

Fujitsu Server PRIMEQUEST Performance Report PRIMEQUEST 3800B2/3800E2

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMEQUEST 3800B2/3800E2.

Explains PRIMEQUEST 3800B2/3800E2 performance data in comparison to other PRIMEQUEST models. In addition to the benchmark results, the explanation for each benchmark and benchmark environment are also included.

Version
1.3
2023-10-17



Contents

- Technical data 3**
 - Suffix of Processor number shows additional feature of Xeon Processor.....6
- SPEC CPU2017 7**
 - Benchmark description7
 - Benchmark environment8
 - Benchmark results..... 10
- STREAM 12**
 - Benchmark description 12
 - Benchmark environment 13
 - Benchmark results..... 14
- SPECjbb2015 16**
 - Benchmark description 16
 - Benchmark environment 17
 - Benchmark results..... 18
- SAP SD 19**
 - Benchmark description 19
 - Benchmark environment 19
 - Benchmark results..... 21
- Disk I/O: Performance of storage media 22**
 - Benchmark description 22
 - Benchmark environment 23
 - Benchmark results..... 25
 - Storage media 26
- OLTP-2 27**
 - Benchmark description 27
 - Benchmark environment 27
 - Benchmark results..... 29
- vServCon 33**
 - Benchmark description 33
 - Benchmark environment 34
 - Benchmark results..... 35
- VMmark V3..... 38**
 - Benchmark description 38
 - Benchmark environment 40
 - Benchmark results..... 43
- Literature 44**

Technical data

PRIMEQUEST 3800B2



PRIMEQUEST 3800E2



Decimal prefixes according to the SI standard are used for measurement units in this white paper (e.g. 1 GB = 10⁹ bytes). In contrast, these prefixes should be interpreted as binary prefixes (e.g. 1 GB = 2³⁰ bytes) for the capacities of caches and memory modules. Separate reference will be made to any further exceptions where applicable.

Model	PRIMEQUEST 3800B2	PRIMEQUEST 3800E2
Form factor	Rack server	
Number of system boards orderable	1 – 4	
Number of I/O units orderable	1 - 2	1 - 4
Number of disk units orderable	0 – 2	0 – 6
Per system board:		
Chipset	Intel C624 Chipset	
Number of sockets	2	
Number of processors orderable	2	
Processor type	2nd Generation Intel Xeon Scalable Processors Family	
Number of memory slots	24 (12 per processor)	
Maximum memory configuration	6,144GB	
Per I/O unit:		
Onboard LAN controller	none	2 x 10GBASE-T LAN ports
PCI slots	4 x PCI-Express 3.0 x 8 4 x PCI-Express 3.0 x16	1 x PCI-Express 3.0 x 8 1 x PCI-Express 3.0 x16
Per disk unit:		
Max. number of internal hard disks	4	

Processors (since system release)								
Processor	cores	Number of threads	L3 Cache [MB]	UPI speed [GT/s]	Rated frequency [Ghz]	Maximum turbo frequency [Ghz]	Maximum memory frequency [MHz]	TDP [Watt]
Xeon Platinum 8280L	28	56	38.5	10.4	2.7	4.0	2,933	205
Xeon Platinum 8280M	28	56	38.5	10.4	2.7	4.0	2,933	205
Xeon Platinum 8280	28	56	38.5	10.4	2.7	4.0	2,933	205
Xeon Platinum 8276L	28	56	38.5	10.4	2.3	4.0	2,933	165
Xeon Platinum 8276M	28	56	38.5	10.4	2.3	4.0	2,933	165
Xeon Platinum 8276	28	56	38.5	10.4	2.3	4.0	2,933	165
Xeon Platinum 8270	26	52	35.8	10.4	2.7	4.0	2,933	205
Xeon Platinum 8268	24	48	35.8	10.4	2.9	3.9	2,933	205
Xeon Platinum 8260L	24	48	35.8	10.4	2.4	3.9	2,933	165
Xeon Platinum 8260M	24	48	35.8	10.4	2.4	3.9	2,933	165
Xeon Platinum 8260	24	48	35.8	10.4	2.4	3.9	2,933	165
Xeon Platinum 8256	4	8	16.5	10.4	3.8	3.9	2,933	105
Xeon Platinum 8253	16	32	22.0	10.4	2.2	3.0	2,933	125
Xeon Gold 6262V	24	48	33.0	10.4	1.9	3.6	2,933	130
Xeon Gold 6254	18	36	24.8	10.4	3.1	4.0	2,933	200
Xeon Gold 6252	24	48	33.0	10.4	2.1	3.7	2,933	150
Xeon Gold 6248	20	40	27.5	10.4	2.5	3.9	2,933	150
Xeon Gold 6246	12	24	24.8	10.4	3.3	4.2	2,933	165
Xeon Gold 6244	8	16	11.0	10.4	3.6	4.4	2,933	150
Xeon Gold 6242	16	32	22.0	10.4	2.8	3.9	2,933	150
Xeon Gold 6240L	18	36	24.8	10.4	2.6	3.9	2,933	150
Xeon Gold 6240M	18	36	24.8	10.4	2.6	3.9	2,933	150
Xeon Gold 6240	18	36	24.8	10.4	2.6	3.9	2,933	140
Xeon Gold 6238M	22	44	30.3	10.4	2.1	3.7	2,933	140
Xeon Gold 6238L	22	44	30.3	10.4	2.1	3.7	2,933	140
Xeon Gold 6238	22	44	30.3	10.4	2.1	3.7	2,933	140
Xeon Gold 6234	8	16	24.8	10.4	3.4	4.0	2,933	130
Xeon Gold 6230	20	40	27.5	10.4	2.1	3.9	2,933	125
Xeon Gold 6226	12	24	19.3	10.4	2.8	3.7	2,933	125
Xeon Gold 6222V	20	40	27.5	10.4	1.8	3.6	2,400	115

Model	PRIMEQUEST 3800B2	PRIMEQUEST 3800E2
Supported CPU	Xeon Platinum 8280L	Xeon Platinum 8280L
	Xeon Platinum 8280M	Xeon Platinum 8280M
	Xeon Platinum 8280	Xeon Platinum 8280
	Xeon Platinum 8276L	Xeon Platinum 8276L
	Xeon Platinum 8276M	Xeon Platinum 8276M
	Xeon Platinum 8276	Xeon Platinum 8276
	Xeon Platinum 8270	Xeon Platinum 8270
	Xeon Platinum 8268	Xeon Platinum 8268
	Xeon Platinum 8260L	Xeon Platinum 8260L
	Xeon Platinum 8260M	Xeon Platinum 8260M
	Xeon Platinum 8260	Xeon Platinum 8260
	Xeon Platinum 8256	Xeon Platinum 8256
	Xeon Platinum 8253	Xeon Platinum 8253
		*Xeon Gold 6262V
		Xeon Gold 6254
		*Xeon Gold 6252
		Xeon Gold 6248
	*Xeon Gold 6246	
	Xeon Gold 6244	
	Xeon Gold 6242	
	*Xeon Gold 6240L	
	*Xeon Gold 6240M	
	Xeon Gold 6240	
	*Xeon Gold 6238M	
	*Xeon Gold 6238L	
	*Xeon Gold 6238	
	*Xeon Gold 6234	
	Xeon Gold 6230	
	*Xeon Gold 6226	
	*Xeon Gold 6222V	

All the processors that can be ordered with the PRIMEQUEST 3800B2/3800E2 support Intel Turbo Boost Technology 2.0. This technology allows you to operate the processor with higher frequencies than the nominal frequency. Listed in the processor table is "Max. Turbo Frequency" for the theoretical maximum frequency with only one active core per processor. The maximum frequency that can actually be achieved depends on the number of active cores, the current consumption, electrical power consumption, and the temperature of the processor.

As a matter of principle, Intel does not guarantee that the maximum turbo frequency can be reached. This is related to manufacturing tolerances, which result in a variance regarding the performance of various examples of a processor model. The range of the variance covers the entire scope between the nominal frequency and the maximum turbo frequency.

The turbo functionality can be set via BIOS option. Fujitsu generally recommends leaving the "Turbo Mode" option set at the standard setting of "Enabled", as performance is substantially increased by the higher frequencies. However, since the higher frequencies depend on general conditions and are not always guaranteed, it can be advantageous to disable the "Turbo Mode" option for application scenarios with intensive use of AVX instructions and a high number of instructions per clock unit, as well as for those that require constant performance or lower electrical power consumption.

Suffix of Processor number shows additional feature of Xeon Processor.

The processors with M/L suffix support larger memory capacity of 2TB/socket(M-suffix) or 4.5TB/socket(L-suffix) whereas normal processors support 1TB/socket memory capacity.

The processors with V suffix are specifically designed to help maximize \$/VM

Suffix	Additional feature
M	Support up to 2TB/socket memory
L	Support up to 4.5TB/socket memory
V	VM Density Optimized

Memory modules (since system release)								
Memory module	Capacity [GB]	Ranks	Bit width of the memory chips	Frequency [MHz]	Load reduced	Registered	NVDIMM	ECC
32 GB (2x16 GB) 1Rx4 DDR4-2933 R ECC	32	1	4	2,933		✓		✓
64 GB (2x32 GB) 2Rx4 DDR4-2933 R ECC	64	2	4	2,933		✓		✓
128 GB (2x64 GB) 2Rx4 DDR4-2933 R ECC	128	2	4	2,933		✓		✓
128 GB (2x64 GB) 4Rx4 DDR4-2933 LR ECC	128	4	4	2,933	✓	✓		✓
256 GB (2x128 GB) 4Rx4 DDR4-2933 LR ECC	256	4	4	2,933	✓	✓		✓
256 GB (2x128 GB) 8Rx4 DDR4-2933 LR 3DS	256	8	4	2,933	✓	✓		✓
512 GB (2x256 GB) 8Rx4 DDR4-2933 LR 3DS	512	8	4	2,933	✓	✓		✓
128GB (1x128GB) DCPMM-2666	128			2,666			✓	✓
256GB (1x256GB) DCPMM-2666	256			2,666			✓	✓
512GB (1x512GB) DCPMM-2666	512			2,666			✓	✓

Power supplies (since system release)	Max. number
Modular PSU 2,200 W 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 3800B2/3800E2.

