THE POSSIBILITIES ARE INFINITE FUJITSU

WHITE PAPER

Performance Report PRIMERGY TX120 S1

Version 1.2
October 2007

Pages 10

Abstract

This document contains a summary of the benchmarks executed for the PRIMERGY TX120 S1.

The PRIMERGY TX120 S1 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.

Contents

Technical Data

The PRIMERGY TX120 S1 is a 1-socket tower server. It includes the Intel 3000 chipset, one Celeron M or Xeon processor, up to 8 GB PC2-4200 DDR2-SDRAM, an 800 MHz (Celeron M) or 1067 MHz (Xeon) front-side bus, a Broadcom BCM5754 1-GBit LAN controller, an LSI1064E 4-port SAS controller with integrated support for RAID-0 and RAID-1, two 2.5" SAS hard disks and three PCI slots (1 x PCI 32-bit/33 MHz, 1 x PCI-Express x1 and 1 x PCI-Express x4).







See http://docs.ts.fujitsu.com/dl.aspx?id=d23cc39b-76f0-4f0a-97b9-ca35b71a54b3 for detailed technical information.

SPECcpu2000*

Benchmark description

Since SPECcpu2000 is nowadays no longer in a position to adequately meet the requirements made of it, SPEC discontinued the maintenance of this benchmark at the end of 2006. Since then it has not been possible to submit any SPECcpu2000 results to SPEC for publication. As the successor to SPECcpu2000 SPEC has developed the benchmark SPECcpu2006.

SPECcpu2000 is a benchmark to measure system efficiency during integer and floating point operations. It consists of an integer test suite containing 12 applications and a floating point test suite containing 14 applications which are extremely computing-intensive and concentrate on the CPU and memory. Other components, such as disk I/O and network, are not measured by this benchmark.

SPECcpu2000 is not bound to a specific operating system. The benchmark is available as source code and is compiled before the actual benchmark, which is why the compiler version and the optimization settings are also integrated into the measurement.

SPECcpu2000 contains two different methods of performance measurement: The first method (SPECint2000 and SPECfp2000) determines the time required to complete a specific task. The second method (SPECint_rate2000 and SPECfp_rate2000) determines the throughput, i.e. how often a task can be completed within a predefined time. Both methods are additionally subdivided into two measuring runs, "base" and "peak", which differ in the way the compiler optimization is used. The "base" values are always used when results are published, the "peak" values are optional.

Benchmark	Arithmetic	Type	Compiler optimization	Measuring result	Application
SPECint2000	integer	peak	aggressive	speed	mono processor
SPECint_base2000	integer	base	conservative		
SPECint_rate2000	integer	peak	aggressive	throughput	mono and multi processor
SPECint_rate_base2000	integer	base	conservative		
SPECfp2000	floating point	peak	aggressive	speed	mono processor
SPECfp_base2000	floating point	base	conservative		
SPECfp_rate2000	floating point	peak	aggressive	throughput	mono and multi processor
SPECfp_rate_base2000	floating point	base	conservative		

The results represent the geometric mean of normalized ratios determined for the individual benchmarks. “Normalized“ means measuring how fast the test system runs in comparison to a reference system. A SPECint_2000 and SPECfp_2000 value of “100“ was determined for the reference system. In case of rate measurement results the determined value is 1.16. If for example the measured system has a SPECint_base2000 value of 200; this means that it has executed this benchmark at least twice as fast as the reference system. The inaccurate term “at least“ is chosen as the geometric mean is used to calculate the result. The effect of this method is a weighting in favor of the lower single results compared with the arithmetical-mean method.

* SPEC®, SPECint®, SPECfp® and the SPEC logo are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).

Benchmark results

Measurements were performed with the processors Xeon 3040 and 3070. The integer test suite of the SPECcpu2000 benchmark programs was compiled with the Intel C++/Fortran compiler 9.1 and run under Windows Server 2003 Enterprise Edition SP1 (32-bit). The floating-point test suite of the SPECcpu2000 benchmark programs was compiled with the Intel C++/Fortran compiler 9.0 and run under SUSE Linux Enterprise Server 10 (64-bit).

