White Paper
FUJITSU Server PRIMERGY
Performance Report PRIMERGY RX4770 M3

This document contains a summary of the benchmarks executed for the FUJITSU Server PRIMERGY RX4770 M3.

The PRIMERGY RX4770 M3 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
1.1
2016-09-08

Contents

Document history

Version 1.0 (2016-07-21)

New:

- Technical data
- SPECcpu2006
  Measurements with Intel® Xeon® Processor E7 v4 Family
- Disk I/O: Performance of RAID controllers
  Measurements with “PRAID CP400i”, “PRAID EP400i” and “PRAID EP420i” controllers
- OLTP-2
  Results for Intel® Xeon® Processor E7 v4 Family
- TPC-E
  Measurement with Xeon E7-8890 v4
- VMmark V2
  “Performance Only“ Measurements with Xeon E7-8890 v4
  „Performance with Server Power“ Measurement with Xeon E7-8890 v4
- STREAM
  Measurements with Intel® Xeon® Processor E7 v4 Family
- LINPACK
  Measurements with Intel® Xeon® Processor E7 v4 Family

Version 1.1 (2016-09-08)

New:

- SAP SD
  Certification number 2016036
- vServCon
  Results for Intel® Xeon® Processor E7 v4 Family

Technical data

PRIMERGY RX4770 M3

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 M3**
Form factor	Rack server
Chipset	Intel® C602 Chipset
Number of sockets	4
Number of processors orderable	2 or 4
Processor type	Intel® Xeon® Processor E7 v4 Family
Number of memory slots	96
Maximum memory configuration	6 TB
Onboard LAN controller	2 × 10 Gbit/s
PCI slots	9 × PCI-Express 3.0 x8 2 × PCI-Express 3.0 x16
Max. number of internal hard disks	12

Processors (since system release)								
Processor	**Cores**	**Threads**	**Cache [MB]**	**QPI Speed [GT/s]**	**Rated Frequency [Ghz]**	**Max. Turbo Frequency [Ghz]**	**Max. Memory Frequency [MHz]**	**TDP [Watt]**
Xeon E7-8893 v4	4	8	60	9.60	3.20	3.50	1866	140
Xeon E7-4809 v4	8	16	20	6.40	2.10	entf.	1866	115
Xeon E7-4820 v4	10	20	25	6.40	2.00	entf.	1866	115
Xeon E7-8891 v4	10	20	60	9.60	2.80	3.50	1866	165
Xeon E7-4830 v4	14	28	35	8.00	2.00	2.80	1866	115
Xeon E7-4850 v4	16	32	40	8.00	2.10	2.80	1866	115
Xeon E7-8860 v4	18	36	45	9.60	2.20	3.20	1866	140
Xeon E7-8867 v4	18	36	45	9.60	2.40	3.30	1866	165
Xeon E7-8870 v4	20	40	50	9.60	2.10	3.00	1866	140
Xeon E7-8880 v4	22	44	55	9.60	2.20	3.30	1866	150
Xeon E7-8890 v4	24	48	60	9.60	2.20	3.40	1866	165

All the processors that can be ordered with the PRIMERGY RX4770 M3, apart from Xeon E7-4809 v4 and Xeon E7-4820 v4, 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 frequency maximum 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 will 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 "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**
16GB (2x8GB) 1Rx4 DDR4-2400 R ECC	16	1	4	2400			✓	✓
32GB (2x16GB) 1Rx4 DDR4-2400 R ECC	32	1	4	2400			✓	✓
64GB (2x32GB) 2Rx4 DDR4-2400 R ECC	64	2	4	2400			✓	✓
128GB (2x64GB) 4Rx4 DDR4-2400 LR ECC	128	4	4	2400		✓	✓	✓

Power supplies (since system release)	**Max. number**
Power Supply Module 1200W w/o power cord	4
Power Supply Module 1600W w/o power cord	4

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

Detailed technical information is available in the data sheet PRIMERGY RX4770 M3.

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 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; the peak values are optional.

Benchmark	Arithmetics	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 favour 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” specify 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 M3
Processor	4 processors of Intel® Xeon® Processor E7 v4 Family
Memory	16 × 32GB (2x16GB) 2Rx4 DDR4-2400 R ECC
Software	
BIOS settings	Energy Performance = Performance Xeon E7-4809 v4, Xeon E7-4820 v4, Xeon E7-4830 v4, Xeon E7-8891 v4, Xeon E7-8893 v4: COD Enable = Disabled Home Dir Snoop with IVT- Style OSB Enable = Enabled Alle others: COD Enable = Enabled Home Dir Snoop with IVT- Style OSB Enable = Disabled
Operating system	SUSE Linux Enterprise Server 12 SP1 (x86_64)
Operating system settings	echo always > /sys/kernel/mm/transparent_hugepage/enabled
Compiler	C/C++: Version 16.0.0.101 of Intel C++ Studio XE for Linux Fortran: Version 16.0.0.101 of Intel Fortran Studio XE 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 on the processor frequency. The number of cores, which are loaded by the benchmark, determines the maximum processor frequency that can be achieved.

Processor	Number of processors	SPECint_rate_base2006	SPECint_rate2006	SPECfp_rate_base2006	SPECfp_rate2006
Xeon E7-8893 v4	4	**944**	**1000**	**827**	**844**
Xeon E7-4809 v4	4	**1100**	**1160**	**982**	**1000**
Xeon E7-4820 v4	4	**1240**	**1290**	**1150**	**1170**
Xeon E7-8891 v4	4	**2140**	**2250**	**1680**	**1710**
Xeon E7-4830 v4	4	**1950**	**2040**	**1550**	**1580**
Xeon E7-4850 v4	4	**2330**	**2440**	**1740**	**1780**
Xeon E7-8860 v4	4	**2870**	**2990**	**2080**	**2130**
Xeon E7-8867 v4	4	**2950**	**3080**	**2130**	**2180**
Xeon E7-8870 v4	4	**3070**	**3200**	**2170**	**2230**
Xeon E7-8880 v4	4	**3320**	**3450**	**2260**	**2320**
Xeon E7-8890 v4	4	**3570**	**3710**	**2360**	**2430**

On 6th June 2016 the PRIMERGY RX4770 M3 with four Xeon E7-8890 v4 processors was ranked first in the 4-socket systems category for the benchmark SPECint_rate_base2006. The current results can be found at http://www.spec.org/cpu2006/results.

The two diagrams below illustrate the throughput of the PRIMERGY RX4770 M3 in comparison to its predecessor PRIMERGY RX4770 M2, in their respective most performant configuration.

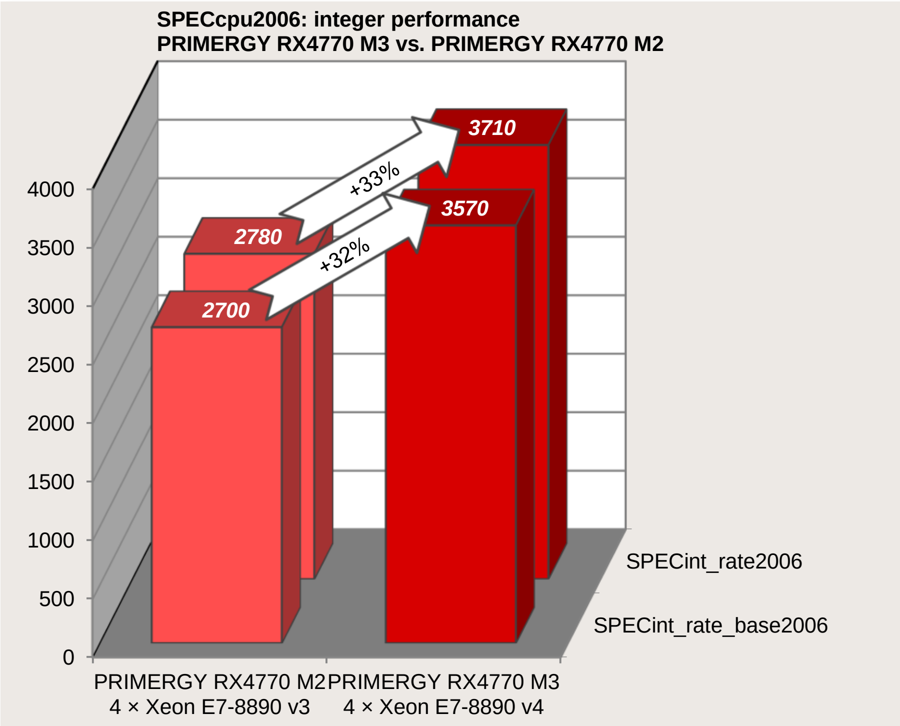

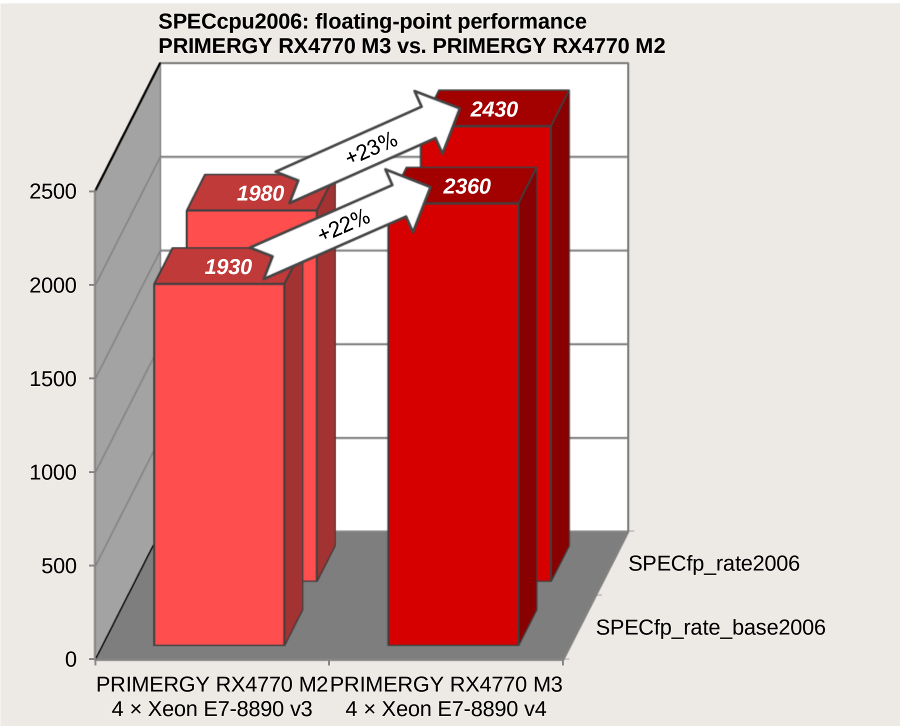

Disk I/O: Performance of RAID controllers

Benchmark description

Performance measurements of disk subsystems for PRIMERGY and PRIMEQUEST servers are used to assess their performance and enable a comparison of the different storage connections for these servers. As standard, these performance measurements are carried out with a defined measurement method, which models the accesses of real application scenarios on the basis of specifications.

The essential specifications are:

- Share of random accesses / sequential accesses
- Share of read / write access types
- Block size (kB)
- Number of parallel accesses (# of outstanding I/Os)

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

Standard load profile	Access	Type of access		Block size [kB]	Application
		read	write		
File copy	random	50%	50%	64	Copying of files
File server	random	67%	33%	64	File server
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 of files

In order to model applications that access in parallel with a different load intensity the "# of Outstanding I/Os" is increased from 1 to 512 (in steps to the power of two).