SPEC CPU2017

Benchmark description

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

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

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

Benchmark	Number of single benchmarks	Arithmetics	Type	Compiler optimization	Measurement result
SPECSpeed2017_int_peak	10	integer	peak	aggressive	Speed
SPECSpeed2017_int_base	10	integer	base	conservative	
SPECrate2017_int_peak	10	integer	peak	aggressive	Throughput
SPECrate2017_int_base	10	integer	base	conservative	
SPECSpeed2017_fp_peak	10	floating point	peak	aggressive	Speed
SPECSpeed2017_fp_base	10	floating point	base	conservative	
SPECrate2017_fp_peak	13	floating point	peak	aggressive	Throughput
SPECrate2017_fp_base	13	floating point	base	conservative	

The measurement results are the geometric average from normalized ratio values which have been determined for individual benchmarks. The geometric average - in contrast to the arithmetic average - means that there is a weighting in favor of the lower individual results. "Normalized" means that the measurement is how fast is the test system compared to a reference system. For example, value "1" was defined for the SPECSpeed2017_int_base, SPECrate2017_int_base, SPECSpeed2017_fp_base, and SPECrate2017_fp_base results of the reference system. A SPECSpeed2017_int_base value of 2 means that the measuring system has handled this benchmark twice as fast as the reference system. A SPECrate2017_fp_base value of 4 means that the measuring system has handled this benchmark about 4/[# base copies] times faster than the reference system. "# base copies" specifies how many parallel instances of the benchmark have been executed.

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

Benchmark environment

System Under Test (SUT)

Hardware

• Model	PRIMEQUEST 3800B2/3800E2
• Processor	2nd Generation Intel Xeon Scalable Processors Family
• Memory	48 × 64 GB (2x32 GB) 2Rx4 PC4-2933Y-R

Software

• BIOS settings	<p>SPECSpeed2017_int_base:</p> <ul style="list-style-type: none"> • Fan Control = Full • Hyper-Threading = Disabled • LLC Deadline Alloc = Disabled • LLC Prefetcher = Enabled • Sub NUMA Clustering = Enabled • UPI Link L0p = Disabled • UPI Link L1 = Disabled <p>SPECSpeed2017_fp_base:</p> <ul style="list-style-type: none"> • Adjacent Cache Line Prefetch = Disabled • Fan Control = Full • Hyper-Threading = Disabled • Intel Virtualization Technology = Disabled • P-State Coordination = SW_ANY • Stale AtoS = Enabled • Sub NUMA Clustering = Enabled • Uncore Frequency Scaling = Disabled <p>SPECrate2017_int_base:</p> <ul style="list-style-type: none"> • DCU Streamer Prefetcher = Disabled • DDR4 Write Data CRC Protection = Disabled*1 • Hardware Prefetcher = Disabled*2 • HWPM Support = Native Mode with no legacy • Stale AtoS = Enabled • Sub Numa Clustering = Enabled*1 • Uncore Frequency Scaling = Disabled • UPI Link L0p = Disabled • XPT Prefetch = Enabled • Fan Control = Full <p>SPECrate2017_fp_base:</p> <ul style="list-style-type: none"> • DCU Streamer Prefetcher = Disabled • DDR4 Write Data CRC Protection = Disabled*1 • Hardware Prefetcher = Disabled*2 • LLC Dead Line Alloc = Enabled • Sub Numa Clustering = Enabled*1 • Uncore Frequency Scaling = Disabled • UPI Link L0p = Disabled • Fan Control = Full <p>*1: E-Model *2: Xeon Platinum 8280L/8280M/8280, Xeon Platinum 8270</p>
-----------------	--

• Operating system	SPECspeed2017_int, SPECrate2017_int , SPECrate2017_fp: SUSE Linux Enterprise Server 15 4.12.14-25.28-default SPECspeed2017_fp: Red Hat Enterprise Linux Server release 7.6 3.10.0-957.el7.x86_64
• Operating system settings	Stack size set to unlimited using "ulimit -s unlimited" Kernel Boot Parameter set with : nohz_full=1-xxx
• Compiler	C/C++: Version 2023.0 of Intel C/C++ Compiler for Linux Fortran: Version 2023.0 of Intel Fortran Compiler for Linux

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

Benchmark results

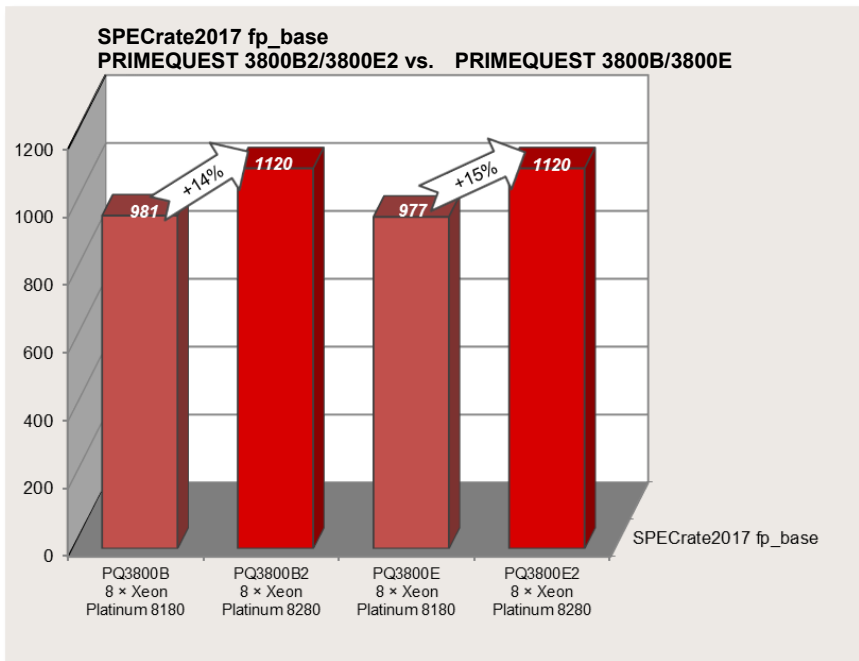
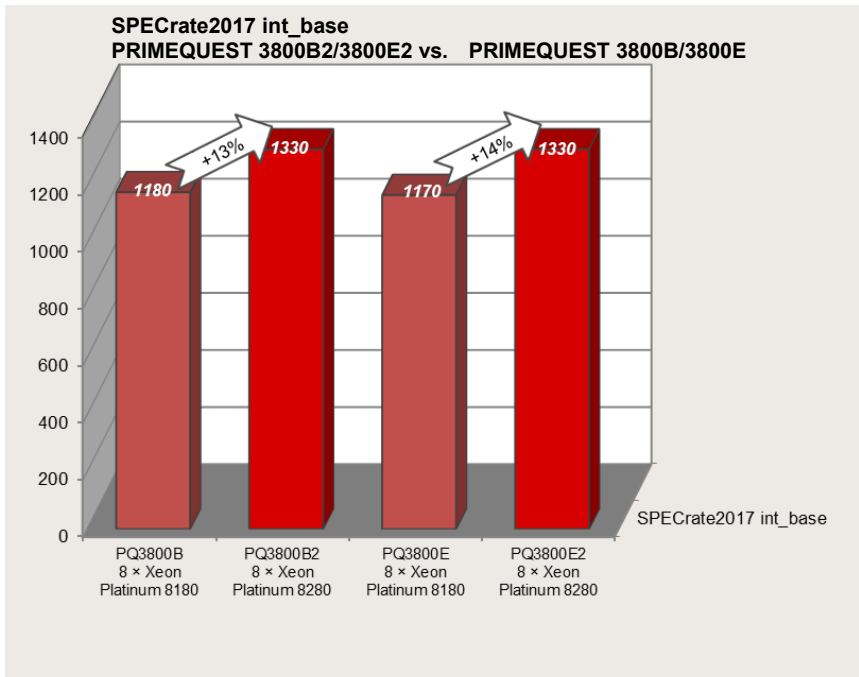
In terms of processors, the benchmark result depends primarily on the size of the processor cache, the support for Hyper-Threading, the number of processor cores, and the processor frequency. In the case of processors with Turbo mode, the number of cores, which are loaded by the benchmark, determines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

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

Processor	Number of processors	PQ3800B2		PQ3800E2	
		SPECrate2017 int_base	SPECrate2017 fp_base	SPECrate2017 int_base	SPECrate2017 fp_base
Xeon Platinum 8280L	8	1,330	1,120	1,330 est.	1,120 est.
Xeon Platinum 8280M	8	1,330 est.	1,120 est.	1,330 est.	1,120 est.
Xeon Platinum 8280	8	1,330 est.	1,120 est.	1,330 est.	1,120 est.
Xeon Platinum 8276L	8	1,188 est.	1,030 est.	1,188 est.	1,030 est.
Xeon Platinum 8276M	8	1,188 est.	1,030 est.	1,188 est.	1,030 est.
Xeon Platinum 8276	8	1,188 est.	1,030 est.	1,188 est.	1,030 est.
Xeon Platinum 8270	8	1,253 est.	1,058 est.	1,253 est.	1,058 est.
Xeon Platinum 8268	8	1,215 est.	1,040 est.	1,215 est.	1,040 est.
Xeon Platinum 8260L	8	1,088 est.	971 est.	1,088 est.	971 est.
Xeon Platinum 8260M	8	1,088 est.	971 est.	1,088 est.	971 est.
Xeon Platinum 8260	8	1,088 est.	971 est.	1,088 est.	971 est.
Xeon Platinum 8256	8	253 est.	307 est.	253 est.	307 est.
Xeon Platinum 8253	8	659 est.	686 est.	659 est.	686 est.
Xeon Gold 6262V	4	Unsupported	Unsupported	472	415
Xeon Gold 6254	4	Unsupported	Unsupported	496	451
Xeon Gold 6252	4	Unsupported	Unsupported	534	479
Xeon Gold 6248	4	Unsupported	Unsupported	481	444
Xeon Gold 6246	4	Unsupported	Unsupported	365	375
Xeon Gold 6244	4	Unsupported	Unsupported	265	297
Xeon Gold 6240L	4	Unsupported	Unsupported	446	419
Xeon Gold 6240M	4	Unsupported	Unsupported	446	419
Xeon Gold 6240	4	Unsupported	Unsupported	446	419
Xeon Gold 6238L	4	Unsupported	Unsupported	484	446
Xeon Gold 6238M	4	Unsupported	Unsupported	484	446
Xeon Gold 6238	4	Unsupported	Unsupported	484	446
Xeon Gold 6222V	4	Unsupported	Unsupported	395	368

Processor	PQ3800B2			
	Cores	Number of Processors	SPECspeed2017 int_base	SPECspeed2017 fp_base
Xeon Platinum 8280L	28	8	10.5	tbd.

