FUJITSU

WHITE PAPER
FUJITSU PRIMERGY SERVERS
PERFORMANCE REPORT PRIMERGY RX200 S6

This document contains a summary of the benchmarks executed for the PRIMERGY RX200 S6.

The PRIMERGY RX200 S6 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
2010-09-09

Contents

Document history

Version 1.0

First report version including the benchmark chapters

- SPECcpu2006
  Measurements with Xeon E5503, E5506, E5507, L5609, L5630, E5620, E5630, E5640, X5667, X5677, L5640, X5650, X5660, X5670 and X5680
- SPECjbb2005
  Measurement with Xeon X5680
- SPECpower_ssj2008
  Measurement with Xeon X5670 and 1 x SSD 2.5“ 64GB
- vServCon
  Measurements with Xeon E5507, L5609, L5630, E5620, E5630, E5640, X5667, X5677, L5640, X5650, X5660, X5670, X5680

Version 1.1

New benchmark chapters:

- VMmark
  Measurement with Xeon X5680

Updated benchmark chapters:

- SPECcpu2006
  All results have been published at http://www.spec.org.

Technical data

With only 1 height unit, the PRIMERGY RX200 S6 is a particularly space-saving dual-socket rack server, which replaces the PRIMERGY RX200 S5. It has an Intel 5500 chip set, two Intel Xeon series 5500 or 5600 processors (Dual-Core, Quad-Core or Hexa-Core), 12 DIMM slots for up to 192 GB DDR3-SDRAM, an onboard 2-port 1-GB Ethernet controller and three PCI slots (one PCI-Express 2.0 x4 and two PCI-Express 2.0 x8). The PRIMERGY RX200 S6 is offered in two versions with different drive bays. The drive bays provide space for either up to 6 or 8 2.5“ drives (SSD, SATA or SAS HDD).

Detailed technical information is available in the data sheet PRIMERGY RX200 S6.

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 results

The PRIMERGY RX200 S6 was measured with Xeon series 5500 and 5600 processors. The benchmark programs were compiled with Intel C++/Fortran Compiler 11.1 and run under SUSE Linux Enterprise Server 11 (64-bit). All results have been published at http://www.spec.org.

Processor	Cores	GHz	L3 cache	Bus	TDP	SPECint_base2006 2 chips	SPECint2006 2 chips
Xeon E5503	2	2	4 MB	800 MHz	80 Watt	**20.8**	**22.4**
Xeon E5506	4	2.13	4 MB	800 MHz	80 Watt	**22.0**	**23.9**
Xeon E5507	4	2.27	4 MB	800 MHz	80 Watt	**23.0**	**25.1**
Xeon L5609	4	1.87	12 MB	800 MHz	40 Watt	**21.2**	**22.7**
Xeon L5630	4	2.13	12 MB	1067 MHz	40 Watt	**26.1**	**28.2**
Xeon E5620	4	2.40	12 MB	1067 MHz	80 Watt	**29.6**	**31.9**
Xeon E5630	4	2.53	12 MB	1067 MHz	80 Watt	**30.7**	**33.2**
Xeon E5640	4	2.67	12 MB	1067 MHz	80 Watt	**32.0**	**34.5**
Xeon X5667	4	3.07	12 MB	1333 MHz	95 Watt	**37.8**	**40.8**
Xeon X5677	4	3.47	12 MB	1333 MHz	130 Watt	**40.1**	**43.5**
Xeon L5640	6	2.27	12 MB	1067 MHz	60 Watt	**30.6**	**33.2**
Xeon X5650	6	2.67	12 MB	1333 MHz	95 Watt	**34.3**	**36.9**
Xeon X5660	6	2.80	12 MB	1333 MHz	95 Watt	**35.4**	**38.3**
Xeon X5670	6	2.93	12 MB	1333 MHz	95 Watt	**36.7**	**39.6**
Xeon X5680	6	3.33	12 MB	1333 MHz	130 Watt	**38.9**	**41.9**

Processor	Cores	GHz	L3 cache	Bus	TDP	SPECint_rate_base2006		SPECint_rate2006	
						1 chip	2 chips	1 chip	2 chips
Xeon E5503	2	2	4 MB	800 MHz	80 Watt	**37.0**	**73.0**	**40.2**	**79.2**
Xeon E5506	4	2.13	4 MB	800 MHz	80 Watt	**71.2**	**139**	**76.2**	**148**
Xeon E5507	4	2.27	4 MB	800 MHz	80 Watt	**74.1**	**144**	**79.2**	**154**
Xeon L5609	4	1.87	12 MB	800 MHz	40 Watt	**70.0**	**135**	**75.7**	**146**
Xeon L5630	4	2.13	12 MB	1067 MHz	40 Watt	**94.5**	**181**	**101**	**193**
Xeon E5620	4	2.40	12 MB	1067 MHz	80 Watt	**107**	**209**	**114**	**223**
Xeon E5630	4	2.53	12 MB	1067 MHz	80 Watt	**112**	**218**	**118**	**232**
Xeon E5640	4	2.67	12 MB	1067 MHz	80 Watt	**116**	**225**	**123**	**239**
Xeon X5667	4	3.07	12 MB	1333 MHz	95 Watt	**137**	**268**	**145**	**284**
Xeon X5677	4	3.47	12 MB	1333 MHz	130 Watt	**145**	**284**	**153**	**300**
Xeon L5640	6	2.27	12 MB	1067 MHz	60 Watt	**144**	**275**	**154**	**295**
Xeon X5650	6	2.67	12 MB	1333 MHz	95 Watt	**165**	**322**	**176**	**344**
Xeon X5660	6	2.80	12 MB	1333 MHz	95 Watt	**170**	**331**	**181**	**354**
Xeon X5670	6	2.93	12 MB	1333 MHz	95 Watt	**174**	**338**	**185**	**362**
Xeon X5680	6	3.33	12 MB	1333 MHz	130 Watt	**182**	**353**	**194**	**379**

Processor	Cores	GHz	L3 cache	Bus	TDP	SPECfp_base2006 2 chips	SPECfp2006 2 chips
Xeon E5503	2	2	4 MB	800 MHz	80 Watt	**24.2**	**26.0**
Xeon E5506	4	2.13	4 MB	800 MHz	80 Watt	**26.3**	**28.3**
Xeon E5507	4	2.27	4 MB	800 MHz	80 Watt	**27.5**	**29.5**
Xeon L5609	4	1.87	12 MB	800 MHz	40 Watt	**25.7**	**27.5**
Xeon L5630	4	2.13	12 MB	1067 MHz	40 Watt	**30.7**	**32.9**
Xeon E5620	4	2.40	12 MB	1067 MHz	80 Watt	**34.2**	**37.0**
Xeon E5630	4	2.53	12 MB	1067 MHz	80 Watt	**35.5**	**38.2**
Xeon E5640	4	2.67	12 MB	1067 MHz	80 Watt	**36.7**	**39.5**
Xeon X5667	4	3.07	12 MB	1333 MHz	95 Watt	**43.4**	**46.9**
Xeon X5677	4	3.47	12 MB	1333 MHz	130 Watt	**45.8**	**49.2**
Xeon L5640	6	2.27	12 MB	1067 MHz	60 Watt	**36.2**	**39.1**
Xeon X5650	6	2.67	12 MB	1333 MHz	95 Watt	**40.3**	**43.3**
Xeon X5660	6	2.80	12 MB	1333 MHz	95 Watt	**41.1**	**44.3**
Xeon X5670	6	2.93	12 MB	1333 MHz	95 Watt	**42.2**	**45.5**
Xeon X5680	6	3.33	12 MB	1333 MHz	130 Watt	**45.0**	**48.3**

Processor	Cores	GHz	L3 cache	Bus	TDP	SPECfp_rate_base2006		SPECfp_rate2006	
						1 chip	2 chips	1 chip	2 chips
Xeon E5503	2	2	4 MB	800 MHz	80 Watt	**37.4**	**72.6**	**38.7**	**75.1**
Xeon E5506	4	2.13	4 MB	800 MHz	80 Watt	**60.5**	**117**	**62.3**	**120**
Xeon E5507	4	2.27	4 MB	800 MHz	80 Watt	**62.2**	**120**	**64.1**	**123**
Xeon L5609	4	1.87	12 MB	800 MHz	40 Watt	**63.1**	**116**	**65.3**	**120**
Xeon L5630	4	2.13	12 MB	1067 MHz	40 Watt	**72.9**	**133**	**75.9**	**139**
Xeon E5620	4	2.40	12 MB	1067 MHz	80 Watt	**83.9**	**162**	**87.3**	**168**
Xeon E5630	4	2.53	12 MB	1067 MHz	80 Watt	**87.0**	**167**	**90.2**	**174**
Xeon E5640	4	2.67	12 MB	1067 MHz	80 Watt	**88.8**	**171**	**92.2**	**177**
Xeon X5667	4	3.07	12 MB	1333 MHz	95 Watt	**105**	**205**	**109**	**212**
Xeon X5677	4	3.47	12 MB	1333 MHz	130 Watt	**110**	**214**	**114**	**222**
Xeon L5640	6	2.27	12 MB	1067 MHz	60 Watt	**107**	**198**	**110**	**205**
Xeon X5650	6	2.67	12 MB	1333 MHz	95 Watt	**119**	**232**	**123**	**241**
Xeon X5660	6	2.80	12 MB	1333 MHz	95 Watt	**122**	**237**	**126**	**245**
Xeon X5670	6	2.93	12 MB	1333 MHz	95 Watt	**123**	**240**	**128**	**249**
Xeon X5680	6	3.33	12 MB	1333 MHz	130 Watt	**128**	**246**	**132**	**255**

The throughput with two processors both with the integer as well as the floating-point test suite is almost twice as large as that with one processor.

