Fujitsu Server PRIMERGY
Performance Report
PRIMERGY TX2550 M7

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMERGY TX2550 M7.

Explaines PRIMERGY TX2550 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
1.0 2023-05-30

Contents

Technical data

PRIMERGY TX2550 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 TX2550 M7
Form factor	Tower / Rack server
Chipset	Intel C741
Number of sockets	2
Number of configurable processors	1 or 2
Processor type	4th Generation Intel Xeon Scalable Processors Family
Number of memory slots	16 (8 per processor)
Maximum memory configuration	4,096 GB
Maximum number of internal storage disks	3.5 inch: 12 2.5 inch: 32
Maximum number of PCI slots	PCI-Express 5.0 (x16 lane): 6 (Low Profile)

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 Platinum 8450H	XCC	28	56	75	16	2.0	3.5	4,800	250
Xeon Platinum 8444H	XCC	16	32	45	16	2.9	4.0	4,800	270
Xeon Gold 6454S	XCC	32	64	60	16	2.2	3.4	4,800	270
Xeon Gold 6448Y	MCC	32	64	60	16	2.1	4.1	4,800	225
Xeon Gold 6442Y	MCC	24	48	60	16	2.6	4.0	4,800	225
Xeon Gold 6438Y+	MCC	32	64	60	16	2.0	4.0	4,800	205
Xeon Gold 6434	MCC	8	16	22.5	16	3.7	4.1	4,800	195
Xeon Gold 6430	XCC	32	64	60	16	2.1	3.4	4,400	270
Xeon Gold 6426Y	MCC	16	32	37.5	16	2.5	4.1	4,800	185
Xeon Gold 5420+	MCC	28	56	52.5	16	2.0	4.1	4,400	205
Xeon Gold 5418Y	MCC	24	48	45	16	2.0	3.8	4,400	185
Xeon Gold 5416S	MCC	16	32	60	16	2.0	4.0	4,400	150
Xeon Gold 5415+	MCC	8	16	22.5	16	2.9	4.1	4,400	150
Xeon Silver 4416+	MCC	20	40	37.5	16	2.0	3.9	4,000	165
Xeon Silver 4410Y	MCC	12	24	30	16	2.0	3.9	4,000	150
Xeon Silver 4410T	MCC	10	20	26.25	16	2.7	4.0	4,000	150
1CPU supported processor									
Xeon Gold 6414U	XCC	32	64	60	-	2.0	3.4	4,800	250
Xeon Gold 5412U	MCC	24	48	45	-	2.1	3.9	4,400	185
Xeon Bronze 3408U	MCC	8	16	22	-	1.8	1.9	4,000	125

All processors that can be ordered with PRIMERGY TX2550 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
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
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	500W platinum PSU	2
	500W titanium PSU	2
	900W platinum PSU	2
	900W titanium PSU	2
	1,600W platinum PSU	2
	1,600W titanium PSU	2
	2,200W platinum PSU	2
	2,400W 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 TX2550 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 TX2550 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 • FAN Control = Full 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 • FAN Control = Full SPECrate2017_int_base: • DCU Streamer Prefetcher = Disabled • Package C State limit = C0 • CPU Performance Boost = Aggressive • SNC(Sub NUMA) =Enable SNC4 • FAN Control = Full 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) • FAN Control = Full
• 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
2CPU Configuration				
Xeon Platinum 8450H	28	2	482 est.	619 est.
Xeon Platinum 8444H	16	2	344 est.	480 est.
Xeon Gold 6454S	32	2	559 est.	679 est.
Xeon Gold 6448Y	32	2	**582**	**700**
Xeon Gold 6442Y	24	2	495 est.	636 est.
Xeon Gold 6438Y+	32	2	552 est.	656 est.
Xeon Gold 6434	8	2	200 est.	289 est.
Xeon Gold 6430	32	2	532 est.	641 est.
Xeon Gold 6426Y	16	2	335 est.	450 est.
Xeon Gold 5420+	28	2	482 est.	598 est.
Xeon Gold 5418Y	24	2	423 est.	528 est.
Xeon Gold 5416S	16	2	284 est.	374 est.
Xeon Gold 5415+	8	2	179 est.	256 est.
Xeon Silver 4416+	20	2	367 est.	461 est.
Xeon Silver 4410Y	12	2	220 est.	325 est.
Xeon Silver 4410T	10	2	213 est.	296 est.
1CPU Configuration				
Xeon Gold 6414U	32	1	269 est.	340 est.
Xeon Gold 5412U	24	1	223 est.	287 est.
Xeon Bronze 3408U	8	1	43.4 est.	76.3 est.

