Fujitsu Server PRIMERGY
Performance Report
PRIMERGY CX2550 M7 / CX2560 M7

FUJITSU

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMERGY CX2550 M7 / CX2560 M7.

Explaines PRIMERGY CX2550 M7 / CX2560 M7 performance data in comparison to other PRIMERGY models. In addition to the benchmark results, the explanation for each benchmark and benchmark environment are also included.

Version
1.2
2023-11-07

Contents

Technical data

PRIMERGY CX400 M7 Chassis

PRIMERGY CX2550 M7

PRIMERGY CX2560 M7

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

Model	PRIMERGY CX2550 M7	PRIMERGY CX2560 M7
Cooling method	Air cooling / Liquid cooling	Air cooling
Form factor	Server node	
Chipset	Intel C741	
Number of sockets	2	
Number of configurable processors	1[*1] or 2	
Processor type	4th Generation Intel Xeon Scalable Processors Family	
Number of memory slots	16 (8 per processor)	24 (12 per processor)
Maximum memory configuration	4,096 GB	6,144 GB
Maximum number of internal storage disks	2.5 inch: 2	2.5 inch: 6
Maximum number of PCI slots	PCI-Express 5.0 (x16 lane): 2 (Low Profile)	

*1 Air cooling only

Processor									
Processor model	Type	Number of cores	Number of threads	L3 Cache	UPI speed	Rated frequency	Maximum turbo frequency	Maximum memory frequency	TDP
				[MB]	[GT/s]	[GHz]	[GHz]	[MHz]	[W]
1CPU and 2CPU supported processor									
Xeon Max 9480 *1	HBM	56	112	112.5	16	1.90	3.50	4,800	350
Xeon Max 9470 *1	HBM	52	104	105	16	2.00	3.50	4,800	350
Xeon Max 9468 *1	HBM	48	96	105	16	2.10	3.50	4,800	350
Xeon Max 9468 *1	HBM	48	96	105	16	2.10	3.50	4,800	350
Xeon Max 9462 *1	HBM	32	64	75	16	2.70	3.50	4,800	350
Xeon Platinum 8490H *1	XCC	60	120	112.5	16	1.90	3.50	4,800	350
Xeon Platinum 8480+ *1	XCC	56	112	105	16	2.00	3.80	4,800	350
Xeon Platinum 8470Q *1	XCC	52	104	105	16	2.10	3.80	4,800	350
Xeon Platinum 8470N *1	XCC	52	104	97.5	16	1.70	3.60	4,800	300
Xeon Platinum 8470 *1	XCC	52	104	105	16	2.00	3.80	4,800	350
Xeon Platinum 8468V *1	XCC	48	96	97.5	16	2.40	3.80	4,800	330
Xeon Platinum 8468 *1	XCC	48	96	105	16	2.10	3.80	4,800	350
Xeon Platinum 8462Y+ *1	MCC	32	64	60	16	2.80	4.10	4,800	300
Xeon Platinum 8460Y+ *1	XCC	40	80	105	16	2.00	3.70	4,800	300
Xeon Platinum 8458P *1	XCC	44	88	82.5	16	2.70	3.80	4,800	350
Xeon Platinum 8452Y *1	XCC	36	72	67.5	16	2.00	3.20	4,800	300
Xeon Gold 6458Q *1	MCC	32	64	60	16	3.10	4.00	4,800	350
Xeon Gold 6454S *1	XCC	32	64	60	16	2.10	3.40	4,800	270
Xeon Gold 6448Y *1	MCC	32	64	60	16	2.10	4.10	4,800	225
Xeon Gold 6444Y *1	MCC	16	32	45	16	3.60	4.10	4,800	270
Xeon Gold 6442Y *1	MCC	24	48	60	16	2.60	4.00	4,800	225
Xeon Gold 6438Y+ *1	MCC	32	64	60	16	2.00	4.00	4,800	205
Xeon Gold 6438N *1	MCC	32	64	60	16	2.00	3.60	4,800	205
Xeon Gold 6438M *1	MCC	32	64	60	16	2.20	3.90	4,800	205
Xeon Gold 6434 *1	MCC	8	16	22.5	16	3.70	4.10	4,800	195
Xeon Gold 6430 *1	XCC	32	64	60	16	1.90	3.40	4,400	270
Xeon Gold 6428N	MCC	32	64	60	16	1.80	3.80	4,000	185
Xeon Gold 6426Y	MCC	16	32	37.5	16	2.50	4.10	4,800	185
Xeon Gold 5420+ *1	MCC	28	56	52.5	16	2.00	4.10	4,400	205
Xeon Gold 5418Y	MCC	24	48	45	16	2.00	3.80	4,400	185
Xeon Gold 5418N	MCC	24	48	45	16	1.80	3.80	4,000	165
Xeon Gold 5416S	MCC	16	32	60	16	2.00	4.00	4,400	150
Xeon Gold 5415+	MCC	8	16	22.5	16	2.90	4.10	4,400	150

Processor model	Type	Number of cores	Number of threads	L3 Cache [MB]	UPI speed [GT/s]	Rated frequency [GHz]	Maximum turbo frequency [GHz]	Maximum memory frequency [MHz]	TDP [W]
1CPU and 2CPU supported processor									
Xeon Silver 4416+	MCC	20	40	37.5	16	2.00	3.90	4,000	165
Xeon Silver 4410Y	MCC	12	24	30	16	2.00	3.90	4,000	150
Xeon Silver 4410T	MCC	10	20	26.25	16	2.70	4.00	4,000	150
1CPU supported processor									
Xeon Gold 5412U *2	MCC	24	48	45	-	2.10	3.90	4,400	185
Xeon Bronze 3408U *2	MCC	8	8	22.5	-	1.80	1.90	4,000	125

*1 Liquid cooling of CX2550 M7 only

*2 Air cooling only

All processors that can be ordered with PRIMERGY CX2550 M7 / CX2560 M7 support Intel Turbo Boost Technology 2.0.
This technology allows you to operate the processor with higher frequencies than the rated frequency. The "maximum turbo frequency" listed in the processor list above is the theoretical maximum frequency when there is only one active core per processor. The maximum frequency that can actually be achieved depends on the number of active cores, current consumption, power consumption, and processor temperature.

As a general rule, Intel does not guarantee that maximum turbo frequencies will be achieved. This is related to manufacturing tolerances, and the performance of each individual processor model varies from each other.

The range of difference covers the range including all of the rated frequency and the maximum turbo frequency.

The turbo function can be set in the BIOS option. Generally, Fujitsu always recommends leaving the [Turbo Mode] option set at the standard setting [Enabled], as performance is substantially increased by the higher frequencies. However, the Turbo Mode frequency depends on the operating conditions mentioned above and is not always guaranteed. The turbo frequency fluctuates in applications where AVX instructions are used intensively and the number of instructions per clock is large. If you need stable performance or want to reduce power consumption, it may be beneficial to set the [Turbo Mode] option to [Disabled] to disable the turbo function.

The processor with the suffix means it is optimized for the following feature.

Suffix	Workload
H	**DB/Analytics** Data analytics and big data usages
M	**Media Transcode** Media, AI, and HPC workloads
N	**Networking** Network and 5G workload environments from edge to the data center
P	**Cloud IaaS** VM environments which require higher frequency
Q	**Liquid Cooled** Environments that require higher core count and higher frequency such as HPC
S	**Storage & HCI** Storage provider and HCI
T	**Long-life Use (IOT)** High reliability and long-life availability usage
U	**1-Socket** Edge server, router, storage and security appliances composed of cost effective 1 socket configuration
V	**Cloud SaaS** VM environments which require power efficiency, higher frequency, and higher core counts
Y	**IaaS, networking, virtualized environments** Environments which require more granular control of CPU performance using Speed Select Technology

Please refer to the below URL for details.

https://www.intel.com/content/www/us/en/support/articles/000059657/processors/intel-xeon-processors.html

Memory modules									
Type	Capacity [GB]	Number of ranks	Bit width of the memory chips	Frequency [MHz]	3DS	Load Reduced	Registered	NVDIMM	ECC
16GB (1x16GB) 1Rx8 DDR5-4800 R ECC	16	1	8	4,800			✓		✓
32GB (1x32GB) 2Rx8 DDR5-4800 R ECC	32	2	8	4,800			✓		✓
32GB (1x32GB) 1Rx4 DDR5-4800 R ECC	32	1	4	4,800			✓		✓
64GB (1x64GB) 2Rx4 DDR5-4800 R ECC	64	2	4	4,800			✓		✓
128GB (1x128GB) 4Rx4 DDR5-4800 R 3DS ECC	128	4	4	4,800	✓		✓		✓
256GB (1x256GB) 8Rx4 DDR5-4800 R 3DS ECC	256	8	4	4,800	✓		✓		✓

Power supplies		Maximum number
Modular redundant PSU	2,600W platinum PSU	2
	2,600W titanium PSU	2

Includes components that will be supported after the system release. Also, some components may not be available in all countries or sales regions.