The following two diagrams illustrate the throughput of the PRIMEQUEST 3800B2/3800E2 in comparison to its predecessor PRIMEQUEST 3800B/3800E, in their respective most performant configuration.



STREAM

Benchmark description

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

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

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

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

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

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

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

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

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

Benchmark environment

System Under Test (SUT)

Hardware

• Model	PRIMEQUEST 3800B2/3800E2
• Processor	2nd Generation Intel Xeon Scalable Processors Family
• Memory	48 × 64 GB (2x32 GB) 2Rx4 PC4-2933Y-R

Software

• BIOS settings	<ul style="list-style-type: none"> • HWPM Support = Disabled • Intel Virtualization Technology = Disabled • Override OS Energy Performance = Enabled • LLC Dead Line Alloc = Disabled • Stale AtoS = Enabled
• Operating system	SUSE Linux Enterprise Server 15
• Operating system settings	<pre>Transparent Huge Pages inactivated sched_cfs_bandwidth_slice_us = 50000 sched_latency_ns = 240000000 sched_migration_cost_ns = 5000000 sched_min_granularity_ns = 100000000 sched_wakeup_granularity_ns = 150000000 cpupower -c all frequency-set -g performance cpupower idle-set -d 1 cpupower idle-set -d 2 cpupower idle-set -d 3 echo 0 > /proc/sys/kernel/numa_balancing echo 1 > /proc/sys/vm/drop_caches ulimit -s unlimited nohz_full = 1-447</pre>
• Compiler	C/C++: Version 2019.3.0.591499 of Intel C/C++ Compiler for Linux
• Benchmark	STREAM Version 5.10

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

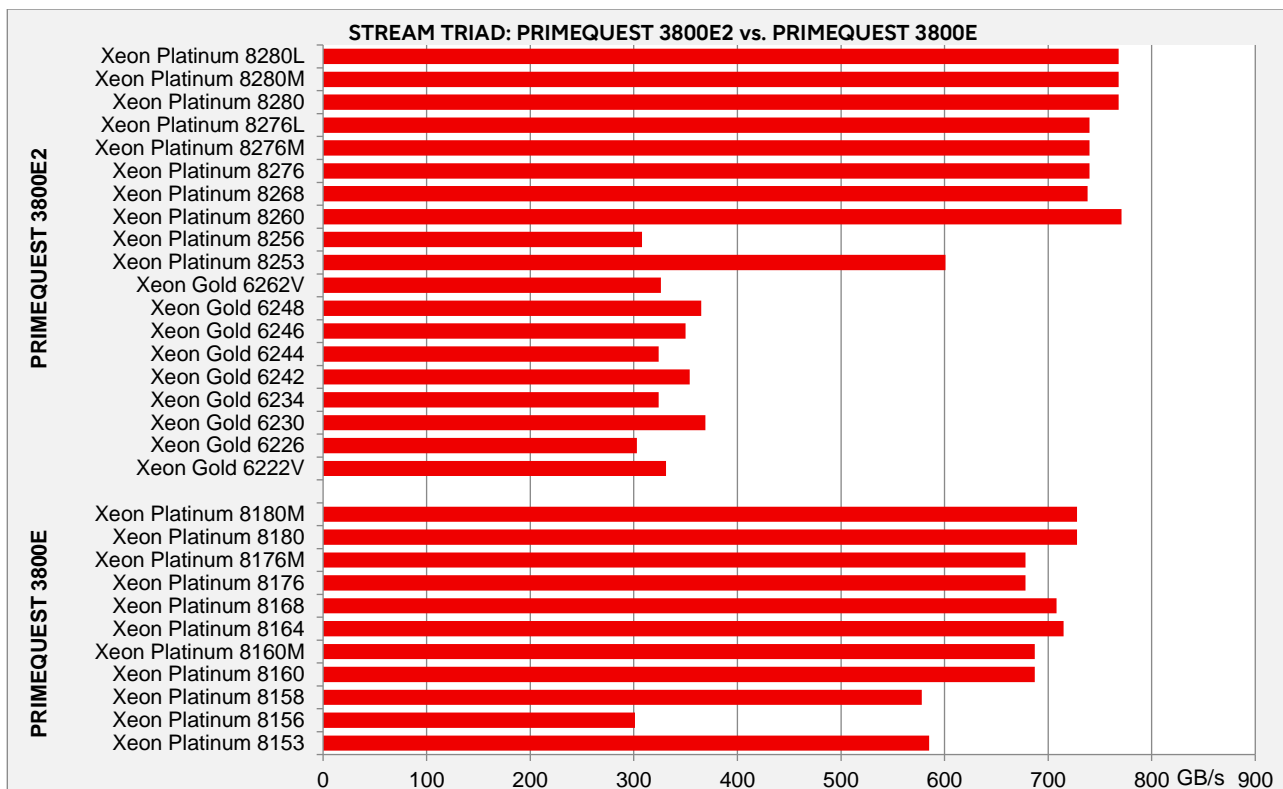
Benchmark results

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

Processor	Memory Frequency	Max. Memory Bandwidth	Cores	Processor Frequency	PQ3800E2 / PQ3800B2	
	[MHz]	[GB/s]		[GHz]	Number of Processors	TRIAD [GB/s]
Xeon Platinum 8280L	2,933	140.8	28	2.7	8	768
Xeon Platinum 8280M	2,933	140.8	28	2.7	8	768 est.
Xeon Platinum 8280	2,933	140.8	28	2.7	8	768 est.
Xeon Platinum 8276L	2,933	140.8	28	2.3	8	740
Xeon Platinum 8276M	2,933	140.8	28	2.3	8	740 est.
Xeon Platinum 8276	2,933	140.8	28	2.3	8	740 est.
Xeon Platinum 8270	2,933	140.8	26	2.7	8	734
Xeon Platinum 8268	2,933	140.8	24	2.9	8	738
Xeon Platinum 8260L	2,933	140.8	24	2.4	8	771 est.
Xeon Platinum 8260M	2,933	140.8	24	2.4	8	771
Xeon Platinum 8260	2,933	140.8	24	2.4	8	771 est.
Xeon Platinum 8256	2,933	140.8	4	3.8	8	308
Xeon Platinum 8253	2,933	140.8	16	2.2	8	601
Xeon Gold 6262V	2,933	140.8	24	1.9	4	326
Xeon Gold 6254	2,933	140.8	18	3.1	4	359
Xeon Gold 6252	2,933	140.8	24	2.1	4	386
Xeon Gold 6248	2,933	140.8	20	2.5	4	365
Xeon Gold 6246	2,933	140.8	12	3.3	4	349
Xeon Gold 6244	2,933	140.8	8	3.6	4	324
Xeon Gold 6242	2,933	140.8	16	2.8	4	354
Xeon Gold 6240L	2,933	140.8	18	2.6	4	358 est.
Xeon Gold 6240M	2,933	140.8	18	2.6	4	358 est.
Xeon Gold 6240	2,933	140.8	18	2.6	4	358
Xeon Gold 6238M	2,933	140.8	22	2.1	4	372 est.
Xeon Gold 6238L	2,933	140.8	22	2.1	4	372 est.
Xeon Gold 6238	2,933	140.8	22	2.1	4	372
Xeon Gold 6234	2,933	140.8	8	3.4	4	324
Xeon Gold 6230	2,933	140.8	20	2.1	4	369
Xeon Gold 6226	2,933	140.8	12	2.8	4	303
Xeon Gold 6222V	2,400	140.8	20	1.8	4	330

Processor	Memory Frequency	Max. Memory Bandwidth	Cores	Processor Frequency	PQ3400E2	
	[MHz]	[GB/s]		[GHz]	Number of Processors	TRIAD [GB/s]
Xeon Platinum 8280L	2,933	140.8	28	2.7	4	392 est.
Xeon Platinum 8280M	2,933	140.8	28	2.7	4	392 est.
Xeon Platinum 8280	2,933	140.8	28	2.7	4	392 est.
Xeon Platinum 8276L	2,933	140.8	28	2.3	4	392 est.
Xeon Platinum 8276M	2,933	140.8	28	2.3	4	392 est.
Xeon Platinum 8276	2,933	140.8	28	2.3	4	392 est.
Xeon Platinum 8268	2,933	140.8	24	2.9	4	386 est.
Xeon Platinum 8256	2,933	140.8	4	3.8	4	198 est.
Xeon Gold 6248	2,933	140.8	20	2.5	4	365 est.
Xeon Gold 6244	2,933	140.8	8	3.6	4	324 est.
Xeon Gold 6242	2,933	140.8	16	2.8	4	354 est.

