

# White Paper

## FUJITSU Server PRIMERGY

### Performance Report PRIMERGY RX4770 M6

This document provides an overview of benchmarks executed on the FUJITSU Server PRIMERGY RX4770 M6.

Explains PRIMERGY RX4770 M6 performance data in comparison to other PRIMERGY models. In addition to the benchmark results, the explanation for each benchmark and benchmark environment are also included.

#### Version

1.2

2021/07/28



## Contents

Document history .....	2
Technical data .....	3
SPECcpu2017 .....	6
SPECpower_ssj2008 .....	10
SAP Sales and Distribution (SD) Standard Application Benchmark .....	15
SAP BW Edition for SAP HANA Standard Application Benchmark .....	18
Disk I/O: Performance of storage media .....	22
OLTP-2 .....	30
VMmark V3 .....	35
STREAM .....	39
LINPACK .....	43
Literature .....	46
Contact .....	48

## Document history

### Version 1.0 (2021/03/19)

#### New:

- Technical data
- SPECcpu2017, STREAM, LINPACK  
Measured and calculated with 3rd Generation Intel Xeon Processor Scalable Family
- SPECpower\_ssj2008  
Measured with Intel Xeon Platinum 8380H
- SAP SD  
Measured with Intel Xeon Platinum 8380HL
- SAP BWH  
Measured with Intel Xeon Platinum 8380HL
- Disc I/O  
Measured with 2.5 inch storage
- OLTP-2  
Calculated with 3rd Generation Intel Xeon Processor Scalable Family
- VMmark V3  
Measured with Intel Xeon Platinum 8380HL

### Version 1.2 (2021/07/28)

#### Updated:

- Contact information and URLs  
Updated to the latest one
- Minor correction

## Technical data

### PRIMERGY RX4770 M6



In this document, the power of 10 (example: 1 GB =  $10^9$  bytes) is used to indicate the capacity of the internal storage, and the power of 2 (example: 1 GB =  $2^{30}$  bytes) is used to indicate the capacity of the cache or memory module. Any other exceptional notation will be specified separately.

Model	PRIMERGY RX4770 M6
Form factor	Rack server
Chipset	Intel C621A
Number of sockets	4
Number of configurable processors	2 or 4
Processor type	3rd Generation Intel Xeon Scalable Processors Family
Number of memory slots	48 (12 per processor)
Maximum memory configuration	12,288 GB
Onboard HDD controller	Controller with RAID (0, 1, 10) function (supports up to 8 SATA HDD/SSD)
PCI slots	PCI-Express 3.0 (x16 lane): 7 (3 (Low Profile) + 4 (Full Height)) PCI-Express 3.0 (x8 lane) : 4 (Low Profile)
Maximum number of internal hard disks	2.5 inches x 24

Processor								
Processor model	Number of cores	Number of threads	Cache [MB]	UPI speed [GT/s]	Rated frequency [GHz]	Maximum turbo frequency [GHz]	Maximum memory frequency [MHz]	TDP [W]
Xeon Platinum 8380H	28	56	38.5	10.4	2.9	4.3	3200	250
Xeon Platinum 8380HL	28	56	38.5	10.4	2.9	4.3	3200	250
Xeon Platinum 8376H	28	56	38.5	10.4	2.6	4.3	3200	205
Xeon Platinum 8376HL	28	56	38.5	10.4	2.6	4.3	3200	205
Xeon Platinum 8360H	24	48	33.0	10.4	3.0	4.2	3200	225
Xeon Platinum 8360HL	24	48	33.0	10.4	3.0	4.2	3200	225
Xeon Platinum 8356H	8	16	35.75	10.4	3.9	4.4	3200	205
Xeon Platinum 8354H	18	32	24.75	10.4	3.1	4.3	3200	205
Xeon Gold 6348H	24	48	33.0	10.4	2.3	4.2	2933	165
Xeon Gold 6330H	24	48	33.0	10.4	2.0	3.7	2933	135
Xeon Gold 6328H	16	32	22.0	10.4	2.8	4.3	2933	165
Xeon Gold 6328HL	16	32	22.0	10.4	2.8	4.3	2933	165
Xeon Gold 5320H	20	40	27.5	10.4	2.4	4.2	2667	150
Xeon Gold 5318H	18	36	24.75	10.4	2.5	3.8	2667	150

All processors that can be ordered with PRIMERGY RX4770 M6 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.

## Additional features of the suffixed processors of Xeon Processor

Processors with L support larger memory capacities. The L model supports 4.5 TB, while the regular model supports 1.12 TB of memory per socket.

Suffix	Additional feature
L	Support up to 4.5TB/socket memory

Memory modules									
Type	Capacity [GB]	Number of ranks	Bit width of the memory chips	Frequency [MHz]	3DS	Load Reduced	Registered	NVDIMM	ECC
8 GB (1x8 GB) 1Rx8 DDR4-3200 R ECC	8	1	8	3200			✓		✓
16 GB (1x16 GB) 2Rx8 DDR4-3200 R ECC	16	2	8	3200			✓		✓
16 GB (1x16 GB) 1Rx4 DDR4-3200 R ECC	16	1	4	3200			✓		✓
32 GB (1x32 GB) 2Rx4 DDR4-3200 R ECC	32	2	4	3200			✓		✓
64 GB (1x64 GB) 2Rx4 DDR4-3200 R ECC	64	2	4	3200			✓		✓
64 GB (1x64 GB) 4Rx4 DDR4-3200 LR ECC	64	4	4	3200		✓	✓		✓
128 GB (1x128 GB) 4Rx4 DDR4-3200 3DS ECC	128	4	4	3200	✓		✓		✓
128 GB (1x128 GB) 4Rx4 DDR4-3200 LR ECC	128	4	4	3200		✓	✓		✓
256 GB (1x256 GB) 8Rx4 DDR4-3200 3DS ECC	256	8	4	3200	✓		✓		✓
128GB (1x128GB) Optane PMem-3200	128			3200				✓	✓
256GB (1x256GB) Optane PMem-3200	256			3200				✓	✓
512GB (1x512GB) Optane PMem-3200	512			3200				✓	✓

Power supplies	Maximum number
Modular PSU 1600W platinum	2
Modular PSU 2200W platinum	2
Modular PSU 900W Titanium	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 RX4770 M6.

## SPECcpu2017

### Benchmark description

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

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

SPECcpu2017 contains two different performance measurement methods. The first method (SPECSpeed 2017 Integer or SPECSpeed 2017 Floating Point) determines the time which is required to process a single task. The second method (SPECrate 2017 Integer or SPECrate 2017 Floating Point) determines the throughput, i.e. the number of tasks that can be handled in parallel. Both methods are also divided into two measurement runs, "base" and "peak." 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 SPECcpu2017 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)	
<b>Hardware</b>	
Model	PRIMERGY RX4770 M6
Processor	4 x 3rd Generation Intel Xeon Scalable Processors Family
Memory	48 x 32 GB (1x32 GB) 2Rx4 PC4-3200AA-R
<b>Software</b>	
BIOS settings	<p>SPECSpeed2017_int:  Hyper Threading = Disabled  Override OS Energy Performance = Enabled  Energy Performance = Performance  HWPM Support = OOB mode  Satel AtoS = Disabled</p> <p>SPECSpeed2017_fp:  Hyper Threading = Disabled  Adjacent Cache Line Prefetch = Disabled  DCU Streamer Prefetcher = Disabled  VT-d = Disabled  Stale AtoS = Disabled  Patrol Scrub = Enabled  FAN Control = Full</p> <p>SPECrate2017_int:  DCU Streamer Prefetcher = Disabled  Intel Virtualization Technology = Disabled  Utilization Profile = Unbalanced  Stle AtoS = Disabled  LLC Dead Line Alloc = Disabled  XPT Prefetch = Enable  Patrol Scrub = Enabled  SNC = Enable SNC2  Fan Control = Full</p> <p>SPECrate2017_fp:  Hyper Threading = Disabled<sup>4</sup>  Utilization Profile = Unbalanced  Stle AtoS = Disabled  LLC Dead Line Alloc = Disabled  SNC = Enable SNC2  Fan Control = Full</p>
Operating system	<p>SPECrate2017_int, SPECrate2017_fp:  SUSE Linux Enterprise Server 15 SP2 5.3.18-22-default</p> <p>SPECSpeed2017_int, SPECSpeed2017_fp:  Red Hat Enterprise Linux Server release 8.2 4.18.0-193.el8.x86_64</p>
Operating system settings	<p>Stack size set to unlimited using "ulimit -s unlimited"</p> <p>SPECSpeed2017_int, SPECrate2017:  Kernel Boot Parameter set with : nohz_full=1-X  (X: logical core number -1)</p> <p>SPECrate2017_int:  sched_min_granularity_ns = 15000000</p> <p>SPECrate2017_fp:  service irqbalance stop</p>
Compiler	<p>C/C++: Version 19.1.1.217 of Intel C/C++ Compiler for Linux</p> <p>Fortran: Version 19.1.1.217 of Intel Fortran Compiler for Linux</p>

