

Fujitsu Server PRIMERGY Performance Report PRIMERGY RX2450 M2

This document provides an overview of benchmarks executed on the Fujitsu Server PRIMERGY RX2450 M2.

Explains PRIMERGY RX2450 M2 performance data in comparison to other PRIMERGY models. In addition to the benchmark results, the explanation for each benchmark and benchmark environment are also included.

Version
1.3
2025-01-28



Contents

- Technical data 3**
- SPEC CPU2017 6**
 - Benchmark description 6
 - Benchmark environment 8
 - Benchmark results 9
- STREAM 12**
 - Benchmark description 12
 - Benchmark environment 13
 - Benchmark results 14
- SPECpower_ssj2008 16**
 - Benchmark description 16
 - Benchmark environment 17
 - Benchmark results (EPYC 9754) 18
 - Benchmark results (EPYC 9845) 19
 - Comparison with prev-generation CPU (EPYC 9845 vs EPYC 9754) 20
- SAP Sales and Distribution (SD) Standard Application Benchmark 21**
 - Description of the benchmark 21
 - Benchmark environment 23
 - Benchmark results 25
- Disk I/O: Performance of storage media 28**
 - Benchmark description 28
 - Benchmark environment 31
 - Benchmark results 36
- VMmark V3 42**
 - Benchmark description 42
 - Benchmark environment 44
 - Benchmark results 46
- VMmark V4 47**
 - Benchmark description 47
 - Benchmark environment 48
 - Benchmark result 50
- Literature 51**

Technical data

PRIMERGY RX2450 M2



In this document, the power of 10 (example: 1 GB = 10⁹ bytes) is used to indicate the capacity of the internal storage, and the power of 2 (example: 1 GB = 2³⁰ bytes) is used to indicate the capacity of the cache or memory module. Any other exceptional notation will be specified separately.

Model	PRIMERGY RX2450 M2
Form factor	Rack server
Number of sockets	2
Number of processors orderable	1 or 2
Processor type	AMD EPYC 9004 Series Processors AMD EPYC 9005 Series Processors
Number of memory slots	24
Maximum memory configuration	6,144 GB
Maximum number of internal storage disks	30
Maximum number of PCI slots	PCI Express 5.0 : 8

Processor							
Model	Number of cores	Number of threads	L3 Cache [MB]	Rated frequency [GHz]	Maximum turbo frequency [GHz]	Maximum memory transfer rate [MT/s]	TDP [W]
AMD EPYC 9004 Series Processors (1CPU and 2CPU supported processor)							
EPYC 9754	128	256	256	2.25	3.10	4,800	360
EPYC 9654	96	192	384	2.40	3.70	4,800	360
EPYC 9634	84	168	384	2.25	3.70	4,800	290
EPYC 9554	64	128	256	3.10	3.75	4,800	360
EPYC 9534	64	128	256	2.45	3.70	4,800	280
EPYC 9454	48	96	256	2.75	3.80	4,800	290
EPYC 9384X	32	64	768	3.10	3.90	4,800	320
EPYC 9354	32	64	256	3.25	3.80	4,800	280
EPYC 9334	32	64	128	2.70	3.90	4,800	210
EPYC 9274F	24	48	256	4.05	4.30	4,800	320
EPYC 9224	24	48	64	2.50	3.70	4,800	200
EPYC 9184X	16	32	768	3.55	4.20	4,800	320
EPYC 9174F	16	32	256	4.10	4.40	4,800	320
EPYC 9124	16	32	64	3.00	3.70	4,800	200
AMD EPYC 9004 Series Processors (1CPU supported processor)							
EPYC 9654P	96	192	384	2.40	3.70	4,800	360
EPYC 9554P	64	128	256	3.10	3.75	4,800	360
EPYC 9454P	48	96	256	2.75	3.80	4,800	290
EPYC 9354P	32	64	256	3.25	3.80	4,800	280
AMD EPYC 9005 Series Processors (1CPU and 2CPU supported processor)							
EPYC 9845	160	320	320	2.10	3.70	6,000	390
EPYC 9745	128	256	256	2.40	3.70	6,000	400
EPYC 9655	96	192	384	2.60	4.50	6,000	400
EPYC 9575F	64	128	256	3.30	5.00	6,000	400
EPYC 9555	64	128	256	3.20	4.40	6,000	360
EPYC 9475F	48	96	256	3.65	4.80	6,000	400
EPYC 9455	48	96	256	3.15	4.40	6,000	300
EPYC 9375F	32	64	256	3.80	4.80	6,000	320
EPYC 9355	32	64	256	3.55	4.40	6,000	280
EPYC 9135	16	32	64	3.65	4.30	6,000	200
EPYC 9115	16	32	64	2.60	4.10	6,000	125
EPYC 9015	8	16	64	3.60	4.10	6,000	125

All processors that can be ordered with the PRIMERGY RX2450 M2 support AMD Turbo Core Technology. This technology allows you to operate the processor with higher frequencies than the rated frequency. The maximum frequency that can actually be achieved depends on the type of applications and the processing load.

The turbo functionality can be set in the BIOS option. Generally, Fujitsu recommends leaving the [Core Performance Boost] option set at the standard setting of [Enabled], as performance is substantially increased by the higher frequencies. However, the higher frequencies depend on the operating conditions mentioned above and is not always guaranteed. If you need stable performance or want to reduce power consumption, it may be beneficial to set [Core Performance Boost] to [Disabled] to disable Turbo function.

Memory modules

Type	Capacity [GB]	Number of ranks	Bit width of the memory chips	Memory transfer rate [MT/s]	3DS	Registered	ECC
DDR5-4800 *1							
16 GB (1x16 GB) 1Rx8 DDR5-4800 RDIMM	16	1	8	4,800		✓	✓
32 GB (1x32 GB) 2Rx8 DDR5-4800 RDIMM	32	2	8	4,800		✓	✓
32 GB (1x32 GB) 1Rx4 DDR5-4800 RDIMM	32	1	4	4,800		✓	✓
64 GB (1x64 GB) 2Rx4 DDR5-4800 RDIMM	64	2	4	4,800		✓	✓
128 GB (1x128 GB) 4Rx4 DDR5-4800 RDIMM 3DS	128	4	4	4,800	✓	✓	✓
256 GB (1x64 GB) 8Rx4 DDR5-4800 RDIMM 3DS	256	8	4	4,800	✓	✓	✓
DDR5-5600 *2							
16 GB (1x16 GB) 1Rx8 DDR5-5600 RDIMM	16	1	8	5,600		✓	✓
32 GB (1x32 GB) 2Rx8 DDR5-5600 RDIMM	32	2	8	5,600		✓	✓
32 GB (1x32 GB) 1Rx4 DDR5-5600 RDIMM	32	1	4	5,600		✓	✓
64 GB (1x64 GB) 2Rx4 DDR5-5600 RDIMM	64	2	4	5,600		✓	✓
96 GB (1x96 GB) 2Rx4 DDR5-5600 RDIMM	96	2	4	5,600		✓	✓
256 GB (1x256 GB) 8Rx4 DDR5-5600 RDIMM 3DS	256	8	4	5,600	✓	✓	✓
DDR5-6400 *2							
16 GB (1x16 GB) 1Rx8 DDR5-6400 RDIMM	16	1	8	6,400		✓	✓
32 GB (1x32 GB) 2Rx8 DDR5-6400 RDIMM	32	2	8	6,400		✓	✓
32 GB (1x32 GB) 1Rx4 DDR5-6400 RDIMM	32	1	4	6,400		✓	✓
64 GB (1x64 GB) 2Rx4 DDR5-6400 RDIMM	64	2	4	6,400		✓	✓
96 GB (1x96 GB) 2Rx4 DDR5-6400 RDIMM	96	2	4	6,400		✓	✓
128 GB (1x128 GB) 2Rx4 DDR5-6400 RDIMM	128	2	4	6,400		✓	✓

*1 Supported by AMD EPYC 9004 Series processors

*2 Supported by AMD EPYC 9005 Series processors

Power supplies	Maximum number	
Modular redundant PSU	900W platinum PSU	2
	900W titanium PSU	2
	1,600W platinum PSU	2
	1,600W titanium PSU	2
	2,200W platinum PSU	2
	2,400W titanium PSU	2
DC PSU	1,300W PSU DC	2
	1,600W PSU HVDC	2

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

Detailed technical information is available in the data sheet PRIMERGY RX2450 M2.

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	Arithmetic	Compiler optimization	Measurement result		
SPECspeed2017_int_peak	10	Integer	Aggressive (peak)	Speed	Performance	
SPECspeed2017_int_energy_peak					Power efficiency	
SPECspeed2017_int_peak	10		Conservative (base)	Speed	Performance	
SPECspeed2017_int_energy_peak					Power efficiency	
SPECspeed2017_int_peak	10		Floating point	Aggressive (peak)	Throughput	Performance
SPECspeed2017_int_energy_peak						Power efficiency
SPECspeed2017_int_peak	10			Conservative (base)	Throughput	Performance
SPECspeed2017_int_energy_peak						Power efficiency
SPECspeed2017_int_peak	10	Aggressive (peak)		Speed	Performance	
SPECspeed2017_int_energy_peak					Power efficiency	
SPECspeed2017_int_peak	10	Conservative (base)		Speed	Performance	
SPECspeed2017_int_energy_peak					Power efficiency	
SPECspeed2017_int_peak	13	Floating point	Aggressive (peak)	Throughput	Performance	
SPECspeed2017_int_energy_peak					Power efficiency	
SPECspeed2017_int_peak	13		Conservative (base)	Throughput	Performance	
SPECspeed2017_int_energy_peak					Power efficiency	

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 $4/[\# \text{ base copies}]$ 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 RX2450 M2
• Processor	2 x AMD EPYC 9004 Series Processors or 1 x AMD EPYC 9004 Series Processors 2 x AMD EPYC 9005 Series Processors
• Memory	<p>AMD EPYC 9004 Series Processors 24 x 32GB (1x32GB) 2Rx8 DDR5-4800 R ECC (2CPU configuration) or 12 x 32GB (1x32GB) 2Rx8 DDR5-4800 R ECC (1CPU configuration)</p> <p>AMD EPYC 9005 Series Processors 24 x 32GB (1x32GB) 2Rx8 DDR5-6400 R ECC or 24 x 64GB (1x64GB) 2Rx4 DDR5-6400 R ECC * * Only measurement of SPECrate2017_int_base with EPYC 9845</p>

Software	
• BIOS settings	Pleaes refer to the site below: https://www.spec.org/cgi-bin/osgresults?conf=cpu2017&op=fetch&field=SYSTEM&pattern=RX2450%20M2
• Operating system	<p>AMD EPYC 9004 Series Processors SPECrate2017_int_base, SPECspeed2017_fp_base: SUSE Linux Enterprise Server 15 SP4 5.14.21-150400.22-default SPECrate2017_fp_base, SPECspeed2017_int_base: Red Hat Enterprise Linux 9.0 (Plow) 5.14.0-70.13.1.el9_0.x86_64</p> <p>AMD EPYC 9005 Series Processors SUSE Linux Enterprise Server 15 SP6 6.4.0-150600.21-default</p>
• Operating system settings	Pleaes refer to the site below: https://www.spec.org/cgi-bin/osgresults?conf=cpu2017&op=fetch&field=SYSTEM&pattern=RX2450%20M2
• Compiler	<p>AMD EPYC 9004 Series Processors C/C++/Fortran: Version 4.0.0 of AOCC</p> <p>AMD EPYC 9005 Series Processors C/C++/Fortran: Version 5.0.0 of AOCC</p>

Benchmark results

For processors, the benchmark results depend 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 mainly load only one core, the maximum processor frequency that can be achieved is higher than with multi-threaded benchmarks.

