Fujitsu Server PRIMERGY
Performance Report
PRIMERGY RX4770 M7

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMERGY RX4770 M7.

Explaines PRIMERGY RX4770 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
1.1 2024-01-12

Contents

Technical data

PRIMERGY RX4770 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 RX4770 M7
Form factor	Rack server
Chipset	Intel C741
Number of sockets	4
Number of configurable processors	2 or 4
Processor type	4th Generation Intel Xeon Scalable Processors Family
Number of memory slots	64 (16 per processor)
Maximum memory configuration	4,096 GB
Maximum number of internal storage disks	2.5 inch: 24
Maximum number of PCI slots	PCI-Express 5.0 (x16 lane): 4 (Low Profile) PCI-Express 5.0 (x8 lane): 2 (Low Profile) PCI-Express 4.0 or 5.0 (x16 lane): 4 (Full Height)

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]
Xeon Platinum 8490H	XCC	60	120	112.5	16	1.90	3.50	4,800	350
Xeon Platinum 8468H	XCC	48	96	105	16	2.10	3.80	4,800	330
Xeon Platinum 8460H	XCC	40	80	105	16	2.20	3.80	4,800	330
Xeon Platinum 8454H	XCC	32	64	82.5	16	2.10	3.40	4,800	270
Xeon Platinum 8450H	XCC	28	56	75	16	2.00	3.50	4,800	250
Xeon Platinum 8444H	XCC	16	32	45	16	2.90	4.00	4,800	270
Xeon Gold 6448H	MCC	32	64	60	16	2.40	4.10	4,800	250
Xeon Gold 6434H	MCC	8	16	22.5	16	3.70	4.10	4,800	195
Xeon Gold 6418H	MCC	24	48	60	16	2.10	4.00	4,800	185
Xeon Gold 6416H	MCC	18	36	45	16	2.20	4.20	4,800	165

All processors that can be ordered with PRIMERGY RX4770 M7 support Intel Turbo Boost Technology 2.0.

This technology allows you to operate the processor with higher frequencies than the rated frequency. The "maximum turbo frequency" listed in the processor list above is the theoretical maximum frequency when there is only one active core per processor. The maximum frequency that can actually be achieved depends on the number of active cores, current consumption, power consumption, and processor temperature.

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

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

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

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

Suffix	Workload
H	**DB/Analytics** Data analytics and big data usages

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	1,600W platinum PSU	3
	1,600W titanium PSU	3
	2,200W platinum PSU	3
	2,400W titanium PSU	3

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 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 RX4770 M7
• Processor	4 x 4th Generation Intel Xeon Scalable Processors Family
• Memory	32 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC
Software	
• BIOS settings	SPECspeed2017_int_base: • Override OS Energy Performance = Enabled • Energy Performance = Balanced Performance • SNC(Sub NUMA) = Enable SNC2 (Disabled when MCC are installed) • CPU Performance Boost = Aggressive • FAN Control = Full SPECSpeed2017_fp_base: • Hyper Threading = Disabled • Package C State limit = C0 • LLC Prefetch = Enabled • Homeless Prefetch = Enabled • DBP-F = Enabled • CPU Performance Boost = Aggressive • FAN Control = Full SPECrate2017_int_base: • DCU Streamer Prefetcher = Disabled • Package C State limit = C0 • LLC Dead Line Alloc = Disabled • CPU Performance Boost = Aggressive • SNC(Sub NUMA) =Enable SNC4(Enable SNC2 when MCC are installed) • FAN Control = Full SPECrate2017_fp_base: • Hyper Threading = Disabled (Enabled when MCC are installed) • Package C State limit = C0 • CPU Performance Boost = Aggressive • SNC (Sub NUMA) =Enable SNC4 (Enable SNC2 when MCC are installed) • FAN Control = Full
• Operating system	SUSE Linux Enterprise Server 15 SP4 5.14.21-150400.22-default
• Operating system settings	Default
• Compiler	C/C++: Version 2023.0 of Intel C/C++ Compiler for Linux Fortran: Version 2023.0 of Intel Fortran Compiler for Linux

Benchmark results

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

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

Processor model	Number of cores	Number of processors	SPECrate2017_int_base	SPECrate2017_fp_base
Xeon Platinum 8490H	60	4	**1,890**	**1,890**
Xeon Platinum 8468H	48	4	1,650 est.	1,730 est.
Xeon Platinum 8460H	40	4	1,490 est.	1,610 est.
Xeon Platinum 8454H	32	4	1,100 est.	1,350 est.
Xeon Platinum 8450H	28	4	942 est.	1,210 est.
Xeon Platinum 8444H	16	4	649 est.	917 est.
Xeon Gold 6448H	32	4	1,230 est.	1,420 est.
Xeon Gold 6434H	8	4	399 est.	570 est.
Xeon Gold 6418H	24	4	886 est.	1,130 est.
Xeon Gold 6416H	18	4	669 est.	902 est.

Processor model	Number of cores	Number of processors	SPECspeed2017_int_base	SPECspeed2017_fp_base
Xeon Platinum 8490H	60	4	-	**385**
Xeon Gold 6416H	16	4	**16.0**	-

The following graphs compare the throughputs of PRIMERGY RX4770 M7 and their older models, PRIMERGY RX4770 M6, with maximum performance configurations.
Showed significant performance improvement over the previous generation.

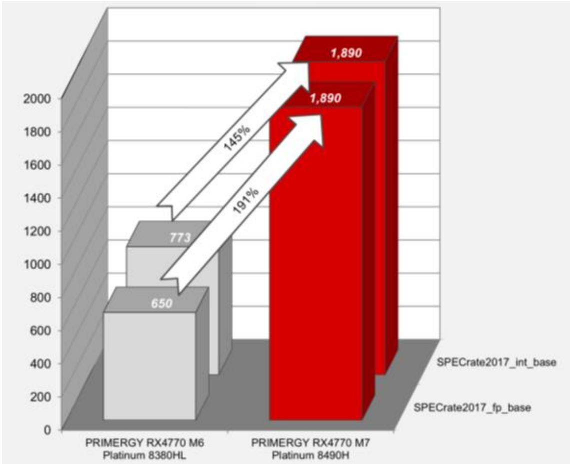

SPECrate2017: Comparison of PRIMERGY RX4770 M6 and PRIMERGY RX4770 M7

STREAM

Benchmark description

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

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

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

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

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

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

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

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

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

Benchmark environment

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY RX4770 M7
• Processor	4 x 4th Generation Intel Xeon Scalable Processors Family
• Memory	32 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC
Software	
• BIOS settings	• DCU Streamer Prefetcher = Disabled • SNC(Sub NUMA) = Enable SNC4 (Enable SNC2 when MCC type installed)) • Intel Virtualization Technology = Disabled • LLC Dead Line Alloc = Disabled • Stale Atos = Enabled
• Operating system	SUSE Linux Enterprise Server 15 SP4 5.14.21-150400.22-default
• Operating system settings	Default
• Compiler	C/C++: Version 2023.0 of Intel C/C++ Compiler for Linux
• Benchmark	STREAM Version 5.10