The following graphs compare the throughputs of PRIMERGY TX2550 M7 and their older models, PRIMERGY TX2550 M5, with maximum performance configurations.

PRIMERGY TX2550 M7 showed significant performance improvements over the previous generation.

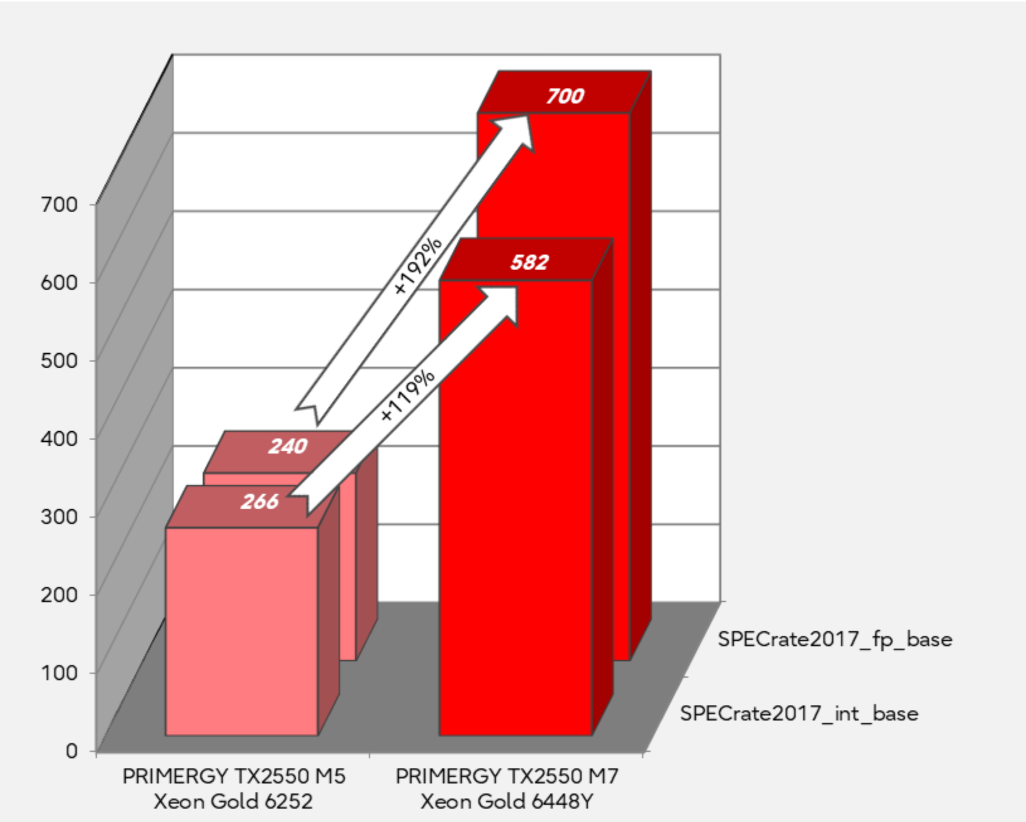

SPECrate2017: Comparison of PRIMERGY TX2550 M5 and PRIMERGY TX2550 M7

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 TX2550 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	Memory frequency [MHz]	Maximum memory bandwidth [GB/s]	Number of cores	Rated frequency [GHz]	Number of processors	TRIAD [GB/s]	
2CPU Configuration							
Xeon Platinum 8450H	4,800	307	28	2.0	2	454	est.
Xeon Platinum 8444H	4,800	307	16	2.9	2	371	est.
Xeon Gold 6454S	4,800	307	32	2.2	2	**428**	
Xeon Gold 6448Y	4,800	307	32	2.1	2	450	est.
Xeon Gold 6442Y	4,800	307	24	2.6	2	425	est.
Xeon Gold 6438Y+	4,800	307	32	2.0	2	446	est.
Xeon Gold 6434	4,800	307	8	3.7	2	219	est.
Xeon Gold 6430	4,400	282	32	2.1	2	404	est.
Xeon Gold 6426Y	4,800	307	16	2.5	2	336	est.
Xeon Gold 5420+	4,400	282	28	2.0	2	403	est.
Xeon Gold 5418Y	4,400	282	24	2.0	2	372	est.
Xeon Gold 5416S	4,400	282	16	2.0	2	274	est.
Xeon Gold 5415+	4,400	282	8	2.9	2	206	est.
Xeon Silver 4416+	4,000	256	20	2.0	2	318	est.
Xeon Silver 4410Y	4,000	256	12	2.0	2	254	est.
Xeon Silver 4410T	4,000	256	10	2.7	2	230	est.
1CPU Configuration							
Xeon Gold 6414U	4,800	307	32	2.0	1	230	est.
Xeon Gold 5412U	4,400	282	24	2.1	1	201	est.
Xeon Bronze 3408U	4,000	256	8	1.8	1	119	est.