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

## Benchmark results

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

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


SPECrate2017				
Processor	Number of cores	Number of processors	SPECrate2017 int_base	SPECrate2017 fp_base
Xeon Platinum 8380H	28	4	773 (est.)	650 (est.)
Xeon Platinum 8380HL	28	4	<b>773</b>	<b>650</b>
Xeon Platinum 8376H	28	4	714 (est.)	604 (est.)
Xeon Platinum 8376HL	28	4	<b>714</b>	<b>604</b>
Xeon Platinum 8360H	24	4	<b>662</b>	<b>586</b>
Xeon Platinum 8360HL	24	4	662 (est.)	586 (est.)
Xeon Platinum 8356H	8	4	<b>300</b>	<b>317</b>
Xeon Platinum 8354H	18	4	<b>541</b>	<b>508</b>
Xeon Gold 6348H	24	4	<b>569</b>	<b>513</b>
Xeon Gold 6330H	24	4	<b>539</b>	<b>496</b>
Xeon Gold 6328H	16	4	<b>457</b>	<b>440</b>
Xeon Gold 6328HL	16	4	457 (est.)	440 (est.)
Xeon Gold 5320H	20	4	<b>514</b>	<b>473</b>
Xeon Gold 5318H	18	4	<b>468</b>	<b>433</b>

SPECspeed2017				
Processor	Number of cores	Number of processors	SPECspeed2017 int_base	SPECspeed2017 fp_base
Xeon Platinum 8380HL	28	4		<b>254</b>
Xeon Platinum 8356H	8	4	<b>12.5</b>	



The following graph compares the throughput of PRIMERGY RX4770 M6 and its older model, PRIMERGY RX4770 M5, with maximum performance configurations.

**SPECrate2017: Comparison of PRIMERGY RX4770 M6 and PRIMERGYRX4770 M5**



On October 28, 2020, PRIMERGY RX4770 M6 with the Xeon Platinum 8380HL processor won first place in the 4-socket Xeon category of the SPECrate2017\_fp\_base benchmark.



On October 28, 2020, PRIMERGY RX4770 M6 with the Xeon Platinum 8380HL processor won first place in the 4-socket Xeon category of the SPECspeed2017\_int\_base benchmark.



On October 28, 2020, PRIMERGY RX4770 M6 with the Xeon Platinum 8380HL processor won first place in the 4-socket Xeon category of the SPECspeed2017\_fp\_base benchmark.

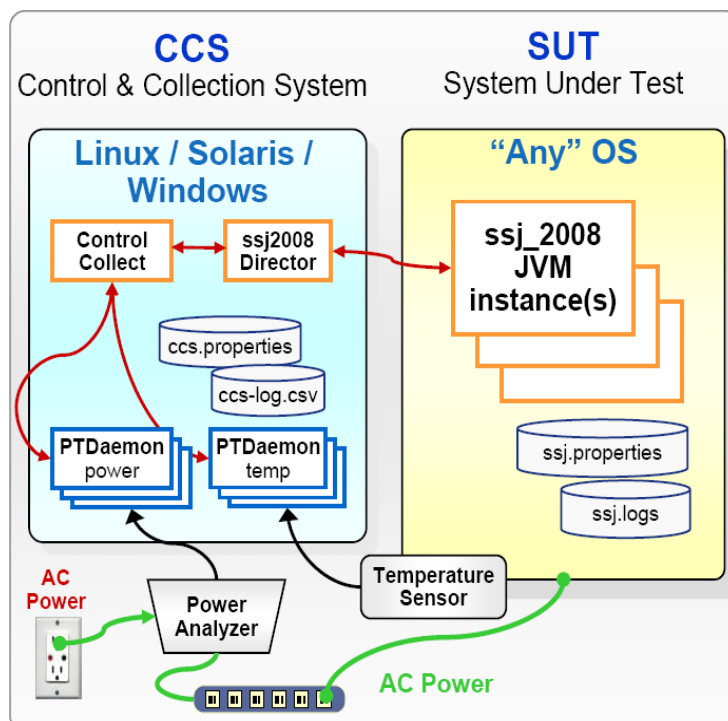

## SPECpower\_ssj2008

### 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)	
<b>Hardware</b>	
Model	PRIMERGY RX4770 M6
Processor	4 x Intel Xeon Platinum 8380H
Memory	24 x 16 GB 2Rx8 PC4-3200Y R
Network interface	1 x PLAN CP I350-T4 4X 1000BASE-T OCPv3
Disk subsystem	1 x SSD M.2 240GB, S26361-F5787-E240
Power Supply Unit	2 x 900 W Titanium S26113-E629-V50-1
<b>Software</b>	
BIOS	R1.1.1
BIOS settings	ASPM Support = L1 only Hardware Prefetcher = Disabled Adjacent Cache Line Prefetch = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Turbo Mode = Disabled Energy Performance = Energy Efficient Override OS Energy Performance = Enabled HWPM Support = Disabled Package C State limit = No limit UPI Link Frequency Select = 9.6GT/s Local x2APIC = Disabled Uncore Frequency Override = Power balanced DDR Performance = Power balanced (effective memory frequency = 2400 MHz) SNC(Sub NUMA) = Enabled SATA Controller = Disabled USB Port Control = Disable all ports Serial Port = Disabled Network Stack = Disabled
Firmware	3.00P
Operating system	SUSE Linux Enterprise Server 15 SP2 5.3.18-19-default

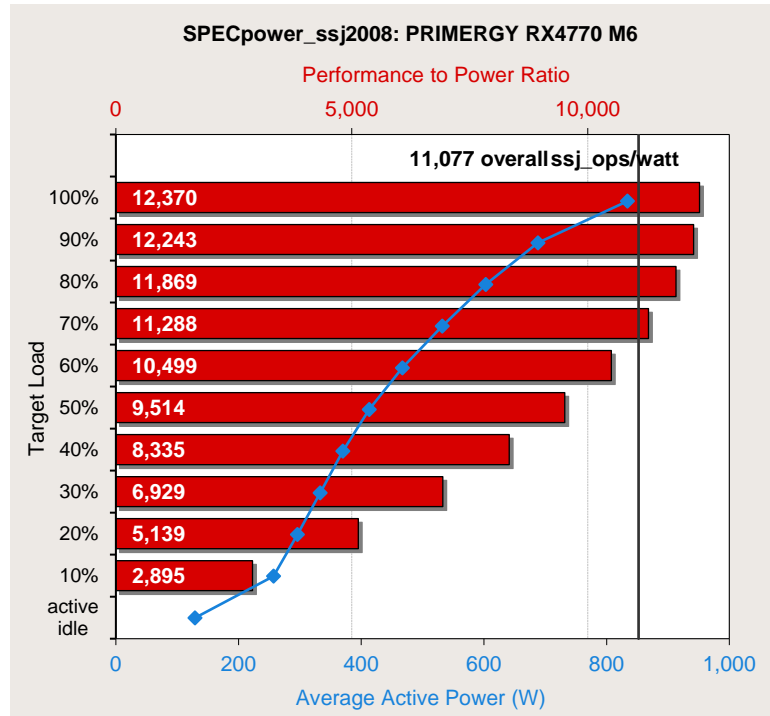
Operating system settings	<p>kernel parameter: pcie_aspm=force pcie_aspm.policy=powersave intel_pstate=disable</p> <p>Benchmark started via ssh</p> <p>modprobe cpufreq_conservative</p> <p>cpupower frequency-set -g conservative</p> <p>echo 1000000 &gt; /sys/devices/system/cpu/cpufreq/conservative/sampling_rate</p> <p>echo 94 &gt; /sys/devices/system/cpu/cpufreq/conservative/up_threshold</p> <p>echo 0 &gt; /sys/devices/system/cpu/cpufreq/conservative/ignore_nice_load</p> <p>echo 1 &gt; /sys/devices/system/cpu/cpufreq/conservative/sampling_down_factor</p> <p>echo 1 &gt; /sys/devices/system/cpu/cpufreq/conservative/freq_step</p> <p>echo 93 &gt; /sys/devices/system/cpu/cpufreq/conservative/down_threshold</p> <p>echo always &gt; /sys/kernel/mm/transparent_hugepage/enabled</p> <p>sysctl -w kernel.sched_migration_cost_ns=6000</p> <p>sysctl -w kernel.sched_min_granularity_ns=10000000</p> <p>sysctl -w kernel.nmi_watchdog=0</p> <p>sysctl -w vm.swappiness=50</p> <p>sysctl -w vm.laptop_mode=5</p> <p>powertop --auto-tune</p> <p>&lt;Yes&gt; The test sponsor attests, as of date of publication, that CVE-2017-5754 (Meltdown) is mitigated in the system as tested and documented.</p> <p>&lt;Yes&gt; The test sponsor attests, as of date of publication, that CVE-2017-5753 (Spectre variant 1) is mitigated in the system as tested and documented.</p> <p>&lt;Yes&gt; The test sponsor attests, as of date of publication, that CVE-2017-5715 (Spectre variant 2) is mitigated in the system as tested and documented.</p>
JVM	Oracle Java HotSpot 64-Bit Server VM 18.9 (build 11.0.8+10-LTS, mixed mode), version 11.0.8
JVM settings	<p>-server -Xmn23000m -Xms25000m -Xmx25000m -XX:SurvivorRatio=1</p> <p>-XX:TargetSurvivorRatio=99 -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4</p> <p>-XX:ParallelGCThreads=14 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270</p> <p>-XX:FreqInlineSize=2500 -XX:+UseLargePages -XX:+UseParallelOldGC</p> <p>-XX:AllocatePrefetchInstr=0 -XX:MinJumpTableSize=18 -XX:UseAVX=0</p> <p>-XX:+UseHugeTLBFS -XX:+UseTransparentHugePages</p>