Detailed technical information is available in the data sheet of PRIMERGY CX2550 M7 / CX2560 M7.

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	PRIMERGY CX2550 M7 / CX2560 M7
• Processor	2 x 4th Generation Intel Xeon Scalable Processors Family or 1 x 4th Generation Intel Xeon Scalable Processors Family
• Memory	16 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (2CPU configuration) or 8 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (1CPU configuration)
Software	
• BIOS settings	SPECspeed2017_int_base: • RdCur for XPT Prefetch = Enable • Adjacent Cache Line Prefetch = Disabled • Package C State limit = C0 • SNC(Sub NUMA) = Enable SNC2 (Disabled when MCC are installed) • HWPM Support = Disabled • AVX P1 = Level2 • CPU Performance Boost = Aggressive SPECSpeed2017_fp_base: • Hyper Threading = Disabled • DCU IP Prefetcher = Disabled • Package C State limit = C0 • LLC Prefetch = Enabled • DBP-F = Enabled • CPU Performance Boost = Aggressive SPECrate2017_int_base: • DCU Streamer Prefetcher = Disabled • Package C State limit = C0 • CPU Performance Boost = Aggressive • SNC(Sub NUMA) =Enable SNC4 SPECrate2017_fp_base: • Hyper Threading = Disabled (Enabled when MCC are installed) • Package C State limit = C0 • CPU Performance Boost = Aggressive • SNC (Sub NUMA) =Enable SNC4 (Enable SNC2 when MCC are installed)
• Operating system	SUSE Linux Enterprise Server 15 SP4 5.14.21-150400.22-default
• Operating system settings	Stack size set to unlimited using "ulimit -s unlimited"
• Compiler	C/C++: Version 2023.0 of Intel C/C++ Compiler for Linux Fortran: Version 2023.0 of Intel Fortran Compiler for Linux

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 model	Number of cores	Number of processors	SPECrate2017_int_base		SPECrate2017_fp_base	
			CX2550 M7	CX2560 M7	CX2550 M7	CX2560 M7
2CPU configuration						
Xeon Max 9480	56	2	**888**	Unsupported	**1120**	Unsupported
Xeon Max 9470	52	2	846 est.	Unsupported	1060 est.	Unsupported
Xeon Max 9468	48	2	784 est.	Unsupported	1010 est.	Unsupported
Xeon Max 9462	32	2	613 est.	Unsupported	846 est.	Unsupported
Xeon Max 9460	40	2	688 est.	Unsupported	932 est.	Unsupported
Xeon Platinum 8490H	60	2	**980**	Unsupported	**986**	Unsupported
Xeon Platinum 8480+	56	2	**933**	Unsupported	**957**	Unsupported
Xeon Platinum 8470Q	52	2	**901**	Unsupported	**939**	Unsupported
Xeon Platinum 8470N	52	2	**819**	Unsupported	**870**	Unsupported
Xeon Platinum 8470	52	2	**892**	Unsupported	**930**	Unsupported
Xeon Platinum 8468V	48	2	**818**	Unsupported	**876**	Unsupported
Xeon Platinum 8468	48	2	**853**	Unsupported	**908**	Unsupported
Xeon Platinum 8462Y+	32	2	667 est.	Unsupported	771 est.	Unsupported
Xeon Platinum 8460Y+	40	2	**691**	Unsupported	**806**	Unsupported
Xeon Platinum 8458P	44	2	807 est.	Unsupported	874 est.	Unsupported
Xeon Platinum 8452Y	36	2	626 est.	Unsupported	724 est.	Unsupported
Xeon Gold 6548Q	32	2	**711**	Unsupported	**807**	Unsupported
Xeon Gold 6454S	32	2	**560**	Unsupported	**684**	Unsupported
Xeon Gold 6448Y	32	2	579 est.	Unsupported	694 est.	Unsupported
Xeon Gold 6444Y	16	2	388 est.	Unsupported	524 est.	Unsupported
Xeon Gold 6442Y	24	2	**490**	Unsupported	**631**	Unsupported
Xeon Gold 6438Y+	32	2	550 est.	Unsupported	650 est.	Unsupported
Xeon Gold 6438N	32	2	547 est.	Unsupported	655 est.	Unsupported
Xeon Gold 6438M	32	2	**558**	Unsupported	**659**	Unsupported
Xeon Gold 6434	8	2	**198**	Unsupported	**283**	Unsupported
Xeon Gold 6430	32	2	529 est.	Unsupported	636 est.	Unsupported
Xeon Gold 6428N	32	2	**501**	**503**	589	591
Xeon Gold 6426Y	16	2	333 est.	335 est.	446 est.	448 est.
Xeon Gold 5420+	28	2	480 est.	Unsupported	593 est.	Unsupported
Xeon Gold 5418Y	24	2	418 est.	**419**	535 est.	537 est.

Processor model	Number of cores	Number of processors	SPECrate2017_int_base		SPECrate2017_fp_base	
			CX2550 M7	CX2560 M7	CX2550 M7	CX2560 M7
2CPU configuration						
Xeon Gold 5418N	24	2	395 est.	**397**	483 est.	**485**
Xeon Gold 5416S	16	2	283 est.	284 est.	385 est.	386 est.
Xeon Gold 5415+	8	2	**178**	**178**	**254**	255 est.
Xeon Silver 4416+	20	2	**362**	364 est.	**457**	459 est.
Xeon Silver 4410Y	12	2	216 est.	**217**	321 est.	**322**
Xeon Silver 4410T	10	2	212 est.	213 est.	293 est.	294 est.
1CPU configuration						
Xeon Gold 5412U	24	1	220 est.	**221**	282 est.	**283**
Xeon Bronze 3408U	8	1	43.2 est.	43.3 est.	75.6 est.	75.9 est.

Processor model	Number of cores	Number of processors	SPECspeed2017_int_base		SPECspeed2017_fp_base	
			CX2550 M7	CX2560 M7	CX2550 M7	CX2560 M7
Xeon Platinum 8490H	60	2	-	Unsupported	**356**	Unsupported
Xeon Platinum 8462Y+	32	2	**16.0**	Unsupported	-	Unsupported
Xeon Gold 6428N	32	2	-	-	**277**	277 est.
Xeon Gold 6426Y	16	2	**15.8**	15.8 est.	-	-

The following graphs compare the throughputs of PRIMERGY CX2550 M7 / CX2560 M7 and their older models, PRIMERGY CX2550 M6 / CX2560 M6, with maximum performance configurations. Both models showed significant performance improvements over the previous generation.

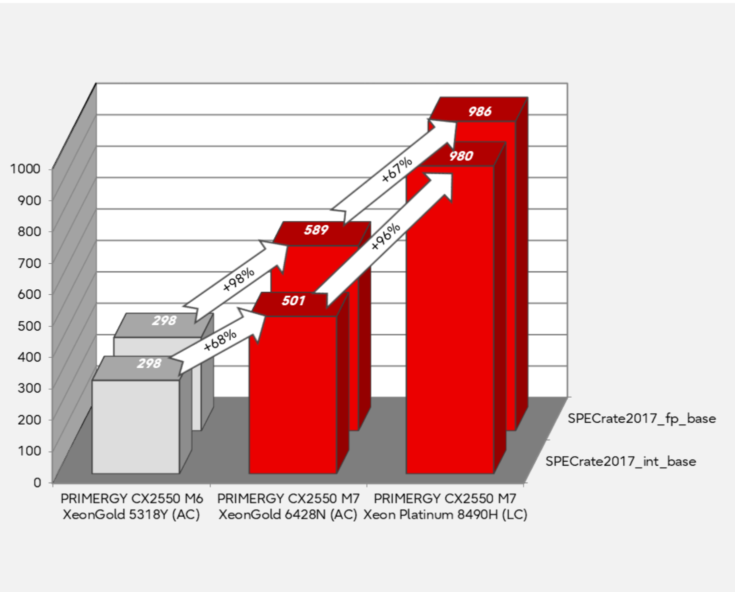

SPECrate2017: Comparison of PRIMERGY CX2550 M6 and PRIMERGY CX2550 M7

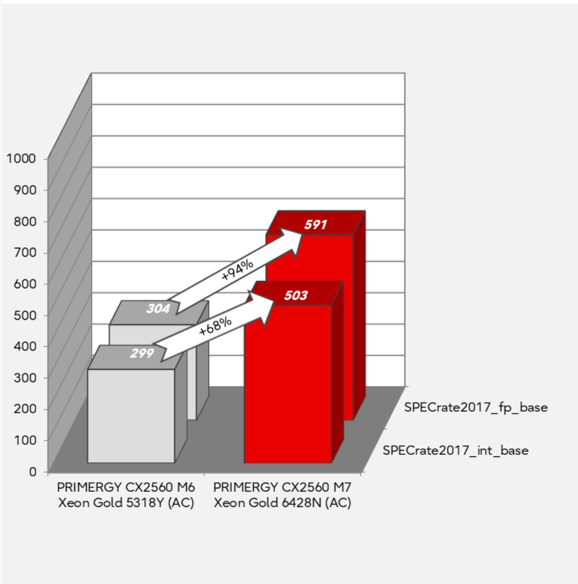

SPECrate2017: Comparison of PRIMERGY CX2560 M6 and PRIMERGY CX2560 M7

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 × 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.