Processor	Cores/Chip	GHz	SLC	FSB	SPECint_rate_base2000	SPECfp_rate_base2000
Celeron M 440	1	2	½ MB	800 MHz	n/a	n/a
Xeon 3040	2	1.87	2 MB per chip	1067 MHz	39.7	35.9
Xeon 3070	2	2.67	4 MB per chip	1067 MHz	59.0	46.7

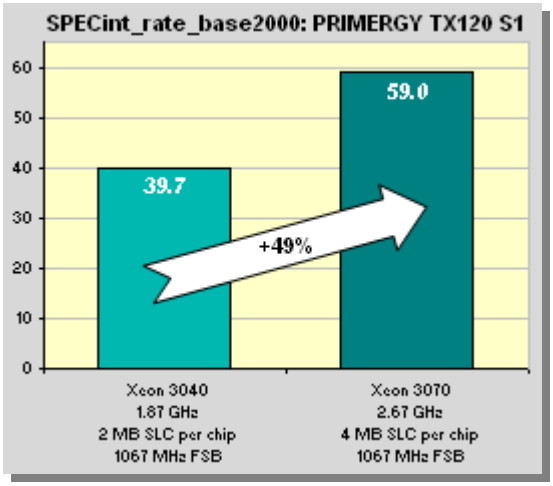

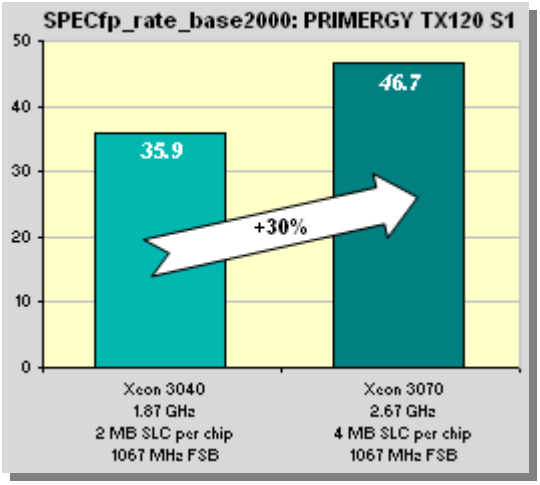

Benchmark environment

All SPECcpu2000 measurements were performed on a PRIMERGY TX120 S1 with the following hardware and software configuration:

Hardware	
Model	PRIMERGY TX120 S1
CPU	Xeon 3040 and 3070
Number of CPUs	1 chip, 2 cores, 2 cores per chip
Primary cache	32 kB instruction + 32 kB data on chip, per core
Secondary cache	Xeon 3040: 2 MB (I+D) on chip, per chip Xeon 3070: 4 MB (I+D) on chip, per chip
Memory	8 GB PC2-4200 DDR2-SDRAM
Software	
Operating system	SPECint2000: Windows Server 2003 Enterprise Edition SP1 (32-bit) SPECfp2000: SUSE Linux Enterprise Server 10 (64-bit)
Compiler	Intel C++/Fortran Compiler 9.1

SPECcpu2006*

Benchmark description

SPECcpu2006 is a benchmark to measure system efficiency during integer and floating point operations. It consists of an integer test suite containing 12 applications and a floating point test suite containing 17 applications which are extremely computing-intensive and concentrate on the CPU and memory. Other components, such as disk I/O and network, are not measured by this benchmark.

SPECcpu2006 is not bound to a specific operating system. The benchmark is available as source code and is compiled before the actual benchmark. Therefore, the compiler version used and its optimization settings have an influence on the measurement result.

SPECcpu2006 contains two different methods of performance measurement: The first method (SPECint2006 and SPECfp2006) determines the time required to complete a single task. The second method (SPECint_rate2006 and SPECfp_rate2006) determines the throughput, i.e. how many tasks can be completed in parallel. Both methods are additionally subdivided into two measuring runs, "base" and "peak", which differ in the way the compiler optimization is used. The "base" values are always used when results are published, the "peak" values are optional.

