White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY RX4770 M4

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY PRIMERGY RX4770 M4.

The PRIMERGY RX4770 M4 performance data are compared with the data of other PRIMERGY models and discussed. In addition to the benchmark results, an explanation has been included for each benchmark and for the benchmark environment.

Version
1.3 2018/11/30

Contents

Document history

Version 1.0 (2018/01/19)

New:

- Technical data
- SPECcpu2006
  Measurements with and results for Intel® Xeon® Processor Scalable Family
- SPECpower_ssj2008
  Measurements with Intel® Xeon® Processor Platinum 8176M.
- OLTP-2
  Results for Intel® Xeon® Processor Scalable Family.
- vServCon
  Measurements with and results for Intel® Xeon® Processor Scalable Family.
- STREAM
  Measurements with and results for Intel® Xeon® Processor Scalable Family.
- LINPACK
  Measurements with and results for Intel® Xeon® Processor Scalable Family.

Version 1.1 (2018/04/10)

New:

- VMmark V3
  "Performance Only" measurement with Intel® Xeon® Platinum 8180
  "Performance with Server Power" measurement with Intel® Xeon® Platinum 8180

Version 1.2(2018/07/24)

Update:

- VMmark V3
  "Performance Only" updates measurement with Intel® Xeon® Platinum 8180
  "Performance with Server Power" updates measurement with Intel® Xeon® Platinum 8180

Version 1.3 (2018/11/30)

New:

- SAP Benchmarks
  Measurement with Intel® Xeon® Platinum 8180

Technical data

PRIMERGY RX4770 M4

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 M4
Model versions	PY RX4770 M4 PY RX4770 M4 LC PY RX4770 M4 Performance
Form factor	Rack server
Chipset	Intel® C624 Chipset
Number of sockets	4
Number of processors orderable	2 or 4
Processor type	Intel® Xeon® Processor Scalable Family
Number of memory slots	48
Maximum memory configuration	6 TB
Onboard LAN controller	DynamicLoM based on Intel® C624 (Intel® X722)
PCI slots	PY RX4770 M4 : PCI-Express 3.0 x16 × 8 PY RX4770 M4 LC : PCI-Express 3.0 x16 × 5 PY RX4770 M4 Performance : PCI-Express 3.0 x16 × 6
Max. number of internal hard disks	PY RX4770 M4 : 2.5" × 16 PY RX4770 M4 LC : 2.5" × 16 PY RX4770 M4 Performance : 2.5" × 8

Processors (since system release)								
Processor	Cores	Threads	Cache [MB]	UPI Speed [GT/s]	Rated Frequency [Ghz]	Max. Turbo Frequency [Ghz]	Max. Memory Frequency [MHz]	TDP [Watt]
Xeon Gold 5115	10	20	13.8	10.4	2.4	3.2	2400	85
Xeon Gold 5118	12	24	16.5	10.4	2.3	3.2	2400	105
Xeon Gold 5120	14	28	19.3	10.4	2.2	3.2	2400	105
Xeon Gold 6130	16	32	22.0	10.4	2.1	3.7	2666	125
Xeon Gold 6140	18	36	24.8	10.4	2.3	3.7	2666	140
Xeon Gold 6138	20	40	27.5	10.4	2.0	3.7	2666	125
Xeon Gold 6148	20	40	27.5	10.4	2.4	3.7	2666	150
Xeon Gold 6152	22	44	30.3	10.4	2.1	3.7	2666	140
Xeon Platinum 8153	16	32	22.0	10.4	2.0	2.8	2666	125
Xeon Platinum 8160	24	48	33.0	10.4	2.1	3.7	2666	150
Xeon Platinum 8164	26	52	35.8	10.4	2.0	3.7	2666	150
Xeon Platinum 8170	26	52	35.8	10.4	2.1	3.7	2666	165
Xeon Platinum 8176	28	56	38.5	10.4	2.1	3.8	2666	165
Xeon Platinum 8180	28	56	38.5	10.4	2.5	3.8	2666	205
Xeon Gold 5122	4	8	16.5	10.4	3.6	3.7	2666	105
Xeon Gold 6128	6	12	19.3	10.4	3.4	3.7	2666	115
Xeon Gold 6134	8	16	24.8	10.4	3.2	3.7	2666	130
Xeon Gold 6144	8	16	24.8	10.4	3.5	4.2	2666	150
Xeon Gold 6126	12	24	19.3	10.4	2.6	3.7	2666	125
Xeon Gold 6136	12	24	24.8	10.4	3.0	3.7	2666	150
Xeon Gold 6146	12	24	24.8	10.4	3.2	4.2	2666	165
Xeon Gold 6132	14	28	19.3	10.4	2.6	3.7	2666	140
Xeon Gold 6142	16	32	22.0	10.4	2.6	3.7	2666	150
Xeon Gold 6150	18	36	24.8	10.4	2.7	3.7	2666	165
Xeon Gold 6154	18	36	24.8	10.4	3.0	3.7	2666	200
Xeon Platinum 8156	4	8	16.5	10.4	3.6	3.7	2666	105
Xeon Platinum 8158	12	24	24.8	10.4	3.0	3.7	2666	150
Xeon Platinum 8168	24	48	33.0	10.4	2.7	3.7	2666	205
Xeon Gold 6134M	8	16	24.8	10.4	3.2	3.7	2666	130
Xeon Gold 6140M	18	36	24.8	10.4	2.3	3.7	2666	140
Xeon Gold 6142M	16	32	22.0	10.4	2.6	3.7	2666	150
Xeon Platinum 8160M	24	48	33.0	10.4	2.1	3.7	2666	150
Xeon Platinum 8170M	26	52	35.8	10.4	2.1	3.7	2666	165
Xeon Platinum 8176M	28	56	38.5	10.4	2.1	3.8	2666	165
Xeon Platinum 8180M	28	56	38.5	10.4	2.5	3.8	2666	205

All the processors that can be ordered with the PRIMERGY RX4770 M4 support Intel® Turbo Boost Technology 2.0. This technology allows you to operate the processor with higher frequencies than the nominal frequency. Listed in the processor table is "Max. Turbo Frequency" for the theoretical maximum frequency with only one active core per processor. The maximum frequency that can actually be achieved depends on the number of active cores, the current consumption, electrical power consumption, and the temperature of the processor.

As a matter of principle, Intel does not guarantee that the maximum turbo frequency can be reached. This is related to manufacturing tolerances, which result in a variance regarding the performance of various examples of a processor model. The range of the variance covers the entire scope between the nominal frequency and the maximum turbo frequency.

The turbo functionality can be set via BIOS option. Fujitsu generally recommends leaving the "Turbo Mode" option set at the standard setting of "Enabled", as performance is substantially increased by the higher frequencies. However, since the higher frequencies depend on general conditions and are not always guaranteed, it can be advantageous to disable the "Turbo Mode" option for application scenarios with intensive use of AVX instructions and a high number of instructions per clock unit, as well as for those that require constant performance or lower electrical power consumption.

Memory modules (since system release)								
Memory module	Capacity [GB]	Ranks	Bit width of the memory chips	Frequency [MHz]	Low voltage	Load reduced	Registered	ECC
8 GB (1x8 GB) 2Rx8 DDR4-2666 R ECC	8	2	8	2666			✓	✓
16 GB (1x16 GB) 2Rx8 DDR4-2666 R ECC	16	2	8	2666			✓	✓
8 GB (1x8 GB) 1Rx4 DDR4-2666 R ECC	8	1	4	2666			✓	✓
16 GB (1x16 GB) 1Rx4 DDR4-2666 R ECC	16	1	4	2666			✓	✓
16 GB (1x16 GB) 2Rx4 DDR4-2666 R ECC	16	2	4	2666			✓	✓
32 GB (1x32 GB) 2Rx4 DDR4-2666 R ECC	32	2	4	2666			✓	✓
64 GB (1x64 GB) 4Rx4 DDR4-2666 3DS ECC	64	4	4	2666			✓	✓
128 GB (1x128 GB) 8Rx4 DDR4-2666 3DS ECC	128	8	4	2666			✓	✓
64 GB (1x64 GB) 4Rx4 DDR4-2666 LR ECC	64	4	4	2666		✓	✓	✓

Power supplies (since system release)	Max. number
Modular PSU 1600 W platinum hp	2

Some components may not be available in all countries or sales regions.
Detailed technical information is available in the data sheet PRIMERGY RX4770 M4.
.

SPECcpu2006

Benchmark description

SPECcpu2006 is a benchmark which measures the system efficiency with integer and floating-point operations. It consists of an integer test suite (SPECint2006) containing 12 applications and a floating-point test suite (SPECfp2006) containing 17 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.

SPECcpu2006 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.

SPECcpu2006 contains two different performance measurement methods: The first method (SPECint2006 or SPECfp2006) determines the time which is required to process a single task. The second method (SPECint_rate2006 or SPECfp_rate2006) 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”, which differ in the use of compiler optimization. When publishing the results, the base values are always used and the peak values are optional.

Benchmark	Arithmetic	Type	Compiler optimization	Measurement result	Application
SPECint2006	integer	peak	aggressive	Speed	single-threaded
SPECint_base2006	integer	base	conservative		
SPECint_rate2006	integer	peak	aggressive	Throughput	multi-threaded
SPECint_rate_base2006	integer	base	conservative		
SPECfp2006	floating point	peak	aggressive	Speed	single-threaded
SPECfp_base2006	floating point	base	conservative		
SPECfp_rate2006	floating point	peak	aggressive	Throughput	multi-threaded
SPECfp_rate_base2006	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. Value “1” was defined for the SPECint_base2006, SPECint_rate_base2006, SPECfp_base2006, and SPECfp_rate_base2006 results of the reference system. For example, a SPECint_base2006 value of 2 means that the measuring system has handled this benchmark twice as fast as the reference system. A SPECfp_rate_base2006 value of 4 means that the measuring system has handled this benchmark some 4/[# base copies] times faster than the reference system. “# base copies” specifies how many parallel instances of the benchmark have been executed.

Not every SPECcpu2006 measurement is submitted by us for publication at SPEC. This is why the SPEC web pages do not have every result. As we archive the log files for all measurements, we can prove the correct implementation of the measurements at any time.

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX4770 M4
Processor	4 × Intel® Xeon® Processor Scalable Family
Memory	48 × 16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC
Software	
BIOS settings	Xeon Platinum 81xx, Gold 61xx: DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Package C State limit = C0 HWPM Support = Disabled Link Frequency Select = 10.4 GT/s Stale AtoS = Enabled LLC Dead Line Alloc = Disabled Patrol Scrub = Disabled IMC Interleaving = 1-way Fan Control = Full
Operating system	SUSE Linux Enterprise Server 12 SP2 (x86_64)
Operating system settings	Int_rate: Stack size set to unlimited using "ulimit -s unlimited" Set Kernel Boot Parameter: nohz_full=1-xx Set CPU frequency governor to maximum performance with: cpupower -c all frequency-set -g performance Set tmpfs filesystem with: mkdir /home/memory mount -t tmpfs -o size=768g,rw tmpfs /home/memory Process tuning settings: echo 10000000 > /proc/sys/kernel/sched_min_granularity_ns echo 15000000 > /proc/sys/kernel/sched_wakeup_granularity_ns echo 0 > /proc/sys/kernel/numa_balancing cpu idle state set with: cpupower idle-set -d 1 cpupower idle-set -d 2 Fp_rate: Stack size set to unlimited using "ulimit -s unlimited" Set Kernel Boot Parameter: nohz_full=1-xx Set CPU frequency governor to maximum performance with: cpupower -c all frequency-set -g performance Set tmpfs filesystem with: mkdir /home/memory mount -t tmpfs -o size=768g,rw tmpfs /home/memory Process tuning settings: echo 10000000 > /proc/sys/kernel/sched_min_granularity_ns echo 15000000 > /proc/sys/kernel/sched_wakeup_granularity_ns cpu idle state set with: cpupower idle-set -d 1 cpupower idle-set -d 2
Compiler	Int_rate: C/C++: Version 18.0.0.128 of Intel C/C++ Compiler for Linux fp_rate: C/C++: Version 17.0.3.191 of Intel C/C++ Compiler for Linux Fortran: Version 17.0.3.191 of Intel Fortran Compiler for Linux

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

Benchmark results

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

The results in italic are estimated values.

Processor	Number of processors	SPECint_rate_base2006	SPECint_rate2006
Xeon Gold 5115	4	*2015*	
Xeon Gold 5118	4	*2385*	
Xeon Gold 5120	4	*2694*	
Xeon Gold 6130	4	*3208*	
Xeon Gold 6140	4	*3619*	
Xeon Gold 6138	4	*3639*	
Xeon Gold 6148	4	*3989*	
Xeon Gold 6152	4	3840	
Xeon Platinum 8153	4	*2817*	
Xeon Platinum 8160	4	*4462*	
Xeon Platinum 8164	4	*4565*	
Xeon Platinum 8170	4	4510	
Xeon Platinum 8176	4	*5017*	
Xeon Platinum 8180	4	5490	
Xeon Gold 5122	4	*1125*	
Xeon Gold 6128	4	1670	
Xeon Gold 6134	4	*2180*	
Xeon Gold 6144	4	2290	
Xeon Gold 6126	4	*2694*	
Xeon Gold 6136	4	*3043*	
Xeon Gold 6146	4	2950	
Xeon Gold 6132	4	3080	
Xeon Gold 6142	4	3370	
Xeon Gold 6150	4	3907	
Xeon Gold 6154	4	4130	
Xeon Platinum 8156	4	1140	
Xeon Platinum 8158	4		
Xeon Platinum 8168	4	4920	
Xeon Gold 6134M	4	*2180*	

Xeon Gold 6140M	4	*3619*	
Xeon Gold 6142M	4	*3370*	
Xeon Platinum 8160M	4	*4462*	
Xeon Platinum 8170M	4	*4510*	
Xeon Platinum 8176M	4	*5017*	
Xeon Platinum 8180M	4	*5490*	

On 16th January 2018 the PRIMERGY RX4770 M4 with four Xeon Platinum 8180 processors was ranked first in the 4-socket systems category for the benchmark SPECint_rate_base2006. The result can be found at https://www.spec.org/cpu2006/results/res2018q1/cpu2006-20171226-51617.html

Processor	Number of processors	SPECfp_rate_base2006	SPECfp_rate_base2006
Xeon Gold 5115	4	*1733*	
Xeon Gold 5118	4	*1969*	
Xeon Gold 5120	4	*2142*	
Xeon Gold 6130	4	*2518*	
Xeon Gold 6140	4	*2737*	
Xeon Gold 6138	4	*2756*	
Xeon Gold 6148	4	*2915*	
Xeon Gold 6152	4	2820	
Xeon Platinum 8153	4	*2320*	
Xeon Platinum 8160	4	*3094*	
Xeon Platinum 8164	4	*3153*	
Xeon Platinum 8170	4	3110	
Xeon Platinum 8176	4	*3332*	
Xeon Platinum 8180	4	3510	
Xeon Gold 5122	4	1059	
Xeon Gold 6128	4	1490	
Xeon Gold 6134	4	*1920*	
Xeon Gold 6144	4	1950	
Xeon Gold 6126	4	*2241*	
Xeon Gold 6136	4	*2439*	
Xeon Gold 6146	4	2420	
Xeon Gold 6132	4	2420	
Xeon Gold 6142	4	2570	
Xeon Gold 6150	4	2836	
Xeon Gold 6154	4	2940	
Xeon Platinum 8156	4	1070	
Xeon Platinum 8158	4		
Xeon Platinum 8168	4	3290	
Xeon Gold 6134M	4	*1920*	
Xeon Gold 6140M	4	*2737*	
Xeon Gold 6142M	4	*2570*	
Xeon Platinum 8160M	4	*3094*	
Xeon Platinum 8170M	4	*3110*	
Xeon Platinum 8176M	4	*3332*	
Xeon Platinum 8180M	4	3510	

The following two diagrams illustrate the throughput of the PRIMERGY RX4770 M4 in comparison to its predecessor PRIMERGY RX4770 M3 in their respective most performant configuration.

