



Using SPEC CPU2006 Benchmark Results to Compare the Compute Performance of Servers

White Paper

June 2010, Revision 1.0



Contents

- 1 Why Should SPEC CPU2006 Matter to You? 3**
 - Integer or Floating-Point Workloads: Which Matters to You?.....3
 - Base or Peak Metrics: Which Matters to You?4
 - Rate or Speed Metrics: Which Matters to You?4
- 2 Interpreting SPEC CPU2006 Scores and Comparing Systems 6**
 - How to Find and Fetch SPEC CPU2006 Scores to Compare6
 - What SPEC CPU2006 Is Not Intended to Measure6
- 3 Conclusion 7**

1 Why Should SPEC CPU2006 Matter to You?

The Standard Performance Evaluation Corporation (SPEC) CPU2006 benchmark is an industry-standard benchmark that is designed to compare the performance of servers. This benchmark can be run cross-platform, so it is possible to compare systems with different architectures running the same workload, such as Intel Xeon, AMD Opteron, Intel Itanium, IBM Power7, and SPARC.

SPEC CPU2006 results are widely available for hundreds of different processor and platform combinations. The benchmark is so common because it is a simplistic benchmark that does not require significant investment or time to run. The workloads do not require extremely large amounts of memory, and since there is essentially no disk or network I/O, little effort is required to build the infrastructure for the benchmark.

By looking at SPEC CPU2006 results, you can immediately tell that not all cores are equal. It becomes clear that number of cores and processor gigahertz speeds do not tell the entire server performance story. Some cores are faster than others, and a system with fewer cores may run faster than another with more cores. SPEC CPU2006 helps you compare the real performance of different platforms regardless of these differences.

Integer or Floating-Point Workloads: Which Matters to You?

The SPEC CPU2006 benchmark measures performance on two different basic workloads: one based on compute-intensive integer math, and one based on compute-intensive floating-point math. In general, the integer workload maps to the performance of business applications found in the data center, and the floating-point workload maps to scientific calculations more often found in an R&D high-performance computing (HPC) environment. Tables 1 and 2 list the SPEC CPU2006 integer and floating-point workloads.

Table 1. Integer Workloads for SPEC CPU2006

400.perlbench: PERL programming language	458.sjeng: Artificial intelligence (chess)
401.bzip2: Compression	462.libquantum: Quantum computing
403.gcc: C compiler	464.h264ref: Video compression
429.mcf: Combinatorial optimization	471.omnetpp: Discrete-event simulation
445.gobmk: Artificial intelligence (go)	473.astar: Path-finding algorithms
456.hmmcr: Search gene sequence	483.xalanbmk: XML processing

Table 2. Floating-Point Workloads for SPEC CPU2006

410.bwaves: Fluid dynamics	450.soplex - Linear Programming
416.gamess: Quantum chemistry	453.povray – Ray-tracing
433.Milc: Quantum chromodynamics	454.calculix – Structural Mechanics
434.zeusmp: Physics	459.GemsFDTD: Computational electromagnetics
435.gromacs: Biochemistry—molecular dynamics	465.tonto: Quantum chemistry
436.cactusADM: Physics—general relativity	470.lbm: Fluid dynamics
437.leslie3d: Fluid dynamics	481.wrf: Weather prediction
444.namd: Biology—molecular dynamics	482.sphinx3: Speech recognition
447.dealll: Finite-element analysis	

Base or Peak Metrics: Which Matters to You?

The SPEC organization describes the SPEC CPU2006 benchmark as “a compute-intensive suite of tests designed to stress a computer’s processor, memory architecture and compilers.” The benchmark provides two kinds of metrics: base scores and peak scores. The base scores are achieved by compiling the benchmark test code using standard compiler options and running the test. The peak scores are achieved by compiling the benchmark test code using super-tuned optimal compiler options and running the test.

When looking at the scores, consider whether you plan to recompile your application code using the exact same compiler used in the test. If you do not plan to recompile your code with this compiler, it is safer to look at the base scores achieved using the standard compiler options. If you plan to obtain the exact compiler used in the benchmark and recompile your application code, then peak scores are for you. Note that peak scores are optional; however, base scores must be published, according to SPEC rules.

Rate or Speed Metrics: Which Matters to You?

In addition to base and peak metrics, SPEC CPU2006 has rate and speed metrics.

The SPEC CPU2006 speed run is designed to measure the speed of a core. This benchmark runs the integer or floating-point workloads end to end in a serial fashion, calculating a score based on the amount of time needed to complete the test. This test is meant to represent a single-threaded application or an application that is not designed to run in a multicore system.

The rate run is designed to measure the throughput of a machine running simultaneous tasks over a certain amount of time. This benchmark runs several instances of the workloads at once and calculates a score based on how much work was done over a certain amount of time. This test represents a multithreaded application designed to run on a modern multicore system.

Table 3. The Eight SPEC CPU2006 Metrics

	Integer Workload	Floating Point Workload
Single-threaded test ("Speed") Designed to measure how long it takes to run the workloads end to end	SPECint_2006 "Peak" metric – using super-tuned compiler options SPECint_base2006 "Base" metric – using "standard" compiler options.	SPECfp_2006 "Peak" metric – using super-tuned compiler options SPECfp_base2006 "Base" metric – using "standard" compiler options.
Multi-threaded test ("Rate") Designed to measure the throughput of a machine running simultaneous tasks over a certain amount of time.	SPECint_rate2006 "Peak" metric – using super-tuned compiler options SPECint_rate_base2006 "Base" metric – using "standard" compiler options.	SPECfp_rate2006 "Peak" metric – using super-tuned compiler options SPECfp_rate_base2006 "Base" metric – using "standard" compiler options.

2 Interpreting SPEC CPU2006 Scores and Comparing Systems

Considering all the permutations, a total of eight possible metrics are used in the SPEC CPU2006 benchmark as summarized in Table 1.

For multithreaded business applications, you should compare SPECint_rate_base2006 scores to measure the difference between platforms. For single-threaded business applications, compare SPECint_base2006 scores.

For multithreaded scientific or HPC applications for which you have source code and that you plan to recompile, compare SPECfp_rate2006 scores. If you do not want to recompile, use SPECfp_rate_base2006. For single-threaded scientific or HPC applications, compare SPECfp_2006 scores if you plan to recompile, and SPECfp_base2006 if you do not.

How to Find and Fetch SPEC CPU2006 Scores to Compare

SPEC CPU2006 scores can be found at the www.spec.org website.

1. From the Results tab at the top of the page, go to CPU2006.
2. Click Search CPU2006 Results. This will take you to the SPEC/OSG Result Search Engine site.
3. Scroll down to Advanced Requests and click Configurable Query Form.
4. From here, you can search the results database based on criteria that you select, such as the name of the processor and number of sockets, system name, or hardware vendor. Enter the criteria that you seek, specify a primary sorting key (typically Hardware Vendor and Ascending, or Baseline and Descending).
5. Click Fetch Results.

What SPEC CPU2006 Is Not Intended to Measure

The SPEC CPU2006 benchmark is designed to measure the performance of a system's processors, memory, chip set and compiler. It will not find the appropriate memory size for a system or indicate the I/O throughput of a system

3 Conclusion

As a multiplatform benchmark, SPEC CPU2006 makes it clear that not all cores are equal and that gigahertz speeds do not consistently translate to system performance. Because the benchmark does not require significant resources to build and run, results can be found for most processor speeds on most platforms. Using the search engine at www.spec.org, you can easily reference these results and compare single-threaded and multithreaded integer and floating-point performance when comparing servers and sizing server workloads.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)