The following graphs compare the throughputs of PRIMERGY TX2550 M7 and their older models, PRIMERGY TX2550 M5, with maximum performance configurations.

PRIMERGY TX2550 M7 showed significant performance improvements over the previous generation.

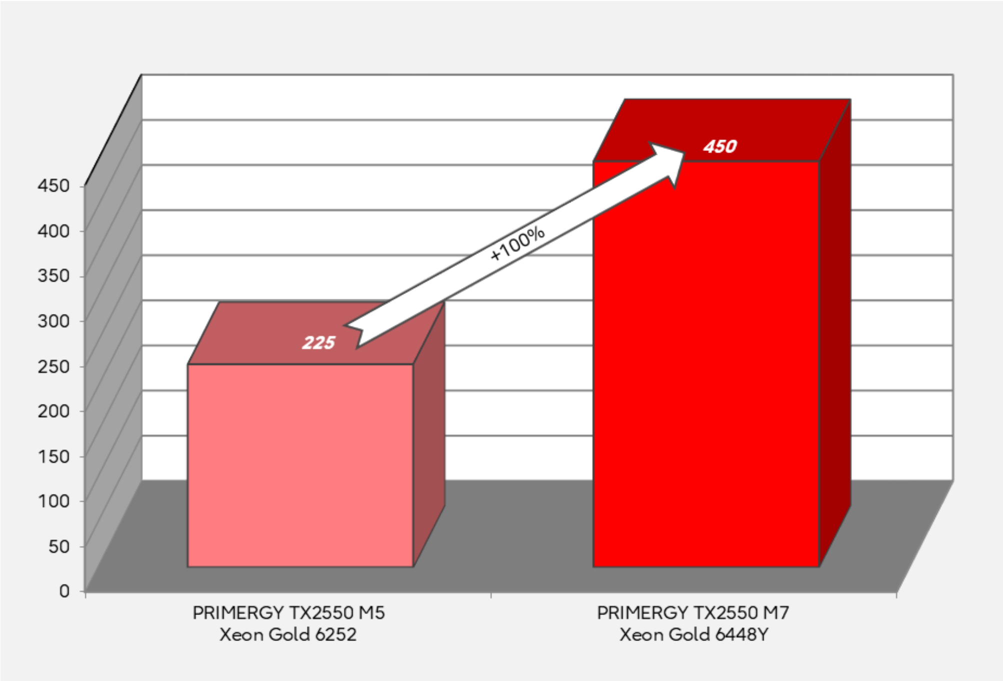

STREAM: Comparison of PRIMERGY TX2550 M5 and PRIMERGY TX2550 M7

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.