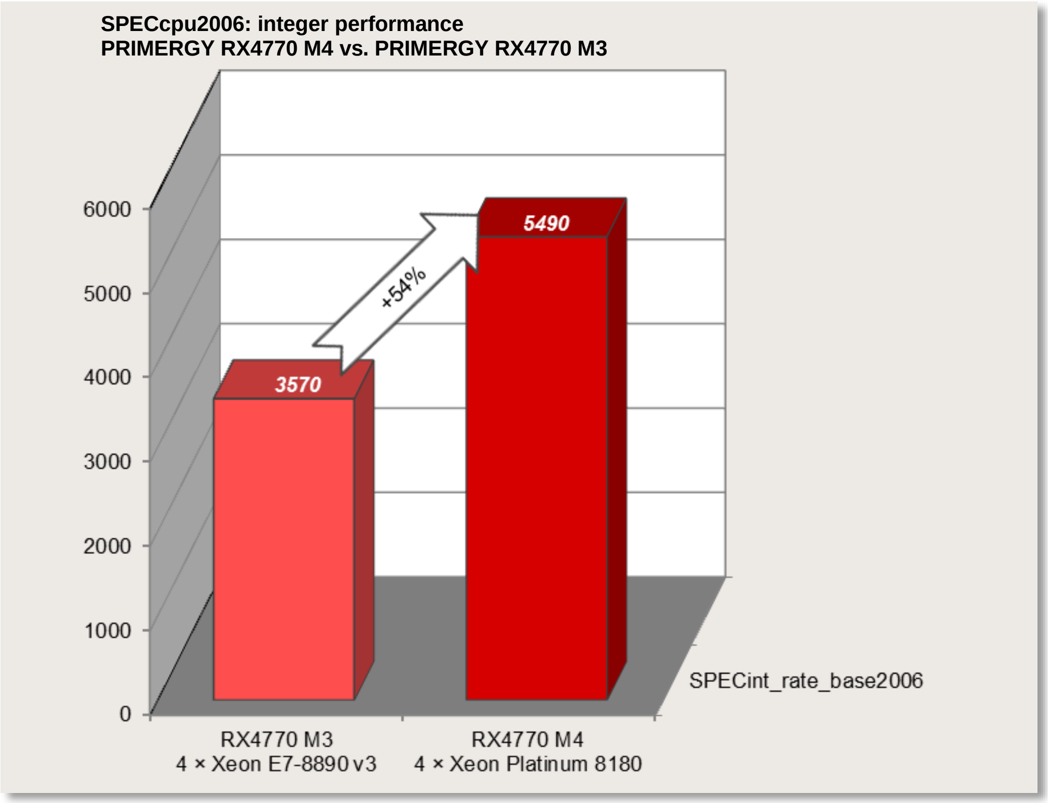

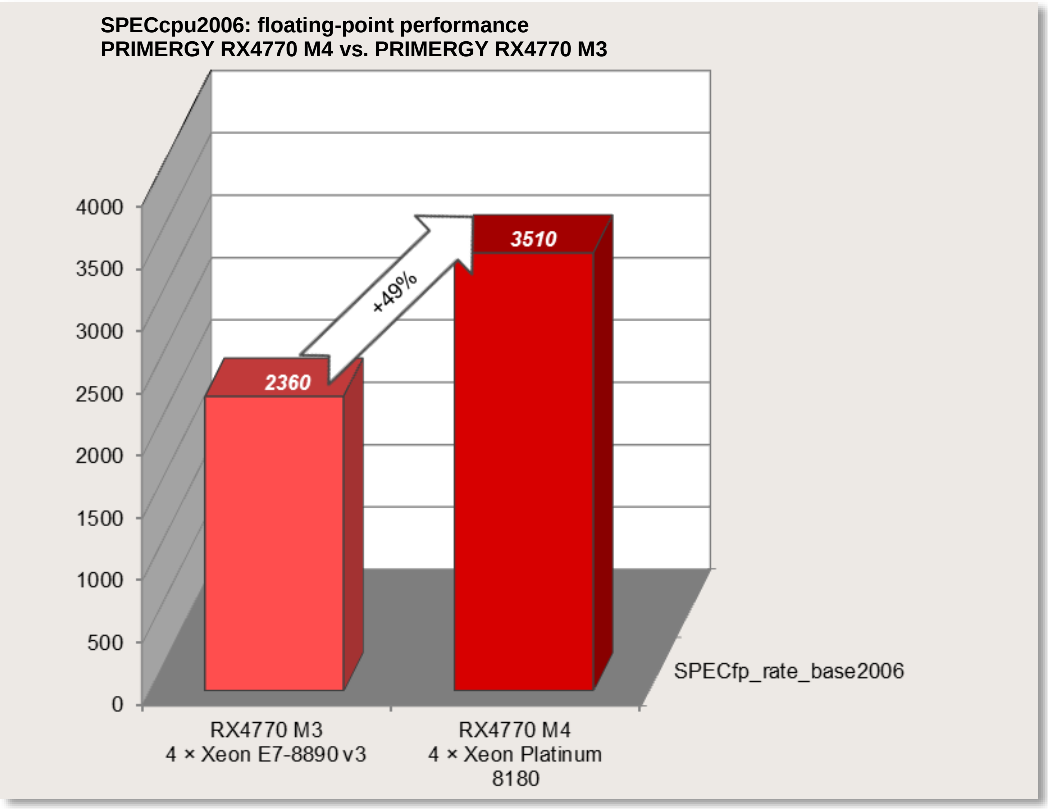

SPECpower_ssj2008

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.

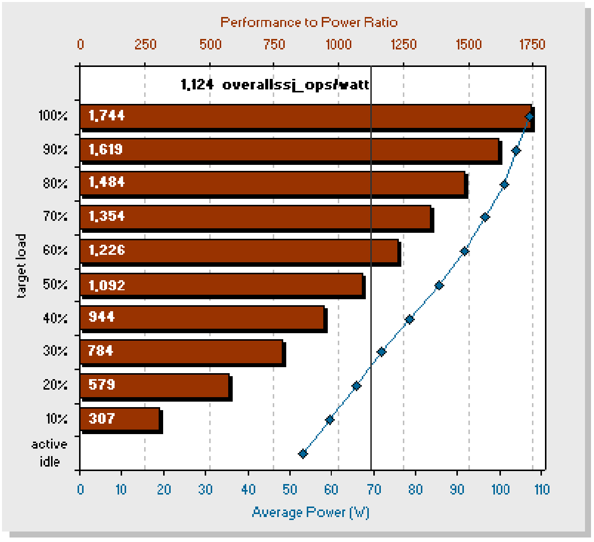

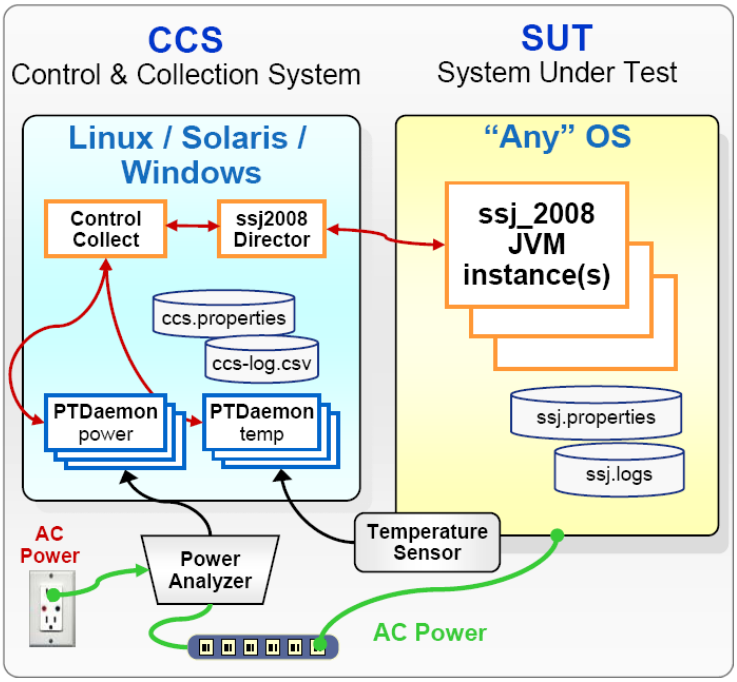

The benchmark runs on a wide variety of operating systems and hardware architectures, and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System 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 M4
Processor	Intel® Xeon® Platinum 8176M
Memory	8 GB (1x8 GB) 2Rx8 DDR4-2666 R ECC × 24
Network interface	1 ×PLAN EM 4 x 1Gb T OCP(Dynamic LOM)
Disk subsystem	Onboard SATA controller 1 ×SSD M.2 SATA 6Gbps 150 GB N H-P
Power Supply Unit	1 x Modular PSU 1600 W platinum hp
Software	
BIOS	R1.7.0
BIOS settings	SATA Controller = Disable Serial Port = Disabled Hardware Prefetcher = Disabled Adjacent Cache Line Prefetch = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Turbo Mode = Disabled Override OS Energy Performance = Enabled Energy Performance = Balanced Energy DDR Performance = Power balanced Autonomous C-state Support = Enabled Package C State limit = No Limit Link Frequency Select = 9.6 GT/s Uncore Frequency Override = Balanced Power IMC Interleaving = 1-way
Firmware	1.10P
Operating system	Microsoft Windows Server 2012 R2 Datacenter
Operating system settings	Set "Turn off hard disk after = 1 Minute" in OS. 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. SPECpower_ssj.props input.load_level.number_warehouses set to 224 due to a known inconsistency in processor reporting with this Java version.
JVM	Oracle Java HotSpot(TM) 64-Bit Server VM (build 24.80-b11, mixed mode), version 1.7.0_80
JVM settings	-server -Xmn1300m -Xms1400m -Xmx1400m -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:LoopUnrollLimit=45 -XX:InitialTenuringThreshold=12 -XX:MaxTenuringThreshold=15 -XX:ParallelGCThreads=2 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:FreqInlineSize=2500 -XX:+AggressiveOpts -XX:+UseLargePages -XX:+UseParallelOldGC

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

Benchmark results

The PRIMERGY RX4770 M4 achieved the following result:

SPECpower_ssj2008 = 12,828 overall ssj_ops/watt

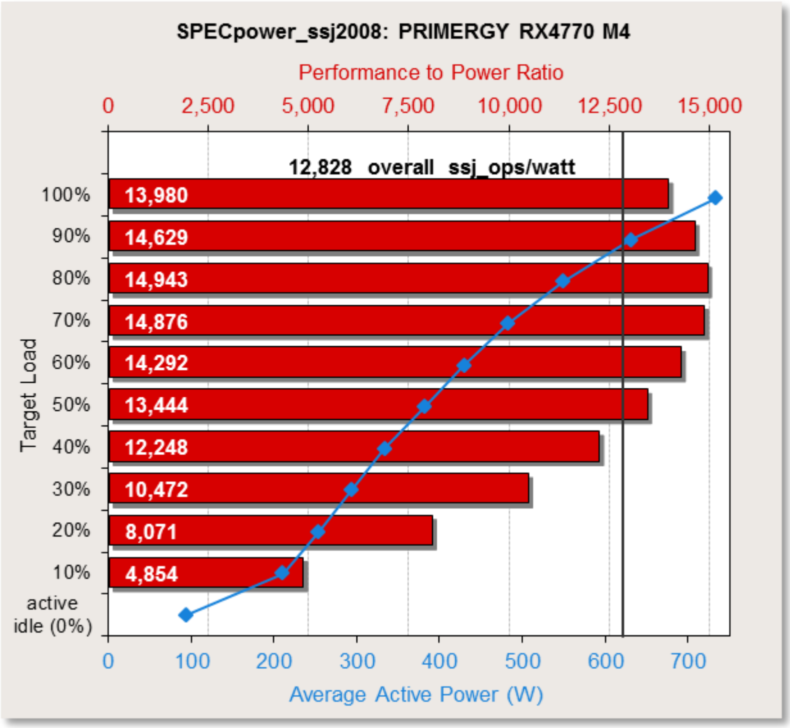

The adjoining diagram shows the result of the configuration described above. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level tagged on the y-axis of the diagram. The blue line shows the run of the curve for the average power consumption (bottom x-axis) at each target load level marked with a small rhomb. The black vertical line shows the benchmark result of 12,828 overall ssj_ops/watt for the PRIMERGY RX4770 M4. 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 %	10,249,832	733	13,980
90 %	9,232,029	631	14,629
80 %	8,208,471	549	14,943
70 %	7,180,090	483	14,876
60 %	6,151,365	430	14,292
50 %	5,128,907	382	13,444
40 %	4,096,830	334	12,248
30 %	3,074,915	294.	10,472
20 %	2,047,755	254	8,071
10 %	1,026,271	211	4,854
Active Idle	0	94.9	0
		∑ssj_ops / ∑power = 12,828	

SAP Benchmarks

Introduction

Since 1993 the SAP Standard Application Benchmarks have been developed by SAP to provide basic information for configuring, sizing and for platform comparison and to conduct platform certification. The first benchmarks available (for SAP R/3 Release 1.1H) were targeted for FI (Financial Accounting), SD (Sales and Distribution), and MM (Materials Management) followed by ATO (Assemble-To-Order), PP (Production Planning), WM (Warehouse Management), BW (Business Information Warehouse) and now many more.

The SAP Benchmark Council (established in 1995 and consisting of representatives of SAP as well as hardware, logo and technology partners involved in benchmarking) define and control the content of the benchmarks and establish rules that encompass the testing procedures. The procedures involve the hardware companies running most of the benchmarks and sending the results to SAP. On request SAP certifies the results.

By far the most popular benchmarks from the many available are the SAP SD as de-facto standard for SAP platforms, the BW Edition for SAP HANA benchmark, gaining more and more relevance with the increasing importance of SAP HANA, the BW Advanced Mixed Load (BW-AML) benchmark and the SAP Server Power benchmark, measuring power efficiency.

