# Fujitsu Server PRIMERGY
## Performance Report
## PRIMERGY RX2530 M7 / RX2540 M7

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMERGY RX2530 M7 / RX2540 M7.

Explaines PRIMERGY RX2530 M7 / RX2540 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 |
| --- |
| 1.3<br>2024-01-12 |

# Contents

# Technical data

## PRIMERGY RX2530 M7



## PRIMERGY RX2540 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 RX2530 M7 | PRIMERGY RX2540 M7 |
|---|---|---|
| Form factor | 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 | 32 (16 per processor) | |
| Maximum memory configuration | 8,192 GB | |
| Maximum number of internal storage disks | 3.5 inch: 4<br>2.5 inch: 10 | 3.5 inch: 10<br>2.5 inch: 30 (Front 24, Rear 6) |
| Maximum number of PCI slots | PCI-Express 4.0 (x16 lane): 1 (Low Profile, dedicated for the modular storage controller)<br>PCI-Express 5.0 (x16 lane): 3 (Low Profile) | PCI Express 4.0 (x16 lane): 1 (Low Profile, dedicated for the modular storage controller)<br>PCI Express 5.0 (x16 lane): 4 (Low Profile)<br>PCI Express 5.0 (x8 lane):   2 (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**

| 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] |
|---|---|---|---|---|---|---|---|---|---|
| Xeon Max 9468 | HBM | 48 | 96 | 105 | 16 | 2.1 | 3.5 | 4,800 | 350 |
| Xeon Max 9462 | HBM | 32 | 64 | 75 | 16 | 2.7 | 3.5 | 4,800 | 350 |
| Xeon Max 9460 | HBM | 40 | 80 | 97.5 | 16 | 2.2 | 3.5 | 4,800 | 350 |
| Xeon Platinum 8490H | XCC | 60 | 120 | 112.5 | 16 | 1.9 | 3.5 | 4,800 | 350 |
| Xeon Platinum 8480+ | XCC | 56 | 112 | 105 | 16 | 2.0 | 3.8 | 4,800 | 350 |
| Xeon Platinum 8470N | XCC | 52 | 104 | 97.5 | 16 | 1.7 | 3.6 | 4,800 | 300 |
| Xeon Platinum 8470 | XCC | 52 | 104 | 105 | 16 | 2.0 | 3.8 | 4,800 | 350 |
| Xeon Platinum 8468V | XCC | 48 | 96 | 97.5 | 16 | 2.4 | 3.8 | 4,800 | 330 |
| Xeon Platinum 8468 | XCC | 48 | 96 | 105 | 16 | 2.1 | 3.8 | 4,800 | 350 |
| Xeon Platinum 8462Y+ | MCC | 32 | 64 | 60 | 16 | 2.8 | 4.1 | 4,800 | 300 |
| Xeon Platinum 8460Y+ | XCC | 40 | 80 | 105 | 16 | 2.0 | 3.7 | 4,800 | 300 |
| Xeon Platinum 8458P | XCC | 44 | 88 | 82.5 | 16 | 2.7 | 3.8 | 4,800 | 350 |
| Xeon Platinum 8452Y | XCC | 36 | 72 | 67.5 | 16 | 2.0 | 3.2 | 4,800 | 300 |
| 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 6444Y | MCC | 16 | 32 | 45 | 16 | 3.6 | 4.1 | 4,800 | 270 |
| 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 6438N | MCC | 32 | 64 | 60 | 16 | 2.0 | 3.6 | 4,800 | 205 |
| Xeon Gold 6438M | MCC | 32 | 64 | 60 | 16 | 2.2 | 3.9 | 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 6428N | MCC | 32 | 64 | 60 | 16 | 1.8 | 3.8 | 4,000 | 185 |
| 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 5418N | MCC | 24 | 48 | 45 | 16 | 1.8 | 3.8 | 4,000 | 165 |
| Xeon Gold 5416S | MCC | 16 | 32 | 30 | 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 |

| 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 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 | 8 | 22.5 | - | 1.8 | 1.9 | 4,000 | 125 |

All processors that can be ordered with PRIMERGY RX2530 M7 / RX2540 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** <br> Data analytics and big data usages |
| M | **Media Transcode** <br> Media, AI, and HPC workloads |
| N | **Networking** <br> Network and 5G workload environments from edge to the data center |
| P | **Cloud IaaS** <br> VM environments which require higher frequency |
| Q | **Liquid Cooled** <br> Environments that require higher core count and higher frequency such as HPC |
| S | **Storage & HCI** <br> Storage provider and HCI |
| T | **Long-life Use (IOT)** <br> High reliability and long-life availability usage |
| U | **1-Socket** <br> Edge server, router, storage and security appliances composed of cost effective 1 socket configuration |
| V | **Cloud SaaS** <br> VM environments which require power efficiency, higher frequency, and higher core counts |
| Y | **IaaS, networking, virtualized environments** <br> 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 |
| DC PSU | 1,300W PSU DC | 2 |
| | 1,600W PSU HVDC | 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 RX2530 M7 / RX2540 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 RX2530 M7 / RX2540 M7 |
| • Processor | 2 x 4th Generation Intel Xeon Scalable Processors Family or<br>1 x 4th Generation Intel Xeon Scalable Processors Family |
| • Memory | 16 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (2CPU configuration) or<br> 8 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (1CPU configuration) |
| **Software** | |
| • BIOS settings | SPECspeed2017_int_base:<br>• RdCur for XPT Prefetch = Enable<br>• Adjacent Cache Line Prefetch = Disabled<br>• Package C State limit = C0<br>• SNC(Sub NUMA) = Enable SNC2 (Disabled when MCC are installed)<br>• HWPM Support = Disabled<br>• AVX P1 = Level2<br>• CPU Performance Boost = Aggressive<br>• FAN Control = Full<br><br>SPECSpeed2017_fp_base:<br>• Hyper Threading = Disabled<br>• DCU IP Prefetcher = Disabled<br>• Package C State limit = C0<br>• LLC Prefetch = Enabled<br>• DBP-F = Enabled<br>• CPU Performance Boost = Aggressive<br>• FAN Control = Full<br><br>SPECrate2017_int_base:<br>• DCU Streamer Prefetcher = Disabled<br>• Package C State limit = C0<br>• CPU Performance Boost = Aggressive<br>• SNC(Sub NUMA) =Enable SNC4<br>• FAN Control = Full<br><br>SPECrate2017_fp_base:<br>• Hyper Threading = Disabled (Enabled when MCC are installed)<br>• Package C State limit = C0<br>• CPU Performance Boost = Aggressive<br>• SNC (Sub NUMA) =Enable SNC4 (Enable SNC2 when MCC are installed)<br>• 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<br>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 | |
|---|---|---|---|---|---|---|
| | | | RX2530 M7 | RX2540 M7 | RX2530 M7 | RX2540 M7 |
| **2CPU configuration** | | | | | | |
| Xeon Max 9468 | 48 | 2 | 756 est. | **773** | 1,020 est. | **1,020** |
| Xeon Max 9462 | 32 | 2 | 591 est. | **604** | 853 est. | **852** |
| Xeon Max 9460 | 40 | 2 | 664 est. | **679** | 940 est. | **938** |
| Xeon Platinum 8490H | 60 | 2 | **943** | **964** | **976** | **974** |
| Xeon Platinum 8480+ | 56 | 2 | **923** | **934** | **956** | **953** |
| Xeon Platinum 8470N | 52 | 2 | 781 est. | **807** | 843 est. | **857** |
| Xeon Platinum 8470 | 52 | 2 | 865 est. | **884** | 924 est. | **923** |
| Xeon Platinum 8468V | 48 | 2 | 792 est. | **809** | 878 est. | **877** |
| Xeon Platinum 8468 | 48 | 2 | 826 est. | **844** | 903 est. | **901** |
| Xeon Platinum 8462Y+ | 32 | 2 | 649 est. | **663** | 766 est. | **765** |
| Xeon Platinum 8460Y+ | 40 | 2 | 675 est. | **690** | 801 est. | **800** |
| Xeon Platinum 8458P | 44 | 2 | 786 est. | **803** | 869 est. | **868** |
| Xeon Platinum 8452Y | 36 | 2 | 609 est. | **623** | 720 est. | **719** |
| Xeon Gold 6454S | 32 | 2 | 541 est. | **553** | 669 est. | **668** |
| Xeon Gold 6448Y | 32 | 2 | 563 est. | **576** | 690 est. | **689** |
| Xeon Gold 6444Y | 16 | 2 | 377 est. | **386** | 521 est. | **520** |
| Xeon Gold 6442Y | 24 | 2 | 480 est. | **490** | 626 est. | **625** |
| Xeon Gold 6438Y+ | 32 | 2 | 535 est. | **547** | 646 est. | **645** |
| Xeon Gold 6438N | 32 | 2 | 532 est. | **544** | 651 est. | **650** |
| Xeon Gold 6438M | 32 | 2 | 543 est. | **555** | 654 est. | **653** |
| Xeon Gold 6434 | 8 | 2 | 194 est. | **198** | 285 est. | **285** |
| Xeon Gold 6430 | 32 | 2 | 515 est. | **526** | 632 est. | **631** |
| Xeon Gold 6428N | 32 | 2 | 488 est. | **499** | 586 est. | **585** |
| Xeon Gold 6426Y | 16 | 2 | 324 est. | **332** | 444 est. | **443** |
| Xeon Gold 5420+ | 28 | 2 | 467 est. | **478** | 590 est. | **589** |
| Xeon Gold 5418Y | 24 | 2 | 409 est. | **418** | 520 est. | **520** |
| Xeon Gold 5418N | 24 | 2 | 386 est. | **395** | 479 est. | **478** |
| Xeon Gold 5416S | 16 | 2 | 275 est. | **282** | 369 est. | **368** |
| Xeon Gold 5415+ | 8 | 2 | 174 est. | **178** | 253 est. | **252** |

| Processor model | Number of cores | Number of processors | SPECrate2017_int_base | | SPECrate2017_fp_base | |
|---|---|---|---|---|---|---|
| | | | RX2530 M7 | RX2540 M7 | RX2530 M7 | RX2540 M7 |
| **2CPU configuration(cont.)** | | | | | | |
| Xeon Silver 4416+ | 20 | 2 | 355 est. | **363** | 454 est. | **453** |
| Xeon Silver 4410Y | 12 | 2 | 213 est. | **218** | 320 est. | **320** |
| Xeon Silver 4410T | 10 | 2 | 207 est. | **211** | 292 est. | **291** |
| **1CPU configuration** | | | | | | |
| Xeon Gold 6414U | 32 | 1 | 261 est. | **267** | 335 est. | **334** |
| Xeon Gold 5412U | 24 | 1 | 216 est. | **221** | 282 est. | **282** |
| Xeon Bronze 3408U | 8 | 1 | 42.0 est. | **42.9** | 75.2 est. | **75.1** |

| Processor model | Number of cores | Number of processors | SPECspeed2017_int_base | | SPECspeed2017_fp_base | |
|---|---|---|---|---|---|---|
| | | | RX2530 M7 | RX2540 M7 | RX2530 M7 | RX2540 M7 |
| Xeon Platinum 8490H | 60 | 2 | - | - | **354** | **355** |
| Xeon Platinum 8462Y+ | 32 | 2 | **16.0** | **15.9** | - | - |

The following graphs compare the throughputs of PRIMERGY RX2530 M7 / RX2540 M7 and their older models, PRIMERGY RX2530 M6 / RX2540 M6, with maximum performance configurations. Both models showed significant performance improvements over the previous generation.
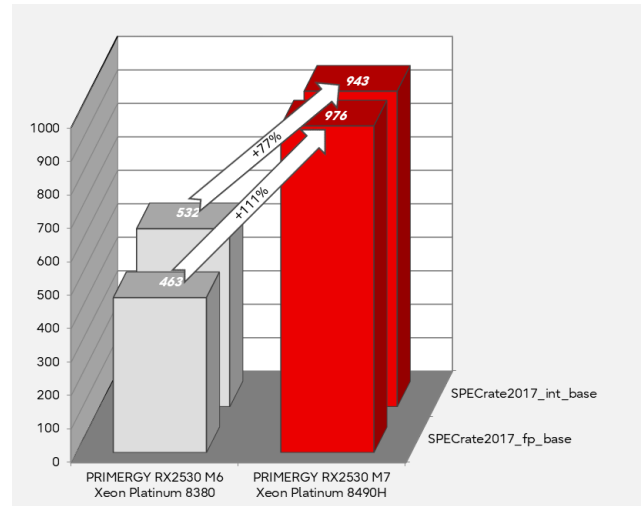


**SPECrate2017**: Comparison of PRIMERGY RX2530 M6 and PRIMERGY RX2530 M7
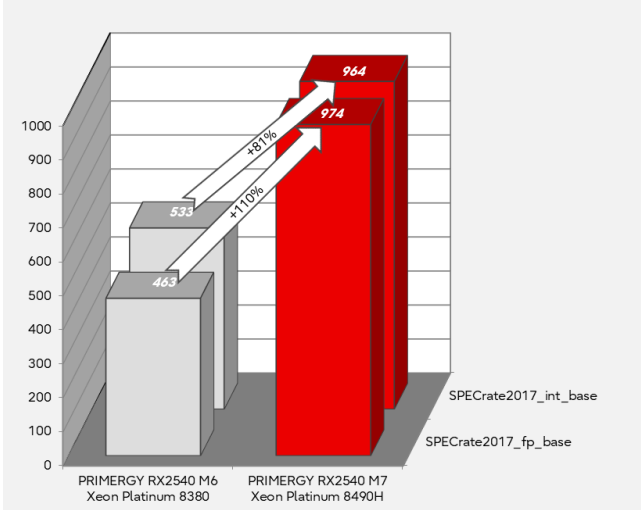


**SPECrate2017**: Comparison of PRIMERGY RX2540 M6 and PRIMERGY RX2540 M7

# 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 | $a(i) = q \times b(i)$ | 16 | 1 |
| SUM | $a(i) = b(i) + c(i)$ | 24 | 1 |
| TRIAD | $a(i) = b(i) + q \times c(i)$ | 24 | 2 |

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

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

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

## *Benchmark environment*

| System Under Test (SUT) |
|---|

### Hardware

| | |
|---|---|
| • Model | PRIMERGY RX2530 M7 / RX2540 M7 |
| • Processor | 2 x 4th Generation Intel Xeon Scalable Processors Family or<br>1 x 4th Generation Intel Xeon Scalable Processors Family |
| • Memory | 16 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (2CPU configuration) or<br> 8 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (1CPU configuration) |