Benchmark results

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

Processor	Memory frequency [MHz]	Maximum memory bandwidth [GB/s]	Number of cores	Rated frequency [GHz]	Number of processors	TRIAD [GB/s]
Xeon Platinum 8490H	4,800	307	60	1.90	4	**908**
Xeon Platinum 8468H	4,800	307	48	2.00	4	**895**
Xeon Platinum 8460H	4,800	307	40	2.00	4	**847**
Xeon Platinum 8454H	4,800	307	32	2.40	4	**821**
Xeon Platinum 8450H	4,800	307	28	2.10	4	**827**
Xeon Platinum 8444H	4,800	307	16	2.80	4	**650**
Xeon Gold 6448H	4,800	307	32	2.40	4	**840**
Xeon Gold 6434H	4,800	307	8	3.70	4	**808**
Xeon Gold 6418H	4,800	307	24	2.10	4	**667**
Xeon Gold 6416H	4,800	307	18	2.20	4	**446**

The following graphs compare the throughputs of PRIMERGY RX4770 M7 and their older models, PRIMERGY RX 4770 M6, with maximum performance configurations.

Showed significant performance improvement over the previous generation.

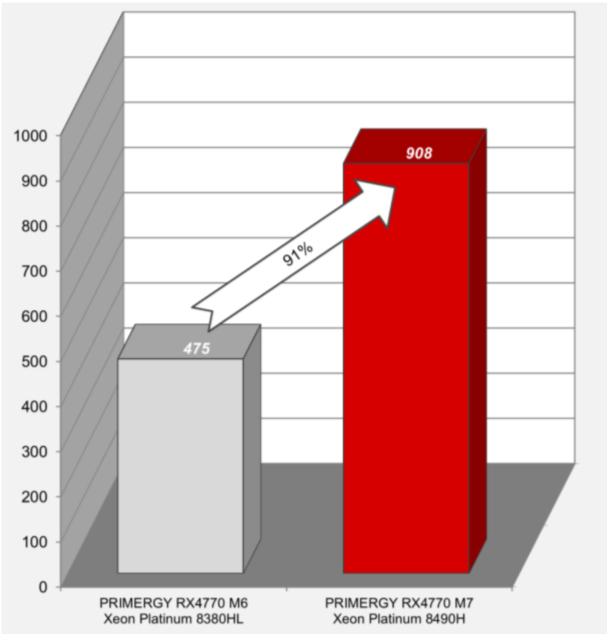

STREAM: Comparison of PRIMERGY RX4770 M6 and PRIMERGY RX4770 M7

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

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY RX4770 M7
• Processor	4 x 4th Generation Intel Xeon Scalable Processors Family
• Memory	32 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC
Software	
• BIOS settings	• HyperThreading = Disabled • CPU Performance Boost = Agressive • Fan Control = Full
• 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	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]	Rmax [GFlops]	Effic.
Xeon Platinum 8490H	60	1.90	4	14,592	**14,622**	100%
Xeon Platinum 8468H	48	2.10	4	12,902	**12,491**	97%
Xeon Platinum 8460H	40	2.20	4	11,264	**11,915**	106%
Xeon Platinum 8454H	32	2.10	4	8,602	**9,019**	105%
Xeon Platinum 8450H	28	2.00	4	7,168	**7,744**	108%
Xeon Platinum 8444H	16	2.90	4	5,939	**6,034**	102%
Xeon Gold 6448H	32	2.40	4	9,830	**10,051**	102%
Xeon Gold 6434H	8	3.70	4	3,789	**3,676**	97%
Xeon Gold 6418H	24	2.10	4	6,451	**7,002**	109%
Xeon Gold 6416H	18	2.20	4	5,069	**5,537**	109%

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 RX4770 M7 and their older models, PRIMERGY RX4770 M6, with maximum performance configurations.
Both models showed significant performance improvements over the previous generation.

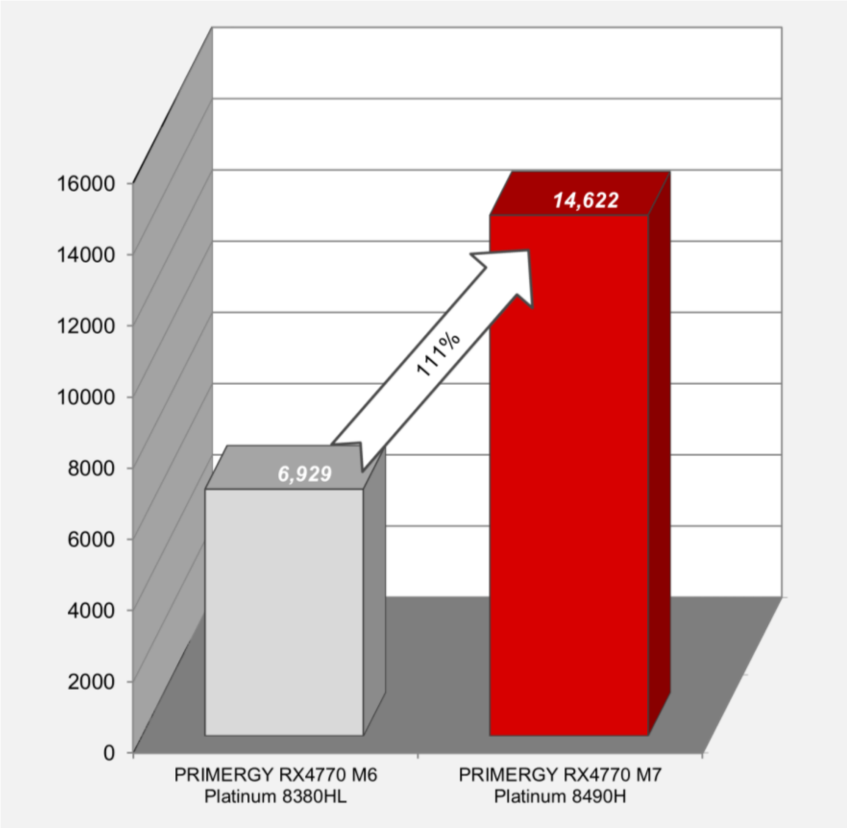

LINPACK: Comparison of PRIMERGY RX4770 M6 and PRIMERGY RX4770 M7

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.