The following diagram illustrates the throughput of the PRIMEQUEST 3800E2 in comparison to its predecessor, the PRIMEQUEST 3800E



SPECjbb2015

Benchmark description

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

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

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


The SPECjbb2015 benchmark has a two performance metrics:

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


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

- SPECjbb2015 Composite
All components run on one JVM running on one host.
- SPECjbb2015 MultijVM
All components are existed on one host, but each runs on a separate JVM.
- SPECjbb2015 Distributed
Back-ends are existed on hosts separated from hosts on which the other components are running. Back-ends and the other components are connected by networks.


Results are not comparable to those in other categories.



(a) example of SPECjbb2015 Composite configuration



(b) example of SPECjbb2015 MultiJVM configuration



(c) example of SPECjbb2015 Distributed configuration

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

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

Benchmark environment

PRIMEQUEST 3800B2 was configured for the SPECjbb2015 MultiJVM benchmark measurement.

System Under Test (SUT)	
Hardware	
Model	PRIMEQUEST 3800B2
Processor	8 × Intel Xeon Platinum 8280L
Memory	48 × 64 GB (2x32 GB) 2Rx4 DDR4-2933 R ECC
Network interface	1 Gbit/s LAN
Disk subsystem	Disk : 1 × SSD SAS 12 Gb/s 2.5" 1.6 TB
Software	
BIOS settings	Stale AtoS set to Enabled Intel Virtualization Technology set to Disabled VT-d set to Disabled XPT Prefetch set to Enabled Energy Performance set to Balanced Performance UPI Link Frequency Select set to 10.4GT/s
Operating system	SUSE Linux Enterprise Server 12 SP4 (x86_64) (4.12.14-94.41-default)
Operating system settings	ulimit -n 1048576 ulimit -i unlimited ulimit -s unlimited ulimit -u unlimited ulimit -l unlimited ulimit -v unlimited ulimit -m unlimited cpupower -c all frequency-set -g performance

	<pre>echo 6000000 > /proc/sys/kernel/sched_min_granularity_ns echo 1000 > /proc/sys/kernel/sched_migration_cost_ns echo 990000 > /proc/sys/kernel/sched_rt_runtime_us tuned-adm profile throughput-performance Add kernel command line "nopti nospec_store_bypass_disable l1tf=off"</pre>
JVM	Oracle Java SE 11.0.2
JVM settings Controller (Ctr)	<pre>-server -Xms2g -Xmx2g -Xmn1536m -XX:+UseLargePages - XX:LargePageSizeInBytes=2m -XX:+UseParallelOldGC -XX:+UseTransparentHugePages</pre>
JVM settings Backends (BE)	<pre>-showversion -server -Xms29g -Xmx29g -Xmn27g -XX:SurvivorRatio=28 -XX:+UseLargePages -XX:LargePageSizeInBytes=2m -XX:+UseParallelOldGC - XX:+AggressiveOpts -XX:+AlwaysPreTouch -XX:-UseAdaptiveSizePolicy -XX:-UsePerfData -XX:TargetSurvivorRatio=95 -XX:ParallelGCThreads=28 -verbose:gc - XX:+PrintGCDetails -XX:-UseBiasedLocking -XX:AllocatePrefetchInstr=2 - XX:+UseTransparentHugePages</pre>
JVM settings Transaction Injector (TxI)	<pre>-server -Xms2g -Xmx2g -Xmn1536m -XX:+UseLargePages - XX:LargePageSizeInBytes=2m -XX:+UseTransparentHugePages -XX:+UseParallelOldGC</pre>
SPECjbb2015 settings	<pre>specjbb.comm.connect.client.pool.size = 224; specjbb.comm.connect.selector.runner.count = 4; specjbb.comm.connect.timeouts.connect = 700000; specjbb.comm.connect.timeouts.read = 700000; specjbb.comm.connect.timeouts.write = 700000; specjbb.comm.connect.worker.pool.max = 112; specjbb.comm.connect.worker.pool.min = 64; specjbb.controller.rtcure.warmup.step = 0.9; specjbb.customerDriver.threads = {probe=64, saturate=64}; specjbb.forkjoin.workers = {Tier1=222, Tier2=10, Tier3=25}; specjbb.group.count = 16; specjbb.mapreducer.pool.size = 4; specjbb.txi.pergroup.count = 1;</pre>

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

Benchmark results

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

686,568 SPECjbb2015-MultiJVM max-jOPS

267,830 SPECjbb2015-MultiJVM critical-jOPS

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

SAP SD

Benchmark description

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

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

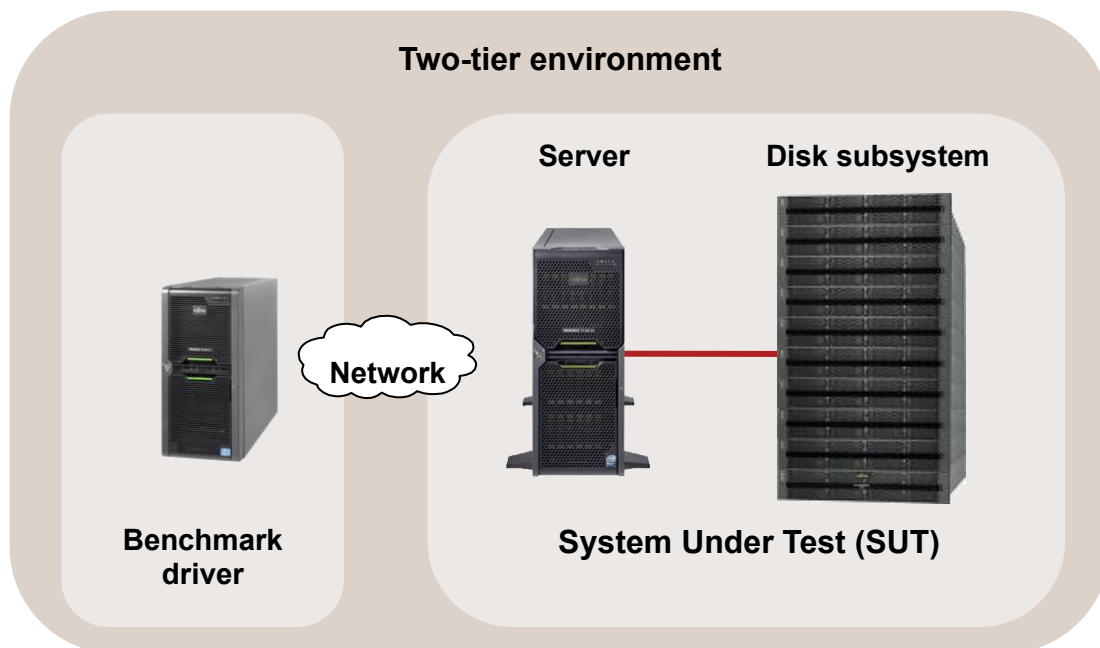
SAP AG has developed SAP Standard Application Benchmarks in order to verify the performance, stability and scaling of a SAP application system. The benchmarks, of which SD Benchmark is the most commonly used and most important, analyze the performance of the entire system and thus measure the quality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

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

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

Benchmark environment

The typical measurement set-up is illustrated below:



System Under Test (SUT)

Hardware

Model	PRIMEQUEST 3800B2
Processor	8 × Xeon Platinum 8280
Memory	48 × 64 GB (2 × 32 GB) 2Rx4 DDR4-2933 R ECC
Network interface	1 × 10 Gbit/s LAN
Disk subsystem	PRIMEQUEST 3800B2: 1 × HDD SAS 12G 600GB 2.5" 15K HOT PL 2.5" 1 × PCIe SSD 2 TB 1 × PRAID EP420i 1 × PRAID EP420e 2 × ETERNUS JX40 S2

Software

Operating system	Microsoft Windows Server 2016
Database	Microsoft SQL Server 2012
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	1 Gbit/s LAN

Software

Operating system	SUSE Linux Enterprise Server 12 SP2
------------------	-------------------------------------

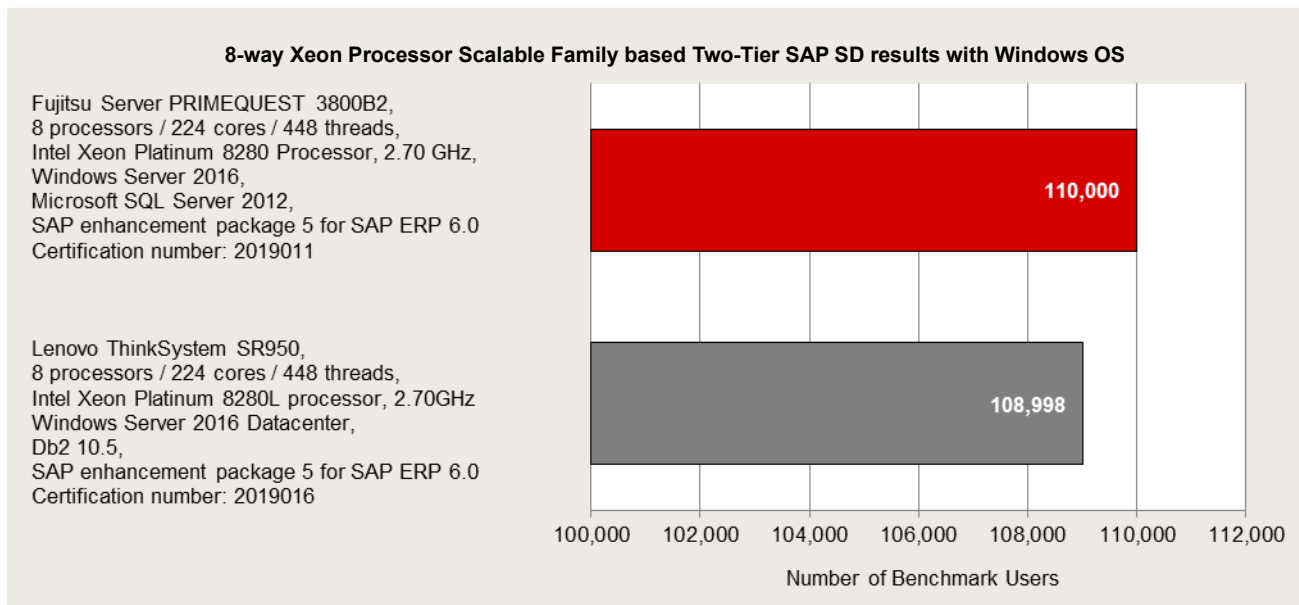
Some components may not be available in all countries or sales regions.