### Software

| | |
|---|---|
| • BIOS settings | • DCU Streamer Prefetcher = Disabled<br>• SNC(Sub NUMA) = Enable SNC4 (Enable SNC2 when MCC type installed))<br>• Intel Virtualization Technology = Disabled<br>• LLC Dead Line Alloc   = Disabled<br>• 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] RX2530 M7 | RX2540 M7 |
|---|---|---|---|---|---|---|---|
| **2CPU configuration** | | | | | | | |
| Xeon Max 9468 | 4,400 | 282 | 48 | 2.1 | 2 | 514 est. | **511** |
| Xeon Max 9462 | 4,400 | 282 | 30 | 2.7 | 2 | 491 est. | **488** |
| Xeon Max 9460 | 4,400 | 282 | 42 | 2.2 | 2 | 514 est. | **512** |
| Xeon Platinum 8490H | 4,800 | 307 | 60 | 1.9 | 2 | **525** | **522** |
| Xeon Platinum 8480+ | 4,800 | 307 | 56 | 2.0 | 2 | 524 est. | **521** |
| Xeon Platinum 8470N | 4,800 | 307 | 52 | 1.7 | 2 | 514 est. | **511** |
| Xeon Platinum 8470 | 4,800 | 307 | 52 | 2.0 | 2 | 513 est. | **511** |
| Xeon Platinum 8468V | 4,800 | 307 | 48 | 2.4 | 2 | 508 est. | **505** |
| Xeon Platinum 8468 | 4,800 | 307 | 48 | 2.1 | 2 | 490 est. | **488** |
| Xeon Platinum 8462Y+ | 4,800 | 307 | 32 | 2.8 | 2 | 477 est. | **474** |
| Xeon Platinum 8460Y+ | 4,800 | 307 | 40 | 2.0 | 2 | 480 est. | **478** |
| Xeon Platinum 8458P | 4,800 | 307 | 44 | 2.7 | 2 | 500 est. | **498** |
| Xeon Platinum 8452Y | 4,800 | 307 | 36 | 2.0 | 2 | 455 est. | **452** |
| Xeon Gold 6454S | 4,800 | 307 | 32 | 2.2 | 2 | 447 est. | **444** |
| Xeon Gold 6448Y | 4,800 | 307 | 32 | 2.1 | 2 | 469 est. | **467** |
| Xeon Gold 6444Y | 4,800 | 307 | 16 | 3.6 | 2 | 385 est. | **383** |
| Xeon Gold 6442Y | 4,800 | 307 | 24 | 2.6 | 2 | 443 est. | **441** |
| Xeon Gold 6438Y+ | 4,800 | 307 | 32 | 2.0 | 2 | 465 est. | **463** |
| Xeon Gold 6438N | 4,800 | 307 | 32 | 2.0 | 2 | 467 est. | **464** |
| Xeon Gold 6438M | 4,800 | 307 | 32 | 2.2 | 2 | 466 est. | **464** |
| Xeon Gold 6434 | 4,800 | 307 | 8 | 3.7 | 2 | 228 est. | **227** |
| Xeon Gold 6430 | 4,400 | 282 | 32 | 2.1 | 2 | 421 est. | **419** |
| Xeon Gold 6428N | 4,000 | 256 | 32 | 1.8 | 2 | 410 est. | **407** |
| Xeon Gold 6426Y | 4,800 | 307 | 16 | 2.5 | 2 | 350 est. | **348** |
| Xeon Gold 5420+ | 4,400 | 282 | 28 | 2.0 | 2 | 420 est. | **418** |
| Xeon Gold 5418Y | 4,400 | 282 | 24 | 2.0 | 2 | 388 est. | **386** |
| Xeon Gold 5418N | 4,000 | 256 | 24 | 1.8 | 2 | 363 est. | **362** |
| Xeon Gold 5416S | 4,400 | 282 | 16 | 2.0 | 2 | 285 est. | **284** |
| Xeon Gold 5415+ | 4,400 | 282 | 8 | 2.9 | 2 | 215 est. | **214** |
| Xeon Silver 4416+ | 4,000 | 256 | 20 | 2.0 | 2 | 331 est. | **330** |
| Xeon Silver 4410Y | 4,000 | 256 | 12 | 2.0 | 2 | 265 est. | **264** |
| Xeon Silver 4410T | 4,000 | 256 | 10 | 2.7 | 2 | 240 est. | **239** |

| Processor | Memory frequency [MHz] | Maximum memory bandwidth [GB/s] | Number of cores | Rated frequency [GHz] | Number of processors | TRIAD [GB/s] | |
|---|---|---|---|---|---|---|---|
| | | | | | | RX2530 M7 | RX2540 M7 |
| **1CPU configuration** | | | | | | | |
| Xeon Gold 6414U | 4,800 | 307 | 32 | 2.0 | 1 | 240 est. | **239** |
| Xeon Gold 5412U | 4,400 | 282 | 24 | 2.1 | 1 | 210 est. | **209** |
| Xeon Bronze 3408U | 4,000 | 256 | 8 | 1.8 | 1 | 124 est. | **123** |

The following graphs compare the throughputs of PRIMERGY RX2530 M7 / RX2540 M7 and their older models, PRIMERGY RX2530 M6 / RX2540 M6, with maximum performance configurations.

Both models showed significant performance improvements over the previous generation.
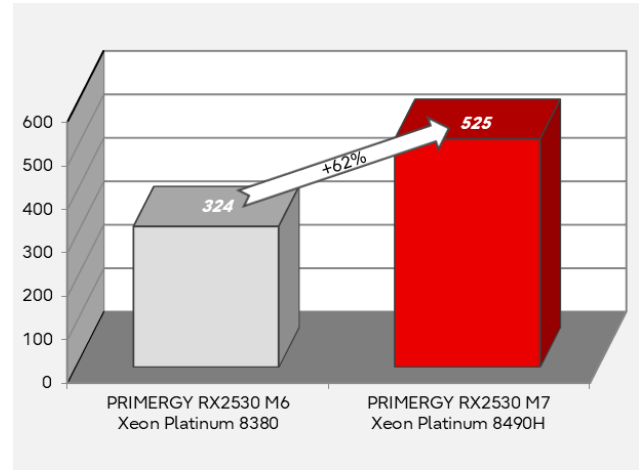


**STREAM**: Comparison of PRIMERGY RX2530 M6 and PRIMERGY RX2530 M7
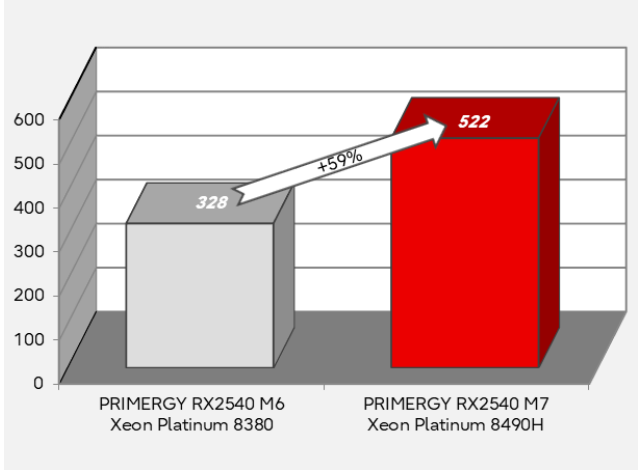


**STREAM**: Comparison of PRIMERGY RX2540 M6 and PRIMERGY RX2540 M7

# 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 https://www.top500.org/. This requires using an HPL-based LINPACK version (see https://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://software.intel.com/en-us/articles/intel-math-kernel-library-linpack-download/.

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

## *Benchmark environment*

| System Under Test (SUT) |
|---|

### Hardware

| | |
|---|---|
| • Model | PRIMERGY RX2530 M7 / RX2540 M7 |
| • Processor | 2 x 4th Generation Intel Xeon Scalable Processors Family or<br>1 x 4th Generation Intel Xeon Scalable Processors Family |
| • Memory | 16 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (2CPU configuration) or<br> 8 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC (1CPU configuration) |

### Software

| | |
|---|---|
| • BIOS settings | • HyperThreading = Disabled<br>• CPU Performance Boost = Agressive<br>• Fan Control = Full |
| • 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<br>(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] | RX2530 M7 | | RX2540 M7 | |
|---|---|---|---|---|---|---|---|---|
| | | | | | Rmax [GFlops] | Effic. | Rmax [GFlops] | Effic. |
| **2CPU configuration** | | | | | | | | |
| Xeon Max 9468 | 48 | 2.1 | 2 | 6,451 | 5,950 est. | 92% | **6,037** | 94% |
| Xeon Max 9462 | 32 | 2.7 | 2 | 5,530 | 4,986 est. | 90% | **5,059** | 91% |
| Xeon Max 9460 | 40 | 2.2 | 2 | 5,632 | 5,479 est. | 97% | **5,559** | 99% |
| Xeon Platinum 8490H | 60 | 1.9 | 2 | 7,296 | **7,279** | 100% | **7,386** | 101% |
| Xeon Platinum 8480+ | 56 | 2.0 | 2 | 7,168 | 7,281 est. | 102% | **7,388** | 103% |
| Xeon Platinum 8470N | 52 | 1.7 | 2 | 5,658 | 6,017 est. | 106% | **6,105** | 108% |
| Xeon Platinum 8470 | 52 | 2.0 | 2 | 6,656 | 6,830 est. | 103% | **6,930** | 104% |
| Xeon Platinum 8468V | 48 | 2.4 | 2 | 7,373 | 6,230 est. | 84% | **6,321** | 86% |
| Xeon Platinum 8468 | 48 | 2.1 | 2 | 6,451 | 6,450 est. | 100% | **6,544** | 101% |
| Xeon Platinum 8462Y+ | 32 | 2.8 | 2 | 5,734 | 5,442 est. | 95% | **5,522** | 96% |
| Xeon Platinum 8460Y+ | 40 | 2.0 | 2 | 5,120 | 5,343 est. | 104% | **5,421** | 106% |
| Xeon Platinum 8458P | 44 | 2.7 | 2 | 7,603 | 6,073 est. | 80% | **6,162** | 81% |
| Xeon Platinum 8452Y | 36 | 2.0 | 2 | 4,608 | 5,100 est. | 111% | **5,175** | 112% |
| Xeon Gold 6454S | 32 | 2.2 | 2 | 4,301 | 4,354 est. | 101% | **4,418** | 103% |
| Xeon Gold 6448Y | 32 | 2.1 | 2 | 4,301 | 4,425 est. | 103% | **4,490** | 104% |
| Xeon Gold 6444Y | 16 | 3.6 | 2 | 3,686 | 3,446 est. | 93% | **3,497** | 95% |
| Xeon Gold 6442Y | 24 | 2.6 | 2 | 3,994 | 3,975 est. | 100% | **4,034** | 101% |
| Xeon Gold 6438Y+ | 32 | 2.0 | 2 | 4,096 | 4,147 est. | 101% | **4,207** | 103% |
| Xeon Gold 6438N | 32 | 2.0 | 2 | 4,096 | 4,249 est. | 104% | **4,311** | 105% |
| Xeon Gold 6438M | 32 | 2.2 | 2 | 4,506 | 4,309 est. | 96% | **4,373** | 97% |
| Xeon Gold 6434 | 8 | 3.7 | 2 | 1,894 | 1,811 est. | 96% | **1,838** | 97% |
| Xeon Gold 6430 | 32 | 2.1 | 2 | 3,891 | 4,257 est. | 109% | **4,320** | 111% |
| Xeon Gold 6428N | 32 | 1.8 | 2 | 3,686 | 3,771 est. | 102% | **3,826** | 104% |
| Xeon Gold 6426Y | 16 | 2.5 | 2 | 2,560 | 2,816 est. | 110% | **2,857** | 112% |
| Xeon Gold 5420+ | 28 | 2.0 | 2 | 3,584 | 3,864 est. | 108% | **3,920** | 109% |
| Xeon Gold 5418Y | 24 | 2.0 | 2 | 3,072 | 3,241 est. | 105% | **3,288** | 107% |
| Xeon Gold 5418N | 24 | 1.8 | 2 | 2,765 | 2,975 est. | 108% | **3,019** | 109% |
| Xeon Gold 5416S | 16 | 2.0 | 2 | 2,048 | 2,194 est. | 107% | **2,226** | 109% |
| Xeon Gold 5415+ | 8 | 2.9 | 2 | 1,485 | 1,495 est. | 101% | **1,517** | 102% |
| Xeon Silver 4416+ | 20 | 2.0 | 2 | 2,560 | 2,880 est. | 113% | **2,923** | 114% |
| Xeon Silver 4410Y | 12 | 2.0 | 2 | 1,536 | 1,852 est. | 121% | **1,879** | 122% |
| Xeon Silver 4410T | 10 | 2.7 | 2 | 1,728 | 1,840 est. | 106% | **1,867** | 108% |

| Processor | Number of cores | Rated frequency [GHz] | Number of processors | Rpeak [GFlops] | RX2530 M7 | | RX2540 M7 | |
|---|---|---|---|---|---|---|---|---|
| | | | | | Rmax [GFlops] | Effic. | Rmax [GFlops] | Effic. |
| **1CPU configuration** | | | | | | | | |
| Xeon Gold 6414U | 32 | 2.0 | 1 | 2,048 | 2,198 est. | 107% | **2,230** | 109% |
| Xeon Gold 5412U | 24 | 2.1 | 1 | 1,613 | 1,844 est. | 114% | **1,871** | 116% |
| Xeon Bronze 3408U | 8 | 1.8 | 1 | 230 | 239 est. | 104% | **242** | 105% |

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 RX2530 M7 / RX2540 M7 and their older models, PRIMERGY RX2530 M6 / RX2540 M6, with maximum performance configurations. Both models showed significant performance improvements over the previous generation.
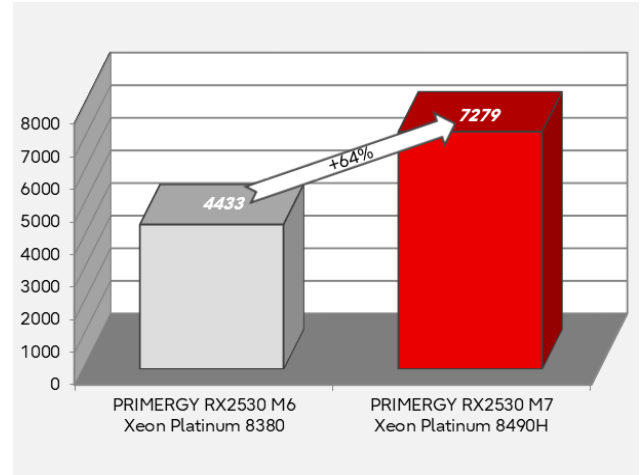


LINPACK: Comparison of PRIMERGY RX2530 M6 and PRIMERGY RX2530 M7
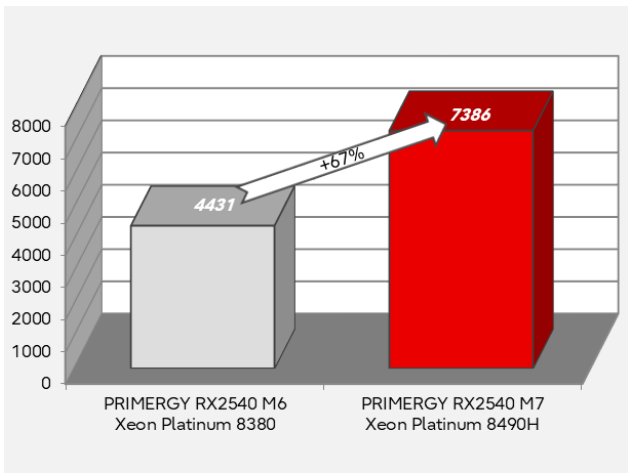


LINPACK: Comparison of PRIMERGY RX2540 M6 and PRIMERGY RX2540 M7

# 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.
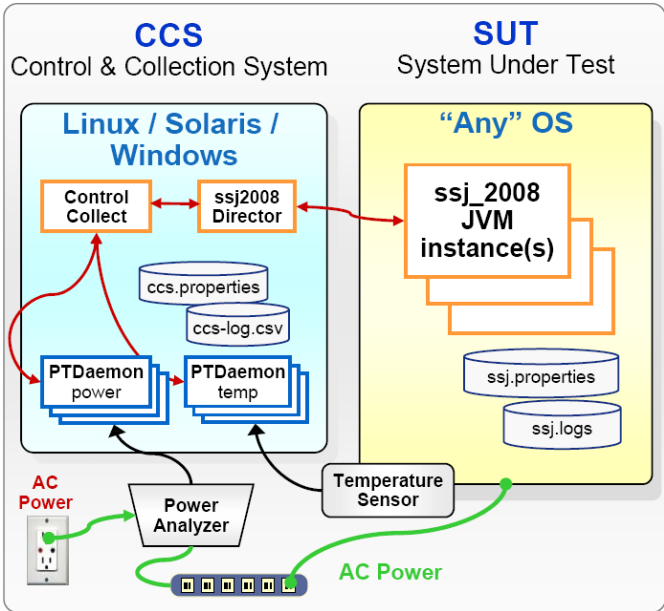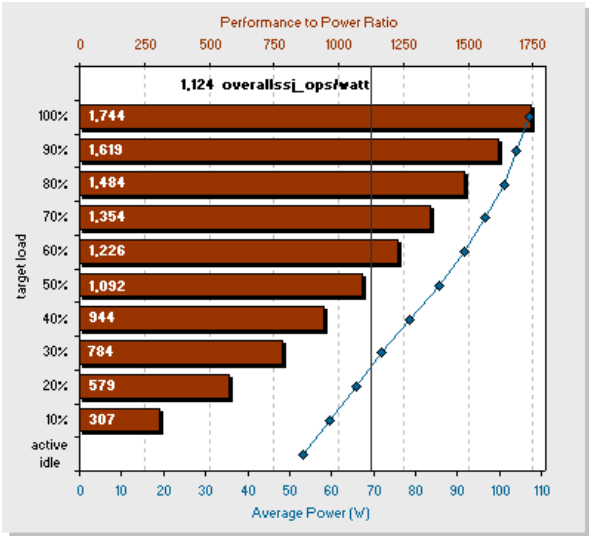