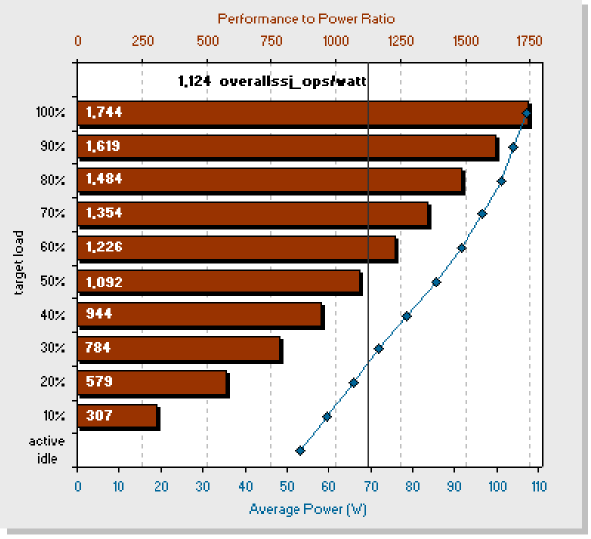

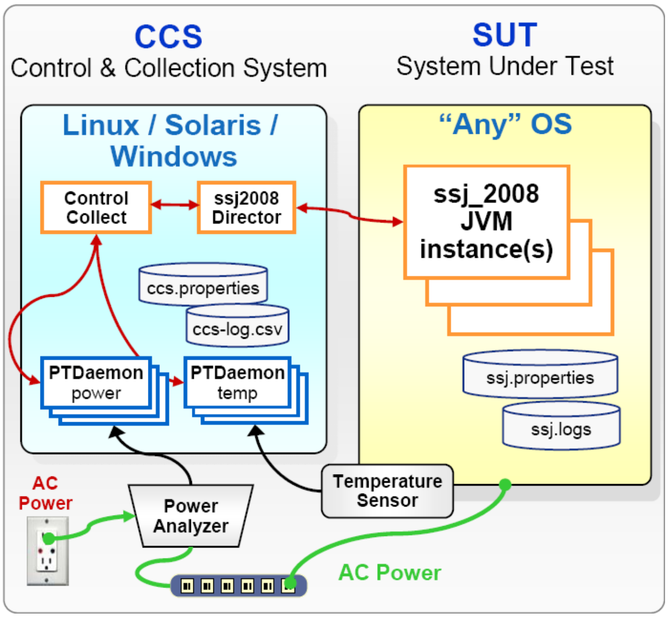

The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System 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 RX4770 M7
• Processor	4 x Xeon Platinum 8490H 60C 1.90GHz 350W
• Memory	32 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 1,600W platinum PSU

Software

• BIOS settings	ASPM Support = Auto Hardware Prefetcher = Disabled Adjacent Cache Line Prefetch = Disabled DCU Streamer Prefetcher = Disabled Intel(R) VT-d = Disabled Package C State limit = No limit Uncore Frequency Scaling = Power balanced CPU Performance Boost = Aggressive SNC(Sub NUMA) = Enable SNC4 SATA Controller = Disabled USB Port Control = Disable all ports Serial Port = Disabled Network Stack = Disabled
• Operating system	Windows Server 2022 Standard
• Operating system settings	Turn off hard disk after = 1 Minute PCI Express Link State Power Management = Maximum power savings Minimum processor state = 0% Maximum processor state = 100% Turn off display after = 1 Minute POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFBOOSTMODE 4 POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFINCTHRESHOLD 95 POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFDECTHRESHOLD 93 POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR PERFDECTIME 1 POWERCFG /SETACVALUEINDEX SCHEME_CURRENT SUB_PROCESSOR IDLESCALING 1 POWERCFG /S SCHEME_CURRENT Using the local security settings console, "lock pages in memory" was enabled for the user running the benchmark. Benchmark was started via Windows Remote Desktop Connection.
• JVM	Oracle Java HotSpot(TM) 64-Bit Server VM 18.9 (build 11.0.16.1+1-LTS, mixed mode)
• JVM settings	-server -Xmn1500m -Xms1625m -Xmx1625m -XX:+UseLargePages -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:MaxTenuringThreshold=15 -XX:ParallelGCThreads=2 -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:-UseAdaptiveSizePolicy -XX:+UseParallelOldGC -XX:FreqInlineSize=2500 -XX:LoopUnrollLimit=45 -XX:InitialTenuringThreshold=12 -XX:-ThreadLocalHandshakes -XX:UseAVX=0

Benchmark results

The PRIMERGY RX4770 M7 in Microsoft Windows Server 2022 Standard achieved the following result:

SPECpower_ssj2008 = 13,261 overall ssj_ops/watt

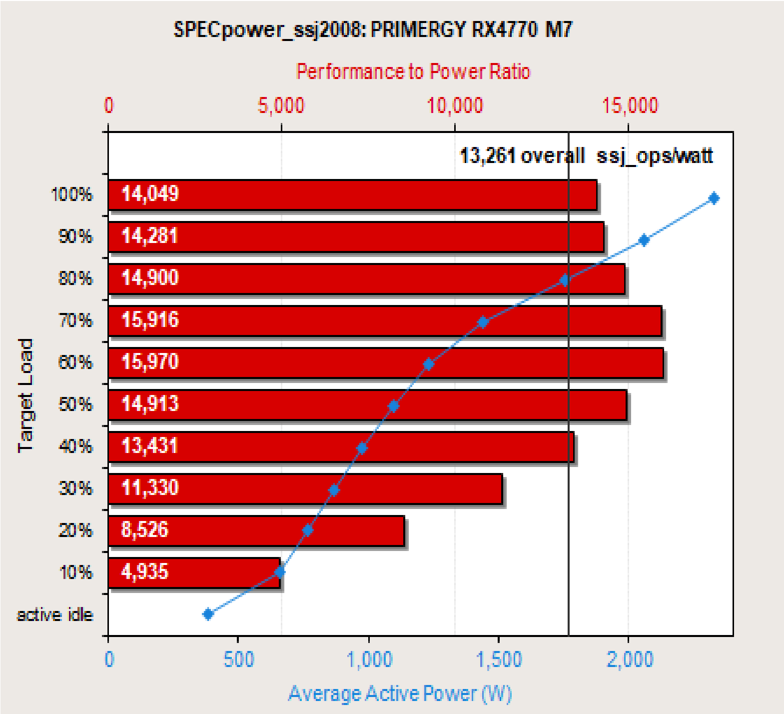

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 13,261 overall ssj_ops/watt for the PRIMERGY RX4770 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
			∑ ssj_ops / ∑ power = 13,704

Comparison with the predecessor

The following diagram shows for each load level (on the x-axis) the throughput (on the left y-axis) and the power consumption (on the right y-axis) of the PRIMERGY RX4770 M7 compared to the predecessor PRIMERGY RX4770 M6.