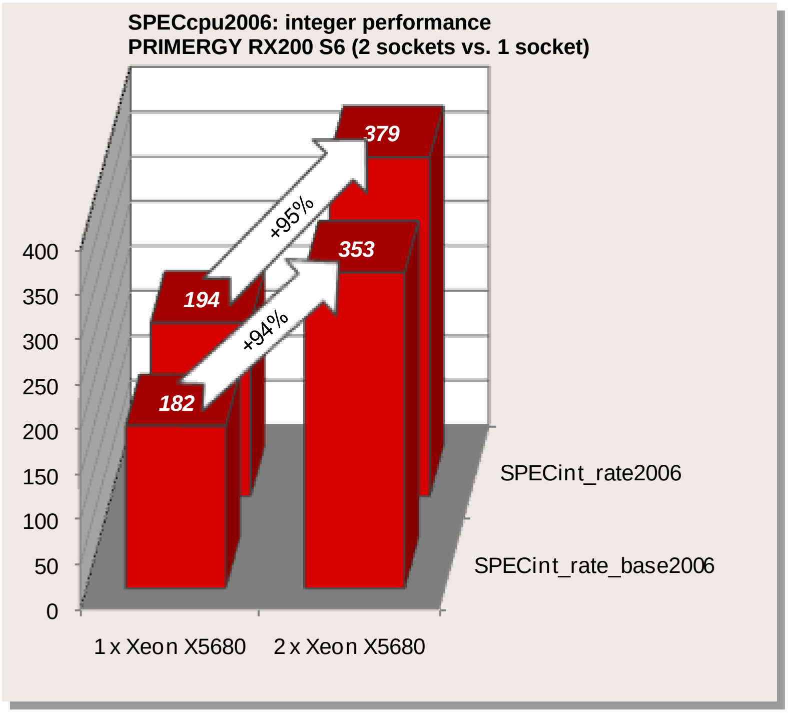

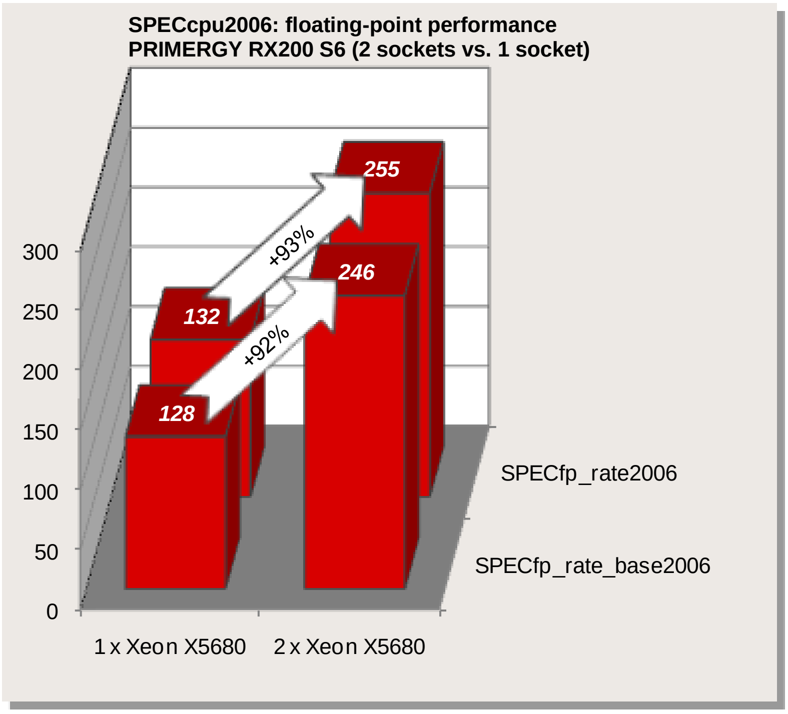

The following diagrams illustrate the throughput of the PRIMERGY RX200 S6 in comparison to its predecessor, the PRIMERGY RX200 S5, in the respective most performant configuration.