SAP Standard Application Benchmark

A SAP Standard Application Benchmark consists of script files that simulate the most typical transactions and workflow of a user scenario. It has a predefined SAP client database that contains sample company data against which the benchmark is run. The benchmark transactions of each component usually reflect the data throughput of an installation (for example orders, number of goods movements, etc.). However benchmark transactions do not reflect reporting because the resource consumption of a customer-defined report depends on the volume of data sought and is therefore not comparable. An exception is the BW benchmark which is mainly (but well defined) reporting activity.

Application components are customized for a benchmark run in such way that the system resource requirements are minimized while still representing an economic reality. Comparable customizing settings (buffer sizes, number of work processes etc.) can be found in live customer installations that need high data throughput.

In general each benchmark user has his own master data such as material, vendor or customer master data to avoid data locking situations. For most benchmarks a maximum of 100 parallel benchmark users can be simulated per client. The multi-tier client/server architecture consists of a database, application and presentation layer. The presentation layer – in real life the PC's of the logged-in users – is handled by one (or occasionally more) machines dubbed 'benchmark driver(s)'. There is no way for the application layer to tell whether it is driven by real users or a simulation environment.

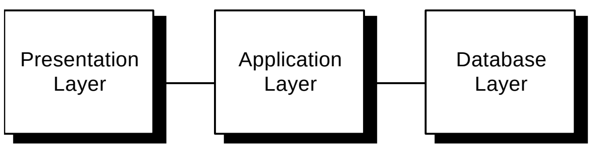

Figure 1: Multi-tier client/server architecture

Possible configurations for a benchmark simulation are:

a) 2-tier configuration architecture

Database and application layer reside on a single system – the simulation is driven by the presentation server (aka benchmark driver).

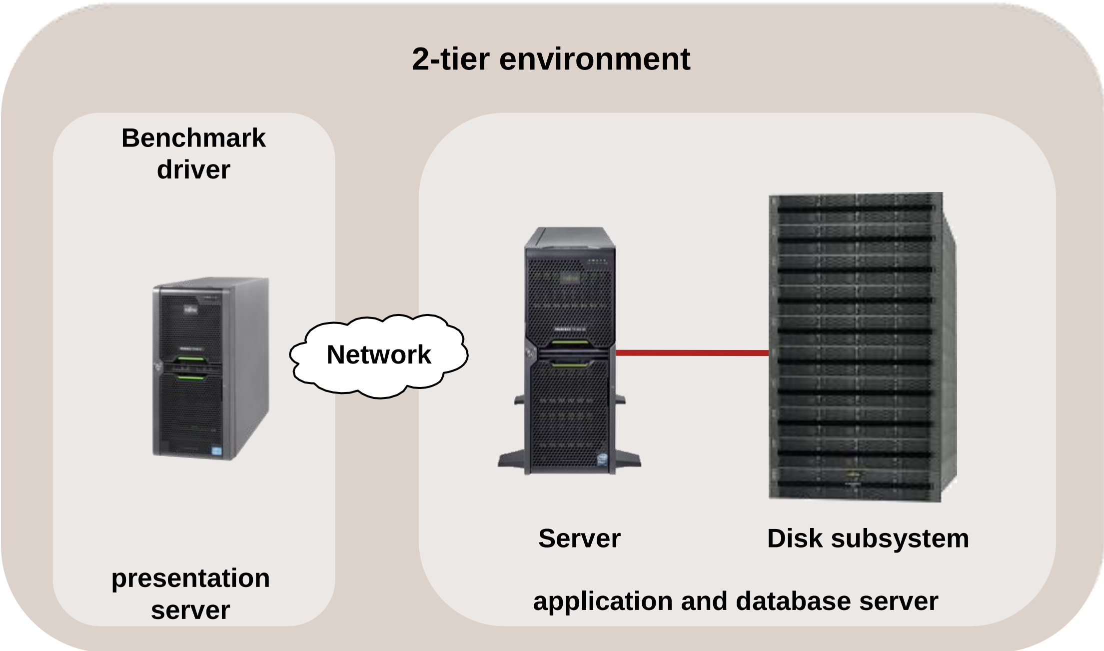

Figure 2: 2-tier configuration

b) 3-tier configuration architecture

Database and application layer reside on separate systems – the simulation is driven by a presentation server (aka benchmark driver).

Based on the architecture possible configurations are:

- 1 database server (or more - using parallel database techniques)
- n application servers with dedicated enqueue, update, message and dialog server
- n presentation servers (benchmark driver)

With such a configuration an impressive degree of scalability can be achieved.

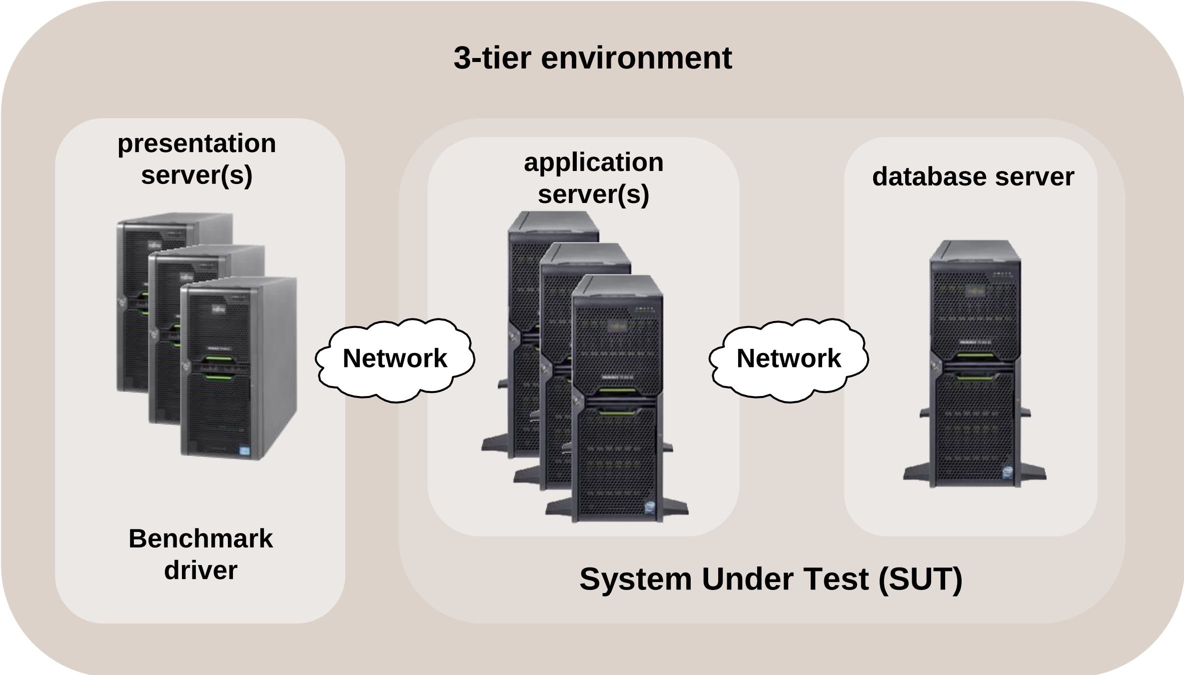

Figure 3: 3-tier configuration

The SAP Standard Application Benchmarks measure all performance relevant metrics such as database request times, wait times, CPU utilization, average dialog response times by a given number of benchmark users (with a fixed think time of 10 seconds between each dialog step) and the achieved throughput. The most significant parameters have to be part of every publication about SAP benchmarks initiated by the platform partners.

The following information must be part of a benchmark press release:

- The SAP® Business Suite component (e.g. “SAP ERP 6.0 with Enhancement Pack 4”)
- RDBMS and Operating System release
- Tested standard SAP® Business Suite components (FI, PP, SD or a combination of these)
- Number of tested benchmark users (if applicable)
- Average dialog response time in ‘n.nn sec’ (if applicable)
- Achieved throughput in dialog steps / hour or business numbers such as “accounts balanced”
- Type of client/server configuration
- A detailed description of hardware configuration (type, size of main memory, average CPU utilization and function for all servers involved in the benchmark)
- Confirmation that the benchmark is certified by SAP (e.g. “This benchmark fully complies with SAP's issued benchmark regulations and has been audited and certified by SAP.”)
- Reference where readers can get more information (e.g. “Details regarding this benchmark are available upon request from hardware partner or SAP AG”).

Benchmark Users and Average Dialog Response Time

A benchmark can only be certified if the average dialog response time is less than a fixed amount of time (think about it as system reaction time). More and more benchmark users are ‘switched’ to the system under test until the mean response time is outside the granted time frame.

Only SAP audited and certified benchmarks may be published by partners to ensure results that can be fairly compared with each other. A typical result would read like ‘2550 SD benchmark users with an average dialog response time of 0.96 seconds’.

Definition of SAPS

SAP has defined a unit for measuring throughput in a SAP Business Suite environment: **SAPS** (SAP Application Benchmark Performance Standard). Since SAPS is a very important unit that is often being used, it is essential to know the definition.

100 SAPS are defined as 2,000 fully processed order items per hour in the **SD standard application benchmark**. This throughput is achieved by processing 6,000 dialog steps (screen changes) and 2,000 postings per hour or processing 2,400 SAP transactions in the SD benchmark.

In the SD standard benchmark ‘fully processed’ means the full workflow of an order item (creating the order, creating a delivery note for this order, displaying the order, changing the delivery, posting a goods issue, listing orders and creating an invoice) has completed.

Benchmark Toolkit

In order to have a benchmark environment which enables fairly easy usage and reproducible results, a continuously maintained and updated toolkit is available.

The SAP Standard Application Benchmark tools and scripts of the toolkit are available for SAP Global Technology Partners only. **Their use by any other party is prohibited**. They are no official SAP product and no official support is available.

The Sales and Distribution (SD) benchmark

The Sales and Distribution benchmark is one of the most CPU consuming and memory demanding 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)

Table 1: Dialog steps of the standard Sales & Distribution (SD) benchmark

Each of the simulated users repeats this series of transactions from the start to the end of a benchmark run. During the so-called ramp-up phase the number of concurrently working users is increased until the expected limit (e.g. 2550) 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) 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.

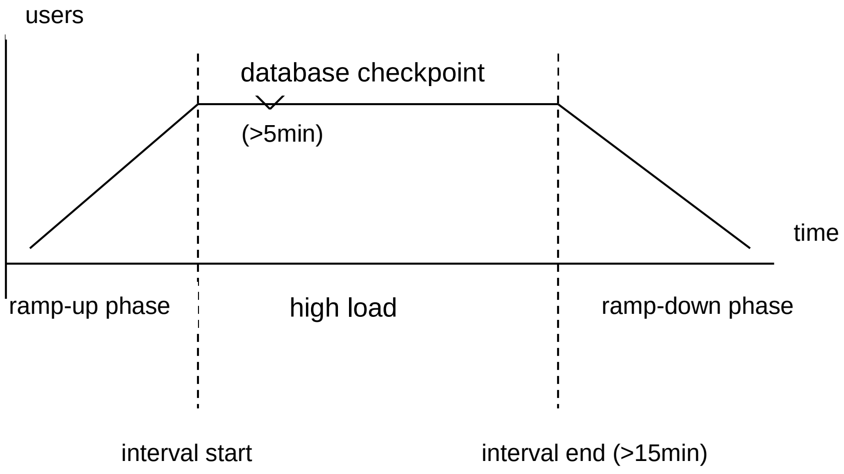

Figure 4: Benchmark run

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX4770 M4
Processor	Intel® Xeon® Platinum 8180 processor, 2.50 GHz, 4 processors
Memory	768 GB (16GB DDR4-2666 ECC × 48)
Network interface	LAN 10 Gbps × 2
Disk Subsystem	RX4770 M4 : on board RAID controller PRAID EP420i 300 GB 10k rpm SAS drive × 1, RAID0 (OS), 450 GB 15k rpm SAS drive × 3, RAID0 (log), PRAID EP420e × 1 JX40 × 1 : 960 GB SSD drive × 8, RAID0 (data) JX40 × 1 : 960 GB SSD drive × 8, RAID0 (data)
Software	
Operating system	Windows Server 2012 R2
Database	Microsoft SQL Server 2012 RTM
SAP Business Suite	SAP enhancement package 5 for SAP ERP 6.0
Benchmark configuration	2-tier

Benchmark results

Results	
Number of SAP SD benchmark users:	56,480
Average dialog response time:	0.98
Throughput:	
Fully processed order line items per hour:	6,175,000
Dialog steps per hour:	18,525,000
SAPS:	308,750
Average database request time (dialog/update):	0.011 sec / 0.019 sec
CPU utilization of central server:	99%

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

Certification date: 2018-01-24
Certification number: 2018004

The SAP BW Edition for SAP HANA Standard Application 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.

Comparability of SAP BW Benchmark Results

The SAP BWH Benchmark goes beyond the scope and features used in the BW-AML benchmark. While BW-AML focuses on traditional BW objects and processes supported by all database platforms, the SAP BWH Benchmark takes advantage of HANA's new features. The workload of the two benchmarks completely differ and thus are not comparable.

Within the SAP BWH Benchmark, both the data model and the query definitions have changed in the course of the development of version 1 to version 3. That is why the results of different versions must not be compared with each other.

BWH Benchmark version 3 for instance is available for SAP HANA 2.0 only and the Data Load KPI (Phase 1) has been changed - the runtime of the latest data set (1,3 billion records) will be the KPI.

In SAP BWH Benchmark version 2 it was allowed to activate a specific SAP HANA performance enhancement function which materializes intermediate query results. Benchmarks that have used the feature cannot be compared with benchmarks that haven't. Whether the function was used can be seen in the details on www.sap.com/benchmark , "Materializing of the intermediate result of the query was enabled (phase 2 and phase 3)" is shown.

In addition, SAP BWH Benchmark results with different number of data sets cannot be compared either.

Currently released version of the benchmark is version 3. Benchmarks with the older versions won't be certified anymore.

Features of the SAP BW Edition for SAP HANA Standard Application Benchmark

The SAP BWH Benchmark consists of 3 phases:

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

Data load phase

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

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

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

One of the central rules of the benchmark stipulates that the memory utilization must be at least 65 percent. The permissible data volumes are a multiple of 1.3 billion initial data records.

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

Query throughput phase

The queries for the throughput phase must be executed via an ABAP program with a variant containing 190 query steps. Users execute the set of navigation steps in random order (via asynchronous RFCs). The queries contain typical query patterns which can be found in BW productive systems of customers.

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

Query runtime phase

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

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

Metrics of the SAP BW Edition for SAP HANA Standard Application Benchmark

The SAP BW edition for SAP HANA Benchmark Version 3 KPIs are:

- Benchmark phase 1 - data load phase:

  Number of initial records

  Runtime of last data set in seconds

- Benchmark phase 2 - query throughput phase:

  Query executions per hour

  CPU utilization of database server in percent

- Benchmark phase 3 - query runtime phase:

  Total runtime of complex query phase in seconds

SAP HANA Scale-up and Scale-out Configuration Architecture

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.

