White Paper
FUJITSU Server PRIMEQUEST Performance Report PRIMEQUEST 3800B/3800E/3400E

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMEQUEST PRIMEQUEST 3800E/3800B/3400E.

The PRIMEQUEST 3800E/3800B/3400E 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
2018/11/30

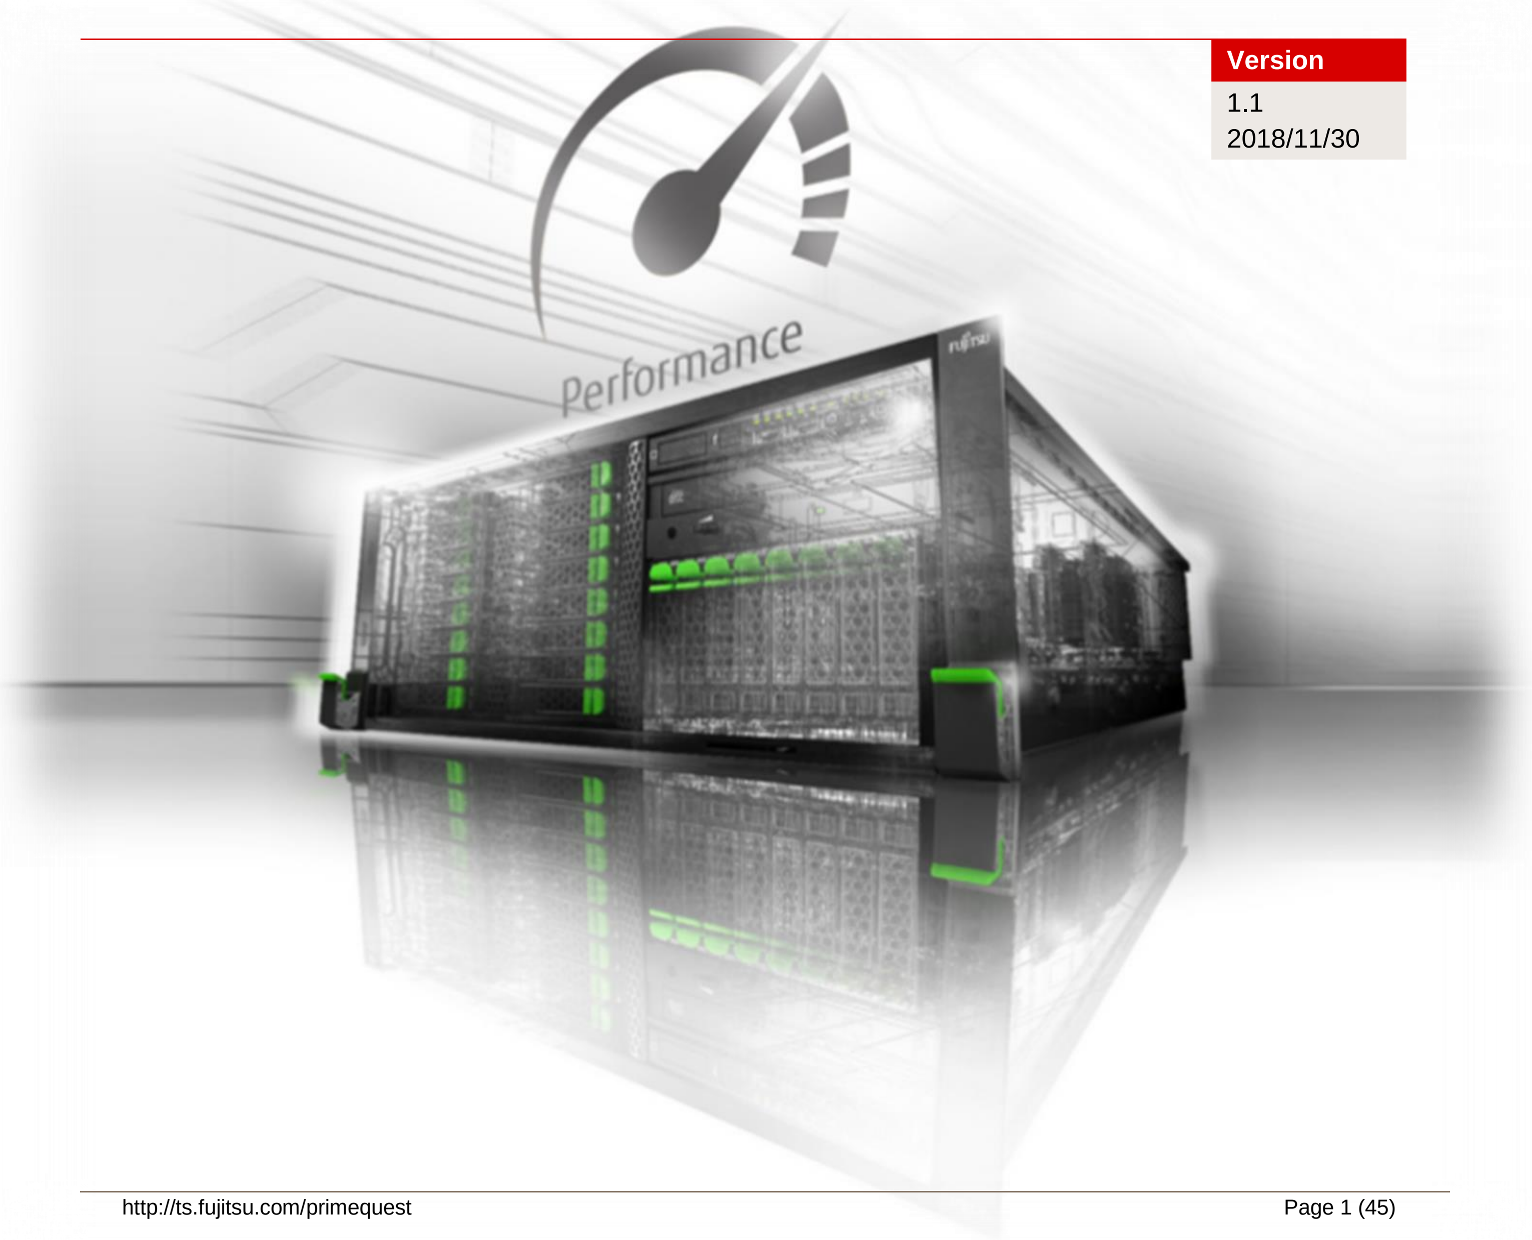

Contents

Document history

Version 1.1 (2018/11/30)

New:

- SPECcpu2017
  Measurements with and results for Intel® Xeon® Processor Scalable Family
- SAP HANA
  Measurements with Intel® Xeon® Processor Platinum 8180

Version 1.0 (2018/02/23)

New:

- Technical data
- SPECcpu2006
  Measurements with and results for Intel® Xeon® Processor Scalable Family
- SAP SD
  Certification number 2017014
- 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

Technical data

PRIMEQUEST 3800B

PRIMEQUEST 3800E/3400E

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>PRIMEQUEST 3800B</th><th>PRIMEQUEST 3800E</th><th>PRIMEQUEST 3400E</th></tr>
<tr><td>Form factor</td><td colspan="3">Rack server</td></tr>
<tr><td>Number of system boards orderable</td><td colspan="2">1 – 4</td><td>1 – 2</td></tr>
<tr><td>Number of I/O units orderable</td><td>2</td><td colspan="2">4</td></tr>
<tr><td>Number of disk units orderable</td><td>0 – 2</td><td colspan="2">0 – 6</td></tr>
<tr><td>Shared LAN controller</td><td>1 x Intel® i210 onboard
10/100/1000 Mbit/s Ethernet</td><td colspan="2">none</td></tr>
<tr><td colspan="4">Per system board:</td></tr>
<tr><td>Chipset</td><td colspan="3">Intel® C621 Chipset</td></tr>
<tr><td>Number of sockets</td><td colspan="3">2</td></tr>
<tr><td>Number of processors orderable</td><td colspan="3">2</td></tr>
<tr><td>Processor type</td><td colspan="3">Intel® Xeon® Processor Scalable Family</td></tr>
<tr><td>Number of memory slots</td><td colspan="3">24 (12 per processor)</td></tr>
<tr><td>Maximum memory configuration</td><td colspan="3">3TB</td></tr>
<tr><td colspan="4">Per I/O unit:</td></tr>
<tr><td>Onboard LAN controller</td><td>none</td><td colspan="2">2 x 10 Gbit/s Ethernet (RJ45)</td></tr>
<tr><td>PCI slots</td><td>4 × PCI-Express 3.0 x 8
4 × PCI-Express 3.0 x16</td><td colspan="2">1 × PCI-Express 3.0 x 8
1 × PCI-Express 3.0 x16</td></tr>
<tr><td colspan="4">Per disk unit:</td></tr>
<tr><td>Max. number of internal hard disks</td><td>8</td><td colspan="2">24</td></tr>
</table>

Processors (since system release)									
Processor	Cores	Threads	Max CPUs	Cache [MB]	UPI Speed [GT/s]	Rated Frequency [GHz]	Max. Turbo Frequency [GHz]	Max. Memory Frequency [MHz]	TDP [Watt]
Xeon Gold 6130	16	32	4	22.0	10.4	2.1	3.7	2666	125
Xeon Gold 6140	18	36	4	24.8	10.4	2.3	3.7	2666	140
Xeon Gold 6138	20	40	4	27.5	10.4	2.0	3.7	2666	125
Xeon Gold 6148	20	40	4	27.5	10.4	2.4	3.7	2666	150
Xeon Gold 6152	22	44	4	30.3	10.4	2.1	3.7	2666	140
Xeon Platinum 8153	16	32	8	22.0	10.4	2.0	2.8	2666	125
Xeon Platinum 8160	24	48	8	33.0	10.4	2.1	3.7	2666	150
Xeon Platinum 8164	26	52	8	35.8	10.4	2.0	3.7	2666	150
Xeon Platinum 8170	26	52	8	35.8	10.4	2.1	3.7	2666	165
Xeon Platinum 8176	28	56	8	38.5	10.4	2.1	3.8	2666	165
Xeon Platinum 8180	28	56	8	38.5	10.4	2.5	3.8	2666	205
Xeon Gold 6128	6	12	4	19.3	10.4	3.4	3.7	2666	115
Xeon Gold 6134	8	16	4	24.8	10.4	3.2	3.7	2666	130
Xeon Gold 6144	8	16	4	24.8	10.4	3.5	3.7	2666	150
Xeon Gold 6126	12	24	4	19.3	10.4	2.6	3.7	2666	125
Xeon Gold 6136	12	24	4	24.8	10.4	3.0	3.7	2666	150
Xeon Gold 6146	12	24	4	24.8	10.4	3.2	3.7	2666	165
Xeon Gold 6132	14	28	4	19.3	10.4	2.6	3.7	2666	140
Xeon Gold 6142	16	32	4	22.0	10.4	2.6	3.7	2666	150
Xeon Gold 6150	18	36	4	24.8	10.4	2.7	3.7	2666	165
Xeon Gold 6154	18	36	4	24.8	10.4	3.0	3.7	2666	200
Xeon Platinum 8156	4	8	8	16.5	10.4	3.6	3.7	2666	105
Xeon Platinum 8158	12	24	8	24.8	10.4	3.0	3.6	2666	150
Xeon Platinum 8168	24	48	8	33.0	10.4	2.7	3.7	2666	205
Xeon Gold 6134M	8	16	4	24.8	10.4	3.2	3.7	2666	130
Xeon Gold 6140M	18	36	4	24.8	10.4	2.3	3.7	2666	140
Xeon Gold 6142M	16	32	4	22.0	10.4	2.6	3.7	2666	150
Xeon Platinum 8160M	24	48	8	33.0	10.4	2.1	3.7	2666	150
Xeon Platinum 8170M	26	52	8	35.8	10.4	2.1	3.7	2666	165
Xeon Platinum 8176M	28	56	8	38.5	10.4	2.1	3.8	2666	165
Xeon Platinum 8180M	28	56	8	38.5	10.4	2.5	3.8	2666	205

Model	PRIMEQUEST 3800B/3800E	PRIMEQUEST 3400E
Supported CPU	Xeon Platinum 8153 Xeon Platinum 8160 Xeon Platinum 8164 Xeon Platinum 8170 Xeon Platinum 8176 Xeon Platinum 8180 Xeon Platinum 8156 Xeon Platinum 8158 Xeon Platinum 8168 Xeon Platinum 8160M Xeon Platinum 8170M Xeon Platinum 8176M Xeon Platinum 8180M	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 8156 Xeon Platinum 8158 Xeon Platinum 8168 Xeon Gold 6134M Xeon Gold 6140M Xeon Gold 6142M Xeon Platinum 8160M Xeon Platinum 8170M Xeon Platinum 8176M Xeon Platinum 8180M

All the processors that can be ordered with the PRIMEQUEST 3800E/3800B/3400E 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 frequency maximum 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 will 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 "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
16GB (2x8GB) 1Rx4 DDR4-2666 R ECC	16	1	4	2666			✓	✓
32GB (2x16GB) 1Rx4 DDR4-2666 R ECC	32	1	4	2666			✓	✓
32GB (2x16GB) 2Rx4 DDR4-2666 R ECC	32	1	4	2666			✓	✓
64GB (2x32GB) 2Rx4 DDR4-2666 R ECC	64	2	4	2666			✓	✓
128GB (2x64GB) 4Rx4 DDR4-2666 LR ECC	128	4	4	2666		✓	✓	✓
128GB (2x64GB) 4Rx4 DDR4-2666 3DS ECC	128	4	4	2666			✓	✓
256GB (2x128GB) 8Rx4 DDR4-2666 3DS ECC	256	8	4	2666			✓	✓

Power supplies (since system release)	Max. number
Modular PSU 2200W platinum hp	4

Some components may not be available in all countries or sales regions.
Detailed technical information is available in the data sheet PRIMEQUEST 3800B/3800E/3400E.