The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System 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 RX2530 M7 / RX2540 M7 |
| • Processor | 2 x Xeon Platinum 8490H |
| • 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<br>Hardware Prefetcher = Disabled<br>Adjacent Cache Line Prefetch = Disabled<br>DCU Streamer Prefetcher = Disabled<br>Intel(R) VT-d = Disabled<br>Package C State limit = No limit<br>Uncore Frequency Scaling = Power balanced<br>CPU Performance Boost = Aggressive<br>SNC(Sub NUMA) = Enable SNC4<br>SATA Controller = Disabled<br>USB Port Control = Disable all ports<br>Serial Port = Disabled<br>Network Stack = Disabled |
| • Operating system | Windows Server 2022 Standard |
| • Operating system settings | Turn off hard disk after = 1 Minute<br>PCI Express Link State Power Management = Maximum power savings<br>Minimum processor state = 0%<br>Maximum processor state = 100%<br>Turn off display after = 1 Minute<br>POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFBOOSTMODE 4<br>POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFINCTHRESHOLD 90<br>POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFDECTHRESHOLD 80<br>POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFDECTIME 1<br>POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR IDLESCALING 1<br>POWERCFG /S SCHEME_CURRENT<br>Using the local security settings console, "lock pages in memory" was enabled for the user running the benchmark.<br>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<br>-XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:InlineSmallCode=3900<br>-XX:MaxInlineSize=270 -XX:MaxTenuringThreshold=15 -XX:ParallelGCThreads=2 -XX:SurvivorRatio=1<br>-XX:TargetSurvivorRatio=99 -XX:-UseAdaptiveSizePolicy -XX:+UseParallelOldGC<br>-XX:FreqInlineSize=2500 -XX:LoopUnrollLimit=45 -XX:InitialTenuringThreshold=12<br>-XX:-ThreadLocalHandshakes -XX:UseAVX=0 |

## *Benchmark results*

### PRIMERGY RX2530 M7

The PRIMERGY RX2530 M7 in Microsoft Windows Server 2022 Standard achieved the following result:

**SPECpower_ssj2008 = 13,704 overall ssj_ops/watt**
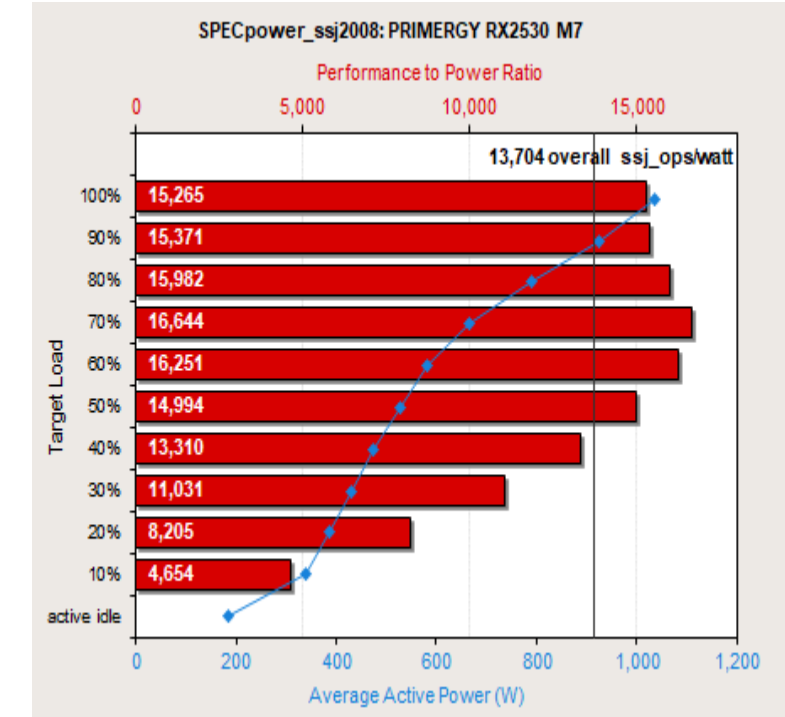


The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 13,704 overall ssj_ops/watt for the PRIMERGY RX2530 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.
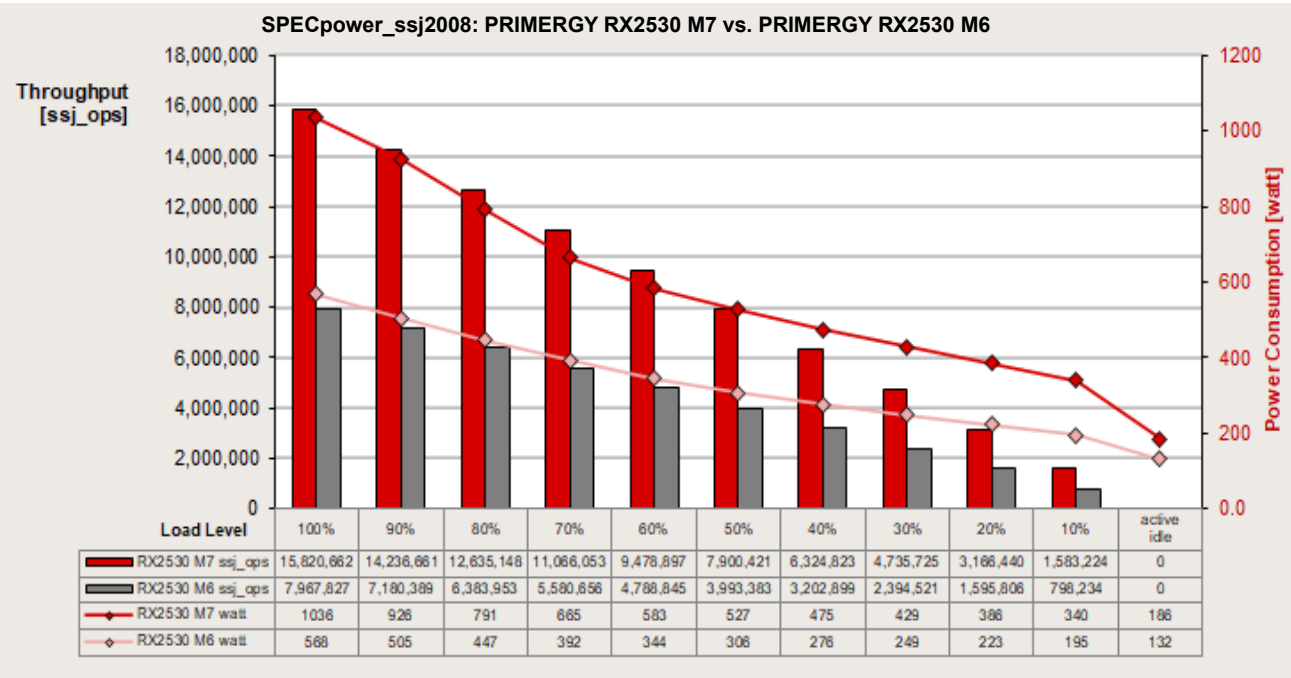
The following table shows the benchmark results for the throughput in ssj_ops, the power consumption in watts and the resulting energy efficiency for each load level.

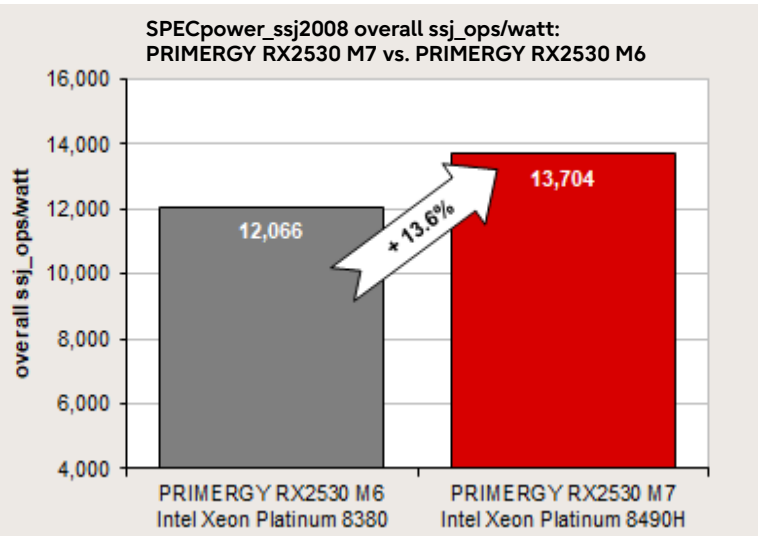| Performance | | Power | Energy Efficiency |
|---|---|---|---|
| Target Load | ssj_ops | Average Power (W) | ssj_ops/watt |
| 100% | 15,820,662 | 1,036 | 15,265 |
| 90% | 14,236,661 | 926 | 15,371 |
| 80% | 12,635,148 | 791 | 15,982 |
| 70% | 11,066,053 | 665 | 16,644 |
| 60% | 9,478,897 | 583 | 16,251 |
| 50% | 7,900,421 | 527 | 14,994 |
| 40% | 6,324,823 | 475 | 13,310 |
| 30% | 4,735,725 | 429 | 11,031 |
| 20% | 3,166,440 | 386 | 8,205 |
| 10% | 1,583,224 | 340 | 4,654 |
| Active Idle | 0 | 186 | 0 |

$\Sigma$ ssj_ops / $\Sigma$ power = 13,704

## Comparison with the predecessor (PRIMERGY RX2530 M7)

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 RX2530 M7 compared to the predecessor PRIMERGY RX2530 M6.



**SPECpower_ssj2008: PRIMERGY RX2530 M7 vs. PRIMERGY RX2530 M6**

| Load Level | 100% | 90% | 80% | 70% | 60% | 50% | 40% | 30% | 20% | 10% | active idle |
|---|---|---|---|---|---|---|---|---|---|---|---|
| RX2530 M7 ssj_ops | 15,820,682 | 14,236,661 | 12,635,148 | 11,066,053 | 9,478,897 | 7,900,421 | 6,324,823 | 4,735,725 | 3,166,440 | 1,583,224 | 0 |
| RX2530 M6 ssj_ops | 7,967,827 | 7,180,389 | 6,383,953 | 5,580,656 | 4,788,845 | 3,993,383 | 3,202,899 | 2,394,521 | 1,595,806 | 798,234 | 0 |
| RX2530 M7 watt | 1036 | 926 | 791 | 665 | 583 | 527 | 475 | 429 | 386 | 340 | 186 |
| RX2530 M6 watt | 568 | 505 | 447 | 392 | 344 | 306 | 276 | 249 | 223 | 195 | 132 |

Thanks to the 4th Generation Intel Xeon Scalable Processors Family, the PRIMERGY RX2530 M7 has a higher throughput. This results in an overall 13.6% increase in energy efficiency in the PRIMERGY RX2530 M7.
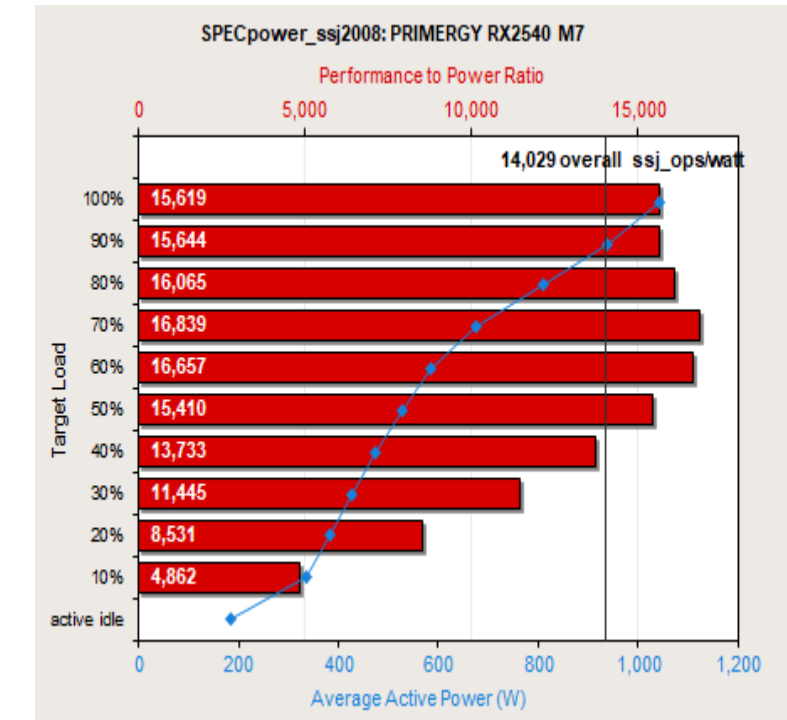


**SPECpower_ssj2008 overall ssj_ops/watt: PRIMERGY RX2530 M7 vs. PRIMERGY RX2530 M6**

PRIMERGY RX2530 M6 Intel Xeon Platinum 8380: 12,066
PRIMERGY RX2530 M7 Intel Xeon Platinum 8490H: 13,704
+ 13.6%

## PRIMERGY RX2540 M7

The PRIMERGY RX2540 M7 in Microsoft Windows Server 2022 Standard achieved the following result:
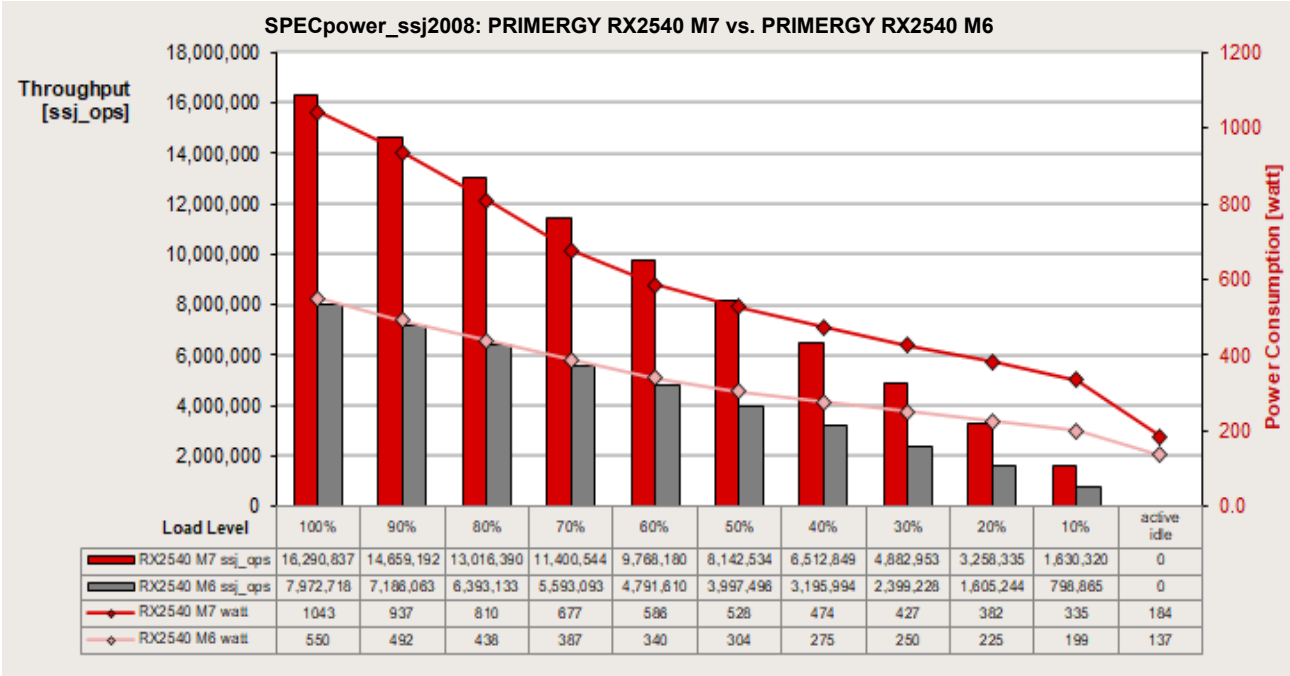
### SPECpower_ssj2008 = 14,029 overall ssj_ops/watt



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 14,029 overall ssj_ops/watt for the PRIMERGY RX2540 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.
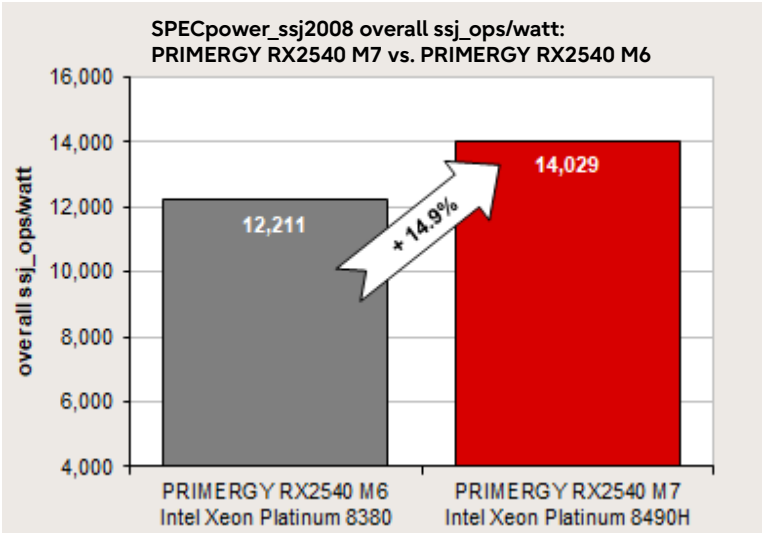
| Performance | | Power | Energy Efficiency |
|---|---|---|---|
| Target Load | ssj_ops | Average Power (W) | ssj_ops/watt |
| 100% | 16,290,837 | 1,043 | 15,619 |
| 90% | 14,659,192 | 937 | 15,644 |
| 80% | 13,016,390 | 810 | 16,065 |
| 70% | 11,400,544 | 677 | 16,839 |
| 60% | 9,768,180 | 586 | 16,657 |
| 50% | 8,142,534 | 528 | 15,410 |
| 40% | 6,512,849 | 474 | 13,733 |
| 30% | 4,882,953 | 427 | 11,445 |
| 20% | 3,258,335 | 382 | 8,531 |
| 10% | 1,630,320 | 335 | 4,862 |
| Active Idle | 0 | 184 | 0 |

$$\Sigma\,ssj\_ops\,/\,\Sigma\,power = 14{,}029$$

## Comparison with the predecessor (PRIMERGY RX2540 M7)

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 RX2540 M7 compared to the predecessor PRIMERGY RX2540 M6.