Results with "est." are estimated values.

Processor model	Number of cores	L3 Cache [MB]	Rated frequency [GHz]	Memory transfer rate [MT/s]	Number of processors	SPECrate2017 int_base	SPECrate2017 fp_base
AMD EPYC 9004 Series Processors (2CPU configuration)							
EPYC 9754	128	256	2.25	4,800	2	1,860	1,440
EPYC 9654	96	384	2.40	4,800	2	1,650	1,410
EPYC 9634	84	384	2.25	4,800	2	1,350	1,250
EPYC 9554	64	256	3.10	4,800	2	1,310	1,220
EPYC 9534	64	256	2.45	4,800	2	1,200	1,140
EPYC 9454	48	256	2.75	4,800	2	1,030	1,070
EPYC 9384X	32	768	3.10	4,800	2	804	939
EPYC 9354	32	256	3.25	4,800	2	742	892
EPYC 9334	32	128	2.70	4,800	2	707	795
EPYC 9274F	24	256	4.05	4,800	2	631	794
EPYC 9224	24	64	2.50	4,800	2	511	577
EPYC 9184X	16	768	3.55	4,800	2	457	584
EPYC 9174F	16	256	4.10	4,800	2	444	600
EPYC 9124	16	64	3.00	4,800	2	355	458
AMD EPYC 9004 Series Processors (1CPU configuration)							
EPYC 9654P	96	384	2.40	4,800	1	831	711
EPYC 9554P	64	256	3.10	4,800	1	657	610
EPYC 9454P	48	256	2.75	4,800	1	518	534
EPYC 9354P	32	256	3.25	4,800	1	372	447
AMD EPYC 9005 Series Processors (2CPU configuration)							
EPYC 9845	160	320	2.10	6,000	2	2,510	2,060
EPYC 9745	128	256	2.40	6,000	2	2,230 est.	1,980
EPYC 9655	96	384	2.60	6,000	2	1,990	1,910
EPYC 9575F	64	256	3.30	6,000	2	1,600	1,680
EPYC 9555	64	128	3.20	6,000	2	1,600	1,690
EPYC 9475F	48	96	3.65	6,000	2	1,340	1,570
EPYC 9455	48	96	3.15	6,000	2	1,230	1,450
EPYC 9375F	32	64	3.80	6,000	2	1,000	1,330
EPYC 9355	32	256	3.55	6,000	2	935	1,250

Processor model	Number of cores	L3 Cache [MB]	Rated frequency [GHz]	Memory transfer rate [MT/s]	Number of processors	SPECrate2017 int_base	SPECrate2017 fp_base
AMD EPYC 9005 Series Processors (2CPU configuration, continued)							
EPYC 9135	16	64	3.65	6,000	2	474	722
EPYC 9115	16	32	2.60	6,000	2	440	593
EPYC 9015	8	64	3.60	6,000	2	239	382

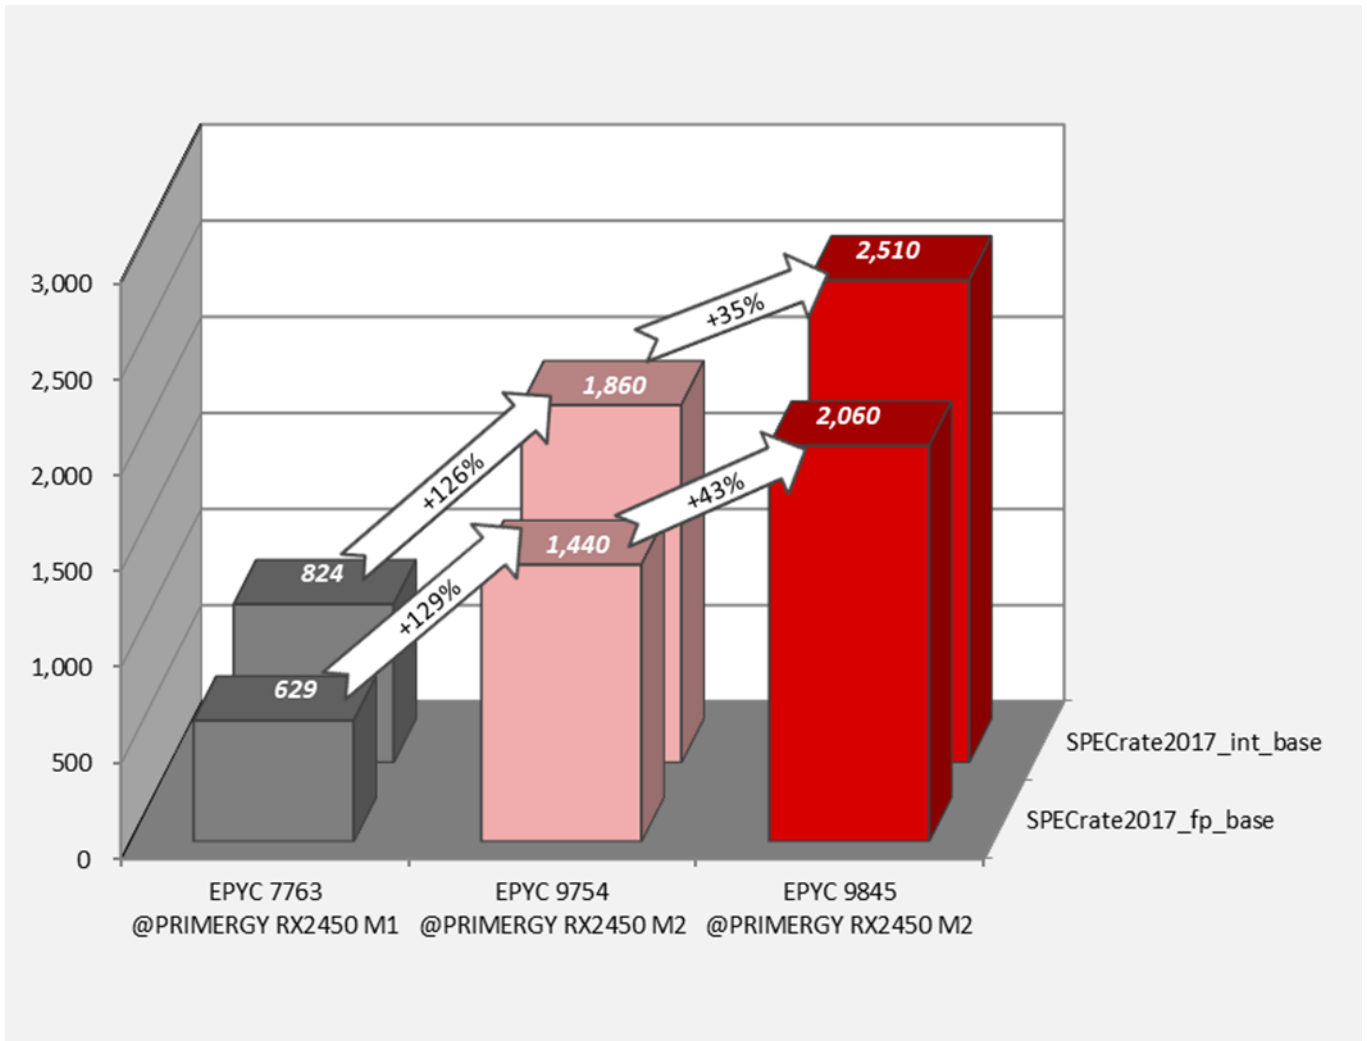
Processor model	Number of cores	L3 Cache [MB]	Rated frequency [GHz]	Memory transfer rate [MT/s]	Number of processors	SPECspeed2017 int_base	SPECspeed2017 fp_base
AMD EPYC 9004 Series Processors							
EPYC 9754	128	256	2.25	4,800	2	-	415
EPYC 9654	96	384	2.40	4,800	2	-	430
EPYC 9174F	16	256	4.10	4,800	2	16.6	-

AMD EPYC 9005 Series Processors							
EPYC 9845	160	320	2.10	6,000	2	-	505
EPYC 9575F	64	256	3.30	6,000	2	21.3	554

The following graphs compare the throughputs of PRIMERGY RX2450 M2 and its older model, PRIMERGY RX2450 M1, with maximum performance configurations.

EPYC 9754, which is the highest performance model in EPYC 9004 series processors, scored 126% (2.26 times) higher on SPECrate2017_int_base and 129% (2.29 times) higher on SPECrate2017_fp_base than EPYC 7763, which is the highest performance model in EPYC 7003 series processors.

Moreover, EPYC 9845, which is the highest performance model in EPYC 9005 series processors supported by PRIMERGY RX2450 M2, scored 35% higher on SPECrate2017_int_base and 43% higher on SPECrate2017_fp_base than EPYC 9754.



SPECrate2017: Comparison of PRIMERGY RX2450 M1 and PRIMERGY RX2450 M2

STREAM

Benchmark description

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

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

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

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

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

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

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

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

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

Benchmark environment

System Under Test (SUT)

Hardware

• Model	PRIMERGY RX2450 M2
• Processor	2 x AMD EPYC 9004 Series Processors or 1 x AMD EPYC 9004 Series Processors 2 x AMD EPYC 9005 Series Processors
• Memory	<p>AMD EPYC 9004 Series Processors</p> <p>32 x 32GB (1x32GB) 2Rx4 DDR5-4800 R ECC (2CPU configuration) or 16 x 32GB (1x32GB) 2Rx4 DDR5-4800 R ECC (1CPU configuration)</p> <p>AMD EPYC 9005 Series Processors</p> <p>32 x 32GB (1x32GB) 2Rx4 DDR5-4800 R ECC</p>

Software

• BIOS settings	<ul style="list-style-type: none"> • SMT Control = Disabled • Power Profile Selection = High Performance • NUMA nodes per socket = NPS4 * <p>* NPS2 in EPYC 9845, EPYC 9015</p>
• Operating system	<p>AMD EPYC 9004 Series Processors</p> <p>Red Hat Enterprise Linux 9.0 (Plow) 5.14.0-70.13.1.el9_0.x86_64</p> <p>AMD EPYC 9005 Series Processors</p> <p>SUSE Linux Enterprise Server 15 SP6 6.4.0-150600.21-default</p>