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 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; the peak values are optional.

Benchmark	Arithmetics	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 favour 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” specify 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	PRIMEQUEST 3800E/3800B/3400E
Processor	Intel® Xeon® Processor Scalable Family
Memory	4 processors: 24 ×16GB (2x16GB) 2Rx4 PC4-2666V R ECC 8 processors: 48 ×16GB (2x16GB) 2Rx4 PC4-2666V R ECC
Software	
BIOS settings	Xeon Platinum 81xx: Intel Virtualization Technology = Disabled HWPM Support = Disabled DCU Streamer Prefetcher = Disabled Stale AtoS = Enabled LLC dead line alloc = Disabled Sub NUMA Clustering = Enabled Fan Control = Full
Operating system	SUSE Linux Enterprise Server 12 SP2 (x86_64)
Operating system settings	Stack size set to unlimited using "ulimit -s unlimited" Set Kernel Boot Parameter : nohz_full=1-xxx isolcpus=1-xxx Set CPU frequency governor to maximum performance with: cpupower -c all frequency-set -g performance Set tmpfs filesystem with: mkdir /home/memory mount -t tmpfs -o size=768g,rw tmpfs /home/memory Process tuning settings: echo 10000000 > /proc/sys/kernel/sched_min_granularity_ns echo 15000000 > /proc/sys/kernel/sched_wakeup_granularity_ns echo always > /sys/kernel/mm/transparent_hugepage/enabled cpu idle state set with: cpupower idle-set -d 2 cpupower idle-set -d 3 set affinity of rcu threads to the cpu0: for i in \`pgrep rcu\` ; do taskset -pc 0 $i ; done int_rate: echo 1 > /proc/sys/kernel/numa_balancing fp_rate: echo 0 > /proc/sys/kernel/numa_balancing
Compiler	C/C++: Version 18.0.0.128 of Intel C/C++ Compiler for Linux Fortran: Version 18.0.0.128 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 result with “est." is an estimated value.

Model Name	PQ3800E		PQ3800B			PQ3400E	
Processor	Number of processors	SPECint_rate_base2006	Number of processors	SPECint_rate_base2006	SPECint_rate2006	Number of processors	SPECint_rate_base2006
Xeon Gold 6130						4	*3146 est.*
Xeon Gold 6140						4	*3644 est.*
Xeon Gold 6138						4	*3779 est.*
Xeon Gold 6148						4	*3893 est.*
Xeon Gold 6152						4	*3987 est.*
Xeon Platinum 8153	8	*5650 est.*	8	**5700**		4	*2850 est.*
Xeon Platinum 8160	8	*8901 est.*	8	**8980**		4	*4490 est.*
Xeon Platinum 8164	8	*9367 est.*	8	**9450**		4	*4725 est.*
Xeon Platinum 8170	8	*9565 est.*	8	**9650**		4	*4825 est.*
Xeon Platinum 8176	8	*10111 est.*	8	**10200**		4	*5100 est.*
Xeon Platinum 8180	8	**11300**	8	**11400**	**11800**	4	*5700 est.*
Xeon Gold 6128						4	*1734 est.*
Xeon Gold 6134						4	*2108 est.*
Xeon Gold 6144						4	*2378 est.*
Xeon Gold 6126						4	*2751 est.*
Xeon Gold 6136						4	*2980 est.*
Xeon Gold 6146						4	*3063 est.*
Xeon Gold 6132						4	*3198 est.*
Xeon Gold 6142						4	*3499 est.*
Xeon Gold 6150						4	*3935 est.*
Xeon Gold 6154						4	*4288 est.*
Xeon Platinum 8156	8	*2230 est.*	8	**2250**		4	*1125 est.*
Xeon Platinum 8158	8	*6126 est.*	8	**6180**		4	*3090 est.*
Xeon Platinum 8168	8	*10210 est.*	8	**10300**		4	*5150 est.*
Xeon Gold 6134M						4	*2108 est.*
Xeon Gold 6140M						4	*3644 est.*
Xeon Gold 6142M						4	*3499 est.*
Xeon Platinum 8160M	8	*8901 est.*	8	*8980 est.*		4	*4490 est.*

Xeon Platinum 8170M	8	*9565 est.*	8	*9650 est.*		4	*4825 est.*
Xeon Platinum 8176M	8	*10111 est.*	8	*10200 est.*		4	*5100 est.*
Xeon Platinum 8180M	8	*11300 est.*	8	*11400 est.*	*11800 est.*	4	*5700 est.*

Model Name	PQ3800E		PQ3800B		PQ3400E	
Processor	Number of processors	SPECfp_rate_base2006	Number of processors	SPECfp_rate_base2006	Number of processors	SPECfp_rate_base2006
Xeon Gold 6130					4	*2444 est.*
Xeon Gold 6140					4	*2688 est.*
Xeon Gold 6138					4	*2796 est.*
Xeon Gold 6148					4	*2850 est.*
Xeon Gold 6152					4	*2860 est.*
Xeon Platinum 8153	8	*4660 est.*	8	**4660**	4	*2330 est.*
Xeon Platinum 8160	8	*6260 est.*	8	**6260**	4	*3130 est.*
Xeon Platinum 8164	8	*6410 est.*	8	**6410**	4	*3205 est.*
Xeon Platinum 8170	8	*6510 est.*	8	**6510**	4	*3255 est.*
Xeon Platinum 8176	8	*6720 est.*	8	**6720**	4	*3360 est.*
Xeon Platinum 8180	8	**7120**	8	**7120**	4	*3560 est.*
Xeon Gold 6128					4	*1511 est.*
Xeon Gold 6134					4	*1846 est.*
Xeon Gold 6144					4	*1978 est.*
Xeon Gold 6126					4	*2211 est.*
Xeon Gold 6136					4	*2353 est.*
Xeon Gold 6146					4	*2454 est.*
Xeon Gold 6132					4	*2454 est.*
Xeon Gold 6142					4	*2607 est.*
Xeon Gold 6150					4	*2809 est.*
Xeon Gold 6154					4	*2982 est.*
Xeon Platinum 8156	8	*2160 est.*	8	**2160**	4	*1080 est.*
Xeon Platinum 8158	8	*4930 est.*	8	**4930**	4	*2465 est.*
Xeon Platinum 8168	8	*6770 est.*	8	**6770**	4	*3385 est.*
Xeon Gold 6134M					4	*1846 est.*
Xeon Gold 6140M					4	*2688 est.*
Xeon Gold 6142M					4	*2607 est.*
Xeon Platinum 8160M	8	*6260 est.*	8	*6260 est.*	4	*3130 est.*
Xeon Platinum 8170M	8	*6510 est.*	8	*6510 est.*	4	*3255 est.*
Xeon Platinum 8176M	8	*6720 est.*	8	*6720 est.*	4	*3360 est.*
Xeon Platinum 8180M	8	*7120 est.*	8	*7120 est.*	4	*3560 est.*

On 3rd Oct 2017 the PRIMEQUEST 3800B with eight Xeon Platinum 8180 processors was ranked first in the 8-socket systems category for the benchmark SPECint_rate_base2006.

On 3rd Oct 2017 the PRIMEQUEST 3800B with eight Xeon Platinum 8180 processors was ranked first in the 8-socket x86 systems category for the benchmark SPECfp_rate_base2006.

On 14th Nov 2017 the PRIMEQUEST 3800E with eight Xeon Platinum 8180 processors was ranked first (tie) in the 8-socket x86 systems category for the benchmark SPECfp_rate_base2006.

The current results can be found at http://www.spec.org/cpu2006/results.

The following two diagrams illustrate the throughput of the PRIMEQUEST 3800E/3800B/3400E in comparison to its predecessor PRIMEQUEST 2800E3/2800B3, in their respective most performant configuration.

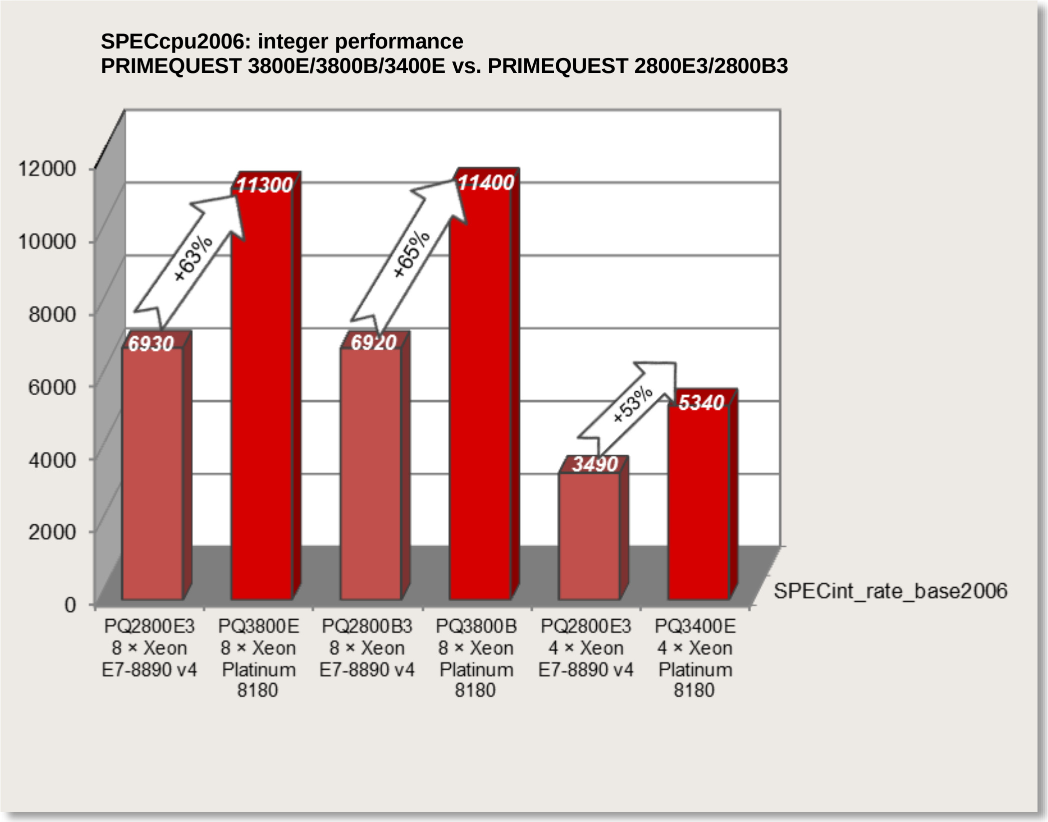

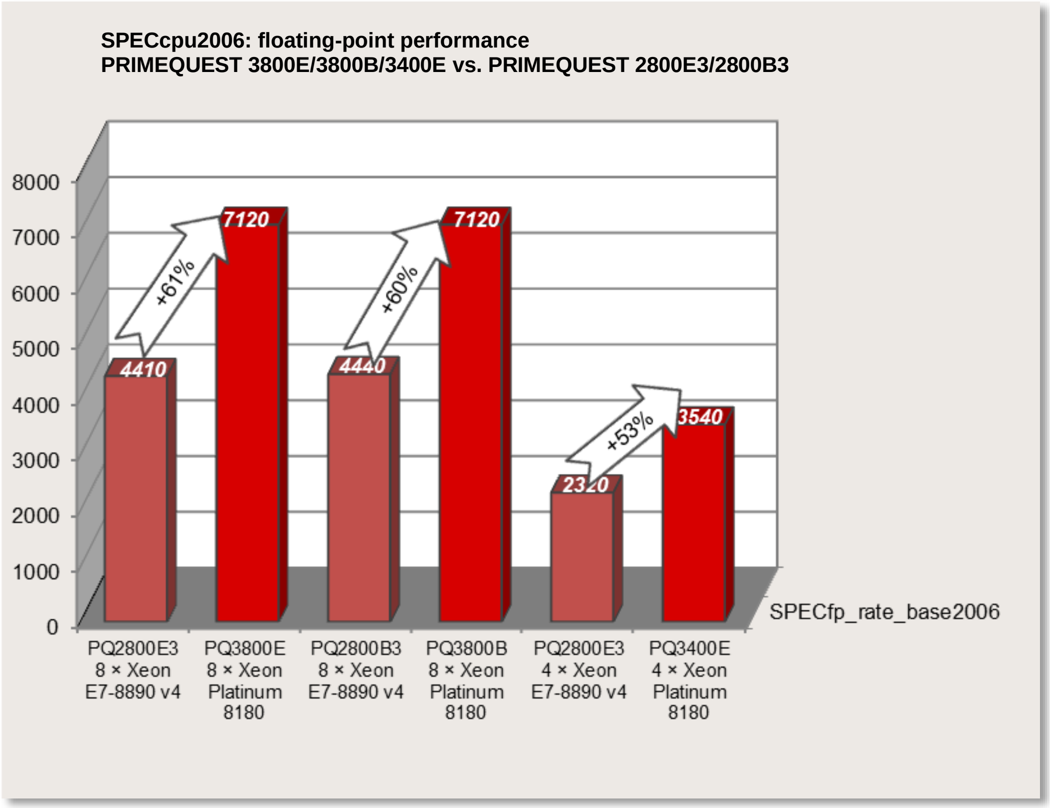

SPECcpu2017

Benchmark description

SPECcpu2017 is a benchmark which measures the system efficiency with integer and floating-point operations. It consists of an integer test suite (SPECrate 2017 Integer, SPECspeed 2017 Integer) containing 10 applications and a floating-point test suite (SPECrate 2017 Floating Point, SPECspeed 2017 Floating Point) containing 14 applications. Both test suites are extremely computing-intensive and concentrate on the CPU and the memory. Other components, such as Disk I/O and network, are not measured by this benchmark.