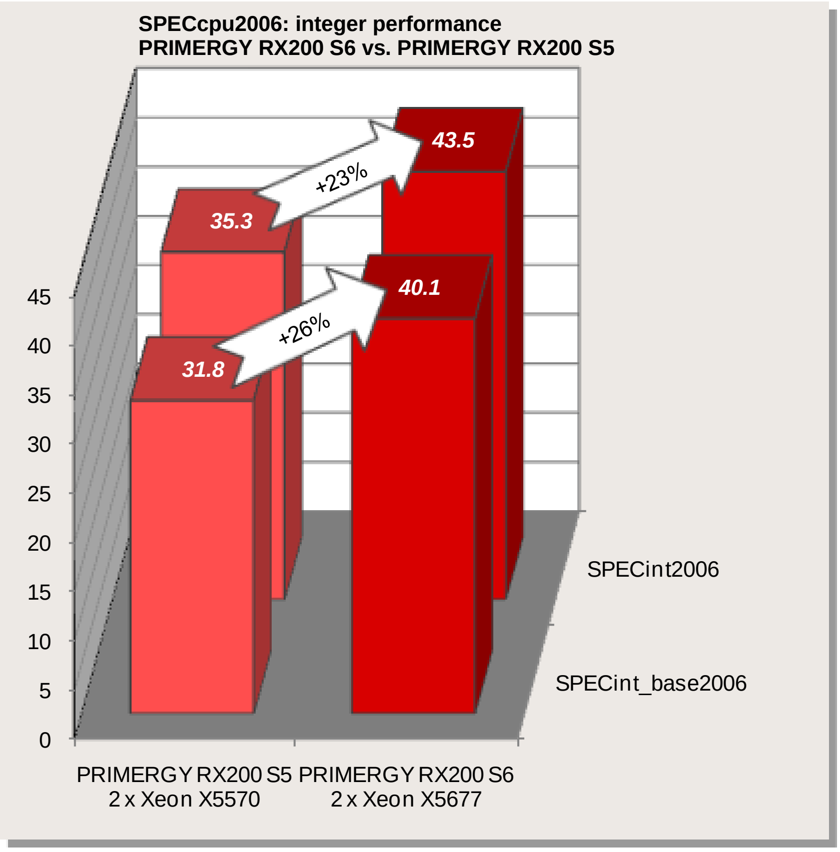

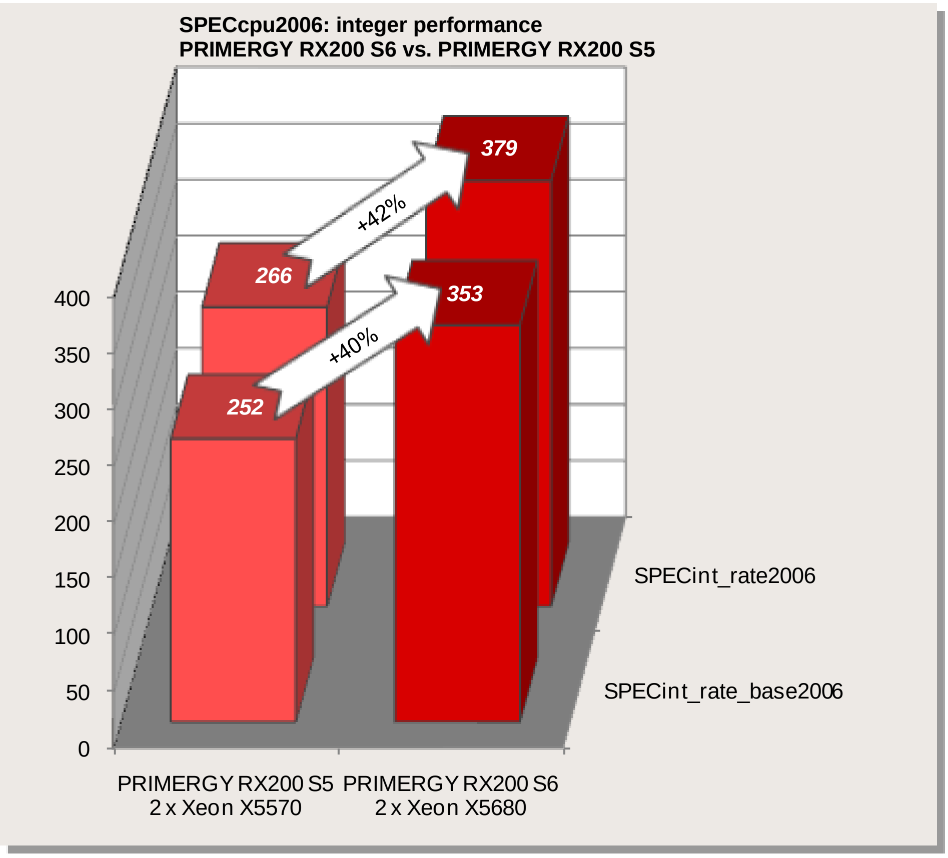

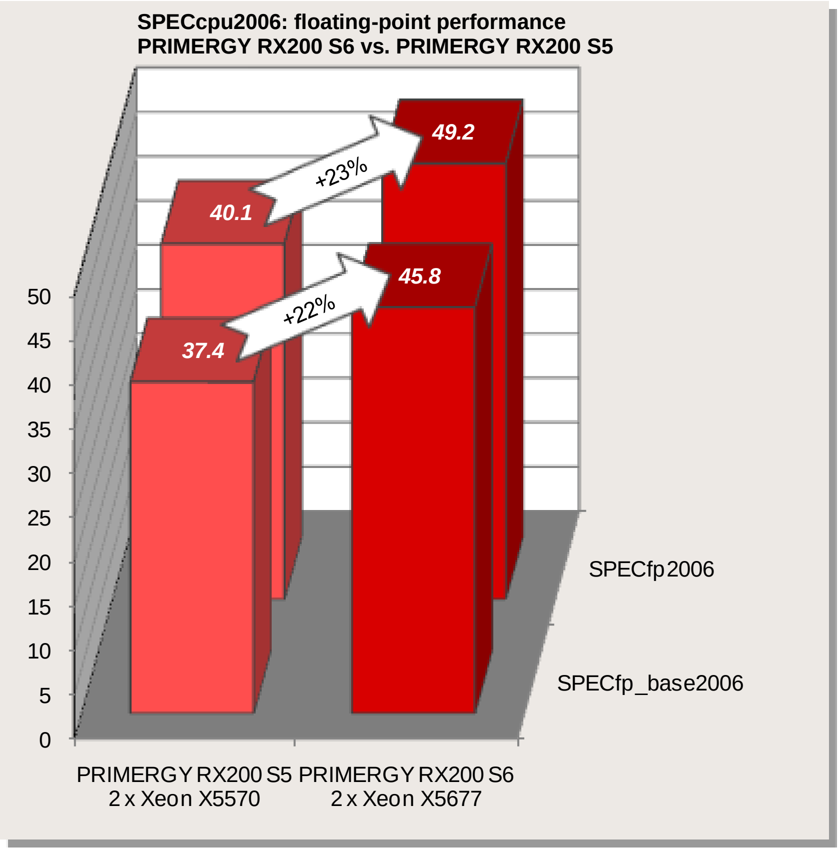
SPECcpu2006: floating-point performance
PRIMERGY RX200 S6 vs. PRIMERGY RX200 S5
49.2
+23%
40.1
45.8
+22%
37.4
0
5
10
15
20
25
30
35
40
45
50
SPECfp2006
SPECfp_base2006
PRIMERGY RX200 S5
2 x Xeon X5570
PRIMERGY RX200 S6
2 x Xeon X5677

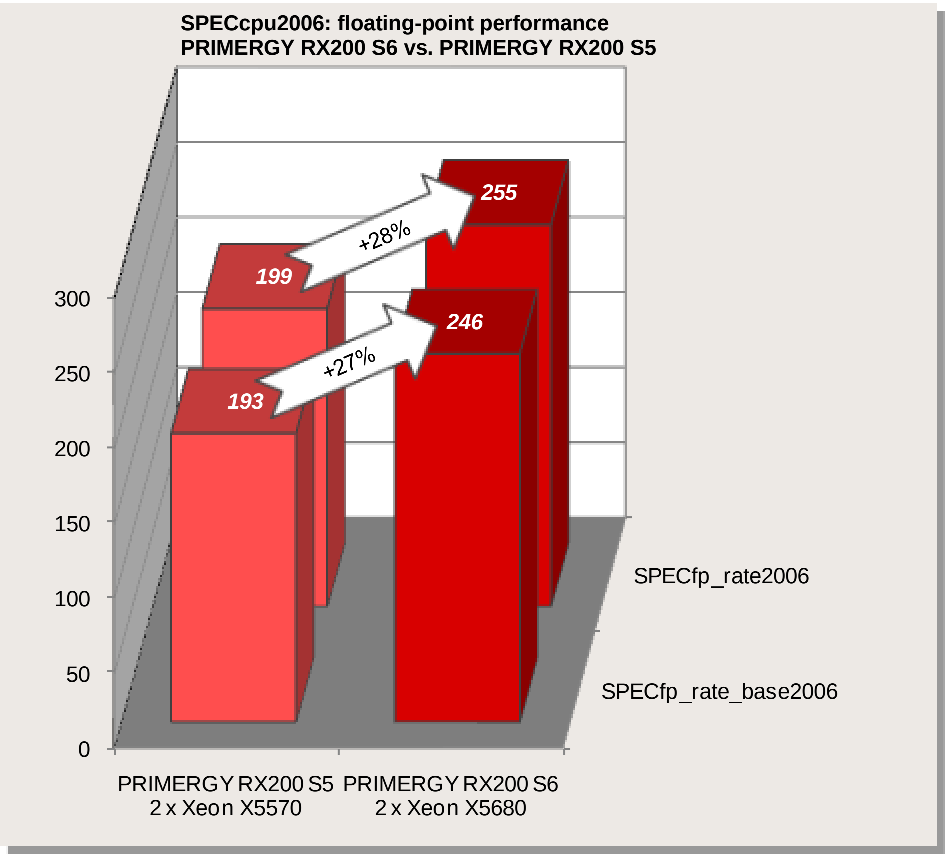
SPECcpu2006: floating-point performance
PRIMERGY RX200 S6 vs. PRIMERGY RX200 S5
255
+28%
199
246
+27%
193
0
50
100
150
200
250
300
SPECfp_rate2006
SPECfp_rate_base2006
PRIMERGY RX200 S5
2 x Xeon X5570
PRIMERGY RX200 S6
2 x Xeon X5680

Benchmark environment

All SPECcpu2006 measurements were made on a PRIMERGY RX200 S6 with the following hardware and software configuration:

Hardware	
Model	PRIMERGY RX200 S6
CPU	Xeon E5503, E5506, E5507, L5609, L5630, E5620, E5630, E5640, X5667, X5677, L5640, X5650, X5660, X5670, X5680
Number of CPUs	1 chip: Xeon E5503: 2 cores Xeon E5506, E5507, L5609, L5630, E5620, E5630, E5640, X5667, X5677: 4 cores all others: 6 cores 2 chips: Xeon E5503: 4 cores Xeon E5506, E5507, L5609, L5630, E5620, E5630, E5640, X5667, X5677: 8 cores all others: 12 cores
Primary cache	32 kB instruction + 32 kB data on chip, per core
Secondary cache	256 kB on chip, per core
Other cache	Xeon E5503, E5506, E5507: 4 MB (I+D) on chip, per chip all others: 12 MB (I+D) on chip, per chip
Software	
Operating System	SUSE Linux Enterprise Server 11 (64-bit)
Compilers	Intel C++/Fortran Compiler 11.1

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

SPECjbb2005

Benchmark description

SPECjbb2005 is a Java business benchmark that focuses on the performance of Java Server platforms. SPECjbb2005 is essentially a modernized SPECjbb2000. The main differences are:

- The transactions have become more complex in order to cover a greater functional scope.
- The working set of the benchmark has been enlarged to the extent that the total system load has increased.
- SPECjbb2000 allows only one active Java Virtual Machine instance (JVM) whereas SPECjbb2005 permits several instances, which in turn achieves greater closeness to reality, particularly with large systems.

On the software side SPECjbb2005 primarily measures the performance of the JVM used with its just-in-time compiler as well as their thread and garbage collection implementation. Some aspects of the operating system used also play a role. As far as hardware is concerned, it measures the efficiency of the CPUs and caches, the memory subsystem and the scalability of shared memory systems (SMP). Disk and network I/O are irrelevant.

SPECjbb2005 emulates a 3-tier client/server system that is typical for modern business process applications with the emphasis on the middle-tier system:

- Clients generate the load, consisting of driver threads, which on the basis of TPC-C benchmark generate OLTP accesses to a database without thinking times.
- The middle tier system implements the business processes and the updating of the database.
- The database takes on the data management and is emulated by Java objects that are in the memory. Transaction logging is implemented on an XML basis.

The major advantage of this benchmark is that it includes all three tiers that run together on a single host. The performance of the middle-tier is measured. Large-scale hardware installations are thus avoided and direct comparisons between the SPECjbb2005 results from the various systems are possible. Client and database emulation are also written in Java.

SPECjbb2005 only needs the operating system as well as a Java Virtual Machine with J2SE 5.0 features.

The scaling unit is a warehouse with approx. 25 MB Java objects. Precisely one Java thread per warehouse executes the operations on these objects. The business operations are assumed by TPC-C:

- New Order Entry
- Payment
- Order Status Inquiry
- Delivery
- Stock Level Supervision
- Customer Report

However, these are the only features SPECjbb2005 and TPC-C have in common. The results of the two benchmarks are not comparable.

SPECjbb2005 has 2 performance metrics:

- bops (business operations per second) is the overall rate of all business operations performed per second.
- bops/JVM is the ratio of the first metrics and the number of active JVM instances.