Benchmark	Arithmetic	Type	Compiler optimization	Measuring result	Application
SPECint2006	integer	peak	aggressive	speed	single threaded
SPECint_base2006	integer	base	conservative		
SPECint_rate2006	integer	peak	aggressive	throughput	multithreaded
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	multithreaded
SPECfp_rate_base2006	floating point	base	conservative		

The results represent the geometric mean of normalized ratios determined for the individual benchmarks. Compared with the arithmetic mean, the geometric mean results in the event of differingly high single results in a weighting in favor of the lower single results. “Normalized“ means measuring how fast the test system runs in comparison to a reference system. The value of “1“ was determined for the SPECint_base2006, SPECint_rate_base2006, SPECfp_base2006 and SPECfp_rate_base2006 results of the reference system. Thus a SPECint_base2006 value of 2 means for example that the measuring system has executed this benchmark approximately twice as fast as the reference system. A SPECfp_rate_base2006 value of 4 means that the measuring system has executed this benchmark about 4/[# base copies] times as fast as the reference system. “# base copies“ here specifies how many parallel instances of the benchmark have been executed.

We do not submit all SPECcpu2006 measurements for publication at SPEC. So not all results appear on SPEC’s web sites. As we archive the log data for all measurements, we are able to prove the correct realization of the measurements any time.

Benchmark results

Measurements were taken with the processor Celeron M 440 and the processors Xeon 3040 and 3070. The benchmark programs were compiled with the Intel C++/Fortran compiler 9.1 and run under SUSE Linux Enterprise Server 10 (64-bit).

All result numbers are published at http://www.spec.org.

* SPEC®, SPECint®, SPECfp® and the SPEC logo are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).

Processor	Cores/Chip	GHz	SLC	FSB	SPECint_base2006	SPECint2006
Celeron M 440	1	2	½ MB	800 MHz	11.1	11.7
Xeon 3040	2	1.87	2 MB per chip	1067 MHz	n/a	n/a
Xeon 3070	2	2.67	4 MB per chip	1067 MHz	n/a	n/a

Processor	Cores/Chip	GHz	SLC	FSB	SPECint_rate_base2006	SPECint_rate2006
Celeron M 440	1	2	½ MB	800 MHz	n/a	n/a
Xeon 3040	2	1.87	2 MB per chip	1067 MHz	21.6	22.6
Xeon 3070	2	2.67	4 MB per chip	1067 MHz	29.3	30.7