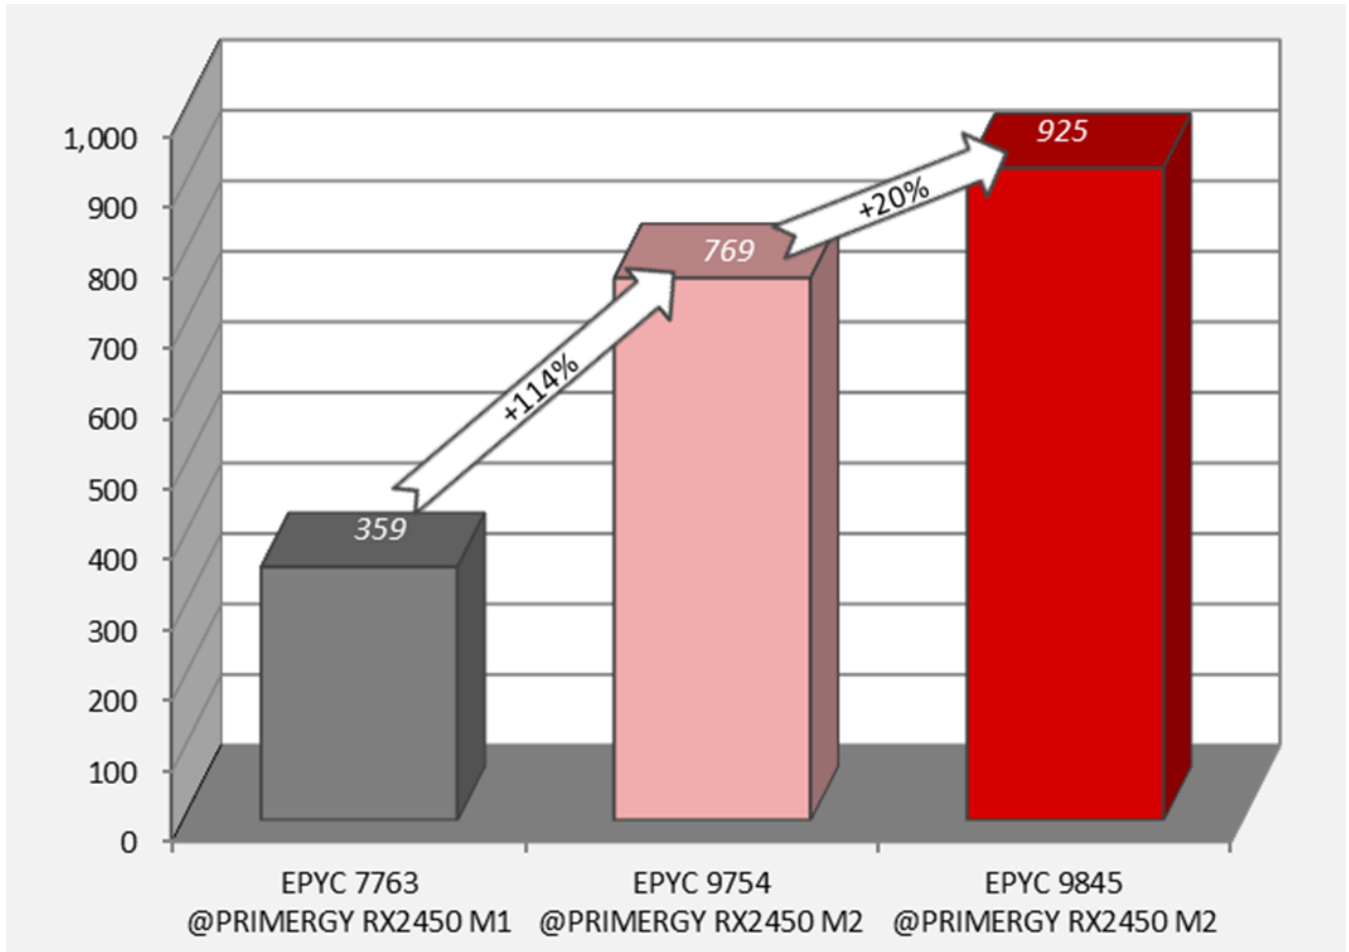
Benchmark results

Processor model	Memory transfer rate [MT/s]	Maximum memory bandwidth [GB/s]	Number of cores	Rated frequency [GHz]	Number of processors	TRIAD [GB/s]
AMD EPYC 9004 Series Processors (2CPU configuration)						
EPYC 9754	4,800	460.8	128	2.25	2	769
EPYC 9654	4,800	460.8	96	2.40	2	766
EPYC 9634	4,800	460.8	84	2.25	2	771
EPYC 9554	4,800	460.8	64	3.10	2	783
EPYC 9534	4,800	460.8	64	2.45	2	782
EPYC 9454	4,800	460.8	48	2.75	2	786
EPYC 9384X	4,800	460.8	32	3.10	2	790
EPYC 9354	4,800	460.8	32	3.25	2	791
EPYC 9334	4,800	460.8	32	2.70	2	715
EPYC 9274F	4,800	460.8	24	4.05	2	789
EPYC 9224	4,800	460.8	24	2.50	2	521
EPYC 9184X	4,800	460.8	16	3.55	2	797
EPYC 9174F	4,800	460.8	16	4.10	2	797
EPYC 9124	4,800	460.8	16	3.00	2	528
AMD EPYC 9004 Series Processors (1CPU configuration)						
EPYC 9654P	4,800	460.8	96	2.40	1	384
EPYC 9554P	4,800	460.8	64	3.10	1	392
EPYC 9454P	4,800	460.8	48	2.75	1	394
EPYC 9354P	4,800	460.8	32	3.25	1	396
AMD EPYC 9005 Series Processors (2CPU configuration)						
EPYC 9845	6,000	576.0	160	2.10	2	925
EPYC 9745	6,000	576.0	128	2.40	2	970
EPYC 9655	6,000	576.0	96	2.60	2	966
EPYC 9575F	6,000	576.0	64	3.30	2	970
EPYC 9555	6,000	576.0	64	3.20	2	970
EPYC 9475F	6,000	576.0	48	3.65	2	965
EPYC 9455	6,000	576.0	48	3.15	2	940
EPYC 9375F	6,000	576.0	32	3.80	2	969
EPYC 9355	6,000	576.0	32	3.55	2	971
EPYC 9135	6,000	576.0	16	3.65	2	884
EPYC 9115	6,000	576.0	16	2.60	2	483
EPYC 9015	6,000	576.0	8	3.60	2	483

The following graphs compare the throughputs of PRIMERGY RX2450 M2 and its older model, PRIMERGY RX2450 M1, with maximum performance configurations.

EPYC 9754, which is the highest performance model in EPYC 9004 series processors, scored 114% (2.14 times) higher than EPYC 7763, which is the highest performance model in EPYC 7003 series processors.

Moreover, EPYC 9845, which is the highest performance model in EPYC 9005 series processors supported by PRIMERGY RX2450 M2, scored 20% higher than EPYC 9754.



STREAM: Comparison of PRIMERGY RX2450 M1 and PRIMERGY RX2450 M2

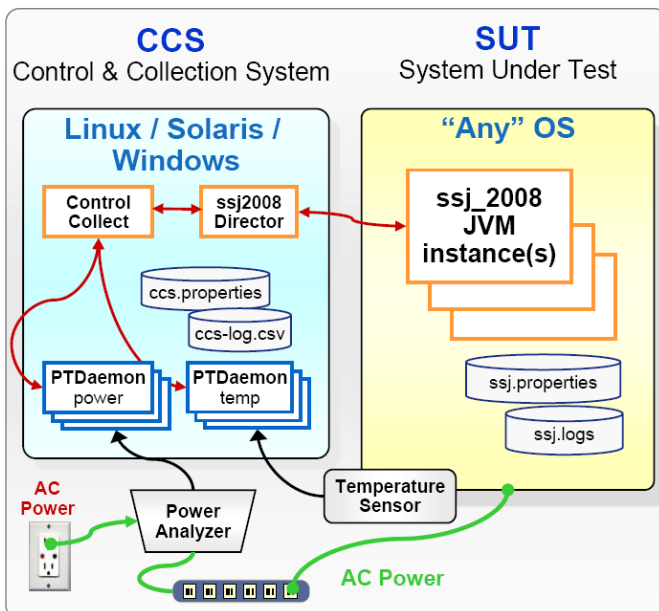

SPECpower_ssj2008

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 RX2450 M2
• Processor	2 x EPYC 9754 2 x EPYC 9845
• Memory	EPYC 9754 24 x 32GB (1x32GB) 2Rx8 DDR5-4800 R ECC EPYC 9845 24 x 32GB (1x32GB) 2Rx8 DDR5-5600 R ECC
• Network interface	1Gbit/s (RJ45) Intel i210 on Motherboard
• Disk subsystem	1 x SSD NVMe M.2 drive for booting, non hot-plug 480GB
• Power Supply Unit	1 x 1,600W titanium PSU

Software

• BIOS settings	See "Details"
• Operating system	Windows Server 2022 Standard
• Operating system settings	See "Details"
• JVM	EPYC 9754 Oracle Java HotSpot(TM) 64-Bit Server VM 18.9 (build 17.0.1+12-LTS-39, mixed mode) EPYC 9845 Oracle Java HotSpot(TM) 64-Bit Server VM (build 17.0.12+8-LTS-286, mixed mode, sharing)
• JVM settings	See "Details"


Details

• See disclosure	EPYC 9754 https://www.spec.org/power_ssj2008/results/res2024q1/power_ssj2008-20240130-01365.html EPYC 9845 https://www.spec.org/power_ssj2008/results/res2024q4/power_ssj2008-20241105-01469.html
------------------	--

Benchmark results (EPYC 9754)

The PRIMERGY RX2450 M2 in Microsoft Windows Server 2022 Standard achieved the following result:

SPECpower_ssj2008 = 30,201 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 30,201 overall ssj_ops/watt for the PRIMERGY RX2450 M2. 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%	28,893,898	810	35,652
90%	26,192,051	749	34,988
80%	23,295,106	670	34,744
70%	20,366,216	583	34,950
60%	17,450,871	523	33,381
50%	14,544,797	475	30,595
40%	11,640,006	427	27,243
30%	8,732,135	375	23,264
20%	5,815,050	318	18,295
10%	2,906,815	257	11,327
Active Idle	0	105	0
Σ ssj_ops / Σ power =			30,201

Benchmark results (EPYC 9845)

The PRIMERGY RX2450 M2 in Microsoft Windows Server 2022 Standard achieved the following result:

SPECpower_ssj2008 = 36,064 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 36,064 overall ssj_ops/watt for the PRIMERGY RX2450 M2. 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%	34,700,332	768	45,162
90%	31,255,927	708	44,159
80%	27,806,478	654	42,536
70%	24,316,399	592	41,065
60%	20,859,745	541	38,553
50%	17,381,787	490	35,467
40%	13,919,885	439	31,688
30%	10,431,534	388	26,914
20%	6,950,358	333	20,852
10%	3,478,272	274	12,696
Active Idle	0	112	0
Σ ssj_ops / Σ power =			36,064

Comparison with prev-generation CPU (EPYC 9845 vs EPYC 9754)


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 RX2450 M2 equipped with EPYC 9845 compared to the PRIMERGY RX2450 M2 equipped with the predecessor EPYC 9754.



The average throughput of the PRIMERGY RX2450 M2 equipped with EPYC 9845 is 19,110,072 ssj_ops, an improvement of 19.6% over the 15,983,695 ssj_ops of the one equipped with EPYC 9754.

On the other hand, the average power consumption of the PRIMERGY RX2450 M2 equipped with EPYC 9845 is 482 W, which is almost same as the 481 W of the one equipped with EPYC 9754.

The energy efficiency of the PRIMERGY RX2450 M2 equipped with EPYC 9845 has improved by 19.4% due to almost the same power consumption and a 19.6% improvement in performance.