**SPECpower_ssj2008: PRIMERGY RX2540 M7 vs. PRIMERGY RX2540 M6**

| Load Level | 100% | 90% | 80% | 70% | 60% | 50% | 40% | 30% | 20% | 10% | active idle |
|---|---|---|---|---|---|---|---|---|---|---|---|
| RX2540 M7 ssj_ops | 16,290,837 | 14,659,192 | 13,016,390 | 11,400,544 | 9,768,180 | 8,142,534 | 6,512,849 | 4,882,953 | 3,258,335 | 1,630,320 | 0 |
| RX2540 M6 ssj_ops | 7,972,718 | 7,186,063 | 6,393,133 | 5,593,093 | 4,791,610 | 3,997,496 | 3,195,994 | 2,399,228 | 1,605,244 | 798,865 | 0 |
| RX2540 M7 watt | 1043 | 937 | 810 | 677 | 586 | 528 | 474 | 427 | 382 | 335 | 184 |
| RX2540 M6 watt | 550 | 492 | 438 | 387 | 340 | 304 | 275 | 250 | 225 | 199 | 137 |

Thanks to the 4th Generation Intel Xeon Scalable Processors Family, the PRIMERGY RX2540 M7 has a higher throughput. This results in an overall 14.9% increase in energy efficiency in the PRIMERGY RX2540 M7.

**SPECpower_ssj2008 overall ssj_ops/watt: PRIMERGY RX2540 M7 vs. PRIMERGY RX2540 M6**

| PRIMERGY RX2540 M6 Intel Xeon Platinum 8380 | PRIMERGY RX2540 M7 Intel Xeon Platinum 8490H |
|---|---|
| 12,211 | 14,029 |

+ 14.9%

# SAP Sales and Distribution (SD) Standard Application Benchmark
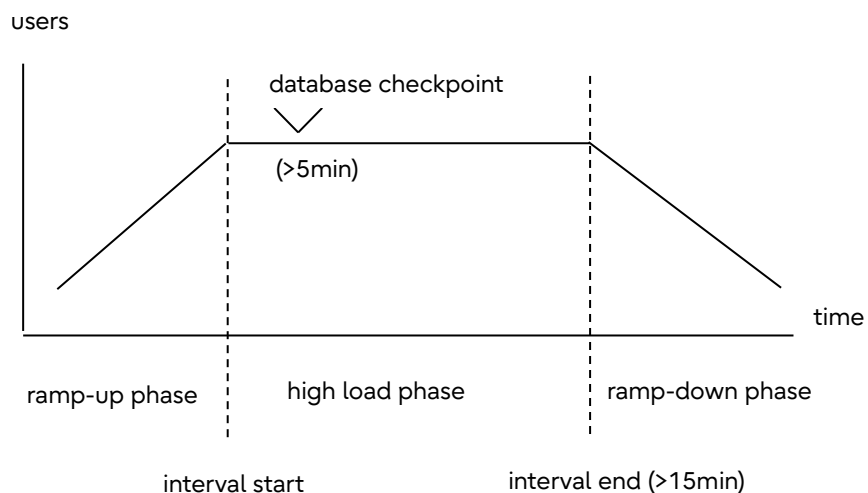
## *Description of the benchmark*

Since 1993 the SAP Standard Application Benchmarks have been developed by SAP in order to verify the performance, stability and scaling of a SAP application system and to provide information for configuring, sizing and for platform comparison. By far the most popular benchmarks from the many available are the SAP SD benchmark and the BW Edition for SAP HANA benchmark (see corresponding section).

The Sales and Distribution benchmark is one of the most CPU consuming benchmarks available and has become a de-facto standard for SAP's platform partners and in the ERP (Enterprise Resource Planning) environment.

During the benchmark a defined sequence of business transactions are run through as shown in the table below. The Sales and Distribution (SD) benchmark covers a sell-from-stock scenario (including a customer order creation, the corresponding delivery with subsequent goods movement and creation of the invoice) and consists of the following SAP transactions:

| |
|---|
| Create an order with five line items (SAP transaction VA01) |
| Create a delivery for this order (SAP transaction VL01N) |
| Display the customer order (SAP transaction VA03) |
| Change the delivery (SAP transaction VL02N) and post goods issue |
| List 40 orders for one sold-to party (SAP transaction VA05) |
| Create an invoice (SAP transaction VF01) |

Each of the simulated users repeats this series of transactions from the start to the end of a benchmark run. The think time between two user actions is 10 seconds. During the so-called ramp-up phase the number of concurrently working users is increased until the expected limit is reached. When all users are active, the test interval starts. This performance level must be maintained for at least 15 minutes (benchmark rule). After at least 5 minutes of the high load phase one or more database checkpoints must be enforced (i.e. all log file data is flushed back to the database within the high load phase) or the amount of created dirty blocks must be written to disk for at least 5 minutes to stress the I/O subsystem in a realistic way (benchmark rule). At the end of the high load phase users are gradually taken off the system until none is active. When the test concludes, all relevant data (some are gathered with a SAP developed Operating System monitor) are then transferred to the presentation server for further evaluation.

users

database checkpoint

∨

(>5min)

time

ramp-up phase

high load phase

ramp-down phase

interval start

interval end (>15min)

A benchmark can only be certified if the average dialog response time is less than 1 second. Certified and published SAP SD Benchmarks are published on SAP's benchmark site [here](here).

## *Benchmark environment*

The benchmark differentiates between a two-tier and a three-tier configuration. The two-tier configuration has the SAP application and database installed on one server. With a three-tier configuration the individual components of the SAP application can be distributed via several servers and an additional server handles the database. The SD benchmark users are simulated by the presentation server aka benchmark driver.
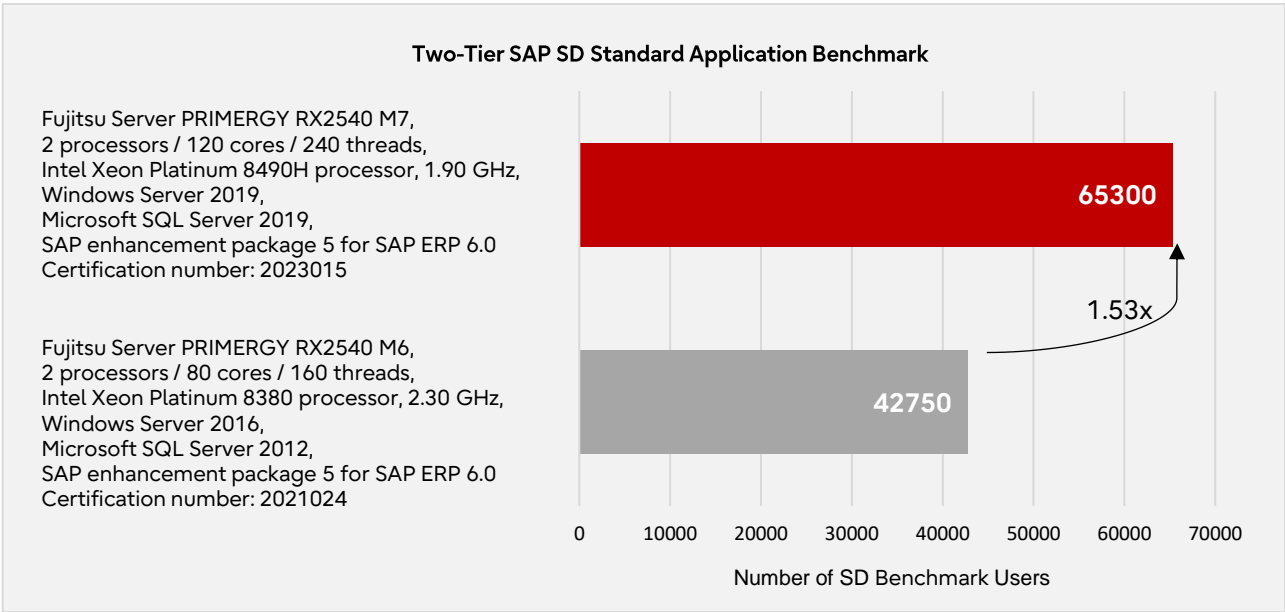
Two SAP SD Benchmark were performed on PRIMERGY RX2540 M7, the first on Windows Server 2019, the second on SUSE Linux Enterprise Server 15, both on a two-tier configuration.
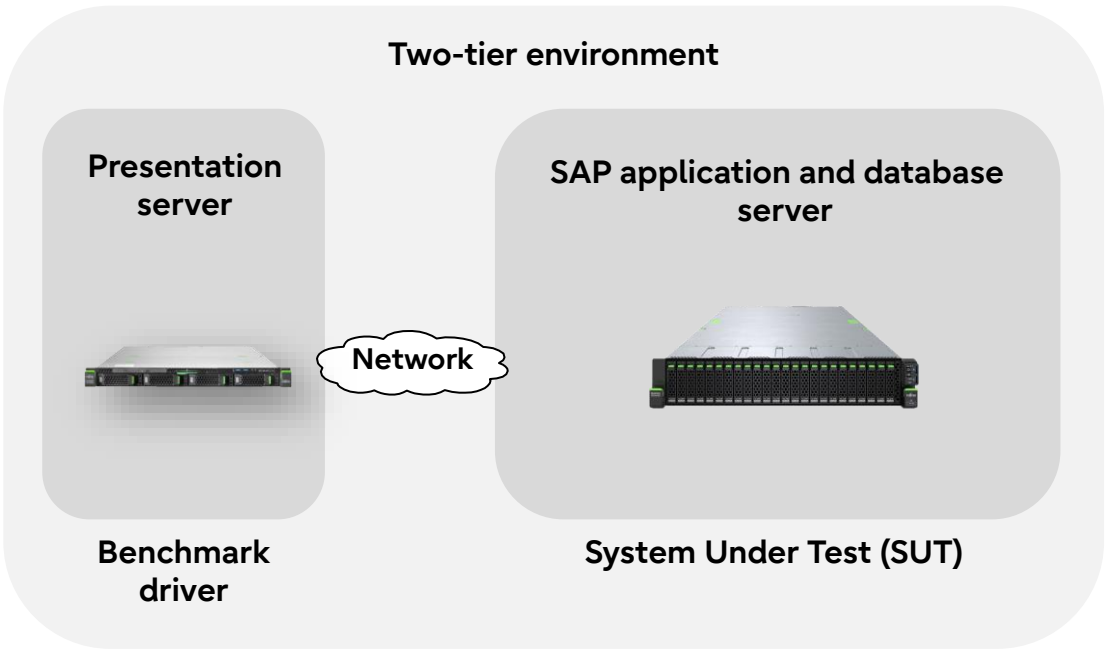


| System Under Test (SUT) – Windows Server 2019 | |
|---|---|
| **Hardware** | |
| • Model | PRIMERGY RX2540 M7 |
| • Processor | 2 x Xeon Platinum 8490H 60C 1.9GHz 350W |
| • Memory | 16 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC |
| • Network interface | 1Gbit/s (RJ45) on Motherboard |
| • Storage subsystem | 3 x PCIe-SSD 2.5" Mixed Use 3.2TB Kioxia CM6-V |
| **Software** | |
| • Operating system | Windows Server 2019 |
| • Database | Microsoft SQL Server 2019 |
| • SAP Business Suite Software | SAP enhancement package 5 for SAP ERP 6.0 |

## System Under Test (SUT) – SUSE Linux Enterprise Server 15

### Hardware

| | |
|---|---|
| • Model | PRIMERGY RX2540 M7 |
| • Processor | 2 x Xeon Platinum 8490H 60C 1.9GHz 350W |
| • Memory | 16 x 128GB (1x128GB) 4Rx4 DDR5-4800 R ECC |
| • Network interface | 1Gbit/s (RJ45) on Motherboard |
| • Storage subsystem | 3 x PCIe-SSD 2.5" Mixed Use 3.2TB Kioxia CM6-V |

### Software

| | |
|---|---|
| • Operating system | SUSE Linux Enterprise Server 15 |
| • Database | SAP ASE 16 |
| • SAP Business Suite Software | SAP enhancement package 5 for SAP ERP 6.0 |

## Benchmark Driver

### Hardware

| | |
|---|---|
| • Model | PRIMERGY RX2530 M1 |
| • Processor | 2 x Xeon E5-2699v3 18C/36T 2.30GHz 45MB 9.6GT/s 2133MHz 145W |
| • Memory | 236 GB |
| • Network interface | 1Gbit/s (RJ45) on Motherboard |

### Software

| | |
|---|---|
| • Operating System | SUSE Linux Enterprise Server 12 SP2 |

## *Benchmark results*

Two SAP SD Benchmark were performed on PRIMERGY RX2540 M7, the first on Windows Server 2019, the second on SUSE Linux Enterprise Server 15, both on a two-tier configuration.

### SAP SD Benchmark on Windows Server 2019

On March 09, 2023, the following SAP Sales and Distribution (SD) Standard Application Benchmark was certified:

| Certification number 2023015 | |
|---|---|
| • Number of SAP SD benchmark users | 65,300 |
| • Average dialog response time | 0.99 seconds |
| • Throughput<br>   Fully processed order line items/hour<br>   Dialog steps/hour<br>   SAPS | <br>7,130,670<br>21,392,000<br>356,530 |
| • Average database request time (dialog/update) | 0.010 sec / 0.011 sec |
| • CPU utilization of central server | 97% |
| • Operating system, central server | Windows Server 2019 |
| • RDBMS | Microsoft SQL Server 2019 |
| • SAP Business Suite software | SAP enhancement package 5 for SAP ERP 6.0 |
| • Configuration Central Server | Fujitsu Server PRIMERGY RX2540 M7,<br>2 processors / 120 cores / 240 threads,<br>Intel Xeon Platinum 8490H processor,<br>1.90 GHz, 80 KB L1 cache and 2,048 KB L2 cache per core,<br>112.5 MB L3 cache per processor, 1,024 GB main memory |

The following chart compares the two-tier SAP SD Standard Application Benchmarks on Windows for PRIMERGY RX2540 M7 and its predecessor RX2540 M6, shown are the number of SD benchmark users.