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

## Benchmark results

The PRIMERGY RX4770 M6 in SUSE Linux Enterprise Server 15 SP2 achieved the following result.

**SPECpower\_ssj2008 = 11,077 overall ssj\_ops/watt**



The graph on the left shows the above measurement results. 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 rhombus. The black vertical line shows the benchmark result of 11,077 overall ssj\_ops/watt for the PRIMERGY RX4770 M6. This is the quotient of the sum of the transaction throughputs for each load level and the sum of the average power consumed for each measurement interval.

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

Performance		Power	Energy Efficiency
Target Load	ssj_ops	Average Power (W)	ssj_ops/watt
100%	9,904,110	834	11,882
90%	8,914,839	688	12,956
80%	7,933,444	603	13,158
70%	6,944,695	532	13,043
60%	5,948,299	467	12,729
50%	4,956,147	413	12,004
40%	3,966,282	370	10,727
30%	2,973,866	333	8,931
20%	1,980,442	296	6,697
10%	990,995	257	3,854
Active Idle	0	129	0
$\sum \text{ssj\_ops} / \sum \text{power} = 11,077$			

**Measurement results of SPECpower\_ssj2008 (November 4, 2020)****11,077 SPECpower\_ssj2008**

On November 4, 2020, PRIMERGY RX4770 M6 with four Xeon Platinum 8380H processors achieved a performance value of 11,077 on the SUSE Linux Enterprise Server 15 SP2 in the SPECpower\_ssj2008 benchmark, in the Linux division of the 3rd Generation Intel Xeon Processor Scalable Family category and won first place in 4-socket server SPECpower\_ssj2008 performance.

For the latest results of the SPECpower\_ssj2008, see [https://www.spec.org/power\\_ssj2008/results/](https://www.spec.org/power_ssj2008/results/).

# SAP Sales and Distribution (SD) Standard Application Benchmark

## Benchmark description

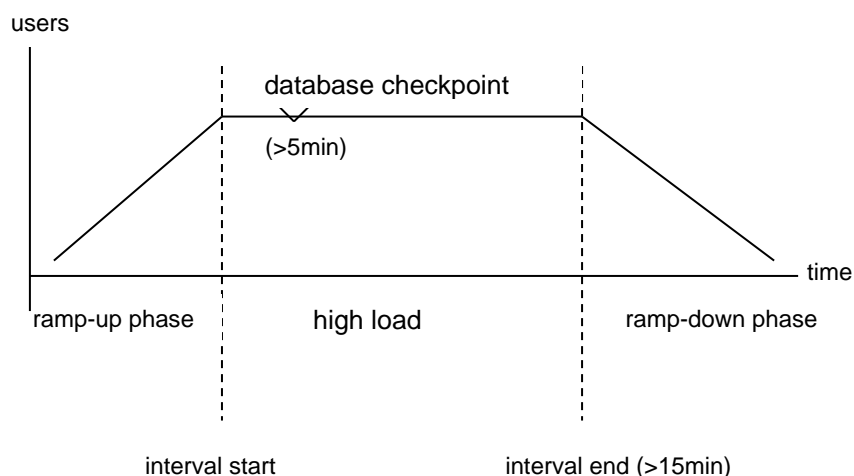
Since 1993 the SAP Standard Application Benchmarks have been developed by SAP in order to verify the performance, stability and scaling of a SAP application system and to provide information for configuring, sizing and for platform comparison. By far the most popular benchmarks from the many available 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.



A benchmark can only be certified if the average dialog response time is less than 1 second.

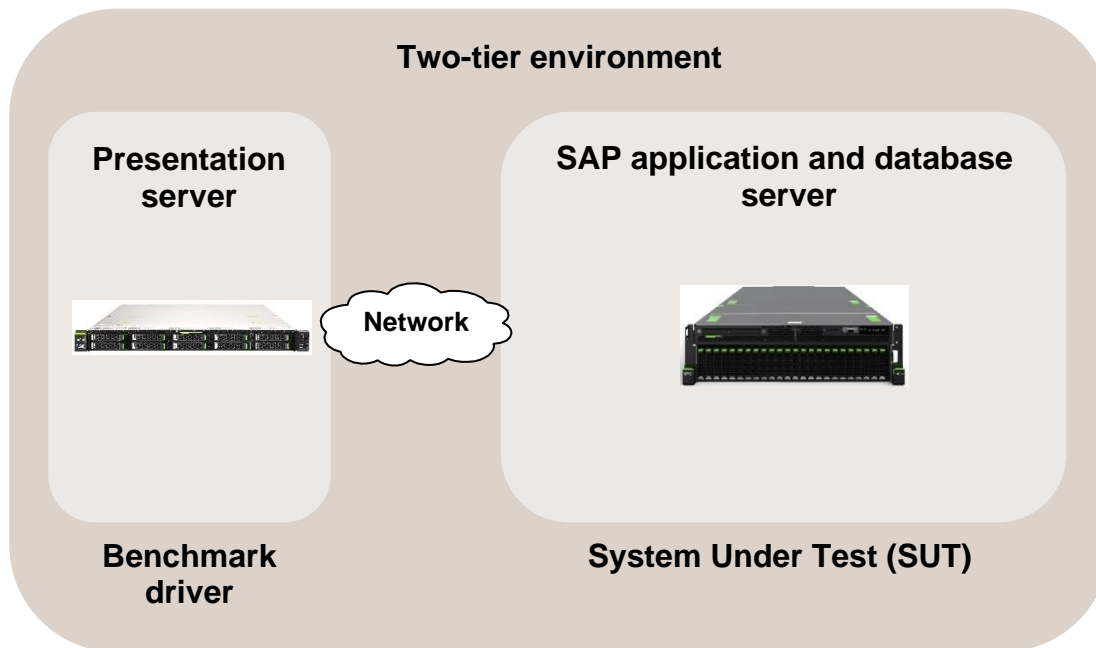
Certified and published SAP SD Benchmarks are published on SAP's benchmark site [here](#).

## 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 SD Benchmark for PRIMERGY RX4770 M6 was performed on a two-tier configuration.