SAP Sales and Distribution (SD) Standard Application Benchmark

Description of the benchmark

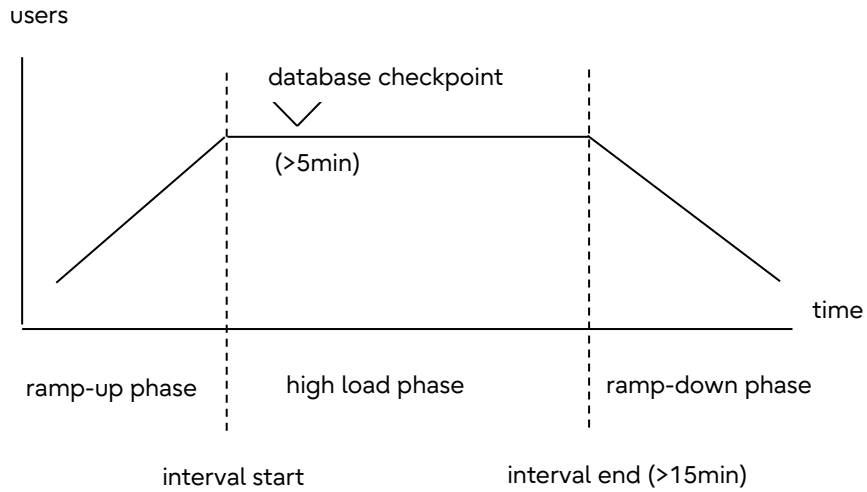
Since 1993 the SAP Standard Application Benchmarks have been developed by SAP in order to verify the performance, stability and scaling of a SAP application system and to provide information for configuring, sizing and for platform comparison. By far the most popular benchmarks from the many available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

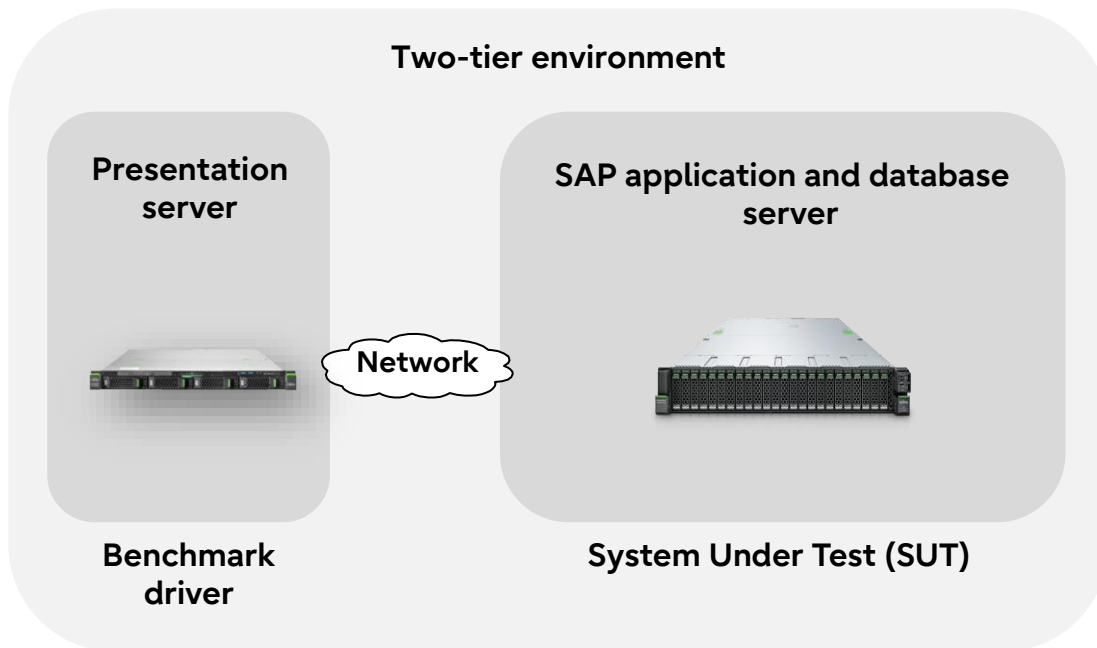


A benchmark can only be certified if the average dialog response time is less than 1 second. Certified and published SAP SD Benchmarks are published on SAP's benchmark site [here](#).

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 RX2450 M2 was performed on a two-tier configuration.



System Under Test (SUT) I: AMD EPYC 9004 Series Processor

Hardware	
• Model	PRIMERGY RX2450 M2
• Processor	2 x EPYC 9654 96C 2.4GHz 360W
• Memory	24 x 64GB (1x64GB) 2Rx4 DDR5-4800 R ECC
• Network interface	PLAN EP X710-DA2 2x10Gb Adapter
• Storage subsystem	1 x PRAID EP640i LP Raid Controller 1 x SSD SAS 2.5" Read Intensive (SFF) 1.92TB Nytro3350 2 x PCIe-SSD 2.5" Mixed Use 3.2TB Kioxia CM7-V
Software	
• Operating system	Windows Server 2022
• Database	Microsoft SQL Server 2019
• SAP Business Suite Software	SAP enhancement package 5 for SAP ERP 6.0

System Under Test (SUT) II: AMD EPYC 9005 Series Processor**Hardware**

• Model	PRIMERGY RX2450 M2
• Processor	2 x EPYC 9845 160C 2.1GHz 390W
• Memory	24 x 256GB (1x256GB) 8Rx4 DDR5-5600 3DS ECC
• Network interface	PLAN EP X710-DA2 2x10Gb Adapter
• Storage subsystem	1 x PRAID EP640i LP Raid Controller 1 x SSD SAS 2.5" Read Intensive (SFF) 1.92TB Nytro3350 2 x PCIe-SSD 2.5" Mixed Use 3.2TB Kioxia CM7-V

Software

• Operating system	SUSE Linux Enterprise Server 15 SP6
• Database	SAP ASE 16
• SAP Business Suite Software	SAP enhancement package 5 for SAP ERP 6.0

Benchmark Driver**Hardware**

• Model	PRIMERGY RX2530 M1
• Processor	2 x Xeon E5-2699v3 18C 2.30GHz 145W
• Memory	236 GB
• Network interface	PLAN EP X710-DA2 2x10Gb Adapter

Software

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

Benchmark results

Two SAP SD Benchmarks were performed on PRIMERGY RX2450 M2, the first on SUT I with AMD EPYC 9004 Series processors 9654. The second on SUT II with AMD EPYC 9005 Series processors 9845.

Result on AMD EPYC 9004 Series Processor (SUT I)

On February 12, 2024, the following SAP Sales and Distribution (SD) Standard Application Benchmark was certified:

Certification number 2024003	
• Number of SAP SD benchmark users	120,000
• Average dialog response time	0.92 seconds
• Throughput	
Fully processed order line items/hour	13,187,670
Dialog steps/hour	39,563,000
SAPS	659,380
• Average database request time (dialog/update)	0.010 sec / 0.016 sec
• CPU utilization of central server	99%
• Operating system, central server	Windows Server 2022
• RDBMS	Microsoft SQL Server 2019
• SAP Business Suite software	SAP enhancement package 5 for SAP ERP 6.0
• Configuration Central Server	Fujitsu PRIMERGY RX2450 M2, 2 processors / 192 cores / 384 threads, EPYC 9654 processor, 2.4 GHz, 64 KB L1 cache and 1,024 KB L2 cache per core, 384 MB L3 cache per processor, 1,536 GB main memory



As of 2024-02-12, the benchmark result of 120,000 SD benchmark users is the best 2 processor SAP SD Benchmark result on Windows. Please read more [here](#).

Result on AMD EPYC 9005 Series Processor (SUT II)

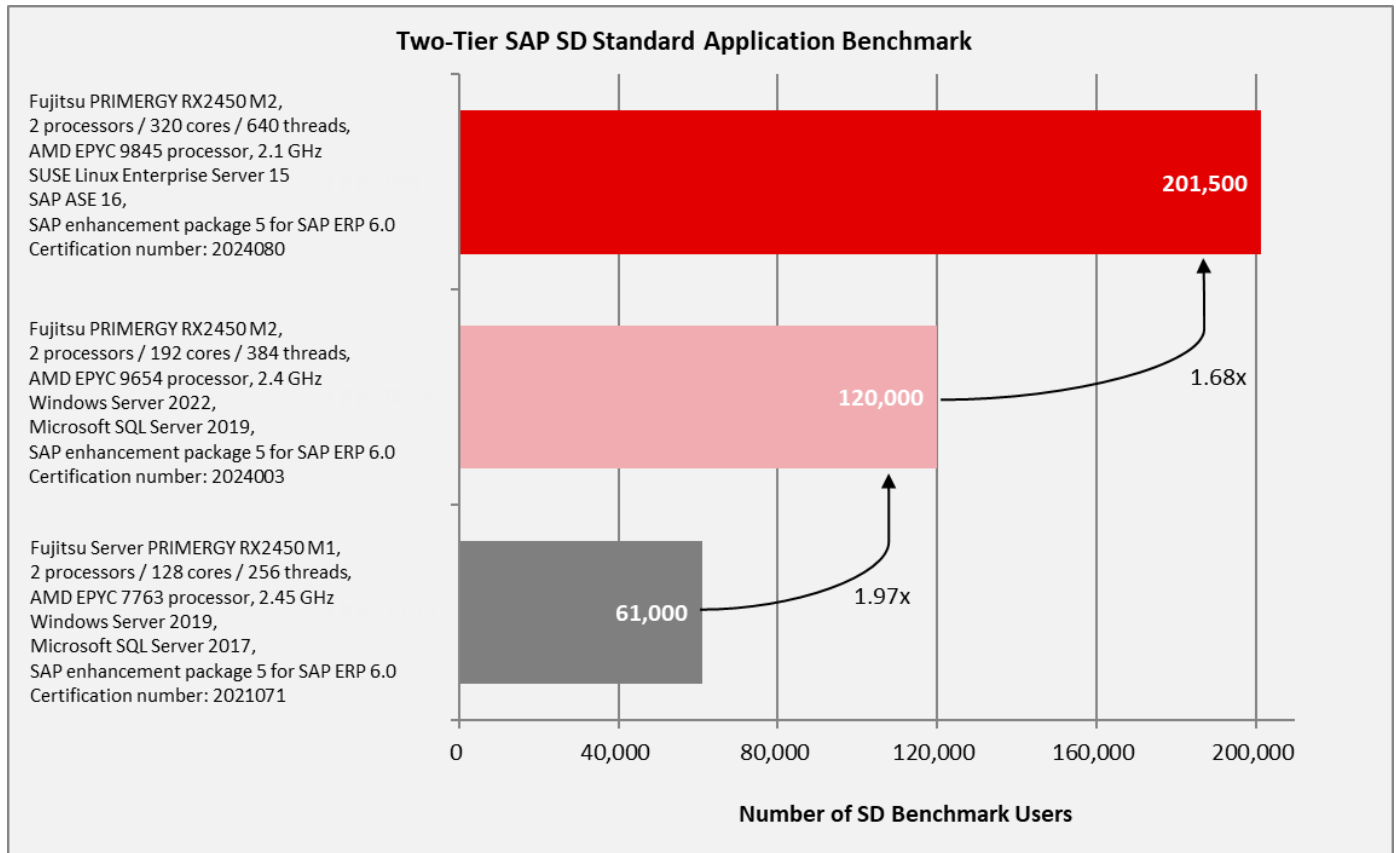
On December 04, 2024, the following SAP Sales and Distribution (SD) Standard Application Benchmark was certified:

Certification number 2024080	
• Number of SAP SD benchmark users	201,500
• Average dialog response time	0.96 seconds
• Throughput	
Fully processed order line items/hour	22,055,330
Dialog steps/hour	66,166,000
SAPS	1,102,770
• Average database request time (dialog/update)	0.005 sec / 0.008 sec
• CPU utilization of central server	90%
• Operating system, central server	SUSE Linux Enterprise Server 15
• RDBMS	SAP ASE 16
• SAP Business Suite software	SAP enhancement package 5 for SAP ERP 6.0
• Configuration Central Server	Fujitsu PRIMERGY RX2450 M2, 2 processors / 320 cores / 640 threads, EPYC 9845 processor, 2.1 GHz, 80 KB L1 cache and 1,024 KB L2 cache per core, 320 MB L3 cache per processor, 6,144 GB main memory



As of 2024-12-04, the benchmark result of 201,500 SD benchmark users is the best 2 processor SAP SD Benchmark result. Please read more [here](#).