**Two-Tier SAP SD Standard Application Benchmark**

Fujitsu Server PRIMERGY RX2540 M7,
2 processors / 120 cores / 240 threads,
Intel Xeon Platinum 8490H processor, 1.90 GHz,
Windows Server 2019,
Microsoft SQL Server 2019,
SAP enhancement package 5 for SAP ERP 6.0
Certification number: 2023015
— **65300**

1.53x

Fujitsu Server PRIMERGY RX2540 M6,
2 processors / 80 cores / 160 threads,
Intel Xeon Platinum 8380 processor, 2.30 GHz,
Windows Server 2016,
Microsoft SQL Server 2012,
SAP enhancement package 5 for SAP ERP 6.0
Certification number: 2021024
— **42750**

Number of SD Benchmark Users

Compared to its predecessor RX2540 M6 with the 3rd Generation Xeon Scalable Familiy processor Intel Xeon Platinum 8380, the new Intel Xeon Platinum 8490H processor achieves factor 1.53 more SD benchmark users.

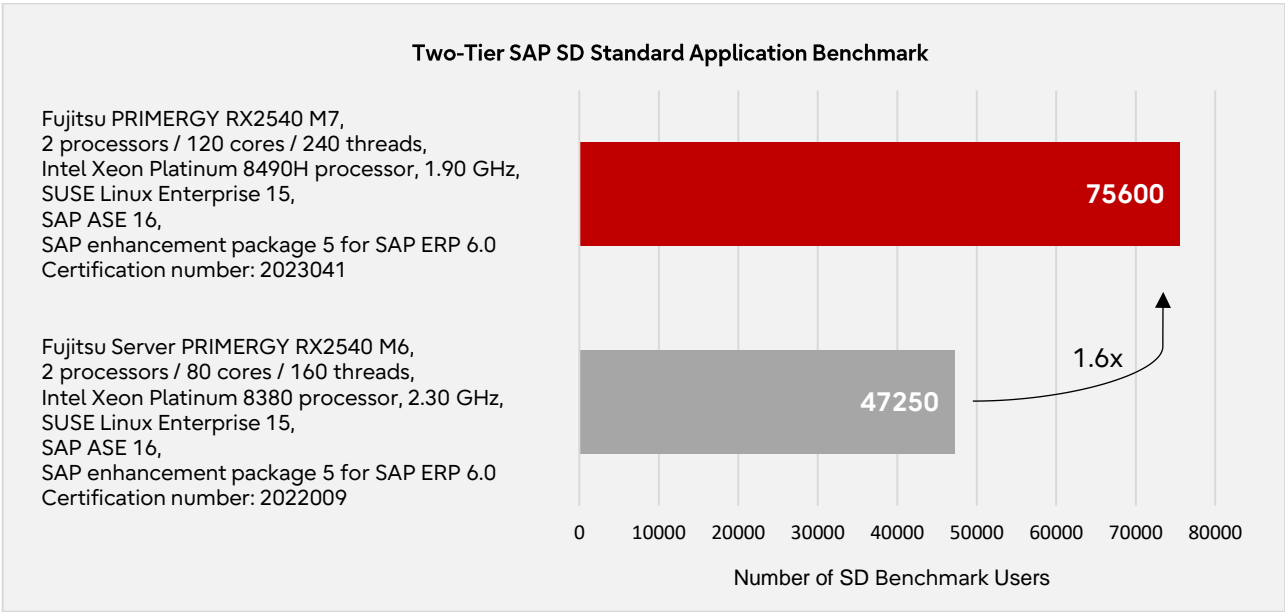The SAP SD Benchmark certificates can be found here: Certification 2023015, Certification 2021024.

**SAP SD Benchmark on SUSE Linux Enterprise Server 15**

On September 14, 2023, the following SAP Sales and Distribution (SD) Standard Application Benchmark was certified:

| Certification number 2023041 | |
| --- | --- |
| • Number of SAP SD benchmark users | 75,600 |
| • Average dialog response time | 0.97 seconds |
| • Throughput<br>Fully processed order line items/hour<br>Dialog steps/hour<br>SAPS | <br>8,269,000<br>24,807,000<br>413,450 |
| • Average database request time (dialog/update) | 0.010 sec / 0.012 sec |
| • CPU utilization of central server | 95% |
| • Operating system, central server | SUSE Linux Enterprise Server 15 |
| • RDBMS | SAP ASE 16 |
| • SAP Business Suite software | SAP enhancement package 5 for SAP ERP 6.0 |
| • Configuration Central Server | Fujitsu PRIMERGY RX2540 M7,<br>2 processors / 120 cores / 240 threads,<br>Intel Xeon Platinum 8490H processor,<br>1.90 GHz, 80 KB L1 cache and 2,048 KB L2 cache per core,<br>112.5 MB L3 cache per processor, 2,048 GB main memory |

The following chart compares the two-tier SAP SD Standard Application Benchmarks on Linux for PRIMERGY RX2540 M7 and its predecessor RX2540 M6, shown are the number of SD benchmark users.

**Two-Tier SAP SD Standard Application Benchmark**

Fujitsu PRIMERGY RX2540 M7,
2 processors / 120 cores / 240 threads,
Intel Xeon Platinum 8490H processor, 1.90 GHz,
SUSE Linux Enterprise 15,
SAP ASE 16,
SAP enhancement package 5 for SAP ERP 6.0
Certification number: 2023041

Fujitsu Server PRIMERGY RX2540 M6,
2 processors / 80 cores / 160 threads,
Intel Xeon Platinum 8380 processor, 2.30 GHz,
SUSE Linux Enterprise 15,
SAP ASE 16,
SAP enhancement package 5 for SAP ERP 6.0
Certification number: 2022009

75600

47250

1.6x

Number of SD Benchmark Users

The 4th Generation Xeon Scalable Family (aka Sapphire Rapids) based RX2540 M7 with Intel Xeon Platinum 8490H delivers an improvement of 1.6x compared to the previous 3rd Generation Xeon Scalable Family (aka Ice Lake) based RX2540 M6 with Intel Xeon Platinum 8380 processor.

The SAP SD Benchmark certificates can be found here: Certification 2023041, Certification 2022009.

# SAP Server Power Standard Application Benchmark

## *Description of the benchmark*

The SAP Server Power Benchmark is based on the SAP Sales & Distribution (SD) benchmark and load profile. The load profile is performed in several load levels. In contrast to the classic SD measurement with only one measurement interval and maximum CPU utilization, the Server Power Benchmark has nine measurement intervals with load levels between Active Idle and 100% (maximum CPU utilization). The diagram below shows these load levels and in particular their sequence. The horizontal sections are the actual measurement intervals. During these phases the QoS (Quality of Service) requirements of the SD benchmark have to be fulfilled - especially the requirement for the average response time per dialog step to be less than one second. As with the SD benchmark, the same requirement of at least 15 minutes also applies for the length of the measurement interval. Between the horizontal phases the number of simulated users is adapted to the required load level in each case.

For the various load levels power measurements with a power meter are conducted.
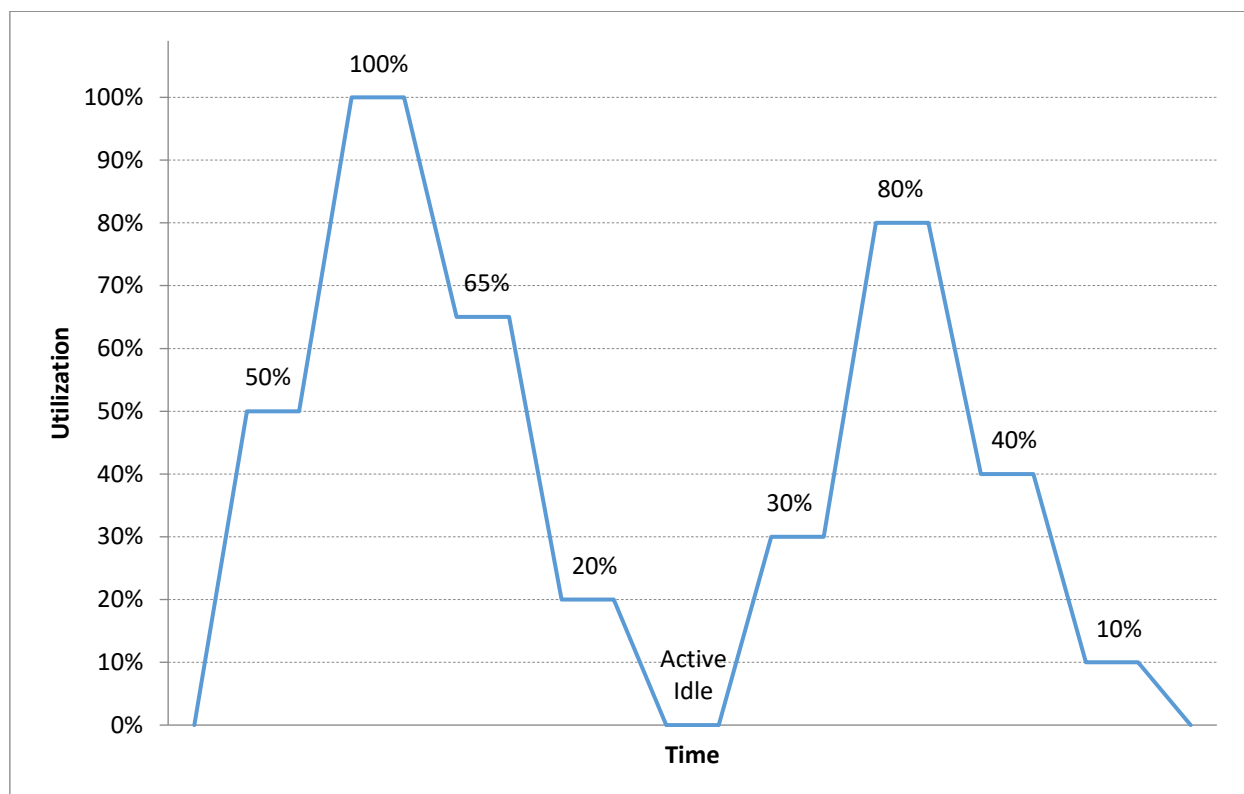
Servers have sophisticated mechanisms for the regulation of power consumption subject to utilization. A particularly effective example is the reduction in CPU frequency at low utilization. The load level methodology in the benchmarks for energy efficiency aims to test the quality of these mechanisms. The second diagram below is taken from the certificate of the SAP Server Power Benchmark on the RX2540 M7 and shows the large range in power consumption between Active Idle and full load.
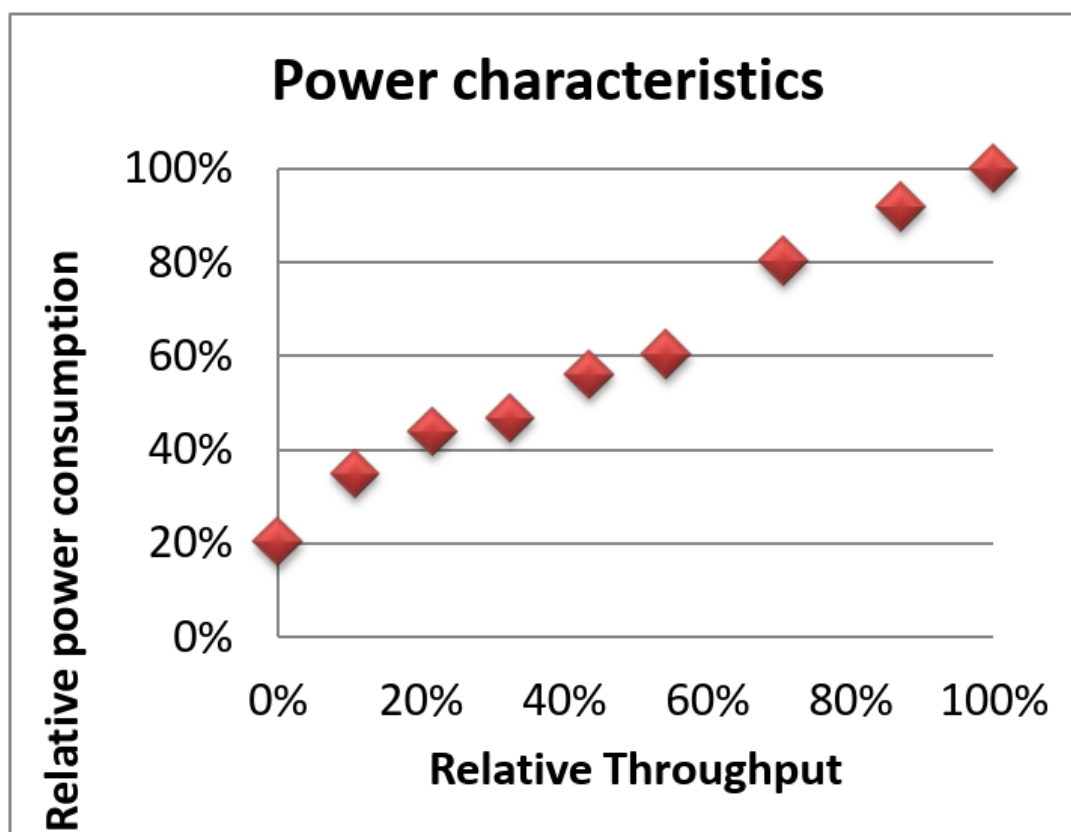
For the sake of completeness it should be mentioned that - apart from power consumption - room temperature is also continuously measured in the vicinity of the measuring configuration during the measurement and may at no point in time be below 20° C. The fan-driven server cooling, which is incorporated in the energy budget, should take place under realistic data center conditions.

The main metric, which is denoted as the Power Efficiency Indicator, is watts/kSAPS. This metric puts power consumption and performance into relation to each other. Performance is expressed by the SAPS throughput measurement. k stands for kilos (1,000). The main metric says how much energy (watts) is needed for the set work quantum of 1,000 SAPS - the lower the value, the better.

The efficiency metric of watts/kSAPS reveals little about which absolute performance level of the servers measured is reached. This is why the arithmetic mean of the nine SAPS values obtained is specified as the second metric: average throughput over all load levels in SAPS.

The SAPS metric is the only performance measure that is specified. The number of respectively configured benchmark users is not mentioned. In this way, the SAP Server Power Benchmark is clearly separated from the SD benchmark, for which the number of users is the primary metric. It does not make sense to compare SAP Server Power measurements and classic publications with the SD benchmark, because the respective goals of maximum energy efficiency and maximum performance are different. Likewise, the secondary SAPS metric of the SD benchmark cannot be compared with the SAPS metric of the power benchmark, because the latter is a mean value taken from nine load levels and the former is a simple value under maximum load.

The minimum room temperature obtained during the measurement is specified as the third metric on the certificate.