System Under Test (SUT)	
<b>Hardware</b>	
Model	PRIMERGY RX4770 M6
Processor	4 × Intel Xeon Platinum 8380HL processor
Memory	48 × 32 GB 2Rx4 DDR4-3200 R ECC
Network interface	1 Gbit LAN
Disk subsystem	PRIMERGY RX4770 M6: 1 × PRAID EP580i RAID Controller 1 × Internal HDD 900GB 1 × NVMe P4610-3.2TB 2 × PRAID EP540e RAID Controller JX40 S2 × 3: 25 × JX40 S2 MLC SSD 960GB 3DWPD
<b>Software</b>	
Operating system	Windows Server 2016 Datacenter
Database	Microsoft SQL Server 2012
SAP Business Suite Software	SAP enhancement package 5 for SAP ERP 6.0




Benchmark driver	
Hardware	
Model	PRIMERGY RX2530 M1
Processor	2 × Intel Xeon E5-2699 v3 processor
Memory	236 GB
Network interface	1 Gbit LAN
Software	
Operating system	SUSE Linux Enterprise Server 12 SP2

## Benchmark results

Certification number 2020051	
Number of SAP SD benchmark users	60,700
Average dialog response time	0.93 seconds
Throughput	
Fully processed order line items/hour	6,665,670
Dialog steps/hour	19,997,000
SAPS	333,280
Average database request time (dialog/update)	0.009 sec / 0.018 sec
CPU utilization of central server	98%
Operating system, central server	Windows Server 2016 Datacenter
RDBMS	Microsoft SQL Server 2012
SAP Business Suite software	SAP enhancement package 5 for SAP ERP 6.0
Configuration Central Server	Fujitsu Server PRIMERGY RX4770 M6, 4 processors / 112 cores / 224 threads, Intel Xeon Platinum 8380HL Processor, 2.90 GHz, 64 KB L1 cache per core and 1,024KB L2 cache per core, 38.5 MB L3 cache per socket, 1,536 GB main memory

The following chart compares the two-tier SAP SD Standard Application Benchmarks for PRIMERGY RX4770 M6 and its predecessor RX4770 M5, shown are the number of benchmark users.



The SAP SD Benchmark certificates can be found here: [Certificate 2020051](#) , [Certificate 2019033](#)

# SAP BW Edition for SAP HANA Standard Application Benchmark

## Benchmark description

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.

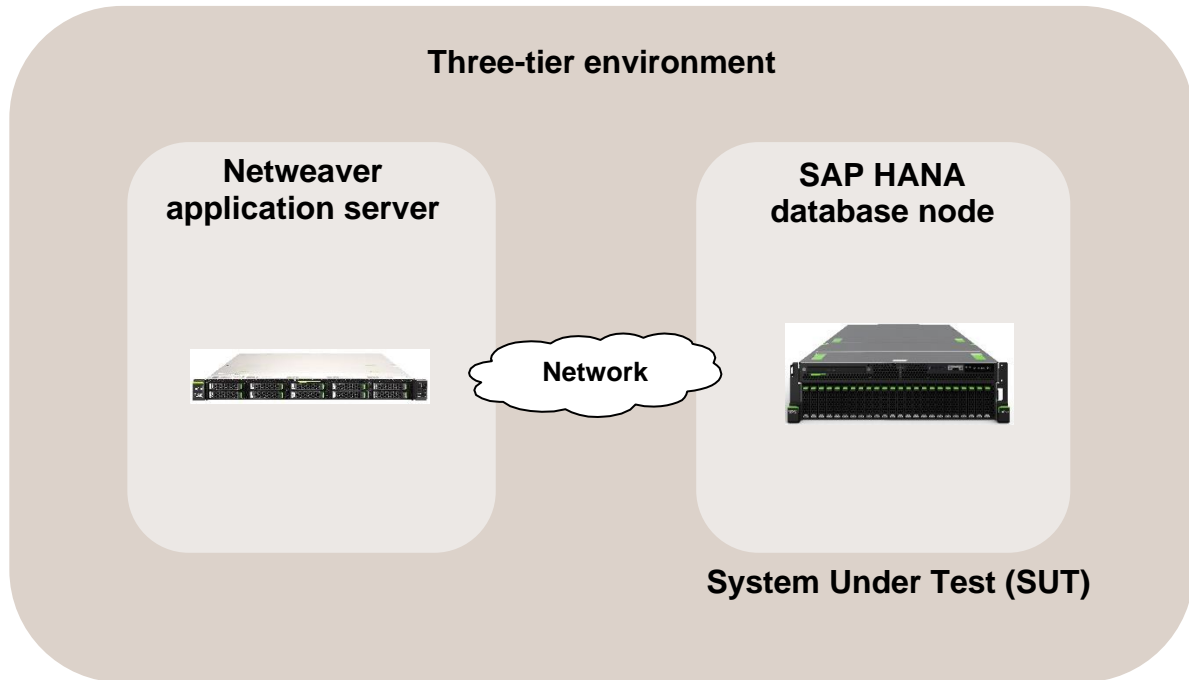
The query runtime phase takes a short time and generates a small load. Only a few processors cores are used, single thread performance is important for short runtimes.

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 RX4770 M6 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)	
<b>Hardware</b>	
Model	PRIMERGY RX4770 M6
Processor	4 × Intel Xeon Platinum 8380HL processor
Memory	<u>Configuration 1 (certification 2020039 below):</u> 24 × 32 GB 2Rx4 DDR4-3200 R ECC and 24 × 128GB Intel Optane persistent memory 200 series <u>Configuration 2 (certification 2020040 below):</u> 24 × 64 GB 2Rx4 DDR4-3200 R ECC and 24 × 128GB Intel Optane persistent memory 200 series
Network interface	1 Gbit LAN
Disk subsystem	PRIMERGY RX4770 M6: 1 × PRAID EP580i RAID Controller 1 × Internal SAS-SSD 1.6TB 2 × PRAID EP540e RAID Controller JX40 S2 × 3: 25 × JX40 S2 MLC SSD 960GB 3DWPDP
<b>Software</b>	
Operating system	SUSE Linux Enterprise Server for SAP Application 15 SP2
Database	SAP HANA 2.0

Application Server	
Hardware	
Model	PRIMERGY RX2540 M5
Processor	2 × Intel Xeon Platinum 8280 processor
Memory	768 GB
Network interface	1 Gbit LAN
Software	
Operating system	SUSE Linux Enterprise Server 12 SP4
Technology platform release	SAP Netweaver 7.50

## Benchmark results

Two SAP BWH Benchmarks were performed on PRIMERGY RX4770 M6, the first with 768 GB DRAM plus 3,072 GB Persistent Memory and 6.5 billion records, the second with 1,536 GB DRAM plus 3,072 GB Persistent Memory and 7.8 billion records.

### 6.5 Billion Records Scenario

On October 23, 2020, the following SAP BW edition for SAP HANA Standard Application Benchmark Version 3 was certified:

Certification number 2020039	
Benchmark Phase 1	
Number of initial records:	6,500,000,000
Runtime of last Data Set (seconds):	12,463
Benchmark Phase 2	
Query Executions per Hour:	5,136
CPU utilization of database server:	94%
Benchmark Phase 3	
Total Runtime of complex query phase (seconds):	96
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 RX4770 M6, 4 processors / 112 cores / 224 threads, Intel Xeon Platinum 8380HL Processor, 2.90 GHz, 64 KB L1 cache per core and 1,024KB L2 cache per core, 38.5 MB L3 cache per socket, <u>768 GB DRAM and 3,072 GB Persistent Memory</u>

The PRIMERGY RX4770 M6, equipped with Intel's 3rd Generation Intel Xeon Scalable Processor 8380HL and Intel Optane Persistent Memory 200 Series (Barlow Pass) achieved the best results in all three phases of the SAP BWH benchmark and thus set a new overall world record in the 6.5 billion initial records category (as of 2020-10-23).

Also very remarkable, that the world record was achieved with 768 GB of DRAM and 3,072 GB Persistent Memory - a ratio of 1:4.

The SAP BWH Benchmark certificate can be found [here](https://www.fujitsu.com/global/products/computing/servers/primergy/).

### 7.8 Billion Records Scenario

On October 28, 2020, the following SAP BW edition for SAP HANA Standard Application Benchmark Version 3 was certified:

Certification number 2020040	
Benchmark Phase 1	
Number of initial records:	7,800,000,000
Runtime of last Data Set (seconds):	12,752
Benchmark Phase 2	
Query Executions per Hour:	4,729
CPU utilization of database server:	95%
Benchmark Phase 3	
Total Runtime of complex query phase (seconds):	90
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 RX4770 M6, 4 processors / 112 cores / 224 threads, Intel Xeon Platinum 8380HL Processor, 2.90 GHz, 64 KB L1 cache per core and 1,024KB L2 cache per core, 38.5 MB L3 cache per socket, <u>1,536 GB DRAM and 3,072 GB Persistent Memory</u>

Right after the first world record with 6.5 billion initial records on October, 23 (see above), RX4770 M6 set another important overall scale-up world record in the next category with 7.8 billion records (as of 2020-10-28).