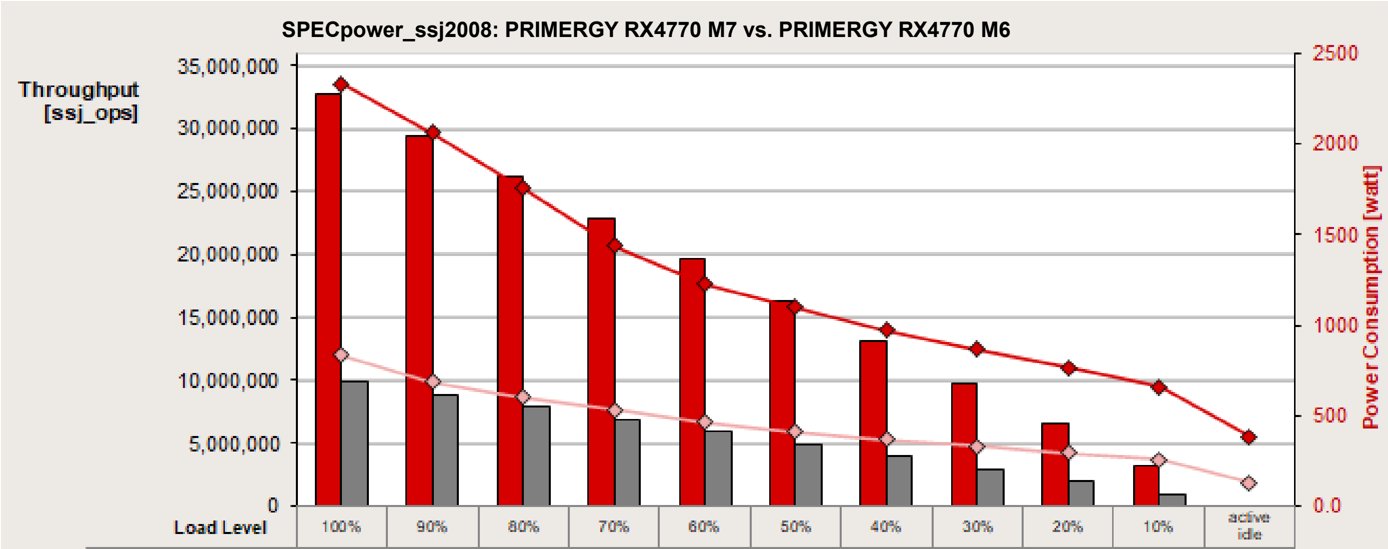

Load Level	100%	90%	80%	70%	60%	50%	40%	30%	20%	10%	active idle
RX4770 M7 ssj_ops	32,704,869	29,437,266	26,146,485	22,870,786	19,618,145	16,353,286	13,068,465	9,796,017	6,539,100	3,269,387	0
RX4770 M6 ssj_ops	9,904,110	8,914,839	7,933,444	6,944,695	5,948,299	4,958,147	3,966,282	2,973,866	1,980,442	990,995	0
RX4770 M7 watt	2328	2061	1755	1437	1228	1097	973	865	767	662	386
RX4770 M6 watt	834	688	603	532	467	413	370	333	296	257	129

Thanks to the 4th Generation Intel Xeon Scalable Processors Family, the PRIMERGY RX4770 M7 has a higher throughput. This results in an overall 19.7% increase in energy efficiency in the PRIMERGY RX4770 M7.

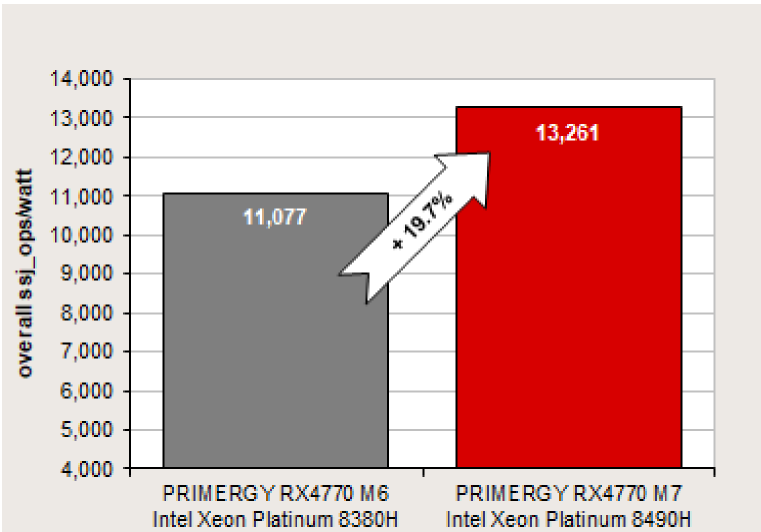

Measurement results of SPECpower_ssj2008 (June 7, 2023)

13,261 SPECpower_ssj2008

On June 7, 2023, PRIMERGY RX4770 M7 with four Xeon Platinum 8490H processors achieved a performance value of 13,261 on the Windows Server 2022 Standard in the SPECpower_ssj2008 benchmark, in the Windows division of the 4th 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/.

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.

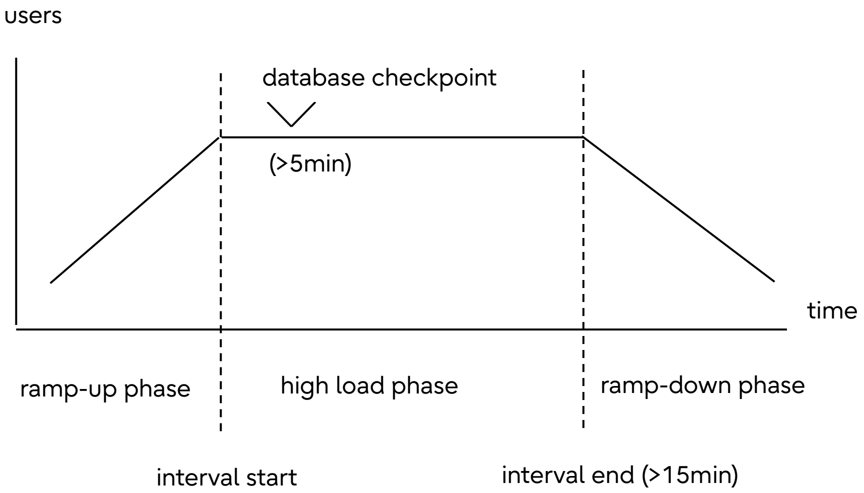

A benchmark can only be certified if the average dialog response time is less than 1 second.
Certified and published SAP SD Benchmarks are published on SAP's benchmark site here.

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 M7 was performed on a two-tier configuration.