The following chart compares the two-tier SAP SD Standard Application Benchmarks for PRIMERGY RX2450 M2 and its predecessor RX2450 M1, shown are the number of SD benchmark users.



The AMD EPYC 9004 Series (aka Genoa) based RX2450 M2 with EPYC 9654 delivers a 1.97 times higher result compared to the previous AMD EPYC 7003 Series (aka Milan) based RX2450 M1 with EPYC 7763, both running Microsoft Windows Server as OS and SQL Server as database system.

Moreover, the AMD EPYC 9005 Series (aka Turin) based RX2450 M2 with EPYC 9845 delivers a 1.68 times higher result compared to the previous AMD EPYC 9004 Series based RX2450 M2 with EPYC 9654. Please note that the RX2450 M2 equipped with EPYC 9845 was running SUSE Linux Enterprise 15 and SAP ASE 16 as database system while the RX2450 M2 with EPYC 9654 was running Windows Server 2022 and SQL Server 2019. The benefit of Linux and ASE over Windows and SQL Server is included in that factor of 1.68.

The SAP SD Benchmark certificates can be found here: Certification [2024080](#), Certification [2024003](#), and Certification [2021071](#).

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 (kiB)
- 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 [kiB]	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.

<i>Data throughput [MiB/s]</i>	$= \text{Transaction rate [IO/s]} \times \text{Block size [MiB]}$
<i>Transaction rate [IO/s]</i>	$= \text{Data throughput [MiB/s]} / \text{Block size [MiB]}$

In this section, a power of 10 (1 TB = 10¹² bytes) is used to indicate the capacity of the hard storage medium, and a power of 2 (1 MiB / s = 2²⁰ 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.](#)"

Controller

PRIMERGY server can use the following controllers.

Controller name	Cache	Supported interfaces			RAID levels
		host	drive	port	
PSAS CP 2100-8i	-	PCIe 3.0 x8	SATA 6G SAS 12G	8	0, 1, 10, 5
PSAS CP 2200-16i	-	PCIe 4.0 x8	SATA 6G SAS 12G	16	0, 1, 10, 5
			PCIe 4.0 x4	4	
PRAID CP600i	-	PCIe 4.0 x8	SATA 6G SAS 12G	8	0, 1, 10
PRAID EP640i	4GB	PCIe 4.0 x8	SATA 6G SAS 12G	8	0, 1, 1E, 10, 5, 50, 6, 60
PRAID EP 3252-8i	2GB	PCIe 4.0 x8	SATA 6G SAS 12G SAS 24G	8	0, 1, 10, 5, 50, 6, 60
PRAID EP 3254-8i	4GB	PCIe 4.0 x8	SATA 6G SAS 12G SAS 24G	8	0, 1, 10, 5, 50, 6, 60
PRAID EP680i	8GB	PCIe 4.0 x8	SATA 6G SAS 12G	16	0, 1, 1E, 10, 5, 50, 6, 60
			PCIe 4.0 x4	4	
PRAID EP 3258-16i	8GB	PCIe 4.0 x8	SATA 6G SAS 12G SAS 24G	16	0, 1, 10, 5, 50, 6, 60
			PCIe 4.0 x4	4	
Retimer card for 2.5" NVMe SSD	-	PCIe 5.0 x16	PCIe 5.0 x4	4	-
M.2 Riser Kit	-	PCIe 5.0 x4	PCIe 5.0 x2	2	-
PDUAL CP300	-	PCIe 4.0 x8	SATA 6G	2	0, 1
			PCIe 4.0 x4	2	

Storage media

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

Model	Storage media type	Interface	Form factor
3.5 inch model ^(*1)	HDD	SAS 12G	3.5 inch
		SATA 6G	3.5 inch
	SSD	SAS 12G / SAS 24G	2.5 inch ^(*2)
		SATA 6G	2.5 inch ^(*2)
2.5 inch model	HDD	SAS 12G	2.5 inch
	SSD	SAS 12G / SAS 24G	2.5 inch
		SATA 6G	2.5 inch
		PCIe 4.0 / PCIe 5.0	2.5 inch
model common	SSD	SATA 6G	M.2
		PCIe 4.0	M.2

(*1) Upgrade kit of Rear 2.5 inch bay enables you to use 2.5 inch model storage.

(*2) It is available with a 3.5 inch cage.

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

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

Cache settings

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

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

Benchmark environment

The following hardware and software components were used for benchmarking.

Hardware

3.5 inch model

Storage media	Category	Drive name
HDD	NL-SAS HDD (SAS 12Gbps, 7.2k rpm) [512e]	MG08SDA600E
		MG08SDA800E
		ST12000NM004J
		ST14000NM004J
		ST16000NM004J
		ST18000NM004J
	NL-SAS HDD (SAS 12Gbps, 7.2k rpm) [512n]	ST2000NM001B
		ST4000NM001B
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512e]	MG08ADA600E
		MG08ADA800E
	BC-SATA HDD (SATA 6Gbps, 7.2k rpm) [512n]	ST2000NM000B
		ST4000NM000B

Storage media	Category	Drive name
SSD	SAS SSD (SAS 12Gbps, Write Intensive)	XS400ME70045 XS800ME70084 XS1600ME70045
	SAS SSD (SAS 12Gbps, Mixed Use)	XS800LE70045 XS1600LE70045 XS3200LE70045
	SAS SSD (SAS 12Gbps, Read Intensive)	XS960SE70045 XS1920SE70045 XS3840SE70045 XS7680SE70045
	SATA SSD (SATA 6Gbps, Mixed Use)	MTFDDAK480TGB MTFDDAK960TGB MTFDDAK1T9TGB MTFDDAK3T8TGB MZ7L3480HELT MZ7L3960HELA MZ7L31T9HENA MZ7L33T8HENA
	SATA SSD (SATA 6Gbps, Read Intensive)	MTFDDAK240TGA MTFDDAK480TGA MTFDDAK960TGA MTFDDAK1T9TGA MTFDDAK3T8TGA MTFDDAK7T6TGA MZ7L3480HEJD MZ7L3960HELT MZ7L31T9HELA MZ7L33T8HELA MZ7L37T6HELA
	SAS SSD (SAS 24Gbps, Write Intensive)	PM7800G10DN PM71T6010DN PM7800G10DF PM71T6010DF
	SAS SSD (SAS 24Gbps, Mixed Use)	PM71T6003DN PM73T2003DN PM76T4003DN
	SAS SSD (SAS 24Gbps, Read Intensive)	PM71T9201DN PM73T8401DN PM77T6801DN PM715T301DN PM77T6801DF PM715T301DF MZILG960HCHQ MZILG1T9HCJR MZILG3T8HCLS MZILG7T6HBLA MZILG15THBLA

2.5 inch model

Storage media	Category	Drive name
HDD	SAS HDD (SAS 12Gbps, 10k rpm) [512e]	AL15SEB18EQ
		AL15SEB24EQ
	SAS HDD (SAS 12Gbps, 10k rpm) [512n]	AL15SEB030N
		AL15SEB060N
		AL15SEB120N

Storage media	Category	Drive name
SSD	SAS SSD (SAS 12Gbps, Write Intensive)	XS400ME70045 XS800ME70084 XS1600ME70045
	SAS SSD (SAS 12Gbps, Mixed Use)	XS800LE70045 XS1600LE70045 XS3200LE70045 XS6400LE70045
	SAS SSD (SAS 12Gbps, Read Intensive)	XS960SE70045 XS1920SE70045 XS3840SE70045 XS7680SE70045
	SATA SSD (SATA 6Gbps, Mixed Use)	MTFDDAK480TGB MTFDDAK960TGB MTFDDAK1T9TGB MTFDDAK3T8TGB MZ7L3480HELT MZ7L3960HELA MZ7L33T8HENA MZ7L31T9HENA
	SATA SSD (SATA 6Gbps, Read Intensive)	MTFDDAK240TGA MTFDDAK480TGA MTFDDAK960TGA MTFDDAK1T9TGA MTFDDAK3T8TGA MTFDDAK7T6TGA MZ7L3480HEJD MZ7L3960HELT MZ7L31T9HELA MZ7L33T8HELA MZ7L37T6HELA
	SAS SSD (SAS 24Gbps, Write Intensive)	PM7800G10DN PM71T6010DN
	SAS SSD (SAS 24Gbps, Mixed Use)	PM71T6003DN PM73T2003DN PM76T4003DN
	SAS SSD (SAS 24Gbps, Read Intensive)	PM71T9201DN PM73T8401DN PM77T6801DN PM715T301DN PM77T6801DF PM715T301DF MZILG960HCHQ MZILG1T9HCJR MZILG3T8HCLS MZILG7T6HBLA MZILG15THBLA

Storage media	Category	Drive name
SSD	PCIe 4.0 SSD (Write Intensive)	SSDPF21Q400GB
		SSDPF21Q800GB
		SSDPF21Q016TB
	PCIe 5.0 SSD (Mixed Use)	KCMY1VUG1T60
		KCMY1VUG3T20
		KCMY1VUG6T40
		KCMY1VUG12T8
	PCIe 5.0 SSD (Read Intensive)	KCMY1RUG1T92
		KCMY1RUG3T84
KCMY1RUG7T68		
KCMY1RUG15T3		

Model common

Storage media	Category	Drive name
M.2 SSD	SATA M.2 drive	MTFDDAV240TGA
		MTFDDAV480TGA
		MTFDDAV960TGA
	PCIe M.2 drive	MTFDKBA960TFR

Software

Operating system	Microsoft Windows Server
Measuring tool	lometer 1.1.0 (icf: benchmark version 3.0)

Logical drive settings to measure

Target Drive	Type RAID 0 logical drive consisting of 1 drive	
Stripe size	HDD : 256KB、SSD : 64 KB	
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.
Number of lometer worker	Sequential Access: 1 Random Access: 1 (except SAS 24G or PCIe 5.0 SSD), 4 (SAS 24G SSD), 16 (PCIe 5.0 SSD)	
Alignment of lometer 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."