SPECcpu2017 is not tied to a special operating system. The benchmark is available as source code and is compiled before the actual measurement. The used compiler version and their optimization settings also affect the measurement result.

SPECcpu2017 contains two different performance measurement methods: The first method (SPECspeed 2017 Integer or SPECspeed 2017 Floating Point) determines the time which is required to process a single task. The second method (SPECrate 2017 Integer or SPECrate 2017 Floating Point) determines the throughput, i.e. the number of tasks that can be handled in parallel. Both methods are also divided into two measurement runs, “base” and “peak”, which differ in the use of compiler optimization. When publishing the results, the base values are always used and the peak values are optional.

Benchmark	Number of single benchmarks	Arithmetics	Type	Compiler optimization	Measurement result
SPECspeed2017_int_peak	10	integer	peak	aggressive	Speed
SPECspeed2017_int_base	10	integer	base	conservative	
SPECrate2017_int_peak	10	integer	peak	aggressive	Throughput
SPECrate2017_int_ base	10	integer	base	conservative	
SPECspeed2017_fp_peak	10	floating point	peak	aggressive	Speed
SPECspeed2017_fp_base	10	floating point	base	conservative	
SPECrate2017_fp_peak	13	floating point	peak	aggressive	Throughput
SPECrate2017_fp_base	13	floating point	base	conservative	

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

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

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMEQUEST 3800E/3800B/3400E
Processor	Intel® Xeon® Processor Scalable Family
Memory	8 processors: 16GB (1x16GB) 2Rx4 PC4-2666V R ECC x 96
Software	
BIOS settings	Xeon Platinum 81xx: DCU Streamer Prefetcher = Disabled Stale AtoS = Enabled LLC dead line alloc = Disabled Sub NUMA Clustering = Enabled Fan Control = Full
Operating system	SUSE Linux Enterprise Server 12 SP2 (x86_64)
Operating system settings	Stack size set to unlimited using "ulimit -s unlimited" Set Kernel Boot Parameter : nohz_full=1-xxx isolcpus=1-xxx Set tmpfs filesystem with: mkdir /home/memory mount -t tmpfs -o size=768g,rw tmpfs /home/memory Process tuning settings: echo 0 > /proc/sys/kernel/numa_balancing cpu idle state set with: cpupower idle-set -d 1 cpupower idle-set -d 2 set affinity of rcu threads to the cpu0: for i in \`pgrep rcu\` ; do taskset -pc 0 $i ; done int_rate: echo always > /sys/kernel/mm/transparent_hugepage/enabled fp_rate: echo never > /sys/kernel/mm/transparent_hugepage/enabled
Compiler	SPECrate2017_fp_base of Xeon Platinum 8180 : C/C++: Version 18.0.3.222 of Intel C/C++ Compiler for Linux Fortran: Version 18.0.3.222 of Intel Fortran Compiler for Linux Others: C/C++: Version 18.0.0.128 of Intel C/C++ Compiler for Linux Fortran: Version 18.0.0.128 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 result with “est." is an estimated value.

Model Name	PQ3800E		PQ3800B		PQ3400E	
Processor	Number of processors	SPECrate2017 int_base	Number of processors	SPECrate2017 int_base	Number of processors	SPECrate2017 int_base
Xeon Gold 6130					4	*322 est.*
Xeon Gold 6140					4	*381 est.*
Xeon Gold 6138					4	*382 est.*
Xeon Gold 6148					4	*424 est.*
Xeon Gold 6152					4	*433 est.*
Xeon Platinum 8153	8	*561 est.*	8	**566**	4	*283 est.*
Xeon Platinum 8160	8	*910 est.*	8	**918**	4	*459 est.*
Xeon Platinum 8164	8	*958 est.*	8	**966**	4	*483 est.*
Xeon Platinum 8170	8	*979 est.*	8	**987**	4	*494 est.*
Xeon Platinum 8176	8	*1041 est.*	8	**1050**	4	*525 est.*
Xeon Platinum 8180	8	**1170**	8	**1180**	4	*590 est.*
Xeon Gold 6128					4	*172 est.*
Xeon Gold 6134					4	*217 est.*
Xeon Gold 6144					4	*237 est.*
Xeon Gold 6126					4	*281 est.*
Xeon Gold 6136					4	*304 est.*
Xeon Gold 6146					4	*321 est.*
Xeon Gold 6132					4	*324 est.*
Xeon Gold 6142					4	*369 est.*
Xeon Gold 6150					4	*409 est.*
Xeon Gold 6154					4	*438 est.*
Xeon Platinum 8156	8	*218 est.*	8	**220**	4	*110 est.*
Xeon Platinum 8158	8	*614 est.*	8	**619**	4	*310 est.*
Xeon Platinum 8168	8	*1051 est.*	8	**1060**	4	*530 est.*
Xeon Gold 6134M					4	*217 est.*
Xeon Gold 6140M					4	*381 est.*
Xeon Gold 6142M					4	*369 est.*
Xeon Platinum 8160M	8	*910 est.*	8	*918 est.*	4	*459 est.*
Xeon Platinum 8170M	8	*979 est.*	8	*987 est.*	4	*494 est.*
Xeon Platinum 8176M	8	*1041 est.*	8	*1050 est.*	4	*525 est.*
Xeon Platinum 8180M	8	*1170 est.*	8	*1180 est.*	4	*590 est.*

Model Name	PQ3800E		PQ3800B		PQ3400E	
Processor	Number of processors	SPECrate2017 fp_base	Number of processors	SPECrate2017 fp_base	Number of processors	SPECrate2017 fp_base
Xeon Gold 6130					4	*325 est.*
Xeon Gold 6140					4	*369 est.*
Xeon Gold 6138					4	*369 est.*
Xeon Gold 6148					4	*395 est.*
Xeon Gold 6152					4	*403 est.*
Xeon Platinum 8153	8	*602 est.*	8	**604**	4	*302 est.*
Xeon Platinum 8160	8	*825 est.*	8	**828**	4	*414 est.*
Xeon Platinum 8164	8	*850 est.*	8	**853**	4	*427 est.*
Xeon Platinum 8170	8	*860 est.*	8	**864**	4	*432 est.*
Xeon Platinum 8176	8	*896 est.*	8	**900**	4	*450 est.*
Xeon Platinum 8180	8	*977 est.*	8	**1060**	4	*491 est.*
Xeon Gold 6128					4	*195 est.*
Xeon Gold 6134					4	*254 est.*
Xeon Gold 6144					4	*264 est.*
Xeon Gold 6126					4	*294 est.*
Xeon Gold 6136					4	*326 est.*
Xeon Gold 6146					4	*335 est.*
Xeon Gold 6132					4	*324 est.*
Xeon Gold 6142					4	*361 est.*
Xeon Gold 6150					4	*387 est.*
Xeon Gold 6154					4	*405 est.*
Xeon Platinum 8156	8	*277 est.*	8	**278**	4	*139 est.*
Xeon Platinum 8158	8	*653 est.*	8	**656**	4	*328 est.*
Xeon Platinum 8168	8	*905 est.*	8	**909**	4	*455 est.*
Xeon Gold 6134M					4	*254 est.*
Xeon Gold 6140M					4	*369 est.*
Xeon Gold 6142M					4	*361 est.*
Xeon Platinum 8160M	8	*825 est.*	8	*838 est.*	4	*414 est.*
Xeon Platinum 8170M	8	*860 est.*	8	*864 est.*	4	*432 est.*
Xeon Platinum 8176M	8	*896 est.*	8	*900 est.*	4	*450 est.*
Xeon Platinum 8180M	8	*977 est.*	8	*981 est.*	4	*491 est.*

On 4th Sep. 2018 the PRIMEQUEST 3800B with eight Xeon Platinum 8180 processors was ranked first in the 8-socket systems category for the benchmark SPECrate2017_fp_base.

The current results can be found at http://www.spec.org/cpu2017/results.

SAP SD

Benchmark description

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

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

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

The benchmark differentiates between a 2-tier and a 3-tier configuration. The 2-tier configuration has the SAP application and database installed on one server. With a 3-tier configuration the individual components of the SAP application can be distributed via several servers and an additional server handles the database.

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

Benchmark environment

The measurement set-up is symbolically illustrated below:

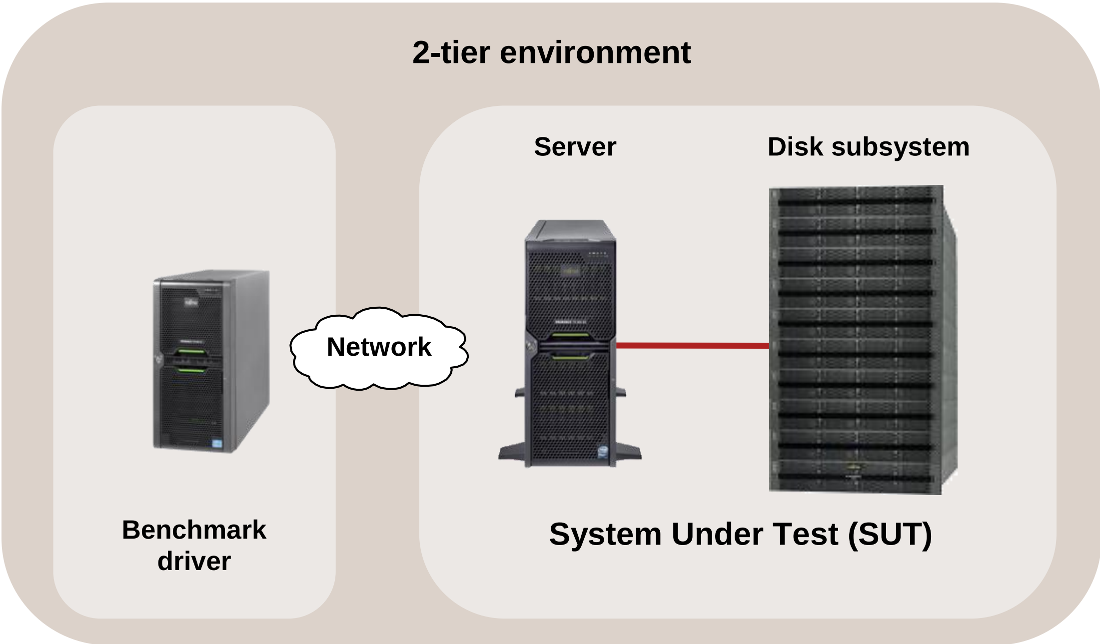

System Under Test (SUT)	
Hardware	
Model	PRIMEQUEST 3800B
Processor	8 × Xeon Platinum 8180
Memory	48 × 64GB (2 x 32GB) 2Rx4 DDR4-2666 R ECC
Network interface	10Gbit/s LAN
Disk subsystem	PRIMEQUEST 3800B: 1 × HD SAS 12G 300GB 10K HOT PL 2.5” 3 × HD SAS 12G 600GB 15K HOT PL 2.5” 1 × PRAID EP420i 1 × PRAID EP420e 2 × ETERNUS JX40 S2
Software	
BIOS settings	Enable SNC
Operating system	Microsoft Windows Server 2012 R2 Standard Edition
Database	Microsoft SQL Server 2012 (64-bit)
SAP Business Suite Software	SAP enhancement package 5 for SAP ERP 6.0

Benchmark driver	
Hardware	
Model	PRIMERGY RX2530 M1
Processor	2 × Xeon E5-2699 v3
Memory	256 GB
Network interface	10Gbit/s LAN
Software	
Operating system	SUSE Linux Enterprise Server 12 SP2

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

Benchmark results

Certification number 2017014	
Number of SAP SD benchmark users	101,000
Average dialog response time	0.78 seconds
Throughput Fully processed order line items/hour Dialog steps/hour SAPS	 11,242,000 33,726,000 562,100
Average database request time (dialog/update)	0.008 sec / 0.014 sec
CPU utilization of central server	94%
Operating system, central server	Windows Server 2012 R2 Standard Edition
RDBMS	SQL Server 2012
SAP Business Suite software	SAP enhancement package 5 for SAP ERP 6.0
Configuration Central Server	Fujitsu PRIMEQUEST 3800B, 8 processors / 224 cores / 448 threads, Intel Xeon Platinum 8180 processor, 2.50 GHz, 64 KB L1 cache and 1024 KB L2 cache per core, 38.5 MB L3 cache per processor, 3072 GB main memory

The PRIMEQUEST 3800B obtained the best two-tier SAP SD Standard Application Benchmark 8-way result on Windows (as of July 11, 2017). The latest SAP SD 2-tier results can be found at http://global.sap.com/solutions/benchmark/sd2tier.epx .