In comparisons of various SPECjbb2005 results, both metrics must be specified.

The following rules, according to which a compliant benchmark run has to be performed, are the basis for these three metrics:

A compliant benchmark run consists of a sequence of measuring points with an increasing number of warehouses (and thus of threads) with the number in each case being increased by one warehouse. The run is started at one warehouse up through 2*MaxWh, but not less than 8 warehouses. MaxWh is the number of warehouses with the highest rate per second the benchmark expects. Per default the benchmark equates MaxWh with the number of CPUs visible by the operating system.

The metric bops is the arithmetic average of all measured operation rates with MaxWh warehouses up to 2*MaxWh warehouses.

Benchmark results

In March 2010 the PRIMERGY RX200 S6 with two Xeon X5680 processors and a memory of 48 GB PC3-10600R DDR3-SDRAM was measured. The measurement was made under Windows Server 2008 R2 Enterprise. Six J9 VM instances from IBM were used.

The following result was obtained:

SPECjbb2005 bops = 927067

SPECjbb2005 bops/JVM = 154511

The following graphics illustrate the throughput of the PRIMERGY RX200 S6 in comparison to its predecessor PRIMERGY RX200 S5, in the respective most performant configuration.

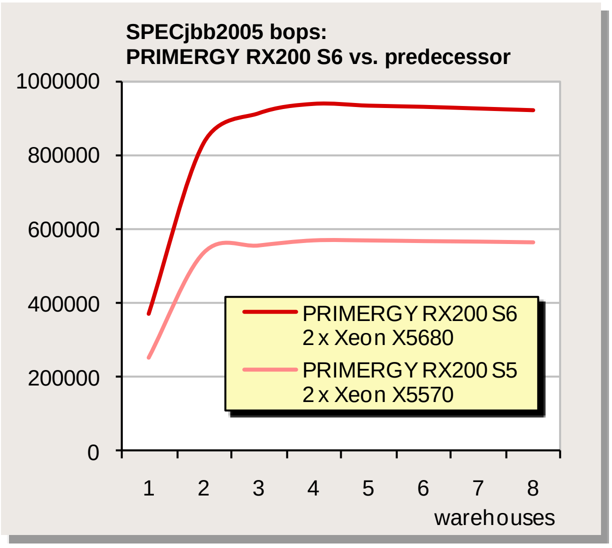

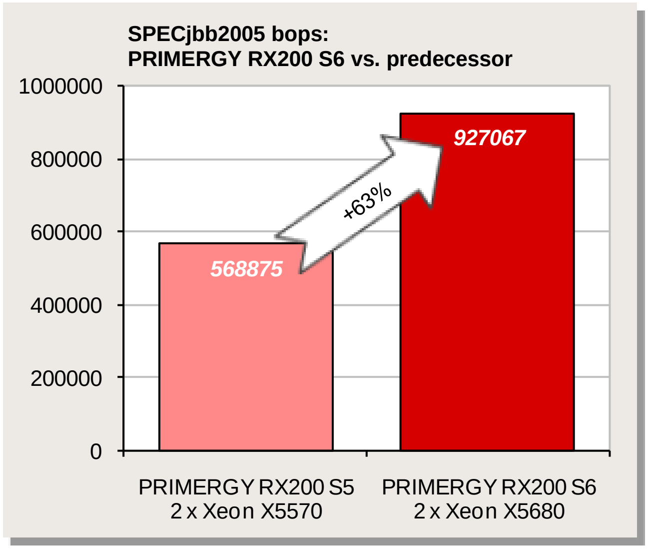

Benchmark environment

The SPECjbb2005 measurement was run on a PRIMERGY RX200 S6 with the following hardware and software configuration:

Hardware	
Model	PRIMERGY RX200 S6
CPU	Xeon X5680
Number of chips	2 chips, 12 cores, 6 cores per chip
Primary cache	32 kB instruction + 32 kB data on chip, per core
Secondary cache	¼ MB (I+D) on chip, per core
Other cache	12 MB (I+D) on chip, per chip
Memory	12 x 4 GB PC3-10600R DDR3-SDRAM
Software	
Operating System	Windows Server 2008 R2 Enterprise
JVM Version	IBM J9 VM (build 2.4, JRE 1.6.0 IBM J9 2.4 Windows Server 2008 amd64-64 jvmwa6460sr6-20090923_42924 (JIT enabled, AOT enabled)

Some components are not available in all countries / sales regions.

SPECpower_ssj2008

Benchmark description

SPECpower_ssj2008 is the first industry-standard SPEC benchmark that evaluates the power and performance characteristics of server class computers. With SPECpower_ssj2008, SPEC has defined server power measurement standards 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 operating environments, and economical to run. It exercises CPUs, caches, memory hierarchy, and the scalability of symmetric multiprocessor systems (SMPs), as well as implementations of the Java Virtual Machine (JVM), Just In Time (JIT) compiler, garbage collection, threads, and some aspects of the operating system.

SPECpower_ssj2008 reports power consumption for servers at different performance levels — from 100-percent to active idle in 10-percent 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 gives information about the energy efficiency of the measured server. Because of its defined measurement standard it allows the customers to compare it to other configurations and servers measured with SPECpower_ssj2008. The adjoining diagram shows a typical graph of a SPECpower_ssj2008 result.

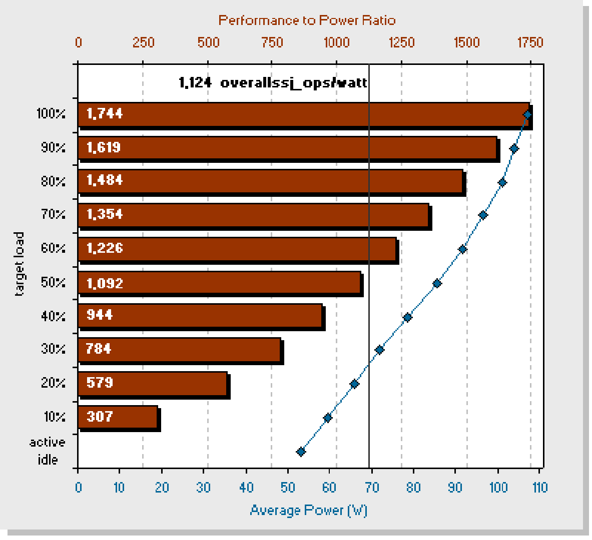

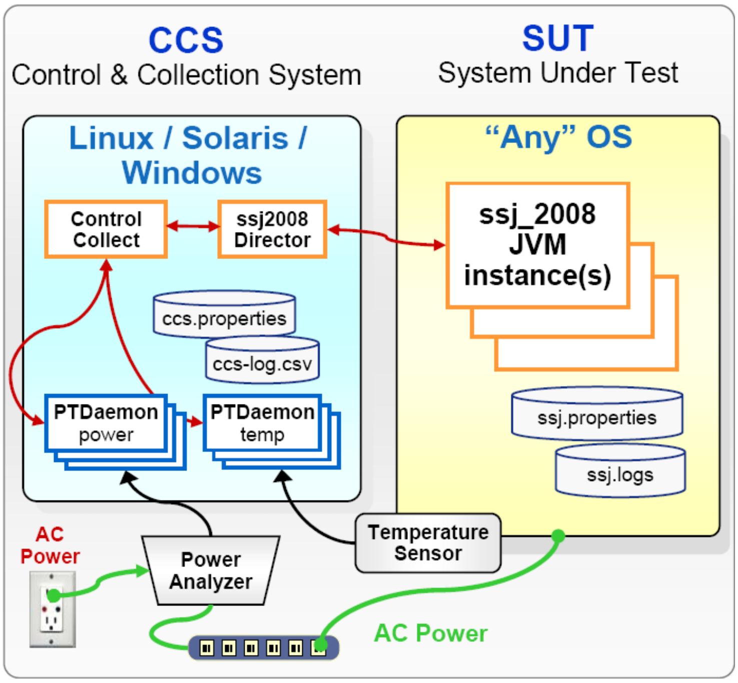

The benchmark runs on a wide variety of operating systems and hardware architectures and does not require extensive client or storage infrastructure. The minimum equipment for SPEC-compliant testing is two networked computers, plus a power analyzer and a temperature sensor. One computer is the System Under Test (SUT) running any of the supported operating systems along with the JVM installed. The JVM provides the environment required to run the SPECpower_ssj2008 workload which is implemented in Java. The other computer is a Collect and Control System (CCS) which controls the operation of the benchmark and captures the power, performance and temperature readings for reporting. The adjoining diagram gives an overview of the different components of this framework.

Benchmark results

In March 2010 the PRIMERGY RX200 S6 was measured with two Intel Xeon X5670 processors and 12 GB of PC3L-10600E DDR3-SDRAM memory. The measurement was taken under Windows Server 2008 Enterprise x64 Edition with SP2 and a J9 2.4 VM by IBM was used.

With the Intel Xeon X5670 processors the PRIMERGY RX200 S6 achieved a result of **2,850 overall ssj_ops/watt**.