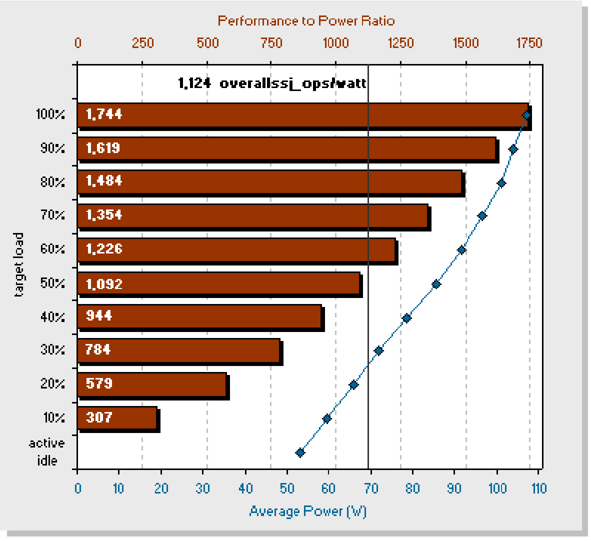

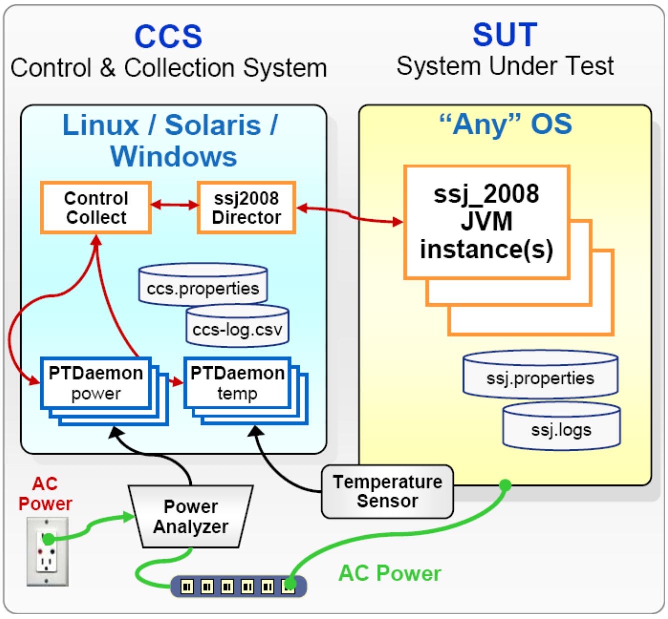

The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System 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

• Model	PRIMERGY TX2550 M7
• Processor	2 x Xeon Gold 6438Y+
• Memory	16 x 32GB (1x32GB) 2Rx8 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
• Power Supply Unit	2 x 900W titanium PSU

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 Serial Port = Disabled 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 90 POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFDECTHRESHOLD 80 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 TX2550 M7 in Microsoft Windows Server 2022 Standard achieved the following result:

SPECpower_ssj2008 = 13,055 overall ssj_ops/watt

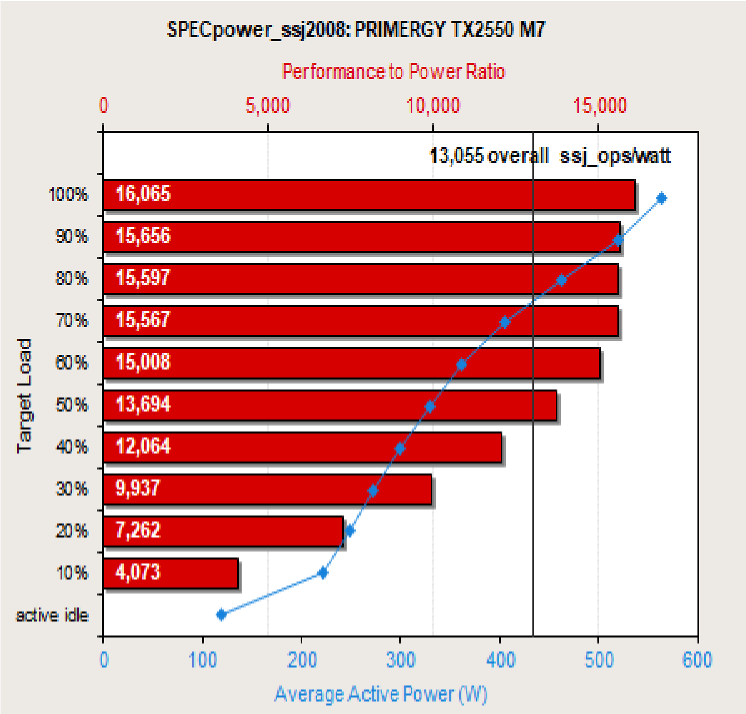

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 13,055 overall ssj_ops/watt for the PRIMERGY TX2550 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%	9,040,604	563	16,065
90%	8,142,124	520	15,656
80%	7,221,043	463	15,597
70%	6,326,586	406	15,567
60%	5,426,798	362	15,008
50%	4,510,954	329	13,694
40%	3,615,218	300	12,064
30%	2,716,354	273	9,937
20%	1,805,349	249	7,262
10%	905,374	222	4,073
Active Idle	0	120	0
			Σ ssj_ops / Σ power = 13,055

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 TX2550 M7 compared to the predecessor PRIMERGY TX2550 M5.