The SAP BW Edition for SAP HANA Standard Application Benchmark

With the increasing importance of SAP HANA and in particular SAP Business Warehouse (SAP BW) on HANA, a new benchmark was introduced in July 2016: the SAP BW Edition for SAP HANA Standard Application Benchmark, referred to as SAP BWH Benchmark in the following.

The benchmark represents a typical mid-size customer scenario and volumes and utilizes the new capabilities of SAP HANA which enable customers to enhance their BW processes.

Since its first edition in 2016, the SAP BWH Benchmark has been further developed and adapted to customer requirements. In the meantime SAP BWH Benchmark version 3 is available.

Comparability of SAP BW Benchmark Results

The SAP BWH Benchmark goes beyond the scope and features used in the BW-AML benchmark. While BW-AML focuses on traditional BW objects and processes supported by all database platforms, the SAP BWH Benchmark takes advantage of HANA's new features. The workload of the two benchmarks completely differ and thus are not comparable.

Within the SAP BWH Benchmark, both the data model and the query definitions have changed in the course of the development of version 1 to version 3. That is why the results of different versions must not be compared with each other.

BWH Benchmark version 3 for instance is available for SAP HANA 2.0 only and the Data Load KPI (Phase 1) has been changed - the runtime of the latest data set (1,3 billion records) will be the KPI.

In SAP BWH Benchmark version 2 it was allowed to activate a specific SAP HANA performance enhancement function which materializes intermediate query results. Benchmarks that have used the feature cannot be compared with benchmarks that haven't. Whether the function was used can be seen in the details on www.sap.com/benchmark , "Materializing of the intermediate result of the query was enabled (phase 2 and phase 3)" is shown.

In addition, SAP BWH Benchmark results with different number of data sets cannot be compared either.

Currently released version of the benchmark is version 3. Benchmarks with the older versions won't be certified anymore.

Features of the SAP BW Edition for SAP HANA Standard Application Benchmark

The SAP BWH Benchmark consists of 3 phases:

- Data load phase
- Query throughput phase
- Query runtime phase

Data load phase

The data flow starts with a data load from the source object into the corporate memory layer. The source object is shipped with the backup.

The source object contains 1.3 billion records (= 1 data set). It is possible to load this data set of 1.3 billion records multiple times.

The data set stored in the source is fetched and propagated through the different layers in 25 load cycles. In other words, 1 load cycle processes 1/25 of the data set.

One of the central rules of the benchmark stipulates that the memory utilization must be at least 65 percent. The permissible data volumes are a multiple of 1.3 billion initial data records.

The data load phase takes several hours and is a combination of CPU- and IO-intensive load. When several HANA nodes are used (see "SAP HANA Scale-up and Scale-out Configuration Architecture" below), significant network load is generated.

Query throughput phase

The queries for the throughput phase must be executed via an ABAP program with a variant containing 190 query steps. Users execute the set of navigation steps in random order (via asynchronous RFCs). The queries contain typical query patterns which can be found in BW productive systems of customers.

The query throughput phase runs one hour and is CPU bound. In a HANA multi-node environment, also significant network load is generated.

Query runtime phase

For the query runtime phase the same ABAP program as for the throughput phase is used with a different variant. The variant contains 10 queries which are executed sequentially. These queries are used to measure the runtime. They contain complex query patterns which are executed in BW productive systems of customers but which are typically not executed by many users in parallel but selectively by some power users. Therefore they are executed sequentially.

The query runtime phase takes a short time and generates a small load. Only a few processors cores are used, single thread performance is important for short runtimes.

Metrics of the SAP BW Edition for SAP HANA Standard Application Benchmark

The SAP BW edition for SAP HANA Benchmark Version 3 KPIs are:

- Benchmark phase 1 - data load phase:

  Number of initial records

  Runtime of last data set in seconds

- Benchmark phase 2 - query throughput phase:

  Query executions per hour

  CPU utilization of database server in percent

- Benchmark phase 3 - query runtime phase:

  Total runtime of complex query phase in seconds

SAP HANA Scale-up and Scale-out Configuration Architecture

In general, a single database node or multiple database nodes can be used for SAP benchmarks to scale the workload. In the context of SAP HANA and particularly the SAP BW Edition for SAP HANA Standard Application Benchmark it is referred to as a scale-up configuration in the case of a single database node and a scale-out configuration in the case of multi database nodes.

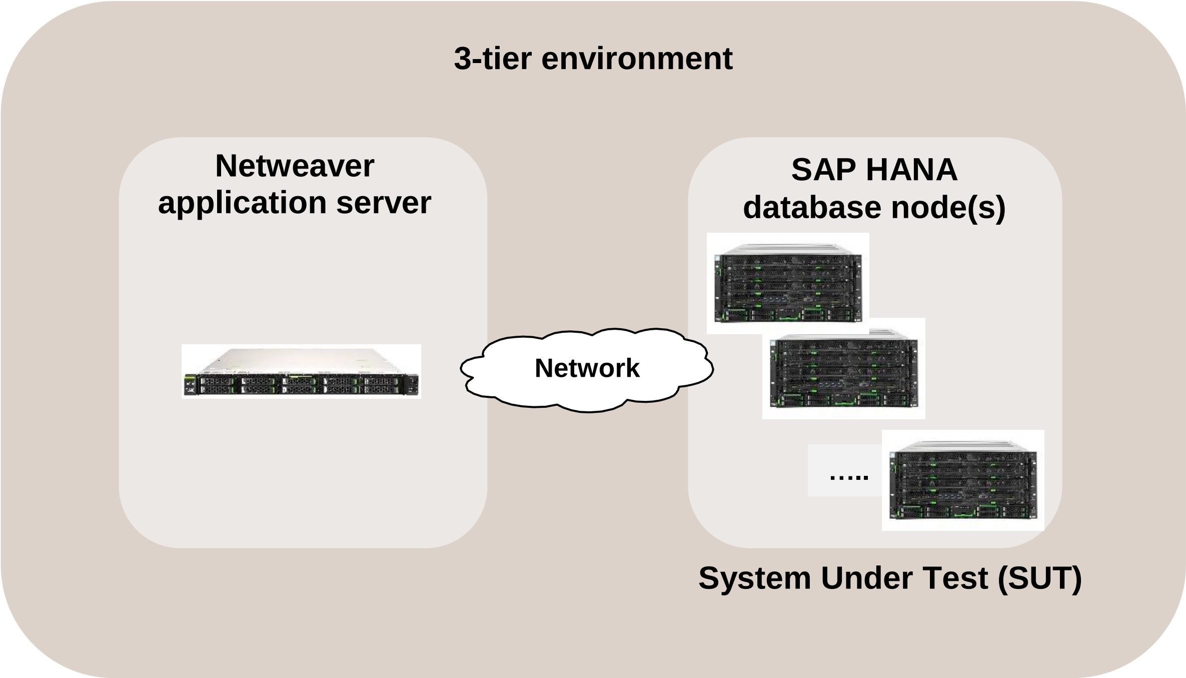

Although an application server is involved in the benchmark, neither performance metrics are measured nor does the server appear on the benchmark certificate.

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMEQUEST 3800B
Processor	Intel® Xeon® Platinum 8180 processor, 2.50 GHz, 8 processors
Memory	3,072 GB (32GB DDR4-2666 ECC × 96)
Network interface	On board LAN 1 Gbps × 1 LAN 10 Gbps × 1
Disk Subsystem	PQ 3800B : on board RAID controller PRAID EP420i 600 GB 10k rpm SAS drive × 1, RAID0 (OS), PRAID EP420e × 2 JX40 S2 × 1 : 960 GB SSD drive × 6, RAID0 (log) JX40 S2 × 1 : 960 GB SSD drive × 8, RAID0 (data) JX40 S2 × 1 : 960 GB SSD drive × 8, RAID0 (data)
Software	
Operating system	SUSE Linux Enterprise Server 12
Database	SAP HANA 2.0
Technology platform release	SAP NetWeaver 7.50
Benchmark version	Version 3

Benchmark results

Result	
Benchmark Phase 1 Number of initial records: Runtime of last Data Set (seconds):	 5,200,000,000 20,836
Benchmark Phase 2 Query Executions per Hour: CPU utilization of database server:	 9,785 97%
Benchmark Phase 3 Total Runtime of complex query phase (seconds):	 119

Certified and published SAP BWH Benchmarks are published on SAP's benchmark site here .
Certification date: 2018-11-06
Certification number: 2018047

Concluding remarks

Available now for over 20 years the CPU and memory demanding SAP SD benchmark has served well to demonstrate the strength or weakness of a hardware platform. SAP certified results will be published on SAP's web site and comprise information on machines, platform partner, SAP® Business Suite component and the number of benchmark users. Each of the tables contains basic information about the certified results. In line with SAP's full disclosure policy more detailed information about each result can be obtained from the respective hardware partner or SAP.

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 measurement set-up is symbolically illustrated below:

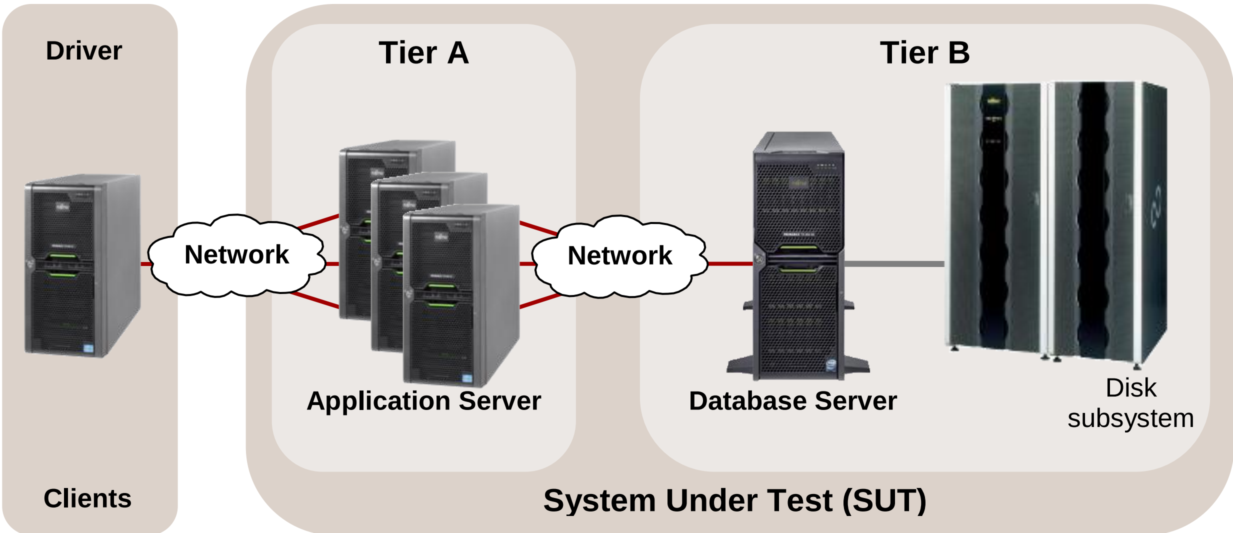

All results were determined by way of example on a PRIMEQUEST 3800E.

Database Server (Tier B)	
Hardware	
Model	PRIMEQUEST 3800E
Processor	Intel® Xeon® Processor Scalable Family
Memory	2 processors: 12 × 128GB (2x64GB) 4Rx4 DDR4-2666 3DS ECC 4 processors: 24 × 128GB (2x64GB) 4Rx4 DDR4-2666 3DS ECC 8 processors: 48 × 128GB (2x64GB) 4Rx4 DDR4-2666 3DS ECC
Network interface	2 × onboard LAN 10 Gb/s
Disk subsystem	PQ 3800E : Onboard RAID controller PRAID EP420i 1 × 300 GB 10k rpm SAS Drive, RAID0 (OS), 6 × PRAID EP420e 6 × JX40: 14 × 960 GB SSD Drive each, RAID0 (data) 5 × JX40: 14 × 400 GB SSD Drive each, RAID0 (temp) 1 × JX40: 14 × 400 GB SSD Drive each, RAID0 (temp) 8 × 900GB 10k rpm SAS Drive, RAID10 (LOG)
Software	
BIOS	PA17124
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, a configuration with a total memory of 3072 GB for the measurements with four processors and a configuration with a total memory of 6144 GB for the measurements with eight processors. The memory configurations had memory access of 2666 MHz.