Benchmark results

Certification number 2019011

Number of SAP SD benchmark users	110,000
Average dialog response time	0.86 seconds
Throughput	
Fully processed order line items/hour	12,150,000
Dialog steps/hour	36,450,000
SAPS	607,500
Average database request time (dialog/update)	0.010 sec / 0.016 sec
CPU utilization of central server	98%
Operating system, central server	Windows Server 2016
RDBMS	Microsoft SQL Server 2012
SAP Business Suite software	SAP enhancement package 5 for SAP ERP 6.0
Configuration Central Server	Fujitsu Server PRIMEQUEST 3800B2, 8 processors / 224 cores / 448 threads, Intel Xeon Platinum 8280 Processor, 2.70 GHz, 64 KB L1 cache per core and 1,024KB L2 cache per core, 38.5 MB L3 cache per processor, 3,072 GB main memory

The following chart shows a comparison of two-tier SAP SD Standard Application Benchmark results for 8-way 2nd Generation Xeon Processor Scalable Family based servers with Windows OS (as of April 2, 2019). The PRIMEQUEST 3800B2 outperforms the comparably configured servers from Lenovo. The latest SAP SD 2-tier results can be found at <https://www.sap.com/dmc/exp/2018-benchmark-directory/>



The PRIMEQUEST 3800B obtained the best two-tier SAP SD Standard Application Benchmark 8-way result on Windows (as of April 2, 2019). The latest SAP SD 2-tier results can be found at <https://www.sap.com/dmc/exp/2018-benchmark-directory/>

Disk I/O: Performance of storage media

Benchmark description

Performance measurements of disk subsystems for PRIMEQUEST servers are carried out with a defined measurement method, which models the accesses of real application scenarios on the basis of specifications.

The essential specifications are as follows.

- Random access / sequential access ratio
- Read / write access ratio
- Block size (kiB)
- Queue Depth (number of IO requests to issue at one time)

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

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

In order to model applications that access in parallel with a different load intensity the Queue Depth is increased from 1 to 512 (in steps to the power of two).

The measurements of this document are based on these standard load profiles.

The main measurement items are as follows.

- Throughput [MiB/s] Throughput in megabytes per second
- Transactions [IO/s] Transaction rate in I/O operations per second
- Latency [ms] Average response time in ms

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

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

In this section, a power of 10 (1 TB = 10¹² bytes) is used to indicate the capacity of the hard storage medium, and a power of 2 (1 MiB / s = 2²⁰ bytes) is used to indicate the capacity of other media, file size, block size, and throughput.

All the details of the measurement method and the basics of disk I/O performance are described in the white paper "Basics of Disk I/O Performance."

Benchmark environment

All the measurement results discussed in this section apply for the hardware and software components listed below:

System Under Test (SUT)

Hardware

Controller: 1x PRAID CP400i		
Storage media	Category	Drive Name
HDD	SAS HDD (SAS 12Gbps, 10k rpm) [512e]	AL15SEB18EQ *2 *3
	SAS HDD (SAS 12Gbps, 10k rpm) [512n]	AL15SEB030N *2 *3
	SAS HDD (SAS 12Gbps, 15k rpm) [512n]	ST300MP0006 *1 *3
SSD	SAS SSD (SAS 12Gbps, 10DWPD)	KPM51MUG400G *2 *3 KPM51MUG800G *2 *3 KPM51MUG1T60 *2 *3
	SAS SSD (SAS 12Gbps, 3DWPD)	WUSTR6440ASS204 *2 *3 WUSTR6480ASS204 *2 *3 WUSTR6416ASS204 *2 *3 WUSTR6432ASS204 *2 *3 WUSTR6464ASS204 *2 *3

Controller: Integrated PCI Express controller
CPU: 2x Intel Xeon Platium 8280L (2.7GHz)

Storage media	Category	Drive Name
SSD	PCIe SSD (3DWPD)	SSDPEDKE020T7C1 *2 *4
		SSDPEDKE040T7C1 *2 *4

Controller: Intel C620 Standard SATA AHCI controller

Storage media	Category	Drive Name
SSD	M.2 Flash device	MTFDDAV240TCB *2 *4
		MTFDDAV480TCB *2 *4

*1) The operating system uses Microsoft Windows Server 2012 Standard R2.

*2) The operating system uses Microsoft Windows Server 2016 Standard.

*3) Measurement area is type 1.

*4) Measurement area is type 2.

Software

Operating system	Microsoft Windows Server 2012 Standard R2 Microsoft Windows Server 2016 Standard	
Benchmark version	3.0	
RAID type	Logical drive of type RAID 0 consisting of 1 hard disk	
Stripe size	Controller default (here 64 kiB)	
Measuring tool	Iometer 1.1.0	
Measurement area	Type1	RAW file system is used. The first 10% of the usable LBA area is used for sequential accesses; the next 25% for random accesses.
	Type2	NTFS file system is used. The 32GiB area is secured for the first of the target drive, and is used for sequential access and random access.
Total number of Iometer workers	1	
Alignment of Iometer accesses	Aligned to whole multiples of 4096 bytes	

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

Benchmark results

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

Controller

The measurements were made using controllers in the table below.

Storage medium	Storage medium	Cache	Supported interfaces		RAID levels
			host	drive	
SSD/HDD	PRAID CP400i	-	PCIe 3.0 x8	SATA 6G SAS 12G	0, 1, 1E, 10, 5, 50
PCIe SSD	Integrated PCI Express controller	-	PCIe 3.0 x4		-
M.2 Flash	C620 Standard SATA AHCI controller	-	DMI 3.0 x4	SATA 6G	-

Storage media

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

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

HDDs and SSDs are operated via host bus adapters, usually RAID controllers, with a SATA or SAS interface. The interface of the RAID controller to the chipset of the systemboard is typically PCIe or, in the case of the integrated onboard controllers, an internal bus interface of the systemboard.

Of all the storage medium types SSDs offer by far the highest transaction rates for random load profiles as well as the shortest access times. In return, however, the price per gigabyte of storage capacity is substantially higher.

Cache settings

In most cases, the cache of HDDs has a great influence on disk-I/O performance. It is frequently regarded as a security problem in case of power failure and is thus switched off. On the other hand, it was integrated by hard disk manufacturers for the good reason of increasing the write performance. For performance reasons it is therefore advisable to enable the hard disk cache. To prevent data loss in case of power failure you are recommended to equip the system with a UPS.

For the purpose of easy and reliable handling of the settings for RAID controllers and hard disks it is advisable to use the RAID-Manager software "ServerView RAID" that is supplied for PRIMEQUEST servers. All the cache settings for controllers and hard disks can usually be made en bloc – specifically for the application – by using the pre-defined modi "Performance" or "Data Protection". The "Performance" mode ensures the best possible performance settings for the majority of the application scenarios.

Performance values

The performance values are summarized in the following tables, in each case specifically for a single storage medium and with various access types and block sizes. The established measurement variables, as already mentioned in the subsection [Benchmark description](#), are used here. Thus, transaction rate is specified for random accesses and data throughput for sequential accesses. To avoid any confusion among the measurement units the tables have been separated for the two access types.

The table cells contain the maximum achievable values. This means that each value is the maximum achievable value of the whole range of load intensities (# of Outstanding I/Os). In order to also visualize the numerical values each table cell is highlighted with a horizontal bar, the length of which is proportional to the numerical value in the table cell. All bars shown in the same scale of length have the same color. In other words, a visual comparison only makes sense for table cells with the same colored bars. Since the horizontal bars in the table cells depict the maximum achievable performance values, they are shown by the color getting lighter as you move from left to right. The light shade of color at the right end of the bar tells you that the value is a maximum value and can only be achieved under optimal prerequisites. The darker the shade becomes as you move to the left, the more frequently it will be possible to achieve the corresponding value in practice.

Storage media

HDDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
1,800	AL15SEB18EQ	SAS 12G	600	512	547	258	255
300	AL15SEB030N	SAS 12G	645	546	568	231	230
300	ST300MP0006	SAS 12G	768	662	472	304	304

SSDs

Capacity [GB]	Storage device	Inter face	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
400	KPM51MUG400G	SAS 12G	84,469	13,329	13,677	1,056	1,041
800	KPM51MUG800G	SAS 12G	99,728	14,549	18,049	1,056	1,042
1,600	KPM51MUG1T60	SAS 12G	108,428	17,243	19,634	1,057	1,042
400	WUSTR6440ASS204	SAS 12G	83,427	14,459	13,924	1,073	626
800	WUSTR6480ASS204	SAS 12G	94,899	22,414	21,187	1,073	1,008
1,600	WUSTR6416ASS204	SAS 12G	97,107	24,053	22,802	1,073	1,029
3,200	WUSTR6432ASS204	SAS 12G	106,745	23,975	22,793	1,073	1,030
6,400	WUSTR6464ASS204	SAS 12G	111,695	23,911	22,639	1,073	1,030
2,000	SSDPEDKE020T7C1	PCIe3 x4	92,172	35,791	32,180	3,080	1,608
4,000	SSDPEDKE040T7C1	PCIe3 x4	111,314	40,984	38,309	3,056	2,157
240	MTFDDAV240TCB	SATA 6G	19,773	3,844	4,968	487	258
480	MTFDDAV480TCB	SATA 6G	22,258	4,935	6,294	509	403

OLTP-2

Benchmark description

OLTP stands for Online Transaction Processing. The OLTP-2 benchmark is based on the typical application scenario of a database solution. In OLTP-2 database access is simulated and the number of transactions achieved per second (tps) determined as the unit of measurement for the system.