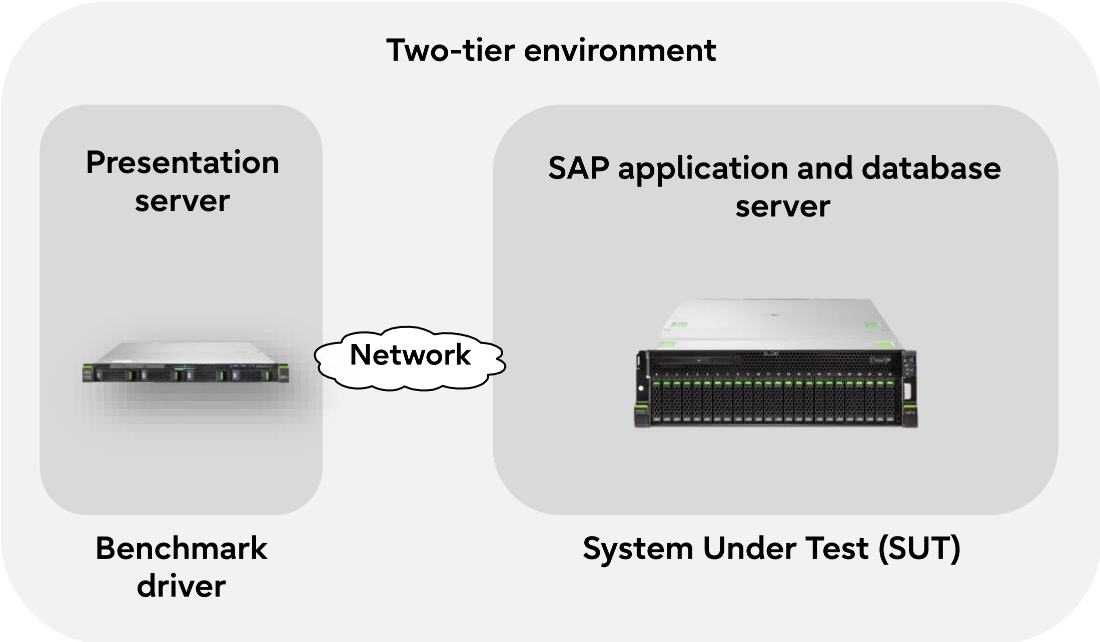

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY RX4770 M7
• Processor	4 x Intel Xeon Platinum 8490H processor 60C 1.9GHz 350W
• Memory	32 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC
• Network interface	PLAN EP X710-DA2 2x10Gb Adapter
• Storage subsystem	1 × PRAID EP640i LP internal RAID controller 1 x SSD SAS 2.5" Mixed Use 1.6TB 1 x PCIe-SSD 2.5" Mixed Use 3.2TB 1 x PRAID EP680e LP external RAID controller 1 x ETERNUS JX40 S2 Enclosure 6 × JX40 S2 TLC SSD 960GB 3DWPD
Software	
• Operating system	Windows Server 2019
• Database	Microsoft SQL Server 2019
• 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	PLAN EP X710-DA2 2x10Gb Adapter

Software

• Operating System	SUSE Linux Enterprise Server 12 SP2

Benchmark results

On November 27, 2023, the following SAP Sales and Distribution (SD) Standard Application Benchmark was certified:

Certification number 2023072	
• Number of SAP SD benchmark users	128,100
• Average dialog response time	0.96 seconds
• Throughput Fully processed order line items/hour Dialog steps/hour SAPS	 14,026,330 42,079,000 701,320
• Average database request time (dialog/update)	0.007 sec / 0.017 sec
• CPU utilization of central server	92%
• 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 PRIMERGY RX4770 M7, 4 processors / 240 cores / 480 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, 2,048 GB main memory

The benchmark result of 128,100 SD benchmark users is the best 4 processor SAP SD Benchmark result on Windows (as of 2023-11-27).

The following chart compares the two-tier SAP SD Standard Application Benchmarks for PRIMERGY RX4770 M7 and its predecessor RX4770 M6, shown are the number of SD benchmark users.

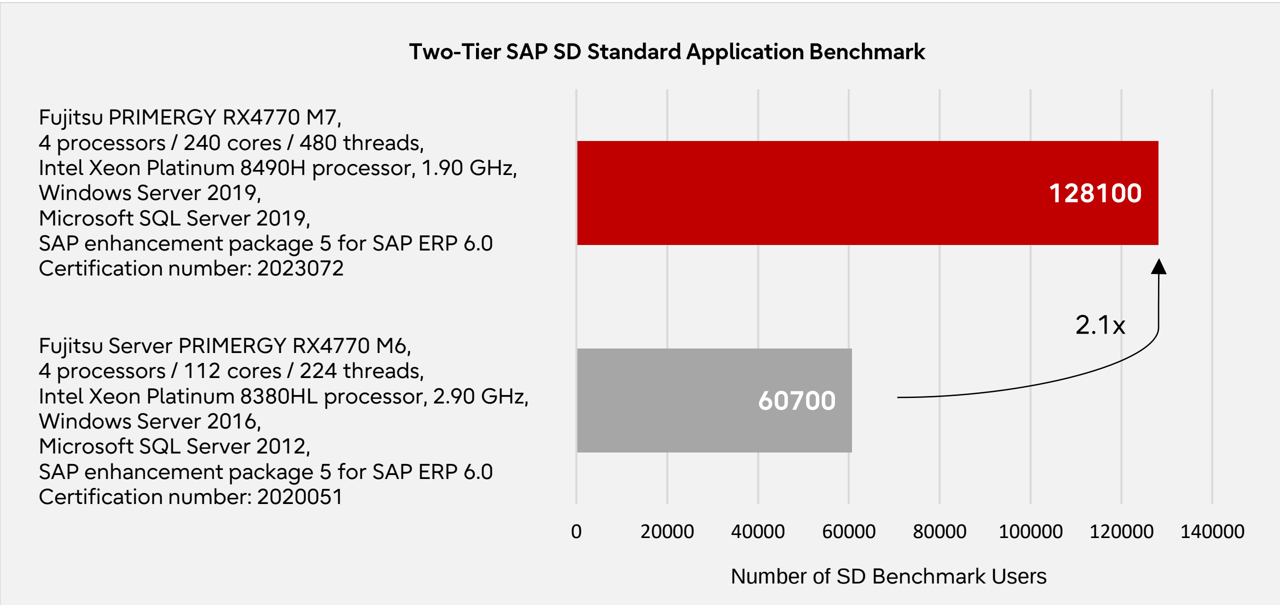

The 4th Generation Xeon Scalable Family (aka Sapphire Rapids) based RX4770 M7 with Intel Xeon Platinum 8490H delivers an improvement of 2.1x compared to the previous 3rd Generation Xeon Scalable Family (aka Cooper Lake) based RX4770 M6 with Intel Xeon Platinum 8380HL processor.

The SAP SD Benchmark certificates can be found here: Certification 2023072, Certification 2020051.

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

<u>Data load phase</u>

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.

<u>Query throughput phase</u>

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.

<u>Query runtime phase</u>

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 <u>here</u>.

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 Benchmark for PRIMERGY RX4770 M7 was performed on a scale-up configuration.

