

Get More out of Your CPU-Intensive Workloads with AWS EC2 M7i Instances



SPECrate[®]2017



**Get Insights in Less Time
by Completing Complex
Scientific Models Faster**



**Improve the
Responsiveness of Your
Artificial Intelligence**

Speed Up AI, Scientific Modeling, and More by Choosing AWS EC2 M7i Instances Enabled by 4th Gen Intel[®] Xeon[®] Scalable Processors

Big problems need big performance from cloud instances to tackle them. If you're in the business of executing complex scientific models, running AI, or compressing video—anything that requires top processing power—choosing the right instance type can make or break your time to insight or workload completion. To see what level of performance you might expect from new AWS M7i instances compared to previous-gen M6i instances, we tested both using two CPU-intensive benchmarks. SPECrate[®]2017 Floating Point is an industry-standard benchmark using floating point workloads that include explosion modeling, fluid dynamics, molecular dynamics, weather forecasting, computational electromagnetics, and regional ocean modeling. SPECrate[®]2017 Integer executes a variety of workloads, including artificial intelligence, general data compression, video compression, discrete event simulation, route planning, and more.¹

Across both benchmarks, new AWS M7i instances with 4th Gen Intel[®] Xeon[®] Scalable processors offered significant throughput improvements, which means that this instance type can help your team finish complex scientific models in less time and give your AI and analytics workloads a boost to deliver a more responsive experience.

Get Answers Faster via Quicker Scientific Modeling

First, we tested with SPECrate[®]2017 Floating Point suite. When an instance's estimated SPECrate[®]2017_fp_base performance score is higher, that means it handled more floating-point computations in a given amount of time. Across five instance sizes, M7i instances outperformed the previous-gen instances, achieving up to 1.43 times better performance (see Figure 1). This means that organizations can expect complex modeling and rendering tasks to finish in significantly less time, so professionals across disciplines can gain insights and act faster.

Normalized Est. SPECrate 2017_fp_base Scores

Score | Higher is better

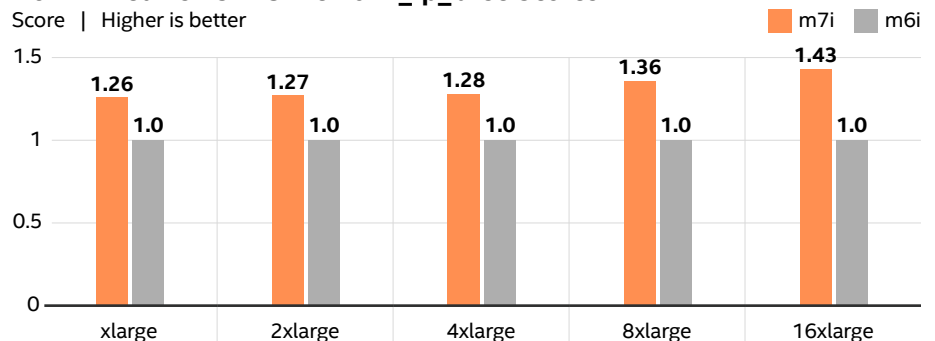


Figure 1. Relative estimated SPECrate[®]2017_int_base performance of M7i instances compared to M6i instances. Higher is better.

Speed up AI

Checking another type of CPU performance with SPECrate®2017 Integer bore similar results. Higher SPECrate®2017_int_base scores indicate that an instance can speed up workloads like artificial intelligence, video compression, and route planning. Again, all five sizes of the M7i instances outperformed the previous-gen instances, achieving up to 1.27 times the performance (see Figure 2). This means that M7i instances can offer a big boost to AI responsiveness and other big tasks such as discrete event simulation—giving you the answers you need faster than ever.

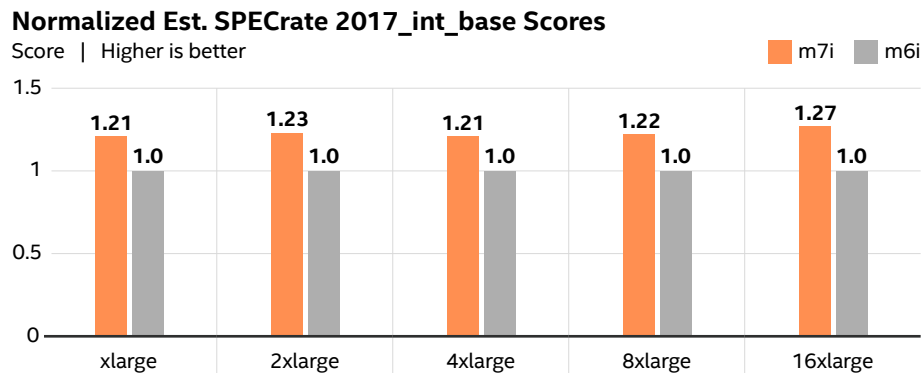


Figure 2. Relative estimated SPECrate®2017_int_base performance of M7i instances compared to M6i instances. Higher is better.

Conclusion

No matter the CPU-intensive workload that your organization runs, higher-performing cloud instances can improve productivity and business agility. Across instance sizes and targeting multiple types of CPU performance, AWS EC2 M7i instances improved workload completion speeds compared to instances built on previous-gen processors. Selecting M7i instances can help your team complete complex scientific rendering tasks and get answers in less time and improve the speed with which your AI models learn and respond—which ultimately helps your business mission. For these and other compute-intensive workloads, choose the superior performance of AWS EC2 M7i instances.

Learn More

To begin running your CPU-intensive workloads on AWS EC2 M7i standard instances featuring 4th Gen Intel® Xeon® Scalable processors, visit <https://aws.amazon.com/ec2/instance-types/m7i/>.

¹ Estimates of SPECrate®2017_int_base and SPECrate®2017_fp_base based on Intel internal measurements. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.

Tests by Intel Aug. 2023 on AWS. Workload: speccpu_2017_v119, Ubuntu 22.04.2 LTS, Binaries: https://af01p-igk.devtools.intel.com/artifactory/platform_hero-repos/hero_features_assets/SpecCpu-2017/FOR-INTEL-cpu2017-1.1.9-ic2023.0-linux-binaries-20221201.tar.xz Compiler: ic2023.0-lin-core-avx512-rate-20221201_intel.cfg, Libraries: cpu2017-1.1.9.iso. All M7i instances: Intel Xeon Platinum 8488C, 4GB RAM/vCPU, Amazon EC2 BIOS, microcode 0x2b000461, OS kernel 5.19.0.1029-aws, 1 CPU socket, 1 NUMA node. All M6i instances: Intel Xeon Platinum 8375C, 4GB RAM/vCPU, Amazon EC2 BIOS, microcode 0xd000390, OS Kernel 5.19.0-1026 AWS (-1025 AWS on 2xlarge), 1 CPU socket, 1 NUMA node.



Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See above for configuration details. No product or component can be absolutely secure. Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Printed in USA 1123/HM/PT/PDF US001