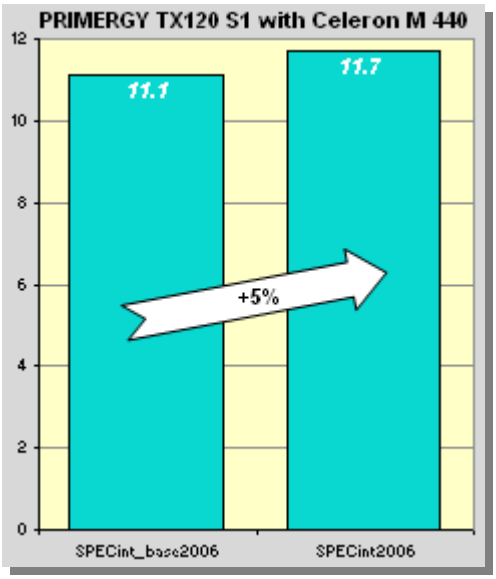

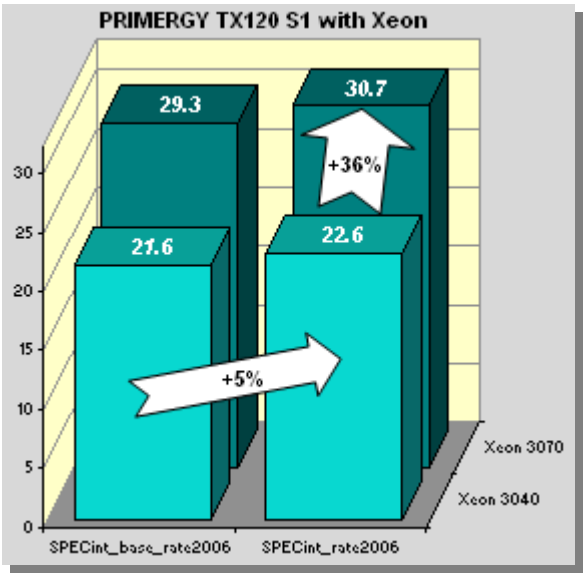

Processor	Cores/Chip	GHz	SLC	FSB	SPECfp_base2006	SPECfp2006
Celeron M 440	1	2	½ MB	800 MHz	11.4	11.5
Xeon 3040	2	1.87	2 MB per chip	1067 MHz	n/a	n/a
Xeon 3070	2	2.67	4 MB per chip	1067 MHz	n/a	n/a

Processor	Cores/Chip	GHz	SLC	FSB	SPECfp_rate_base2006	SPECfp_rate2006
Celeron M 440	1	2	½ MB	800 MHz	n/a	n/a
Xeon 3040	2	1.87	2 MB per chip	1067 MHz	19.8	20.4
Xeon 3070	2	2.67	4 MB per chip	1067 MHz	24.8	25.6

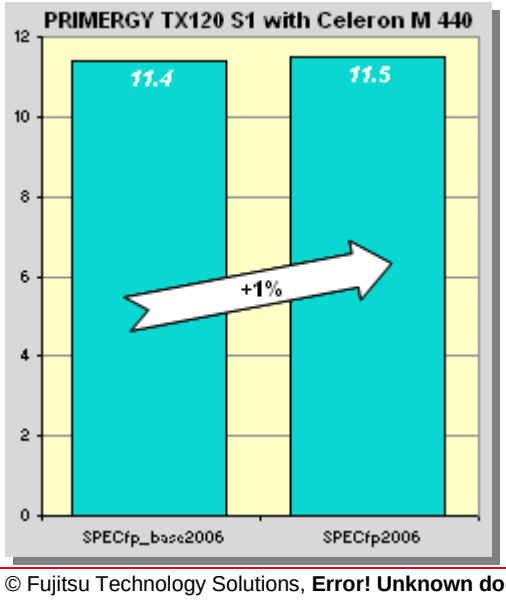

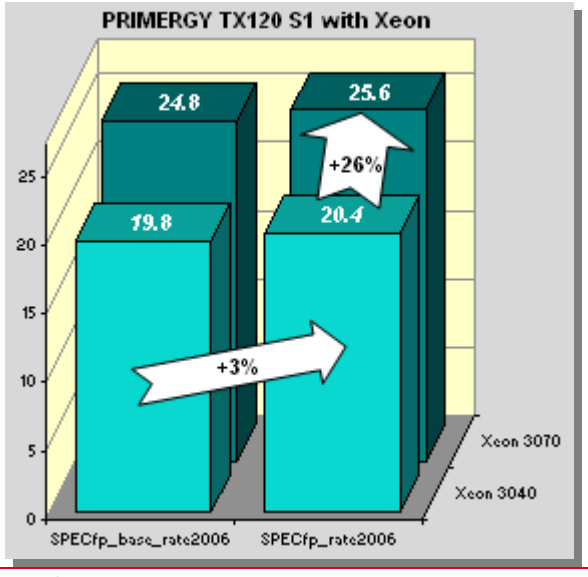

The measured SPECint_rate_2006 results are generally 4-5% above the SPECint_rate_base2006 results. The measured SPECfp_rate_2006 results are generally 3% above the SPECfp_rate_base2006 results.

Benchmark environment

All SPECcpu2006 measurements were performed on a PRIMERGY TX120 S1 with the following hardware and software configuration:

Hardware	
Model	PRIMERGY TX120 S1
CPU	Celeron M 440, Xeon 3040 and 3070
Number of CPUs	1
Primary Cache	32 kB instruction + 32 kB data on chip, per core
Secondary Cache	Celeron M 440: ½ MB (I+D) on chip, per chip Xeon 3040: 2 MB (I+D) on chip, per chip Xeon 3070: 4 MB (I+D) on chip, per chip
Memory	8 GB PC2-4200 DDR2-SDRAM
Software	
Operating System	SUSE Linux Enterprise Server 10 (64-bit)
Compiler	Intel C++/Fortran Compiler 9.1