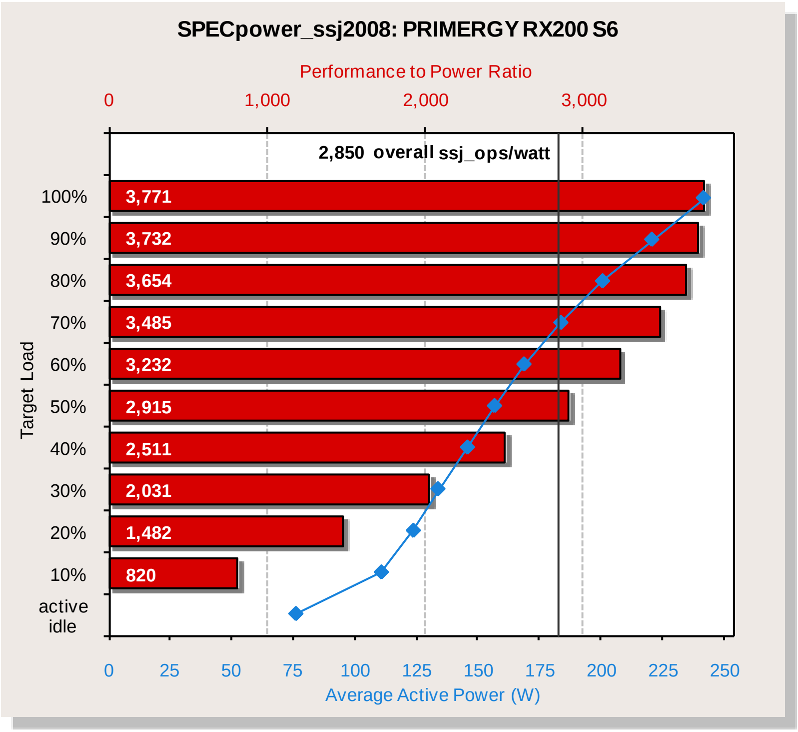

The adjoining diagram shows the result graph of the configuration described above, measured with the PRIMERGY RX200 S6. The red horizontal bars show the performance to power ratio in ssj_ops/watt (upper x-axis) for each target load level which are tagged on the y-axis of the diagram. The blue line shows 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 2,850 overall ssj_ops/watt for the PRIMERGY RX200 S6. This is the quotient of the sum of the transaction throughputs for each measurement interval and the sum of the average power consumed for each measurement interval.

The following table shows the detailed information for the throughput in ssj_ops, the average power consumption in watts and the resulting energy efficiency for each load level of the measured result.

Performance		**Power**	**Energy Efficiency**
Target Load	ssj_ops	Average Power (W)	ssj_ops/watt
100%	913,790	242	3,771
90%	825,810	221	3,732
80%	735,628	201	3,654
70%	640,194	184	3,485
60%	547,810	169	3,232
50%	459,091	157	2,915
40%	365,763	146	2,511
30%	273,004	134	2,031
20%	183,518	124	1,482
10%	91,022	111	820
Active Idle	0	76.2	0
		∑ssj_ops / ∑power =	**2,850**

The configuration was tuned to get the best possible result for this server in terms of performance per watt. The memory configuration with 6 x 2 GB was selected to meet the criteria of best performance at lowest power consumption by populating only one slot of each available memory channel with Low Voltage DIMMs. This configuration enables the benchmark to use the full potential of the available memory bandwidth and at the same time consumes less power than a 2 or 3 DIMMs/channel configuration. The most important factor in the hardware configuration is the right choice of the processor. Processors are the part of a server which consumes the most power beside the memory subsystem. For the PRIMERGY RX200 S6 the hexa-core Intel Xeon X5670 processors with a Thermal Design Power (TDP) of 95 watt showed the best efficiency score.

Benchmark environment

The SPECpower_ssj2008 measurement presented here was performed on a PRIMERGY RX200 S6 with the following hardware and software configuration using the ZES Zimmer LMG95 power analyzer:

Hardware	
Model	PRIMERGY RX200 S6
Processor (TDP)	Intel Xeon X5670 (95 W)
Number of chips	2 chips, 6 cores per chip, 2 threads per core
Primary Cache	32 KB instruction + 32 KB data on chip, per core
Secondary Cache	256 KB (I+D) on chip, per core
Tertiary Cache	12 MB (I+D) on chip, per chip
Memory	6 x 2 GB PC3L-10600E DDR3-SDRAM
Network Interface	2 x 1 GBit LAN Intel 82575EB Gigabit Network Connection (onboard)
Disk Subsystem	1 x Integrated SATA controller 1 x 2.5” SATA SSD, 64 GB, JBOD
Power Supply Unit	1 x 450 W Fujitsu Technology Solutions DPS-400AB-10 A
Software	
Operating System	Windows Server 2008 Enterprise x64 Edition + SP2
JVM Version	IBM J9 VM (build 2.4, JRE 1.6.0 IBM J9 2.4 Windows Server 2008 amd64-64 jvmwa6460sr6-20090923_42924 (JIT enabled, AOT enabled)
JVM affinity	start /affinity [0xF,0xF0,0xF00,0xF000,0xF0000,0xF00000]
JVM options	-Xaggressive -Xcompressedrefs -Xgcpolicy:gencon -Xmn1400m -Xms1550m -Xmx1550m -XlockReservation -Xnoloa -XtlhPrefetch -Xlp

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

vServCon

Benchmark description

vServCon is a benchmark used by Fujitsu Technology Solutions 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.

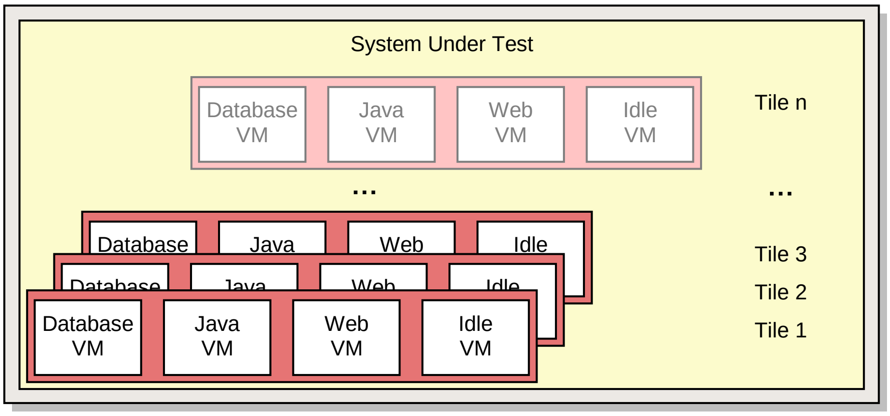

Each of the three vServCon application scenarios provides a specific benchmark result in the form of application-specific transaction rates for the respective VM. In order to derive a normalized score, the individual benchmark results for one tile are put in relation to the respective results of a reference 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 available in the document: Benchmark Overview vServCon.

Benchmark results

The PRIMERGY RX200 S6 is 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 50% higher virtualization performance can be achieved (measured in vServCon score) as 6-Core processors are also available. On the basis of the previously described vServCon profile almost optimal utilization of the CPU system resources is possible with 27 real application VMs (equivalent to nine tiles) if the system is fully assembled with two such processors.

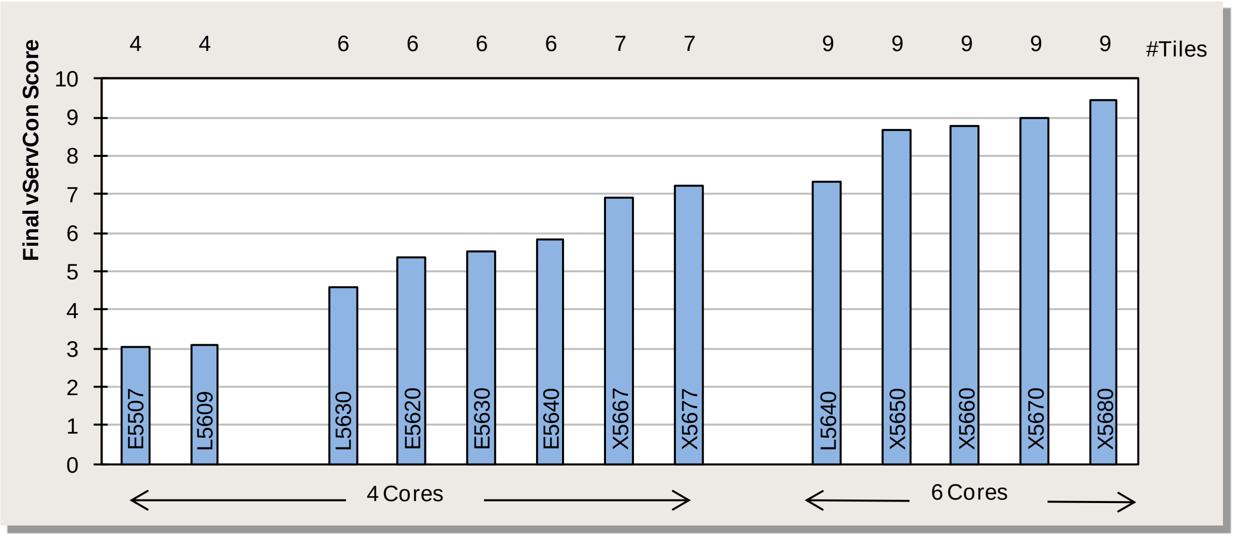

The first diagram compares the virtualization performance values that can be achieved with the individual processors. A large selection of released system processors with four or six cores was considered.