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	PRIMERGY CX2550 M7 / CX2560 M7
• Processor	2 x 4th Generation Intel Xeon Scalable Processors Family or 1 x 4th Generation Intel Xeon Scalable Processors Family
• Memory	16 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (2CPU configuration) or 8 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (1CPU configuration)
Software	
• BIOS settings	• DCU Streamer Prefetcher = Disabled • SNC(Sub NUMA) = Enable SNC4 (Enable SNC2 when MCC type installed)) • Intel Virtualization Technology = Disabled • LLC Dead Line Alloc = Disabled • Stale Atos = Enabled
• Operating system	SUSE Linux Enterprise Server 15 SP4 5.14.21-150400.22-default
• Operating system settings	Default
• Compiler	C/C++: Version 2023.0 of Intel C/C++ Compiler for Linux
• Benchmark	STREAM Version 5.10

Benchmark results

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

Processor	Maximum Memory frequency [MHz]	Maximum memory bandwidth [GB/s]	Number of cores	Rated frequency [GHz]	Number of processors	TRIAD [GB/s]	
						CX2550 M7	CX2560 M7
2CPU configuration							
Xeon Max 9480	4,800	307	56	1.90	2	**516**	Unsupported
Xeon Max 9470	4,800	307	52	2.00	2	513 est.	Unsupported
Xeon Max 9468	4,800	307	48	2.10	2	514 est.	Unsupported
Xeon Max 9462	4,800	307	32	2.70	2	491 est.	Unsupported
Xeon Max 9460	4,800	307	40	2.20	2	514 est.	Unsupported
Xeon Platinum 8490H	4,800	307	60	1.90	2	**523**	Unsupported
Xeon Platinum 8480+	4,800	307	56	2.00	2	**518**	Unsupported
Xeon Platinum 8470Q	4,800	307	52	2.10	2	**492**	Unsupported
Xeon Platinum 8470N	4,800	307	52	1.70	2	**487**	Unsupported
Xeon Platinum 8470	4,800	307	52	2.00	2	**511**	Unsupported
Xeon Platinum 8468V	4,800	307	48	2.40	2	**490**	Unsupported
Xeon Platinum 8468	4,800	307	48	2.10	2	**485**	Unsupported
Xeon Platinum 8462Y+	4,800	307	32	2.80	2	474 est.	Unsupported
Xeon Platinum 8460Y+	4,800	307	40	2.00	2	**469**	Unsupported
Xeon Platinum 8458P	4,800	307	44	2.70	2	498 est.	Unsupported
Xeon Platinum 8452Y	4,800	307	36	2.00	2	452 est.	Unsupported
Xeon Gold 6548Q	4,800	307	32	3.10	2	**444**	Unsupported
Xeon Gold 6454S	4,800	307	32	2.10	2	**445**	Unsupported
Xeon Gold 6448Y	4,800	307	32	2.10	2	467 est.	Unsupported
Xeon Gold 6444Y	4,800	307	16	3.60	2	383 est.	Unsupported
Xeon Gold 6442Y	4,800	307	24	2.60	2	**454**	Unsupported
Xeon Gold 6438Y+	4,800	307	32	2.00	2	463 est.	Unsupported
Xeon Gold 6438N	4,800	307	32	2.00	2	464 est.	Unsupported
Xeon Gold 6438M	4,800	307	32	2.20	2	464 est.	Unsupported
Xeon Gold 6434	4,800	307	8	3.70	2	**234**	Unsupported
Xeon Gold 6430	4,400	282	32	1.90	2	419 est.	Unsupported
Xeon Gold 6428N	4,000	256	32	1.80	2	**407**	**408**
Xeon Gold 6426Y	4,800	307	16	2.50	2	348 est.	349 est.
Xeon Gold 5420+	4,400	282	28	2.00	2	418 est.	Unsupported
Xeon Gold 5418Y	4,400	282	24	2.00	2	387 est.	**387**
Xeon Gold 5418N	4,000	256	24	1.80	2	361 est.	**361**
Xeon Gold 5416S	4,400	282	16	2.00	2	287 est.	**288**

Processor	Maximum Memory frequency [MHz]	Maximum memory bandwidth [GB/s]	Number of cores	Rated frequency [GHz]	Number of processors	TRIAD [GB/s] CX2550 M7		TRIAD [GB/s] CX2560 M7	
2CPU configuration									
Xeon Gold 5415+	4,400	282	8	2.90	2	**215**		215	est.
Xeon Silver 4416+	4,000	256	20	2.00	2	**331**		332	est.
Xeon Silver 4410Y	4,000	256	12	2.00	2	266	est.	**267**	
Xeon Silver 4410T	4,000	256	10	2.70	2	239	est.	239	est.
1CPU configuration									
Xeon Gold 5412U	4,400	282	24	2.10	1	207	est.	**207**	
Xeon Bronze 3408U	4,000	256	8	1.80	1	123	est.	123	est.

The following graphs compare the throughputs of PRIMERGY CX2550 M7 / CX2560 M7 and their older models, PRIMERGY CX2550 M6 / CX2560 M6, with maximum performance configurations.

Both models showed significant performance improvements over the previous generation.

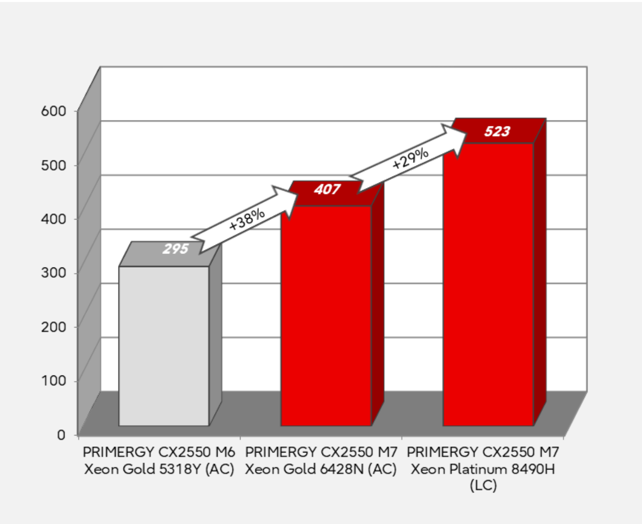

STREAM: Comparison of PRIMERGY CX2550 M6 and PRIMERGY CX2560 M7

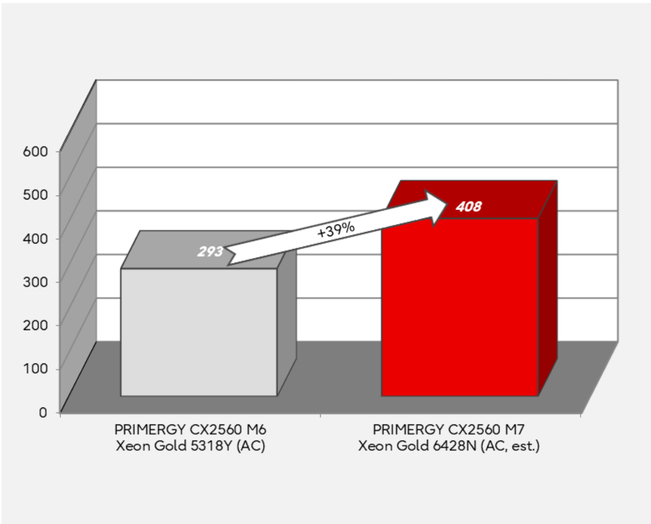

STREAM: Comparison of PRIMERGY CX2550 M6 and PRIMERGY CX2560 M7

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. The description can be found in the following document.

https://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 x 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 x n matrix the number of arithmetic operations required for the solution is $2/3n^3 + 2n^2$. Thus, the choice of n determines the duration of the measurement. In other words, if n is doubled, the measurement time will be approximately eight times longer. The size of n also has an influence on the measurement result itself. As n increases, the measured value asymptotically approaches its 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. he 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: 1 billion floating point operations/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

x Number of processor cores of the computer

x 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/. This requires using an HPL-based LINPACK version (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 https://www.intel.com/content/www/us/en/developer/articles/technical/onemkl-benchmarks-suite.html.