The memory configuration was adjusted for the higher amount of data. While 768 GB DRAM and 3,072 GB Persistent Memory (ratio 1:4) were used for the 6.5 billion records scenario, the 7.8 billion records scenario ran with 1,536 GB DRAM and 3,072 GB Persistent Memory (ratio 1:2).

The SAP BWH Benchmark certificate can be found [here](#) .

## Disk I/O: Performance of storage media

### Benchmark description

Performance measurements of disk subsystems for PRIMERGY servers are used for performance evaluation. It is possible to compare different storage connections as well. As standard, these performance measurements 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)
- Number of parallel accesses (number of outstanding I/Os)

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 "number of Outstanding I/Os" 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 [MB/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.

<i>Data throughput [MB/s]</i>	<i>= Transaction rate [IO/s] x Block size [MB]</i>
<i>Transaction rate [IO/s]</i>	<i>= Data throughput [MB/s] / Block size [MB]</i>

In this section, a power of 10 (1 TB = 10<sup>12</sup> bytes) is used to indicate the capacity of the hard storage medium, and a power of 2 (1 MB / s = 2<sup>20</sup> bytes) is used to indicate the capacity of other media, file size, block size, and throughput.

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

## Benchmark environment

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

System Under Test (SUT)			
Hardware			
Controller:		PRAID EP540i	
Storage media	Category	Drive name	
HDD	SAS HDD(SAS 12Gbps, 15krpm)[512n]	ST300MP0006 ST600MP0006	
	SAS HDD(SAS 12Gbps, 10krpm)[512e]	AL15SEB06EQ AL15SEB12EQ	
	SAS HDD(SAS 12Gbps, 10krpm)[512n]	AL15SEB060N AL15SEB120N	
	NL-SAS HDD (SAS 12Gbps, 7.2krpm)[512n]	ST2000NX0433	
	BC-SATA HDD(SATA 6Gbps, 7.2krpm)[512e]	ST1000NX0313	
	BC-SATA HDD(SATA 6Gbps, 7.2krpm)[512n]	ST1000NX0423	
SSD	SAS SSD(SAS 12Gbps, Write Intensive)	KPM51MUG400G KPM51MUG800G KPM51MUG1T60	
	SAS SSD(SAS 12Gbps, Mixed Use)	WUSTR6440ASS204 WUSTR6480ASS204 WUSTR6416ASS204 WUSTR6432ASS204	
	SAS SSD(SAS 12Gbps, Read Intensive)	WUSTR1548ASS204 WUSTR1596ASS204 WUSTR1519ASS204 WUSTR1538ASS204 WUSTR1576ASS204	
	SATA SSD(SATA 6Gbps, Mixed Use)	MTFDDAK240TDT MTFDDAK480TDT MTFDDAK960TDT MTFDDAK1T9TDT MTFDDAK3T8TDT	
	SATA SSD(SATA 6Gbps, Read Intensive)	MTFDDAK240TDS MTFDDAK480TDS MTFDDAK960TDS MTFDDAK1T9TDS MTFDDAK3T8TDS MTFDDAK7T6TDS	
	PCIe SSD(Write intensive)	SSDPE21K750GA	
	PCIe SSD(Mixed Use)	SSDPE2KE016T8 SSDPE2KE032T8 SSDPE2KE064T8	
	PCIe SSD(Read intensive)	SSDPE2KX010T8 SSDPE2KX020T8 SSDPE2KX040T8	

Controller: Intel C620 Standard SATA AHCI controller			
	Storage media	Category	Drive name
	SSD	M.2 Flash module	MTFDDAV240TDS MTFDDAV480TDS

Software		
Operating system		Microsoft Windows Server 2016 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 (other than 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

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



## Benchmark results

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

### Controller

The measurements were made using controllers in the table below.

Storage 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
M.2 Flash	C620 Standard SATA AHCI controller	-	DMI 3.0 x4	SATA 6G	-

### Storage media

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

Storage media type	Interface	Form factor
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 3.0	2.5 inch

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

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

### Cache settings

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

For the purpose of easy and reliable handling of the settings for RAID controllers and hard disks it is advisable to use the RAID-Manager software "ServerView RAID" that is supplied for 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. To avoid any confusion among the measurement units the tables have been separated for the two access types.

The table cells contain the maximum achievable values. This means that each value is the maximum achievable value of the whole range of load intensities (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

### HDDs

Random accesses (units: IO/s)




































































































Capacity [GB]	Storage device	Interface	Transactions [IO/s]		
			Database	Fileserver	Filecopy
300	ST300MP0006	SAS 12G	790	696	666
600	ST600MP0006	SAS 12G	736	651	601
600	AL15SEB06EQ	SAS 12G	670	570	603
1,200	AL15SEB12EQ	SAS 12G	734	630	637
600	AL15SEB060N	SAS 12G	684	565	584
1,200	AL15SEB120N	SAS 12G	746	609	604
2,000	ST2000NX0433	SAS 12G	492	405	388
1,000	ST1000NX0313	SATA 6G	393	340	337
1,000	ST1000NX0423	SATA 6G	399	344	346

Sequential accesses (units: MB/s)



































































Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
300	ST300MP0006	SAS 12G	304	304
600	ST600MP0006	SAS 12G	301	300
600	AL15SEB06EQ	SAS 12G	260	259
1,200	AL15SEB12EQ	SAS 12G	260	256
600	AL15SEB060N	SAS 12G	231	232
1,200	AL15SEB120N	SAS 12G	230	219
2,000	ST2000NX0433	SAS 12G	132	132
1,000	ST1000NX0313	SATA 6G	134	134
1,000	ST1000NX0423	SATA 6G	134	134

## SSDs

Random accesses (units: IO/s)

Capacity [GB]	Storage device	Interface	Transactions [IO/s]		
			Database	Fileserver	Filecopy
400	KPM51MUG400G	SAS 12G	 104,184	 14,603	 14,780
800	KPM51MUG800G	SAS 12G	 123,267	 18,333	 20,740
1,600	KPM51MUG1T60	SAS 12G	 123,420	 20,967	 22,941
400	WUSTR6440ASS204	SAS 12G	 110,903	 18,389	 15,114
800	WUSTR6480ASS204	SAS 12G	 138,943	 24,921	 24,087
1,600	WUSTR6416ASS204	SAS 12G	 137,178	 24,932	 27,138
3,200	WUSTR6432ASS204	SAS 12G	 144,367	 24,924	 27,671
6,400	WUSTR6464ASS204	SAS 12G	 148,587	 24,891	 27,847
480	WUSTR1548ASS204	SAS 12G	 90,423	 12,036	 10,782
960	WUSTR1596ASS204	SAS 12G	 126,502	 19,057	 17,422
1,920	WUSTR1519ASS204	SAS 12G	 129,456	 21,970	 22,623
3,840	WUSTR1538ASS204	SAS 12G	 134,032	 24,860	 26,851
7,680	WUSTR1576ASS204	SAS 12G	 138,973	 24,938	 27,233
240	MTFDDAK240TDT	SATA 6G	 46,406	 5,989	 6,121
480	MTFDDAK480TDT	SATA 6G	 49,138	 6,383	 6,600
960	MTFDDAK960TDT	SATA 6G	 50,488	 6,970	 7,136
1,920	MTFDDAK1T9TDT	SATA 6G	 50,669	 7,183	 7,336
3,840	MTFDDAK3T8TDT	SATA 6G	 49,490	 7,115	 7,208
240	MTFDDAK240TDS	SATA 6G	 42,594	 5,435	 5,510
480	MTFDDAK480TDS	SATA 6G	 47,577	 6,109	 6,310
960	MTFDDAK960TDS	SATA 6G	 50,134	 6,633	 6,852
1,920	MTFDDAK1T9TDS	SATA 6G	 50,638	 7,078	 7,286
3,840	MTFDDAK3T8TDS	SATA 6G	 49,542	 7,097	 7,196
7,680	MTFDDAK7T6TDS	SATA 6G	 47,200	 7,134	 7,563
750	SSDPE21K750GA	PCIe3 x4	 194,085	 37,392	 36,626
1,600	SSDPE2KE016T8	PCIe3 x4	 276,785	 45,739	 40,923
3,200	SSDPE2KE032T8	PCIe3 x4	 306,446	 53,059	 50,093
6,400	SSDPE2KE064T8	PCIe3 x4	 297,505	 56,338	 56,632
1,000	SSDPE2KX01	PCIe3 x4	 153,263	 25,891	 21,942
2,000	SSDPE2KX02	PCIe3 x4	 237,530	 38,336	 34,740
4,000	SSDPE2KX04	PCIe3 x4	 242,546	 39,242	 38,151
240	MTFDDAV240TDS	SATA 6G	 32,805	 5,482	 5,518
480	MTFDDAV480TDS	SATA 6G	 39,927	 6,384	 6,575

