SPEC CPU 2017 & CHANGING PERFORMANCE

UNDERSTANDING SPEC CPU CHANGES & AMD'S NEW PERFORMANCE COMPILER & LIBRARIES

EXECUTIVE SUMMARY

Benchmarks help determine the relative performance of different platforms, neutralizing technical differences so IT can better ascertain the actual workload impact. The SPEC CPU 2006 benchmark has helped IT generate a relative performance baseline for CPUs across different platforms for over ten years. But being more than a decade old means the benchmark no longer provides IT with the most effective comparisons, as platform dynamics have changed greatly. SPEC's new SPEC CPU 2017 benchmark is designed to reflect the changing times.

Compilers and optimizations also play a major role in benchmark performance. AMD has adopted a new performance compiler, AMD Optimizing C/C++ Compiler (AOCC), which leverages the open source LLVM compiler. AMD has moved to an open source library strategy to help provide better alignment with today's server workloads and measurement tools.

Moor Insights & Strategy (MI&S) believes that the new SPEC CPU 2017 benchmark, along with AMD's move to the AOCC performance compiler and open source library strategy, will provide server customers with better tools for gauging and optimizing the CPU performance of a platform, especially with advanced workloads.

TODAY'S CUSTOMER NEEDS

With billions of dollars of IT purchases on the line each year, organizations need a fair and unbiased way to cut through the many vendor performance claims to determine the best servers for their specific applications. Many server purchases today are made based on general-purpose needs, requiring standardization on a single platform that must do well in a variety of environments. CPU-specific benchmarks can be one of the best ways for companies to ascertain how servers will perform.

The Standard Performance Evaluation Corporation (SPEC) is one of the most prominent benchmark providers and many server customers and vendors use SPEC as their primary measure for CPU performance. As an industry consortium, SPEC represents multiple companies with multiple viewpoints, bringing one of the more balanced views of performance with a better "apples to apples" comparison. As a consortium, the participants all have a stake in keeping the playing field level.

SPEC CPU 2006 is a good benchmark for comparing CPU / platform differences at the base level. Many businesses use SPEC CPU as a starting point to ascertain a minimum level of base performance, enabling them to narrow the decision. After this performance level has been determined, evaluators will often move on to more workload-related benchmarks. SPEC CPU provides a good baseline comparison to "start the conversation", creating a minimum point of entry for evaluation.

SPEC has recently announced that the SPEC CPU 2006 benchmark will be retired and replaced with SPEC CPU 2017, which will help in analyzing platforms against today's workloads. This replacement will happen against a backdrop of rapidly changing applications and environments, along with the new server platforms destined for the market in 2017—all resulting in potential for confusion.

How SPEC CPU is Changing in 2017

For any benchmark, one of the greatest challenges is the applicability / correlation to the underlying workload tasks. Some benchmarks, like TPC (for database) or SPECjbb 2015 (for Java applications) take on a very specific role in the evaluation process. However, as arbiter of a more generic and wide-ranging performance analysis, SPEC CPU faces a larger challenge. Many businesses will standardize on a few models of servers to help streamline their procurement process; specific benchmarks become less relevant, and a more CPU-centric benchmark rises to the top.

By comprehending both performance and variances across a wide range of platforms, CPU benchmarks can be representative (though not 100% definitive) of actual performance. The SPEC CPU 2006 benchmark suite uses a variety of real world applications, combining those scores into a result. It was released in 2006 when Gartner and IDC referred to virtualization as a "hot trend" with an installed base creeping towards 40%—meaning that most servers being deployed at that time were still standalone, not virtualized. Cloud technology was not even a consideration, with containers and microservices even further away from contemplation.

The benchmark's two main elements are SPECspeed, which measures single tasks, and SPECrate, which measures multiple simultaneous tasks. While the speed tests of 2006 may still accurately reflect single-threaded performance today, the rate tests of 2006 are less optimized for the environments of both today and the future.

Changes in SPEC CPU 2017 include updates and improvements that better reflect the size and complexity of modern workloads. The addition of OpenMP to the SPECspeed metrics (an API for parallel processing popular in the HPC community) enables better

performance measurement of today's highly parallelized systems that support more cores and threads. An additional, experimental metric enables measurement of power consumption while running the benchmark, giving users insight into the relationship between performance and power. The benchmark still has two primary sub-components: SPECint 2017 for measuring integer performance and SPECfp 2017 for floating point. The workloads included in these two components are based on real-world problems that businesses face today and are portable across a wide range of platforms and operating environments.

After the SPEC CPU 2017 benchmark is released (June 2017), there will be a 6-month window where platform results can still be published under the old benchmark. However, after this window closes, vendors will only be able to publish through SPEC with the 2017 version. For this reason, MI&S believes the most applicable performance measurements for today's modern systems will come from the new 2017 benchmark, due to its future longevity and better applicability to modern workloads. After the older benchmark is retired, there will be more clarity in the market, but in the near term, there will be potential for confusion. It will be important for businesses to understand which version of the benchmark they are reviewing to make proper comparisons.

WORKLOAD DIFFERENCES ARE RAISING BENCHMARK COMPLEXITY

Much has changed in the world of enterprise technology over the last decade, transforming the datacenter and the applications that run there. There is far more reliance on threading and less reliance on single-threaded performance in the datacenter as applications become more concurrent. With the advent of cloud technology, hybrid cloud, and hybrid IT, the next few years will see even more dramatic IT changes as well. This fluid enterprise environment not only makes older benchmarks less relevant, it also makes comparisons to previous generations of systems difficult, leaving IT unclear about the exact uplift from one generation to the next.