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

Benchmark environment

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY CX2550 M7 / CX2560 M7
• Processor	2 x 4th Generation Intel Xeon Scalable Processors Family or 1 x 4th Generation Intel Xeon Scalable Processors Family
• Memory	16 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (2CPU configuration) or 8 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (1CPU configuration)
Software	
• BIOS settings	• HyperThreading = Disabled • CPU Performance Boost = Agressive
• Operating system	SUSE Linux Enterprise Server 15 SP4 5.14.21-150400.22-default
• Operating system settings	Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1)
• Compiler	C/C++: Version 2023.0 of Intel C/C++ Compiler for Linux
• Benchmark	Intel Optimized MP LINPACK Benchmark for Clusters

Benchmark results

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

Processor	Number of cores	Rated frequency [GHz]	Number of processors	Rpeak [GFlops]	CX2550 M7 Rmax [GFlops]		CX2550 M7 Effic.	CX2560 M7 Rmax [GFlops]		CX2560 M7 Effic.
2CPU configuration										
Xeon Max 9480	56	1.90	2	6,810	**6,781**		100%	Unsupported		
Xeon Max 9470	52	2.00	2	6,656	6,706	est.	101%	Unsupported		
Xeon Max 9468	48	2.10	2	6,451	6,616	est.	103%	Unsupported		
Xeon Max 9462	32	2.70	2	5,530	5,545	est.	123%	Unsupported		
Xeon Max 9460	40	2.20	2	5,632	6,092	est.	108%	Unsupported		
Xeon Platinum 8490H	60	1.90	2	7,296	**7,584**		100%	Unsupported		
Xeon Platinum 8480+	56	2.00	2	7,168	**7,293**		102%	Unsupported		
Xeon Platinum 8470Q	52	2.10	2	6,989	**7,051**		101%	Unsupported		
Xeon Platinum 8470N	52	1.70	2	5,658	**6,264**		111%	Unsupported		
Xeon Platinum 8470	52	2.00	2	6,656	**7,051**		106%	Unsupported		
Xeon Platinum 8468V	48	2.40	2	7,373	**7,022**		95%	Unsupported		
Xeon Platinum 8468	48	2.10	2	6,451	**6,698**		104%	Unsupported		
Xeon Platinum 8462Y+	32	2.80	2	5,734	5,810	est.	101%	Unsupported		
Xeon Platinum 8460Y+	40	2.00	2	5,120	**5,538**		108%	Unsupported		
Xeon Platinum 8458P	44	2.70	2	7,603	6,483	est.	85%	Unsupported		
Xeon Platinum 8452Y	36	2.00	2	4,608	5,444	est.	118%	Unsupported		
Xeon Gold 6548Q	32	3.10	2	6,349	**6,160**		97%	Unsupported		
Xeon Gold 6454S	32	2.10	2	4,301	**4,667**		109%	Unsupported		
Xeon Gold 6448Y	32	2.10	2	4,301	4,724	est.	110%	Unsupported		
Xeon Gold 6444Y	16	3.60	2	3,686	3,679	est.	100%	Unsupported		
Xeon Gold 6442Y	24	2.60	2	3,994	**4,133**		103%	Unsupported		
Xeon Gold 6438Y+	32	2.00	2	4,096	4,426	est.	108%	Unsupported		
Xeon Gold 6438N	32	2.00	2	4,096	4,536	est.	111%	Unsupported		
Xeon Gold 6438M	32	2.20	2	4,506	4,600	est.	102%	Unsupported		
Xeon Gold 6434	8	3.70	2	1,894	1,933	est.	102%	Unsupported		
Xeon Gold 6430	32	1.90	2	3,891	4,544	est.	117%	Unsupported		
Xeon Gold 6428N	32	1.80	2	3,686	**4,025**		109%	4,132	est.	112%
Xeon Gold 6426Y	16	2.50	2	2,560	3,006	est.	117%	3,086	est.	121%
Xeon Gold 5420+	28	2.00	2	3,584	4,125	est.	115%	Unsupported		
Xeon Gold 5418Y	24	2.00	2	3,072	3,452	est.	112%	**3,544**		115%
Xeon Gold 5418N	24	1.80	2	2,765	3,080	est.	111%	3,162	est.	114%
Xeon Gold 5416S	16	2.00	2	2,048	2,303	est.	112%	2,364	est.	115%

Processor	Number of cores	Rated frequency [GHz]	Number of processors	Rpeak [GFlops]	CX2550 M7 Rmax [GFlops]		CX2550 M7 Effic.	CX2560 M7 Rmax [GFlops]		CX2560 M7 Effic.
2CPU configuration										
Xeon Gold 5415+	8	2.90	2	1,485	**1,590**		107%	1,632	est.	110%
Xeon Silver 4416+	20	2.00	2	2,560	**2,940**		115%	3,018	est.	118%
Xeon Silver 4410Y	12	2.00	2	1,536	1,833	est.	119%	1,881	est.	122%
Xeon Silver 4410T	10	2.70	2	1,728	1,964	est.	114%	2,016	est.	117%
1CPU configuration										
Xeon Gold 5412U	24	2.10	1	1,613	1,840	est.	114%	**1,888**		117%
Xeon Bronze 3408U	8	1.80	1	230	255	est.	111%	261	est.	113%

Rpeak values in the table above were calculated by the base frequency of each processor. Since we enabled Turbo mode in the measurements, the average Turbo frequency exceeded the base frequency for some processors.

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

The following graphs compare the throughputs of PRIMERGY CX2550 M7 / CX2560 M7 and their older models, PRIMERGY CX2550 M6 / CX2560 M6, with maximum performance configurations. Both models showed significant performance improvements over the previous generation.