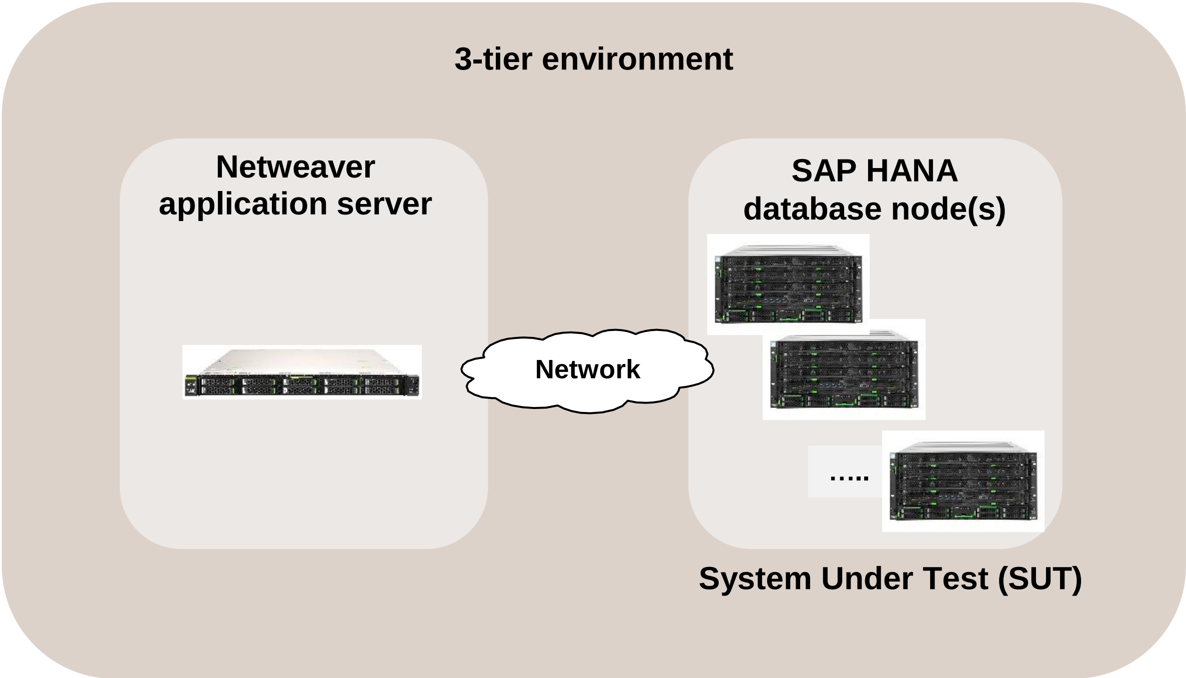

Although an application server is involved in the benchmark, neither performance metrics are measured nor does the server appear on the benchmark certificate.

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX4770 M4
Processor	Intel® Xeon® Platinum 8180 processor, 2.50 GHz, 4 processors
Memory	1,536 GB (32GB DDR4-2666 ECC × 48)
Network interface	On board LAN 1 Gbps × 1 LAN 10 Gbps × 1
Disk Subsystem	RX4770 M4 : On board RAID controller PRAID EP420i 600 GB 10k rpm SAS drive × 1, RAID0 (OS), PRAID EP420e × 3 JX40 × 1 : 960 GB SSD drive × 6, RAID0 (log) JX40 × 1 : 960 GB SSD drive × 8, RAID0 (data) JX40 × 1 : 960 GB SSD drive × 8, RAID0 (data)
Software	
Operating system	SUSE Linux Enterprise Server 12
Database	SAP HANA 2.0
Technology platform release	SAP NetWeaver 7.50
Benchmark version	Version 3

Benchmark results

Results	
Benchmark Phase 1 Number of initial records: Runtime of last Data Set (seconds):	 1,300,000,000 15,836
Benchmark Phase 2 Query Executions per Hour: CPU utilization of database server:	 10,898 96%
Benchmark Phase 3 Total Runtime of complex query phase (seconds):	 85

Certified and published SAP BWH Benchmarks are published on SAP's benchmark site here .
Certification date: 2018-11-09
Certification number: 2018048

Concluding remarks

Available now for over 20 years the CPU and memory demanding SAP SD benchmark has served well to demonstrate the strength or weakness of a hardware platform. SAP certified results will be published on SAP's web site and comprise information on machines, platform partner, SAP® Business Suite component and the number of benchmark users. Each of the tables contains basic information about the certified results. In line with SAP's full disclosure policy more detailed information about each result can be obtained from the respective hardware partner or SAP.

OLTP-2

Benchmark description

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

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

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

Further information can be found in the document Benchmark Overview OLTP-2.

Benchmark environment

The typical measurement set-up is illustrated below:

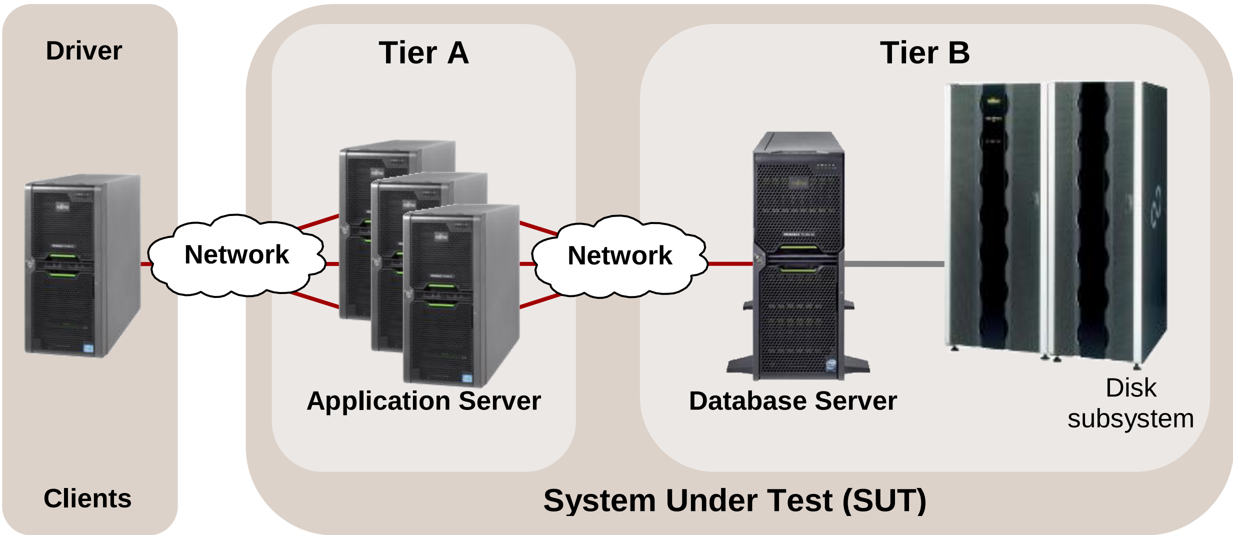

Database Server (Tier B)	
Hardware	
Model	PRIMERGY RX4770 M4
Processor	Intel® Xeon® Processor Scalable Family
Memory	2 processors: 24 × 64 GB (1x64 GB) 4Rx4 DDR4-2666 3DS ECC 4 processors: 48 × 64 GB (1x64 GB) 4Rx4 DDR4-2666 3DS ECC
Network interface	2 × onboard LAN 10 Gb/s
Disk subsystem	RX4770 M4: Onboard RAID controller PRAID EP420i 1 × 300 GB 10k rpm SAS Drive, RAID0 (OS), 6 × PRAID EP420e 6 × JX40: 14 × 960 GB SSD Drive, RAID0 (data) 5 × JX40: 14 × 400 GB SSD Drive, RAID0 (temp) 1 × JX40: 14 × 400 GB SSD Drive, RAID0 (temp) 8 × 900 GB 10k rpm SAS Drive, RAID10 (Log)
Software	
BIOS	Version R1.5.0
Operating system	Microsoft Windows Server 2016 Standard
Database	Microsoft SQL Server 2017 Enterprise

Application Server (Tier A)	
Hardware	
Model	1 × PRIMERGY RX2530 M2
Processor	2 × Xeon E5-2690 v4
Memory	128 GB, 2400 MHz registered ECC DDR4
Network interface	2 × onboard LAN 10 Gb/s 1 × Dual Port LAN 1 Gb/s
Disk subsystem	2 × 300 GB 10k rpm SAS Drive
Software	
Operating system	Microsoft Windows Server 2012 R2 Standard

Client	
Hardware	
Model	1 × PRIMERGY RX2530 M2
Processor	2 × Xeon E5-2667 v4
Memory	128 GB, 2400 MHz registered ECC DDR3
Network interface	1 × onboard Quad Port LAN 1 Gb/s
Disk subsystem	1 × 300 GB 10k rpm SAS Drive
Software	
Operating system	Microsoft Windows Server 2012 R2 Standard
Benchmark	OLTP-2 Software EGen version 1.14.0

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

Benchmark results

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

A guideline in the database environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses. This why a configuration with a total memory of 3072 GB was considered for the measurements with four processors and a configuration with a total memory of 1536 GB for the measurements with two processor. Both memory configurations have memory access of 2666 MHz..

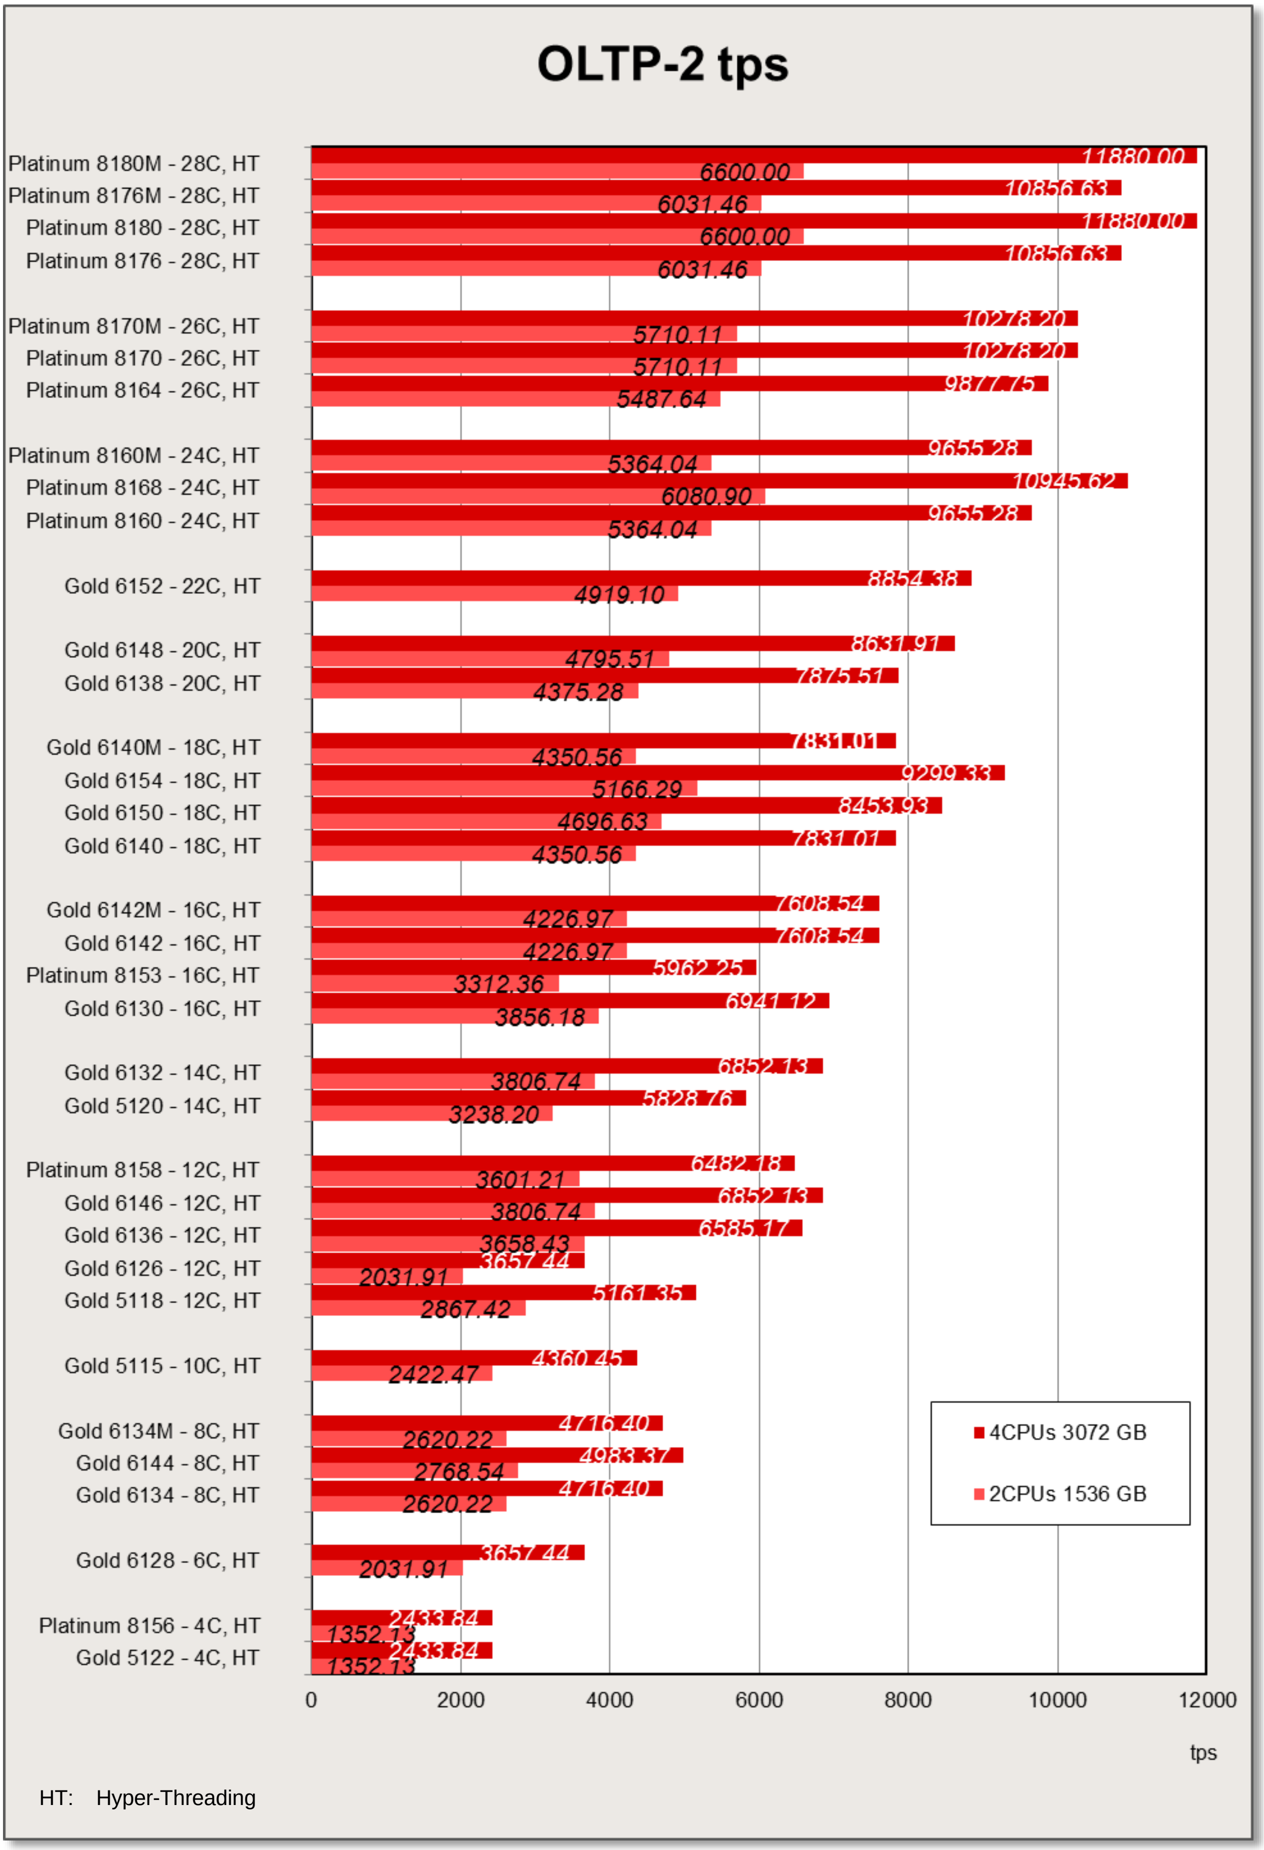
OLTP-2 tps
Platinum 8180M - 28C, HT
11880.00
6600.00
Platinum 8176M - 28C, HT
10856.63
6031.46
Platinum 8180 - 28C, HT
11880.00
6600.00
Platinum 8176 - 28C, HT
10856.63
6031.46
Platinum 8170M - 26C, HT
10278.20
5710.11
Platinum 8170 - 26C, HT
10278.20
5710.11
Platinum 8164 - 26C, HT
9877.75
5487.64
Platinum 8160M - 24C, HT
9655.28
5364.04
Platinum 8168 - 24C, HT
10945.62
6080.90
Platinum 8160 - 24C, HT
9655.28
5364.04
Gold 6152 - 22C, HT
8854.38
4919.10
Gold 6148 - 20C, HT
8631.91
4795.51
Gold 6138 - 20C, HT
7875.51
4375.28
Gold 6140M - 18C, HT
7831.01
4350.56
Gold 6154 - 18C, HT
9299.33
5166.29
Gold 6150 - 18C, HT
8453.93
4696.63
Gold 6140 - 18C, HT
7831.01
4350.56
Gold 6142M - 16C, HT
7608.54
4226.97
Gold 6142 - 16C, HT
7608.54
4226.97
Platinum 8153 - 16C, HT
5962.25
3312.36
Gold 6130 - 16C, HT
6941.12
3856.18
Gold 6132 - 14C, HT
6852.13
3806.74
Gold 5120 - 14C, HT
5828.76
3238.20
Platinum 8158 - 12C, HT
6482.18
3601.21
Gold 6146 - 12C, HT
6852.13
3806.74
Gold 6136 - 12C, HT
6585.17
3658.43
Gold 6126 - 12C, HT
3657.44
2031.91
Gold 5118 - 12C, HT
5161.35
2867.42
Gold 5115 - 10C, HT
4360.45
2422.47
Gold 6134M - 8C, HT
4716.40
2620.22
Gold 6144 - 8C, HT
4983.37
2768.54
Gold 6134 - 8C, HT
4716.40
2620.22
Gold 6128 - 6C, HT
3657.44
2031.91
Platinum 8156 - 4C, HT
2433.84
1352.13
Gold 5122 - 4C, HT
2433.84
1352.13
0
2000
4000
6000
8000
10000
12000
tps
4CPUs 3072 GB
2CPUs 1536 GB
HT: Hyper-Threading