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.

Values in rows with "est." are predicted values.

3.5 inch model

HDDs

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
□ NL-SAS 12Gbps HDD 7.2krpm [512e]							
6,000	MG08SDA600E	SAS 12G	545	455	442	234	234
8,000	MG08SDA800E	SAS 12G	542	449	449	250	250
12,000	ST12000NM004J	SAS 12G	609	578	534	266	266
14,000	ST14000NM004J	SAS 12G	616	589	524	270	269
16,000	ST16000NM004J	SAS 12G	610	586	548	270	270
18,000	ST18000NM004J	SAS 12G	603	578	522	265	262
□ NL-SAS 12Gbps HDD 7.2krpm [512n]							
2,000	ST2000NM001B	SAS 12G	489	431	428	200	200
4,000	ST4000NM001B	SAS 12G	541	486	471	239	239
□ BC-SATA HDD 7.2krpm [512e]							
6,000	MG08ADA600E	SATA 6G	497	452	447	239	239
8,000	MG08ADA800E	SATA 6G	477	429	430	248	248
□ BC-SATA HDD 7.2krpm [512n]							
2,000	ST2000NM000B	SATA 6G	415	366	389	197	196
4,000	ST4000NM000B	SATA 6G	468	422	435	236	236

SSDs

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
□ SAS 12Gbps SSD (WI)							
400	XS400ME70045	SAS 12G	129,340	20,008	18,833	1,043	1,064
800	XS800ME70084	SAS 12G	123,357	23,784	19,429	1,052	872
1,600	XS1600ME70045	SAS 12G	129,571	20,051	20,744	1,043	1,069
□ SAS 12Gbps SSD (MU)							
800	XS800LE70045	SAS 12G	129,684	20,041	18,837	1,043	1,065
1,600	XS1600LE70045	SAS 12G	129,624	20,040	20,803	1,043	1,068
3,200	XS3200LE70045	SAS 12G	129,555	20,045	20,735	1,043	1,063
6,400	XS6400LE70045	SAS 12G	129,369	20,052	20,767	1,043	1,066
□ SAS 12Gbps SSD (RI)							
960	XS960SE70045	SAS 12G	129,567	20,014	18,805	1,043	1,064
1,920	XS1920SE70045	SAS 12G	129,710	20,045	20,819	1,043	1,068
3,840	XS3840SE70045	SAS 12G	129,573	20,023	20,733	1,043	1,064
7,680	XS7680SE70045	SAS 12G	129,861	20,046	20,762	1,043	1,068
□ SATA SSD (MU)							
480	MTFDDAK480TGB	SATA 6G	43,705	5,729	5,839	491	449
960	MTFDDAK960TGB	SATA 6G	43,732	6,155	6,257	491	449
1,920	MTFDDAK1T9TGB	SATA 6G	43,735	6,394	6,513	490	449
3,840	MTFDDAK3T8TGB	SATA 6G	43,415	6,576	6,636	483	446
480	MZ7L3480HELT	SATA 6G	55,649	8,269	8,169	531	495
960	MZ7L3960HELA	SATA 6G	55,915	8,285	8,193	531	495
1,920	MZ7L31T9HENA	SATA 6G	55,829	8,283	8,190	531	495
3,840	MZ7L33T8HENA	SATA 6G	55,728	8,278	8,184	531	495
□ SATA SSD (RI)							
240	MTFDDAK240TGA	SATA 6G	41,808	5,120	5,293	480	360
480	MTFDDAK480TGA	SATA 6G	43,618	5,625	5,761	490	450
960	MTFDDAK960TGA	SATA 6G	43,631	5,878	6,033	484	449
1,920	MTFDDAK1T9TGA	SATA 6G	43,688	6,334	6,447	491	450
3,840	MTFDDAK3T8TGA	SATA 6G	43,392	6,539	6,626	483	445
7,680	MTFDDAK7T6TGA	SATA 6G	42,940	7,065	7,278	491	446
480	MZ7L3480HELD	SATA 6G	51,235	7,353	6,805	531	373
960	MZ7L3960HELT	SATA 6G	55,918	8,287	8,193	531	496
1,920	MZ7L31T9HELA	SATA 6G	55,918	8,288	8,195	531	496
3,840	MZ7L33T8HELA	SATA 6G	55,725	8,280	8,184	531	495
7,680	MZ7L37T6HELA	SATA 6G	55,798	8,253	8,169	531	496
□ SAS 24Gbps SSD (WI)							
800	PM7800G10DN	SAS 12G	168,061	20,678	23,006	1,070	1,076
		SAS 24G	204,529	25,996	25,095	1,960	1,603
1,600	PM71T6010DN	SAS 12G	173,094	22,676	26,505	1,070	1,076
		SAS 24G	208,291	26,190	24,674	1,960	1,319
800	PM7800G10DF	SAS 12G	168,061	20,678	23,006	1,070	1,076
		SAS 24G	204,452	26,015	25,115	1,960	1,603
1,600	PM71T6010DF	SAS 12G	173,094	22,676	26,505	1,070	1,076
		SAS 24G	208,527	26,219	24,674	1,960	1,318
□ SAS 24Gbps SSD (MU)							
1,600	PM71T6003DN	SAS 12G	168,200	20,700	22,800	1,070	1,076
		SAS 24G	204,400	26,000	25,100	1,963	1,603
3,200	PM73T2003DN	SAS 12G	173,000	22,600	26,500	1,070	1,076
		SAS 24G	208,200	26,100	24,600	1,960	1,318
6,400	PM76T4003DN	SAS 12G	171,200	21,400	23,200	1,070	1,076
		SAS 24G	190,700	23,900	22,500	1,963	1,175

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
□ SAS 24Gbps SSD (RI)							
1,920	PM71T9201DN	SAS 12G	168,283	20,710	22,880	1,070	1,076
		SAS 24G	204,491	26,066	25,188	1,963	1,603
3,840	PM73T8401DN	SAS 12G	173,000	22,600	26,500	1,070	1,076
		SAS 24G	208,200	26,100	24,600	1,960	1,318
7,680	PM77T6801DN	SAS 12G	171,279	21,408	23,284	1,070	1,076
		SAS 24G	190,784	23,941	22,542	1,963	1,175
15,360	PM715T301DN	SAS 12G	167,002	20,281	20,643	1,070	1,070
		SAS 24G	146,385	18,465	17,688	1,963	974
7,680	PM77T6801DF	SAS 12G	174,701	21,398	23,308	1,070	1,076
		SAS 24G	190,770	23,946	22,551	1,963	1,175
15,360	PM715T301DF	SAS 12G	167,002	20,281	20,643	1,070	1,070
		SAS 24G	146,212	18,490	17,662	1,963	974
960	MZILG960HCHQ	SAS 12G	168,803	20,856	25,149	1,075	1,062
		SAS 24G	180,407	33,362	25,733	2,049	1,290
1,920	MZILG1T9HC.R	SAS 12G	186,387	25,625	27,424	1,075	1,062
		SAS 24G	222,003	44,525	47,753	2,049	1,916
3,840	MZILG3T8HCLS	SAS 12G	187,273	25,619	33,914	1,075	1,062
		SAS 24G	244,313	45,247	54,927	2,049	1,912
7,680	MZILG7T6HBLA	SAS 12G	187,183	25,636	33,940	1,075	1,062
		SAS 24G	246,245	45,190	54,780	2,049	1,913
15,360	MZILG15THBLA	SAS 12G	187,298	25,653	33,915	1,075	1,062
		SAS 24G	229,145	45,292	54,954	2,048	1,909

**2.5 inch model
HDDs**

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
□ SAS 12Gbps HDD 10krpm [512e]							
1,800	AL15SEB18EQ	SAS 12G	767	631	624	255	249
2,400	AL15SEB24EQ	SAS 12G	754	620	617	264	260
□ SAS 12Gbps HDD 10krpm [512n]							
300	AL15SEB030N	SAS 12G	641	547	557	231	230
600	AL15SEB060N	SAS 12G	682	558	568	232	231
1,200	AL15SEB120N	SAS 12G	732	603	593	230	225

SSDs

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
□ SAS 12Gbps SSD (WI)							
400	XS400ME70045	SAS 12G	129,340	20,008	18,833	1,043	1,064
800	XS800ME70084	SAS 12G	123,357	23,784	19,429	1,052	872
1,600	XS1600ME70045	SAS 12G	129,571	20,051	20,744	1,043	1,069
□ SAS 12Gbps SSD (MU)							
800	XS800LE70045	SAS 12G	129,684	20,041	18,837	1,043	1,065
1,600	XS1600LE70045	SAS 12G	129,624	20,040	20,803	1,043	1,068
3,200	XS3200LE70045	SAS 12G	129,555	20,045	20,735	1,043	1,063
6,400	XS6400LE70045	SAS 12G	129,369	20,052	20,767	1,043	1,066
□ SAS 12Gbps SSD (RI)							
960	XS960SE70045	SAS 12G	129,567	20,014	18,805	1,043	1,064
1,920	XS1920SE70045	SAS 12G	129,710	20,045	20,819	1,043	1,068
3,840	XS3840SE70045	SAS 12G	129,573	20,023	20,733	1,043	1,064
7,680	XS7680SE70045	SAS 12G	129,861	20,046	20,762	1,043	1,068
□ SATA SSD (MU)							
480	MTFDDAK480TGB	SATA 6G	43,705	5,729	5,839	491	449
960	MTFDDAK960TGB	SATA 6G	43,732	6,155	6,257	491	449
1,920	MTFDDAK1T9TGB	SATA 6G	43,735	6,394	6,513	490	449
3,840	MTFDDAK3T8TGB	SATA 6G	43,415	6,576	6,636	483	446
480	MZ7L3480HELT	SATA 6G	55,649	8,269	8,169	531	495
960	MZ7L3960HELA	SATA 6G	55,915	8,285	8,193	531	496
1,920	MZ7L33T8HENA	SATA 6G	55,829	8,283	8,190	531	496
3,840	MZ7L31T9HENA	SATA 6G	55,728	8,278	8,184	531	495
□ SATA SSD (RI)							
240	MTFDDAK240TGA	SATA 6G	41,808	5,120	5,293	480	360
480	MTFDDAK480TGA	SATA 6G	43,618	5,625	5,761	490	450
960	MTFDDAK960TGA	SATA 6G	43,631	5,878	6,033	484	449
1,920	MTFDDAK1T9TGA	SATA 6G	43,688	6,334	6,447	491	450
3,840	MTFDDAK3T8TGA	SATA 6G	43,392	6,539	6,626	483	445
7,680	MTFDDAK7T6TGA	SATA 6G	42,940	7,065	7,278	491	446
480	MZ7L3480HEJD	SATA 6G	51,235	7,353	6,805	531	373
960	MZ7L3960HELT	SATA 6G	55,918	8,287	8,193	531	496
1,920	MZ7L31T9HELA	SATA 6G	55,918	8,288	8,195	531	496
3,840	MZ7L33T8HELA	SATA 6G	55,725	8,280	8,184	531	495
7,680	MZ7L37T6HELA	SATA 6G	55,798	8,253	8,169	531	496
□ SAS 24Gbps SSD (WI)							
800	PM7800G10DN	SAS 12G	168,061	20,678	23,006	1,070	1,076
		SAS 24G	204,529	25,996	25,095	1,960	1,603
1,600	PM71T6010DN	SAS 12G	173,094	22,676	26,505	1,070	1,076
		SAS 24G	208,291	26,190	24,674	1,960	1,319
□ SAS 24Gbps SSD (MU)							
1,600	PM71T6003DN	SAS 12G	168,200	20,700	22,800	1,070	1,076
		SAS 24G	204,400	26,000	25,100	1,963	1,603
3,200	PM73T2003DN	SAS 12G	173,000	22,600	26,500	1,070	1,076
		SAS 24G	208,200	26,100	24,600	1,960	1,318
6,400	PM76T4003DN	SAS 12G	171,200	21,400	23,200	1,070	1,076
		SAS 24G	190,700	23,900	22,500	1,963	1,175