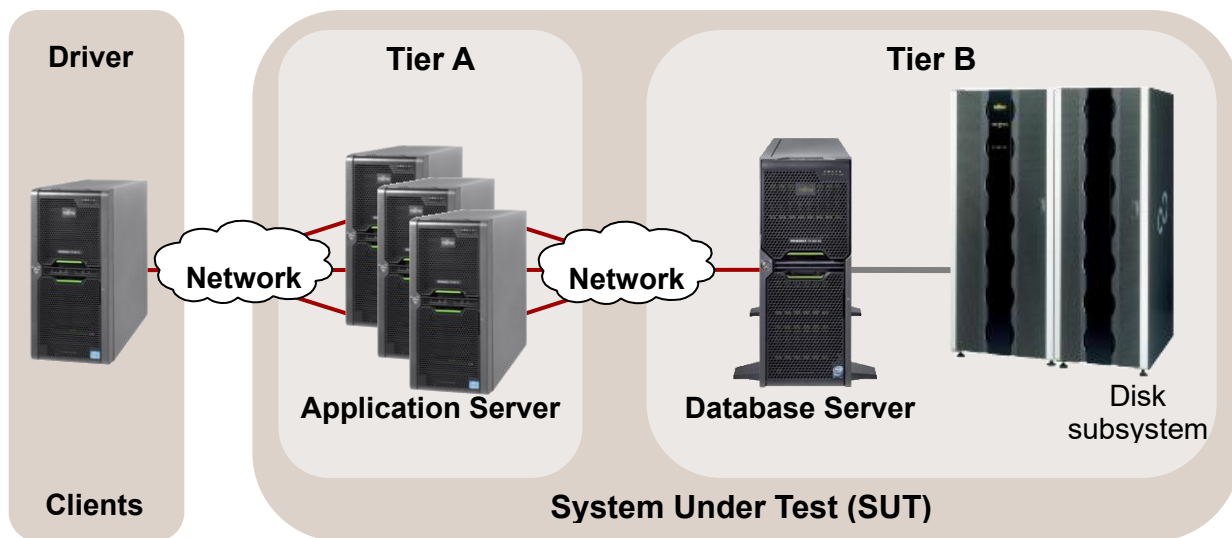
In contrast to benchmarks such as SPECint and TPC-E, which were standardized by independent bodies and for which adherence to the respective rules and regulations are monitored, OLTP-2 is an internal benchmark of Fujitsu. OLTP-2 is based on the well-known database benchmark TPC-E. OLTP-2 was designed in such a way that a wide range of configurations can be measured to present the scaling of a system with regard to the CPU and memory configuration.

Even if the two benchmarks OLTP-2 and TPC-E simulate similar application scenarios using the same load profiles, the results cannot be compared or even treated as equal, as the two benchmarks use different methods to simulate user load. OLTP-2 values are typically similar to TPC-E values. A direct comparison, or even referring to the OLTP-2 result as TPC-E, is not permitted, especially because there is no price-performance calculation.

Further information can be found in the document [Benchmark Overview OLTP-2](#).

Benchmark environment

The typical measurement set-up is illustrated below:



All OLTP-2 results were Calculated based on the configuration of the next following pages of PRIMEQUEST 3800E2.

Database Server (Tier B)

Hardware

Model	PRIMEQUEST 3800E2
Processor	2nd Generation Intel Xeon Scalable Processors Family
Memory	2 processors:12 ×128 GB (2x64 GB) 2Rx4 DDR4-2933 ECC 4 processors:24 ×128 GB (2x64 GB) 2Rx4 DDR4-2933 ECC 8 processors:48 ×128 GB (2x64 GB) 2Rx4 DDR4-2933 ECC
Network interface	1 × Dual port onboard LAN 10 Gb/s
Disk subsystem	PRIMEQUEST 3800E2 : Onboard RAID controller PRAID EP420i 2 × 300 GB 10k rpm SAS Drive, RAID 1 (OS), 6 × 1.6 TB SSD, RAID 10 (LOG) 3 × 1.6 TB SSD, RAID 10 (temp) 10 × PRAID EP540e 10 × JX40 S2:9 × 1.6 TB SSD Drive each, RAID5 (data)

Software

BIOS	PB19033
Operating system	Microsoft Windows Server 2016 Standard + KB4462928
Database	Microsoft SQL Server 2017 Enterprise + KB4341265

Application Server (Tier A)

Hardware

Model	1 × PRIMERGY RX2530 M4
Processor	2 × Xeon Platinum 8180
Memory	192 GB, 2666 MHz Registered ECC DDR4
Network interface	1 × Dual port onboard LAN 10 Gb/s 1 × Dual Port LAN 1 Gb/s
Disk subsystem	2 × 300 GB 10k rpm SAS Drive

Software

Operating system	Microsoft Windows Server 2016 Standard
------------------	--

Client

Hardware

Model	1 × PRIMERGY RX2530 M2
Processor	2 × Xeon E5-2667 v4
Memory	128 GB, 2400 MHz registered ECC DDR4
Network interface	1 × onboard Quad Port LAN 1 Gb/s
Disk subsystem	1 × 300 GB 10k rpm SAS Drive

Software

Operating system	Microsoft Windows Server 2012 R2 Standard
Benchmark	OLTP-2 Software EGen version 1.14.0

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

Benchmark results

Database performance greatly depends on the configuration options with CPU, memory and on the connectivity of an adequate disk subsystem for the database. In the following scaling considerations for the processors we assume that both the memory and the disk subsystem has been adequately chosen and is not a bottleneck.

A guideline in the database environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses. This is why a configuration with a total memory of 6144 GB was considered for the measurements with eight processors and a configuration with a total memory of 3072 GB for the measurements with four processors and a configuration with a total memory of 1536 GB for the measurements with two processors. Both memory configurations have memory access of 2933 MHz..


PRIMEQUEST 3800E2 and 3800B2 are equivalent in performance.

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

Processor	Cores	Threads	8CPU Score	4CPU Score	2CPU Score
Xeon Platinum 8280L	28	56	19,455 est.	11,512 est.	6,732 est.
Xeon Platinum 8280M	28	56	19,455 est.	11,512 est.	6,732 est.
Xeon Platinum 8280	28	56	19,455 est.	11,512 est.	6,732 est.
Xeon Platinum 8276L	28	56	16,900 est.	10,000 est.	5,848 est.
Xeon Platinum 8276M	28	56	16,900 est.	10,000 est.	5,848 est.
Xeon Platinum 8276	28	56	16,900 est.	10,000 est.	5,848 est.
Xeon Platinum 8270	26	52	18,014 est.	10,659 est.	6,233 est.
Xeon Platinum 8268	24	48	17,293 est.	10,233 est.	5,984 est.
Xeon Platinum 8260L	24	48	15,721 est.	9,302 est.	5,440 est.
Xeon Platinum 8260M	24	48	15,721 est.	9,302 est.	5,440 est.
Xeon Platinum 8260	24	48	15,721 est.	9,302 est.	5,440 est.
Xeon Platinum 8256	4	8	3,799 est.	2,248 est.	1,315 est.
Xeon Platinum 8253	16	32	9,629 est.	5,698 est.	3,332 est.
Xeon Gold 6262V	24	48		8,217 est.	4,805 est.
Xeon Gold 6254	18	36		8,527 est.	4,987 est.
Xeon Gold 6252	24	48		8,876 est.	5,191 est.
Xeon Gold 6248	20	40		8,140 est.	4,760 est.
Xeon Gold 6246	12	24		tbd	tbd
Xeon Gold 6244	8	16		4,612 est.	2,697 est.
Xeon Gold 6242	16	32		7,093 est.	4,148 est.
Xeon Gold 6240L	18	36		7,442 est.	4,352 est.
Xeon Gold 6240M	18	36		7,442 est.	4,352 est.
Xeon Gold 6240	18	36		7,442 est.	4,352 est.
Xeon Gold 6238L	22	44		7,752 est.	4,533 est.
Xeon Gold 6238M	22	44		7,752 est.	4,533 est.
Xeon Gold 6238	22	44		7,752 est.	4,533 est.
Xeon Gold 6234	8	16		4,302 est.	2,516 est.

Xeon Gold 6230	20	40		7,326 est.	4,284 est.
Xeon Gold 6226	12	24		5,465 est.	3,196 est.
Xeon Gold 6222V	20	40		6,628 est.	3,876 est.

