.fujitsu.com/dl.aspx?id=a31cff74-84d4-4b6d-b699-7e36448fe9ae
CX2560 M4: http://docs.ts.fujitsu.com/dl.aspx?id=73f89221-8a8a-4a20-ab04-c0655fafe716
CX2570 M4: http://docs.ts.fujitsu.com/dl.aspx?id=a85b3de7-cd93-44e7-92fc-5ca526323494

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

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

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

2018/10/23 WW EN