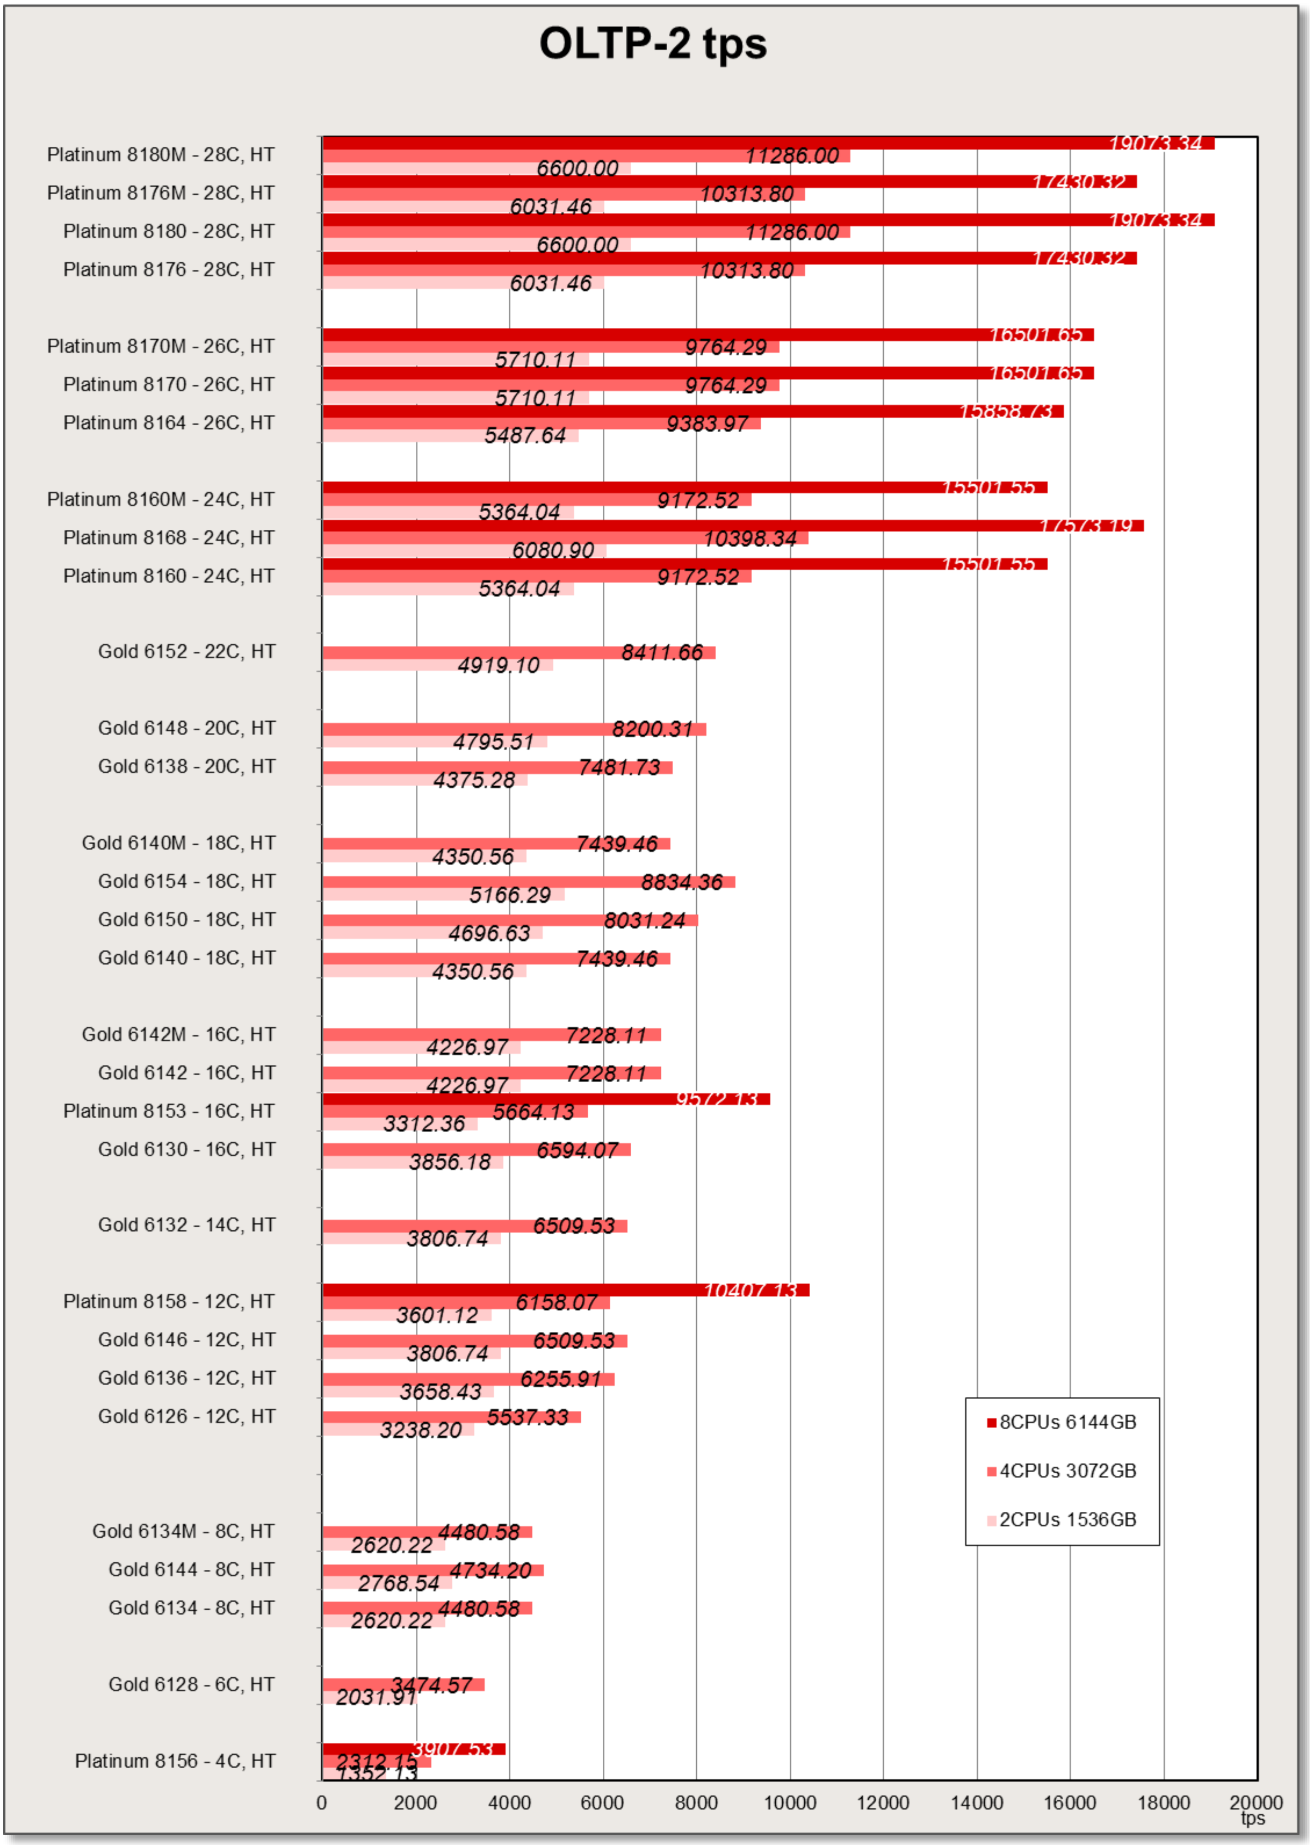
OLTP-2 tps
Platinum 8180M - 28C, HT
19073.34
11286.00
6600.00
Platinum 8176M - 28C, HT
17430.32
10313.80
6031.46
Platinum 8180 - 28C, HT
19073.34
11286.00
6600.00
Platinum 8176 - 28C, HT
17430.32
10313.80
6031.46
Platinum 8170M - 26C, HT
16501.65
9764.29
5710.11
Platinum 8170 - 26C, HT
16501.65
9764.29
5710.11
Platinum 8164 - 26C, HT
15858.73
9383.97
5487.64
Platinum 8160M - 24C, HT
15501.55
9172.52
5364.04
Platinum 8168 - 24C, HT
17573.19
10398.34
6080.90
Platinum 8160 - 24C, HT
15501.55
9172.52
5364.04
Gold 6152 - 22C, HT
8411.66
4919.10
Gold 6148 - 20C, HT
8200.31
4795.51
Gold 6138 - 20C, HT
7481.73
4375.28
Gold 6140M - 18C, HT
7439.46
4350.56
Gold 6154 - 18C, HT
8834.36
5166.29
Gold 6150 - 18C, HT
8031.24
4696.63
Gold 6140 - 18C, HT
7439.46
4350.56
Gold 6142M - 16C, HT
7228.11
4226.97
Gold 6142 - 16C, HT
7228.11
4226.97
Platinum 8153 - 16C, HT
9572.13
5664.13
3312.36
Gold 6130 - 16C, HT
6594.07
3856.18
Gold 6132 - 14C, HT
6509.53
3806.74
Platinum 8158 - 12C, HT
10407.13
6158.07
3601.12
Gold 6146 - 12C, HT
6509.53
3806.74
Gold 6136 - 12C, HT
6255.91
3658.43
Gold 6126 - 12C, HT
5537.33
3238.20
Gold 6134M - 8C, HT
4480.58
2620.22
Gold 6144 - 8C, HT
4734.20
2768.54
Gold 6134 - 8C, HT
4480.58
2620.22
Gold 6128 - 6C, HT
3474.57
2031.91
Platinum 8156 - 4C, HT
3907.53
2312.15
1352.13
0
2000
4000
6000
8000
10000
12000
14000
16000
18000
20000
tps
8CPUs 6144GB
4CPUs 3072GB
2CPUs 1536GB

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.

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 performance.

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

If you compare the maximum achievable OLTP-2 values of the current system generation with the values that were achieved on the predecessor systems, the result is an increase of about 83%.

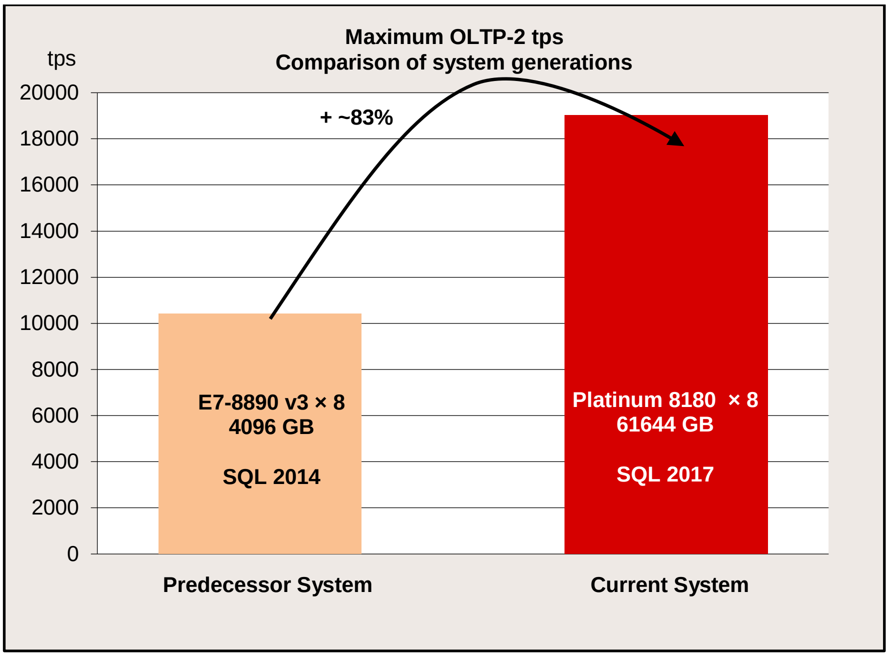

現行モデル	PQ3800B	PQ3800E
旧モデル	PQ2800B3	PQ2800E3

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). Add to these a fourth machine, the so-called idle VM. 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.