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

The main results of a measurement are:

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

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

Data throughput [MB/s]	*= Transaction rate [IO/s] × Block size [MB]*
Transaction rate [IO/s]	*= Data throughput [MB/s] / Block size [MB]*

This section specifies capacities of storage media on a basis of 10 (1 TB = 10^{12} bytes) while all other capacities, file sizes, block sizes and throughputs are specified on a basis of 2 (1 MB/s = 2^{20} bytes/s).

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

Benchmark environment

All the measurement results discussed in this chapter were determined using the hardware and software components listed below:

System Under Test (SUT)		
Hardware		
Processor	4 × Xeon E7-8867 v4 @ 2.40GHz	
Controller	1 × “PRAID CP400i”, “PRAID EP400i”, “PRAID EP420i”: Driver name: megasas2.sys, Driver version: 6.706.06 Firmware package: 24.7.0-0061	
Storage media	SSDs	HDDs
	Toshiba PX02SMF040	HGST HUC156045CSS204
Software		
BIOS settings	Intel Virtualization Technology = Disabled VT-d = Disabled Energy Performance = Performance Utilization Profile = Unbalanced CPU C6 Report = Disabled	
Operating system	Microsoft Windows Server 2012 R2 Standard	
Operating system settings	Choose or customize a power plan: High performance For the processes that create disk I/Os: set the AFFINITY to the CPU node to which the PCIe slot of the RAID controller is connected	
Administration software	ServerView RAID Manager 6.2.1	
Benchmark version	3.0	
Stripe size	Controller default	
Measuring tool	Iometer 1.1.0	
Measurement area	The first 10% of the usable LBA area is used for sequential accesses; the next 25% for random accesses.	
File system	raw	
Total number of Iometer workers	1	
Alignment of Iometer accesses	Aligned to whole multiples of 4096 bytes	

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

Benchmark results

The results presented here are designed to help you choose the right solution from the various configuration options of the PRIMERGY RX4770 M3 in the light of disk-I/O performance. Various combinations of RAID controllers and storage media will be analyzed below. Information on the selection of storage media themselves is to be found in the section “Disk I/O: Performance of storage media”.

Hard disks

The hard disks are the first essential component. If there is a reference below to “hard disks”, this is meant as the generic term for HDDs (“hard disk drives”, in other words conventional hard disks) and SSDs (“solid state drives”, i.e. non-volatile electronic storage media).

Mixed drive configurations of SAS and SATA hard disks in one system are permitted, unless they are excluded in the configurator for special hard disk types.

More detailed performance statements about hard disk types are available in the section “Disk I/O: Performance of storage media” in this performance report.

Model versions

The maximum number of hard disks in the system depends on the system configuration. The following table lists the essential cases. Only the highest supported version is named for all the interfaces we have dealt with in this section.

Form factor	Interface	Connection type	Number of PCIe controllers	Maximum number of hard disks
2.5"	SATA 6G, SAS 12G	direct	1	8

RAID controller

In addition to the hard disks the RAID controller is the second performance-determining key component. In the case of these controllers the “modular RAID” concept of the PRIMERGY servers offers a plethora of options to meet the various requirements of a wide range of different application scenarios.

The following table summarizes the most important features of the available RAID controllers of the PRIMERGY RX4770 M3. A short alias is specified here for each controller, which is used in the subsequent list of the performance values.

Controller name	Alias	Cache	Supported interfaces		In the system		FBU
					Max. # disks per controller	RAID levels	
PRAID CP400i	PRAID CP400i	-	SATA 6G SAS 12G	PCIe 3.0 x8	8 × 2.5"	0, 1, 1E, 5, 10, 50	-
PRAID EP400i	PRAID EP400i	1 GB	SATA 6G SAS 12G	PCIe 3.0 x8	8 × 2.5"	0, 1, 1E, 5, 6, 10, 50, 60	✓
PRAID EP420i	PRAID EP420i	2 GB	SATA 6G SAS 12G	PCIe 3.0 x8	8 × 2.5"	0, 1, 1E, 5, 6, 10, 50, 60	✓

System-specific interfaces

The interfaces of a controller in CPU direction (DMI or PCIe) and in the direction of hard disks (SAS or SATA) have in each case specific limits for data throughput. These limits are listed in the following table. The minimum of these two values is a definite limit, which cannot be exceeded. This value is highlighted in bold in the following table.

Controller alias	Effective in the configuration				Connection via expander
	# Disk-side data channels	Limit for throughput of disk interface	# CPU-side data channels	Limit for throughput of CPU-side interface	
PRAID CP400i	8 × SAS 12G	8240 MB/s	8 × PCIe 3.0	**6761 MB/s**	-
PRAID EP400i	8 × SAS 12G	8240 MB/s	8 × PCIe 3.0	**6761 MB/s**	-
PRAID EP420i	8 × SAS 12G	8240 MB/s	8 × PCIe 3.0	**6761 MB/s**	-

More details about the RAID controllers of the PRIMERGY systems are available in the white paper “RAID Controller Performance”.

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.

In the case of controllers with a cache there are several parameters that can be set. The optimal settings can depend on the RAID level, the application scenario and the type of data medium. In the case of RAID levels 5 and 6 in particular (and the more complex RAID level combinations 50 and 60) it is obligatory to enable the controller cache for application scenarios with write share. If the controller cache is enabled, the data temporarily stored in the cache should be safeguarded against loss in case of power failure. Suitable accessories are available for this purpose (e.g. a FBU).

For the purpose of easy and reliable handling of the settings for RAID controllers and hard disks it is advisable to use the software "ServerView RAID Manager" that is supplied for the server. All the cache settings for controllers and hard disks can usually be made en bloc – specifically for the application – by using the pre-defined modi "Performance", "Data Protection" or "Fast Path optimum". The "Performance" mode ensures the best possible performance settings for the majority of the application scenarios with HDDs. In connection with the "FastPath" RAID controller option, the "Fast Path optimum" mode should be selected if maximum transaction rates are to be achieved with SSDs for random accesses with small blocks (≤ 8 kB, e. g. OLTP operation of databases).

More information about the setting options of the controller cache is available in the white paper "RAID Controller Performance".

Performance values

In general, disk-I/O performance of a logical drive depends on the type and number of hard disks, on the RAID level and on the RAID controller. If the limits of the system-specific interfaces are not exceeded, the statements on disk-I/O performance are therefore valid for all PRIMERGY systems. This is why all the performance statements of the document "RAID Controller Performance" also apply for the PRIMERGY RX4770 M3 if the configurations measured there are also supported by this system.

The performance values of the PRIMERGY RX4770 M3 are listed in table form below, specifically for different RAID levels, access types and block sizes. Substantially different configuration versions are dealt with separately. The established measurement variables, as already mentioned in the subsection Benchmark description, are used here. Thus, transaction rate is specified for random accesses and data throughput for sequential accesses. To avoid any confusion among the measurement units the tables have been separated for the two access types.

The table cells contain the maximum achievable values. This has three implications: On the one hand hard disks with optimal performance were used (the components used are described in more detail in the subsection Benchmark environment). Furthermore, cache settings of controllers and hard disks, which are optimal for the respective access scenario and the RAID level, are used as a basis. And ultimately each value is the maximum value for the entire load intensity range (# 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.

2.5" - Random accesses (maximum performance values in IO/s):

PRIMERGY RX4770 M3							
Configuration version			RAID level	HDDs random 8 kB blocks 67% read [IO/s]	HDDs random 64 kB blocks 67% read [IO/s]	SSDs random 8 kB blocks 67% read [IO/s]	SSDs random 64 kB blocks 67% read [IO/s]
RAID Controller	Hard disk type	#Disks					
PRAID CP400i	HUC156045CSS204 SAS HDD PX02SMF040 SAS SSD	2	1	1153	971	65174	9478
		8	10	4294	2229	117403	26336
		8	0	4784	2531	164298	38806
		8	5	2435	1382	28732	17952
PRAID EP400i	HUC156045CSS204 SAS HDD PX02SMF040 SAS SSD	2	1	1590	888	62721	9386
		8	10	4623	3265	183026	28011
		8	0	5291	3784	237818	45574
		8	5	3008	2097	132953	16056
PRAID EP420i	HUC156045CSS204 SAS HDD PX02SMF040 SAS SSD	2	1	1544	994	63189	9993
		8	10	4616	3213	194669	25960
		8	0	5230	3729	248073	44424
		8	5	2970	2039	133538	16243

2.5" - Sequential accesses (maximum performance values in MB/s):

PRIMERGY RX4770 M3							
Configuration version							
RAID Controller	Hard disk type	#Disks	RAID level	HDDs sequential 64 kB blocks 100% read [MB/s]	HDDs sequential 64 kB blocks 100% write [MB/s]	SSDs sequential 64 kB blocks 100% read [MB/s]	SSDs sequential 64 kB blocks 100% write [MB/s]
PRAID CP400i	HUC156045CSS204 SAS HDD PX02SMF040 SAS SSD	2	1	327	230	1910	422
		8	10	1051	892	5873	1640
		8	0	1795	1779	5851	3297
		8	5	1579	1558	5840	1760
PRAID EP400i	HUC156045CSS204 SAS HDD PX02SMF040 SAS SSD	2	1	350	232	1877	406
		8	10	1148	945	5825	1486
		8	0	1874	1892	5817	3081
		8	5	1648	1658	5871	2661
PRAID EP420i	HUC156045CSS204 SAS HDD PX02SMF040 SAS SSD	2	1	375	232	1887	400
		8	10	1174	942	5795	1665
		8	0	1871	1892	5819	3077
		8	5	1648	1650	5869	2569

Conclusion

At full configuration with powerful hard disks the PRIMERGY RX4770 M3 achieves a throughput of up to 5873 MB/s for sequential load profiles and a transaction rate of up to 248073 IO/s for typical, random application scenarios.

To operate SSDs within the maximum performance range the PRAID CP400i is already suited for the simpler RAID levels 0, 1 and 10, and a PRAID controller with cache is to be preferred for RAID 5.

In the event of HDDs the controller cache for random load profiles with a significant write share has performance advantages for all RAID levels.

SAP SD

Benchmark description

The SAP application software consists of modules used to manage all standard business processes. These include modules for ERP (Enterprise Resource Planning), such as Assemble-to-Order (ATO), Financial Accounting (FI), Human Resources (HR), Materials Management (MM), Production Planning (PP) plus Sales and Distribution (SD), as well as modules for SCM (Supply Chain Management), Retail, Banking, Utilities, BI (Business Intelligence), CRM (Customer Relation Management) or PLM (Product Lifecycle Management).

The application software is always based on a database so that a SAP configuration consists of the hardware, the software components operating system, zhe database and the SAP software itself.

SAP AG has developed SAP Standard Application Benchmarks in order to verify the performance, stability and scaling of a SAP application system. The benchmarks, of which SD Benchmark is the most commonly used and most important, analyze the performance of the entire system and thus measure the quality of the integrated individual components.

The benchmark differentiates between a 2-tier and a 3-tier configuration. The 2-tier configuration has the SAP application and database installed on one server. With a 3-tier configuration the individual components of the SAP application can be distributed via several servers and an additional server handles the database.

The entire specification of the benchmark developed by SAP AG, Walldorf, Germany can be found at: http://www.sap.com/benchmark.