The relatively large performance differences between the processors as seen in the diagram can be better understood when considering the technical data in the subsequent table. The two processors in the group on the left are entry-models. When moving to the middle group, Hyper-Threading and turbo mode plays a role, and the processor-related maximum memory clock rate increases to 1066 MHz. (More information about this topic can be found in the White Paper Memory performance of XEON 5600 (Westmere-EP)-based systems.) The performance increase here is thus large. A further clear increase in performance is seen between the Xeon E5640 and X5667 processors which is explained by the increase of the processor-related memory clock rate to 1333 MHz. The right-hand group consists of the six-core processors, which - as expected - have almost 50% more performance than the corresponding four-core versions - otherwise with the same features.

Processor	#cores	L3 cache	Processor frequency	Memory clock rate	Hyper-Threading	Turbo mode	#Tiles	Score
E5507	4	4 MB	2.27 GHz	800 MHz	-	-	4	3.02
L5609	4	12 MB	1.87 GHz	800 MHz	-	-	4	3.08
L5630	4	12 MB	2.13 GHz	1066 MHz	✓	✓	6	4.59
E5620	4	12 MB	2.40 GHz	1066 MHz	✓	✓	6	5.38
E5630	4	12 MB	2.53 GHz	1066 MHz	✓	✓	6	5.54
E5640	4	12 MB	2.67 GHz	1066 MHz	✓	✓	6	5.81
X5667	4	12 MB	3.07 GHz	1333 MHz	✓	✓	7	6.93
X5677	4	12 MB	3.47 GHz	1333 MHz	✓	✓	7	7.22
L5640	6	12 MB	2.27 GHz	1333 MHz	✓	✓	9	7.34
X5650	6	12 MB	2.67 GHz	1333 MHz	✓	✓	9	8.68
X5660	6	12 MB	2.80 GHz	1333 MHz	✓	✓	9	8.79
X5670	6	12 MB	2.93 GHz	1333 MHz	✓	✓	9	9.00
X5680	6	12 MB	3.33 GHz	1333 MHz	✓	✓	9	9.43

The next diagram illustrates the virtualization performance of the PRIMERGY RX200 S6 for increasing numbers of VMs based on the Xeon E5640 (4-core) and X5650 (6-core) processors. The respective CPU loads of the host have also been entered. The number of tiles with optimal CPU load is typically at about 90%; beyond that you have overload, which is where virtualization performance no longer increases, or sinks again.

In addition to the increased number of physical cores, Hyper-Threading is an additional reason for the high number of operable VMs. 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.

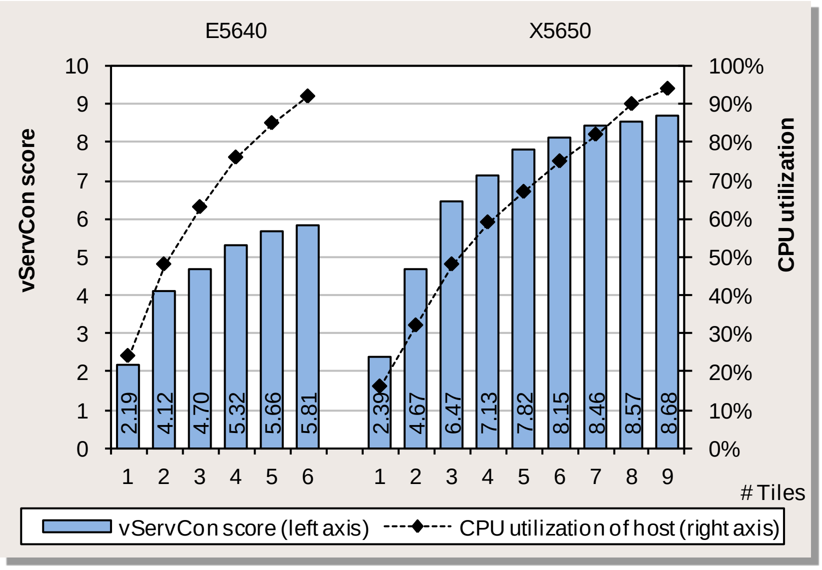

The scaling curves for the number of tiles as seen in the previous diagram are specifically for systems with Hyper-Threading. 12 physical and thus 24 logical cores are available with the Xeon X5650 processor; approximately four of them are used per tile (see Benchmark description). This means that a parallel use of the same physical cores by several VMs is avoided up to a maximum of about three 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.

A guideline in the virtualization environment for selecting main memory is that sufficient quantity is more important than the speed of the memory accesses.

Previously, the virtualization performance of the system was analyzed as a whole. Below, performance is also to be discussed from the viewpoint of an individual application VM in the described virtualized environment. The following uses the system with the Xeon X5650 processor as an example.

If the number of application VMs is optimal as far as the overall performance is concerned, the performance of an individual VM is already notably lower than in operational low-load situations. The next diagram illustrates this via the relative performance in relation to the reference system with one individual application VM of each of the three types for increasing VM numbers. The first column of a group views one VM in the array of a total of three application VMs (1 tile), the second one is for the array of 6 application VMs (2 tiles), etc. The values are presented - both individually and in total for all VMs of the respective type - through the height of the stacked columns.

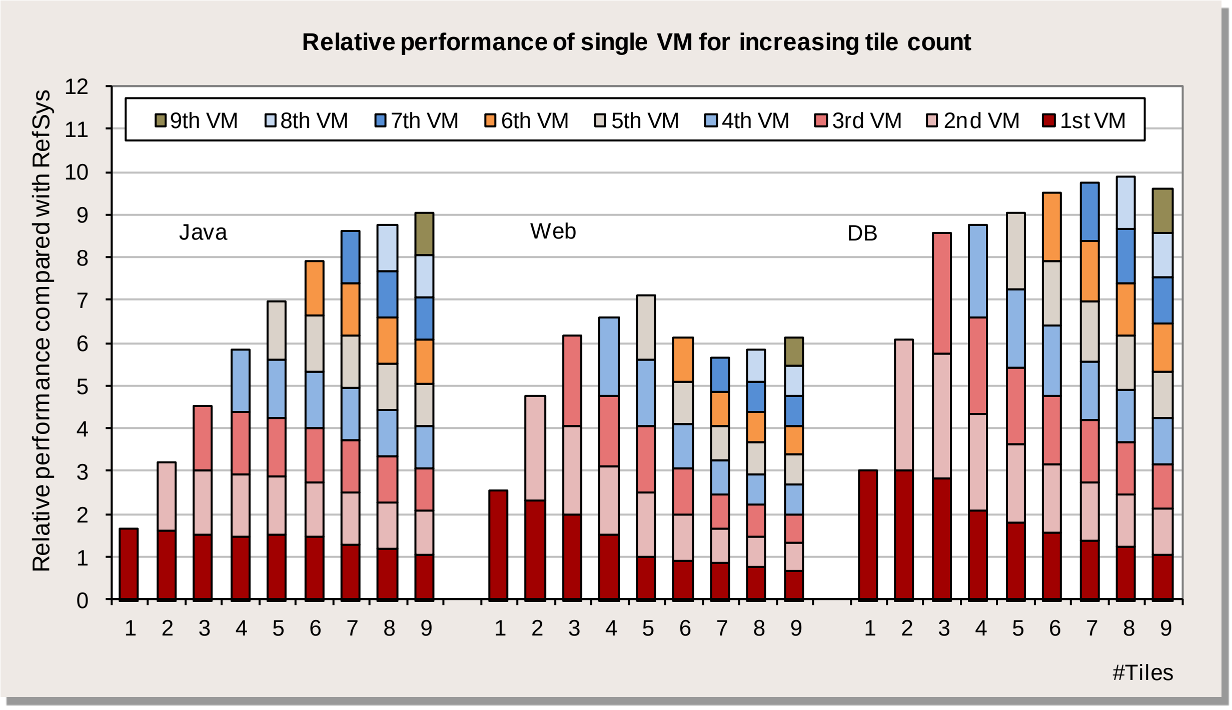

With regard to the VM numbers on a virtualization host it is necessary in a specific case to weigh up the performance requirements of an individual application against the overall requirements.

The virtualization-relevant progress in processor technology since 2008 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. Four systems are compared each with 2.67 GHz processor frequency: a system from 2008 with 2 × Xeon E5430, a system from 2009 with 2 × Xeon X5550 and a PRIMERGY RX200 S6 with 2 × Xeon E5640 or 2 × Xeon X5650.