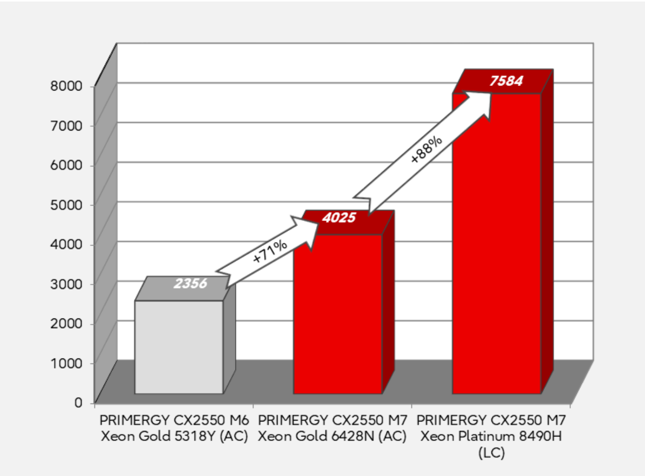

LINPACK: Comparison of PRIMERGY CX2550 M6 and PRIMERGY CX2550 M7

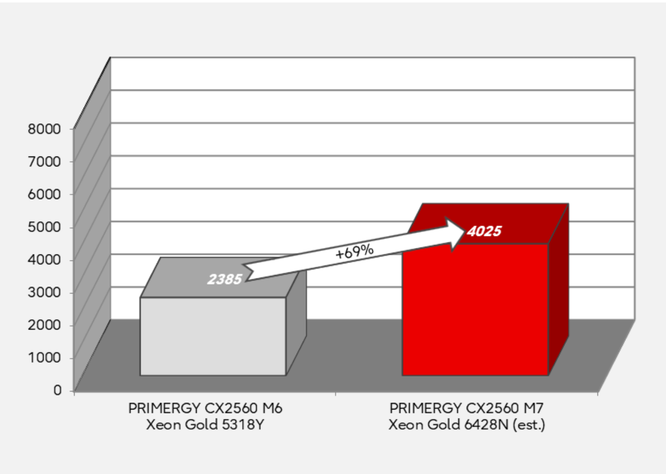

LINPACK: Comparison of PRIMERGY CX2560 M6 and PRIMERGY CX2560 M7

SPECpower_ssj2008

Benchmark description

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

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

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

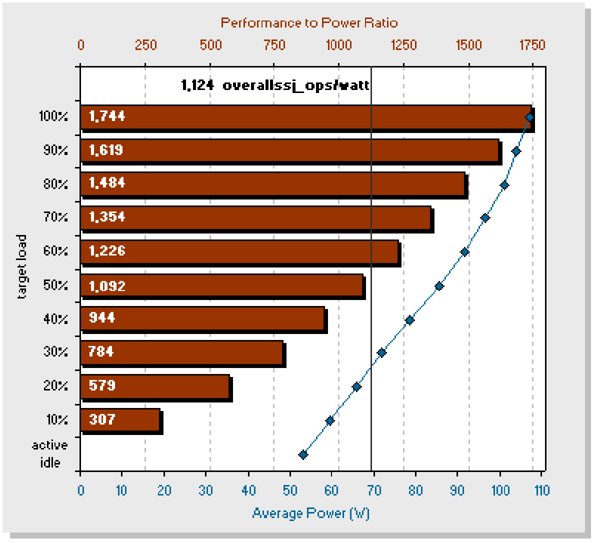

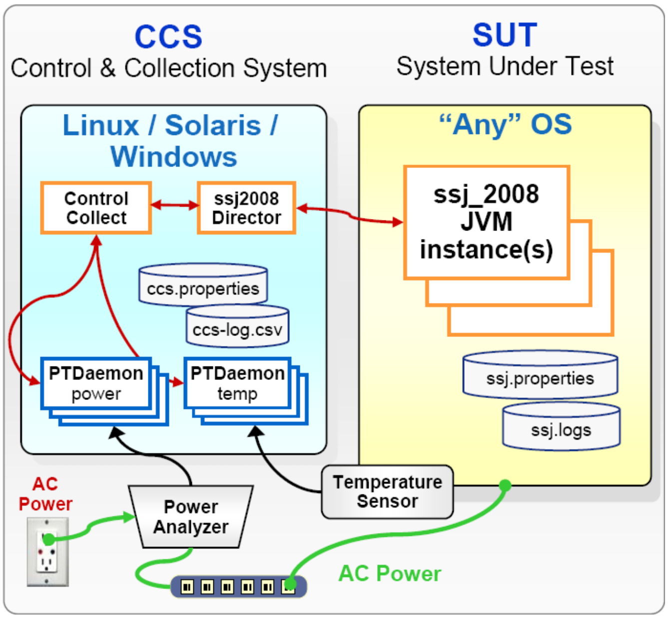

The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System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

Benchmark environment

System Under Test (SUT)	
Hardware (chassis)	
• Enclosure	PRIMERGY CX400 M7
• Power Supply Unit	2 x 2600W platinum PSU
• Number of servers	4
• Model	PRIMERGY CX2560 M7
Hardware (per node)	
• Processor	2 x Xeon Gold 6428N 32C 1.80GHz 185W
• Memory	16 x 16GB (1x16GB) 1Rx8 DDR5-4800 R ECC
• Network interface	1Gbit/s (RJ45) on Motherboard
• Disk subsystem	1 x SSD SATA M.2 drive for booting, non hot-plug 240GB
Software	
• BIOS settings	ASPM Support = Auto Hardware Prefetcher = Disabled Adjacent Cache Line Prefetch = Disabled DCU Streamer Prefetcher = Disabled Intel(R) VT-d = Disabled Package C State limit = No limit Uncore Frequency Scaling = Power balanced CPU Performance Boost = Aggressive SNC(Sub NUMA) = Enable SNC2 SATA Controller = Disabled USB Port Control = Disable all ports Network Stack = Disabled
• Operating system	Windows Server 2022 Standard
• Operating system settings	Turn off hard disk after = 1 Minute PCI Express Link State Power Management = Maximum power savings Minimum processor state = 0% Maximum processor state = 100% Turn off display after = 1 Minute POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFBOOSTMODE 4 POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFINCTHRESHOLD 95 POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFDECTHRESHOLD 93 POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFDECTIME 1 POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR IDLESCALING 1 POWERCFG /S SCHEME_CURRENT Using the local security settings console, "lock pages in memory" was enabled for the user running the benchmark. Benchmark was started via Windows Remote Desktop Connection.
• JVM	Oracle Java HotSpot(TM) 64-Bit Server VM 18.9 (build 11.0.16.1+1-LTS, mixed mode)
• JVM settings	-server -Xmn1500m -Xms1625m -Xmx1625m -XX:+UseLargePages -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:MaxTenuringThreshold=15 -XX:ParallelGCThreads=2 -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:-UseAdaptiveSizePolicy -XX:+UseParallelOldGC -XX:FreqInlineSize=2500 -XX:LoopUnrollLimit=45 -XX:InitialTenuringThreshold=12 -XX:-ThreadLocalHandshakes -XX:UseAVX=0

Benchmark results

The PRIMERGY CX2560 M7 in Microsoft Windows Server 2022 Standard achieved the following result:

SPECpower_ssj2008 = 11,580 overall ssj_ops/watt

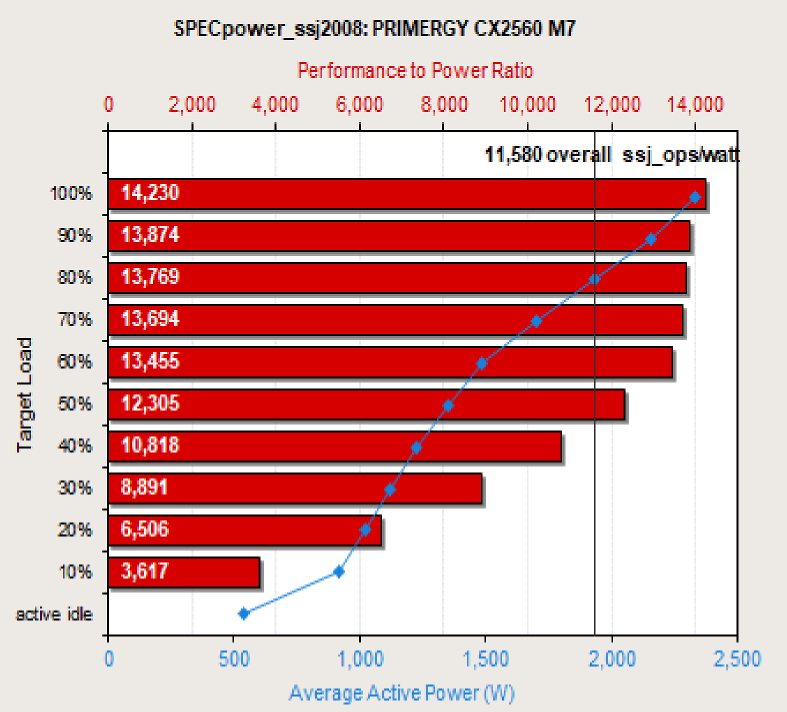

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 11,580 overall ssj_ops/watt for the PRIMERGY CX2560 M7. This is the quotient of the sum of the transaction throughputs for each load level and the sum of the average power consumed for each measurement interval.

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

Performance		Power	Energy Efficiency
Target Load	ssj_ops	Average Power (W)	ssj_ops/watt
100%	33,244,766	2,336	14,230
90%	29,919,914	2,157	13,874
80%	26,629,892	1,934	13,769
70%	23,273,792	1,700	13,694
60%	19,973,092	1,484	13,455
50%	16,626,128	1,351	12,305
40%	13,293,306	1,229	10,818
30%	9,985,041	1,123	8,891
20%	6,659,775	1,024	6,506
10%	3,322,452	919	3,617
Active Idle	0	541	0
			Σ ssj_ops / Σ power = 11,580

Comparison with the predecessor

The following diagram shows for each load level (on the x-axis) the throughput (on the left y-axis) and the power consumption (on the right y-axis) of the PRIMERGY CX2560 M7 compared to the predecessor PRIMERGY CX2560 M6.