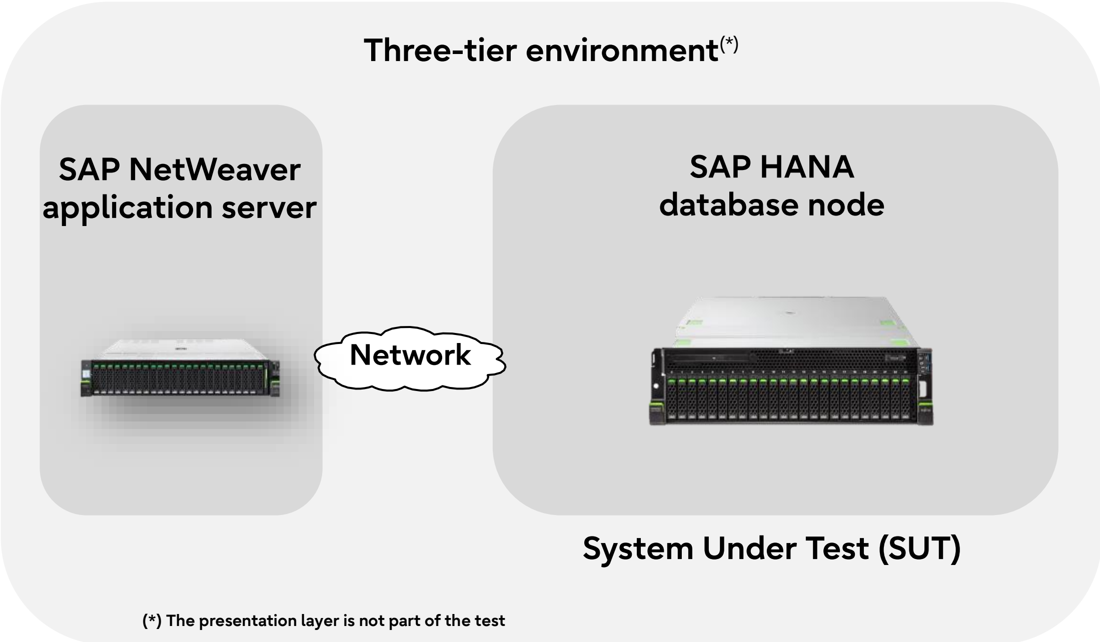

Although an application server is involved in the benchmark, neither performance metrics are measured nor does the server appear on the benchmark certificate.

System Under Test (SUT)	
Hardware	
• Model	PRIMERGY RX4770 M7
• Processor	4 x Intel Xeon Platinum 8490H processor 60C 1.9GHz 350W
• Memory	64 x 64 GB (1x64GB) 2Rx4 DDR5-4800 R ECC
• Network interface	1 Gbit/s (RJ45) on Motherboard
• Storage subsystem	1 × PRAID EP640i LP internal RAID controller 1 x SSD SAS 2.5" Mixed Use 1.6TB 2 x PRAID EP680e LP external RAID controller 3 x ETERNUS JX40 S2 Enclosure 25 × JX40 S2 TLC SSD 960GB 3DWPD
Software	
• Operating system	SUSE Linux Enterprise Server 15 SP4
• Database	SAP HANA 2.0

Application Server

Hardware

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

Software

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

Benchmark results

On June 13, 2023, the following SAP BW edition for SAP HANA Standard Application Benchmark Version 3 was certified:

Certification number 2023025	
• Benchmark Phase 1 Number of initial records Runtime of last Data Set (seconds)	 7,800,000,000 8,562
• Benchmark Phase 2 Query Executions per Hour CPU utilization of database server	 11,493 97%
• Benchmark Phase 3 Total Runtime of complex query phase (seconds)	 88
• 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 M7, 4 processors / 240 cores / 480 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, 4,096 GB DRAM

The benchmark set an overall world record on the SAP BW Edition for SAP HANA Standard Application Benchmark Version 3 on scale-up configurations with 7.8 billion initial records (as of 2023-06-13).

The SAP BWH Benchmark certificate can be found here: Certification 2023025.

The following charts compare the results of all three benchmark phases of the SAP BW Edition for SAP HANA Standard Application Benchmark for PRIMERGY RX4770 M7 and its predecessor PRIMERGY RX4770 M6.

On PRIMERGY RX4770 M6 the benchmark was also conducted with 7.8 billion initial records. Here are the details of the benchmark:

FUJITSU Server PRIMERGY RX4770 M6, 4 processors / 112 cores / 224 threads Intel Xeon Platinum 8380HL processor, 1536 GB DRAM, 3072 GB Persistent Memory, SAP BW edition for SAP HANA Standard Application Benchmark Version 3, 3-tier scale-up configuration, Bare-Metal environment, SAP HANA 2.0, SAP NetWeaver 7.50, SUSE Linux Enterprise Server 15, 7.8 billion initial records, 12752 seconds runtime of last data set, 4729 query executions/h, 90 seconds total runtime of complex query phase. Certification 2020040.

Benchmark phase 2, the CPU bound query throughput phase, is best to compare the performance of the 4th Generation Intel Xeon Scalable Processor aka Sapphire Rapids against the 3rd Generation Xeon Scalable Processor aka Cooper Lake.