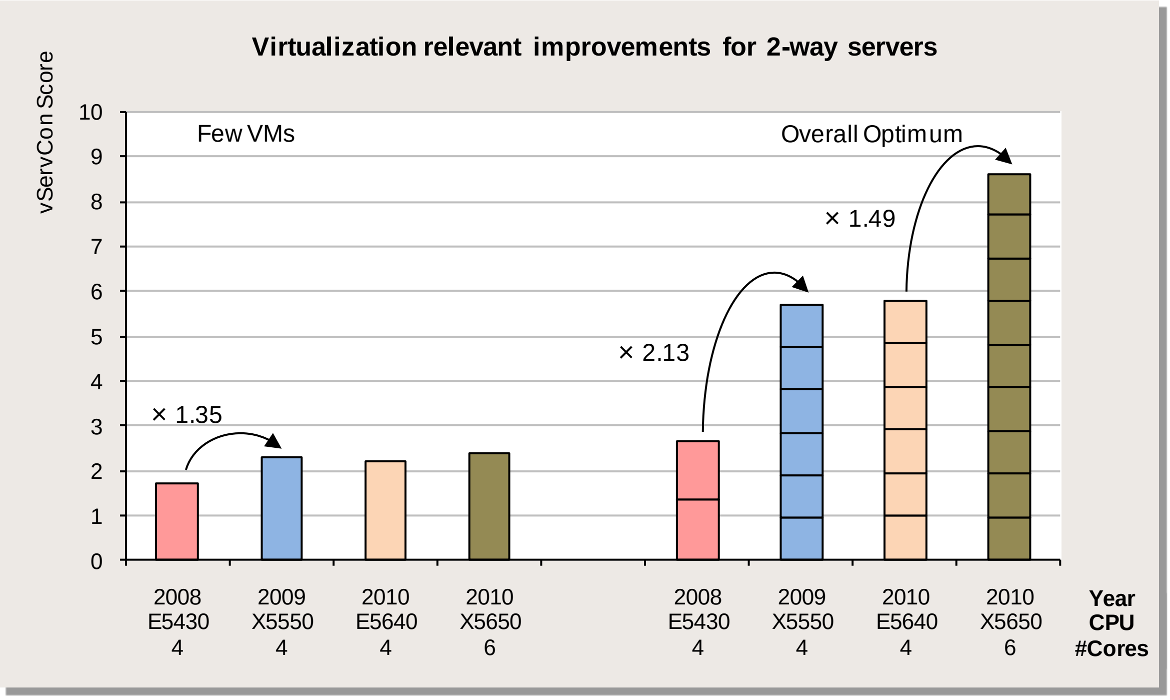

The clearest performance improvements arose from 2008 to 2009 with the introduction of the Xeon 5500 processor generation (e. g. via the feature “Extended Page Tables” (EPT)[1]). One sees an increase of the vServCon score by a factor of 1.35 with a few VMs (one tile).

With full utilization of the systems with VMs there was an increase by a factor of 2.13. The one reason was the performance increase that could be achieved for an individual VM (see score for a few VMs). The other reason was that more VMs were possible with total optimum (via Hyper-Threading and the increase in the number of physical cores). However, it can be seen that the optimum was "bought" with a triple number of VMs with a reduced performance of the individual VM.

Where exactly is the technology progress between 2009 and 2010? The performance for an individual VM in low-load situations has basically remained the same for the processors compared here with the same clock frequency but with different cache size and speed of memory connection. The progress is, on the one hand, the higher number of cores and, on the other hand, the reduced amount of electrical power consumption at full utilization. The guideline is usually seen in TDP (Thermal Design Power). The Xeon X5550 processor used in 2009 has TDP of 95 W, whereas TDP for the Xeon E5640 processor currently in use is only 80 W. The system with the more recent processor still attains the same vServCon score with full utilization with VMs (6 tiles).

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. More than approximately 30% to 50% increased throughput compared to an identically clocked processor of the Xeon 5400 generation from 2008 is not possible here. Performance increases in the virtualization environment since 2009 are mainly achieved by increased VM numbers due to the increased number of available logical or physical cores.

[1] EPT accelerates memory virtualization via hardware support for the mapping between host and guest memory addresses.

Benchmark environment

The measurements were made with the environment described below:

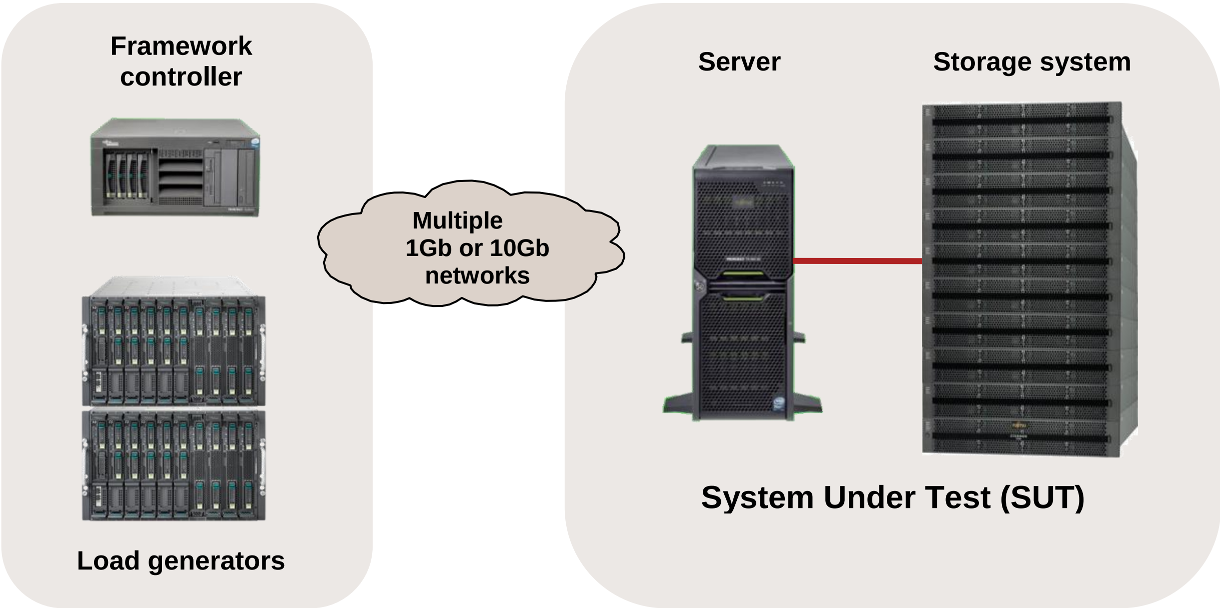

SUT hardware		
Model	PRIMERGY RX200 S6	
Processor	2 × Xeon E5507 (4C, 2.27 GHz) 2 × Xeon L5609 (4C, 1.87 GHz) 2 × Xeon L5630 (4C, 2.13 GHz) 2 × Xeon E5620 (4C, 2.40 GHz) 2 × Xeon E5630 (4C, 2.53 GHz) 2 × Xeon E5640 (4C, 2.67 GHz) 2 × Xeon X5667 (4C, 3.07 GHz)	2 × Xeon X5677 (4C, 3.47 GHz) 2 × Xeon L5640 (6C, 2.27 GHz) 2 × Xeon X5650 (6C, 2.67 GHz) 2 × Xeon X5660 (6C, 2.80 GHz) 2 × Xeon X5670 (6C, 2.93 GHz) 2 × Xeon X5680 (6C, 3.33 GHz)
Memory	96 GB (a PC3-10600R each, 8 GB, in DIMM-1A until DIMM-1F and in DIMM-2A until DIMM-2F)	
Network interface	2 × 1-GBit LAN; one for load (via 2 LAN adapters), one for control.	
Disk subsystem	No internal hard disks were used, but only one storage system FibreCAT CX500. One 50 GB LUN per tile for the “virtual disk files” of the VMs. Each LUN is a RAID 0 array consisting of 5 Seagate ST373454 disks (15 krpm)	
Storage connection	Via FC controller Qlogic QLE 2460	
SUT software		
Operating System	Hypervisor VMware ESX Server	
Release	Version 4.0 U1; Build 236512	
BIOS	Version 6.00 R1.02.3031; deviations from default: Adjacent Cache Line Prefetch: Disabled; Hardware Prefetch: Disabled DCU Streamer Prefetch: Disabled; Data Reuse Optimization: Disabled	
SUT: virtualization-specific details		
ESX settings	Default	
General details	Described in the Benchmark Overview vServCon.	
Load generator hardware		
Model	2 server blades per tile in PRIMERGY BX600 S3 chassis	
Processor	2 × Xeon 5130 each, 2000 MHz	
Memory	1 – 2 GB	
Network interface	2 × 1 GBit LAN each	
Operating System	W2K3 EE	

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

VMmark

Benchmark description

VMmark is a benchmark developed by VMware to compare server configurations with hypervisor solutions from VMware regarding their suitability for server consolidation. VMmark is currently the only established benchmark for this purpose. In addition to the software for load generation, the benchmark consists of a defined load profile and binding regulations. Benchmark results achieved with VMmark can be submitted to VMware and are published on their Internet site after a successful review process.

For a benchmark like VMmark to fulfil its objective, it must map the real world of a data center regarding server consolidation; in other words it must consider existing servers with those application scenarios that are normally virtualized. These servers have weak utilization levels and the aim is thus to consolidate as many of them as possible as VMs. Therefore, such a benchmark must assess for a virtualization host both the suitably determined overall throughput across the various application VMs as well as the number of efficiently operable VMs.