The following diagram shows the OLTP-2 transaction rates that can be achieved with one and two processors of the Intel® Xeon® Processor Scalable Family.

It is evident that a wide performance range is covered by the variety of released processors. The features of the processors are summarized in the section “Technical data”.

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

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

If you compare the maximum achievable OLTP-2 values of the current system generation with the values that were achieved on the predecessor systems, the result is an increase of about 38%.

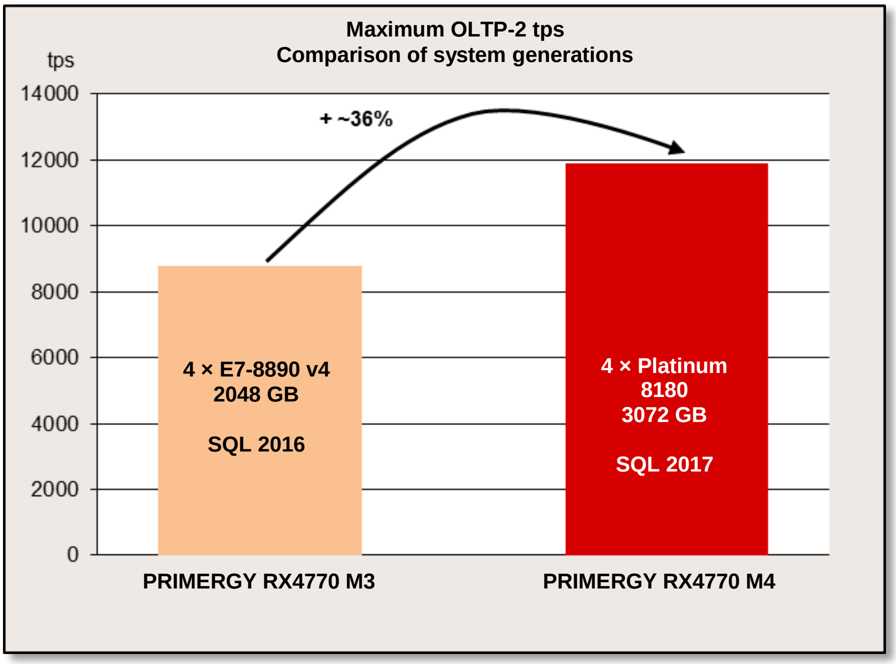

vServCon

Benchmark description

vServCon is a benchmark used by Fujitsu to compare server configurations with hypervisor with regard to their suitability for server consolidation. This allows both the comparison of systems, processors and I/O technologies as well as the comparison of hypervisors, virtualization forms, and additional drivers for virtual machines.

vServCon is not a new benchmark in the true sense of the word. It is more a framework that combines already established benchmarks (or in modified form) as workloads in order to reproduce the load of a consolidated and virtualized server environment. Three proven benchmarks are used which cover the application scenarios database, application server, and web server.

Application scenario	Benchmark	No. of logical CPU cores	Memory
Database	Sysbench (adapted)	2	1.5 GB
Java application server	SPECjbb (adapted, with 50% - 60% load)	2	2 GB
Web server	WebBench	1	1.5 GB

Each of the three application scenarios is allocated to a dedicated virtual machine (VM). A fourth machine, the so-called idle VM, is added to these. These four VMs make up a “tile”. Depending on the performance capability of the underlying server hardware, you may as part of a measurement also have to start several identical tiles in parallel in order to achieve a maximum performance score.

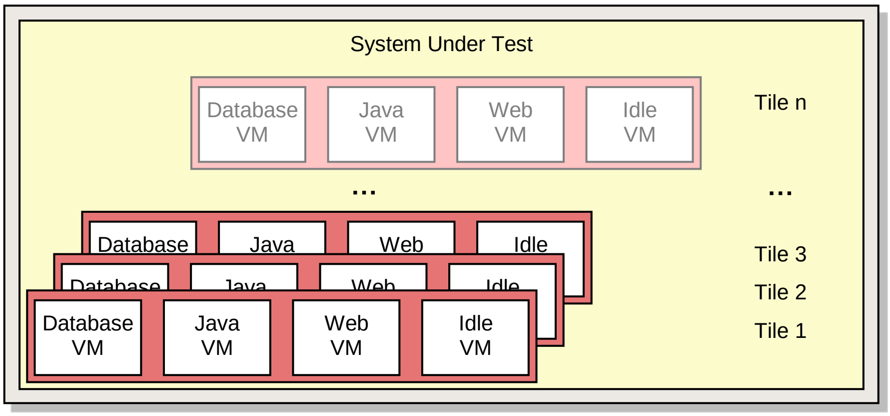

Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score, the individual benchmark result for one tile is put in relation to the respective result of a reference system. The resulting relative performance value is then suitably weighted and finally added up for all VMs and tiles. The outcome is a score for this tile number.

As a general rule, start with one tile, and this procedure is performed for an increasing number of tiles until no further significant increase in this vServCon score occurs. The final vServCon score is then the maximum of the vServCon scores for all tile numbers. This score thus reflects the maximum total throughput that can be achieved by running the mix defined in vServCon that consists of numerous VMs up to the possible full utilization of CPU resources. This is why the measurement environment for vServCon measurements is designed in such a way that only the CPU is the limiting factor and that no limitations occur as a result of other resources.

The progression of the vServCon scores for the tile numbers provides useful information about the scaling behavior of the “System under Test”.

A detailed description of vServCon is in the document: Benchmark Overview vServCon.

Benchmark environment

The measurement set-up is symbolically illustrated below:

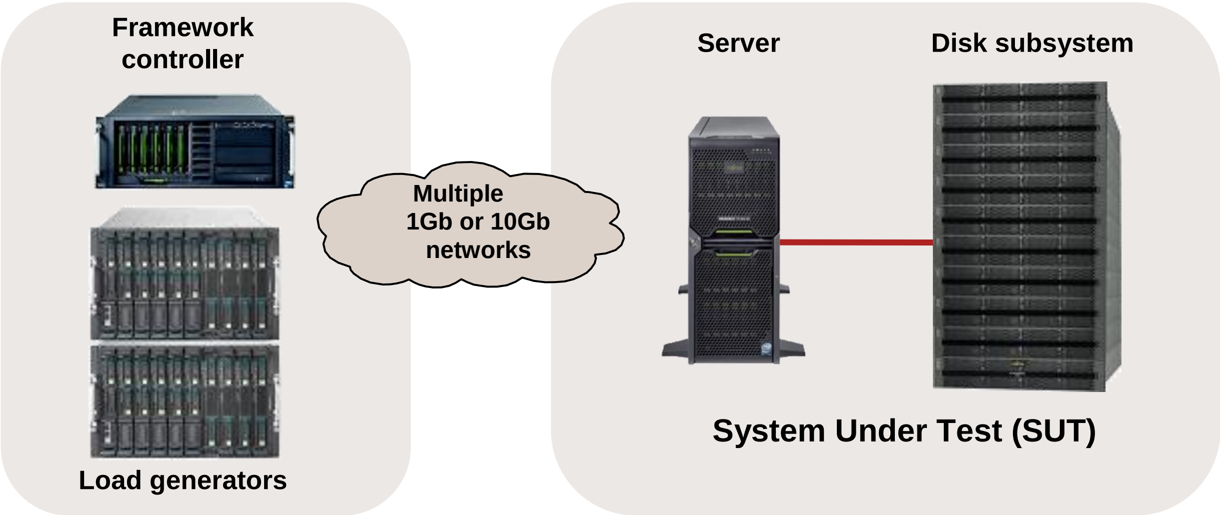

System Under Test (SUT)	
Hardware	
Processor	4 × Intel® Xeon® Processor Scalable Family
Memory	48 × 32 GB (1x32 GB) 2Rx4 DDR4-2666 R ECC
Network interface	1 × Emulex OneConnect OCe14000 Dual Port Adapter with 10Gb SFP+ DynamicLoM interface module
Disk subsystem	1 ×dual-channel FC controller Emulex LPe160021 LINUX/LIO based flash storage system
Software	
Operating system	VMware ESXi 6.5.0b Build 5146846

Load generator (incl. Framework controller)	
Hardware (Shared)	
Enclosure	5 x PRIMERGY RX2530 M2
Hardware	
Processor	2 × XeonE5-2683 v4
Memory	128 GB
Network interface	3 × 1 Gbit LAN
Software	
Operating system	VMware ESXi 6.0.0 U1b Build 3380124

Load generator VM (VMs on various servers)	
Hardware	
Processor	1 × logical CPU
Memory	4048 MB
Network interface	2 × 1 Gbit/s LAN
Software	
Operating system	Microsoft Windows Server 2008 Standard Edition 32bit

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

Benchmark results

The PRIMERGY dual-socket rack and tower systems dealt with here are based on processors of the Intel® Xeon® Processor Scalable Family. The features of the processors are summarized in the section "Technical data".

The available processors of these systems with their results can be seen in the following table.

Processor			Score	#Tiles
Intel® Xeon® Processor Scalable Family	4 Cores Hyper-Threading, Turbo-Modus	Gold 5122	22.2	14
		Platinum 8156	22.9	14
	6 Cores Hyper-Threading, Turbo-Modus	Gold 6128	33.1	22
	8 Cores Hyper-Threading, Turbo-Modus	Gold 6134	43.2	24
		Gold 6144	46.1	25
		Gold 6134M	46.1	24
	10 Cores Hyper-Threading, Turbo-Modus	Gold 5115	40.1	34
	12 Cores Hyper-Threading, Turbo-Modus	Gold 5118	47.4	37
		Gold 6126	60.4	38
		Gold 6136	53.5	40
		Gold 6146	55.2	40
		Platinum 8158	61.8	40
	14 Cores Hyper-Threading, Turbo-Modus	Gold 5120	53.5	35
		Gold 6132	61.5	36
	16 Cores Hyper-Threading, Turbo-Modus	Gold 6130	65.2	47
		Platinum 8153	54.7	46
		Gold 6142	68.5	48
		Gold 6142M	68.5	48
	18 Cores Hyper-Threading, Turbo-Modus	Gold 6140	71.8	52
		Gold 6150	77.6	53
		Gold 6154	83.6	54
		Gold 6140M	83.6	54
	20 Cores Hyper-Threading, Turbo-Modus	Gold 6138	72.3	48
		Gold 6148	78.9	49
	22 Cores Hyper-Threading, Turbo-Modus	Gold 6152	79.6	57
	24 Cores Hyper-Threading, Turbo-Modus	Platinum 8160	88.5	60
		Platinum 8168	88.5	60
		Platinum 8160M	101.4	62

	26 Cores Hyper-Threading, Turbo-Modus	Platinum 8164	90.6	60
		Platinum 8170	95.7	60
		Platinum 8170M	95.7	60
	28 Cores Hyper-Threading, Turbo-Modus	Platinum 8176	99.6	66
		Platinum 8180	113.6	67
		Platinum 8176M	99.6	66
		Platinum 8180M	113.6	67

These PRIMERGY dual-socket rack and tower systems are very suitable for application virtualization owing to the progress made in processor technology. Compared with a system based on the previous processor generation, approximately 50% higher virtualization performance can be achieved (measured in vServCon score in their maximum configuration).