Benchmark environment

The measurement set-up is symbolically illustrated below:

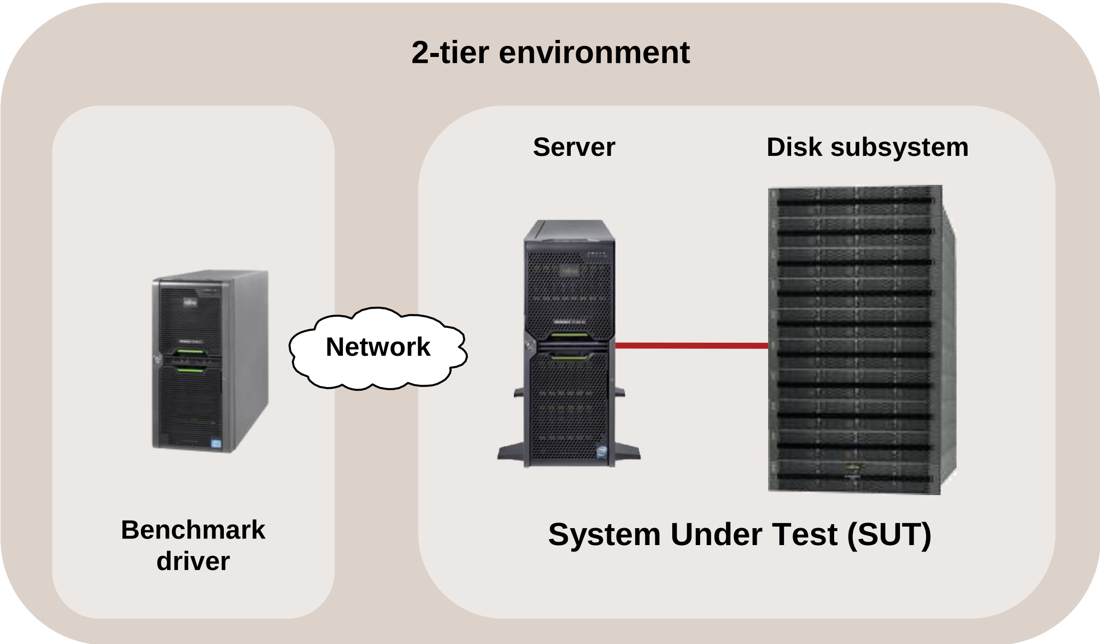

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX4770 M3
Processor	4 × Xeon E7-8890 v4
Memory	32 × 32GB (2x16GB) 2Rx4 DDR4-2400 R ECC
Network interface	1Gbit/s LAN
Disk subsystem	PRIMERGY RX4770 M3: 4 × HD SAS 12G 300GB 15K HOT PL 2.5" EP 1 × PRAID EP420i 1 × PRAID CP400e FH 2 × Eternus JX40
Software	
BIOS settings	COD Enable = Enabled
Operating system	Microsoft Windows Server 2012 R2 Standard Edition
Database	Microsoft SQL Server 2012 (64-bit)
SAP Business Suite Software	SAP enhancement package 5 for SAP ERP 6.0

Benchmark driver	
Hardware	
Model	PRIMERGY RX300 S4
Processor	2 × Xeon X5460
Memory	32 GB
Network interface	1Gbit/s LAN
Software	
Operating system	SUSE Linux Enterprise Server 11 SP1

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

Benchmark results

Certification number 2016036	
Number of SAP SD benchmark users	39,500
Average dialog response time	0.98 seconds
Throughput Fully processed order line items/hour Dialog steps/hour SAPS	 4,315,330 12,946,000 215,770
Average database request time (dialog/update)	0.010 sec / 0.030 sec
CPU utilization of central server	98%
Operating system, central server	Windows Server 2012 R2 Standard Edition
RDBMS	SQL Server 2012
SAP Business Suite software	SAP enhancement package 5 for SAP ERP 6.0
Configuration Central Server	Fujitsu PRIMERGY RX4770 M3 4 processors / 96 cores / 192 threads Intel Xeon Processor E7-8890 v4, 2.20 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 60 MB L3 cache per processor 1024 GB main memory

The following chart shows a comparison of two-tier SAP SD Standard Application Benchmark results for 4-way Xeon E7 v4 based servers with Windows OS and SQL Server database (as of July 19, 2016). The PRIMERGY RX4770 M3 outperforms the comparably configured servers from HP. The latest SAP SD 2-tier results can be found at http://www.sap.com/solutions/benchmark/sd2tier.epx.

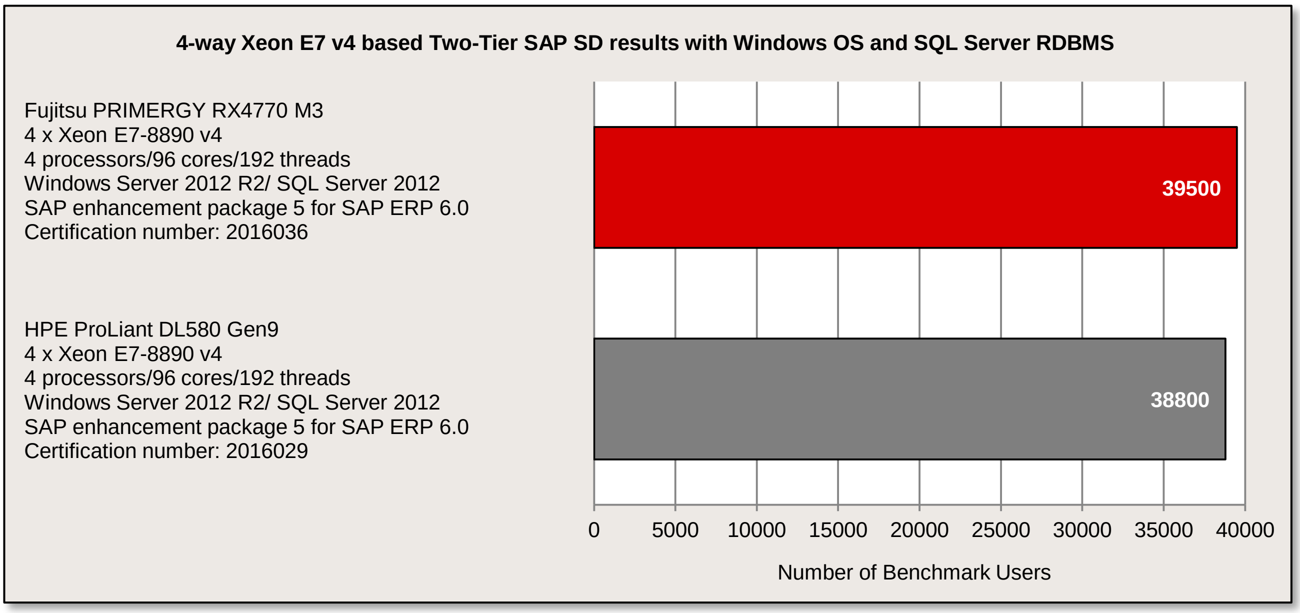

The following diagram illustrates the throughput of the PRIMERGY RX4770 M3 in comparison to its predecessor, the PRIMERGY RX4770 M2, in the respective most performant configuration.

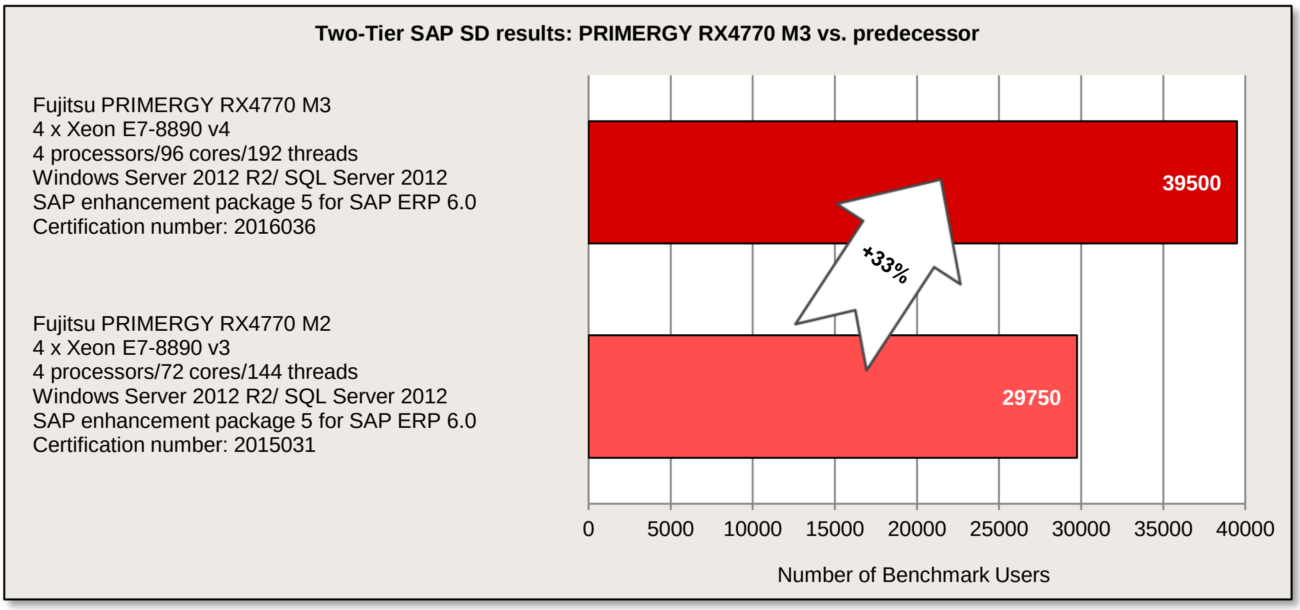

OLTP-2

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 measurement set-up is symbolically illustrated below:

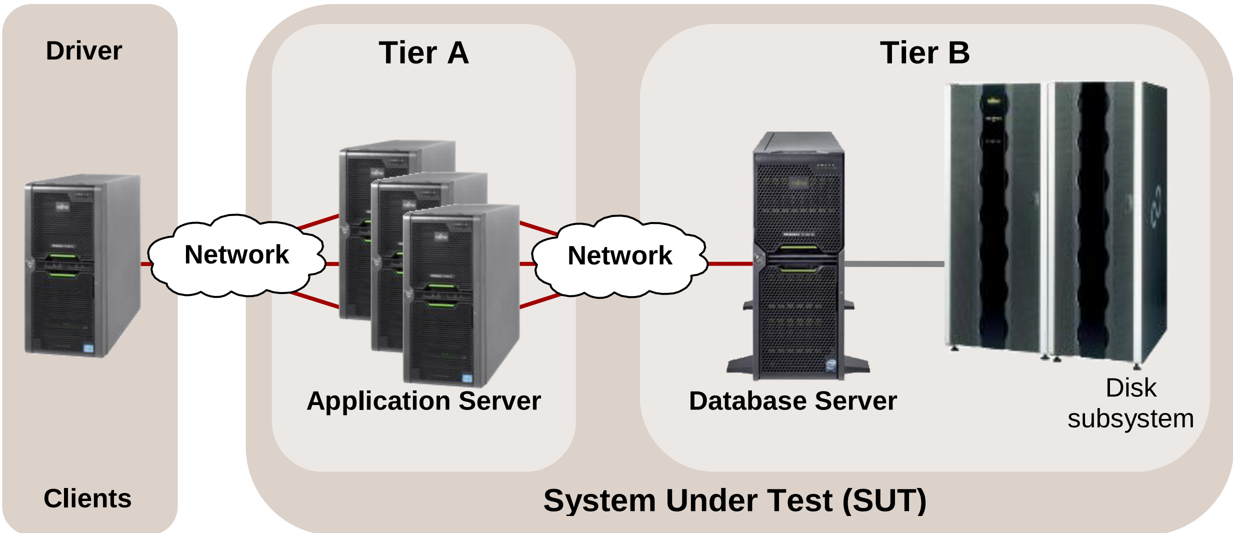

Database Server (Tier B)	
Hardware	
Model	PRIMERGY RX4770 M3
Processor	Intel® Xeon® Processor E7 v4 Family
Memory	2048 GB: 32 × 64GB (2x32GB) 2Rx4 DDR4-2400 R ECC 1024 GB: 16 × 64GB (2x32GB) 2Rx4 DDR4-2400 R ECC
Network interface	2 × onboard LAN 10 Gb/s
Disk subsystem	RX4770 M3: PRAID EP400i 2 × 300 GB 15k rpm SAS Drives, RAID1 (OS), 6 × 600 GB 15k rpm SAS Drives, RAID10 (LOG) 7 × PRAID EP420e 7 × JX40 S2: Je 24 × 400 GB SSD Drive, RAID5 (data)