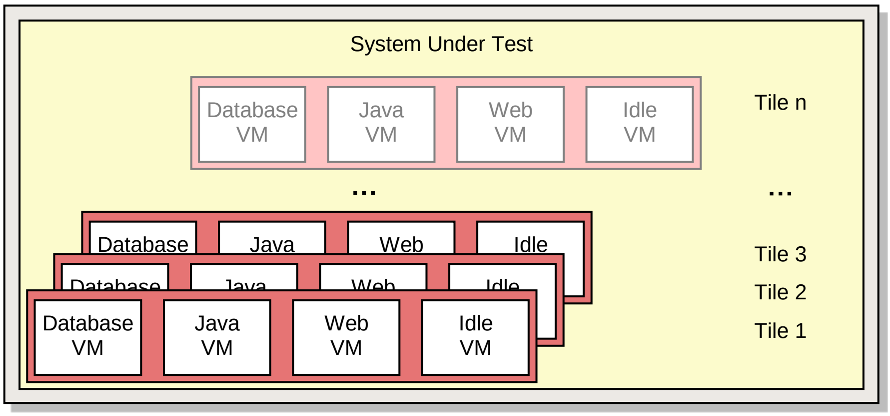

Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score, the individual benchmark results for one tile are put in relation to the respective results of a reference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

Starting as a rule with one tile, 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 measurement set-up is symbolically illustrated below:

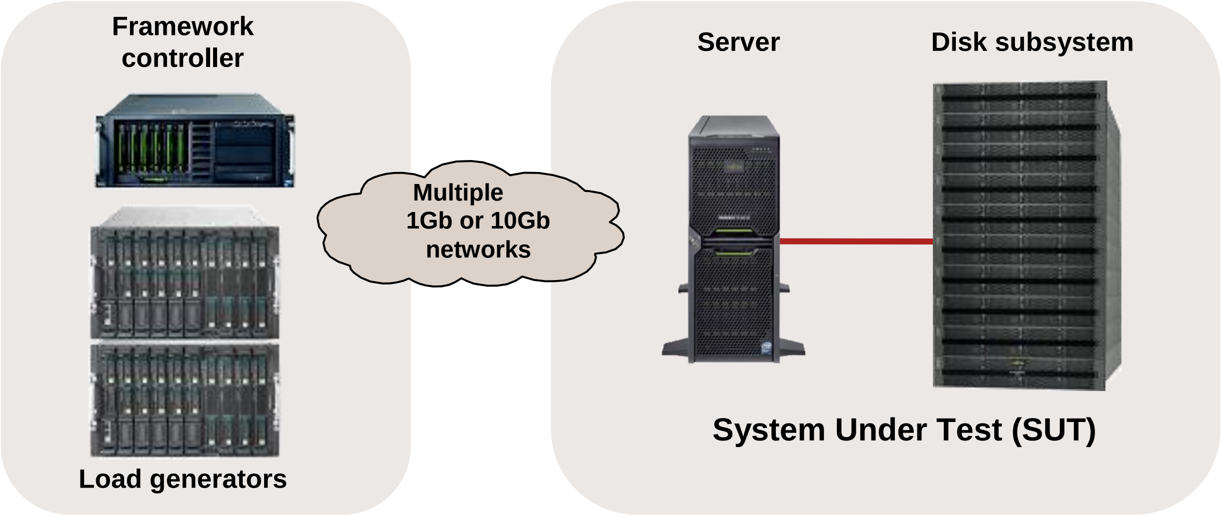

All results were determined by way of example on a PRIMEQUEST 3800B.

System Under Test (SUT)	
Hardware	
Processor	Intel® Xeon® Processor Scalable Family
Memory	4 processors: 12 × 64GB (2x32GB) 2Rx4 PC4-2666V R ECC 8 processors: 24 × 64GB (2x32GB) 2Rx4 PC4-2666V 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	256 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 PRIMEQUEST 8-socket or 4-sokcet rack 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.

This result of cursive is an estimated value.

Processor			PQ3800B				PQ3800E				PQ3400E	
			4cpus		8cpus		4cpus		8cpus		4cpus	
			Score	#Tiles	Score	#Tiles	Score	#Tiles	Score	#Tiles	Score	#Tiles
Intel® Xeon® Processor Scalable Family	4 Cores Hyper-Threading, Turbo-Modus	Platinum 8156	*23.2*	*15*	*45.6*	*29*	23.2	15	45.6	29	*23.2*	*15*
	6 Cores Hyper-Threading, Turbo-Modus	Gold 6128									*34.8*	*22*
	8 Cores Hyper-Threading, Turbo-Modus	Gold 6134									*45.4*	*24*
		Gold 6144									*48.4*	*25*
		Gold 6134M									*45.4*	*24*
	12 Cores Hyper-Threading, Turbo-Modus	Gold 6126									*63.5*	*38*
		Gold 6136									*56.3*	*40*
		Gold 6146									*64.9*	*40*
		Platinum 8158	*62.3*	*42*	*122.9*	*82*	62.3	42	122.9	82	*62.3*	*42*
	14 Cores Hyper-Threading, Turbo-Modus	Gold 6132									*64.6*	*35*
	16 Cores Hyper-Threading, Turbo-Modus	Gold 6130									*68.5*	*47*
		Platinum 8153	*58.4*	*42*	*114.5*	*96*	58.4	42	114.5	96	*58.4*	*42*
		Gold 6142									*72.0*	*48*
		Gold 6142M									*72.0*	*48*
	18 Cores Hyper-Threading, Turbo-Modus	Gold 6140									*75.5*	*52*
		Gold 6150									*81.5*	*53*
		Gold 6154									*87.9*	*54*
		Gold 6140M									*75.5*	*52*
	20 Cores Hyper-Threading, Turbo-Modus	Gold 6138									*75.9*	*48*
		Gold 6148									*82.9*	*49*
	22 Cores Hyper-Threading, Turbo-Modus	Gold 6152									*83.7*	*57*

	24 Cores Hyper-Threading, Turbo-Modus	Platinum 8160	93.0	65	182.7	123	93.0	65	182.7	123	93.0	65
		Platinum 8168	104.6	66	208.2	125	104.6	66	208.2	125	104.6	66
		Platinum 8160M	93.0	65	182.7	123	93.0	65	182.7	123	93.0	65
	26 Cores Hyper-Threading, Turbo-Modus	Platinum 8164	95.2	58	187.0	122	95.2	58	187.0	122	95.2	58
		Platinum 8170	102.7	58	200.0	122	102.7	58	200.0	122	102.7	58
		Platinum 8170M	102.7	58	200.0	122	102.7	58	200.0	122	102.7	58
	28 Cores Hyper-Threading, Turbo-Modus	Platinum 8176	104.7	68	205.5	132	104.7	68	205.5	132	104.7	68
		Platinum 8180	117.2	69	229.2	134	117.2	69	229.2	134	117.2	69
		Platinum 8176M	104.7	68	205.5	132	104.7	68	205.5	132	104.7	68
		Platinum 8180M	117.2	69	229.2	134	117.2	69	229.2	134	117.2	69

These PRIMEQUEST 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.