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



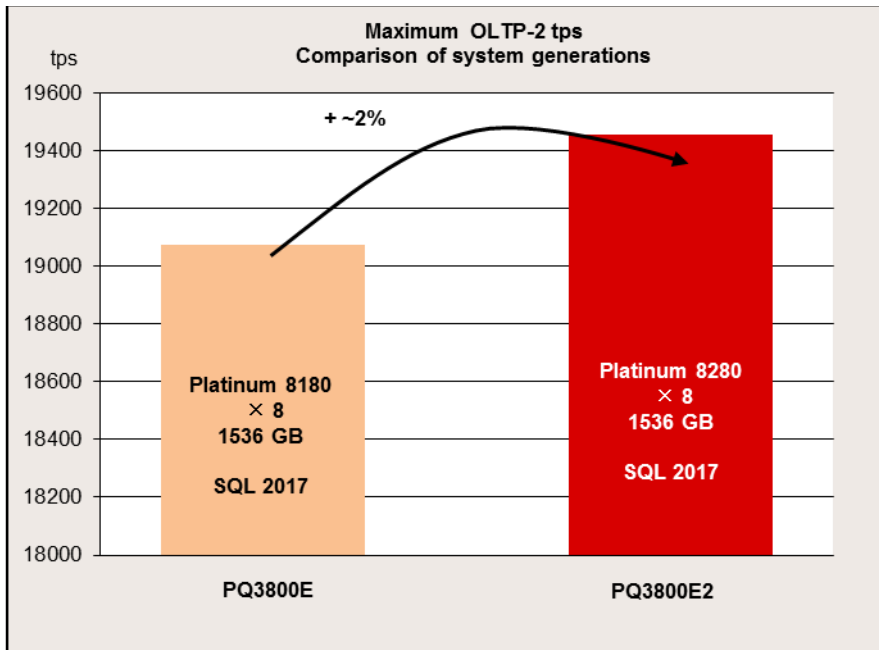
It is evident that a wide performance range is covered by the variety of released processors. If you compare the OLTP-2 value of the processor with the lowest performance (Xeon Platinum 8256) with the value of the processor with the highest performance (Xeon Platinum 8280), the result is an 5-fold increase in performance.

The features of the processors are summarized in the section "Technical data".

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

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

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



vServCon

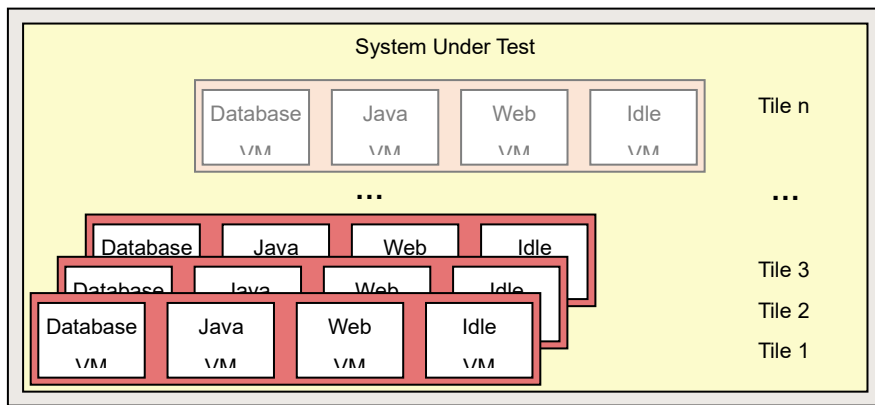
Benchmark description

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

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

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

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



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


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

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

A detailed description of vServCon is in the document: [Benchmark Overview vServCon](#).

Benchmark environment

The typical measurement set-up is illustrated below:



All vServCon results were Calculated based on the configuration of the next following pages of PRIMEQUEST 3800E2.

System Under Test (SUT)

Hardware

Processor	2nd Generation Intel Xeon Scalable Processors Family
Memory	48 × 64 GB (2x32 GB) 2Rx4 PC4-2933Y-R
Network interface	2 × Intel Ethernet Controller X710 for 10GbE SFP+
Disk subsystem	1 × dual-channel FC controller Emulex LPe160021 LINUX/LIO based flash storage system

Software

Operating system	VMware ESXi 6.7 EP06 Build 11675023
------------------	-------------------------------------

Load generator (incl. Framework controller)

Hardware (Shared)

Enclosure	4 × PRIMERGY RX2530 M2
-----------	------------------------

Hardware

Processor	2 × XeonE5-2683 v4
Memory	128 GB
Network interface	3 × 1 Gbit LAN

Software

Operating system	VMware ESXi 6.0.0 U2 Build 3620759
------------------	------------------------------------

Load generator VM (on various servers)

Hardware

Processor	1 × logical CPU
Memory	4048 MB
Network interface	2 × 1 Gbit/s LAN

Software

Operating system	Microsoft Windows Server 2008 Standard Edition 32 bit
------------------	---

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

Benchmark results

The PRIMEQUEST rack systems dealt with here are based on processors of the 2nd Generation 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.

PRIMEQUEST 3800E2 and 3800B2 are equivalent in performance.


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

Processor	Co res	Thre ads	Number of process ors	#Tiles	Score	Number of process ors	#Tiles	Score
Xeon Platinum 8280L	28	56	8	140 est.	237.9 est.	4	70 est.	121.4 est.
Xeon Platinum 8280M	28	56	8	140 est.	237.9 est.	4	70 est.	121.4 est.
Xeon Platinum 8280	28	56	8	140 est.	237.9 est.	4	70 est.	121.4 est.
Xeon Platinum 8276L	28	56	8	138 est.	206.7 est.	4	69 est.	105.7 est.
Xeon Platinum 8276M	28	56	8	138 est.	206.7 est.	4	69 est.	105.7 est.
Xeon Platinum 8276	28	56	8	138 est.	206.7 est.	4	69 est.	105.7 est.
Xeon Platinum 8270	26	52	8	126 est.	220.3 est.	4	63 est.	112.4 est.
Xeon Platinum 8268	24	48	8	126 est.	211.5 est.	4	63 est.	107.9 est.
Xeon Platinum 8260L	24	48	8	122 est.	192.3 est.	4	61 est.	98.1 est.
Xeon Platinum 8260M	24	48	8	122 est.	192.3 est.	4	61 est.	98.1 est.
Xeon Platinum 8260	24	48	8	122 est.	192.3 est.	4	61 est.	98.1 est.
Xeon Platinum 8256	4	8	8	36 est.	46.5 est.	4	18 est.	23.7 est.
Xeon Platinum 8253	16	32	8	100 est.	117.8 est.	4	50 est.	60.1 est.
Xeon Gold 6262V	24	48		Unsupported	Unsupported	4	56 est.	86.6 est.
Xeon Gold 6254	18	36		Unsupported	Unsupported	4	55 est.	89.9 est.
Xeon Gold 6252	24	48		Unsupported	Unsupported	4	57 est.	93.6 est.
Xeon Gold 6248	20	40		Unsupported	Unsupported	4	52 est.	85.8 est.
Xeon Gold 6246	12	24		Unsupported	Unsupported	4	tbd	tbd
Xeon Gold 6244	8	16		Unsupported	Unsupported	4	29 est.	48.6 est.
Xeon Gold 6242	16	32		Unsupported	Unsupported	4	52 est.	74.8 est.
Xeon Gold 6240L	18	36		Unsupported	Unsupported	4	53 est.	78.4 est.
Xeon Gold 6240M	18	36		Unsupported	Unsupported	4	53 est.	78.4 est.
Xeon Gold 6240	18	36		Unsupported	Unsupported	4	53 est.	78.4 est.
Xeon Gold 6238M	22	44		Unsupported	Unsupported	4	57 est.	81.7 est.
Xeon Gold 6238L	22	44		Unsupported	Unsupported	4	57 est.	81.7 est.
Xeon Gold 6238	22	44		Unsupported	Unsupported	4	57 est.	81.7 est.
Xeon Gold 6234	8	16		Unsupported	Unsupported	4	29 est.	45.4 est.
Xeon Gold 6230	20	40		Unsupported	Unsupported	4	52 est.	77.2 est.

Xeon Gold 6226	12	24	Unsupported	Unsupported	4	42 est.	57.6 est.
Xeon Gold 6222V	20	40	Unsupported	Unsupported	4	50 est.	69.9 est.

These PRIMEQUEST rack systems are very suitable for application virtualization owing to the progress made in processor technology. Compared with a system based on the previous processor generation, approximately 3.6% higher virtualization performance can be achieved (measured in vServCon score in their maximum configuration).

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



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. Also, 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.

In VMmark V3 there 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 is the maximum sum of the benefits of server aggregation and is used as a comparison criterion for different hardware platforms.

This score is determined from the individual results of the VMs and an infrastructure components 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 components workload.

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”.

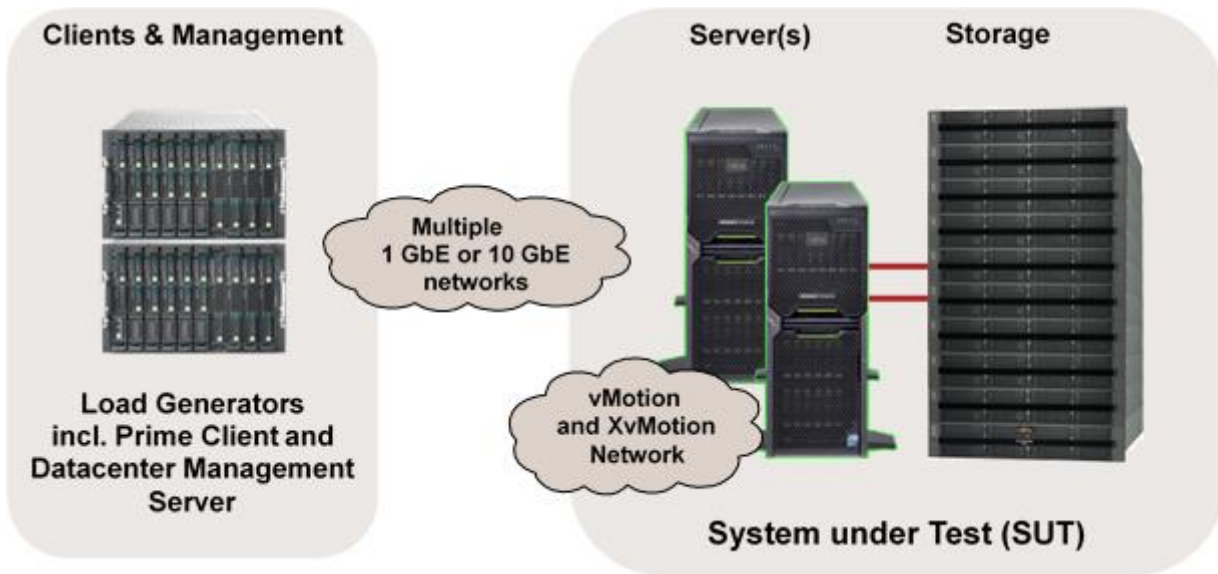
In the case of the two test types “Performance with Server Power” and “Performance with Server and Storage Power”, a so-called “Server PPKW Score” and “Server and Storage PPKW Score” are determined. These are the performance scores divided by the average power consumption in kilowatts (PPKW = performance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

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

A detailed description of VMmark V3 is available in the document [Benchmark Overview VMmark V3](#).