Software	
BIOS	Version R1.0.0
Operating system	Microsoft Windows Server 2012 R2 Standard
Database	Microsoft SQL Server 2016 Enterprise

Application Server (Tier A)	
Hardware	
Model	1 × PRIMERGY RX2530 M1
Processor	2 × Xeon E5-2697 v3
Memory	64 GB, 2133 MHz registered ECC DDR4
Network interface	2 × onboard LAN 10 Gb/s 1 × Dual Port LAN 1Gb/s
Disk subsystem	2 × 300 GB 15k rpm SAS Drive
Software	
Operating system	Microsoft Windows Server 2012 R2 Standard

Client	
Hardware	
Model	1 × PRIMERGY RX300 S7
Processor	2 × Xeon E5-2667 v2
Memory	64 GB, 1600 MHz registered ECC DDR3
Network interface	2 × onboard LAN 1 Gb/s 1 × Dual Port LAN 1Gb/s
Disk subsystem	1 × 300 GB 15k 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 2048 GB was considered for the measurements with four processors and a configuration with a total memory of 1024 GB for the measurements with two processors. The OLTP-2 scaling measurements presented here were all performed with a memory access speed – depending on the processor type – of at most 1600 MHz. Further information about memory performance can be found in the White Paper Memory performance of Xeon E7 v4 (Broadwell-EX)-based systems.

The following diagram shows the OLTP-2 transaction rates that can be achieved with two and four processors of the Intel® Xeon® Processor E7 v4 Product Family.

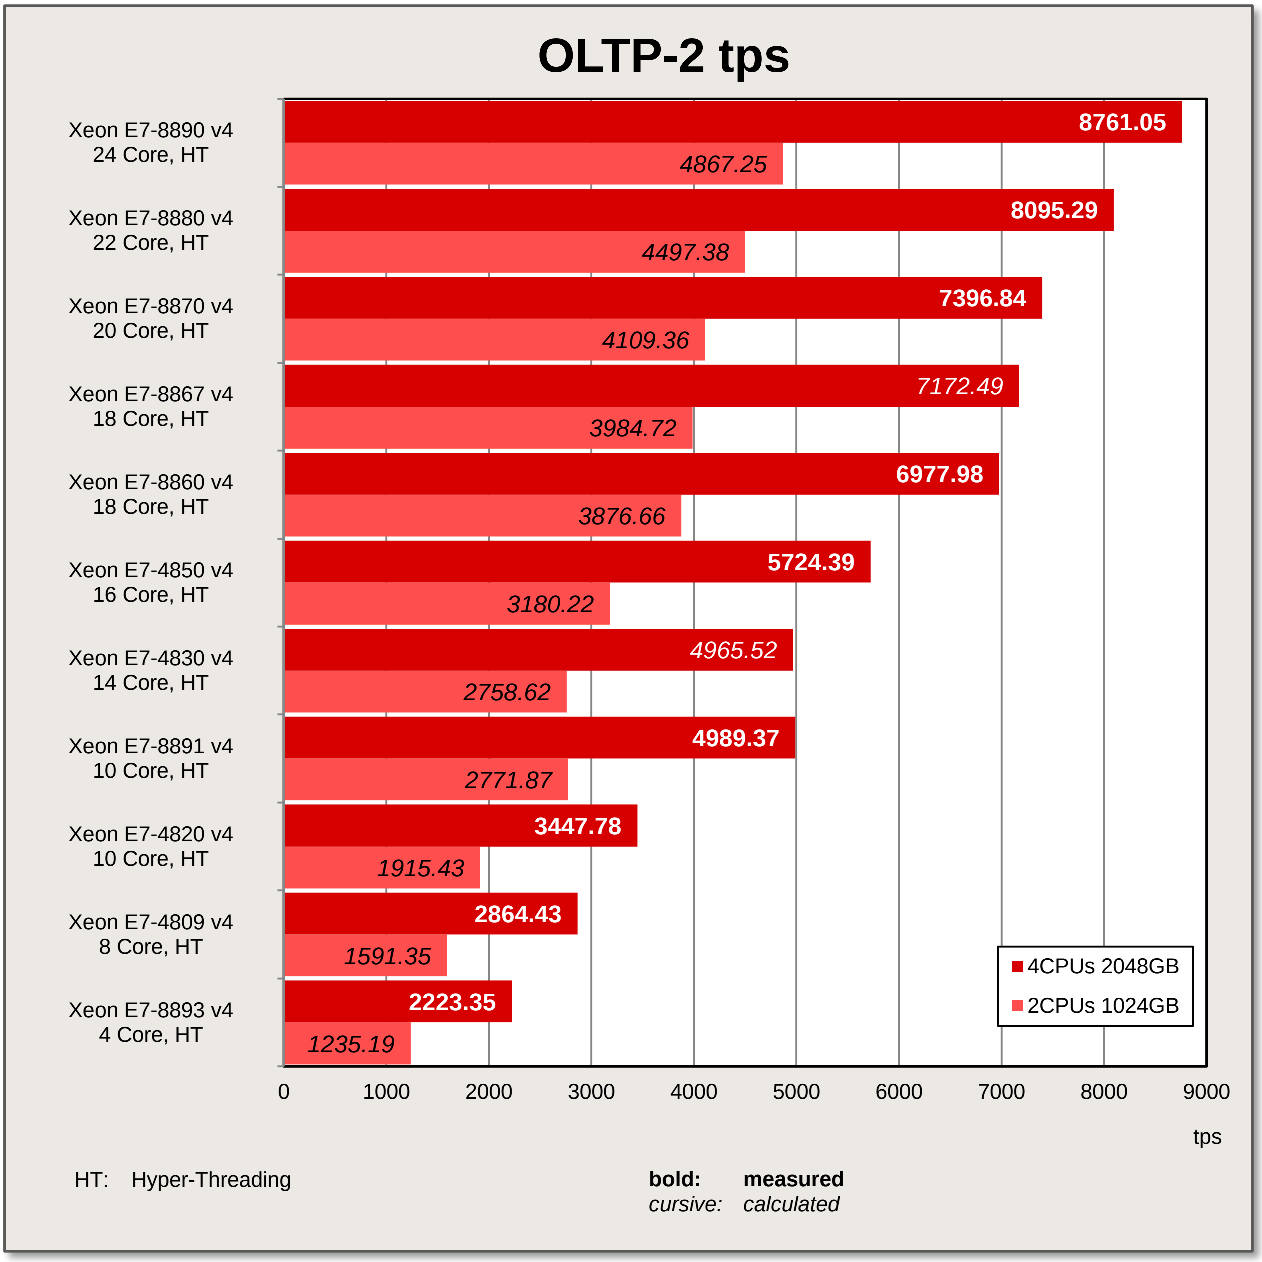

It is evident that a wide performance range is covered by the variety of released processors. If you compare the OLTP-2 value of the processor with the lowest performance (Xeon E7-8893 v4) with the value of the processor with the highest performance (Xeon E7-8890 v4), the result is a 3.9-fold increase in performance.

Based on the results achieved the processors can be divided into different performance groups:

The start is made with Xeon E7-8893 v4 and E7-4809 v4 as processors with four or eight cores. The Xeon E7-4820 v4 processor with its ten cores has the next best performance to offer. Due to its high clock frequency and the high QPI speed of 9.60 GT/s a higher throughput rate is achieved with the performance-optimized 10-core processor Xeon E7-8891 v4.

The groups of 14-, 16-, 18 and 20-core processors offer in this processor series a medium-range OLTP-2 performance. Due to the various technical features of the processors in these groups (see “Technical data”) it is possible to choose the right CPU depending on the usage scenario.

The processors with 22 or 24 cores are to be found at the upper end of the performance scale. An OLTP performance of between 8170.13 tps (4 × Xeon E7-8880 v4) and 8761.05 tps (4 × Xeon E7-8890 v4) is achieved.

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 27%.

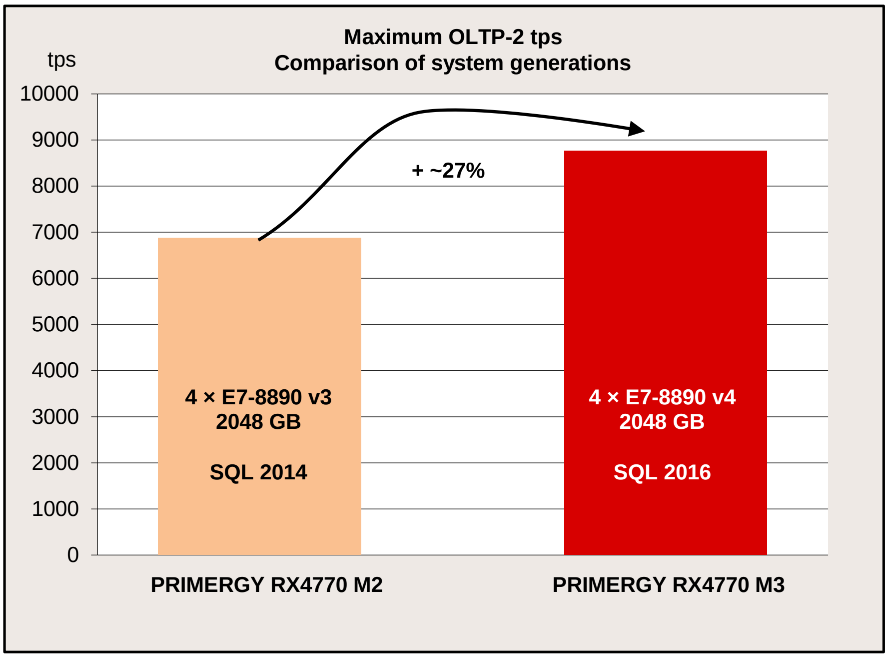

TPC-E

Benchmark description

The TPC-E benchmark measures the performance of online transaction processing systems (OLTP) and is based on a complex database and a number of different transaction types that are carried out on it. TPC-E is not only a hardware-independent but also a software-independent benchmark and can thus be run on every test platform, i.e. proprietary or open. In addition to the results of the measurement, all the details of the systems measured and the measuring method must also be explained in a measurement report (Full Disclosure Report or FDR). Consequently, this ensures that the measurement meets all benchmark requirements and is reproducible. TPC-E does not just measure an individual server, but a rather extensive system configuration. Keys to performance in this respect are the database server, disk I/O and network communication.

The performance metric is tpsE, where tps means transactions per second. tpsE is the average number of Trade-Result-Transactions that are performed within a second. The TPC-E standard defines a result as the tpsE rate, the price per performance value (e.g. $/tpsE) and the availability date of the measured configuration.

Further information about TPC-E can be found in the overview document Benchmark Overview TPC-E.

Benchmark results

In July 2016 Fujitsu submitted a TPC-E benchmark result for the PRIMERGY RX4770 M3 with the 24-core processor Intel Xeon E7-8890 v4 and 2048 GB memory.

The results show an enormous increase in performance compared with the PRIMERGY RX4770 M2 with a simultaneous reduction in costs.