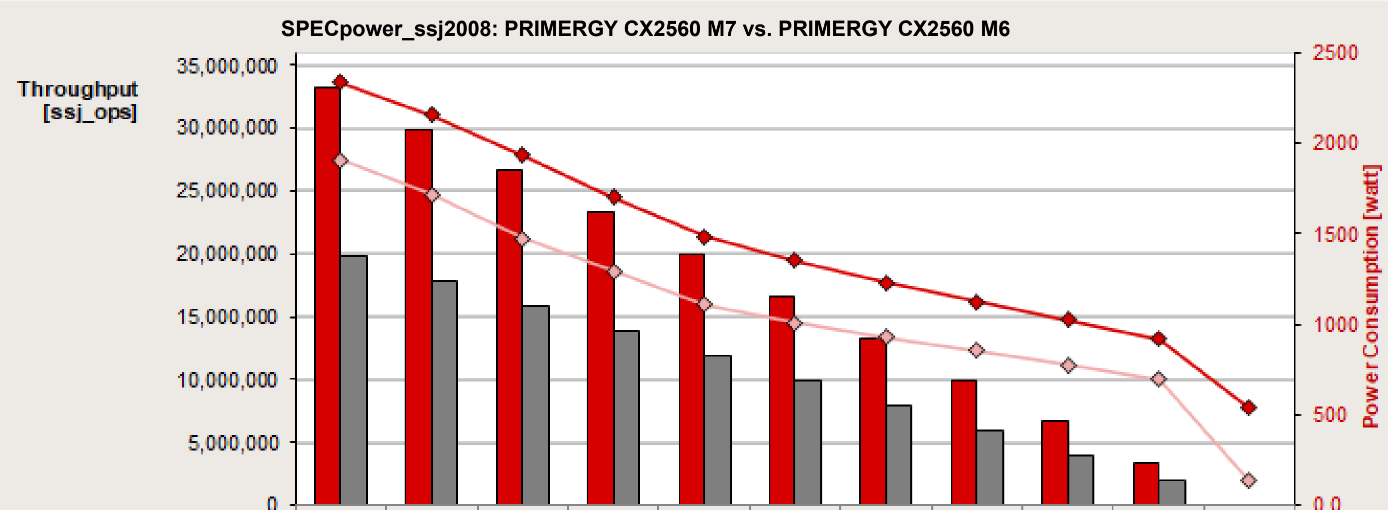

Load Level	100%	90%	80%	70%	60%	50%	40%	30%	20%	10%	active idle
CX2560 M7 ssj_ops	33,244,766	29,919,914	26,629,892	23,273,792	19,973,092	16,626,128	13,293,306	9,985,041	6,659,775	3,322,452	0
CX2560 M6 ssj_ops	19,762,291	17,822,302	15,860,766	13,904,810	11,895,706	9,899,337	7,920,473	5,940,102	3,961,149	1,980,785	0
CX2560 M7 watt	2336	2157	1934	1700	1484	1351	1229	1123	1024	919	541
CX2560 M6 watt	1906	1715	1474	1292	1109	1009	929	856	778	696	135

Thanks to the 4th Generation Intel Xeon Scalable Processors Family, the PRIMERGY CX2560 M7 has a higher throughput. This results in an overall 30.1% increase in energy efficiency in the PRIMERGY CX2560 M7.

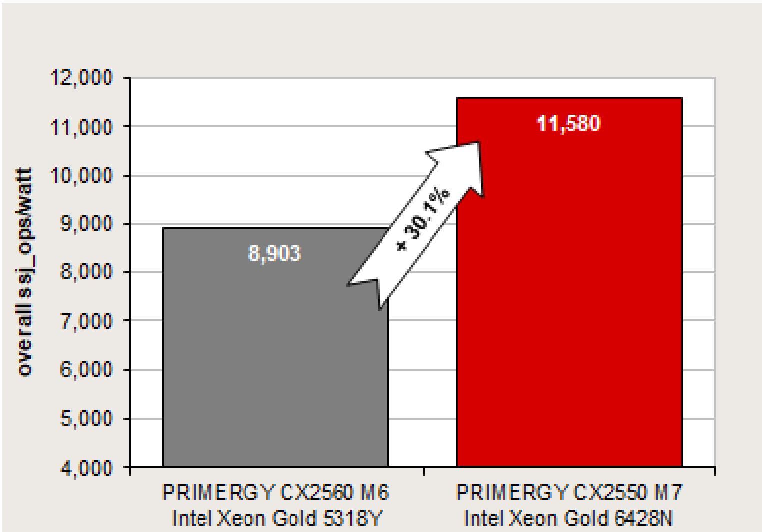

Disk I/O: Performance of storage media

Benchmark description

Performance measurements of disk subsystems for PRIMERGY 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 (kB)
- 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 [kB]	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 = 10^{12} bytes) is used to indicate the capacity of the hard storage medium, and a power of 2 (1 MiB / s = 2^{20} 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: PRAID EP540i

Storage media	Category	Drive name
HDD	SAS HDD (SAS 12Gbps, 15k rpm) [512n]	ST300MP0006 * ST600MP0006 *
	SAS HDD (SAS 12Gbps, 10k rpm) [512e]	AL15SEB18EQ * AL15SEB24EQ *
	SAS HDD (SAS 12Gbps, 10k rpm) [512n]	AL15SEB060N * AL15SEB120N *
	NL-SAS HDD (SAS 12Gbps, 7.2k rpm) [512n	ST2000NX0433 *
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512n]	ST1000NX0423 ST2000NX0403
SSD	SAS SSD (SAS 12Gbps, Write Intensive)	XS400ME70084 * XS800ME70084 * XS1600ME70084 *
	SAS SSD (SAS 12Gbps, Mixed Use)	XS800LE70084 * XS1600LE70084 * XS3200LE70084 *
	SAS SSD (SAS 12Gbps, Read Intensive)	XS960SE70084 * XS1920SE70084 * XS3840SE70084 * XS7680SE70084 *
	SATA SSD (SATA 6Gbps, Mixed Use)	MTFDDAK480TDT MTFDDAK960TDT MTFDDAK1T9TDT MTFDDAK3T8TDT MZ7L3480HBLT MZ7L3960HBLT MZ7L31T9HBNA MZ7L33T8HBNA

* Not Supported drives on CX2550 M7.

Controller: PRAID EP540i

Storage media	Category	Drive name
	SATA SSD (SATA 6Gbps, Read Intensive)	MTFDDAK240TDS
		MTFDDAK480TDS
		MTFDDAK960TDS
		MTFDDAK1T9TDS
		MTFDDAK3T8TDS
		MTFDDAK7T6TDS
		MZ7L3240HCHQ
		MZ7L3480HCHQ
		MZ7L3960HCJR
		MZ7L31T9HBLT
		MZ7L33T8HBLT
		MZ7L37T6HBLA

Controller: PRAID EP680i

Storage media	Category	Drive name
SSD	PCIe SSD (Write intensive)	SSDPF21Q400GB
		SSDPF21Q800GB
		SSDPF21Q016TB
	PCIe SSD (Mixed Use)	KCM61VUL1T60
		KCM61VUL3T20
		KCM61VUL6T40
	PCIe SSD (Read intensive)	KCM61RUL960G
		KCM61RUL1T92
		KCM61RUL3T84
		KCM61RUL7T68

Controller: Intel C741 Standard SATA AHCI controller

Storage media	Category	Drive name
SSD	M.2 Flash module	MTFDDAV240TDS
		MTFDDAV480TDS

Controller: Intel C741 Standard NVM Express controller

Storage media	Category	Drive name
SSD	M.2 Flash module (NVMe)	MTFDKBA480TFR
		MTFDKBA960TFR

Software		
Operating system		Microsoft Windows Server 2019 Standard
Benchmark version		3.0
RAID type		Type RAID 0 logical drive consisting of 1 hard disk
Stripe size		HDD: 256KB, SSD: 64 KB
Measuring tool		Iometer 1.1.0
Measurement area	HDD, SSD (Except M.2)	RAW file system is used. The first 32GB of available LBA space is used for sequential access. The following 64GB is used for random access.
	SSD (M.2)	NTFS file system is used. The first 32GB of available LBA space is used for sequential access. The following 64GB is used for random access.
Total number of Iometer worker		1
Alignment of Iometer accesses		Aligned to access block size

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 media	Controller name	Cache	Supported interfaces		RAID levels
			host	drive	
SSD/HDD	PRAID EP540i	-	PCIe 3.0 x8	SATA 6G SAS 12G PCIe 3.0 x16	0, 1, 1E, 10, 5, 50
PCIe SSD 2.5"	PRAID EP680i	-	PCIe 4.0 x8	SATA 6G SAS 12G PCIe 4.0 x16	0, 1, 1E, 10, 5, 50
M.2 Flash	C741 Standard SATA AHCI controller	-	DMI 3.0 x4	SATA 6G	-
M.2 Flash (NVMe)	C741 Standard NVM Express controller	-	DMI 3.0 x4	PCIe 3.0 x2	

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 used for PRIMERGY servers.

Model	Storage media type	interface	Form factor
2.5 inch model	HDD	SAS 12G	2.5 inch
		SATA 6G	2.5 inch
	SSD	SAS 12G	2.5 inch
		SATA 6G	2.5 inch or M.2
		PCIe 4.0	2.5 inch
		PCIe 4.0	M.2

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 system board is typically PCIe or, in the case of the integrated onboard controllers, an internal bus interface of the system board.