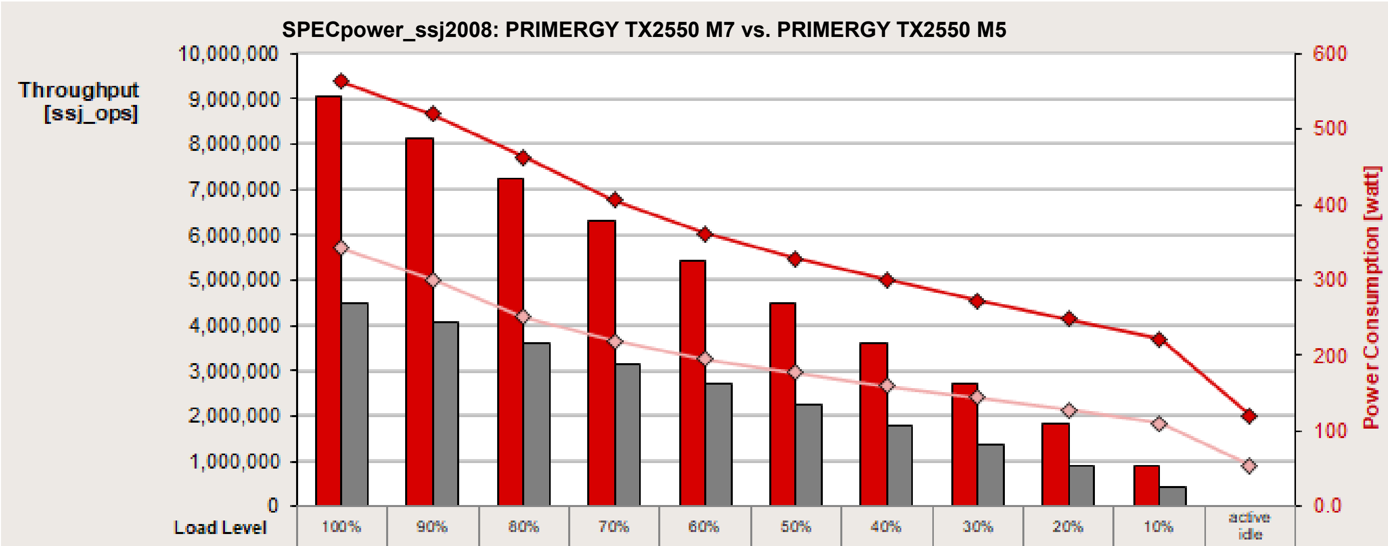

Load Level	100%	90%	80%	70%	60%	50%	40%	30%	20%	10%	active idle
TX2550 M7 ssj_ops	9,040,604	8,142,124	7,221,043	6,326,586	5,426,798	4,510,954	3,615,218	2,716,354	1,805,349	905,374	0
TX2550 M5 ssj_ops	4,507,558	4,057,819	3,600,840	3,157,941	2,701,446	2,255,622	1,801,241	1,353,108	902,810	451,236	0
TX2550 M7 watt	563	520	463	406	362	329	300	273	249	222	120
TX2550 M5 watt	342	301	251	219	195	177	160	145	128	110	55

Thanks to the 4th Generation Intel Xeon Scalable Processors Family, the PRIMERGY TX2550 M7 has a higher throughput. This results in an overall 9.28% increase in energy efficiency in the PRIMERGY TX2550 M7.