<table>
<tr><td>FUJITSU</td><td colspan="3">FUJITSU Server
PRIMERGY RX4770 M3</td><td colspan="2">TPC-E 1.14.0
TPC Pricing 1.7.0

Report Date
July 12, 2016</td></tr>
<tr><td>TPC-E Throughput
8,796.47 tpsE</td><td colspan="2">Price/Performance
$ 116.62 USD per tpsE</td><td>Availability Date
July 31, 2016</td><td colspan="2">Total System Cost
$ 1,025,815 USD</td></tr>
<tr><td colspan="6">Database Server Configuration</td></tr>
<tr><td>Operating System
Microsoft Windows Server 2012 R2 Standard Edition</td><td colspan="2">Database Manager
Microsoft SQL Server 2016 Enterprise Edition</td><td colspan="2">Processors/Cores/Threads
4/96/192</td><td>Memory
2048 GB</td></tr>
<tr><td colspan="3">SUT
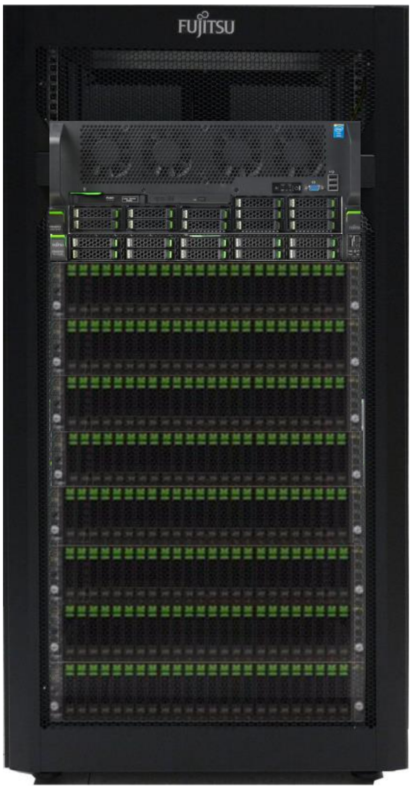</td><td colspan="3">Tier A
PRIMERGY RX2530 M1
2x Intel Xeon E5-2697 v3 2.60 GHz
64 GB Memory
2x 300 GB 15k rpm SAS Drive
2x onboard LAN 10 Gb/s
1x Dual Port LAN 1 Gb/s
1x SAS RAID controller

Tier B
PRIMERGY RX4770 M3
4x Intel Xeon E7-8890 v4 2.20 GHz
2048 GB Memory
2x 300 GB 15k rpm SAS Drives
6x 600 GB 15k rpm SAS Drives
2x onboard LAN 10 Gb/s
8x SAS RAID Controller

Storage
1x PRIMECENTER Rack
7x ETERNUS JX40
168x 400 GB SSD Drives</td></tr>
<tr><td>Initial Database Size
36,951 GB</td><td colspan="3">Redundancy Level 1
RAID-5 data and RAID-10 log</td><td colspan="2">Storage
168 x 400 GB SSD
6 x 600 GB 15k rpm HDD</td></tr>
</table>

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

More details about this TPC-E result, in particular the Full Disclosure Report, can be found via the TPC web page http://www.tpc.org/tpce/results/tpce_result_detail.asp?id=116071201.

In July 2016, Fujitsu is represented with seven results in the TPC-E list (without historical results).

System and Processors	Throughput	Price / Performance	Availability Date
PRIMERGY RX300 S8 with 2 × Xeon E5-2697 v2	2472.58 tpsE	$135.14 per tpsE	September 10, 2013
PRIMEQUEST 2800E with 2 × Xeon E7-8890 v2	8582.52 tpsE	$205.43 per tpsE	Mai 1, 2014
PRIMERGY RX2540 M1 with 2 × Xeon E5-2699 v3	3772.08 tpsE	$130.44 per tpsE	December 1, 2014
PRIMERGY RX4770 M2 with 4 × Xeon E7-8890 v3	6904.53 tpsE	$126.49 per tpsE	June 1, 2015
PRIMEQUEST 2800E2 with 8 × Xeon E7-8890 v3	10058.28 tpsE	$187.53 per tpsE	November 11, 2015
PRIMERGY RX2540 M2 with 2 × Xeon E5-2699 v4	4734.87 tpsE	$111.65 per tpsE	July 31, 2016
PRIMERGY RX4770 M3 with 4 × Xeon E7-8890 v4	8796.47 tpsE	$116.62 per tpsE	July 31, 2016

See the TPC web site for more information and all the TPC-E results (including historical results) (http://www.tpc.org/tpce).

The following overview, sorted according to price/performance, shows the best TPC-E price per performance ratios (as of July 12th, 2016, without historical results) and the corresponding TPC-E throughputs. PRIMERGY RX4770 M3 with a price per performance ratio of $116.62/tpsE has the best price/performance value of all 4-socket systems.

System		Processor type processors/ cores/threads	tpsE (higher is better)	$/tpsE (lower is better)	availability date
Fujitsu	PRIMERGY RX2540 M2	2 × Intel Xeon E5-2699 v4	4,734.87	111.65	2016-07-31
Fujitsu	**PRIMERGY RX4770 M3**	4 × Intel Xeon E7-8890 v4	**8,796.47**	**116.62**	**2016-07-31**
Lenovo	System x3650 M5	2 × Intel Xeon E5-2699 v4	4,938.14	117.91	2016-07-31
Fujitsu	PRIMERGY RX4770 M2	4 × Intel Xeon E7-8890 v3	6,904.53	126.94	2015-01-06
Fujitsu	PRIMERGY RX2540 M1	2 × Intel Xeon E5-2699 v3	3,772.08	130.44	2014-12-01
Fujitsu	PRIMERGY RX300 S8	2 × Intel Xeon E5-2697 v2	2,472.58	135.14	2013-09-10
Lenovo	System x3850 X6	4 × Intel Xeon E7-8890 v4	9,068.00	139.85	2016-07-31
Lenovo	System x3950 X6	8 × Intel Xeon E7-8890 v3	11,058.99	143.91	2015-12-17
IBM	System x3650 M4	2 × Intel Xeon E5-2697 v2	2,591.00	150.00	2013-09-10
Fujitsu	PRIMEQUEST 2800E2	8 × Intel Xeon E7-8890 v3	10,058.28	187.53	2015-11-11

See the TPC web site for more information and all the TPC-E results (including historical results) (http://www.tpc.org/tpce).

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). Add to these a fourth machine, the so-called idle VM. 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.

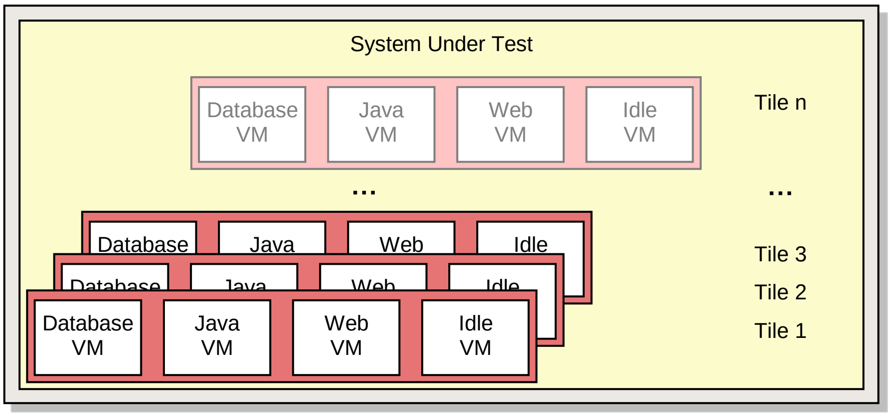

Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score, the individual benchmark results for one tile are put in relation to the respective results of a reference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

Starting as a rule with one tile, 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".

Moreover, vServCon also documents the total CPU load of the host (VMs and all other CPU activities) and, if possible, electrical power consumption.

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

Benchmark environment

The measurement set-up is symbolically illustrated below:

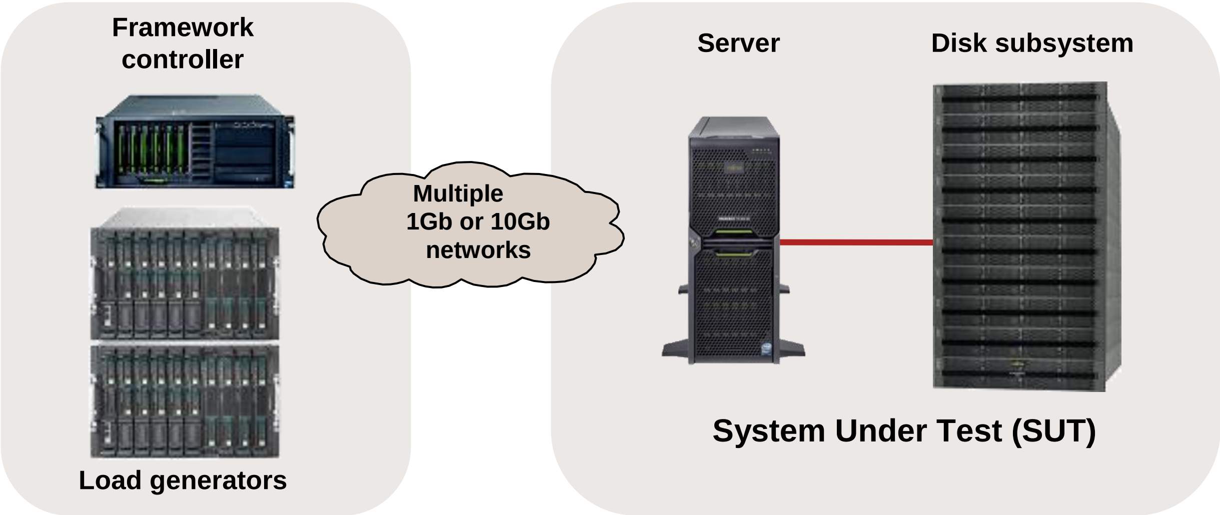

System Under Test (SUT)	
Hardware	
Model	PRIMERGY RX4770 M3
Processor	Intel® Xeon® Processor E7 v4 Family
Memory	1 TB: 32 × 32GB (2x16GB) 2Rx4 DDR4-2400 R ECC
Network interface	1 × dual port 1GbE adapter 1 × dual port 10GbE server adapter
Disk subsystem	1 × dual-channel FC-Controller Emulex LPe16002 LINUX/LIO based flash storage system
Software	
Operating system	VMware ESXi 6.0.0 U2 Build 3620759

Load generator (incl. Framework controller)	
Hardware (Shared)	
Enclosure	PRIMERGY BX900
Hardware	
Model	18 × PRIMERGY BX920 S1 server blades
Processor	2 × Xeon X5570
Memory	12 GB
Network interface	3 × 1 Gbit/s LAN
Software	
Operating system	Microsoft Windows Server 2003 R2 Enterprise with Hyper-V

Load generator VM (per tile 3 load generator VMs on various server blades)	
Hardware	
Processor	1 × logical CPU
Memory	512 MB
Network interface	2 × 1 Gbit/s LAN
Software	
Operating system	Microsoft Windows Server 2003 R2 Enterprise Edition

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

Benchmark results

The quad-socket system PRIMERGY RX4770 M3 dealt with here are based on processors of the Intel® Xeon® Processor E7 v4 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 E7 v4 Family	4 Cores Hyper-Threading, Turbo-Mode	E7-8893 v4	17.3	9
	8 Cores Hyper-Threading	E7-4809 v4	21.5	15
	10 Cores Hyper-Threading	E7-4820 v4	26.1	19
	10 Cores Hyper-Threading, Turbo-Mode	E7-8891 v4	40.9	22
	14 Cores Hyper-Threading, Turbo-Mode	E7-4830 v4	39.3	26
	16 Cores Hyper-Threading, Turbo-Mode	E7-4850 v4	47.4	30
	18 Cores Hyper-Threading, Turbo-Mode	E7-8860 v4	58.5	35
		E7-8867 v4	59.9	35
	20 Cores Hyper-Threading, Turbo-Mode	E7-8870 v4	63.5	39
	22 Cores Hyper-Threading, Turbo-Mode	E7-8880 v4	69.5	42
	24 Cores Hyper-Threading, Turbo-Mode	E7-8890 v4	75.6	44

These PRIMERGY quad-socket systems are very suitable for application virtualization thanks to the progress made in processor technology. Compared with a system based on the previous processor generation an approximate 30% higher virtualization performance can be achieved (measured in vServCon score in their maximum configuration).