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 PRIMERGY 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 mode "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.

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 (number 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 performance

HDDs

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
☐ SAS HDD 15krpm [512n]							
300	ST300MP0006	SAS 12G	790	696	666	304	304
600	ST600MP0006	SAS 12G	736	651	601	301	300
☐ SAS HDD 10krpm [512e]							
1,800	AL15SEB18EQ	SAS 12G	767	631	624	255	249
2,400	AL15SEB24EQ	SAS 12G	754	620	617	264	260
☐ SAS HDD 10krpm [512n]							
600	AL15SEB060N	SAS 12G	698	586	600	232	232
1,200	AL15SEB120N	SAS 12G	732	604	615	230	226
☐ NL-SAS HDD 7.2krpm [512e]							
2,000	ST2000NX0433	SAS 12G	489	403	388	132	132
☐ BC-SATA HDD 7.2krpm [512n]							
1,000	ST1000NX0423	SATA 6G	415	350	349	131	131
2,000	ST2000NX0403	SATA 6G	459	379	385	133	133

SSDs

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
☐ SAS SSD (WI)							
400	XS400ME70084	SAS 12G	122,956	22,969	19,438	1,052	872
800	XS800ME70084	SAS 12G	123,848	23,784	19,435	1,052	874
1,600	XS1600ME70084	SAS 12G	123,277	23,725	19,270	1,051	884
☐ SAS SSD (MU)							
800	XS800LE70084	SAS 12G	121,914	23,707	19,257	1,052	871
1,600	XS1600LE70084	SAS 12G	122,949	23,771	19,455	1,052	874
3,200	XS3200LE70084	SAS 12G	123,090	22,816	19,418	1,051	872
☐ SAS SSD (RI)							
960	XS960SE70084	SAS 12G	123,014	23,678	19,424	1,052	870
1,920	XS1920SE70084	SAS 12G	123,093	23,760	19,423	1,052	874
3,840	XS3840SE70084	SAS 12G	122,810	22,949	19,406	1,051	871
7,680	XS7680SE70084	SAS 12G	123,461	22,899	19,516	1,051	880
☐ SATA SSD (MU)							
480	MTFDDAK480TDT	SATA 6G	49,138	6,383	6,600	508	437
960	MTFDDAK960TDT	SATA 6G	50,488	6,970	7,136	508	486
1,920	MTFDDAK1T9TDT	SATA 6G	50,669	7,183	7,336	508	487
3,840	MTFDDAK3T8TDT	SATA 6G	49,490	7,115	7,208	493	474
480	MZ7L3480HBLT	SATA 6G	52,039	8,009	7,952	521	487
960	MZ7L3960HBLT	SATA 6G	51,997	8,006	7,968	519	487
1,920	MZ7L31T9HBNA	SATA 6G	51,907	8,026	7,971	520	487
3,840	MZ7L33T8HBNA	SATA 6G	51,799	7,955	7,931	518	487

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
☐ SATA SSD (RI)							
240	MTFDDAK240TDS	SATA 6G	42,594	5,435	5,510	508	301
480	MTFDDAK480TDS	SATA 6G	47,577	6,109	6,310	508	401
960	MTFDDAK960TDS	SATA 6G	50,134	6,633	6,852	506	480
1,920	MTFDDAK1T9TDS	SATA 6G	50,638	7,078	7,286	508	488
3,840	MTFDDAK3T8TDS	SATA 6G	49,542	7,097	7,196	495	477
7,680	MTFDDAK7T6TDS	SATA 6G	47,200	7,134	7,563	508	487
240	MZ7L3240HCHQ	SATA 6G	52,340	8,048	7,958	526	383
480	MZ7L3480HCHQ	SATA 6G	52,168	8,083	8,012	526	487
960	MZ7L3960HCJR	SATA 6G	52,372	8,094	8,023	526	488
1,920	MZ7L31T9HBLT	SATA 6G	52,329	8,072	8,021	526	488
3,840	MZ7L33T8HBLT	SATA 6G	52,229	8,058	8,001	526	487
7,680	MZ7L37T6HBLA	SATA 6G	51,917	8,008	7,957	524	487
☐ PCIe SSD (MU)							
1,600	KCM61VUL1T60	PCIe4 x4	272,211	49,350	47,236	6,649	2,740
3,200	KCM61VUL3T20	PCIe4 x4	314,143	72,898	75,032	6,649	4,062
6,400	KCM61VUL6T40	PCIe4 x4	305,271	67,808	71,273	6,649	3,853
☐ PCIe SSD (RI)							
960	KCM61RUL960G	PCIe4 x4	77,623	9,719	6,428	6,633	1,400
1,920	KCM61RUL1T92	PCIe4 x4	180,706	19,204	12,678	6,649	2,730
3,840	KCM61RUL3T84	PCIe4 x4	315,657	72,526	75,132	6,649	4,048
7,680	KCM61RUL7T68	PCIe4 x4	311,548	68,020	71,191	6,649	3,853
☐ M.2 SATA SSD							
240	MTFDDAV240TDS	SATA 6G	31,923	5,489	5,512	504	299
480	MTFDDAV480TDS	SATA 6G	39,553	6,331	6,516	501	394
☐ M.2 NVMe SSD							
480	MTFDKBA480TFR	PCIe3 x2	74,947	15,849	12,564	1,644	685
960	MTFDKBA960TFR	PCIe3 x2	147,206	31,459	25,928	1,644	1,381

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 SPEC CPU 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:

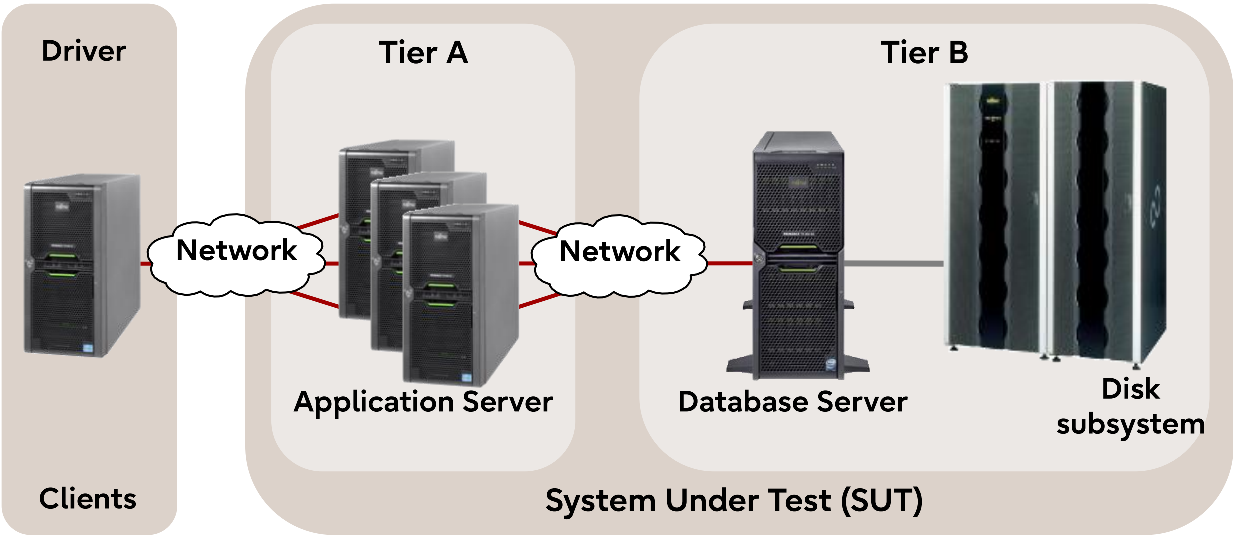

All OLTP-2 results were calculated based on the configuration of the next following pages of PRIMERGY RX2540 M7.