The following diagram compares the virtualization performance values that can be achieved with the processors reviewed here.

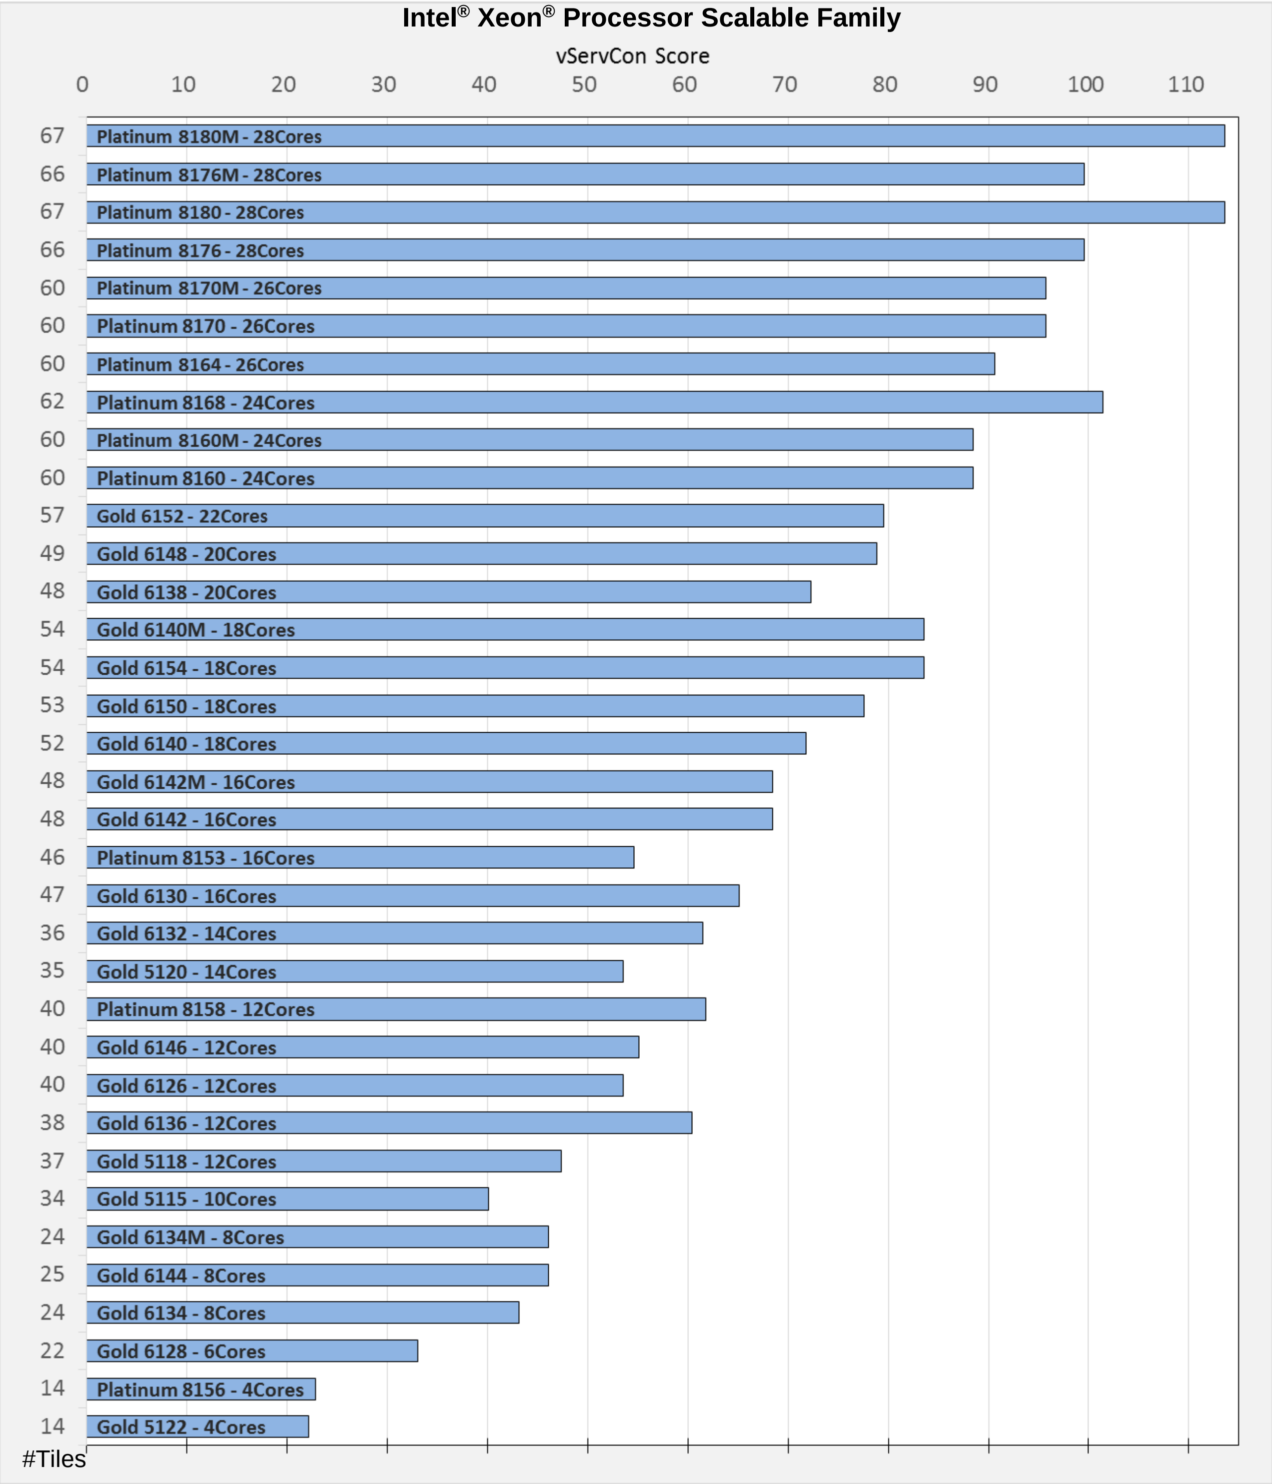

The relatively large performance differences between the processors can be explained by their features. The values scale on the basis of the number of cores, the size of the L3 cache and the CPU clock frequency and as a result of the features of Hyper-Threading and turbo mode, which are available in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

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

As a matter of principle, the memory access speed also influences performance. A guideline in the virtualization environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses. The vServCon scaling measurements presented here were all performed with a memory access speed – depending on the processor type – of at most 2666 MHz.

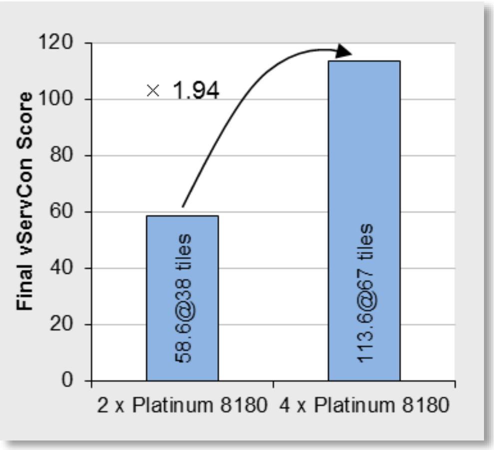

Until now, we have looked at the virtualization performance of a fully configured system. However, with a server with four sockets, the question also arises as to how good performance scaling is from two to four processors. The better the scaling, the lower the overhead usually caused by the shared use of resources within a server. The scaling factor also depends on the application. If the server is used as a virtualization platform for server consolidation, the system scales with a factor of 1.94. When operated with four processors, the system thus achieves twice the performance as with two processors, as is illustrated in this diagram using the processor version Xeon Platinum 8180 as an example.

The next diagram illustrates the virtualization performance for increasing numbers of VMs based on the Xeon Gold 6132 (14-Core) processors.

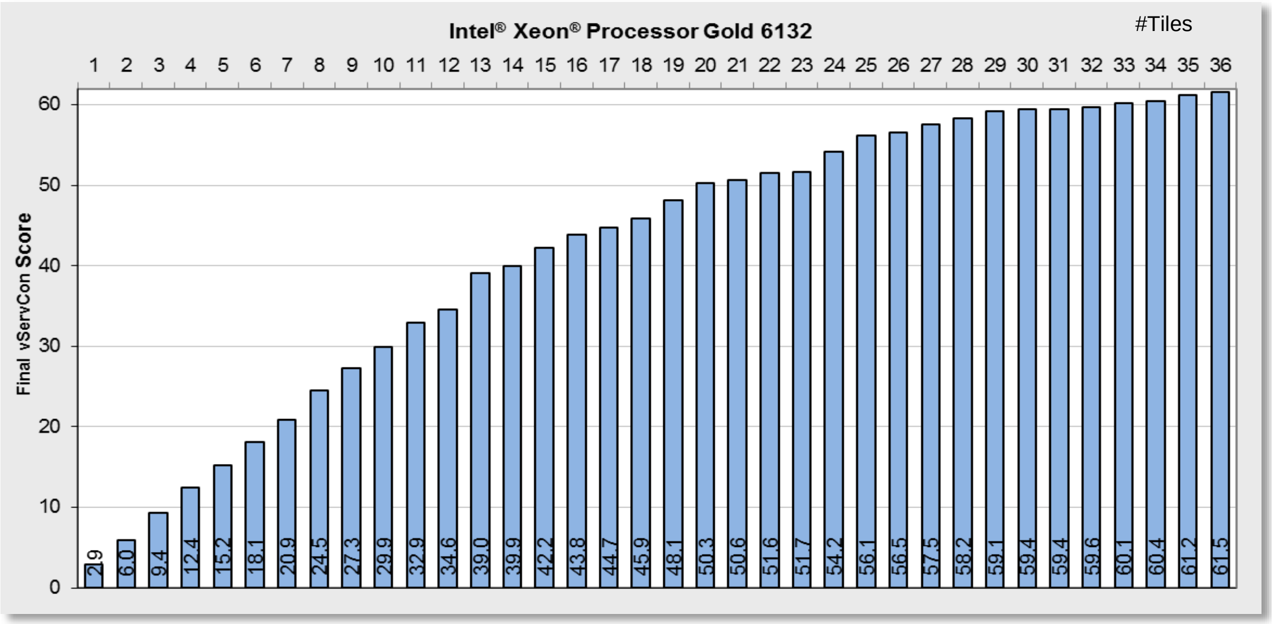

The previous diagram examined the total performance of all application VMs of a host. However, studying the performance from an individual application VM viewpoint is also interesting. This information is in the previous diagram. For example, the total optimum is reached in the above Xeon Gold 6132 situation with 108 application VMs (36 tiles, not including the idle VMs) The low load case is represented by three application VMs (one tile, not including the idle VM). Remember, the vServCon score for one tile is an average value across the three application scenarios in vServCon. This average performance of one tile drops when changing from the low load case to the total optimum of the vServCon score – from 2.9 to 61.5/36=1.7, i.e. to 58%. The individual types of application VMs can react very differently in the high load situation. It is thus clear that in a specific situation the performance requirements of an individual application must be balanced against the overall requirements regarding the numbers of VMs on a virtualization host.

The virtualization-relevant progress in processor technology since 2009 has an effect on the one hand on an individual VM and, on the other hand, on the possible maximum number of VMs up to CPU full utilization. The following comparison shows the proportions for both types of improvements. Seven systems are compared with the best processors each (see table opposite) for few VMs and for highest maximum performance.

	System	Best Performance Few VMs	vServCon Score 1 Tile	Best Maximum Performance	vServCon Score max.
2009	RX600 S4	X7460	1.69	X7460	6.66@ 6 tiles
2010	RX600 S5	X7542	2.07	X7560	17.8@18 tiles
2011	RX600 S6	E7-8837	2.24	E7-4870	23.9@24 tiles
2014	RX4770 M1	E7-4890 v2	2.76	E7-4890 v2	44.9@28 tiles
2015	RX4770 M2	E7-8893 v3	3.04	E7-8890 v3	58.2@34 tiles
2016	RX4770 M3	E7-8893 v4	3.18	E7-8890 v4	75.6@44 tiles
2017	RX4770 M4	Gold 6144	3.58	Platinum 8180	113.6@68tiles

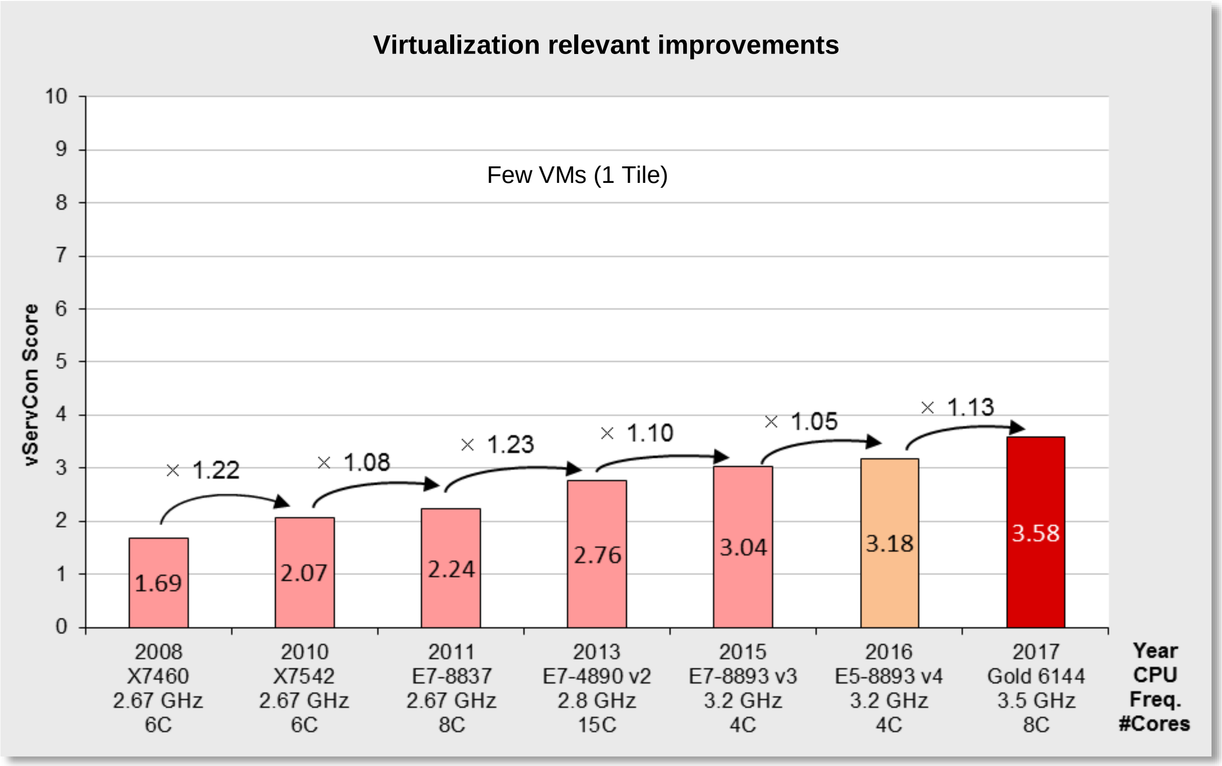

The performance for an individual VM in low-load situations has only slightly increased for the processors compared here with the highest clock frequency per core. We must explicitly point out that the increased virtualization performance as seen in the score cannot be completely deemed as an improvement for one individual VM.