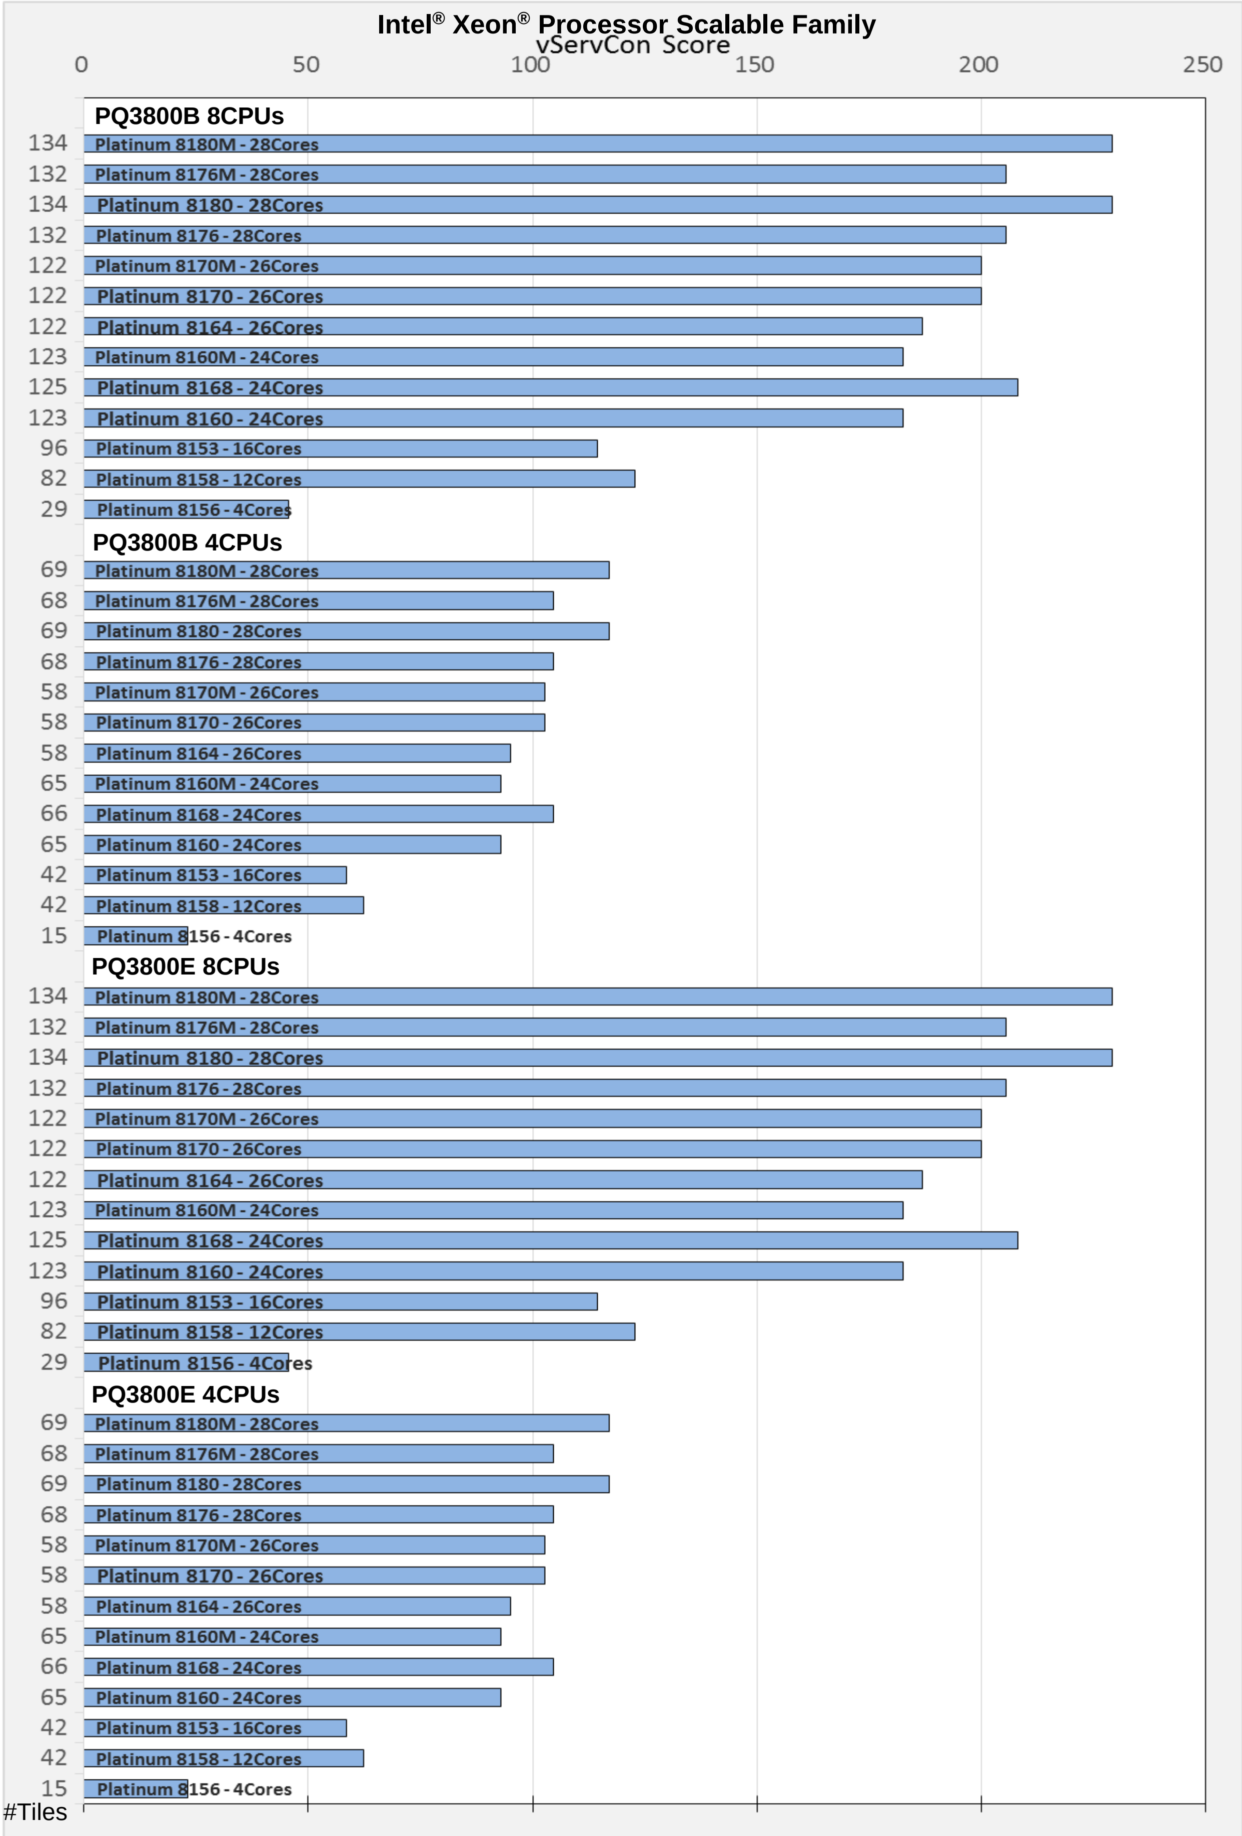

Intel® Xeon® Processor Scalable Family
vServCon Score
0
50
100
150
200
250
PQ3800B 8CPUs
134 Platinum 8180M - 28Cores
132 Platinum 8176M - 28Cores
134 Platinum 8180 - 28Cores
132 Platinum 8176 - 28Cores
122 Platinum 8170M - 26Cores
122 Platinum 8170 - 26Cores
122 Platinum 8164 - 26Cores
123 Platinum 8160M - 24Cores
125 Platinum 8168 - 24Cores
123 Platinum 8160 - 24Cores
96 Platinum 8153 - 16Cores
82 Platinum 8158 - 12Cores
29 Platinum 8156 - 4Cores
PQ3800B 4CPUs
69 Platinum 8180M - 28Cores
68 Platinum 8176M - 28Cores
69 Platinum 8180 - 28Cores
68 Platinum 8176 - 28Cores
58 Platinum 8170M - 26Cores
58 Platinum 8170 - 26Cores
58 Platinum 8164 - 26Cores
65 Platinum 8160M - 24Cores
66 Platinum 8168 - 24Cores
65 Platinum 8160 - 24Cores
42 Platinum 8153 - 16Cores
42 Platinum 8158 - 12Cores
15 Platinum 8156 - 4Cores
PQ3800E 8CPUs
134 Platinum 8180M - 28Cores
132 Platinum 8176M - 28Cores
134 Platinum 8180 - 28Cores
132 Platinum 8176 - 28Cores
122 Platinum 8170M - 26Cores
122 Platinum 8170 - 26Cores
122 Platinum 8164 - 26Cores
123 Platinum 8160M - 24Cores
125 Platinum 8168 - 24Cores
123 Platinum 8160 - 24Cores
96 Platinum 8153 - 16Cores
82 Platinum 8158 - 12Cores
29 Platinum 8156 - 4Cores
PQ3800E 4CPUs
69 Platinum 8180M - 28Cores
68 Platinum 8176M - 28Cores
69 Platinum 8180 - 28Cores
68 Platinum 8176 - 28Cores
58 Platinum 8170M - 26Cores
58 Platinum 8170 - 26Cores
58 Platinum 8164 - 26Cores
65 Platinum 8160M - 24Cores
66 Platinum 8168 - 24Cores
65 Platinum 8160 - 24Cores
42 Platinum 8153 - 16Cores
42 Platinum 8158 - 12Cores
15 Platinum 8156 - 4Cores
#Tiles

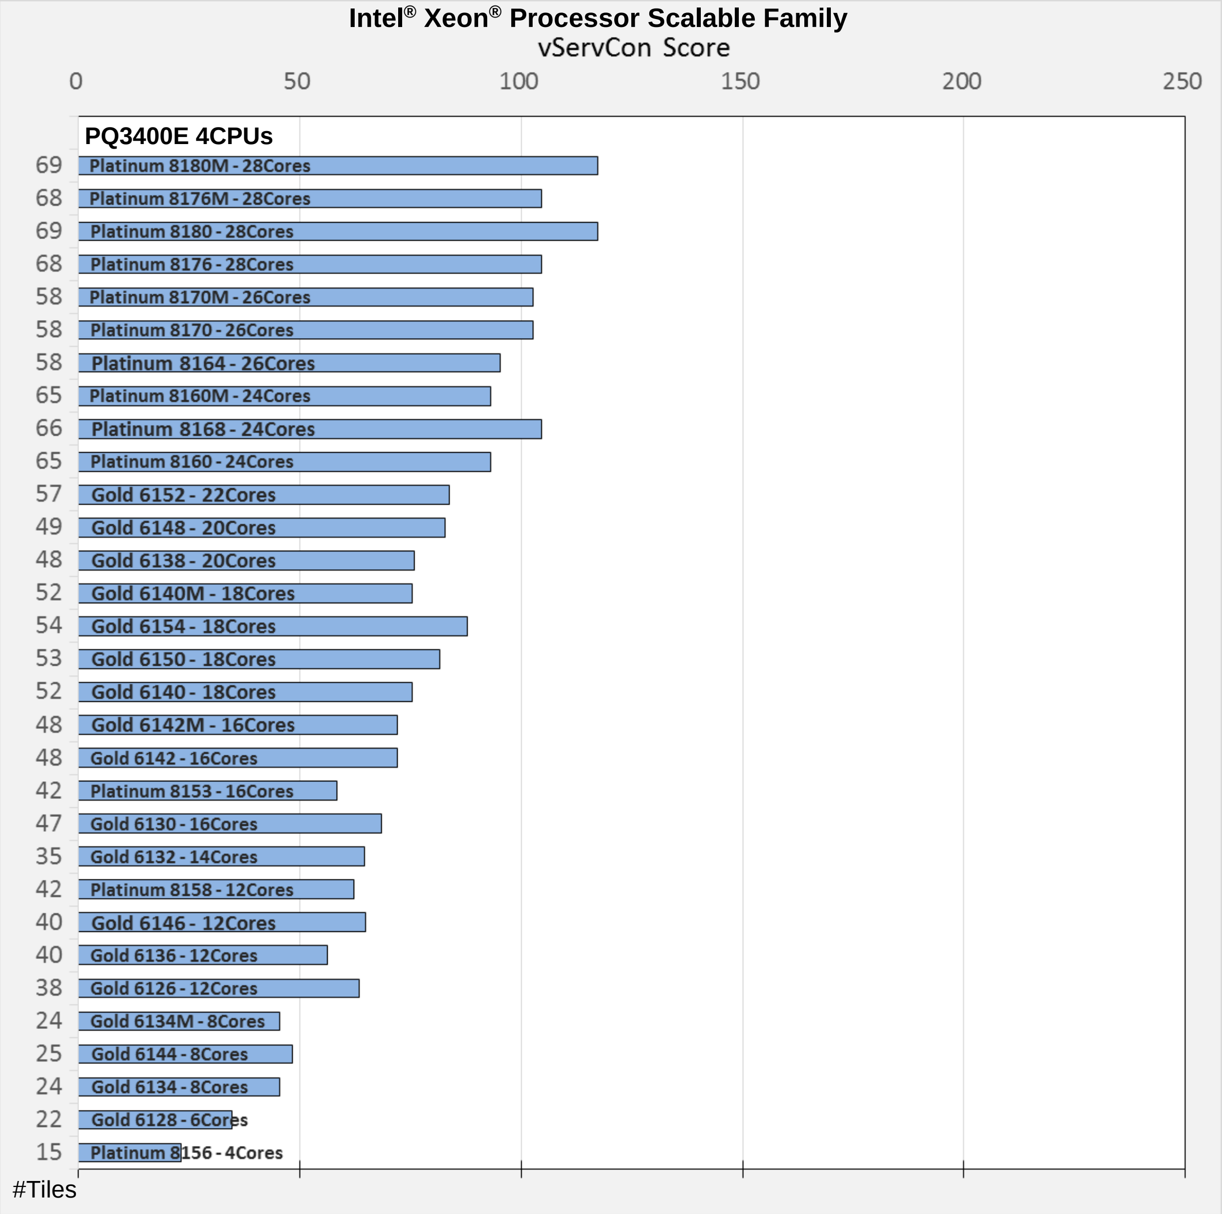

The relatively large performance differences between the processors can be explained by their features. The values scale on the basis of the number of cores, the size of the L3 cache and the CPU clock frequency and as a result of the features of Hyper-Threading and turbo mode, which are available in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

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

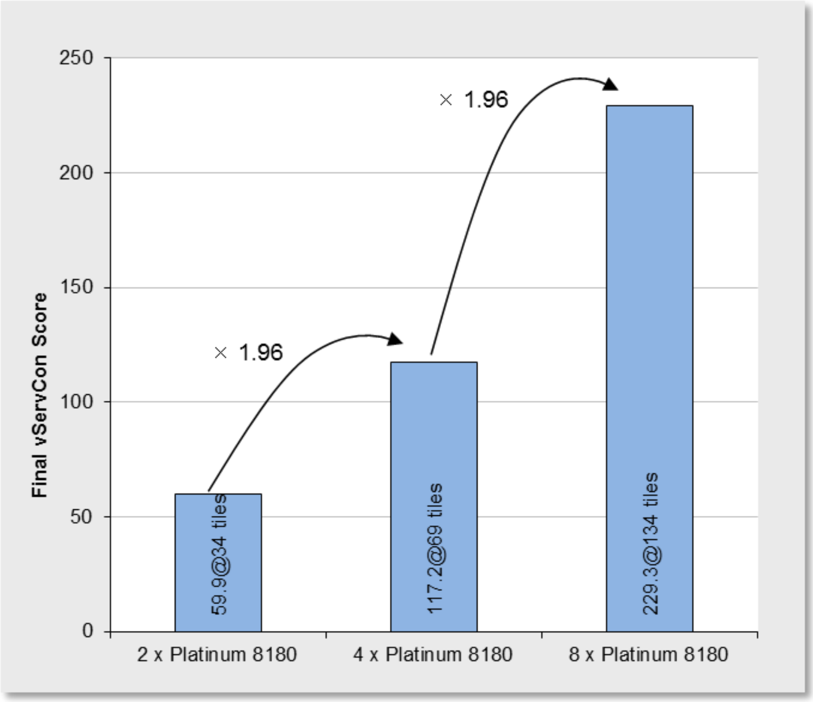

Until now we have looked at the virtualization performance of a fully configured system. However, with a server with eight sockets the question also arises as to how good performance scaling is from two to four processors. The better the scaling, the lower the overhead usually caused by the shared use of resources within a server. The scaling factor also depends on the application. If the server is used as a virtualization platform for server consolidation, the system scales with a factor of 1.96. When operated with four or eight processors, the system thus achieves twice the performance as with two or four processors, as is illustrated in the diagram opposite 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 Platinum 8156 (4-Core) processor.
In addition to the increased number of physical cores, Hyper-Threading, which is supported by all Intel® Xeon® Processor Scalable Family processors, 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.

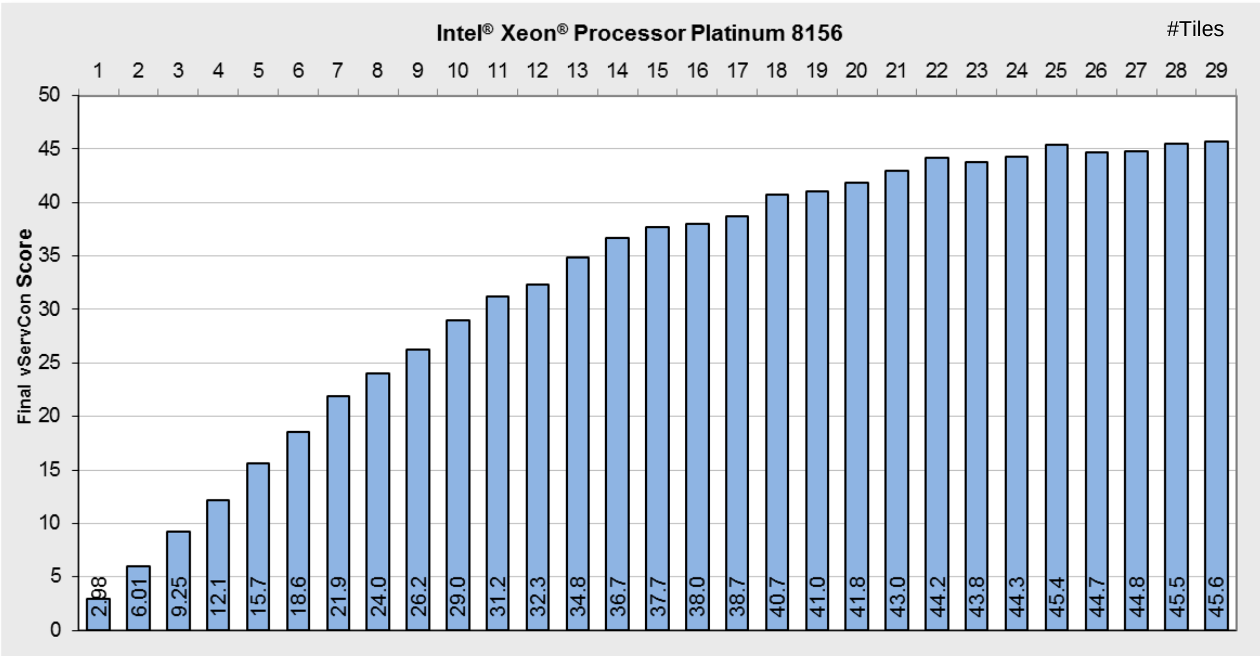

The scaling curves for the number of tiles as seen in the previous diagram are specifically for systems with Hyper-Threading. 4 physical and thus 8 logical cores are available with the Xeon Platinum 8156 processors; approximately four of them are used per tile (see Benchmark description). So, with a server with eight Xeon Platinum 8156 Processors can use 64 logical cores. This means that a parallel use of the same physical cores by several VMs is avoided up to a maximum of about 16 tiles. That is why the performance curve in this range scales almost ideal. For the quantities above the growth is flatter up to CPU full utilization.