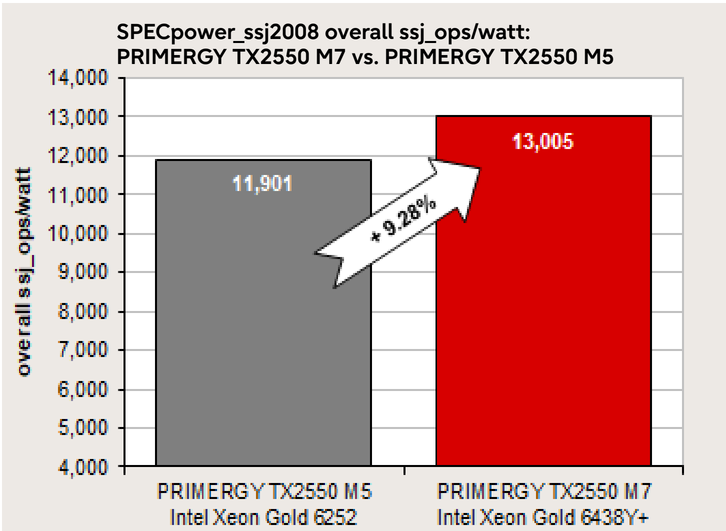

Disk I/O: Performance of storage media

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

3.5 inch model

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) [512e]	ST12000NM004J ST18000NM004J
	NL-SAS HDD (SAS 12Gbps, 7.2k rpm) [512n]	ST2000NM003A ST4000NM003A
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512e]	ST12000NM000J ST18000NM000J
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512n]	ST2000NM000A ST4000NM000A
SSD	SATA SSD (SATA 6Gbps, Mixed Use)	MTFDDAK480TDT MTFDDAK960TDT MTFDDAK1T9TDT MTFDDAK3T8TDT
		MZ7L3480HBLT MZ7L3960HBLT MZ7L31T9HBNA MZ7L33T8HBNA
	SATA SSD (SATA 6Gbps, Read Intensive)	MTFDDAK240TDS MTFDDAK480TDS MTFDDAK960TDS MTFDDAK1T9TDS MTFDDAK3T8TDS MTFDDAK7T6TDS
		MZ7L3240HCHQ MZ7L3480HCHQ MZ7L3960HCJR MZ7L31T9HBLT
		MZ7L33T8HBLT MZ7L37T6HBLA

Controller:	Intel C741 Standard SATA AHCI controller	
Storage media	Category	Drive name
SSD	M.2 Flash module	MTFDDAK240TDS MTFDDAV480TDS

Controller:	Intel C741 Standard NVM Express controller	
Storage media	Category	Drive name
SSD	M.2 Flash module (NVMe)	MTFDKBA480TFR MTFDKBA960TFR

2.5 inch model

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, 7.2krpm)[512n]	ST1000NX0423 ST2000NX0403

Controller:	PRAID EP540i	
Storage media	Category	Drive name
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
	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	MTFDDAK240TDS 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 x16	0, 1, 1E, 10, 5, 50
PCIe SSD 2.5"	PRAID EP680i	-	PCIe 4.0 x8	SATA 6G SAS 12G PCIe 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
3.5 inch model	HDD	SAS 12G	3.5 inch or 2.5 inch1)
		SATA 6G	3.5 inch
	SSD	SATA 6G	2.5 inch1) or M.2
		PCIe 4.0	M.2
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

1) It is available with a 3.5 inch cage.

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

3.5 inch model

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]							
12,000	ST12000NM004J	SAS 12G	506	577	525	266	265
18,000	ST18000NM004J	SAS 12G	500	579	515	265	262
☐ NL-SAS HDD 7.2krpm [512n]							
2,000	ST2000NM003A	SAS 12G	378	343	336	237	237
4,000	ST2000NM004A	SAS 12G	369	333	330	214	215
☐ BC-SATA HDD 7.2krpm [512e]							
12,000	ST12000NM000J	SATA 6G	523	497	463	263	263
18,000	ST18000NM000J	SATA 6G	517	518	484	271	269
☐ BC-SATA HDD 7.2krpm [512n]							
2,000	ST2000NM000A	SATA 6G	331	304	313	230	207
4,000	ST4000NM000A	SATA 6G	313	290	297	211	210

SSDs

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	filecopy	Streaming	Restore
☐ 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
☐ 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
☐ 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

2.5 inch model

HDDs

Capacity [GB]	Storage device	Interface	Transactions [IO/s] Database	Transactions [IO/s] Fileserver	Transactions [IO/s] filecopy	Throughput [MiB/s] Streaming	Throughput [MiB/s] 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] Database	Transactions [IO/s] Fileserver	Transactions [IO/s] filecopy	Throughput [MiB/s] Streaming	Throughput [MiB/s] 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 (WI)							
400	SSDPF21Q400GB	PCIe4 x4	303,783	91,576	84,727	6,693	4,562
800	SSDPF21Q800GB	PCIe4 x4	290,266	99,852	94,882	6,738	4,512
1,600	SSDPF21Q016TB	PCIe4 x4	304,687	108,995	110,292	6,682	4,382
- 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:

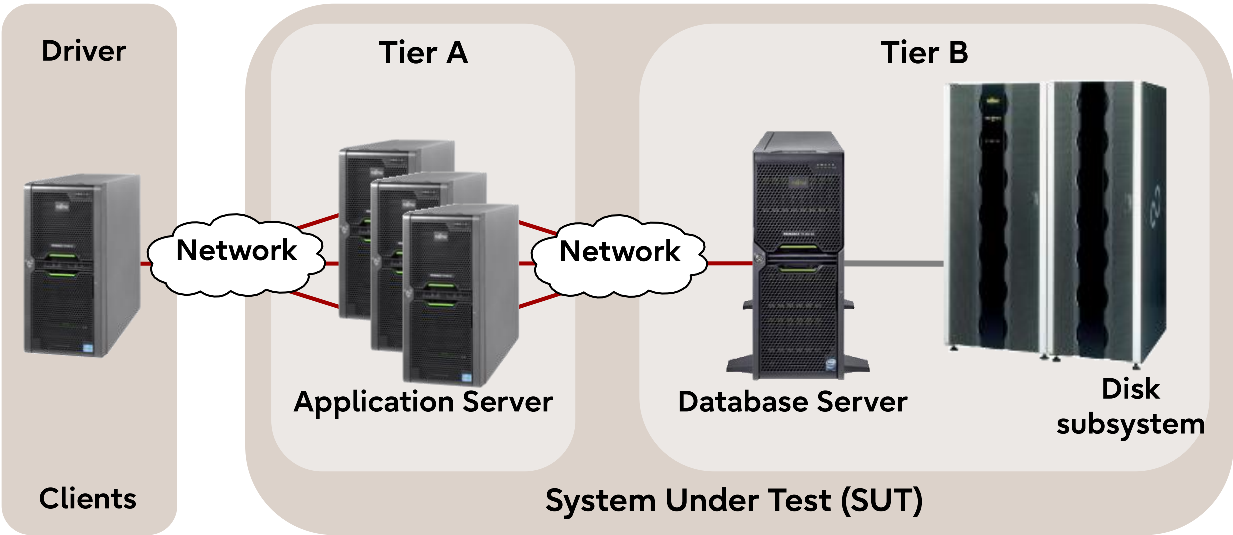

All OLTP-2 results were calculated based on the configuration of the next following pages of PRIMERGY RX2540 M7.

Database Server (Tier B)	
Hardware	
• Model	PRIMERGY RX2540 M7
• Processor	4th Generation Intel Xeon Processor Scalable Family
• Memory	2 processor: 32 x 64 GB (1x64 GB) 2Rx4 DDR5-4800 ECC
• Network interface	1 x Dual port LAN 10 Gbps 1 x Quad port OCPv3 LAN 1 Gbps
• 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
• Memory	192 GB, 2666 MHz Registered ECC DDR4
• Network interface	1 x Dual port LAN 10 Gbps 1 x Dual port onboard LAN 1 Gbps
• Disk subsystem	2 x 300 GB 10k rpm SAS drive
Software	
• Operating system	Microsoft Windows Server 2016 Standard

Client	
Hardware	
• Model	1 x PRIMERGY RX2530 M2
• Processor	2 x Xeon E5-2667 v4
• Memory	128 GB, 2400 MHz Registered ECC DDR4
• Network interface	1 x Quad port onboard LAN 1 Gbps
• Disk subsystem	1 x 300 GB 10k rpm SAS drive
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 measurements 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 Platinum 8450H	28	56	6,355	est.
Xeon Platinum 8444H	16	32	4,536	est.
Xeon Gold 6454S	32	64	7,533	est.
Xeon Gold 6448Y	32	64	7,361	est.
Xeon Gold 6442Y	24	48	6,579	est.
Xeon Gold 6438Y+	32	64	7,086	est.
Xeon Gold 6434	8	16	2,766	est.
Xeon Gold 6430	32	64	7,186	est.
Xeon Gold 6426Y	16	32	4,361	est.
Xeon Gold 5420+	28	56	6,505	est.
Xeon Gold 5418Y	24	48	5,745	est.
Xeon Gold 5416S	16	32	3,705	est.
Xeon Gold 5415+	8	16	2,467	est.
Xeon Silver 4416+	20	40	4,822	est.
Xeon Silver 4410Y	12	24	3,041	est.
Xeon Silver 4410T	10	20	2,903	est.