Performance increases in the virtualization environment since 2010 are mainly achieved by increases in the maximum number of VMs that can be operated.

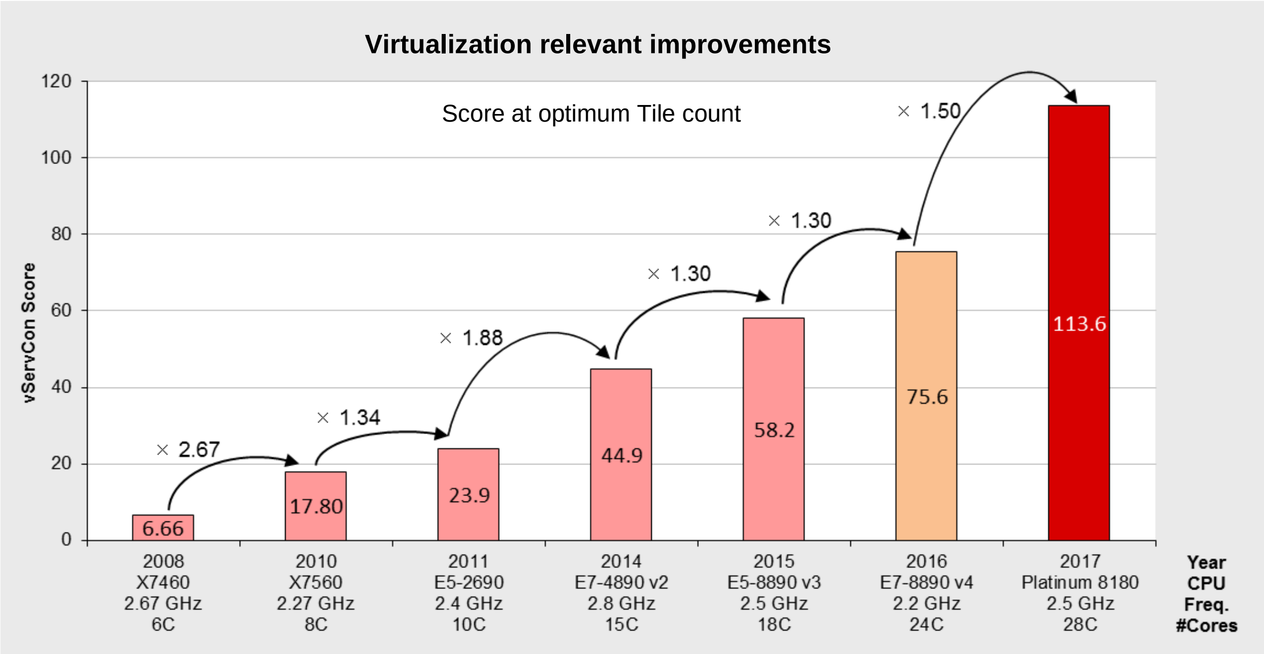

VMmark V3

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 and 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.

A new feature of VMmark V3 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 reflects the maximum total consolidation benefit of all VMs for a server configuration with hypervisor and is used as a comparison criterion of various hardware platforms.

This score is determined from the individual results of the VMs and an infrastructure 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 workload components.

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, which 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 measurement set-up is symbolically illustrated below:

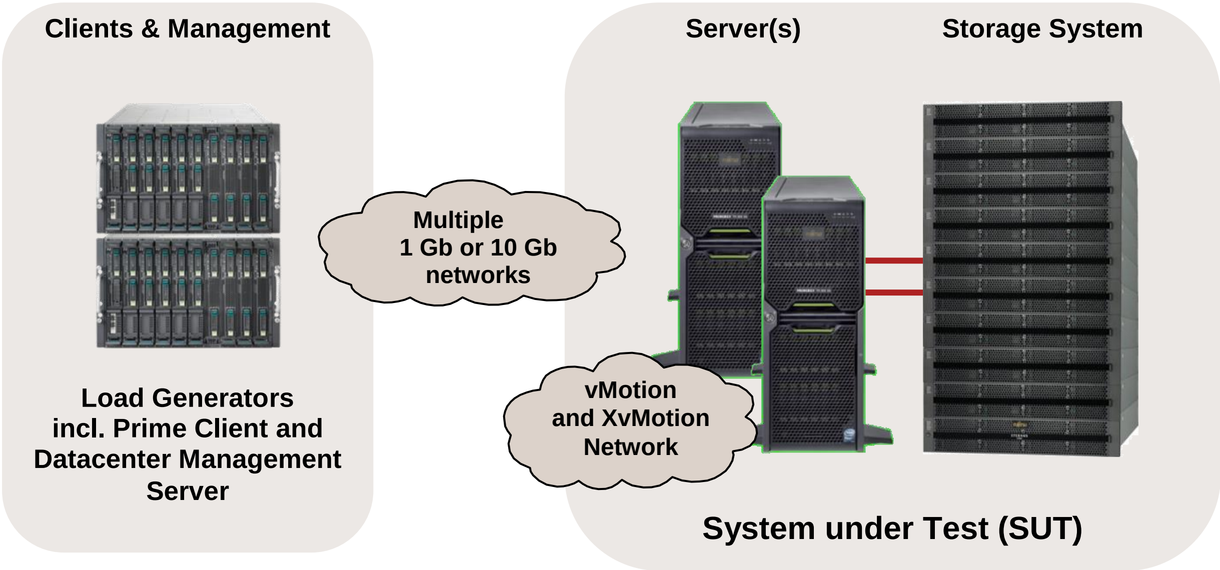

System Under Test (SUT)	
Hardware	
Number of servers	2
Model	PRIMERGY RX4770 M4
Processor	4 × Xeon Platinum 8180
Memory	1536 GB: 48 × 32 GB (1x32 GB) 2Rx4 DDR4-2666 R ECC
Network interface	3 × Emulex OneConnect OCe14000 Dual Port 10GbE Adapter
Disk subsystem	2 × Dual port PFC EP LPe31002 3 × PRIMERGY RX2540 M2 configured as Fibre Channel target: 1 × SAS-SSD (400 GB) 1 × Fusion-io ioMemory PX600(1.3 TB) 3 × Fusion-io ioMemory PX600(2.6 TB) RAID 0 with several LUNs Total: 28.5 TB 1 × PRIMERGY TX2560 M2 configured as Fibre Channel target: 1 × SAS-HDD (500 GB) 2 × Fusion-io ioMemory PX600(2.6 TB) RAID 0 with several LUNs Total: 5.7 TB 1 × PRIMEQUEST 2800E3 configured as Fibre Channel target: 1 × SAS-HDD (600 GB) 3 × Intel P3700 (800 GB) RAID 0 with several LUNs Total: 3.0 TB
Software	
BIOS	R1.7.0
BIOS settings	See details
Operating system	VMware ESXi 6.5.0 U1g Build 7967591
Operating system settings	ESX settings: see details

Details	
See disclosure	http://www.vmware.com/a/assets/vmmark/pdf/2018-07-18-Fujitsu-RX4770M4.pdf http://www.vmware.com/a/assets/vmmark/pdf/2018-07-18-Fujitsu-RX4770M4-serverPPKW.pdf

Datacenter Management Server (DMS)	
Hardware	
Model	1 × PRIMERGY RX2530 M2
Processor	2 × Intel Xeon E5-2698
Memory	64 GB
Network interface	1 × Emulex One Connect Oce14000 1GbE Dual Port Adapter
Software	
Operating system	VMware ESXi 6.0.0 U2 Build 3620759
Datacenter Management Server (DMS) VM	
Hardware	
Processor	8 × logical CPU
Memory	32 GB
Network interface	1 × 1 Gbit/s LAN
Software	
Operating system	Microsoft Windows Server 2008 R2 Standard x64 Edition

Load generator	
Hardware	
Model	5 × PRIMERGY RX2530 M2
Processor	Client Host 1-4: 2 × Xeon E5-2699 v4 Client Host 5 : 2 × Xeon E5-2699A v4
Memory	256 GB
Network interface	1 × Emulex One Connect Oce14000 1GbE Dual Port Adapter 1 × Emulex One Connect Oce14000 10GbE Dual Port Adapter
Software	
Operating system	VMware ESXi 6.5.0 U1 Build 5969303

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

Benchmark results

"Performance Only" measurement result (July 18 2018)

On July 18, 2018 Fujitsu achieved with a PRIMERGY RX4770 M4 with Xeon Platinum 8180 processors and VMware ESXi 6.5.0 U1g a VMmark V3 score of "14.86@15 tiles" in a system configuration with a total of 4 × 56 processor cores and when using two identical servers in the "System under Test" (SUT). With this result the PRIMERGY RX4770 M4 is in the official VMmark V3 "Performance Only" ranking the most powerful four-socket server in a "matched pair" configuration consisting of two identical hosts (valid as of benchmark results publication date).

All comparisons for the competitor products reflect the status of July 18 2018. The current VMmark V3 "Performance Only" results as well as the detailed results and configuration data are available at https://www.vmware.com/products/vmmark/results3x.html .

The diagram shows the "Performance Only" result of the PRIMERGY RX4770 M4 in comparison with the best four-socket systems in a "matched pair" configuration.

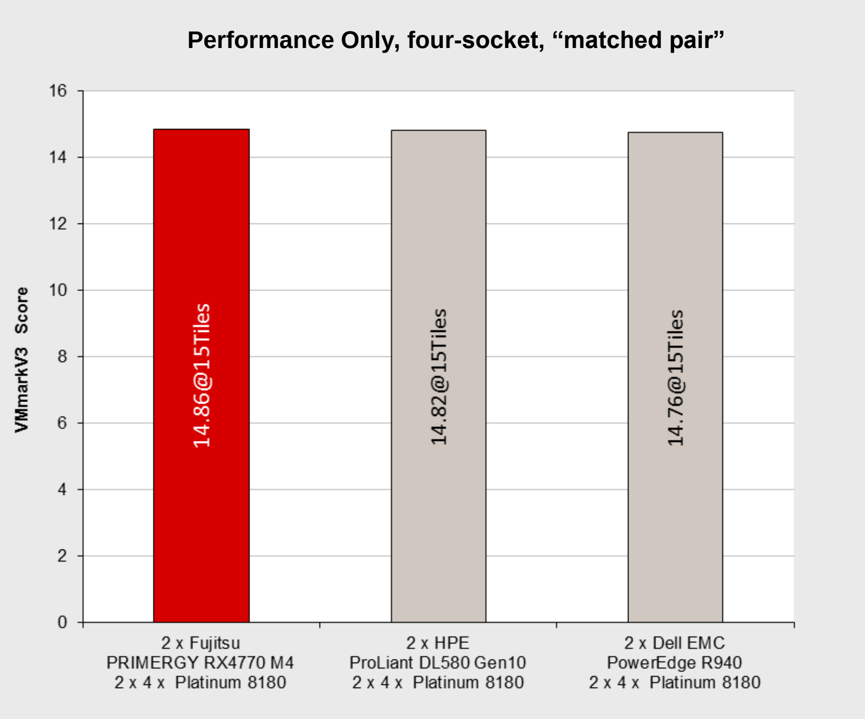

The processors used, which with a good hypervisor setting could make optimal use of their processor features, were the essential prerequisites for achieving the PRIMERGY RX4770 M4 result. These features include Hyper-Threading. All this has a particularly positive effect during virtualization.

four-socket systems, "matched pair"	VMmark V3 Score	Difference
Fujitsu PRIMERGY RX4770 M4	14.86@15 tiles	
HPE ProLiant DL580 Gen10	14.82@15 tiles	0.27%
Dell EMC PowerEdge R940	14.76@15 tiles	0.68%

All VMs, their application data, the host operating system as well as additionally required data were on a powerful Fibre Channel disk subsystem. As far as possible, the configuration of the disk subsystem takes the specific requirements of the benchmark into account. The use of flash technology in the form of SAS SSDs and PCIe-SSDs in the powerful Fibre Channel disk subsystem resulted in further advantages in response times of the storage medium used.

The network connection to the load generators and the infrastructure-workload connection between the hosts were implemented via 10GbE LAN ports.

All the components used were optimally attuned to each other.

"Performance with Server Power" measurement result (July 18 2018)

On July 18, 2018 Fujitsu achieved with a PRIMERGY RX4770 M4 with Xeon Platinum 8180 processors and VMware ESXi 6.5.0 U1g a VMmark V3 "Server PPKW Score" of "5.9120@15 tiles" in a system configuration with a total of 4 × 56 processor cores and when using two identical servers in the "System under Test" (SUT). With this result the PRIMERGY RX4770 M4 is in the official VMmark V3 "Performance with Server Power" ranking the most powerful four-socket server in a "matched pair" configuration consisting of two identical hosts (valid as of benchmark results publication date).

The current VMmark V3 "Performance with Server Power" results as well as the detailed results and configuration data are available at https://www.vmware.com/products/vmmark/results3x.html .

VMmark® is a product of VMware, Inc.

STREAM

Benchmark description

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

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

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

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

During implementation the defined data area, consisting of 8 byte elements, it is successively copied to four types, and arithmetic calculations are also performed to some extent.

Type	Execution	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.

This chapter specifies throughputs on a basis of 10 (1 GB/s = 10^9 Byte/s).

Benchmark environment

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX4770 M4
Processor	4 × Intel® Xeon® Processor Scalable Family
Memory	48 × 16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC
Software	
BIOS settings	Xeon Platinum 81xx, Gold 61xx: DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Power Technology = Custom HWPM Support = Disabled UPI Link Frequency Select = 10.4 GT/s Sub NUMA Clustering = Enabled Stale Atos = Enabled LLC dead line alloc = Disabled
Operating system	SUSE Linux Enterprise Server 12 SP2 (x86_64)
Operating system settings	Transparent Huge Pages inactivated sched_cfs_bandwidth_slice_us = 50000 sched_latency_ns = 240000000 sched_migration_cost_ns = 5000000 sched_min_granularity_ns = 100000000 sched_wakeup_granularity_ns = 150000000 aio-max-nr = 1048576 cpupower -c all frequency-set -g performance cpupower idle-set -d 1 cpupower idle-set -d 2 cpupower idle-set -d 3 echo 0 > /proc/sys/kernel/numa_balancing echo 1 > /proc/sys/vm/drop_caches ulimit -s unlimited
Compiler	C/C++: Version 17.0.0 20160721 of Intel C++ Compiler for Linux
Benchmark	Stream.c Version 5.10

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

Benchmark results

This results in italic are estimated values.