## *Benchmark environment*

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

The SAP Server Power Benchmark for PRIMERGY RX2540 M7 was performed on a two-tier configuration.
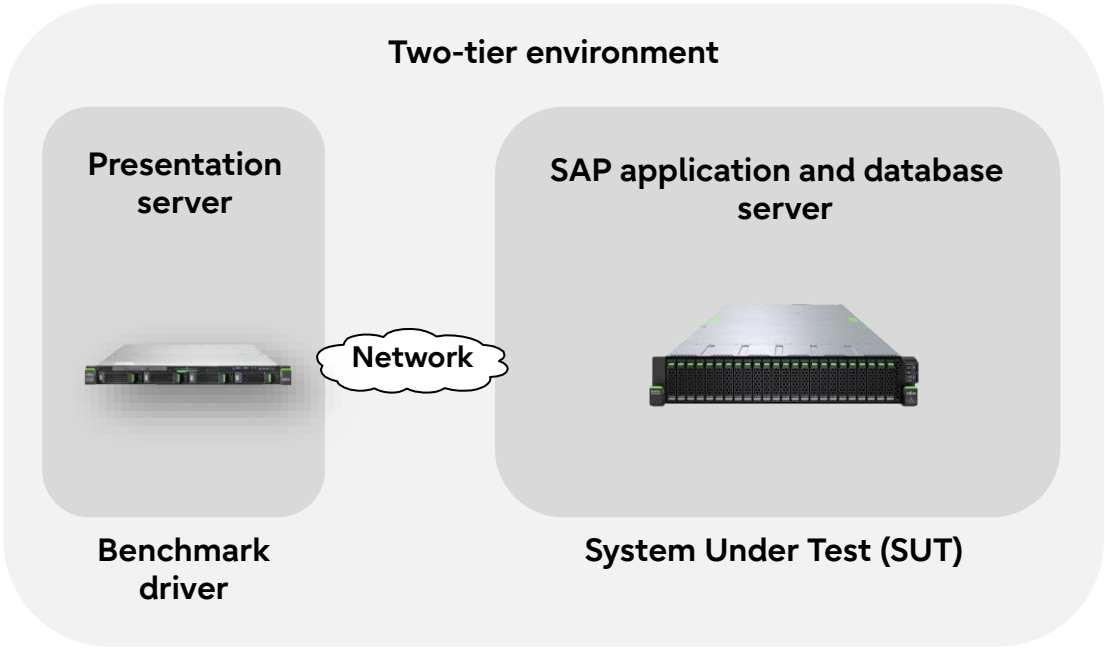


| System Under Test (SUT) | |
|---|---|
| **Hardware** | |
| • Model | PRIMERGY RX2540 M7 |
| • Processor | 2 x Xeon Platinum 8490H 60C 1.9GHz 350W |
| • Memory | 32 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC |
| • Network interface | 1Gbit/s (RJ45) on Motherboard |
| • Storage subsystem | 1 x PCIe-SSD 2.5" Mixed Use 6.4TB Kioxia CM6-V |
| • Power supply | 2 x 1600W titanium PSU |
| **Software** | |
| • Operating system | SUSE Linux Enterprise Server 15 |
| • Database | SAP ASE 16 |
| • SAP Business Suite Software | SAP enhancement package 5 for SAP ERP 6.0 |

## Benchmark Driver

### Hardware

| | |
|---|---|
| • Model | PRIMERGY RX2530 M2 |
| • Processor | 2 x Xeon E5-2699v4 22C/44T 2.20GHz 55MB 9.6GT/s 2400MHz 145W |
| • Memory | 256 GB |
| • Network interface | 1Gbit/s (RJ45) on Motherboard |

### Software

| | |
|---|---|
| • Operating System | SUSE Linux Enterprise Server 12 SP2 |

### Benchmark Driver

## Benchmark results

On November 21, 2023, the following SAP Server Power Standard Application Benchmark was certified:

| Certification number 2023071 | |
|---|---|
| • Power Efficiency Indicator – Server (watts/kSAPS) | 3.89 |
| • Average throughput over all load levels (SAPS) | 180.730 |
| • Minimum ambient temperature (degrees Celsius) | 21.5 |
| • Operating system, central server | SUSE Linux Enterprise Server 15 |
| • RDBMS | SAP ASE 16 |
| • SAP Business Suite software | SAP enhancement package 5 for SAP ERP 6.0 |
| • Configuration Central Server | Fujitsu Server PRIMERGY RX2540 M7, 2 processors / 120 cores / 240 threads, Intel Xeon Platinum 8490H processor, 1.90 GHz, 80 KB L1 cache and 2,048 KB L2 cache per core, 112.5 MB L3 cache per processor, 2,024 GB main memory, 2 x 1600W S26113-E652-V60-1 Titanium, 1 x SSD PCIe4 |

The following diagram shows the range in power consumption between Active Idle and full load:
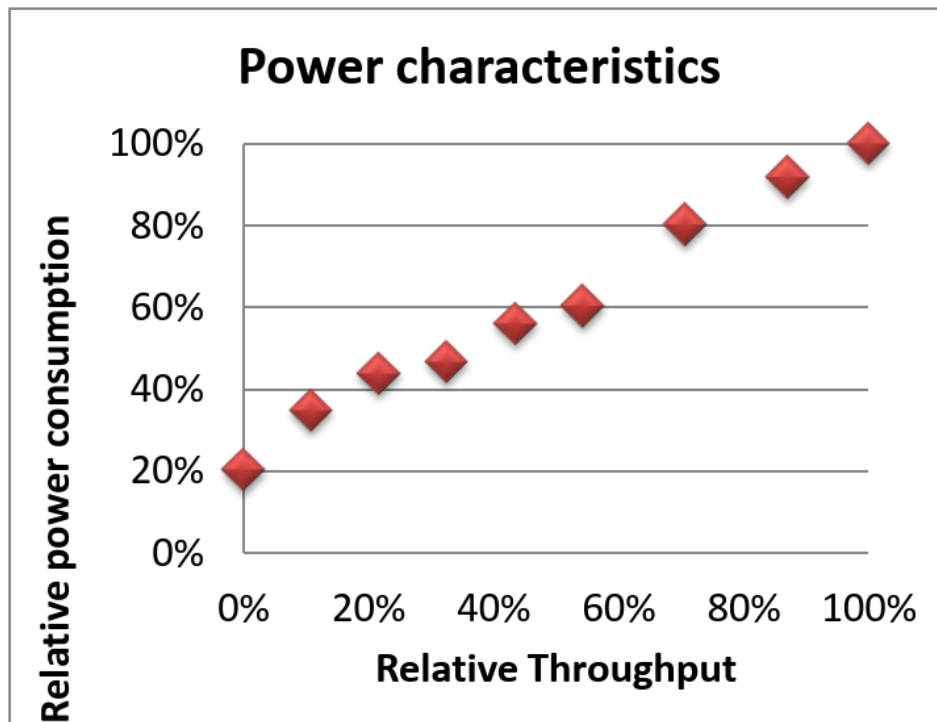


The SAP Server Power Benchmark certificate can be found here: Certification 2023071 .

# SAP BW Edition for SAP HANA Standard Application Benchmark

## *Description of the benchmark*

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

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

Since its first edition in 2016, the SAP BWH Benchmark has been further developed and adapted to customer requirements. In the meantime, SAP BWH Benchmark version 3 is available. Benchmarks with the older versions won't be certified anymore. The results of different versions must not be compared with each other.

The SAP BWH Benchmark consists of 3 phases:

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

### Data load phase

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

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

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

The permissible data volumes are a multiple of 1.3 billion initial data records. The minimum number of data sets to be loaded is dependent on the size of the main memory.

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

### Query throughput phase

The queries for the throughput phase must be executed via an ABAP program with a variant containing 380 queries. Users execute the set of navigation steps in random order (via asynchronous RFCs). The queries contain typical query patterns which can be found in BW productive systems of customers.

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

### Query runtime phase

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

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

Certified and published SAP BWH Benchmarks are published on SAP's benchmark site here.

## Benchmark environment

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

The SAP BWH Benchmarks for PRIMERGY RX2540 M7 were performed on a scale-up configuration.
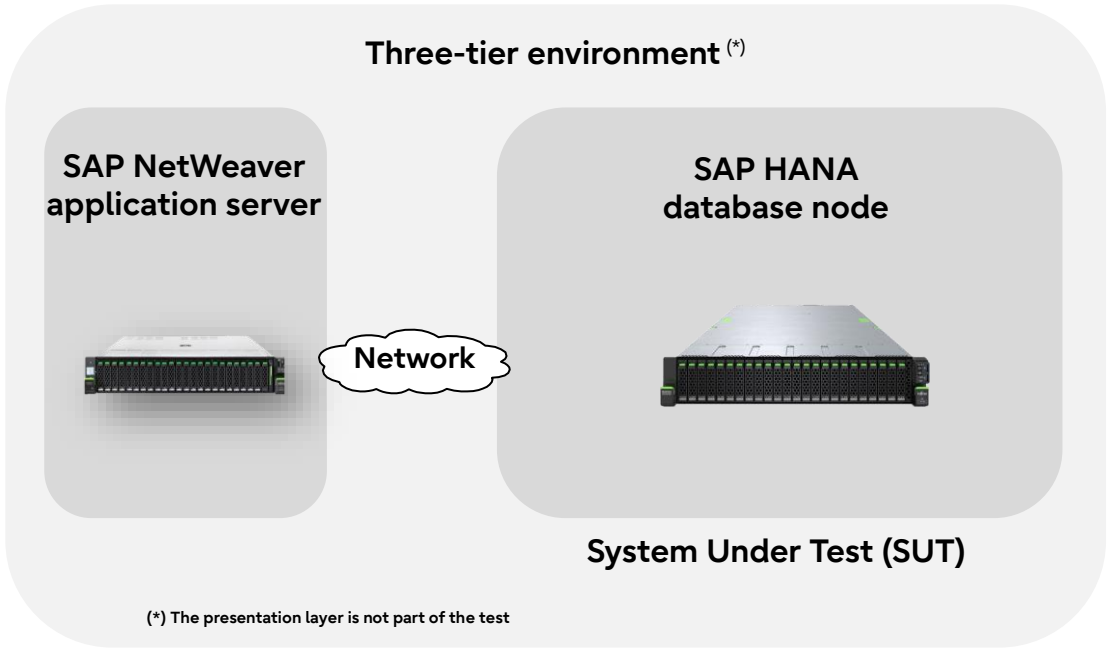


Although an application server is involved in the benchmark, neither performance metrics are measured nor does the server appear on the benchmark certificate.

| System Under Test (SUT) | |
|---|---|
| **Hardware** | |
| • Model | PRIMERGY RX2540 M7 |
| • Processor | 2 x Xeon Platinum 8480+ 56C 2.0GHz 350W |
| • Memory | 32 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC |
| • Network interface | 1Gbit/s (RJ45) on Motherboard |
| • Storage subsystem | 3 x PCIe-SSD 2.5" Mixed Use 3.2TB Kioxia CM6-V |
| **Software** | |
| • Operating system | SUSE Linux Enterprise Server 15 |
| • Database | SAP HANA 2.0 |

## Application Server

### Hardware

| | |
|---|---|
| • Model | PRIMERGY RX2540 M5 |
| • Processor | 2 x Xeon Platinum 8280L 28C 2.7GHz 205W |
| • Memory | 12 x 64GB (1x64GB) 2Rx4 DDR4-2933 R ECC |
| • Storage subsystem | 2 x HDD SAS 2.5" 15K 600GB, 3 x SSD 1.5TB<br>1 x PACC EP P4800X AIC PCIe-SSD 750GB<br>1 x PRAID EP420i RAID Controller |
| • Network interface | 1Gbit/s (RJ45) on Motherboard |

### Software

| | |
|---|---|
| • Operating System | SUSE Linux Enterprise Server 15 |
| • Technology platform release | SAP Netweaver 7.50 |

# Benchmark results

Two SAP BWH Benchmarks were performed on PRIMERGY RX2540 M7, the first with 2.6 billion records, the second with 3.9 billion records.

## 2.6 Billion Records Scenario

On January 10, 2023, the following SAP BW edition for SAP HANA Standard Application Benchmark Version 3 was certified:

| Certification number 2023001 | |
|---|---|
| • Benchmark Phase 1<br>  Number of initial records<br>  Runtime of last Data Set (seconds) | 2,600,000<br>8,750 |
| • Benchmark Phase 2<br>  Query Executions per Hour<br>  CPU utilization of database server | 11,560<br>99% |
| • Benchmark Phase 3<br>  Total Runtime of complex query phase (seconds) | 78 |
| • Operating system | SUSE Linux Enterprise Server 15 |
| • Database | SAP HANA 2.0 |
| • Technology platform release | SAP Netweaver 7.50 |
| • Configuration<br>  Database Server | Fujitsu Server PRIMERGY RX2540 M7,<br>2 processors / 112 cores /224 threads,<br>Intel Xeon Platinum 8480+ processor,<br>2.00 GHz, 80 KB L1 cache and 2,048 KB L2 cache per core,<br>105 MB L3 cache per processor,<br>2,048 GB DRAM |

The benchmark was published along with Intel's launch of the 4th Generation Xeon Scalable Family processors, aka Sapphire Rapids SP and set an overall world record (as of 2023-01-10) on the SAP BW Edition for SAP HANA Standard Application Benchmark Version 3 in the 2.6 billion initial records category.

The SAP BWH Benchmark certificate can be found here: Certification 2023001.

## 3.9 Billion Records Scenario

On January 20, 2023, the following SAP BW edition for SAP HANA Standard Application Benchmark Version 3 was certified:

| Certification number 2023006 | |
|---|---|
| • Benchmark Phase 1<br>   Number of initial records<br>   Runtime of last Data Set (seconds) | 3,900,000<br>9,500 |
| • Benchmark Phase 2<br>   Query Executions per Hour<br>   CPU utilization of database server | 9,251<br>99% |
| • Benchmark Phase 3<br>   Total Runtime of complex query phase (seconds) | 84 |
| • Operating system | SUSE Linux Enterprise Server 15 |
| • Database | SAP HANA 2.0 |
| • Technology platform release | SAP Netweaver 7.50 |
| • Configuration Database Server | Fujitsu Server PRIMERGY RX2540 M7,<br>2 processors / 112 cores /224 threads,<br>Intel Xeon Platinum 8480+ processor,<br>2.00 GHz, 80 KB L1 cache and 2,048 KB L2 cache per core,<br>105 MB L3 cache per processor,<br>2,048 GB DRAM |

Shortly after the first SAP BW Edition for SAP HANA Benchmark world record with 2.6 billion records, PRIMERGY RX2540 M7 set a second overall world record (as of 2023-01-20) on the SAP BW Edition for SAP HANA Standard Application Benchmark Version 3 in the 3.9 billion initial records category.

The SAP BWH Benchmark certificate can be found here: Certification 2023006.

# 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 MB / 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<br>ST600MP0006 |
| | SAS HDD (SAS 12Gbps, 10k rpm) [512e] | AL15SEB18EQ<br>AL15SEB24EQ |
| | SAS HDD (SAS 12Gbps, 10k rpm) [512n] | AL15SEB060N<br>AL15SEB120N |
| | NL-SAS HDD (SAS 12Gbps, 7.2k rpm) [512e] | ST12000NM004J<br>ST18000NM004J |
| | BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512e] | ST12000NM000J<br>ST18000NM000J |
| | BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512n] | ST2000NM000A<br>ST4000NM000A |
| SSD | SAS SSD (SAS 12Gbps, Write Intensive) | XS400ME70084<br>XS800ME70084<br>XS1600ME70084 |
| | SAS SSD (SAS 12Gbps, Mixed Use) | XS800LE70084<br>XS1600LE70084<br>XS3200LE70084 |
| | SAS SSD (SAS 12Gbps, Read Intensive) | XS960SE70084<br>XS1920SE70084<br>XS3840SE70084<br>XS7680SE70084 |
| | SATA SSD (SATA 6Gbps, Mixed Use) | MTFDDAK480TDT<br>MTFDDAK960TDT<br>MTFDDAK1T9TDT<br>MTFDDAK3T8TDT |
| | | MZ7L3480HBLT<br>MZ7L3960HBLT<br>MZ7L31T9HBNA<br>MZ7L33T8HBNA |

**Controller:**          PRAID EP540i

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

**Controller:**          Intel C741 Standard SATA AHCI controller

| Storage media | Category | Drive name |
|---|---|---|
| SSD | M.2 Flash module | MTFDDAK240TDS<br>MTFDDAV480TDS |

**Controller:**          Intel C741 Standard NVM Express controller

| Storage media | Category | Drive name |
|---|---|---|
| SSD | M.2 Flash module (NVMe) | MTFDKBA480TFR<br>MTFDKBA960TFR |

## 2.5 inch model

**Controller:**          PRAID EP540i

| Storage media | Category | Drive name |
|---|---|---|
| HDD | SAS HDD (SAS 12Gbps, 15k rpm) [512n] | ST300MP0006<br>ST600MP0006 |
| | SAS HDD (SAS 12Gbps, 10k rpm) [512e] | AL15SEB18EQ<br>AL15SEB24EQ |
| | SAS HDD (SAS 12Gbps, 10k rpm) [512n] | AL15SEB060N<br>AL15SEB120N |
| | NL-SAS HDD (SAS 12Gbps, 7.2k rpm) [512n | ST2000NX0433 |
| | BC-SATA HDD(SATA 6Gbps, 7.2krpm)[512n] | ST1000NX0423<br>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 inch[1] |
| | | SATA 6G | 3.5 inch |
| | SSD | SAS 12G | 2.5 inch[1] |
| | | SATA 6G | 2.5 inch[1] 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 |
| **- 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 |
| **- 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 |

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

# VMmark V3

## *Benchmark description*

VMmark V3 is a benchmark developed by VMware to compare server configurations with hypervisor solutions from VMware regarding their suitability for server consolidation. In addition to the software for load generation, the benchmark consists of a defined load profile and binding regulations. The benchmark results can be submitted to VMware and are published on their Internet site after a successful review process. After the discontinuation of the proven benchmark "VMmark V2" in September 2017, it has been succeeded by "VMmark V3". VMmark V2 required a cluster of at least two servers and covers data center functions, like Cloning and Deployment of virtual machines (VMs), Load Balancing, as well as the moving of VMs with vMotion and also Storage vMotion. VMmark V3 covers the moving of VMs with XvMotion in addition to VMmark V2. Also, changes application architecture to more scalable workloads.