## Sequential accesses (units: MB/s)

Capacity [GB]	Storage device	Interface	Throughput [MB/s]	
			Streaming	Restore
400	KPM51MUG400G	SAS 12G		
800	KPM51MUG800G	SAS 12G		
1,600	KPM51MUG1T60	SAS 12G		
400	WUSTR6440ASS204	SAS 12G		
800	WUSTR6480ASS204	SAS 12G		
1,600	WUSTR6416ASS204	SAS 12G		
3,200	WUSTR6432ASS204	SAS 12G		
6,400	WUSTR6464ASS204	SAS 12G		
480	WUSTR1548ASS204	SAS 12G		
960	WUSTR1596ASS204	SAS 12G		
1,920	WUSTR1519ASS204	SAS 12G		
3,840	WUSTR1538ASS204	SAS 12G		
7,680	WUSTR1576ASS204	SAS 12G		
240	MTFDDAK240TDT	SATA 6G		
480	MTFDDAK480TDT	SATA 6G		
960	MTFDDAK960TDT	SATA 6G		
1,920	MTFDDAK1T9TDT	SATA 6G		
3,840	MTFDDAK3T8TDT	SATA 6G		
240	MTFDDAK240TDS	SATA 6G		
480	MTFDDAK480TDS	SATA 6G		
960	MTFDDAK960TDS	SATA 6G		
1,920	MTFDDAK1T9TDS	SATA 6G		
3,840	MTFDDAK3T8TDS	SATA 6G		
7,680	MTFDDAK7T6TDS	SATA 6G		
750	SSDPE21K750GA	PCIe3 x4		
1,600	SSDPE2KE016T8	PCIe3 x4		
3,200	SSDPE2KE032T8	PCIe3 x4		
6,400	SSDPE2KE064T8	PCIe3 x4		
1,000	SSDPE2KX01	PCIe3 x4		
2,000	SSDPE2KX02	PCIe3 x4		
4,000	SSDPE2KX04	PCIe3 x4		
240	MTFDDAV240TDS	SATA 6G		
480	MTFDDAV480TDS	SATA 6G		

## OLTP-2

### Benchmark description

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

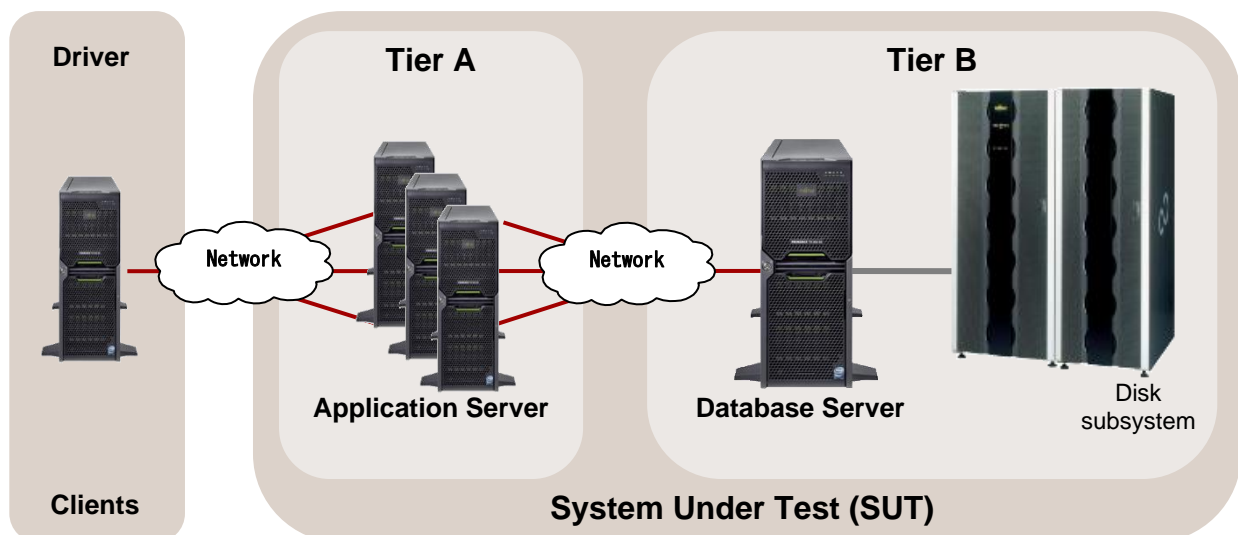
In contrast to benchmarks such as SPEC CPU and TPC-E, which were standardized by independent bodies and for which adherence to the respective rules and regulations are monitored, OLTP-2 is an internal benchmark of Fujitsu. OLTP-2 is based on the well-known database benchmark TPC-E. OLTP-2 was designed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

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

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

### Benchmark environment

The typical measurement set-up is illustrated below.



All OLTP-2 results were Calculated based on the configuration of the following pages of PRIMERGY RX4770 M6.

Database Server (Tier B)	
<b>Hardware</b>	
Model	PRIMERGY RX4770 M6
Processor	3rd Generation Intel Xeon Processor Scalable Family
Memory	2 processors: 24 x 64 GB (1x64 GB) 2Rx4 DDR4-3200 ECC 4 processors: 48 x 64 GB (1x64 GB) 2Rx4 DDR4-3200 ECC
Network interface	2 x Dual port LAN 10 Gbps 1 x Quad port onboard LAN 1 Gbps
Disk subsystem	RX4770 M6: RAID controller PRAID EP540i 6 x 1.6 TB SSD drive, RAID10 (LOG), 5 x RAID controller PRAID EP540e 10 x JX40 S2: 4 x 1.6 TB SSD drive, RAID10 (temp), 120 x 1.6 TB SSD drive, RAID5 (data)
<b>Software</b>	
BIOS	Version R1.5.0
Operating system	Microsoft Windows Server 2016 Standard
Database	Microsoft SQL Server 2017 Enterprise + KB4341265
Application Server (Tier A)	
<b>Hardware</b>	
Model	1 x PRIMERGY RX2530 M4
Processor	2 x Xeon Platinum 8180
Memory	192 GB, 2666 MHz Registered ECC DDR4
Network interface	2 x Dual port LAN 10 Gbps 1 x Dual port onboard LAN 1 Gbps
Disk subsystem	2 x 300 GB 10k rpm SAS drive
<b>Software</b>	
Operating system	Microsoft Windows Server 2016 Standard
Clients	
<b>Hardware</b>	
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
<b>Software</b>	
Operating system	Microsoft Windows Server 2012 R2 Standard
Benchmark	OLTP-2 Software EGen version 1.14.0

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

## Benchmark results

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


A guideline in the database environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses. This why a configuration with a total memory of 3072 GB was considered for the measurements with four processors and a configuration with a total memory of 1536 GB for the measurements with two processors. Both memory configurations have memory access of 3200 MHz.

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

Processor	Number of cores	Number of threads	4CPU Score	2CPU Score
Xeon Platinum 8380HL	28	56	12,410 (est.)	6,963 (est.)
Xeon Platinum 8380H	28	56	12,410 (est.)	6,963 (est.)
Xeon Platinum 8376HL	28	56	11,456 (est.)	6,428 (est.)
Xeon Platinum 8376H	28	56	11,456 (est.)	6,428 (est.)
Xeon Platinum 8360HL	24	48	10,885 (est.)	6,107 (est.)
Xeon Platinum 8360H	24	48	10,885 (est.)	6,107 (est.)
Xeon Platinum 8356H	8	16	5,148 (est.)	2,888 (est.)
Xeon Platinum 8354H	18	36	9,017 (est.)	4,954 (est.)
Xeon Gold 6348H	24	48	9,785 (est.)	5,490 (est.)
Xeon Gold 6330H	24	48	9,278 (est.)	5,206 (est.)
Xeon Gold 6328HL	16	32	7,878 (est.)	4,420 (est.)
Xeon Gold 6328H	16	32	7,878 (est.)	4,420 (est.)
Xeon Gold 5320H	20	40	9,017 (est.)	5,059 (est.)
Xeon Gold 5318H	18	36	8,233 (est.)	4,619 (est.)