SPECjbb2005*

Benchmark description

SPECjbb2005 is a Java business benchmark that focuses on the performance of Java server platforms. It is essentially a modernized version of SPECjbb2000 with the main differences being:

- 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 measures the implementations of the JVM, JIT (Just-In-Time) compiler, garbage collection, threads and some aspects of the operating system. 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 emphasis on the middle tier system:

- Clients generate the load, consisting of driver threads, which on the basis of the 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, thus avoiding large-scale hardware installations and making direct comparisons possible between SPECjbb2005 results of different systems. 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 it is necessary to state both metrics.

The following rules, according to which a compliant benchmark run has to be performed, are the basis for these 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*MaxWhm but not less than 8 warehouses. MaxWhm is the number of warehouses with the highest operation rate per second the benchmark expects. Per default the benchmark equates MaxWH with the number of CPUs visible by the operating system.
The metrics bops is the arithmetic average of all measured operation rates with between MaxWhm warehouses and 2*MaxWhm warehouses.

* SPEC®, SPECjbb® and the SPEC logo are registered trademarks of the Standard Performance Evaluation Corporation (SPEC).

Benchmark results

In December 2006 the PRIMERGY TX120 S1 was measured with the Xeon 3070 processor and a memory of 8 GB PC2-4200 DDR2-SDRAM. The measurement was taken under Windows Server 2003 Enterprise x64 Edition SP1. As JVM, one instance of JRockit(R) 5.0 P27.1.0 (build P27.1.0-7-71488-1.5.0_08-20061103-1228-windows-x86_64) by BEA was used.

The following result was achieved:

SPECjbb2005 bops	= 68534
SPECjbb2005 bops/JVM	= 68534

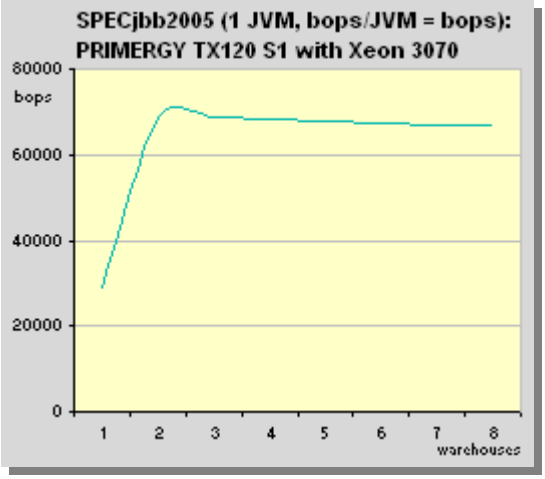

The benchmark result includes all measuring results of between 2 and 4 warehouses.

Benchmark environment

The SPECjbb2005 measurements were performed on a PRIMERGY TX120 S1 with the following hardware and software configuration:

Hardware	
Model	PRIMERGY TX120 S1
CPU	Xeon 3070
Number of chips	1 chip, 2 cores, 2 cores per chip
Primary Cache	32 kB instruction + 32 kB data on chip, per core
Secondary Cache	4 MB (I+D) on chip, per chip
Other Cache	None
Memory	4 x 2 GB PC2-4200 DDR2-SDRAM
Software	
Operating System	Windows Server 2003 Enterprise x64 Edition SP1
JVM Version	BEA JRockit(R) 5.0 P27.1.0 (build P27.1.0-7-71488-1.5.0_08-20061103-1228-windows-x86_64)

Literature

PRIMERGY Systems	http://ts.fujitsu.com/primergy
PRIMERGY Performance	http://ts.fujitsu.com/products/standard_servers/primergy_bov.html
SPECcpu2000	http://www.spec.org/osg/cpu2000
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

Contact

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

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

Published by department:

Enterprise Products
PRIMERGY Server
PRIMERGY Performance Lab
mailto:primergy.benchmark@ts.fujitsu.com

Internet:
http://ts.fujitsu.com/primergy

Extranet:
http://partners.ts.fujitsu.com/com/products/servers/primergy