FIGURE 1: CPU REQUIREMENTS SHIFT OVER TIME

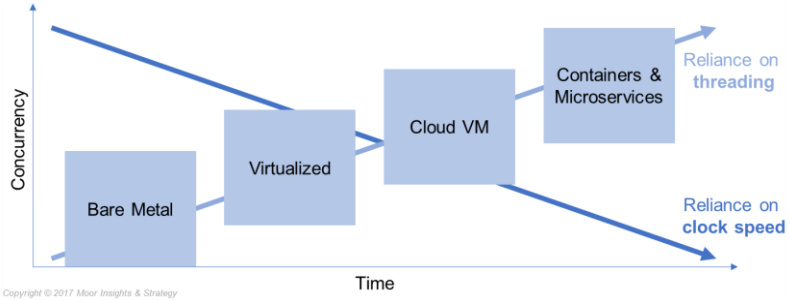

Source: Moor Insights & Strategy

As workloads changed, CPUs followed the trend. In 2006, CPUs were very different than they are today. There were no quad-core CPUs, much slower memory channels, significantly slower CPU-CPU communication paths, and immature SMT that shared too many resources, preventing true scaling.

TABLE 1: CPU COMPARISONS

	2006	2017	Increase
Cores	2	32	1,600%
Threads	2	64	3,200%
Memory Speed	400MHz	2666MHz	566%
Memory Bandwidth*	6,400MB/s	170,400MB/s	2,563%

**Assumes 2 memory channels for 2006 CPUs and 8 memory channels for 2017 CPUs*

Source: AMD

COMPILER DIFFERENCES CAN IMPACT BENCHMARK RESULTS

Compilers are tools that enable a developer to create (compile) an application from lines of code. The compiler can comprehend differences in the underlying platform, allowing some features to be exposed to the application while hiding others that may not be supported. There are a number of compilers in the market today, both open source and proprietary. Some of the leading compilers are Clang (LLVM project), the GNU Compiler Collection (GCC), Intel C++ Compiler (ICC), and Visual C++ (Microsoft). The choice of the compiler can have a significant impact on the overall performance of the compiled benchmark. An un-optimized benchmark or application may underperform, as certain CPU-supported instructions may not be flagged for use.

AMD'S COMPILER & LIBRARY STRATEGY

AMD has embraced an open source strategy. It has focused its optimization work on the latest benchmark versions in support of its new generation of platforms developed around the Zen core. This open source strategy is important for AMD, because it both differentiates from the more proprietary approach of Intel as well as lowers the barriers to acceptance by ISVs and other software developers. With open source, there is also an opportunity for feedback and collaboration with the wider developer community. Historically, AMD leveraged Open64 for its performance compiler, but with the advent of SPEC CPU 2017, AMD will be introducing the AOCC compiler. AMD's new performance compiler leverages the open source LLVM compiler, which is used by the high-performance computing market. There is an active LLVM community, and AMD believes this change in alliances will provide customers with a better-optimized compiler solution for AMD server products going forward. AMD will work closely with the open

source community and will provide performance optimization contributions to the LLVM community. AMD will also maintain the separate AOCC branded compiler, which will be available for any customer to use as well.

In another key strategy shift, AMD will retire its AMD Core Math Library (ACML) in favor of the open source BLIS, libFlame, and FFTW libraries. AMD is actively engaged in these communities and has made significant contributions. For functions that are not widely supported in open source, such as Random Number Generation (RNG and PRNG), and Core Math (LibM), AMD will maintain proprietary libraries and provide AMD-optimized binaries to the developer community. AMD's goal is to develop a core set of horizontal libraries that can be leveraged across multiple environments, especially within HPC applications. These libraries will enable software developers to work faster, delivering better performance through more standardized and competitive libraries that are optimized for AMD platforms and work well on all platforms. By using an open source approach, developers have a strong base to start from, requiring less work than starting from scratch.

CALL TO ACTION

Benchmarks can be a convenient and accurate way to discern performance differences between disparate platforms. Any comparison should focus on the newest benchmark version, as that will most likely have the best comparison relative to the current and future workloads that an enterprise would be running. We believe the SPEC CPU benchmark has always provided a healthy comparison at a base level between platforms and can be a proxy for performance between platforms with respect to their workloads. The SPEC CPU 2017 benchmark will be the best proxy for current and future performance, more accurately reflecting the most modern workloads.

MI&S recommends that businesses use this as a guide in making system decisions, although we also recognize that performance is just one of the many criteria that a business will consider in enterprise purchase decisions. AMD's focus on the compiler and library should give customers and developers important tools that can assist the complex task of gauging and optimizing CPU performance. The fact that AMD's approach is tied to open source and open standards should make it welcome in a market where proprietary solutions hold less value.

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHOR
John Fruehe, Senior Analyst at Moor Insights & Strategy

PUBLISHER
Patrick Moorhead, Founder, President, & Principal Analyst at Moor Insights & Strategy

EDITOR / DESIGN
Scott McCutcheon, Director of Research at Moor Insights & Strategy

INQUIRIES
Contact us if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS
This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

DISCLOSURES
This paper was commissioned by Advanced Micro Devices. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER
The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.