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
□ SAS 24Gbps SSD (RI)							
1,920	PM71T9201DN	SAS 12G	168,283	20,710	22,880	1,070	1,076
		SAS 24G	204,491	26,066	25,188	1,963	1,603
3,840	PM73T8401DN	SAS 12G	173,000	22,600	26,500	1,070	1,076
		SAS 24G	208,200	26,100	24,600	1,960	1,318
7,680	PM77T6801DN	SAS 12G	171,279	21,408	23,284	1,070	1,076
		SAS 24G	190,784	23,941	22,542	1,963	1,175
15,360	PM715T301DN	SAS 12G	167,002	20,281	20,643	1,070	1,070
		SAS 24G	146,385	18,465	17,688	1,963	974
7,680	PM77T6801DF	SAS 12G	174,701	21,398	23,308	1,070	1,076
		SAS 24G	190,770	23,946	22,551	1,963	1,175
15,360	PM715T301DF	SAS 12G	167,002	20,281	20,643	1,070	1,070
		SAS 24G	146,212	18,490	17,662	1,963	974
960	MZILG960HCHQ	SAS 12G	168,803	20,856	25,149	1,075	1,062
		SAS 24G	180,407	33,362	25,733	2,049	1,290
1,920	MZILG1T9HC.R	SAS 12G	186,387	25,625	27,424	1,075	1,062
		SAS 24G	222,003	44,525	47,753	2,049	1,916
3,840	MZILG3T8HCLS	SAS 12G	187,273	25,619	33,914	1,075	1,062
		SAS 24G	244,313	45,247	54,927	2,049	1,912
7,680	MZILG7T6HBLA	SAS 12G	187,183	25,636	33,940	1,075	1,062
		SAS 24G	246,245	45,190	54,780	2,049	1,913
15,360	MZILG15THBLA	SAS 12G	187,298	25,653	33,915	1,075	1,062
		SAS 24G	229,145	45,292	54,954	2,048	1,909
□ PCIe SSD (WI)							
400	SSDPF21Q400GB	PCIe4 x4	303,783	91,576	84,727	6,693	4,562
800	SSDPF21Q800GB	PCIe4 x4	290,266	99,852	94,882	6,738	4,512
1,600	SSDPF21Q016TB	PCIe4 x4	304,687	108,995	110,292	6,682	4,382
□ PCIe SSD (MU) (*1)							
1,600	KCMY1VUG1T60	PCIe4 x4	431,300	57,900	50,400	7,204	3,430
3,200	KCMY1VUG3T20	PCIe4 x4	557,564	95,486	86,744	7,219	5,079
6,400	KCMY1VUG6T40	PCIe4 x4	557,874	109,610	102,691	7,219	5,013
12,800	KCMY1VUG12T8	PCIe4 x4	558,473	103,865	98,998	6,728	5,111
□ PCIe SSD (RI) (*1)							
1,920	KCMY1RUG1T92	PCIe4 x4	431,394	57,935	50,484	7,204	3,430
3,840	KCMY1RUG3T84	PCIe4 x4	557,352	95,493	86,690	6,963	4,406
7,680	KCMY1RUG7T68	PCIe4 x4	609,834	107,833	98,803	7,041	4,416
15,360	KCMY1RUG15T3	PCIe4 x4	557,277	103,784	100,005	7,183	4,429

(*1) Performance value for PRAID EP680i connection. The drive supports PCIe 5.0, but the interface operates at PCIe 4.0.

Model common

Capacity [GB]	Storage device	Interface	Transactions [IO/s]			Throughput [MiB/s]	
			Database	Fileserver	Filecopy	Streaming	Restore
□ M.2 SATA SSD (PDUAL CP300)							
240	MTFDDAV240TGA	SATA 6G	45,009	5,324	5,490	474	353
480	MTFDDAV480TGA	SATA 6G	48,771	5,870	6,022	501	484
960	MTFDDAV960TGA	SATA 6G	51,373	6,252	6,429	471	486
□ M.2 NVMe SSD (PDUAL CP300)							
960	MTFDKBA960TFR	PCIe4 x4	139,598	31,160	25,761	4,923	1,380
□ M.2 NVMe SSD (M.2 Riser Kit)							
960	MTFDKBA960TFR	PCIe4 x2	147,090	31,373	25,555	3,482	1,382

VMmark V3

Benchmark description

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

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

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

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

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

In VMmark V3 there is an 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”.

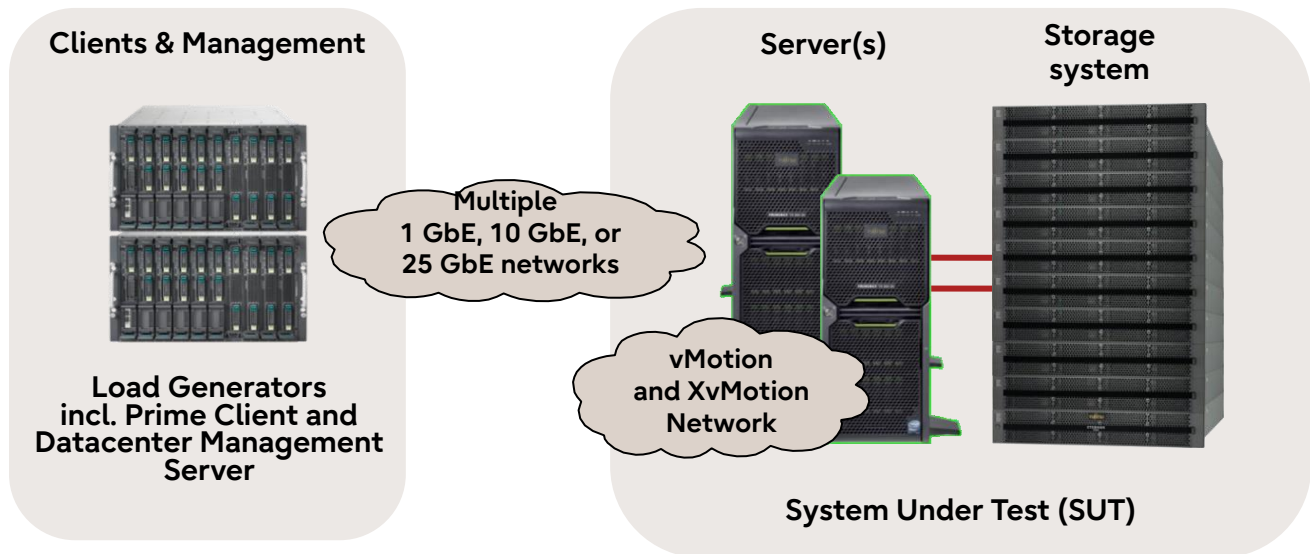
In the case of the two test types “Performance with Server Power” and “Performance with Server and Storage Power”, a so-called “Server PPKW Score” and “Server and Storage PPKW Score” are determined. These are the performance scores divided by the average power consumption in kilowatts (PPKW = performance 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 RX2450 M2
• Processor	2 x EPYC 9754 128C 360W
• Memory	6144 GB: 24 x 256 GB (1x256 GB) 8Rx4 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 QLE2772 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 M6 : 6 x Intel P4800X PCIe SSD (750 GB) 1 x PRIMERGY RX2540 M6 : 5 x Intel P4800X PCIe SSD (750 GB)

Software

• BIOS settings	See "Details"
• Operating system	VMware ESXi 8.0 Update 1a, Build 21813344
• Operating system settings	ESX settings: see "Details"

Details

See disclosure	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2024-02-06-Fujitsu-PRIMERGY-RX2450M2.pdf https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2024-02-06-Fujitsu-PRIMERGY-RX2450M2-serverPPKW.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

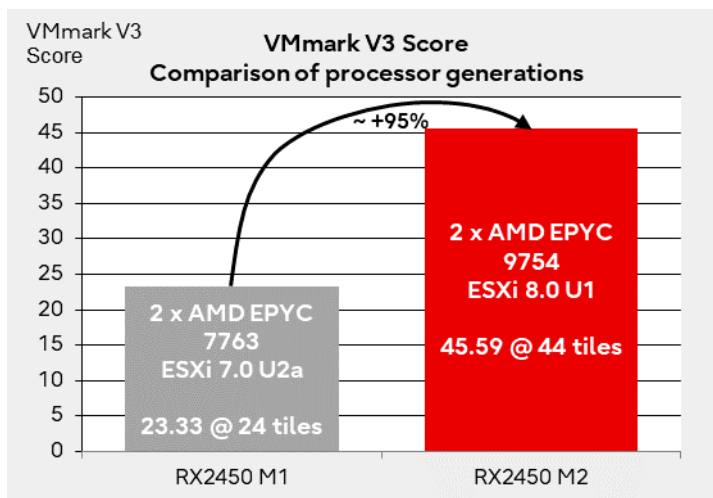
Hardware	
• Model	6 x PRIMERGY RX2530 M2 2 x PRIMERGY RX4770 M4
• Processor	2 x Intel Xeon E5-2699 v4 (4 x PRIMERGY RX2530 M2) 2 x Intel Xeon E5-2699A v4 (2 x PRIMERGY RX2530 M2) 4 x Intel Xeon Platinum 8180M (2 x PRIMERGY RX4770 M4)
• Memory	256 GB (6 x PRIMERGY RX2530 M2) 736 GB (1 x PRIMERGY RX4770 M4) 1504 GB (1 x PRIMERGY RX4770 M4)
• Network interface	6 x PRIMERGY RX2530 M2 1 x Emulex One Connect Oce14000 1GbE dual port PCIe adapter 1 x Emulex One Connect Oce14000 10GbE dual port PCIe adapter 1 x PRIMERGY RX4770 M4 1 x Intel OCP 1Gb quad port 1 x Emulex One Connect Oce14000 10GbE dual port PCIe adapter 1 x PRIMERGY RX4770 M4 1 x Intel OCP 10Gb dual port 1 x Intel X710-DA2 10GbE dual port PCIe adapter
Software	
• Operating system	VMware ESXi 7.0 Update 3c, Build 19193900

Benchmark results

"Performance Only" measurement results (February 6, 2024)

On February 6th, 2024, Fujitsu achieved a VMmark V3.1.1 score of "45.59@44 tiles" using PRIMERGY RX2450 M2 with EPYC 9754 processors and VMware ESXi 8.0 Update 1a. At this time, the system configuration had a total of 2 x 256 processor cores, and two identical servers were used for the "System Under Test" (SUT).