Database Server (Tier B)	
Hardware	
• Model	PRIMERGY RX2540 M7
• Processor	4th Generation Intel Xeon Processor Scalable Family
• Memory	32 x 64 GB (1x64 GB) 2Rx4 DDR5-4800 ECC
• Network interface	1 x PLAN EP X710-DA2 2x10Gb SFP+ 1 x PLAN CP I350-T4 4X 1000BASE-T OCPv3 PT
• Disk subsystem	RX2540 M7 : 1 x RAID controller (internal, 4GB cache) 6 x 1.6 TB SSD drive, RAID10 (log) 5 x RAID controller (external, 4GB cache) 10 x JX40 S2 : 4 x 1.6 TB SSD drive, RAID10（temp） 49 x 1.6 TB SSD drive, RAID5 (data) 30 x 960 GB SSD drive, RAID (data)
Software	
• Operating system	Microsoft Windows Server 2022 Standard
• Database	Microsoft SQL Server 2022 Enterprise

Application Server (Tier A)	
Hardware	
• Model	1 x PRIMERGY RX2530 M4
• Processor	2 x Xeon Platinum 8180 28C 2.5GHz 205W
• Memory	12 x 16GB (1x16GB) 1Rx4 DDR4-2666 R ECC
• Network interface	1 x PLAN EP X710-DA2 2x10Gb SFP+ LP 2 x RJ45 1GbE ports on systemboard
• Disk subsystem	2 x HDD SAS 2.5" 10K 512n (SFF) 300GB
Software	
• Operating system	Microsoft Windows Server 2016 Standard

Client	
Hardware	
• Model	1 x PRIMERGY RX2530 M2
• Processor	2 x Xeon E5-2667v4 8C/16T 3.20GHz 25MB 9.6GT/s 2400MHz 135W
• Memory	8 x 16GB (1x16GB) 2Rx4 DDR4-2400 R ECC
• Network interface	1 x PLAN EM 4x1Gb T interface card
• Disk subsystem	1 x HDD SAS 2.5" 10K 512n (SFF) 300GB
Software	
• Operating system	Microsoft Windows Server 2012 R2 Standard
• Benchmark	OLTP-2 Software EGen version 1.14.0

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 important. This why a configuration with a total memory of 1024 GB was considered for the estimation with two processors. The memory configurations have memory access of 4800 MHz.

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

Processor	Cores	Threads	2CPU Score
Xeon Gold 6428N	32	64	7,311 est.
Xeon Gold 6426Y	16	32	4,362 est.
Xeon Gold 5418Y	24	48	6,026 est.
Xeon Gold 5418N	24	48	5,788 est.
Xeon Gold 5416S	16	32	3,705 est.
Xeon Gold 5415+	8	16	2,422 est.
Xeon Silver 4416+	20	40	4,823 est.
Xeon Silver 4410Y	12	24	3,041 est.
Xeon Silver 4410T	10	20	2,960 est.

The following graph shows the OLTP-2 transaction rates obtained with the 4th Generation Intel Xeon Processor Scalable Family.

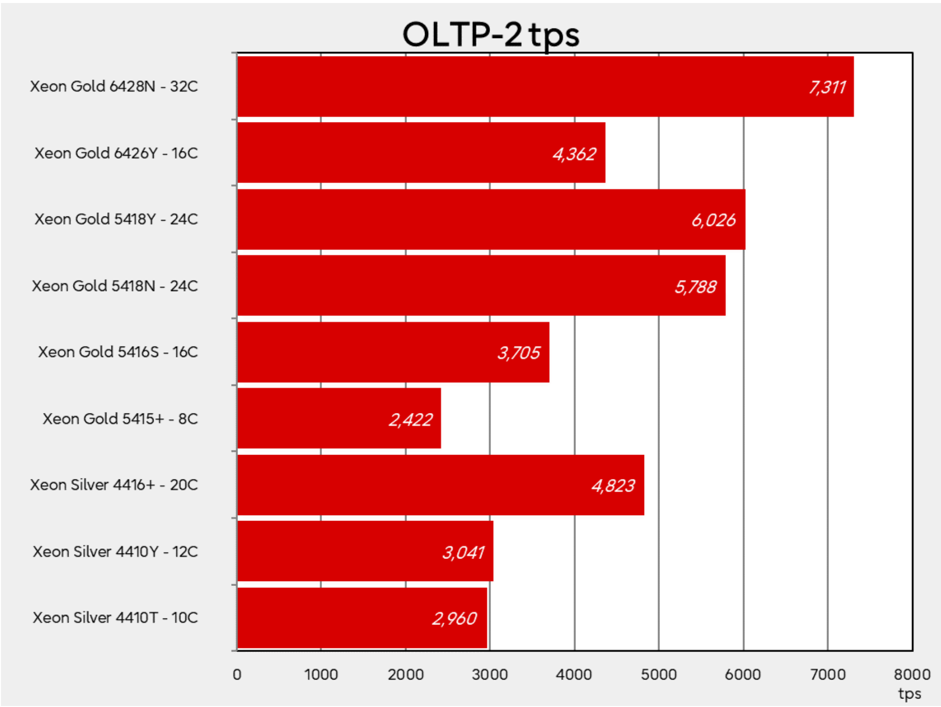

The features of the processors are summarized in the section "Technical data."

In general, 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.

The highest value for OLTP-2 on the current PRIMERGY model is about 24% higher than the highest value on the previous model.

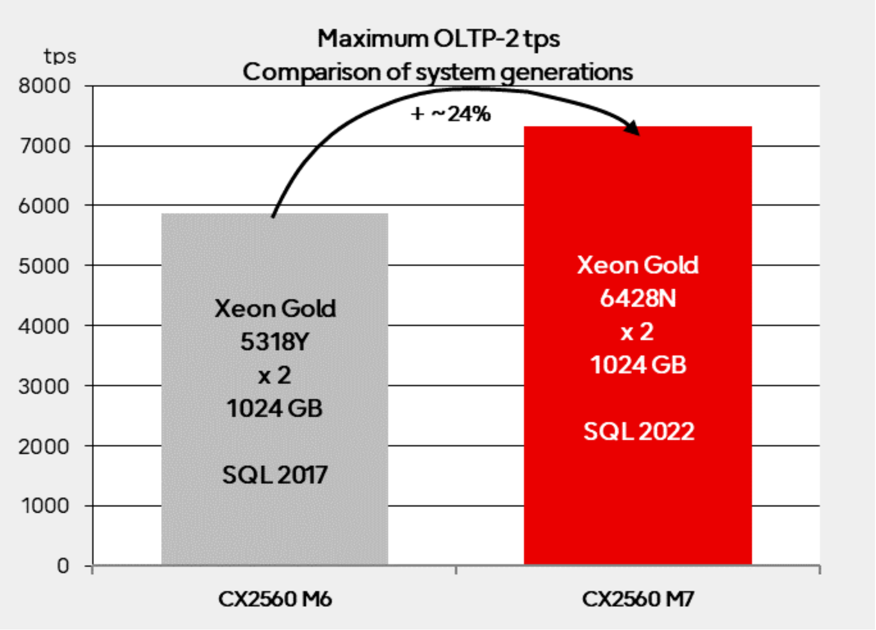

Literature

PRIMERGY Servers

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

PRIMERGY CX2550 M7 / CX2560 M7

This Whitepaper

https://docs.ts.fujitsu.com/dl.aspx?id=b5a565d4-3621-4ee4-abd8-5520f477c68d

https://docs.ts.fujitsu.com/dl.aspx?id=ff75ca36-4545-45dc-9856-7f38c1f5b85e

Data sheet

CX400 M7: https://docs.ts.fujitsu.com/dl.aspx?id=b4d5aa5c-bc24-46c7-8bdb-341d3022b2d0

CX2550 M7: https://docs.ts.fujitsu.com/dl.aspx?id=75d4b451-3e2b-49ba-b986-1ddd7902e904

CX2560 M7: https://docs.ts.fujitsu.com/dl.aspx?id=75238bfe-708f-44fe-b654-e789712a0ce6

PRIMERGY 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/

LINPACK

The LINPACK Benchmark: Past, Present, and Future

https://www.netlib.org/utk/people/JackDongarra/PAPERS/hplpaper.pdf

TOP500

https://www.top500.org/

HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers

https://www.netlib.org/benchmark/hpl/

Intel Math Kernel Library – LINPACK Download

https://www.intel.com/content/www/us/en/developer/articles/technical/onemkl-benchmarks-suite.html

SPECpower_ssj2008

https://www.spec.org/power_ssj2008

Benchmark Overview SPECpower_ssj2008

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

OLTP-2

Benchmark Overview OLTP-2

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

Document change history

Version	Date	Description
1.2	2023-11-07	Update: • Technical data • SPEC CPU2017, STREAM, LINPACK Measured and calculated additionally with 4th Generation Intel Xeon Processor Scalable Family
1.1	2023-07-25	Update: • Technical data, • SPEC CPU2017, STREAM, LINPACK Update Supported CPU
1.0	2023-07-04	New: • Technical data • SPEC CPU2017, STREAM, LINPACK Measured and calculated with 4th Generation Intel Xeon Processor Scalable Family • SPECpower_ssj2008 Measured with Intel Xeon Gold 6428N • Disk I/O Measured with 2.5 inch model • OLTP-2 Calculated with 4th 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