Benchmark environment

The typical measurement set-up is illustrated below:



System Under Test (SUT)

Hardware

• Number of servers	2
• Model	PRIMEQUEST 3800B2/3800E2
• Processor	8 × Intel Xeon Platinum 8280
• Memory	3,072 GB: 48 × 64 GB (2x32 GB) 2Rx4 DDR4-2933 R ECC
• Network interface	4 × Intel Ethernet Controller X710 for 10GbE SFP+ 1 × Intel Ethernet Controller 10G X550
• Disk subsystem	4 × Dual port PFC EP LPe31002 6 × PRIMERGY RX2540 M4, 2 × PRIMERGY RX2540 M5, 2 × PRIMERGY RX2540M2 configured as Fibre Channel target 4 × PRIMERGY RX2540 M4 2 × Micron MTFDDAK480 TDC SATA-SSD (480 GB) 2 × Intel P4800X 750GB PCIe SSD (750 GB) 1 × Intel P4600 2TB PCIe SSD (2 TB) 1 × Intel P4600 4TB PCIe SSD (4 TB) 1 × PRIMERGY RX2540 M4 2 × Micron MTFDDAK480 TDC SATA-SSD (480 GB) 2 × Intel P4800X 750GB PCIe SSD (750 GB) 2 × Intel P4600 2TB PCIe SSD (2 TB) SAS-SSD1, 2(RAID 1) 1 × PRIMERGY RX2540 M4 1 × Micron MTFDDAK480 TDC SATA-SSD (960 GB) 3 × Intel P4800X 750GB PCIe SSD (750 GB) 1 × Intel P4600 2TB PCIe SSD (2 TB) 1 × PRIMERGY RX2540 M5 1 × Micron MTFDDAK240 TCB SATA-SSD (240 GB) 3 × Intel P4800X 750GB PCIe SSD (750 GB) 1 × Intel P4600 2TB PCIe SSD (2 TB) 1 × PRIMERGY RX2540 M5 1 × Micron MTFDDAK480 TDC SATA-SSD (240 GB) 2 × Intel P4800X 750GB PCIe SSD (750 GB)

	1 × PRIMERGY RX2540 M2 1 × 400GB SAS-SSD Toshiba PX02SMF040 3 × Fusion ioMemory PX600 2.6TB PCIe SSD 1 × Fusion ioMemory PX600 1.3TB PCIe SSD 1 × PRIMERGY RX2540 M2 1 × 400GB SAS-SSD Toshiba PX02SMF040 3 × Fusion ioMemory PX600 2.6TB PCIe SSD
Software	
• BIOS	V1.0.0.0 R1.10.0 for D3858-B1x
• BIOS settings	See details
• Operating system	VMware ESXi 6.7 EP 06, Build 11675023
• Operating system settings	ESX settings: see details

Detail

See disclosure	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmark/2019-06-14-Fujitsu-PRIMEQUEST3800E2.pdf
----------------	---

Datacenter Management Server (DMS)**Hardware**

• Model	1 × PRIMERGY RX2530 M2
• Processor	1 × Intel Xeon E5-2698 v4
• Memory	64 GB
• Network interface	1 × Emulex One Connect Oce14000 1 GbE Dual Port Adapter

Software

• Operating system	VMware ESXi 6.7 EP 02a Build 9214924
--------------------	--------------------------------------

Datacenter Management Server (DMS) VM**Hardware**

• Processor	4 × logical CPU
• Memory	16 GB
• Network interface	1 × 1 Gbit/s LAN

Software

• Operating system	VMware vCenter Server Appliance 6.7.0d Build 9451876
--------------------	--

Load generator**Hardware**

• Model	6 × PRIMERGY RX2530 M2, 1 × PRIMERGY RX4770 M4 4 × PRIMERGY RX2530M2 Intel Xeon E5-2699 v4 1 × Emulex One Connect OCe14000 1GbE Dual Port Adapter 1 × Emulex One Connect OCe14000 10GbE Dual Port Adapter 2 × PRIMERGY RX2530M2 Intel Xeon E5-2699A v4 1 × Emulex One Connect OCe14000 1GbE Dual Port Adapter 1 × Emulex One Connect OCe14000 10GbE Dual Port Adapter 1 × PRIMERGY RX4770M4 Intel Xeon Platinum 8180M 1 × Intel Ethernet Connection X722 for 1GbE
---------	--

	2 × Emulex OneConnect OCe14000 10GbE Dual Port Adapter
• Processor	7 × Xeon E5-2699 v4
• Memory	6 × RX2530M2 256GB 1 × RX4770M4 768GB
• Network interface	1 × Emulex One Connect Oce14000 1GbE Dual Port Adapter 1 × Emulex One Connect Oce14000 10GbE Dual Port Adapter
Software	
• Operating system	VMware ESXi 6.7 U1 Build 10302608

Benchmark results

"Performance Only" measurement result (June 14 2019)



On June 14, 2019 Fujitsu achieved with a PRIMEQUEST 3800E2 with Xeon Platinum 8280 processors and VMware VMware ESXi 6.7 EP 06 a VMmark V3 score of "33.04@35 tiles" in a system configuration with a total of 8 × 56 processor cores and when using two identical servers in the "System under Test" (SUT). With this result the PRIMEQUEST 3800E2 is in the official VMmark V3 "Performance Only" ranking the most powerful two-socket server in a "matched pair" configuration consisting of two identical hosts (valid as of benchmark results publication date).

All comparisons for the competitor products reflect the status of June 14, 2019. The current VMmark V3 "Performance Only" results as well as the detailed results and configuration data are available at <https://www.vmware.com/products/vmmark/results3x.html>.

The processors used, which with a good hypervisor setting could make optimal use of their processor features, were the essential prerequisites for achieving the PRIMEQUEST 3800E2 result. These features include Hyper-Threading. All this has a particularly positive effect during virtualization.

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

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

All the components used were optimally attuned to each other.

VMmark is a product of VMware, Inc.

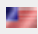
Literature


PRIMEQUEST Servers

<https://www.fujitsu.com/global/products/computing/servers/mission-critical/>

PRIMEQUEST 3800B2/3800E2

This Whitepaper

 <https://docs.ts.fujitsu.com/dl.aspx?id=4d21ef80-aeaa-4c0f-9e69-22b565852c76>

 <https://docs.ts.fujitsu.com/dl.aspx?id=19528606-66f4-4d69-9ecc-210e074040ee>

Data sheet

<https://docs.ts.fujitsu.com/dl.aspx?id=23b40e45-2e45-4bce-9b8e-5f4ed96e6d36> (PQ3800B2)

<https://docs.ts.fujitsu.com/dl.aspx?id=5430abaf-b293-43cc-babf-dcc741b0c2ce> (PQ3800E2)

PRIMEQUEST Performance

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

SPEC CPU2017

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

Benchmark Overview SPECcpu2017

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

STREAM

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

SPECjbb2015

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

SAP SD

<https://www.sap.com/benchmark>

Benchmark overview SAP SD:

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

OLTP-2

Benchmark Overview OLTP-2

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

vServCon

Benchmark Overview vServCon

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

VMmark V3

<https://www.vmware.com/products/vmmark.html>

Benchmark Overview VMmark V3

<https://docs.ts.fujitsu.com/dl.aspx?id=e6f9973c-90d6-47c6-b317-e388a978bfb7>

Document change history

Version	Date	Description
1.3	2023-10-10	Updated: <ul style="list-style-type: none"> • New Visual Identity format
1.2	2021-07-28	Updated: <ul style="list-style-type: none"> • Contact information and URLs Updated to the latest one • Minor correction
1.1	2019-10-04	New: <ul style="list-style-type: none"> • Disk I/O: Performance of storage media Results for 2.5" storage media • VMmark V3 "Performance Only" measurement with Intel Xeon Platinum 8280 • LINPACK Measured with 2nd Generation Intel Xeon Processor Scalable Family Updated: <ul style="list-style-type: none"> • SPECcpu2017 Measured additionally with 2nd Generation Intel Xeon Processor Scalable Family
1.0	2019-04-30	New: <ul style="list-style-type: none"> • Technical data • SPECcpu2017 Measurements with 2nd Generation Intel Xeon Processor Scalable Family • SPECjbb2015 Measurement with Intel Xeon Platinum 8280L • SAP SD Certification number 2019011 • OLTP-2 Calculated with 2nd Generation Intel Xeon Processor Scalable Family • vServCon Calculated with 2nd Generation Intel Xeon Processor Scalable Family

Contact

Fujitsu

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

PRIMERGY Performance and Benchmarks

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

© Fujitsu 2023. All rights reserved. Fujitsu and Fujitsu logo are trademarks of Fujitsu Limited registered in many jurisdictions worldwide. Other product, service and company names mentioned herein may be trademarks of Fujitsu or other companies. This document is current as of the initial date of publication and subject to be changed by Fujitsu without notice. This material is provided for information purposes only and Fujitsu assumes no liability related to its use.