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 ("QPI Speed") also determines performance. 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.

More information about the topic "Memory Performance" can be found in the White Paper Memory performance of Xeon E7 v4 (Broadwell-EX)-based systems.

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

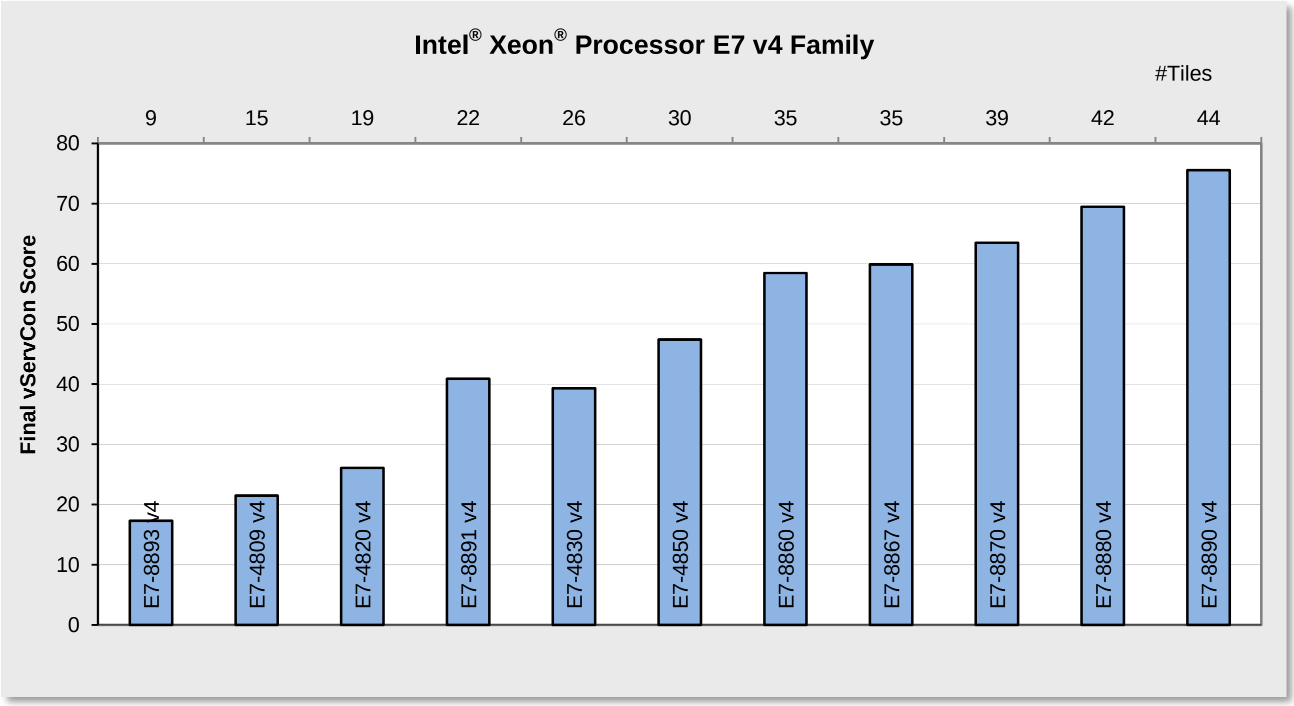

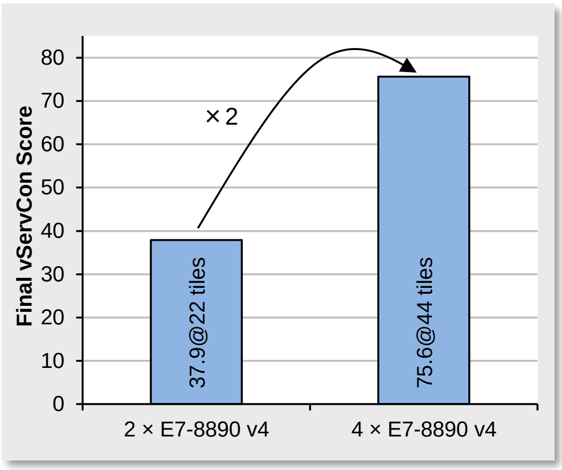

Until now we have looked at the virtualization performance of a fully configured system. However, with a server with four sockets the question also arises as to how good performance scaling is from two to four processors. The better the scaling, the lower the overhead usually caused by the shared use 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 2. When operated with four processors, the system thus achieves twice the performance as with two processors, as is illustrated in the diagram opposite using the processor version Xeon E7-8890 v4 as an example.

The next diagram illustrates the virtualization performance for increasing numbers of VMs based on the Xeon E7-4830 v3 (14-Core) processor.

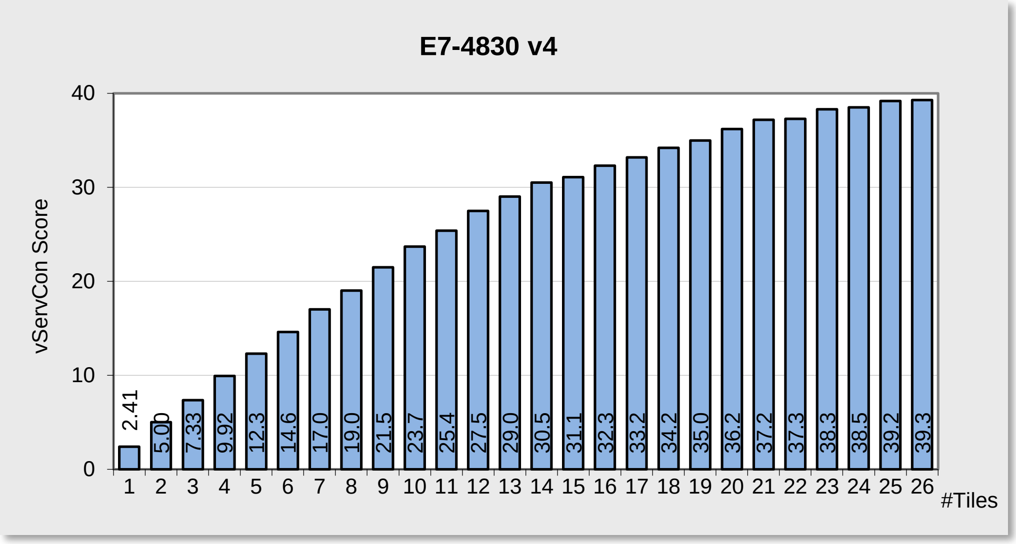

In addition to the increased number of physical cores, Hyper-Threading, which is supported by all Xeon E7 processors, is an additional reason for the high number of VMs that can be operated. As is known, a physical processor core is consequently divided into two logical cores so that the number of cores available for the hypervisor is doubled. This standard feature thus generally increases the virtualization performance of a system.

The scaling curves for the number of tiles as seen in the previous diagram are specifically for systems with Hyper-Threading. 56 physical and thus 112 logical cores are available with the Xeon E7-4830 v4 processors; approximately four of them are used per tile (see Benchmark description). This means that a parallel use of the same physical cores by several VMs is avoided up to a maximum of about fourteen tiles. That is why the performance curve in this range scales almost ideal. For the quantities above the growth is flatter up to CPU full utilization.

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 E7-4830 v4 situation with 78 application VMs (26 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 2.41 to 39.3/26=1.50, i.e. to 59%. 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. Five systems are compared: a system from 2009, a system from 2010, a system from 2011, a system from 2014, a system from 2015 and a current system 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

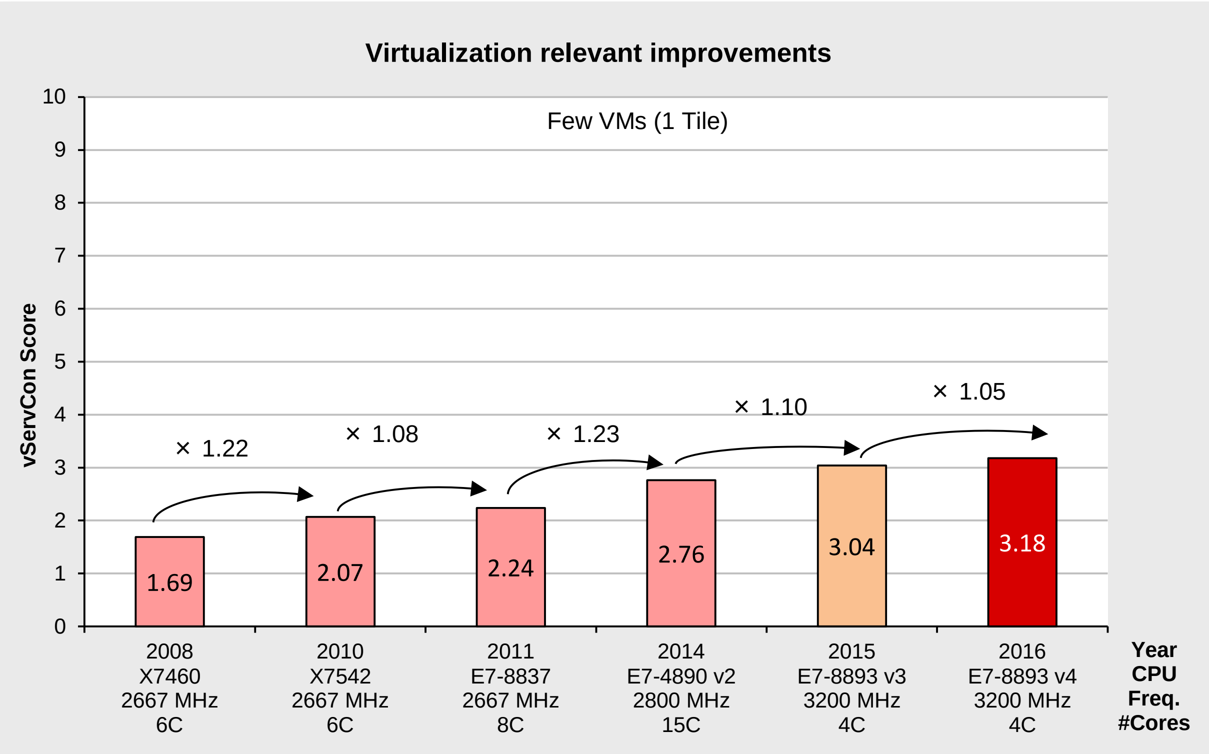

The performance for an individual VM in low-load situations has only slightly increased for the processors compared here with the highest clock frequency per core. We must explicitly point out that the increased virtualization performance as seen in the score cannot be completely deemed as an improvement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

