



0

0



SIMCENTER[™] STAR-CCM+[™] ON AMAZON EC2 HPC7A INSTANCES COMPUTATIONAL FLUID DYNAMICS

Powered by 4th Gen AMD EPYC[™] Processors

October 2023

AT A GLANCE

Amazon Elastic Compute Cloud Hpc7a instances powered by 4th Gen AMD EPYC[™] 9004 processors deliver a superb generational uplift vs prior-generation Hpc6a instances running Simcenter[™] STAR-CCM+[™] and strong scale-out performance.¹

PERFORMANCE HIGHLIGHTS

A single 192-core Amazon EC2 Hpc7a instance delivers generational performance uplifts of ~2.72x to ~2.88x with a geomean uplift of ~2.79x compared to a single 96-core Amazon EC2 Hpc6a instance running a variety of Simcenter STAR-CCM+ benchmarks. Amazon EC2 Hpc7a instances also offer strong scaling uplifts to support realistic production-size models.



AMAZON EC2 HPC7A INSTANCES

According to Amazon, Amazon Elastic Compute Cloud (Amazon EC2) Hpc7a instances, powered by 4th Gen AMD EPYC processors, deliver up to 2.5x better performance across various HPC workloads compared to Amazon EC2 Hpc6a instances. Hpc7a instances feature 2x higher core density (up to 192 cores), 2.1x higher memory bandwidth throughput, 2x memory (768 GB), and 3x higher network bandwidth compared to Hpc6a instances. These instances offer 300 Gbps of Elastic Fabric Adapter (EFA) network bandwidth, powered by the AWS Nitro System, for fast and low-latency inter-node communications. Hpc7a instances feature Double Data Rate 5 (DDR5) memory, which provides 50% higher memory bandwidth compared to DDR4 memory to enable high-speed access to data in memory. These instances are ideal for compute-intensive, latency-sensitive HPC workloads, helping you scale more efficiently on fewer nodes compared to Hpc6a instances.²

Hpc7a instances are powered by 24-, 48-, 96-, or 192 4th Gen AMD EPYC processor cores. All Hpc7a instance sizes have identical shared resources (memory, storage, EFA network bandwidth, and network bandwidth). The variety of Hpc7a instance sizes helps you optimize the number of cores per instance to meet your specific workload needs. See <u>Amazon EC2 Hpc7a Instances</u>* for complete details.

IN THIS BRIEF

- Scaling Results.....Page 2
- AMD EPYC 9004 Series Processors.....Page 2
- Siemens[®] Simcenter STAR-CCM+Page 2

| • | System Configuration | Page 3 |
|---|----------------------|--------|
| • | Test Methodology | Page 3 |



SCALING RESULTS

Figure 1 shows excellent scalability running Simcenter STAR-CCM+ from one to 16 Amazon EC2 hpc7a.96xlarge instances. A production-sized model such as LeMans Segregated 100M delivers a super-linear scale-out speedup of ~9.41x at eight Amazon EC2 Hpc7a instances with 1536 total cores and a speedup of ~19.50x at 16 nodes with 3072 total cores, both relative to a single Hpc7a instance with 192 cores. Similarly, a larger model such as VTM Benchmark 178M shows a super-linear speedup of ~8.17x at eight Amazon EC2 Hpc7a instance nodes with 1536 total cores and a speedup of ~15.54x at 16 Amazon EC2 Hpc7a instances relative to a single Amazon EC2 Hpc7a instance with 192 cores.³



Figure 1: Simcenter STAR-CCM+ scaling from one to eight Amazon EC2 hpc7a.96xlarge instances

AMD EPYC 9004 SERIES PROCESSORS

AMD EPYC 9004 Series Processors continue to redefine the standards for modern datacenters. 4th Gen AMD EPYC processors are built on the innovative x86 architecture and "Zen 4" core. 4th Gen AMD EPYC processors deliver efficient, optimized performance by combining higr frequencies, the largest-available L3 cache, 128 lanes of PCIe[®] 5 I/O, and synchronized fabric and memory clock speeds, plus support for up to 6 TB of DDR5-4800 memory. Built-in security features, such as AMD Infinity Fabric[™] technology, Secure Memory Encryption (SME), and Secure Encrypted Virtualization (SEV-SNP) help protect data while it is in use.⁴

SIMCENTER[™] STAR-CCM+

Engineers need to predict the consequences of any design changes on the real-world performance of their product, for better or for worse. Historically those predictions came from hand calculations or from experimenting with physical prototypes. Modern engineering simulation applications offer comprehensive predictions that are usually more accurate and less expensive than experimental testing. Simcenter[™] STAR-CCM+[™] is a multiphysics Computational Fluid Dynamics (CFD) application that offers a single integrated CAD, automated meshing, multiphysics CFD, sophisticated postprocessing, and design exploration environment.⁵ The design exploration and optimization capabilities of Simcenter STAR-CCM+ allow engineers to efficiently explore the entire design space and make better design decisions faster by simulating products operating under real-world conditions.

Unlike calculations or prototypes, engineering simulation applications allow engineers to explore product performance over the full range of possible operating conditions instead of just a handful of carefully selected "design points." Per Siemens, the additional insight gained by using Simcenter STAR-CCM+ to guide your design process ultimately leads to higher quality and more robust products that better fulfill customer expectations.