In addition to the "Performance Only" result, alternatively measure the electrical power consumption and publish it as a "Performance with Server Power" result (power consumption of server systems only) and/or "Performance with Server and Storage Power" result (power consumption of server systems and all storage components).

VMmark V3 is not a new benchmark in the actual sense. It is in fact a framework that consolidates already established benchmarks, as workloads in order to simulate the load of a virtualized consolidated server environment. Two proven benchmarks, which cover the application scenarios Scalable web system and E commerce system were integrated in VMmark V3.

| Application scenario | Load tool | # VMs |
|---|---|---|
| Scalable web system | Weathervane | 14 |
| E-commerce system | DVD Store 3 client | 4 |
| Standby system | | 1 |

Each of the three application scenarios is assigned to a total of 18 dedicated virtual machines. Then add to these an 19th VM called the "standby server". These 19 VMs form a "tile". Because of the performance capability of the underlying server hardware, it is usually necessary to have started several identical tiles in parallel as part of a measurement in order to achieve a maximum overall performance.

In VMmark V3 there is an an infrastructure component, which is present once for every two hosts. It measures the efficiency levels of data center consolidation through VM Cloning and Deployment, vMotion, XvMotion and Storage vMotion. The Load Balancing capacity of the data center is also used (DRS, Distributed Resource Scheduler).

The result of VMmark V3 for test type "Performance Only" is a number, known as a "score", which provides information about the performance of the measured virtualization solution. The score is the maximum sum of the benefits of server aggregation and is used as a comparison criterion for different hardware platforms.

This score is determined from the individual results of the VMs and an infrastructure components result. Each of the five VMmark V3 application or front-end VMs provides a specific benchmark result in the form of application-specific transaction rates for each VM. In order to derive a normalized score, the individual benchmark result for each tile is put in relation to the respective results of a reference system. The resulting dimensionless performance values are then averaged geometrically and finally added up for all VMs. This value is included in the overall score with a weighting of 80%. The infrastructure workload is only present in the benchmark once for every two hosts; it determines 20% of the result. The number of transactions per hour and the average duration in seconds respectively are determined for the score of the infrastructure components workload.

In addition to the actual score, the number of VMmark V3 tiles is always specified with each VMmark V3 score. The result is thus as follows: "Score@Number of Tiles", for example "8.11@8 tiles".

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

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

A detailed description of VMmark V3 is available in the document Benchmark Overview VMmark V3.

## *Benchmark environment*

The typical measurement set-up is illustrated below:
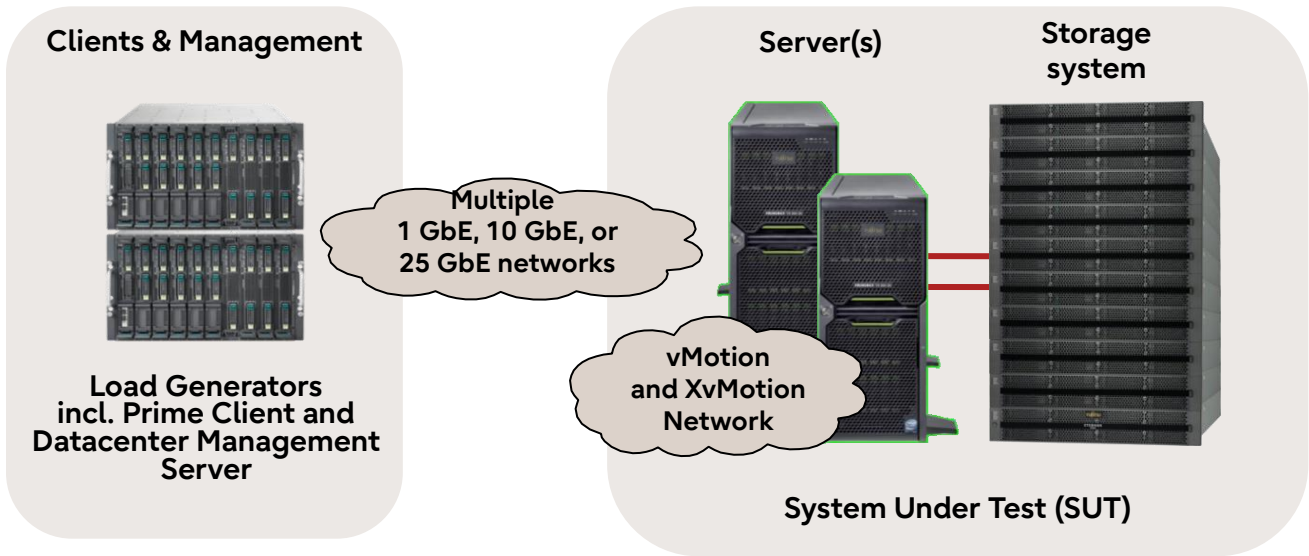


All the benchmark results were measured with the following environment:

| System Under Test (SUT) | |
|---|---|
| **Hardware** | |
| • Number of servers | 2 |
| • Model | PRIMERGY RX2540 M7 |
| • Processor | 2 x Xeon Platinum 8490H |
| • Memory | 2048 GB : 32 x 64 GB (1x64 GB) 2Rx4 DDR5-4800 R ECC |
| • Network interface | 2 x Mellanox MCX4121A-ACAT dual port 25Gb SFP28 PCIe adapter<br>1 x 1Gbit/s (RJ45) on Motherboard |
| • Disk subsystem | 2 x Emulex LPe35002 dual port 32Gb PCIe adapter<br><br>8 x PRIMERGY RX2540 M4 & M5 configured as Fibre Channel targets<br>4 x PRIMERGY RX2540 M4 :<br>   3 x Intel P4800X PCIe SSD (750 GB)<br>   1 x Intel P4600 PCIe SSD (4 TB)<br>1 x PRIMERGY RX2540 M4 :<br>   3 x Intel P4800X PCIe SSD (750 GB)<br>   1 x Intel P4600 PCIe SSD (2 TB)<br>2 x PRIMERGY RX2540 M5 :<br>   3 x Intel P4800X PCIe SSD (750 GB)<br>   1 x Intel P4610 PCIe SSD (3.2 TB)<br>1 x PRIMERGY RX2540 M5 :<br>   2 x Intel P4800X PCIe SSD (750 GB)<br>   1 x Intel P4610 PCIe SSD (3.2 TB) |
| **Software** | |
| • BIOS settings | See "Details" |
| • Operating system | VMware ESXi 8.0 GA, Build 20513097 |
| • Operating system settings | ESX settings: see "Details" |

| Detail | |
|---|---|
| See disclosure | https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2023-03-21-Fujitsu-PRIMERGY-RX2540M7.pdf |
| | https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2023-03-21-Fujitsu-PRIMERGY-RX2540M7-serverPPKW.pdf |
| | https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2023-03-21-Fujitsu-PRIMERGY-RX2540M7-serverstoragePPKW.pdf |

| Datacenter Management Server (DMS) | |
|---|---|
| **Hardware** | |
| • Model | 1 x PRIMERGY RX2530 M2 |
| • Processor | 1 x Intel Xeon E5-2698 v4 |
| • Memory | 80 GB |
| • Network interface | 1 x Emulex One Connect Oce14000 1GbE dual port PCIe adapter |
| **Software** | |
| • Operating system | VMware ESXi 7.0 Update 3c, Build 19193900 |
| **Datacenter Management Server (DMS) VM** | |
| **Hardware** | |
| • Processor | 4 x Logical CPU |
| • Memory | 21 GB |
| • Network interface | 1 x 1 Gbit/s LAN |
| **Software** | |
| • Operating system | VMware vCenter Server Appliance 8.0 GA, Build 20519528 |

| Load generator | |
|---|---|
| **Hardware** | |
| • Model | PRIMERGY RX2530 M2 × 6 |
| • Processor | 4 x PRIMERGY RX2530 M2 : |
| |    2 x Intel Xeon E5-2699 v4 |
| | 2 x PRIMERGY RX2530 M2 : |
| |    2 x Intel Xeon E5-2699A v4 |
| • Memory | 256 GB |
| • Network interface | 1 x Emulex One Connect Oce14000 1GbE dual port PCIe adapter |
| | 1 x Emulex One Connect Oce14000 10GbE dual port PCIe adapter |
| **Software** | |
| • Operating system | VMware ESXi 7.0 Update 3c, Build 19193900 |

## *Benchmark results*

## "Performance Only" measurement results (March 21, 2023)

On March 21, 2023, Fujitsu achieved a VMmark V3.1.1 score of "23.38@23 tiles" using PRIMERGY RX2540 M7 with Xeon Platinum 8490H processors and VMware ESXi 8.0 GA. At this time, the system configuration had a total of 2 x 120 processor cores, and two identical servers were used for the "System Under Test" (SUT). Based on the above results, PRIMERGY RX2540 M7 is rated as the most powerful 2-socket Intel processor based rack server in a "matched pair" configuration with two identical hosts in the official VMmark V3 "Performance Only" ranking (as of the date the benchmark results were published).

All comparisons for the competitor products reflect the status of the date of the publication. For the latest VMmark V3 "Performance Only" results, as well as detailed results and configuration data, see https://www.vmware.com/products/vmmark/results3x.html.

All VMs, their application data, the host operating system, and any additional data needed are stored in a powerful Fiber Channel disk subsystem. This disk subsystem uses fast PCIe SSDs such as Intel Optane to improve storage media response time. Network connectivity with host-side load generators and infrastructure load connectivity between hosts are implemented using 25GbE LAN ports.
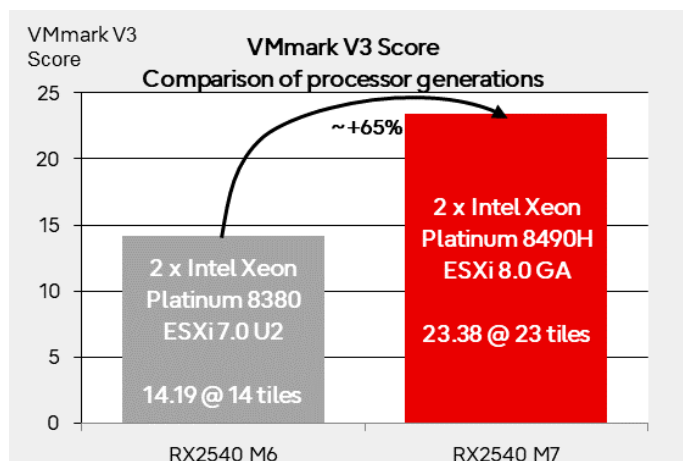


The graph on the left compares the VMmark V3 scores of the PRIMEGY RX2540 M7 and the previous generation PRIMEGY RX2540 M6.

The PRIMERGY RX2540 M7 achieved a 65% improvement in score compared to the previous generation PRIMERGY RX2540 M6. This is due to the improved performance of the 4th generation Intel Xeon scalable processor and the effective use of the capabilities of the VMware ESXi hypervisor.

## "Performance with Server Power" measurement results (March 21, 2023)
## "Performance with Server and Storage Power" measurement results (March 21, 2023)

On March 21, 2023, Fujitsu achieved a VMmark V3.1.1 "Server PPKW" score of "9.7059@22 tiles" using PRIMERGY RX2540 M7 with Xeon Platinum 8490H processors and VMware ESXi 8.0 GA. At the same time, it also achieved a VMmark V3.1.1 "Server and Storage PPKW" score of "4.8019@22 tiles". These were system configurations with a total of 2 x 120 processor cores, and two identical servers were used for the "System Under Test" (SUT). Based on the above results, PRIMERGY RX2540 M7 is rated as the most energy efficient Intel processor based virtual server in the world in the official VMmark V3 "Performance with Server Power" ranking and "Performance with Server and Storage Power" ranking (as of the date the benchmark results were published).

For the latest VMmark V3 "Performance with Server Power" results, detailed results, and configuration data, see https://www.vmware.com/products/vmmark/results3x.1.html.

For the latest VMmark V3 "Performance Server and Storage Power" results, detailed results, and configuration data, see https://www.vmware.com/products/vmmark/results3x.2.html.

# 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:
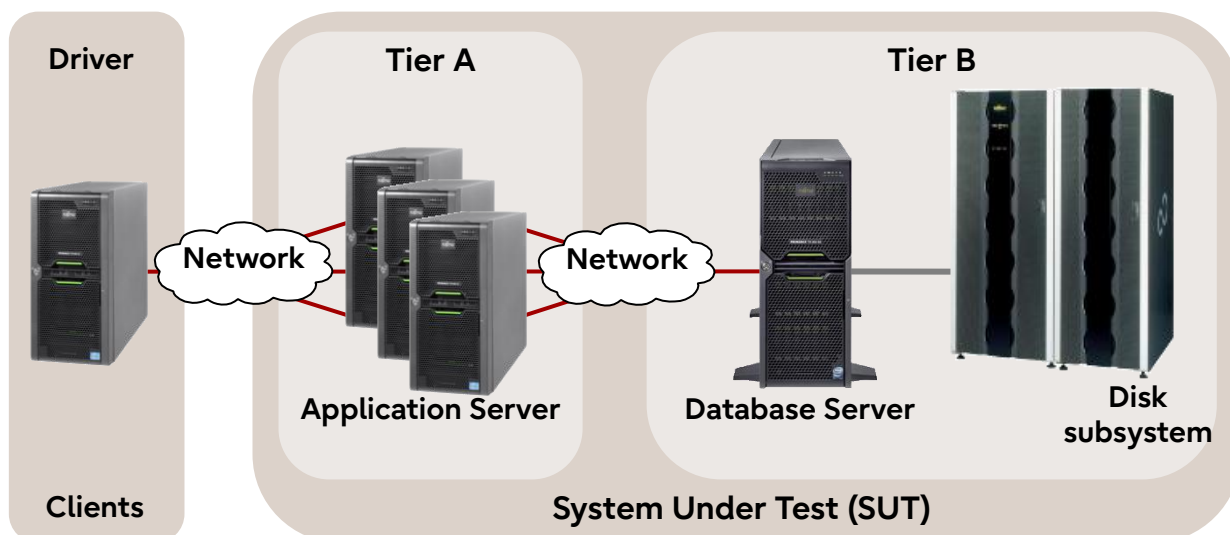


All OLTP-2 results were measured or calculated based on the configuration of the next following pages of PRIMERGY RX2540 M7.