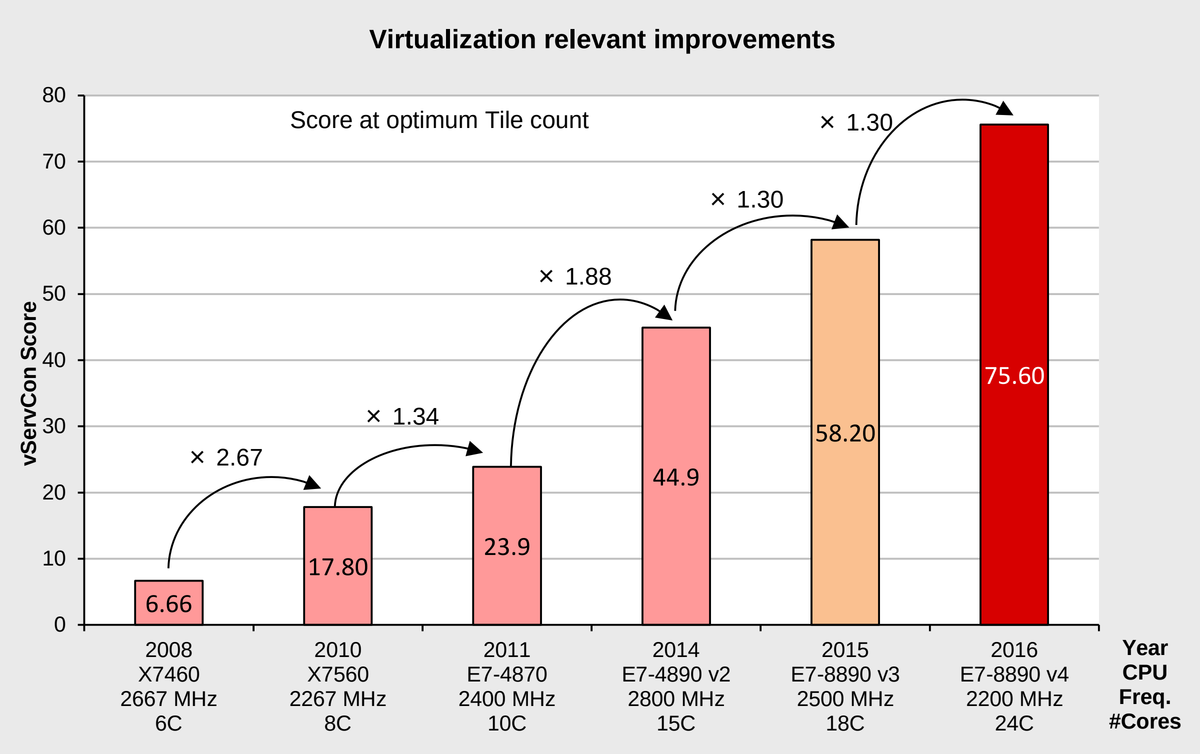

VMmark V2

Benchmark description

VMmark V2 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 V1" in October 2010, it has been succeeded by "VMmark V2", which 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 and also Storage vMotion.

In addition to the "Performance Only" result, it is also possible from version 2.5 of VMmark to 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 V2 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. Three proven benchmarks, which cover the application scenarios mail server, Web 2.0, and e-commerce were integrated in VMmark V2.

Application scenario	Load tool	# VMs
Mail server	LoadGen	1
Web 2.0	Olio client	2
E-commerce	DVD Store 2 client	4
Standby server	(IdleVMTest)	1

Each of the three application scenarios is assigned to a total of seven dedicated virtual machines. Then add to these an eighth VM called the "standby server". These eight 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 V2 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 and Storage vMotion. The Load Balancing capacity of the data center is also used (DRS, Distributed Resource Scheduler).

The result of VMmark V2 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 V2 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 results for one tile are 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 V2 tiles is always specified with each VMmark V2 score. The result is thus as follows: "Score@Number of Tiles", for example "4.20@5 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" is determined, which is the performance score 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 V2 is available in the document Benchmark Overview VMmark V2.

Benchmark environment

The measurement set-up is symbolically illustrated below:

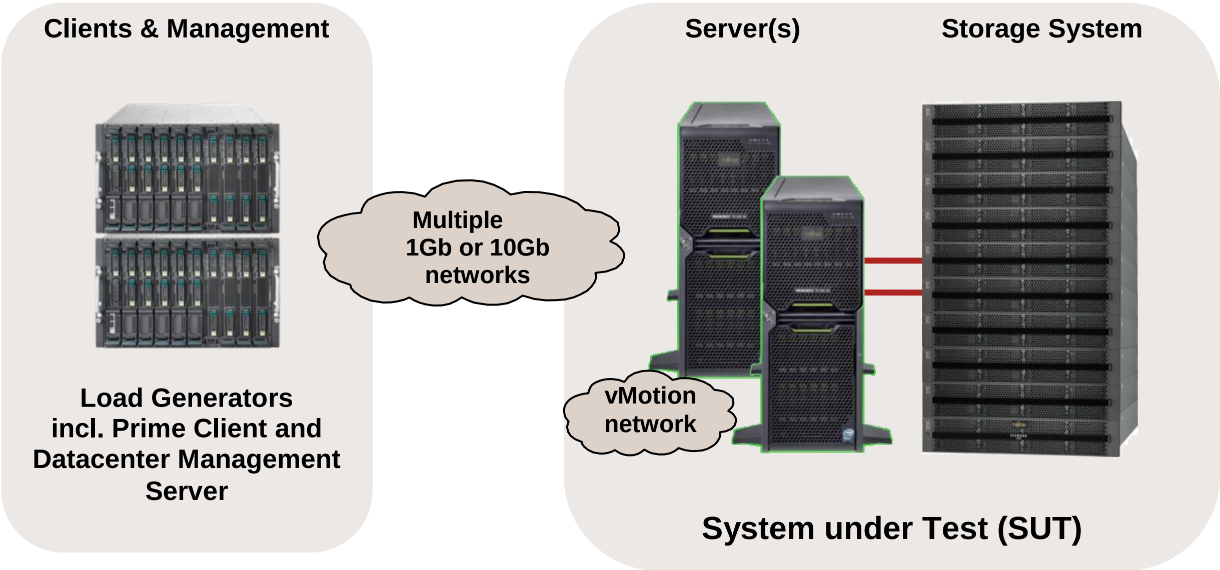

System Under Test (SUT)	
Hardware	
Number of servers	2
Model	PRIMERGY RX4770 M3
Processor	4 × Xeon E7-8890 v4
Memory	1024 GB: 16 × 64GB (2x32GB) 2Rx4 DDR4-2400 R ECC
Network interface	Intel X540-AT2 Dual Port 10GbE Onboard Adapter 1 × Emulex OneConnect OCe14000 Dual 10GbE Port Adapter
Disk subsystem	Dual port PFC EP LPe16002 3 × PRIMERGY RX300 S8 configured as Fibre Channel target Details see disclosures
Software	
BIOS	Version V5.0.0.11 R1.0.0
BIOS settings	See details
Operating system	VMware ESXi 6.0.0 U2 Build 3620759
Operating system settings	ESXi settings: see details

Datacenter Management Server (DMS)	
Hardware (Shared)	
Enclosure	PRIMERGY BX600
Network Switch	1 × PRIMERGY BX600 GbE Switch Blade 30/12
Hardware	
Model	1 × server blade PRIMERGY BX620 S5
Processor	2 × Xeon X5570
Memory	24 GB
Network interface	6 × 1 Gbit/s LAN
Software	
Operating system	VMware ESXi 5.1.0 Build 799733

Datacenter Management Server (DMS)VM	
Hardware	
Processor	4 × logical CPU
Memory	10 GB
Network interface	2 × 1 Gbit/s LAN
Software	
Operating system	Microsoft Windows Server 2008 R2 Enterprise x64 Edition

Prime Client	
Hardware (Shared)	
Enclosure	PRIMERGY BX600
Network Switch	1 × PRIMERGY BX600 GbE Switch Blade 30/12
Hardware	
Model	1 × server blade PRIMERGY BX620 S5
Processor	2 × Xeon X5570
Memory	12 GB
Network interface	6 × 1 Gbit/s LAN
Software	
Operating system	Microsoft Windows Server 2008 Enterprise x64 Edition SP2

Load generator	
Hardware	
Model	3 × PRIMERGY RX600 S6
Processor	4 × Xeon E7-4870
Memory	512 GB
Network interface	6 × 1 Gbit/s LAN
Software	
Operating system	VMware ESX 4.1.0 U2 Build 502767
Load generator VM (per tile 1 load generator VM)	
Hardware	
Processor	4 × logical CPU
Memory	4 GB
Network interface	1 × 1 Gbit/s LAN
Software	
Operating system	Microsoft Windows Server 2008 Enterprise x64 Edition SP2

Details	
See disclosure	http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2016-06-21-Fujitsu-RX4770M3.pdf http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vmmark/2016-06-21-Fujitsu-RX4770M3-serverPPKW.pdf

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

Benchmark results

“Performance Only” measurement result

On June 21, 2016 Fujitsu achieved with two PRIMERGY RX4770 M3 systems with Xeon E7-8890 v4 processors and VMware ESXi 6.0.0 U2 a VMmark V2 score of “61.32@52 tiles” in a system configuration with a total of 2 × 96 processor cores and when using two identical servers/partitions in the “System under Test” (SUT). With this result the PRIMERGY RX4770 M3 is in the official VMmark V2 “Performance Only” ranking the most powerful 4-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 21st June 2016. The current VMmark V2 results as well as the detailed results and configuration data are available at http://www.vmware.com/products/vmmark/results.html.

The diagram shows the result of the PRIMERGY RX4770 M3 in comparison with the best 4-socket systems in a “matched pair” configuration.

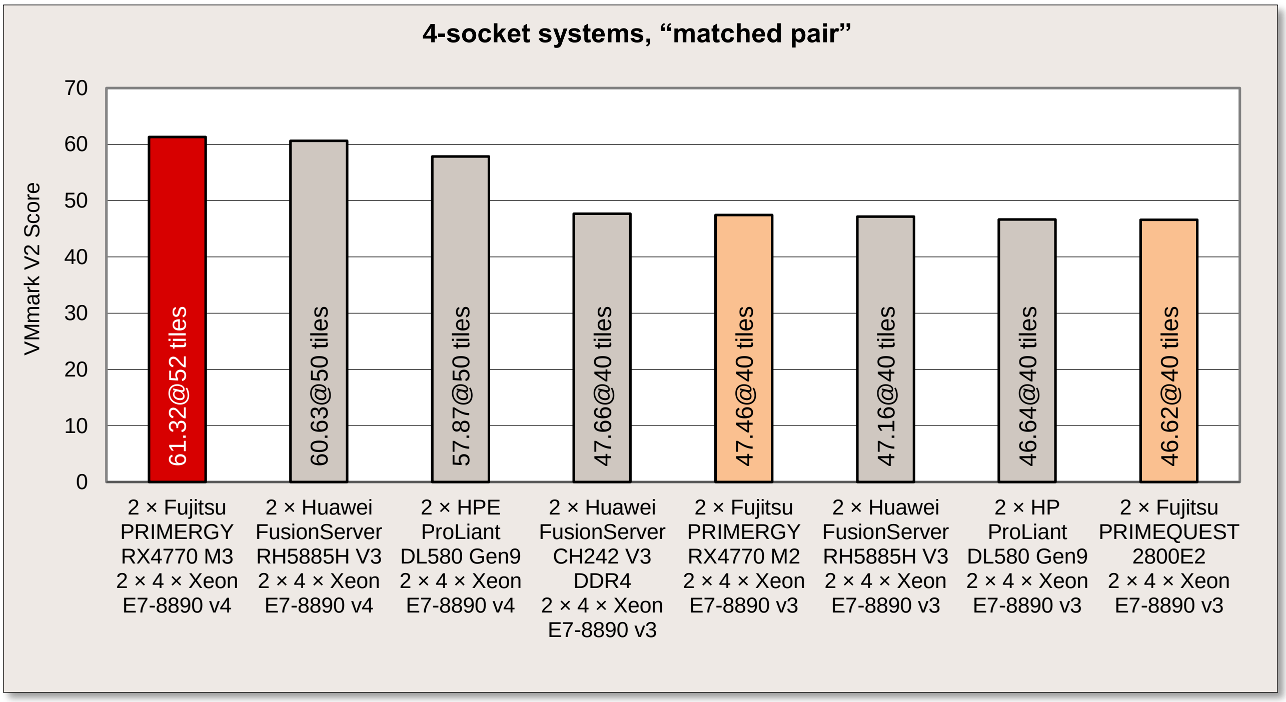

The table opposite shows the difference in the score (in %) between the Fujitsu system and the other 4-socket systems (“matched pair”).