The following solution concept has been established for these two objectives: a representative group of application scenarios is selected in the benchmark. They are started simultaneously as VMs on a virtualization host when making a measurement. Each of these VMs is operated with a suitable load tool at a defined lower load level. Such a group of VMs is known as a "tile".

A tile in VMmark consists of six VMs; five of them are allocated to the selected application scenarios on a dedicated basis. A sixth is added, the so-called standby VM. VMmark mandatorily allocates to each VM certain resources with regard to logical processors, memory and hard disk space. The table describes these six VMs and the load tools used to measure them.

Application scenario	Load tool
Database server	Sysbench
File server	Dbench (modified)
Java application server	SPECjbb2005 (modified)
Mail server	Loadsim 2003
Web server	SPECweb2005 (modified)
Standby server	-

Depending on the performance capability of the underlying server hardware, you will - as part of a measurement - mostly have to start several identical tiles in parallel in order to achieve a maximum overall performance.

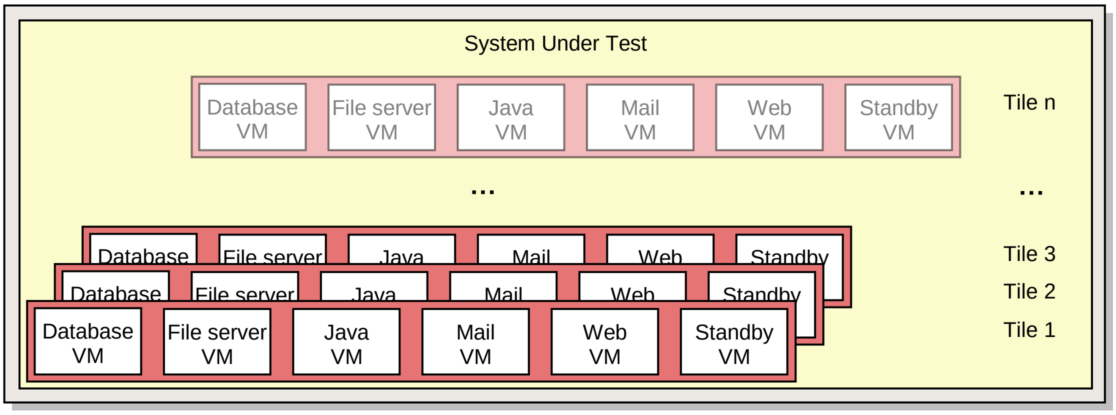

Each of the five VMmark application scenarios provides a specific result for each VM. In order to derive a score the individual results are appropriately summarized for all VMs. The outcome is the VMmark score for this tile number, that is why - in addition to the actual score - the number of tiles is always specified, e.g. "12.34@5 tiles".

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

Benchmark results

On 10th August 2010 Fujitsu achieved a VMmark score of "35.09@27 tiles" in a system configuration with a total of 12 processor cores with a PRIMERGY RX200 S6 and VMware ESX 4.0 Update 2.

This score as well as the detailed results and configuration data can be seen at http://www.vmware.com/products/vmmark/results.html.

With the score of "35.09@27 tiles" the PRIMERGY RX200 S6 is from the VMmark viewpoint the most powerful rack server with twelve cores and is at the same time second in the VMmark ranking for servers of the twelve-core category (valid as of benchmark results publication date).

The main prerequisites in attaining this result were the processor, the frequency-optimized 6-core processor Xeon X5680 and the hypervisor version which optimally uses the processor features. These features include the extended page tables (EPT) [2] and Hyper-Threading. All this has a particularly positive effect during virtualization.

On account of the requirements made by the benchmark the memory had to be extended to 192 GB (12 × 16 GB) when operating 27 tiles. As a result of the memory architecture of the 16 GB DIMMs it is then reduced to the speed of 800 MHz.

All VMs, their application data, the host operating system as well as additionally required data were on a powerful fibre channel disk subsystem from ETERNUS DX80 systems with a total of 41 LUNs.

All the components used were optimally attuned to each other.

[2] EPT accelerates memory virtualization via hardware support for the mapping between host and guest memory addresses.

Benchmark environment

The measurement set-up is symbolically illustrated below:

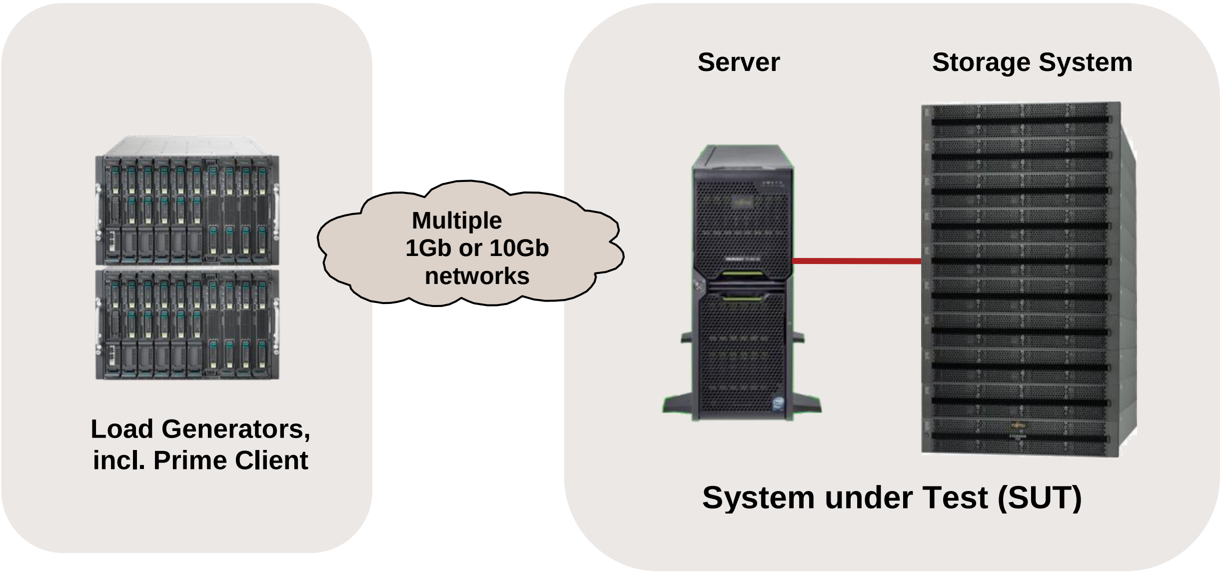

SUT hardware	
Model	PRIMERGY RX200 S6
Processor	2 × Xeon X5680 (6-Core, 3.33 GHz)
Memory	192 GB (12 × 16 GB per DIMM), 1333 MHz registered ECC DDR3
Network interface	1 × integrated Intel 82575EB dual port 1GbE adapter 1 × Intel 10 Gigabit X520-DA2 dual port server adapter
Disk subsystem	No internal hard disks were used. 9 ETERNUS DX80 storage systems: a total of 196 hard disks in several RAID-0 arrays.
Storage connection	1 × dual-channel Emulex LPe12002
SUT software	
Operating system	Hypervisor VMware ESX Server
ESX version	VMware ESX v4.0 Update 2; Build 261974
BIOS version	Rev R1.02.3031
Load generator hardware	
Model	27 × server blade PRIMERGY BX620 S4
Processor	2 × Intel Xeon 5130, 2 GHz
Memory	3 GB
Network interface	1 × 1 GBit LAN each
Operating system	Microsoft Windows Server 2003 R2 Enterprise, aktualisiert mit SP2 und KB955839
Details	
See disclosures	http://www.vmware.com/files/pdf/vmmark/VMmark-Fujitsu-2010-08-09-RX200S6.pdf

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

Literature

PRIMERGY Systems

http://ts.fujitsu.com/primergy

PRIMERGY RX200 S6

Data sheet
http://docs.ts.fujitsu.com/dl.aspx?id=fc72a397-be21-4e91-b795-6480200f7c6c

Memory performance of Xeon 5600 (Westmere-EP)-based systems
http://docs.ts.fujitsu.com/dl.aspx?id=f622cc5b-c6f4-41c5-ae86-a642b4d5d255

PRIMERGY Performance

http://ts.fujitsu.com/products/standard_servers/primergy_bov.html

SPECcpu2006

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

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

SPECjbb2005

http://www.spec.org/jbb2005

Benchmark overview SPECjbb2005
http://docs.ts.fujitsu.com/dl.aspx?id=5411e8f9-8c56-4ee9-9b3b-98981ab3e820

SPECpower_ssj2008

http://www.spec.org/power_ssj2008

Benchmark Overview SPECpower_ssj2008
http://docs.ts.fujitsu.com/dl.aspx?id=166f8497-4bf0-4190-91a1-884b90850ee0

VMmark

Benchmark overview VMmark
http://docs.ts.fujitsu.com/dl.aspx?id=5cbc7fe2-e493-479d-b52a-f88dc46efe21

VMmark
http://www.vmmark.com

VMmark results
http://www.vmware.com/products/vmmark/results.html

vServCon

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

Contact

FUJITSU Technology Solutions
Website: http://ts.fujitsu.com

PRIMERGY Product Marketing
mailto:Primergy-PM@ts.fujitsu.com

PRIMERGY Performance and Benchmarks
mailto:primergy.benchmark@ts.fujitsu.com