## 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<br>1 x Quad port OCPv3 LAN 1 Gbps |
| • Disk subsystem | RX2540 M7：	1 x RAID controller (internal, 4GB cache)<br>	6 x 1.6 TB SSD drive, RAID10 (log)<br>	5 x RAID controller (external, 4GB cache)<br><br>10 x JX40 S2：	4 x 1.6 TB SSD drive, RAID10（temp）<br>	49 x 1.6 TB SSD drive, RAID5 (data)<br>	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<br>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

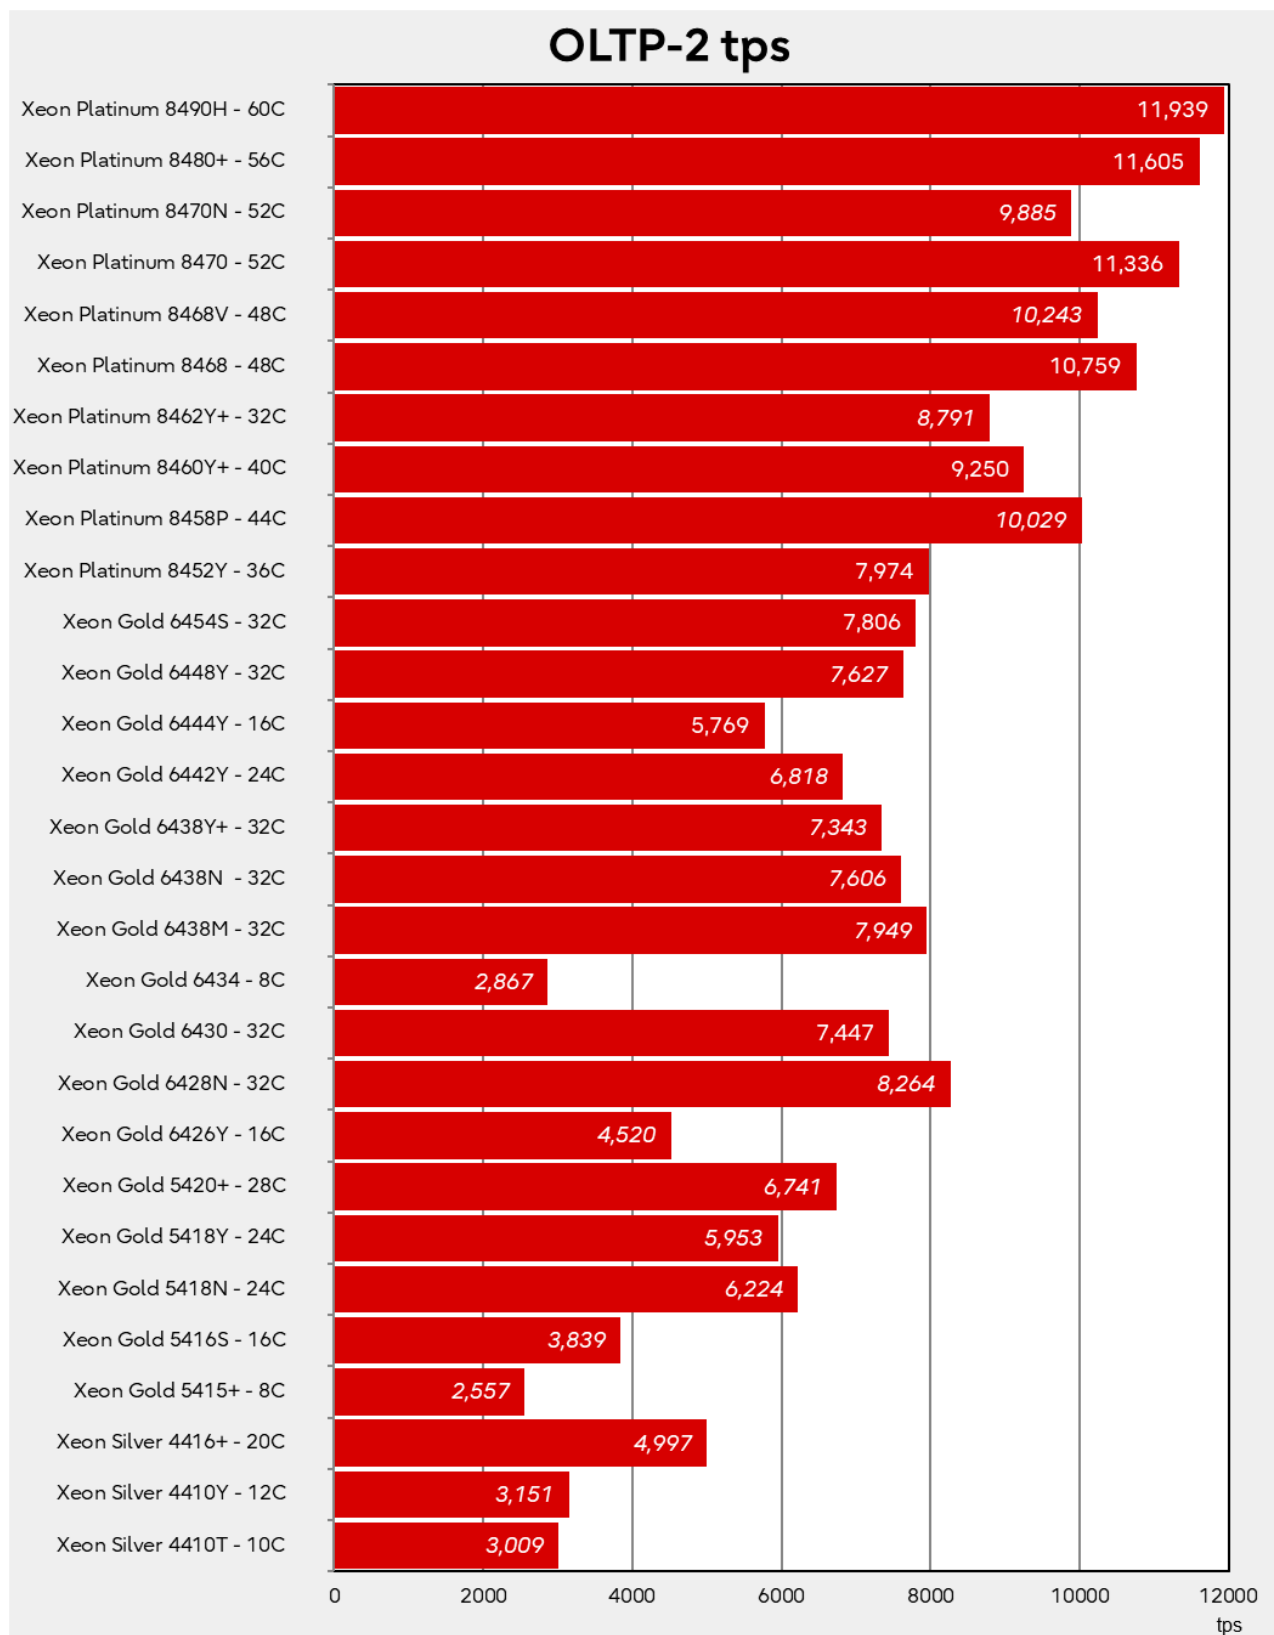| | |
|---|---|
| • Operating system | Microsoft Windows Server 2012 R2 Standard |
| • Benchmark | OLTP-2 Software EGen version 1.14.0 |

## 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 2048 GB was considered for the measurements with two processors. The memory configurations have memory access of 4400 MHz.

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

| Processor | Cores | Threads | 2CPU Score | |
|---|---|---|---|---|
| Xeon Platinum 8490H | 60 | 120 | 11,939 | |
| Xeon Platinum 8480+ | 56 | 112 | 11,605 | |
| Xeon Platinum 8470N | 52 | 104 | 9,885 | est. |
| Xeon Platinum 8470 | 52 | 104 | 11,336 | |
| Xeon Platinum 8468V | 48 | 96 | 10,243 | est. |
| Xeon Platinum 8468 | 48 | 96 | 10,759 | |
| Xeon Platinum 8462Y+ | 32 | 64 | 8,791 | est. |
| Xeon Platinum 8460Y+ | 40 | 80 | 9,250 | |
| Xeon Platinum 8458P | 44 | 88 | 10,029 | est. |
| Xeon Platinum 8452Y | 36 | 72 | 7,974 | est. |
| Xeon Gold 6454S | 32 | 64 | 7,806 | |
| Xeon Gold 6448Y | 32 | 64 | 7,627 | est. |
| Xeon Gold 6444Y | 16 | 32 | 5,769 | |
| Xeon Gold 6442Y | 24 | 48 | 6,818 | est. |
| Xeon Gold 6438Y+ | 32 | 64 | 7,343 | est. |
| Xeon Gold 6438N | 32 | 64 | 7,606 | est. |
| Xeon Gold 6438M | 32 | 64 | 7,949 | est. |
| Xeon Gold 6434 | 8 | 16 | 2,867 | est. |
| Xeon Gold 6430 | 32 | 64 | 7,447 | |
| Xeon Gold 6428N | 32 | 64 | 8,264 | est. |
| Xeon Gold 6426Y | 16 | 32 | 4,520 | est. |
| Xeon Gold 5420+ | 28 | 56 | 6,741 | est. |
| Xeon Gold 5418Y | 24 | 48 | 5,953 | est. |
| Xeon Gold 5418N | 24 | 48 | 6,224 | est. |
| Xeon Gold 5416S | 16 | 32 | 3,839 | est. |
| Xeon Gold 5415+ | 8 | 16 | 2,557 | est. |
| Xeon Silver 4416+ | 20 | 40 | 4,997 | est. |
| Xeon Silver 4410Y | 12 | 24 | 3,151 | est. |
| Xeon Silver 4410T | 10 | 20 | 3,009 | est. |

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

## OLTP-2 tps

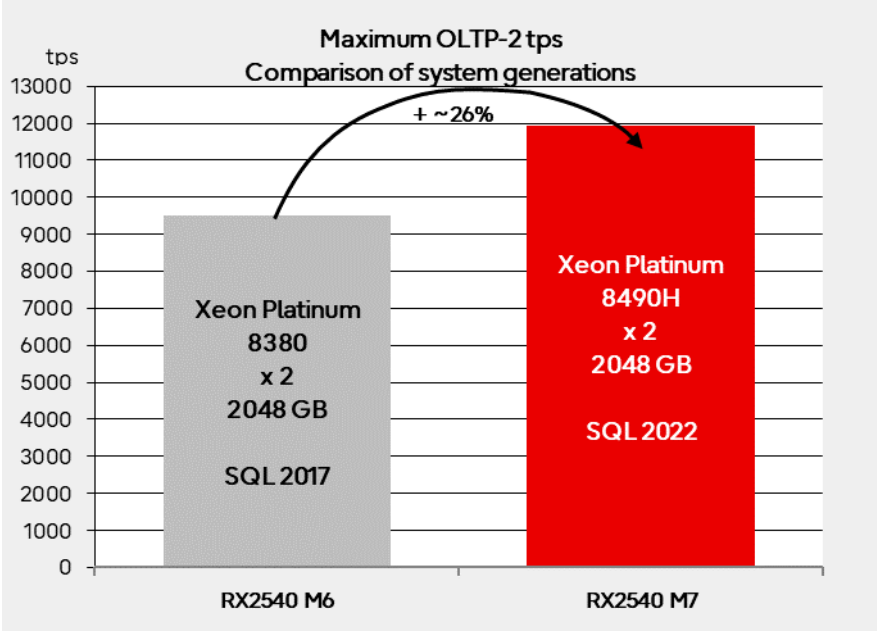| Processor | tps |
|---|---|
| Xeon Platinum 8490H - 60C | 11,939 |
| Xeon Platinum 8480+ - 56C | 11,605 |
| Xeon Platinum 8470N - 52C | 9,885 |
| Xeon Platinum 8470 - 52C | 11,336 |
| Xeon Platinum 8468V - 48C | 10,243 |
| Xeon Platinum 8468 - 48C | 10,759 |
| Xeon Platinum 8462Y+ - 32C | 8,791 |
| Xeon Platinum 8460Y+ - 40C | 9,250 |
| Xeon Platinum 8458P - 44C | 10,029 |
| Xeon Platinum 8452Y - 36C | 7,974 |
| Xeon Gold 6454S - 32C | 7,806 |
| Xeon Gold 6448Y - 32C | 7,627 |
| Xeon Gold 6444Y - 16C | 5,769 |
| Xeon Gold 6442Y - 24C | 6,818 |
| Xeon Gold 6438Y+ - 32C | 7,343 |
| Xeon Gold 6438N - 32C | 7,606 |
| Xeon Gold 6438M - 32C | 7,949 |
| Xeon Gold 6434 - 8C | 2,867 |
| Xeon Gold 6430 - 32C | 7,447 |
| Xeon Gold 6428N - 32C | 8,264 |
| Xeon Gold 6426Y - 16C | 4,520 |
| Xeon Gold 5420+ - 28C | 6,741 |
| Xeon Gold 5418Y - 24C | 5,953 |
| Xeon Gold 5418N - 24C | 6,224 |
| Xeon Gold 5416S - 16C | 3,839 |
| Xeon Gold 5415+ - 8C | 2,557 |
| Xeon Silver 4416+ - 20C | 4,997 |
| Xeon Silver 4410Y - 12C | 3,151 |
| Xeon Silver 4410T - 10C | 3,009 |

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 26% higher than the highest value on the previous model.

# Literature

| **PRIMERGY Servers** |
|---|
| https://www.fujitsu.com/global/products/computing/servers/primergy/ |

| **PRIMERGY RX2530 M7 / RX2540 M7** |
|---|
| This Whitepaper |
| 🇺🇸 https://docs.ts.fujitsu.com/dl.aspx?id=39539f52-34e1-4f06-96d5-f10e3c471dea |
| 🔴 https://docs.ts.fujitsu.com/dl.aspx?id=6099b0c3-30f8-42bc-8e7d-d65349bdd981 |
| Data sheet |
| RX2530 M7: https://docs.ts.fujitsu.com/dl.aspx?id=79c296f8-af2c-45dd-af82-d0f3f9cc06fd |
| RX2540 M7: https://docs.ts.fujitsu.com/dl.aspx?id=b066247f-ad57-44bb-9dcd-afd03e637c3c |

| **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 |

| **SAP SD / Server Power / BWH** |
|---|
| https://www.sap.com/benchmark |
| Benchmark results |
| SAP SD:              https://www.sap.com/dmc/exp/2018-benchmark-directory/#/sd |
| SAP Server Power:    https://www.sap.com/dmc/exp/2018-benchmark-directory/#/power |
| SAP BWH:             https://www.sap.com/dmc/exp/2018-benchmark-directory/#/bwh |
| Benchmark overview |
| SAP SD:        http://docs.ts.fujitsu.com/dl.aspx?id=0a1e69a6-e366-4fd1-a1a6-0dd93148ea10 |
| SAP Server Power: |
|                 http://docs.ts.fujitsu.com/dl.aspx?id=9a500709-589c-4a36-9a5d-bcf28debabd7 |
| SAP BWH:       http://docs.ts.fujitsu.com/dl.aspx?id=70a4c869-586c-49f3-a6a4-47f188dd72b3 |

| **OLTP-2** |
|---|
| Benchmark Overview OLTP-2 |
| https://docs.ts.fujitsu.com/dl.aspx?id=e6f7a4c9-aff6-4598-b199-836053214d3f |

| **VMmark V3** |
|---|
| https://www.vmware.com/products/vmmark.html |
| Benchmark Overview VMmark V3 |
| https://docs.ts.fujitsu.com/dl.aspx?id=e6f9973c-90d6-47c6-b317-e388a978bfb7 |

| **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

## Document change history

| Version | Date | Description |
|---------|------|-------------|
| 1.3 | 2024-01-12 | Update:<br><br>• SAP SD Standard Application Benchmark<br>  Measured on SUSE Linux Enterprise Server 15 with Intel Xeon Platinum 8490H<br>• SAP Server Power Standard Application Benchmark<br>  Measured with Intel Xeon Platinum 8490H |
| 1.2 | 2023-11-09 | Update:<br><br>• Technical data<br>• SPEC CPU2017, STREAM, LINPACK<br>  Measured and calculated additionally with 4th Generation Intel Xeon Processor Scalable Family<br>• OLTP-2<br>  Measured and calculated additionally with 4th Generation Intel Xeon Processor Scalable Family |
| 1.1 | 2023-05-30 | Update:<br><br>• Technical data<br>• SPEC CPU2017, STREAM, LINPACK<br>  Measured and calculated with 4th Generation Intel Xeon Processor Scalable Family<br>• OLTP-2<br>  Measured and calculated with 4th Generation Intel Xeon Processor Scalable Family |
| 1.0 | 2023-04-28 | New:<br><br>• Technical data<br>• SPEC CPU2017, STREAM, LINPACK<br>  Measured and calculated with 4th Generation Intel Xeon Processor Scalable Family<br>• SPECpower_ssj2008<br>  Measured with Intel Xeon Platinum 8490H<br>• SAP SD Standard Application Benchmark<br>  Measured with Intel Xeon Platinum 8490H<br>• SAP BWH Standard Application Benchmark<br>  Measured with Intel Xeon Platinum 8480+<br>• Disk I/O<br>  Measured with 2.5 / 3.5 inch model<br>• VMmark V3<br>  Measured with Intel Xeon Platinum 8490H |

**Contact**

**Fujitsu**
 Web site: https://www.fujitsu.com
**PRIMERGY Performance and Benchmarks**
 mailto:fj-benchmark@dl.jp.fujitsu.com