


# Boost Floating-Point Performance with AWS m6i Instances Featuring 3<sup>rd</sup> Gen Intel<sup>®</sup> Xeon<sup>®</sup> Scalable Processors



SPECrate<sup>®</sup>2017



Achieve up to 18% more est. SPECrate<sup>®</sup>2017 **\_fp\_base** performance with 4-vCPU m6i instances

*vs. m6a instances*



Achieve up to 40% more est. SPECrate<sup>®</sup>2017 **\_fp\_base** performance with 8-vCPU m6i instances

*vs. m6a instances*



Achieve up to 58% more est. SPECrate<sup>®</sup>2017 **\_fp\_base** performance with 16-vCPU m6i instances

*vs. m6a instances*

## Achieve 58% Better Estimated SPECrate<sup>®</sup>2017\_**\_fp\_base** Performance than m6a Instances with 3<sup>rd</sup> Gen AMD EPYC<sup>™</sup> Processors

When an organization relies heavily on complex workloads involving floating-point calculations, it's imperative that those workloads get the performance they need to run smoothly. The cloud offers advantages in flexibility and manageability, but there are hundreds of possible cloud solutions—which is the right fit for such a workload?

Intel tested the estimated floating-point performance of two sets of instances on Amazon Web Services (AWS):

- m6i instances featuring 3<sup>rd</sup> Gen Intel Xeon Scalable processors
- m6a instances featuring 3<sup>rd</sup> Gen AMD EPYC processors

Intel used the SPECrate<sup>®</sup>2017 Floating Point suite of benchmarks for these measurements, with higher estimated SPECrate<sup>®</sup>2017\_**\_fp\_base** performance indicating that the instance could handle more floating point computations in the same amount of time. To highlight performance advantages as the instances scaled, Intel tested both sets of instances with 4 vCPUs, 8 vCPUs, and 16 vCPUs. Across all three sizes, the m6i instances provided greater estimated SPECrate<sup>®</sup>2017\_**\_fp\_base** performance, indicating that m6i instances may be a wiser choice for complex floating-point workloads.

## See Better Estimated SPECrate<sup>®</sup>2017\_**\_fp\_base** Performance on Small Instances

At the smallest size Intel tested, the m6i instances with 3<sup>rd</sup> Gen Intel Xeon Scalable processors delivered 18% more estimated SPECrate<sup>®</sup>2017\_**\_fp\_base** performance than the m6a instances with 3<sup>rd</sup> Gen AMD EPYC processors.

### Normalized Estimated SPECrate<sup>®</sup>2017\_**\_fp\_base** Scores on 4 vCPUs

Higher is better




Figure 1. Relative est. SPECrate<sup>®</sup>2017\_**\_fp\_base** performance of 4-vCPU m6i instances vs. 4-vCPU m6a instances. Higher numbers are better.

## See Better Estimated SPECrate®2017\_fp\_base Performance on Medium-Size Instances

Figure 2 highlights the increase in estimated SPECrate®2017\_fp\_base performance on 8-vCPU m6i instances—enabled by new 3<sup>rd</sup> Gen Intel® Xeon® Scalable processors—compared to equivalent m6a instances with AMD EPYC processors.




Figure 2. Relative est. SPECrate®2017\_fp\_base performance of 8-vCPU m6i instances vs. 8-vCPU m6a instances. Higher numbers are better.

## See Better Estimated SPECrate®2017\_fp\_base Performance on Larger Instances

As Figure 3 shows, the most significant performance boost for m6i instances came at the largest size instance, 16 vCPUs. In this test, the m6i instances featuring 3<sup>rd</sup> Gen Intel Xeon Scalable processors outperformed the m6a instances with 3<sup>rd</sup> Gen AMD EPYC processors by 58%.




Figure 3. Relative est. SPECrate®2017\_fp\_base performance of 16-vCPU m6i instances vs. 16-vCPU m6a instances. Higher numbers are better.

As the size of the instance increased, so did the difference in floating-point performance between m6i instances with new Intel processors and m6a instances with AMD processors. Particularly for complex workloads running on larger instances, organizations could benefit from selecting m6i instances enabled by 3<sup>rd</sup> Gen Intel Xeon Scalable processors.

### Learn More

To get started running your floating-point workloads on AWS m6i instances enabled by 3<sup>rd</sup> Gen Intel Xeon Scalable processors, go to <https://aws.amazon.com/ec2/instance-types/m6i/>.

Tests performed by Intel on Oct. 2021-Jan. 2022. All Intel configs: Intel Xeon Platinum 8375C CPU @ 2.90GHz, AWS us-east-2, EBS 512GB, up to 10Gbps network BW, up to 12.5Gbps storage BW, Ubuntu 20.04.3 LTS Kernel 5.11.0-1022-aws, cpu2017 v1.1.8, ICC 2021.4 revB\_8GBqkmallo, -xCORE-AVX512, ic2021.1-lin-core-avx512-rate-20201113\_revB.cfg. All AMD configs: AMD EPYC 7R13 Processor, up to 12.5Gbps network BW, up to 6.6Gbps storage BW, Ubuntu 20.04.3 LTS kernel 5.11.0-1022-aws, AOCC 3.0, -march=znver3, aocc3.0-lin-znver3-rate-published-20210317. cfg. xlarge VMs: 4 cores, 16GB RAM, 4 workload copies; 2xlarge VMs: 8 cores, 32GB RAM, 8 workload copies; 4xlarge VMs: 16 cores, 64GB RAM, 16 workload copies.



Performance varies by use, configuration and other factors. Learn more at [www.intel.com/PerformanceIndex](http://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 backup 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 0422/JO/PT/PDF US001