The following graph shows the OLTP-2 transaction rate obtained with the Intel Xeon Processor Scalable Family (2 or 4).




It is evident that a wide performance range is covered by the variety of released processors. If you compare the OLTP-2 value of the processor with the lowest performance (Xeon Platinum 8356H) with the value of the processor with the highest performance (Xeon Platinum 8380H) the OLTP-2 value increased by a factor of 2.4. The features of the processors are summarized in the section "Technical data."

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

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

The highest value for OLTP-2 on the current PRIMERGY model is about 3% higher than the highest value on the previous model.



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

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

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

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

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

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

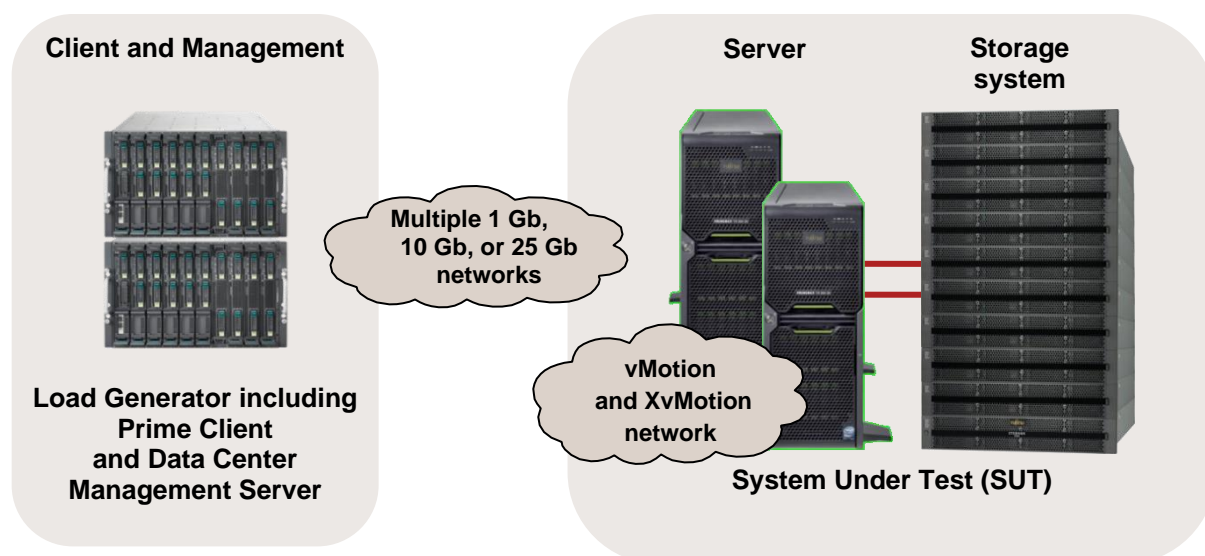
In the case of the two test types "Performance with Server Power" and "Performance with Server and Storage Power," a so-called "Server PPKW Score" and "Server and Storage PPKW Score" are determined. These are the performance scores divided by the average power consumption in kilowatts (PPKW = performance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

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

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

## Benchmark environment

The typical measurement set-up is illustrated below.



System Under Test (SUT)	
<b>Hardware</b>	
Number of servers	2
Model	PRIMERGY RX4770 M6
Processor	4 x Intel Xeon Platinum 8380HL
Memory	3072 GB : 48 x 64 GB (1x64 GB) 2Rx4 DDR4-3200 R ECC
Network interface	2 x Mellanox MCX4121A-ACAT dual port 25Gb SFP28 PCIe adapter 1 x Intel I350-T2 1Gb quad port OCPv3
Disk subsystem	2 x Qlogic QLE2772 dual port 32Gb PCIe adapter 7 x PRIMERGY RX2540 M4 & M5 configured as Fiber Channel targets 4 x PRIMERGY RX2540 M4: 2 x Micron MTFDDAK480TDC SATA SSD (480 GB, RAID1) 3 x Intel P4800X PCIe SSD (750 GB) 1 x Intel P4600 PCIe SSD (4 TB) 1 x PRIMERGY RX2540 M4: 1 x Micron MTFDDAK480TDC SATA SSD (480 GB) 3 x Intel P4800X PCIe SSD (750 GB) 1 x Intel P4600 PCIe SSD (2 TB) 1 x PRIMERGY RX2540 M5: 1 x Samsung MZ7KH480HAHQ SATA SSD (480 GB) 3 x Intel P4800X PCIe SSD (750 GB) 3 x Intel P4610 PCIe SSD (3.2 TB) 1 x PRIMERGY RX2540 M5: 1 x Samsung MZ7KH480HAHQ SATA SSD (480 GB) 2 x Intel P4800X PCIe SSD (750 GB) 2 x Intel P4610 PCIe SSD (3.2 TB)
<b>Software</b>	
BIOS	V1.0.0.0 R1.1.0 for D33892-A1x
BIOS settings	See "Details"
Operating system	VMware ESXi 7.0 U1, Build 16850804

Operating system settings	ESX settings: see "Details"
---------------------------	-----------------------------

Details	
Public URLs	<a href="https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2020-10-22-Fujitsu-RX4770M6.pdf">https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2020-10-22-Fujitsu-RX4770M6.pdf</a> <a href="https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2020-10-22-Fujitsu-RX4770M6-serverPPKW.pdf">https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2020-10-22-Fujitsu-RX4770M6-serverPPKW.pdf</a> <a href="https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2020-10-22-Fujitsu-RX4770M6-serverstoragePPKW.pdf">https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2020-10-22-Fujitsu-RX4770M6-serverstoragePPKW.pdf</a>

Datacenter Management Server (DMS)	
Hardware	
Model	1 x PRIMERGY RX2540 M2
Processor	1 x Intel Xeon E5-2698 v4
Memory	64 GB
Network interface	1 x Emulex One Connect Oce14000 1GbE dual port PCIe adapter
Software	
Operating system	VMware ESXi 6.7 EP 02a Build 9214924
Datacenter Management Server (DMS) VM	
Hardware	
Processor	4 x Logical CPU
Memory	19 GB
Network interface	1 x 1 Gbit/s LAN
Software	
Operating system	VMware vCenter Server Appliance 7.0 U1 Build 16860138

Load generator	
Hardware	
Model	6 x PRIMERGY RX2530 M2
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 6.7 EP 08 Build 13473784

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


## Benchmark results

### "Performance Only" measurement results (October 22, 2020)



On October 22, 2020, Fujitsu achieved a VMmark V3.1.1 score of "19.72@20 tiles" using PRIMERGY RX4770 M6 with a Xeon Platinum 8380HL processor and VMware ESXi 7.0 U1. At this time, the system configuration had a total of 2 x 112 processor cores, and two identical servers were used for the "System Under Test" (SUT). Based on the above results, PRIMERGY RX4770 M6 is rated as the most powerful 4-socket 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 October 22, 2020. 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 PRIMERGY RX4770 M6 and the previous generation PRIMERGY RX4770 M5. The PRIMERGY RX4770 M6 achieved a 14% improvement in score compared to the previous generation PRIMERGY RX4770 M5. This is due to the improved performance of the 3rd generation Intel Xeon scalable processor and the effective use of the capabilities of the VMware ESXi hypervisor.

### "Performance with Server Power" measurement results (October 22, 2020)

### "Performance with Server and Storage Power" measurement results (October 22, 2020)



On October 22, 2020, Fujitsu achieved a VMmark V3.1.1 "Server PPKW" score of "6.3410@20 tiles" using PRIMERGY RX4770 M6 with a Xeon Platinum 8380HL processor and VMware ESXi 7.0 U1. At the same time, also achieved a VMmark V3.1.1 "Server and Storage PPKW" score of "3.7609@20 tiles." These were system configurations with a total of 2 x 112 processor cores, and two identical servers were used for the "System Under Test" (SUT). Based on the above results, PRIMERGY RX4770 M6 is rated as the most energy efficient 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>.