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

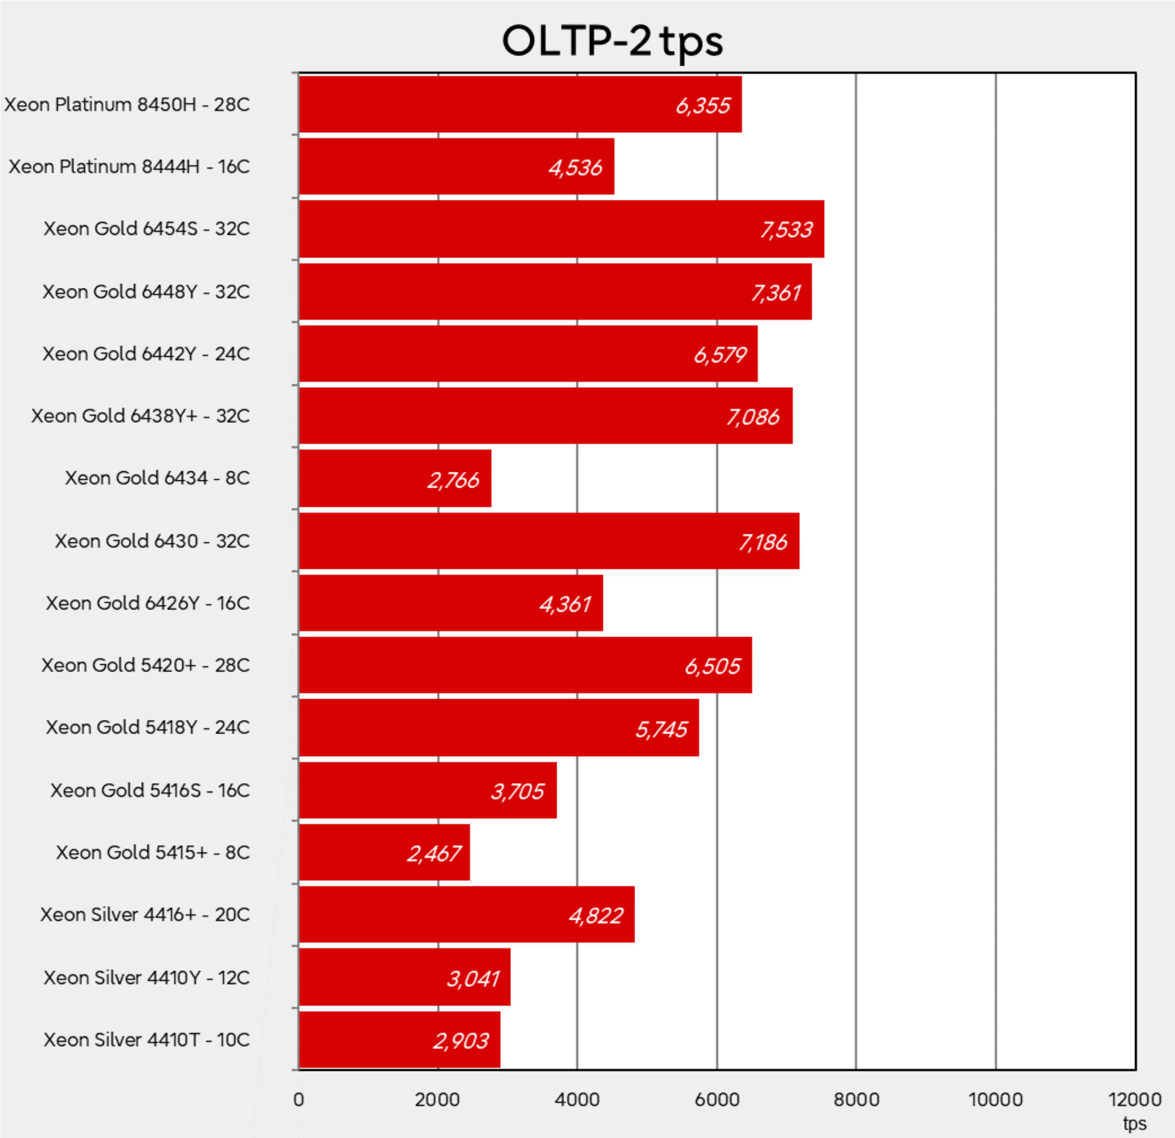

The features of the processors are summarized in the section "Technical data."

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.

Literature

PRIMERGY Servers

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

PRIMERGY TX2550 M7

This Whitepaper

https://docs.ts.fujitsu.com/dl.aspx?id=6dc6f483-c50e-407d-ab7f-3fe75062f440

https://docs.ts.fujitsu.com/dl.aspx?id=bedd549c-e005-4966-ae76-0539b82c8293

Data sheet

TX2550 M7: https://docs.ts.fujitsu.com/dl.aspx?id=28d1e428-c1ff-4d2e-b477-e7cfe4da59eb

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

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

STREAM

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

Document change history

Version	Date	Description
1.0	2023-05-30	New: • Technical data • SPEC CPU2017, STREAM Measured and calculated with 4th Generation Intel Xeon Processor Scalable Family • SPECpower_ssj2008 Measured with Intel Xeon Platinum 6438Y+ • Disk I/O Measured with 2.5 / 3.5 inch model • OLTP2 Measured and 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