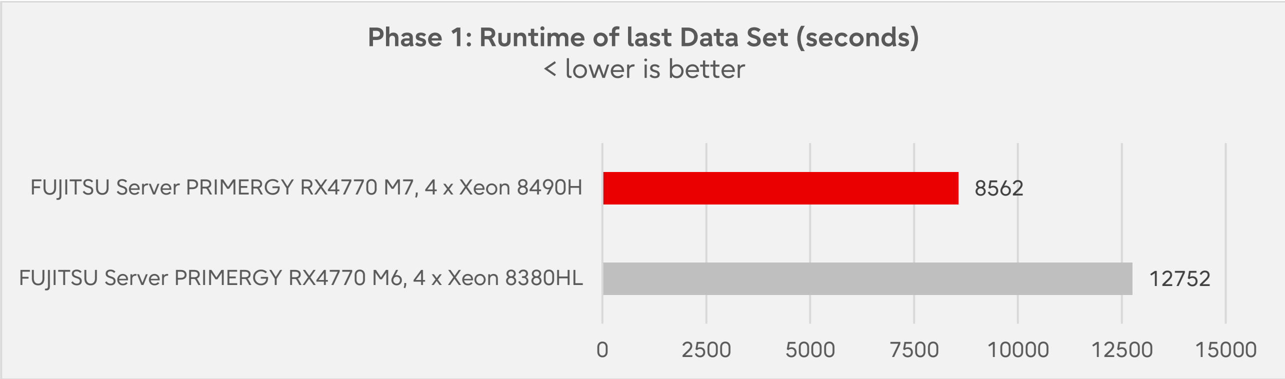

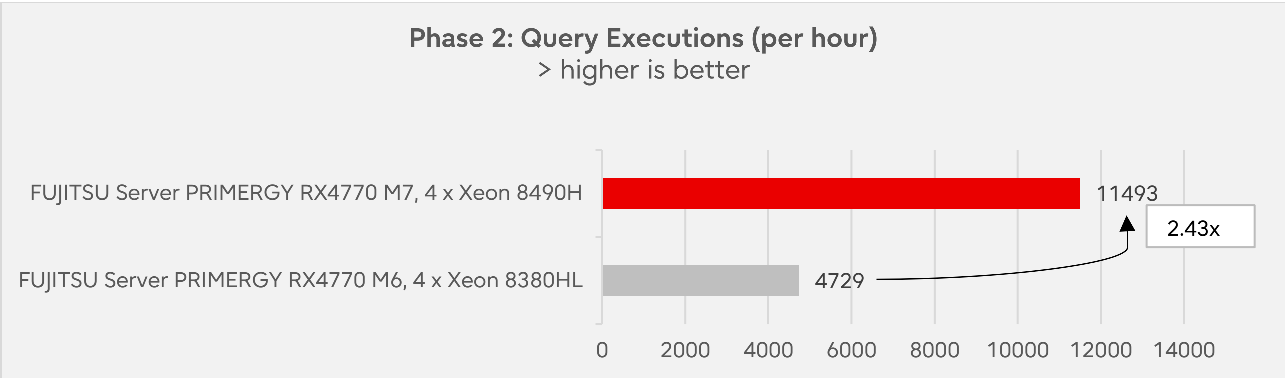

Compared to its predecessor RX4770 M6 with Intel Xeon Platinum 8380HL processor (4 processors, 112 cores, 224 threads, 3rd Generation Xeon Scalable Processor aka Cooper Lake), RX4770 M7 increases the query throughput in benchmark phase 2 by an impressive factor of 2.43.

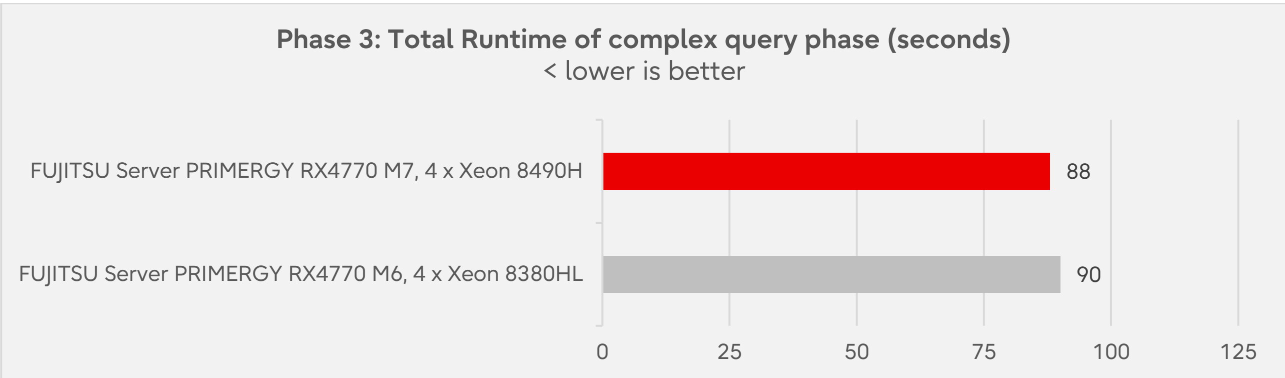

Disk I/O: Performance of storage media

Benchmark description

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

The essential specifications are as follows.

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

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

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

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

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

The main measurement items are as follows.

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

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

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

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

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

Benchmark environment

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

System Under Test (SUT)

Hardware

Controller: PRAID EP540i

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

Controller:	PRAID EP540i	
Storage media	Category	Drive name
	SATA SSD (SATA 6Gbps, Read Intensive)	MTFDDAK240TDS
		MTFDDAK480TDS
		MTFDDAK960TDS
		MTFDDAK1T9TDS
		MTFDDAK3T8TDS
		MTFDDAK7T6TDS
		MZ7L3240HCHQ
		MZ7L3480HCHQ
		MZ7L3960HCJR
		MZ7L31T9HBLT
		MZ7L33T8HBLT
		MZ7L37T6HBLA

Controller:	PRAID EP680i	
Storage media	Category	Drive name
SSD	PCIe SSD (Write intensive)	SSDPF21Q400GB
		SSDPF21Q800GB
		SSDPF21Q016TB
	PCIe SSD (Mixed Use)	KCM61VUL1T60
		KCM61VUL3T20
		KCM61VUL6T40
	PCIe SSD (Read intensive)	KCM61RUL960G
		KCM61RUL1T92
		KCM61RUL3T84
		KCM61RUL7T68

Controller:	Intel C741 Standard SATA AHCI controller	
Storage media	Category	Drive name
SSD	M.2 Flash module	MTFDDAV240TDS
		MTFDDAV480TDS

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

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

Benchmark results

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

Controller

The measurements were made using controllers in the table below.

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

Storage media

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

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

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

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

Cache settings

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

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

Performance values

The performance values are summarized in the following tables. In each case specifically for a single storage medium and with various access types and block sizes. The established measurement variables, as already mentioned in the subsection "Benchmark description" are used here. Thus, transaction rate is specified for random accesses and data throughput for sequential accesses.

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

Storage media performance

HDDs

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

SSDs

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

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

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).

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

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.

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:

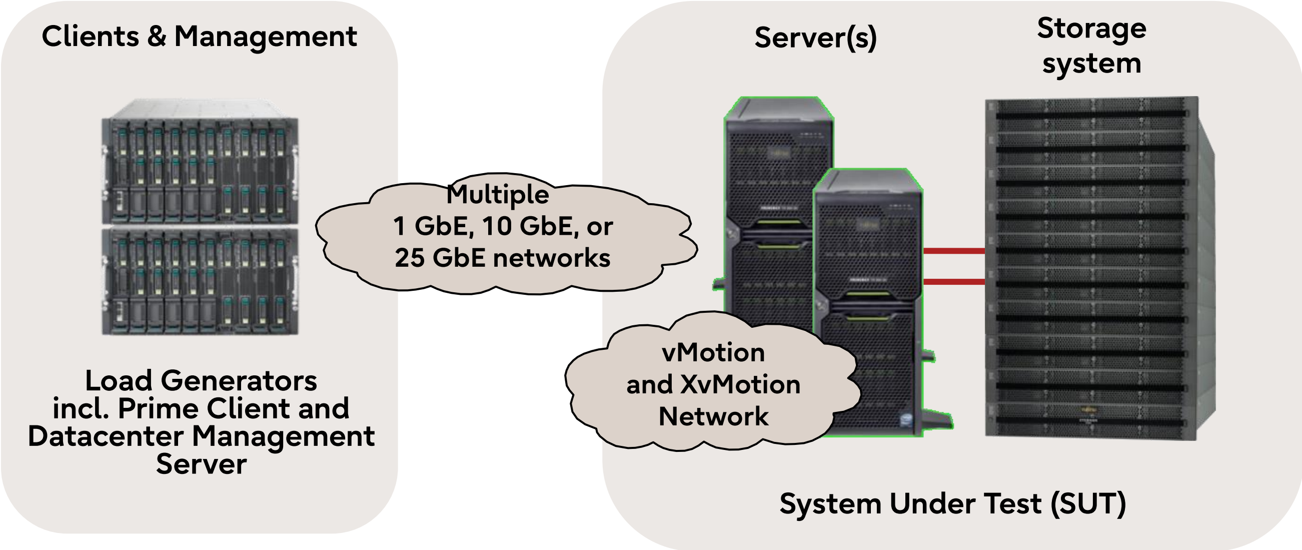

All the benchmark results were measured with the following environment:

System Under Test (SUT)	
Hardware	
• Number of servers	2
• Model	PRIMERGY RX4770 M7
• Processor	4 x Xeon Platinum 8490H 60C 1.9GHz 350W
• Memory	4096 GB : 32 x 128 GB (1x128 GB) 4x4 DDR5-4800 R ECC
• Network interface	2 x PLAN EP E810-XXVDA2 2X 25Gb SFP28 LP 1 x 1Gbit/s (RJ45) on Motherboard
• Disk subsystem	2 x PFC EP LPe35002 2X 32GFC PCIe v4 LP 14 x PRIMERGY RX2540 M4, M5 & M6 configured as Fibre Channel targets 4 x PRIMERGY RX2540 M4 : 3 x Intel P4800X PCIe SSD (750 GB) 1 x Intel P4600 PCIe SSD (4 TB) 1 x PRIMERGY RX2540 M4 : 3 x Intel P4800X PCIe SSD (750 GB) 1 x Intel P4600 PCIe SSD (2 TB) 3 x PRIMERGY RX2540 M5 : 3 x Intel P4610 PCIe SSD (3.2 TB) 1 x PRIMERGY RX2540 M5 : 1 x Intel P4610 PCIe SSD (3.2 TB) 4 x PRIMERGY RX2540 M5 : 6 x Intel P4800X PCIe SSD (750 GB) 1 x PRIMERGY RX2540 M5 : 5 x Intel P4800X PCIe SSD (750 GB)
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/2023-08-22-Fujitsu-PRIMERGY-RX4770M7.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	6 x PRIMERGY RX2530 M2 2 x PRIMERGY RX4770 M4
• Processor	2 x Intel Xeon E5-2699 v4 (4 x PRIMERGY RX2530 M2) 2 x Intel Xeon E5-2699A v4 (2 x PRIMERGY RX2530 M2) 4 x Intel Xeon Platinum 8180M (2 x PRIMERGY RX4770 M4)
• Memory	256 GB (6 x PRIMERGY RX2530 M2) 736 GB (1 x PRIMERGY RX4770 M4) 1504 GB (1 x PRIMERGY RX4770 M4
• 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 (August 22, 2023)

On August 22, 2023, Fujitsu achieved a VMmark V3.1.1 score of "44.69@44 tiles" using PRIMERGY RX4770 M7 with Xeon Platinum 8490H processors and VMware ESXi 8.0 GA. At this time, the system configuration had a total of 2 x 240 processor cores, and two identical servers were used for the "System Under Test" (SUT). Based on the above results, PRIMERGY RX4770 M7 is rated as the most powerful 4-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.

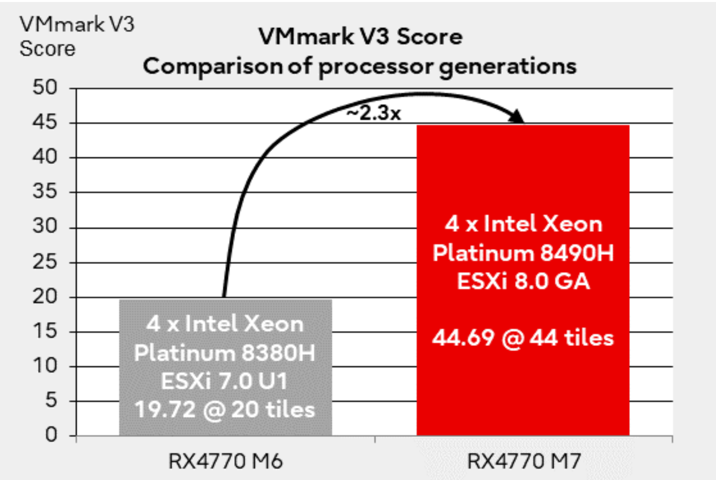

The graph on the left compares the VMmark V3 scores of the PRIMEGY RX4770 M7 and the previous generation PRIMEGY RX4770 M6.

The PRIMERGY RX4770 M7 achieved about 2.3 times improvement in score compared to the previous generation PRIMERGY RX4770 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.

Literature

PRIMERGY Servers

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

PRIMERGY RX4770 M7

This Whitepaper

https://docs.ts.fujitsu.com/dl.aspx?id=f0a61d60-af69-41b1-9ba3-feaceeecfb63

https://docs.ts.fujitsu.com/dl.aspx?id=5345eea5-9d26-447a-9ebe-ee3d2c8a9ee5

PRIMERGY Performance

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

SPEC CPU2017

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

Benchmark Overview SPECcpu2017

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

STREAM

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

LINPACK

The LINPACK Benchmark: Past, Present, and Future

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

TOP500

https://www.top500.org/

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

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

Intel Math Kernel Library – LINPACK Download

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

SPECpower_ssj2008

https://www.spec.org/power_ssj2008

Benchmark Overview SPECpower_ssj2008

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

SAP SD / BWH

https://www.sap.com/benchmark

Benchmark results

SAP SD: https://www.sap.com/dmc/exp/2018-benchmark-directory/#/sd

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 BWH: http://docs.ts.fujitsu.com/dl.aspx?id=70a4c869-586c-49f3-a6a4-47f188dd72b3

Document change history

Version	Date	Description
1.1	2024-01-12	Update: • SAP SD Standard Application Benchmark Measured with Intel Xeon Platinum 8490H • VMmark V3 Measured with Intel Xeon Platinum 8490H
1.0	2023-07-04	New: • Technical data • SPEC CPU2017, STREAM, LINPACK Measured and calculated with 4th Generation Intel Xeon Processor Scalable Family • SPECpower_ssj2008 Measured with Intel Xeon Platinum 8490H • SAP BWH Standard Application Benchmark Measured with Intel Xeon Platinum 8490H • Disk I/O Measured with 2.5 inch model

Contact

Fujitsu

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

PRIMERGY Performance and Benchmarks

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