SYSTEM CONFIGURATION

| | ΗΡር6Α INSTANCE | HPC7A INSTANCE |
|-------------------|-----------------------------|-----------------------------|
| Instance Type | hpc6a.48xlarge | hpc7a.96xlarge |
| # of Instances | 1 | 1, 2, 4, 8, and 16 |
| Testing Timeframe | July and August, 2022 | July and August, 2022 |
| Solver | Simcenter STAR-CCM+ 2020.3 | Simcenter STAR-CCM+ 2302 |
| MPI | OpenMPI 4.0.1 | IntelMPI 2021.9 |
| OS | Amazon Linux [®] 2 | Amazon Linux [®] 2 |

Table 1: Instance and software configurations

TEST METHODOLOGY

Simcenter STAR-CCM+ provides a standard set of benchmarks across different model sizes that represent typical usage. The benchmarks were run for the following models and sizes for both Hpc7a and Hpc6a instances: LeMans Poly 17M, EMP HydroCyclone 30M, LeMans Segregated 100M, and VTM Benchmark 178M. Simcenter STAR-CCM+ reports average elapsed time (sec) as the standard metric for evaluating the performance of a run. The benchmarks were run 3 times for each data point, and the average of the 3 elapsed times from each run was calculated while ensuring that the runtime variability between the individual runs was <3%. The single node uplift/speedup is calculated as the ratio of systems under test (Hpc7a instance type) to the reference system (Hpc6a instance type). Table 1 provides the instance and software configuration for the results reported in this Performance Brief.

FOR ADDITIONAL INFORMATION

Please see the following additional resources for more information about 4th Gen AMD EPYC features, architecture, and available models:

• <u>AMD EPYC[™] 9004 Series Processors</u>

AMD EPYC[™] Products

• <u>AMD EPYC[™] Tuning Guides</u>

REFERENCES

- Testing performed in July and August, 2023. Cloud performance results presented are based on the test date in the configuration and are in alignment with AMD internal bare-metal testing factoring in cloud service provider overhead. Results may vary due to changes to the underlying configuration, and other conditions such as the placement of the VM and its resources, optimizations by the cloud service provider, accessed cloud regions, co-tenants, and the types of other workloads exercised at the same time on the system.
- Information obtained from <u>https://aws.amazon.com/ec2/instance-types/hpc7a/</u>*.
- 3. AMD defines "linear scaling" as an equal and proportionate application performance uplift relative to single node performance; that is, when scaling out to 2 nodes results in 2x the performance of a single node, scaling out to 4 nodes results in 4x the performance of a single node, and so forth. "Super-linear" scaling is when the performance uplift achieved by adding one or more node(s) is greater than linear. AMD allows a +/- of 2% margin of error when claiming linear or super linear scaling. GD-205
- 4. AMD Infinity Guard features vary by EPYC[®] Processor generations. Infinity Guard security features must be enabled by server OEMs and/or Cloud Service Providers to operate. Check with your OEM or provider to confirm support of these features. Learn more about Infinity Guard at <u>https://www.amd.com/en/technologies/infinity-guard</u>. GD-183
- 5. Please see https://plm.sw.siemens.com/en-US/simcenter/fluids-thermal-simulation/star-ccm/*for detailed information about Simcenter STAR-CCM+.



AUTHORS

Shashank Kashyap, Sai Kovouri, and Sylvester Rajasekaran contributed to this Performance Brief.

RELATED LINKS

- <u>Amazon EC2 Hpc7a Instances</u>*
- <u>Elastic Fabric Adapter</u>*
- <u>Amazon FSx for Lustre</u>*
- <u>AWS ParallelCluster</u>*
- AWS Nitro System*
- <u>Siemens Software</u>*, <u>Simcenter</u>* products, and <u>Simcenter STAR-CCM+</u>*
- <u>AMD EPYC[™] Processors</u> and the <u>AMD Documentation Hub</u>

*Links to third party sites are provided for convenience and unless explicitly stated, AMD is not responsible for the contents of such linked sites and no endorsement is implied.

BOOST PERFORMANCE WITH 4TH GEN AMD EPYC

Amazon EC2 Hpc7a instances are powered by 4th Gen AMD EPYC processors and offer up to 192 CPU cores, 768 GB of high-bandwidth DDR5 memory, and 300 Gbps of Elastic Fabric Adapter (EFA) network bandwidth for fast, lowlatency inter-node communications. These instances offer efficient scaling and are ideal for computeintensive, latency-sensitive HPC workloads.

AMD EPYC 9004 FOR HPC

4th Gen AMD EPYC processors deliver blazing per-core performance thanks to fast CPU frequencies, lower latency memory, and a unified cache structure. AMD EPYC processors provide high bandwidth between nodes with support for PCIe[®] Gen 5 network devices and accelerators that greatly benefit HPC applications.

SIEMENS

Siemens focuses on industry, infrastructure, transport, and healthcare. Siemens creates technologies for more resourceefficient factories, resilient supply chains, and smarter buildings and grids, cleaner and more comfortable transportation, and advanced healthcare to add real value for customers.

SIMCENTER STAR-CCM+

Simcenter STAR-CCM+ is a leading computational fluid dynamics (CFD) software application for fast, accurate simulation of almost any engineering problem that involves the fluids, structures, and all the associated physics. The single integrated environment includes CAD, automated meshing, multiphysics CFD, sophisticated post-processing, and design exploration.

DISCLAIMERS

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

COPYRIGHT NOTICE

©2023 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, Infinity Fabric, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Siemens, the Siemens logo, Simcenter and Simcenter StarCCM+ are trademarks or registered trademarks of Siemens Industry Software Inc., or its subsidiaries or affiliates, in the United States and in other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.