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 RX2450 M2 and the previous generation PRIMERGY RX2450 M1.

The PRIMERGY RX2450 M2 achieved a 95% improvement in score compared to the previous generation PRIMERGY RX2450 M1. This is due to the improved performance of the 4th generation AMD EPYC processor and the effective use of the capabilities of the VMware ESXi hypervisor.

For the latest VMmark V3 "Performance Only" results, detailed results, and configuration data, see

<https://www.vmware.com/products/vmmark/results3x.0.html>.

"Server Power-Performance" measurement results (February 6, 2024)



On February 6th, 2024, Fujitsu achieved a VMmark V3.1.1 "Server PPKW" score of "14.2452@44 tiles" using PRIMERGY RX2450 M2 with EPYC 9754 processors and VMware ESXi 8.0 Update 1a. It was the system configurations with a total of 2 x 256 processor cores, and two identical servers were used for the "System Under Test" (SUT). Based on the above results, PRIMERGY RX2450 M2 is rated as the most energy efficient two-socket rack-mount server in the world in the official VMmark V3 "Server Power-Performance" ranking (as of the date the benchmark results were published).

All comparisons for the competitor products reflect the status of the date of publication. For the latest VMmark V3 "Server Power-Performance" results, detailed results, and configuration data, see <https://www.vmware.com/products/vmmark/results3x.1.html>.

VMmark V4

Benchmark description

VMmark V4 is a benchmark developed by VMware to compare server configurations with hypervisor solutions from VMware regarding their suitability for server consolidation. After the discontinuation of the proven benchmark “VMmark V3” in July 2024, it has been succeeded by “VMmark V4”. In addition to the software for load generation, the benchmark consists of a defined load profile and binding regulations, as in VMmark V3. The benchmark results can be submitted to VMware and are published on their Internet site after a successful review process. VMmark V4 also requires 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, Storage vMotion, and XvMotion.

VMmark V4 is, like VMmark V3, a framework that consolidates already established benchmarks, as workloads in order to simulate the load of a virtualized consolidated server environment. On the other hand, VMmark 4 added application scenarios to better represent the complex workloads in today's enterprise data centers. As a result, four proven benchmarks, which cover the application scenarios Scalable web, E-commerce, Social Network, and Database, have been integrated in VMmark V4. In addition, new container technology, which configures today's data center workloads, was incorporated as well as traditional VM.

Application scenario	Load tool	# VMs
Scalable web simulation	Weathervane	13
E-commerce simulation	DVD Store 3.5	4
Social Network	DeathStarBench	1
Database simulation	NoSQLBench	3
Standby		1

Each of the four application scenarios is assigned to a total of 21 dedicated virtual machines. Then add to these a 22nd VM called the “standby server”. These 22 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 V4 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 V4 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.

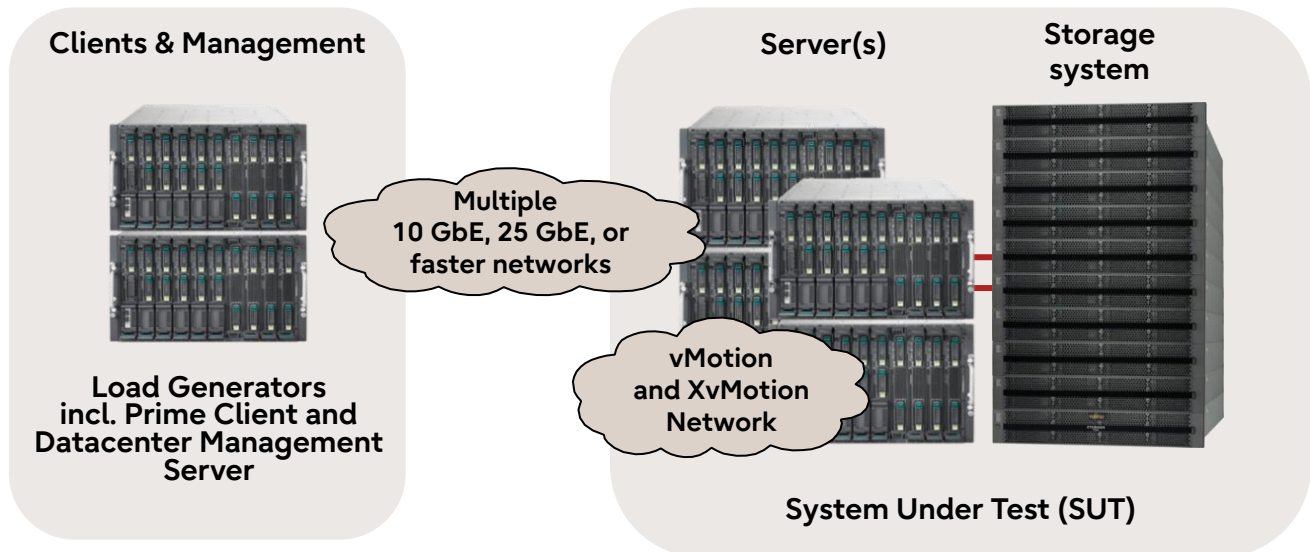
This score is determined from the individual results of the VMs and an infrastructure components result. Each of the nine VMmark V4 application or front-end VMs provides a specific benchmark result in the form of application-specific transaction rates for each VM. In order to derive a normalized score, 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 V4 tiles is always specified with each VMmark V4 score. The result is thus as follows: “Score@Number of Tiles”, for example “5.00@6 tiles”.

Note that the results of VMmark V4 cannot be compared with those of earlier versions such as 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 RX2450 M2
• Processor	2 x EPYC 9845 160C 390W
• Memory	6,144 GB: 24 x 256 GB (1x256 GB) 8Rx4 DDR5-5600 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 QLE2772 2X 32GFC PCIe v4 LP 5 x PRIMERGY RX2540 M6 configured as Fibre Channel targets 2 x PRIMERGY RX2540 M6 : 6 x Intel P4800X PCIe SSD (750 GB) 1 x Kioxia CM6-V PCIe SSD (3.2 TB) 2 x PRIMERGY RX2540 M6 : 3 x Intel P4800X PCIe SSD (750 GB) 1 x Kioxia CM6-V PCIe SSD (3.2 TB) 1 x PRIMERGY RX2540 M6 : 5 x Intel P4800X PCIe SSD (750 GB) 1 x Kioxia CM6-V PCIe SSD (3.2 TB)
Software	
• BIOS settings	See "Details"
• Operating system	VMware ESXi 8.0 Update 3, Build 24022510
• Operating system settings	ESX settings: see "Details"

Details

See disclosure <https://www.vmware.com/docs/2024-11-12-Fujitsu-PRIMERGY-RX2450M2>

Datacenter Management Server (DMS)

Hardware

• Model	1 x PRIMERGY TX2550 M7
• Processor	1 x Intel Xeon Gold 6554S
• Memory	384 GB
• Network interface	1 x Intel I350 1Gb onboard

Software

• Operating system	VMware ESXi 8.0 Update 2, Build 22380479
--------------------	--

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 8.0 U3, Build 24022515
--------------------	--

Load generator

Hardware

• Model	2 x PRIMERGY RX2540 M7 1 x PRIMERGY RX4770 M7 2 x PRIMERGY RX2530 M6
• Processor	2 x Intel Xeon Platinum 8480+ (2 x PRIMERGY RX2540 M7) 4 x Intel Xeon Platinum 8490H (1 x PRIMERGY RX4770 M7) 2 x Intel Xeon Platinum 8368 (2 x PRIMERGY RX2530 M6)
• Memory	512 GB (2 x PRIMERGY RX2540 M7) 2048 GB (1 x PRIMERGY RX4770 M7) 2048 GB (2 x PRIMERGY RX2530 M6)
• Network interface	2 x PRIMERGY RX2540 M7 1 x Intel I210 1Gb onboard 1 x Intel X710-DA2 10Gb dual port PCIe adapter 1 x PRIMERGY RX4770 M7 1 x Intel I210 1Gb onboard 2 x Intel X710-DA2 10Gb dual port PCIe adapter 2 x PRIMERGY RX2530 M6 1 x Intel I350 1Gb quad port OCPv3 1 x Intel X710-DA2 10GbE dual port PCIe adapter

Software

• Operating system	VMware ESXi 8.0 Update 2, Build 22380479
--------------------	--

Benchmark result

On November 12th, 2024, Fujitsu achieved a VMmark V4.0.1 score of "5.00@6 tiles" using PRIMERGY RX2450 M2 with EPYC 9845 processors and VMware ESXi 8.0 Update 3. Two identical servers were used for the "System Under Test" (SUT) and the system configuration for the measurement had a total of 2 x 320 processor cores.

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 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 results of VMmark V4 are not comparable to those of the previous versions, such as VMmark V3. For the latest VMmark V4 results, detailed results, and configuration data, see <https://www.vmware.com/products/vmmark/results4x>.


Literature


PRIMERGY Servers

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

PRIMERGY RX2450 M2

This Whitepaper

 <https://docs.ts.fujitsu.com/dl.aspx?id=1f65fb4f-107c-4a8b-8ebf-2a41d9843103>

 <https://docs.ts.fujitsu.com/dl.aspx?id=74eb84bf-a965-4fa6-9223-3b2045f7b479>

Data sheet:

<https://docs.ts.fujitsu.com/dl.aspx?id=de08aeef-312a-4876-9a2e-6852dc15353e>

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/>

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>

Benchmark overview

SAP SD: <http://docs.ts.fujitsu.com/dl.aspx?id=0a1e69a6-e366-4fd1-a1a6-0dd93148ea10>

VMmark

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

Benchmark Overview VMmark V3

<https://docs.ts.fujitsu.com/dl.aspx?id=e6f9973c-90d6-47c6-b317-e388a978bfb7>

Document change history

Version	Date	Description
1.3	2025-01-28	Update: <ul style="list-style-type: none"> • Technical data • SPEC CPU2017, STREAM Measured and calculated with AMD EPYC 9005 Series Processor
1.2	2025-01-14	Update: <ul style="list-style-type: none"> • SAP SD Standard Application Benchmark Measured with EPYC 9845
1.1	2024-11-26	Update: <ul style="list-style-type: none"> • Technical data • SPEC CPU2017, STREAM Measured with AMD EPYC 9005 Series Processor • SPECpower_ssj2008 Measured with EPYC 9845 • Disk I/O Measured with 2.5 / 3.5 inch model • VMmark V4 Measured with EPYC 9845
1.0	2024-04-30	New: <ul style="list-style-type: none"> • Technical data • SPEC CPU2017, STREAM Measured with AMD EPYC 9004 Series Processor • SPECpower_ssj2008 Measured with EPYC 9754 • SAP SD Standard Application Benchmark Measured with EPYC 9654 • Disk I/O Measured with 2.5 / 3.5 inch model • VMmark V3 Measured with EPYC 9754

Contact

Fujitsu

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

PRIMERGY Performance and Benchmarks

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

© Fujitsu 2024. All rights reserved. Fujitsu and Fujitsu logo are trademarks of Fujitsu Limited registered in many jurisdictions worldwide. Other product, service and company names mentioned herein may be trademarks of Fujitsu or other companies. This document is current as of the initial date of publication and subject to be changed by Fujitsu without notice. This material is provided for information purposes only and Fujitsu assumes no liability related to its use.