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 Platinum 8156 situation with 87 application VMs (29 tiles, not including the idle VMs); the low load case is represented by three application VMs (one tile, not including the idle VM). Remember: the vServCon score for one tile is an average value across the three application scenarios in vServCon. This average performance of one tile drops when changing from the low load case to the total optimum of the vServCon score – from 2.98 to 45.6/29=1.57, i.e. to 53%. 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.

STREAM

Benchmark description

STREAM is a synthetic benchmark that has been used for many years to determine memory throughput and which 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, 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	PRIMEQUEST 3800E/3800B/3400E
Processor	Intel® Xeon® Processor Scalable Family
Memory	4 processors: 24 ×32GB (2x16GB) 2Rx4 PC4-2666V R ECC 8 processors: 48 ×32GB (2x16GB) 2Rx4 PC4-2666V R ECC
Software	
BIOS settings	Intel Virtualization Technology = Disabled Enable CPU HWPM = Disabled Override OS Energy Performance = Enabled Utilization Profile = Unbalanced 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 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/soft_watchdog echo 0 > /proc/sys/kernel/numa_balancing echo 1 > /proc/sys/vm/drop_caches ulimit -s unlimited
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 result of cursive is an estimated value from the result of PRIMEQUEST3800B.

Processor	Memory Frequency [MHz]	Max. Memory Bandwidth [GB/s]	Cores	Processor Frequency [GHz]	PQ3800E PQ3800B Number of Processors	PQ3800E PQ3800B TRIAD [GB/s]	PQ3400E Number of Processors	PQ3400E TRIAD [GB/s]
Xeon Gold 6130	2666	128.0	16	2.1			4	*355*
Xeon Gold 6140	2666	128.0	18	2.3			4	*362*
Xeon Gold 6138	2666	128.0	20	2.0			4	*372*
Xeon Gold 6148	2666	128.0	20	2.4			4	*372*
Xeon Gold 6152	2666	128.0	22	2.1			4	*376*
Xeon Platinum 8153	2666	128.0	16	2.0	8	585	4	*357*
Xeon Platinum 8160	2666	128.0	24	2.1	8	687	4	*381*
Xeon Platinum 8164	2666	128.0	26	2.0	8	715	4	*383*
Xeon Platinum 8170	2666	128.0	26	2.1	8	687	4	*381*
Xeon Platinum 8176	2666	128.0	28	2.1	8	678	4	*381*
Xeon Platinum 8180	2666	128.0	28	2.5	8	728	4	381
Xeon Gold 6128	2666	128.0	6	3.4			4	*282*
Xeon Gold 6134	2666	128.0	8	3.2			4	*331*
Xeon Gold 6144	2666	128.0	8	3.5			4	*325*
Xeon Gold 6126	2666	128.0	12	2.6			4	*340*
Xeon Gold 6136	2666	128.0	12	3.0			4	*366*
Xeon Gold 6146	2666	128.0	12	3.2			4	*366*
Xeon Gold 6132	2666	128.0	14	2.6			4	*337*
Xeon Gold 6142	2666	128.0	16	2.6			4	*347*
Xeon Gold 6150	2666	128.0	18	2.7			4	*364*
Xeon Gold 6154	2666	128.0	18	3.0			4	*362*
Xeon Platinum 8156	2666	128.0	4	3.6	8	301	4	*158*
Xeon Platinum 8158	2666	128.0	12	3.0	8	578	4	*302*
Xeon Platinum 8168	2666	128.0	24	2.7	8	708	4	*378*
Xeon Gold 6134M	2666	128.0	8	3.2			4	*331*
Xeon Gold 6140M	2666	128.0	18	2.3			4	*362*
Xeon Gold 6142M	2666	128.0	16	2.6			4	*347*
Xeon Platinum 8160M	2666	128.0	24	2.1	8	*687*	4	*381*
Xeon Platinum 8170M	2666	128.0	26	2.1	8	*687*	4	*381*
Xeon Platinum 8176M	2666	128.0	28	2.1	8	*678*	4	*381*
Xeon Platinum 8180M	2666	128.0	28	2.5	8	*728*	4	*381*

The following diagram illustrates the throughput of the PRIMEQUEST 3800E in comparison to its predecessor, the PRIMEQUEST 2800E3

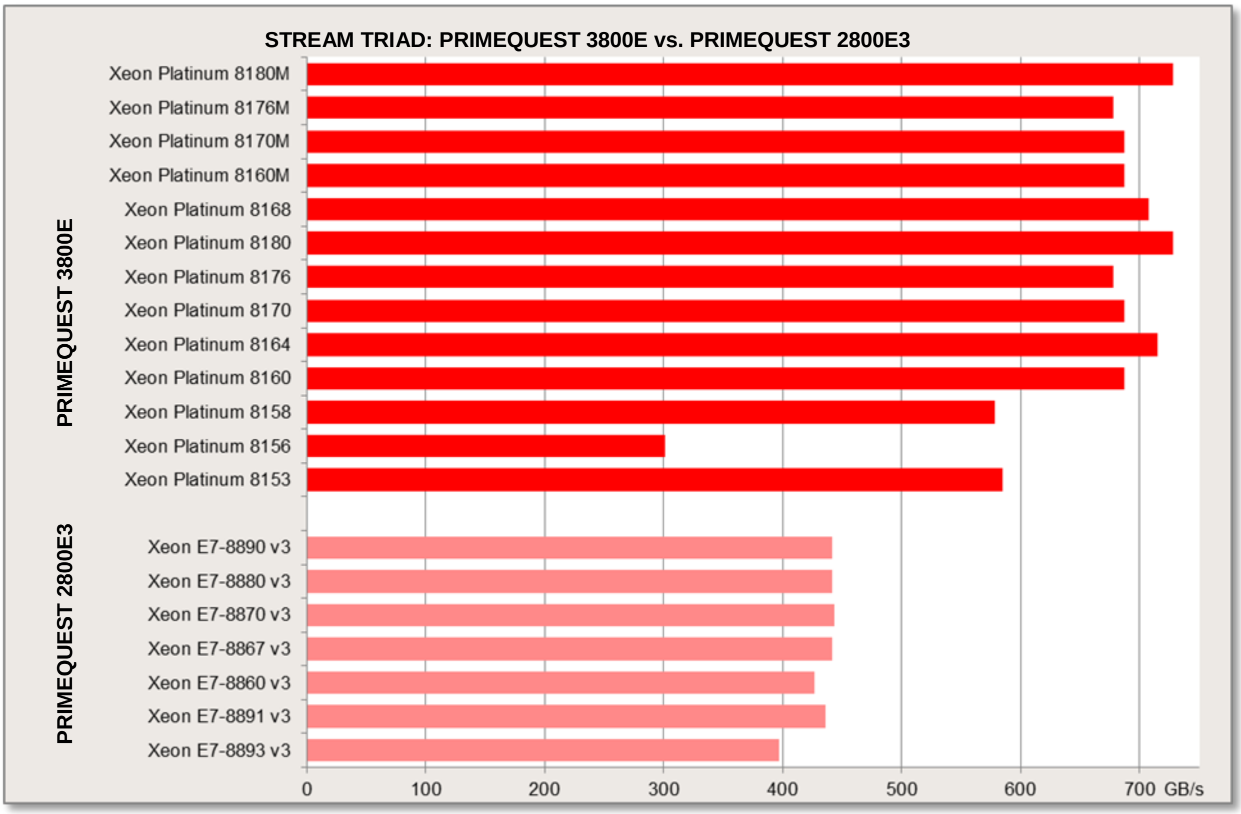

LINPACK

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 $^2/_3 n^3 + 2n^2$. 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	PRIMEQUEST 3800E/3800B/3400E
Processor	Intel® Xeon® Processor Scalable Family
Memory	4 processors: 24 ×32GB (2x16GB) 2Rx4 PC4-2666V R ECC 8 processors: 48 ×32GB (2x16GB) 2Rx4 PC4-2666V R ECC
Software	
BIOS settings	HyperThreading = Disabled Intel Virtualization Technology = Disabled Enable CPU HWPM = Disabled DCU Streamer Prefetcher = Disabled LLC Dead Line Alloc = Disabled Stale AtoS = Enabled Sub NUMA Clustering = Disabled Fan Control = Full
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 from the result of PRIMEQUEST 3800B.

Processor	Cores	Processor Frequency [GHz]	Number of Processors	Rpeak [GFlops]	PQ3800E/PQ3800B	
					Rmax [GFlops]	Efficiency
Xeon Platinum 8153	16	2.0	8	8,192	**6,195**	76%
Xeon Platinum 8160	24	2.1	8	12,904	**9,494**	74%
Xeon Platinum 8164	26	2.0	8	13,312	**9,773**	73%
Xeon Platinum 8170	26	2.1	8	13,976	**10,288**	74%
Xeon Platinum 8176	28	2.1	8	15,052	**10,946**	72%
Xeon Platinum 8180	28	2.5	8	17,920	**12,922**	72%
Xeon Platinum 8156	4	3.6	8	3,686	**2,968**	81%
Xeon Platinum 8158	12	3.0	8	9,216	**7,123**	77%
Xeon Platinum 8168	24	2.7	8	16,588	**11,523**	69%
Xeon Platinum 8160M	24	2.1	8	12,904	**9,494**	74%
Xeon Platinum 8170M	26	2.1	8	13,976	**10,288**	73%
Xeon Platinum 8176M	28	2.1	8	15,052	**10,946**	72%
Xeon Platinum 8180M	28	2.5	8	17,920	**12,922**	72%

Processor	Cores	Processor Frequency [GHz]	Number of Processors	Rpeak [GFlops]	PQ3400E	
					Rmax [GFlops]	Efficiency
Xeon Gold 6130	16	2.1	4	4,300	***3,541***	82%
Xeon Gold 6140	18	2.3	4	5,300	***3,952***	75%
Xeon Gold 6138	20	2.0	4	5,120	***3,776***	74%
Xeon Gold 6148	20	2.4	4	6,144	***4,324***	70%
Xeon Gold 6152	22	2.1	4	5,914	***4,265***	72%
Xeon Platinum 8153	16	2.0	4	4,096	***3,025***	74%
Xeon Platinum 8160	24	2.1	4	6,452	***4,637***	72%
Xeon Platinum 8164	26	2.0	4	6,656	***4,840***	73%
Xeon Platinum 8170	26	2.1	4	6,988	***5,326***	76%
Xeon Platinum 8176	28	2.1	4	7,526	***5,437***	72%
Xeon Platinum 8180	28	2.5	4	8,960	***6,525***	73%
Xeon Gold 6128	6	3.4	4	2,612	***1,937***	74%
Xeon Gold 6134	8	3.2	4	3,276	***2,485***	76%
Xeon Gold 6144	8	3.5	4	3,584	***2,543***	71%
Xeon Gold 6126	12	2.6	4	3,994	***3,052***	76%

Xeon Gold 6136	12	3.0	4	4,608	***3,483***	76%
Xeon Gold 6146	12	3.2	4	4,916	***3,678***	75%
Xeon Gold 6132	14	2.6	4	4,660	***3,698***	79%
Xeon Gold 6142	16	2.6	4	5,324	***4,089***	77%
Xeon Gold 6150	18	2.7	4	6,220	***4,383***	70%
Xeon Gold 6154	18	3.0	4	6,912	***5,283***	76%
Xeon Platinum 8156	4	3.6	4	1,843	***1,499***	81%
Xeon Platinum 8158	12	3.0	4	4,608	***3,597***	78%
Xeon Platinum 8168	24	2.7	4	8,294	***5,375***	65%
Xeon Gold 6134M	8	3.2	4	3,276	***2,485***	76%
Xeon Gold 6140M	18	2.3	4	5,300	***3,952***	75%
Xeon Gold 6142M	16	2.6	4	5,324	***4,089***	77%
Xeon Platinum 8160M	24	2.1	4	6,452	***4,637***	72%
Xeon Platinum 8170M	26	2.1	4	6,988	***5,326***	76%
Xeon Platinum 8176M	28	2.1	4	7,526	***5,437***	72%
Xeon Platinum 8180M	28	2.5	4	8,960	***6,525***	73%

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 does not as a matter of principle 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

PRIMEQUEST Servers

http://ts.fujitsu.com/primequest

PRIMEQUEST 3800E/3800B/3400E

This White Paper:
http://docs.ts.fujitsu.com/dl.aspx?id=984d7a1d-15f3-4d94-9d31-5b0f6af37e4c

Data sheet
PQ3800E: http://docs.ts.fujitsu.com/dl.aspx?id=e9e76d18-d315-4208-846b-1ea19b12b2d8
PQ3800B: http://docs.ts.fujitsu.com/dl.aspx?id=5987ba8e-e786-449a-bc90-725e3a239fc8
PQ3400E: http://docs.ts.fujitsu.com/dl.aspx?id=757a1ab9-c35c-4ab7-8dcd-9e8941922b3c

PRIMEQUEST 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

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