4-socket systems, “matched pair”	VMmark V2 score	Difference
Fujitsu PRIMERGY RX4770 M3	61.32@52 tiles	
Huawei FusionServer RH5885H V3	60.63@50 tiles	1.14%
HPE ProLiant DL580 Gen9	57.87@50 tiles	5.96%
Huawei FusionServer CH242 V3 DDR4	47.66@40 tiles	28.66%
Fujitsu PRIMERGY RX4770 M2	47.46@40 tiles	29.20%
Huawei FusionServer RH5885H V3	47.16@40 tiles	30.03%
HP ProLiant DL580 Gen9	46.64@40 tiles	31.48%
Fujitsu PRIMEQUEST 2800E2	46.62@40 tiles	31.53%

The processors used, which with a good hypervisor setting could make optimal use of their processor features, were the essential prerequisites for achieving all PRIMERGY RX4770 M3 result. These features include Hyper-Threading. All this has a particularly positive effect during virtualization.

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 was implemented via 10Gb LAN ports. The infrastructure-workload connection between the hosts was by means of 1Gb LAN ports.

All the components used were optimally attuned to each other.

„Performance with Server Power“ measurement result

On June 21, 2016 Fujitsu achieved with two PRIMERGY RX4770 M3 systems with Xeon E7-8890 v4 processors and VMware ESXi 6.0.0 U2 a VMmark V2 “Server PPKW Score” of “30.4559@52 tiles” in a system configuration with a total of 2 × 96 processor cores and when using two identical servers/partitions in the “System under Test” (SUT). With this result the PRIMERGY RX4770 M3 is in the official VMmark V2 “Performance with Server Power” ranking the most energy-efficient 4-socket server worldwide (valid as of benchmark results publication date).

All comparisons for the competitor products reflect the status of 21st June 2016. The current VMmark V2 “Performance with Server Power” results as well as the detailed results and configuration data are available at http://www.vmware.com/products/vmmark/results.1.html.

The diagram shows all VMmark V2 “Performance with Server Power“ results.

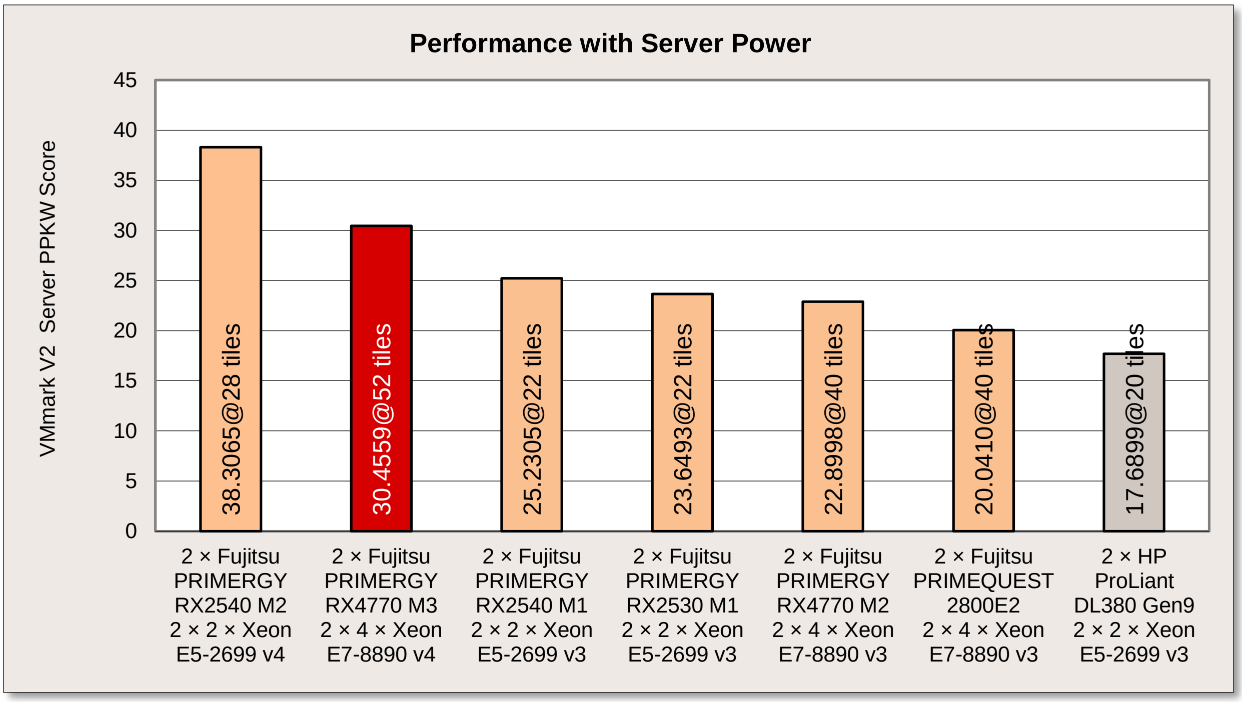

STREAM

Benchmark description

STREAM is a synthetic benchmark that has been used for many years to determine memory throughput and which 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, 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 M3
Processor	4 processors of Intel® Xeon® Processor E7 v4 Family
Memory	16 × 32GB (2x16GB) 2Rx4 DDR4-2400 R ECC
Software	
BIOS settings	Energy Performance = Performance COD Enable = Disabled Home Dir Snoop with IVT- Style OSB Enable = Enabled HyperThreading = Disabled
Operating system	SUSE Linux Enterprise Server 12 SP1 (x86_64)
Operating system settings	Transparent Huge Pages inactivated
Compiler	Intel C++ Composer XE 2016 for Linux
Benchmark	STREAM Version 5.10

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

Benchmark results

Processor	Memory Frequency [MHz]	Max. Memory Bandwidth [GB/s]	Cores	Processor Frequency [GHz]	Number of Processors	TRIAD [GB/s]
Xeon E7-8893 v4	1600	102	4	3.20	4	
Xeon E7-4809 v4	1600	102	8	2.10	4	175
Xeon E7-4820 v4	1600	102	10	2.00	4	198
Xeon E7-8891 v4	1600	102	10	2.80	4	260
Xeon E7-4830 v4	1600	102	14	2.00	4	234
Xeon E7-4850 v4	1600	102	16	2.10	4	239
Xeon E7-8860 v4	1600	102	18	2.20	4	265
Xeon E7-8867 v4	1600	102	18	2.40	4	265
Xeon E7-8870 v4	1600	102	20	2.10	4	269
Xeon E7-8880 v4	1600	102	22	2.20	4	264
Xeon E7-8890 v4	1600	102	24	2.20	4	266

Further information about memory performance can be found in the White Paper Memory performance of Xeon E7 v4 (Broadwell-EX) based systems.

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 M3
Processor	4 processors of Intel® Xeon® Processor E7 v4 Family
Memory	16 × 32GB (2x16GB) 2Rx4 DDR4-2400 R ECC
Software	
BIOS settings	Energy Performance = Performance COD Enable = Disabled Home Dir Snoop with IVT- Style OSB Enable = Enabled HyperThreading = Disabled All processors apart from Xeon E7-4809 v4: Turbo Mode = Enabled (default) = Disabled
Operating system	SUSE Linux Enterprise Server 12 SP1 (x86_64)
Benchmark	HPL version: Intel Optimized MP LINPACK Benchmark for Clusters 11.3

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

Benchmark results

Processor	Cores	Rated Frequency [Ghz]	Number of processors	Rpeak [GFlops]	Rmax (without Turbo Mode) [GFlops]	Efficiency	Rmax (with Turbo Mode) [GFlops]	Efficiency
Xeon E7-8893 v4	4	3.20	4	819				
Xeon E7-4809 v4	8	2.10	4	1075	**1020**	**95%**		
Xeon E7-4820 v4	10	2.00	4	1280				
Xeon E7-8891 v4	10	2.80	4	1792	**1704**	**95%**	1938	108%
Xeon E7-4830 v4	14	2.00	4	1792	**1701**	**95%**	1863	104%
Xeon E7-4850 v4	16	2.10	4	2150	**2027**	**94%**	2225	103%
Xeon E7-8860 v4	18	2.20	4	2534	**2384**	**94%**	2504	99%
Xeon E7-8867 v4	18	2.40	4	2764	**2600**	**94%**	2722	98%
Xeon E7-8870 v4	20	2.10	4	2688	**2533**	**94%**	2677	100%
Xeon E7-8880 v4	22	2.20	4	3097	**2915**	**94%**	3004	97%
Xeon E7-8890 v4	24	2.20	4	3379	**3176**	**94%**	3338	99%

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 does not as a matter of principle 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 M3

This White Paper:
- http://docs.ts.fujitsu.com/dl.aspx?id=8b5a6911-0329-477b-a46b-8891da2b627a
- http://docs.ts.fujitsu.com/dl.aspx?id=1370ebd1-e19a-43c7-91e2-a9c142b3b0b7
- http://docs.ts.fujitsu.com/dl.aspx?id=7984ec7d-3891-4a43-8769-71820068ec18

Data sheet
http://docs.ts.fujitsu.com/dl.aspx?id=0769aca2-f7b3-4afe-8ad1-3e6a9a4d0ce6

PRIMERGY Performance

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

Performance of Server Components

http://www.fujitsu.com/fts/products/computing/servers/mission-critical/benchmarks/x86-components.html

BIOS optimizations for Xeon E5 v4 & E7 v4 based systems
http://docs.ts.fujitsu.com/dl.aspx?id=eb90c352-8d98-4f5a-9eed-b5aade5ccae1

Memory performance of Xeon E7 v4 (Broadwell-EX)-based systems
http://docs.ts.fujitsu.com/dl.aspx?id=7bd26a0c-a46c-4717-be6d-78abebba56b2

RAID Controller Performance
http://docs.ts.fujitsu.com/dl.aspx?id=9845be50-7d4f-4ef7-ac61-bbde399c1014

Disk I/O: Performance of storage media and RAID controllers

Basics of Disk I/O Performance
http://docs.ts.fujitsu.com/dl.aspx?id=65781a00-556f-4a98-90a7-7022feacc602

Information about Iometer
http://www.iometer.org

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/

OLTP-2

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

SAP SD

http://www.sap.com/benchmark

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

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/

TPC-E

http://www.tpc.org/tpce

Benchmark Overview TPC-E
http://docs.ts.fujitsu.com/dl.aspx?id=da0ce7b7-3d80-48cd-9b3a-d12e0b40ed6d

VMmark V2

Benchmark Overview VMmark V2
http://docs.ts.fujitsu.com/dl.aspx?id=2b61a08f-52f4-4067-bbbf-dc0b58bee1bd

VMmark V2
http://www.vmmark.com

vServCon

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

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

2016-09-08 WW EN