Processor	Memory Frequency [MHz]	Max. Memory Bandwidth [GB/s]	Cores	Processor Frequency [GHz]	Number of Processors	TRIAD [GB/s]
Xeon Gold 5115	2400	115.2	10	2.4	4	*280*
Xeon Gold 5118	2400	115.2	12	2.3	4	*308*
Xeon Gold 5120	2400	115.2	14	2.2	4	*331*
Xeon Gold 6130	2666	128.0	16	2.1	4	*379*
Xeon Gold 6140	2666	128.0	18	2.3	4	*386*
Xeon Gold 6138	2666	128.0	20	2.0	4	*397*
Xeon Gold 6148	2666	128.0	20	2.4	4	*397*
Xeon Gold 6152	2666	128.0	22	2.1	4	397
Xeon Platinum 8153	2666	128.0	16	2.0	4	*381*
Xeon Platinum 8160	2666	128.0	24	2.1	4	*406*
Xeon Platinum 8164	2666	128.0	26	2.0	4	*408*
Xeon Platinum 8170	2666	128.0	26	2.1	4	399
Xeon Platinum 8176	2666	128.0	28	2.1	4	*406*
Xeon Platinum 8180	2666	128.0	28	2.5	4	406
Xeon Gold 5122	2666	128.0	4	3.6	4	*198*
Xeon Gold 6128	2666	128.0	6	3.4	4	294
Xeon Gold 6134	2666	128.0	8	3.2	4	*346*
Xeon Gold 6144	2666	128.0	8	3.5	4	362
Xeon Gold 6126	2666	128.0	12	2.6	4	*362*
Xeon Gold 6136	2666	128.0	12	3.0	4	*390*
Xeon Gold 6146	2666	128.0	12	3.2	4	403
Xeon Gold 6132	2666	128.0	14	2.6	4	373
Xeon Gold 6142	2666	128.0	16	2.6	4	394
Xeon Gold 6150	2666	128.0	18	2.7	4	*388*
Xeon Gold 6154	2666	128.0	18	3.0	4	398
Xeon Platinum 8156	2666	128.0	4	3.6	4	212
Xeon Platinum 8158	2666	128.0	12	3.0	4	*322*
Xeon Platinum 8168	2666	128.0	24	2.7	4	397
Xeon Gold 6134M	2666	128.0	8	3.2	4	*353*
Xeon Gold 6140M	2666	128.0	18	2.3	4	*386*
Xeon Gold 6142M	2666	128.0	16	2.6	4	*394*
Xeon Platinum 8160M	2666	128.0	24	2.1	4	*406*
Xeon Platinum 8170M	2666	128.0	26	2.1	4	*399*
Xeon Platinum 8176M	2666	128.0	28	2.1	4	*406*
Xeon Platinum 8180M	2666	128.0	28	2.5	4	*406*

The following diagram illustrates the throughput of the PRIMERGY RX4770 M4 in comparison to its predecessor, the PRIMERGY RX4770 M3

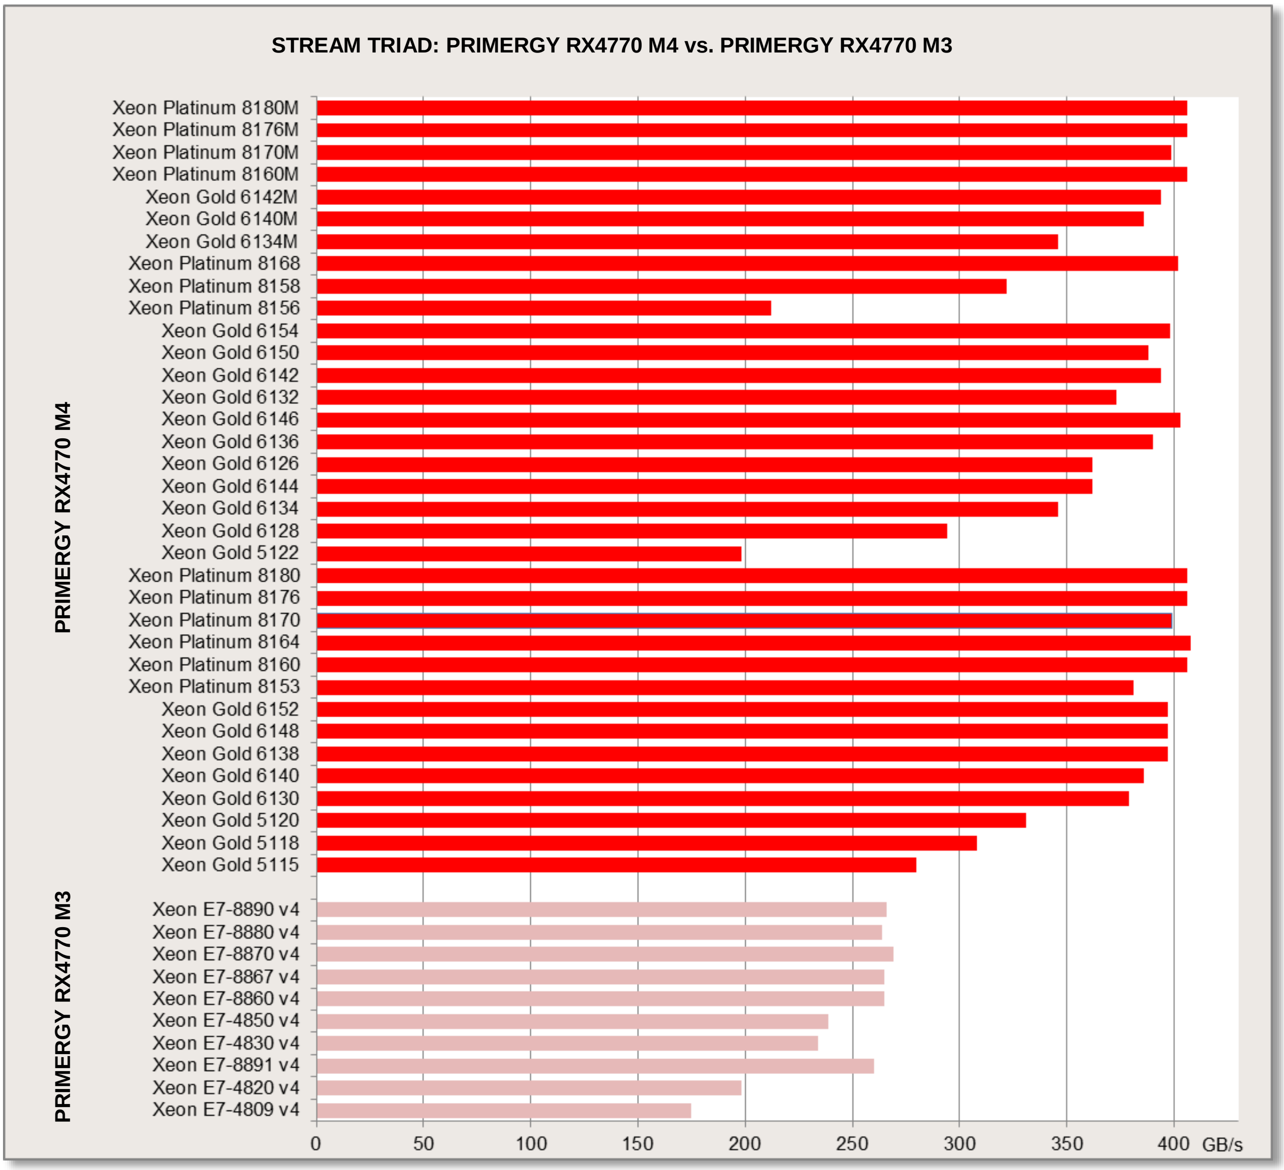

LINPACK

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. A description can be found in the document http://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 × 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 × n matrix the number of arithmetic operations required for the solution is ${}^{2}/{}_{3}n^3 + 2n^2$. Thus, the choice of n determines the duration of the measurement: a doubling of n results in an approximately eight-fold increase in the duration of the measurement. The size of n also has an influence on the measurement result itself. As n increases, the measured value asymptotically approaches a limit. The size of the matrix is therefore usually adapted to the amount of memory available. Furthermore, the memory bandwidth of the system only plays a minor role for the measurement result, but a role that cannot be fully ignored. The processor performance is the decisive factor for the measurement result. Since the algorithm used permits parallel processing, in particular the number of processors used and their processor cores are - in addition to the clock rate - of outstanding significance.

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

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
× Number of processor cores of the computer
× 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 http://www.top500.org/. The use of a LINPACK version based on HPL is prerequisite for this (see http://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 http://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 M4
Processor	4 × Intel® Xeon® Processor Scalable Family
Memory	48 ×16 GB (1x16 GB) 2Rx4 PC4-2666V R ECC
Software	
BIOS settings	Xeon Platinum 81xx, Gold 61xx: Hyper-Threading = Disabled DCU Streamer Prefetcher = Disabled Intel Virtualization Technology = Disabled Package C State limit = C0 HWPM Support = Disabled Link Frequency Select = 10.4 GT/s Stale AtoS = Enabled LLC Dead Line Alloc = Disabled Patrol Scrub = Disabled IMC Interleaving = 2-way Sub NUMA Clustering = Disabled Fan Control = Full
Operating system	SUSE Linux Enterprise Server 12 SP2 (x86_64)
Operating system settings	run with avx512 cpupower -c all frequency-set -g performance sched_cfs_bandwidth_slice_us = 50000 sched_latency_ns = 240000000 sched_migration_cost_ns = 5000000 sched_min_granularity_ns = 100000000 sched_wakeup_granularity_ns = 150000000 aio-max-nr = 1048576
Benchmark	MPI version: Intel® Math Kernel Library Benchmarks for Linux OS (l_mklb_p_2017.3.017)

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

Benchmark results

This results in italic are estimated values.

Processor	Cores	Processor Frequency [GHz]	Number of Processors	Rpeak [GFlops]	Rmax [GFlops]	Efficiency
Xeon Gold 5115	10	2.4	4	1536	*1245*	81%
Xeon Gold 5118	12	2.3	4	1766	*1527*	86%
Xeon Gold 5120	14	2.2	4	1972	*1279*	65%
Xeon Gold 6130	16	2.1	4	4300	*3298*	77%
Xeon Gold 6140	18	2.3	4	5300	*3681*	69%
Xeon Gold 6138	20	2.0	4	5120	*3517*	69%
Xeon Gold 6148	20	2.4	4	6144	*4027*	66%
Xeon Gold 6152	22	2.1	4	5914	3619	61%
Xeon Platinum 8153	16	2.0	4	4096	*2817*	69%
Xeon Platinum 8160	24	2.1	4	6452	*4319*	67%
Xeon Platinum 8164	26	2.0	4	6656	*4508*	68%
Xeon Platinum 8170	26	2.1	4	6988	4277	61%
Xeon Platinum 8176	28	2.1	4	7526	*5064*	67%
Xeon Platinum 8180	28	2.5	4	8960	6077	68%
Xeon Gold 5122	4	3.6	4	1844	*1341*	73%
Xeon Gold 6128	6	3.4	4	2612	1409	54%
Xeon Gold 6134	8	3.2	4	3276	*2314*	71%
Xeon Gold 6144	8	3.5	4	3584	1982	55%
Xeon Gold 6126	12	2.6	4	3994	*2843*	71%
Xeon Gold 6136	12	3.0	4	4608	*3244*	70%
Xeon Gold 6146	12	3.2	4	4916	3050	62%
Xeon Gold 6132	14	2.6	4	4660	3041	65%
Xeon Gold 6142	16	2.6	4	5324	*3517*	66%
Xeon Gold 6150	18	2.7	4	6220	4082	66%
Xeon Gold 6154	18	3.0	4	6912	4440	64%
Xeon Platinum 8156	4	3.6	4	1843	879	48%
Xeon Platinum 8158	12	3.0	4	4608	*3350*	73%
Xeon Platinum 8168	24	2.7	4	8294	4955	60%
Xeon Gold 6134M	8	3.2	4	3276	*2314*	71%
Xeon Gold 6140M	18	2.3	4	5300	*3681*	69%
Xeon Gold 6142M	16	2.6	4	5324	*5320*	72%
Xeon Platinum 8160M	24	2.1	4	6452	*4319*	67%
Xeon Platinum 8170M	26	2.1	4	6988	*4960*	71%
Xeon Platinum 8176M	28	2.1	4	7526	*5064*	67%
Xeon Platinum 8180M	28	2.5	4	8960	*6077*	68%

Rmax = Measurement result

Rpeak = Maximum number of floating point operations per clock cycle
× Number of processor cores of the computer
× Rated frequency [GHz]

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 cases, you should disable the turbo functionality via BIOS option.

Literature

PRIMERGY Servers

http://primergy.com/

PRIMERGY RX4770 M4

This White Paper:
http://docs.ts.fujitsu.com/dl.aspx?id=222a4455-b070-483a-b155-0b87eb086b0c,
http://docs.ts.fujitsu.com/dl.aspx?id=140122ac-01bf-4785-87b0-e05e1110734c

Data sheet
http://docs.ts.fujitsu.com/dl.aspx?id=bbc16505-0e26-4bd8-963b-fd982e99aff2

PRIMERGY Performance

http://www.fujitsu.com/fts/x86-server-benchmarks

OLTP-2

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

SPECcpu2006

http://www.spec.org/osg/cpu2006

Benchmark overview SPECcpu2006
http://docs.ts.fujitsu.com/dl.aspx?id=1a427c16-12bf-41b0-9ca3-4cc360ef14ce

STREAM

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

vServCon

Benchmark Overview vServCon
http://docs.ts.fujitsu.com/dl.aspx?id=b953d1f3-6f98-4b93-95f5-8c8ba3db4e59

LINPACK

The LINPACK Benchmark: Past, Present, and Future
http://www.netlib.org/utk/people/JackDongarra/PAPERS/hplpaper.pdf

TOP500
http://www.top500.org/

HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers
http://www.netlib.org/benchmark/hpl/

Intel Math Kernel Library – LINPACK Download
http://software.intel.com/en-us/articles/intel-math-kernel-library-linpack-download/

SAP

SAP Standard Application Benchmarks
http://www.sap.com/benchmark
https://www.sap.com/dmc/exp/2018-benchmark-directory/#/bwh

Contact

FUJITSU

Website: http://www.fujitsu.com/

PRIMERGY Product Marketing

mailto:Primergy-PM@ts.fujitsu.com

PRIMERGY Performance and Benchmarks

mailto:primergy.benchmark@ts.fujitsu.com

2018/11/30 WW EN