VMmark is a product of VMware, Inc.

# 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)	
<b>Hardware</b>	
Model	PRIMERGY RX4770 M6
Processor	4 x 3rd Generation Intel Xeon Scalable Processors Family
Memory	48 x 32 GB (1x32 GB) 2Rx4 PC4-3200AA-R
<b>Software</b>	
BIOS settings	SNC = Enable SNC2 Override OS Energy Performance = Enabled Energy Performance = Performance HWPM Support = Disabled Intel Virtualization Technology = Disabled LLC Dead Line Alloc = Disabled Stale AtoS = Enabled
Operating system	Red Hat Enterprise Linux Server release 8.2 4.18.0-193.el8.x86_64
Operating system settings	Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1) Transparent Huge Pages inactivated sched_cfs_bandwidth_slice_us = 50000 sched_migration_cost_ns = 5000000 cpupower -c all frequency-set -g performance cpupower idle-set -d 1 cpupower idle-set -d 2 cpupower idle-set -d 3 echo 0 > /proc/sys/kernel/numa_balancing echo 1 > /proc/sys/vm/drop_caches ulimit -s unlimited
Compiler	C/C++: Version 19.1.1.217 of Intel C/C++ Compiler for Linux
Benchmark	STREAM Version 5.10

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



## Benchmark results


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

Processor	Memory frequency [MHz]	Maximum memory bandwidth *1 [GB/s]	Number of cores	Processor frequency [GHz]	Number of processors	TRIAD [GB/s]
Xeon Platinum 8380H	3200	154	28	2.9	4	475 (est.)
Xeon Platinum 8380HL	3200	154	28	2.9	4	<b>475</b>
Xeon Platinum 8376H	3200	154	28	2.6	4	476 (est.)
Xeon Platinum 8376HL	3200	154	28	2.6	4	476 (est.)
Xeon Platinum 8360H	3200	154	24	3.0	4	<b>464</b>
Xeon Platinum 8360HL	3200	154	24	3.0	4	464 (est.)
Xeon Platinum 8356H	3200	154	8	3.9	4	<b>369</b>
Xeon Platinum 8354H	3200	154	18	3.1	4	<b>424</b>
Xeon Gold 6348H	3200	141	24	1.9	4	<b>429</b>
Xeon Gold 6330H	3200	141	24	3.1	4	<b>436</b>
Xeon Gold 6328H	3200	141	16	2.1	4	406 (est.)
Xeon Gold 6328HL	3200	141	16	2.5	4	<b>406</b>
Xeon Gold 5320H	3200	128	20	3.3	4	<b>386</b>
Xeon Gold 5318H	3200	128	18	3.6	4	<b>389</b>

\*1: Value per Processor

The following diagram illustrates the throughput of the RX4770 M6 in comparison to its predecessor, the RX4770 M5.

### STREAM TRIAD: Comparison of PRIMERGY RX4770 M6 and PRIMERGY RX4770 M5



# 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  $n \times n$  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  $n \times n$  matrix the number of arithmetic operations required for the solution is  $\frac{2}{3}n^3 + 2n^2$ . 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. The processor performance is the decisive factor for the measurement result. Since the algorithm used permits parallel processing, in particular the number of processors used and their processor cores are - in addition to the clock rate - of outstanding significance.

LINPACK is used to measure how many floating point operations were carried out per second. The result is referred to as **Rmax** and specified in GFlops (Giga Floating Point Operations per Second: 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.

$$R_{peak} = \text{Maximum number of floating point operations per clock cycle} \\ \times \text{Number of processor cores of the computer} \\ \times \text{Rated processor frequency [GHz]}$$

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/content/www/us/en/develop/articles/intel-mkl-benchmarks-suite.html>.

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

## Benchmark environment

System Under Test (SUT)	
<b>Hardware</b>	
Model	PRIMERGY RX4770 M6
Processor	4 x 3rd Generation Intel Xeon Scalable Processors Family
Memory	48 x 32 GB (1x32 GB) 2Rx4 PC4-3200AA-R
<b>Software</b>	
BIOS settings	HyperThreading = Disabled Link Frequency Select = 10.4 GT/s HWPM Support = Disabled Intel Virtualization Technology = Disabled LLC Dead Line Alloc = Disabled Stale AtoS = Enabled Fan Control = Full
Operating system	Red Hat Enterprise Linux Server release 8.2 4.18.0-193.el8.x86_64
Operating system settings	Kernel Boot Parameter set with : nohz_full=1-X (X: logical core number -1) cpupower -c all frequency-set -g performance run with avx512
Compiler	C/C++: Version 19.1.1.217 of Intel C/C++ Compiler for Linux
Benchmark	Intel Optimized MP LINPACK Benchmark for Clusters

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

## Benchmark results

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

Processor	Number of cores	Processor frequency [GHz]	Number of processors	Rpeak [GFlops]	Rmax [GFlops]	Efficiency
Xeon Platinum 8380H	28	2.9	4	10,394	6,929 (est.)	67%
Xeon Platinum 8380HL	28	2.9	4	10,394	<b>6,929</b>	67%
Xeon Platinum 8376H	28	2.6	4	9,318	6,009 (est.)	64%
Xeon Platinum 8376HL	28	2.6	4	9,318	<b>6,009</b>	64%
Xeon Platinum 8360H	24	3.0	4	9,216	<b>5,949</b>	65%
Xeon Platinum 8360HL	24	3.0	4	9,216	5,949 (est.)	65%
Xeon Platinum 8356H	8	3.9	4	3,994	<b>2,992</b>	75%
Xeon Platinum 8354H	18	3.1	4	7,142	<b>4,825</b>	68%
Xeon Gold 6348H	24	1.9	4	7,066	<b>4,534</b>	64%
Xeon Gold 6330H	18	3.1	4	6,144	<b>4,020</b>	65%
Xeon Gold 6328H	24	2.1	4	5,734	<b>3,929</b>	69%
Xeon Gold 6328HL	20	2.5	4	5,734	3,929 (est.)	69%
Xeon Gold 5320H	12	3.3	4	6,144	<b>4,039</b>	66%
Xeon Gold 5318H	8	3.6	4	5,760	<b>3,757</b>	65%

*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.


## Literature


### PRIMERGY Servers

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

### PRIMERGY RX4770 M6

This White Paper:

 <https://docs.ts.fujitsu.com/dl.aspx?id=1924db9e-fb0e-4ee0-8608-da78acb9fd5f>

 <https://docs.ts.fujitsu.com/dl.aspx?id=b4bf88e9-201c-4408-b7c2-7fd4e3f2b61c>

Data sheet

<https://docs.ts.fujitsu.com/dl.aspx?id=c73aa330-ea96-4e4a-bfcc-5ac79a09ed7d>

### PRIMERGY Performance

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

### SPECcpu2017

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

Benchmark Overview SPECcpu2017

<https://docs.ts.fujitsu.com/dl.aspx?id=0f641c7e-bb5e-45e4-854f-cdd31faf5343>

### SPECpower\_ssj2008

[https://www.spec.org/power\\_ssj2008/](https://www.spec.org/power_ssj2008/)

Benchmark Overview SPECpower\_ssj2008

<https://docs.ts.fujitsu.com/dl.aspx?id=a133cf86-63be-4b5a-8b0f-a27621c8d3c5>

### SPECjbb2015

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

### SAP SD

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

<https://www.sap.com/dmc/exp/2018-benchmark-directory/#/sd>

Benchmark overview SAP SD

<https://docs.ts.fujitsu.com/dl.aspx?id=ab13a8c0-44d8-40ee-9415-695d372e2e7b>

### SAP BWH

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

<https://www.sap.com/dmc/exp/2018-benchmark-directory/#/bwh>

Benchmark overview SAP BWH

<https://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=9775e8b9-d222-49db-98b1-4796fbc6d7a>

### vServCon

Benchmark Overview vServCon

<https://docs.ts.fujitsu.com/dl.aspx?id=c3d5ce5d-5610-43c6-86b4-051549940a71>

### VMmark V3

VMmark

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

**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://software.intel.com/content/www/us/en/develop/articles/intel-mkl-benchmarks-suite.html>

## Contact

### FUJITSU

Website: <https://www.fujitsu.com/>

### PRIMERGY Product Marketing

<mailto:Primergy-PM@ts.fujitsu.com>

### PRIMERGY Performance and Benchmarks

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