

# ANSYS® FLUENT® ON AMAZON EC2 HPC7A INSTANCES COMPUTATIONAL FLUID DYNAMICS

Powered by 4th Gen AMD EPYC™ Processors

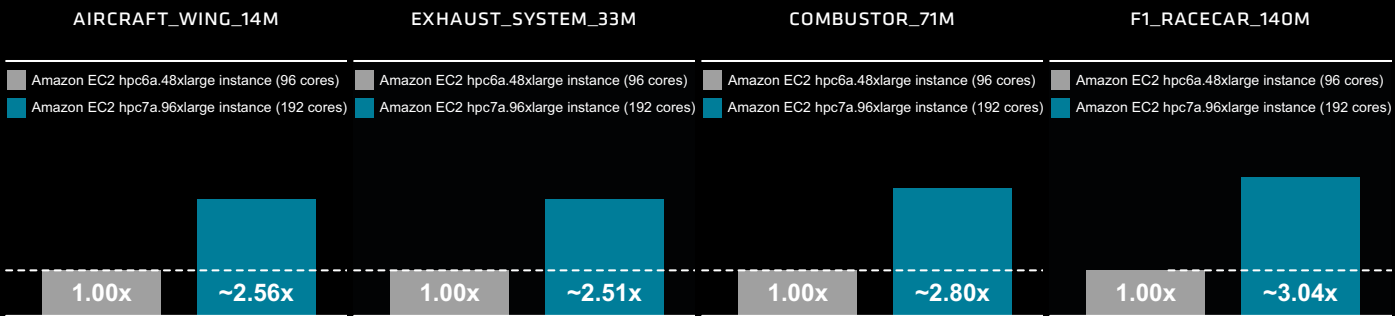
October 2023

## AT A GLANCE

Amazon Elastic Compute Cloud Hpc7a instances powered by 4th Gen AMD EPYC™ 9004 processors deliver an exceptional generational uplift vs prior-generation Hpc6a instances running Ansys® Fluent® along with outstanding scale-out performance.<sup>1</sup>

## PERFORMANCE HIGHLIGHTS

A single 192-core Amazon EC2 Hpc7a instance delivers generational performance uplifts of ~2.50x to ~3.04x with a geomean uplift of ~2.72x compared to a single 96-core Amazon EC2 Hpc6a instance running a variety of Ansys Fluent 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.<sup>2</sup>

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 [Amazon EC2 Hpc7a Instances\\*](#) for complete details.

### IN THIS BRIEF

- Scaling Results ..... [Page 2](#)
- AMD EPYC 9004 Series Processors ..... [Page 2](#)
- Ansys Fluent ..... [Page 2](#)
- System Configuration ..... [Page 3](#)
- Test Methodology ..... [Page 3](#)

## SCALING RESULTS

Figures 1 and 2 below showcase the excellent scalability of Amazon EC2 hpc7a.96xlarge instances for large production-sized models, such as combustor\_71m and f1\_racecar\_140m. A production-sized model such as f1\_racecar\_140m shows a speedup of ~14.96x at 16 instances (3072 cores) relative to a single node with 192 cores. The open\_racecar\_280m model scales in multiples of two nodes because this model cannot run on a single node due to memory constraints. Here, hpc7a.96xlarge instances scale at ~7.95x at 16 instances (3072 cores) relative to two nodes with 384 cores. Linear scaling is defined as eight nodes in this case because of the aforementioned memory constraints.<sup>3</sup>

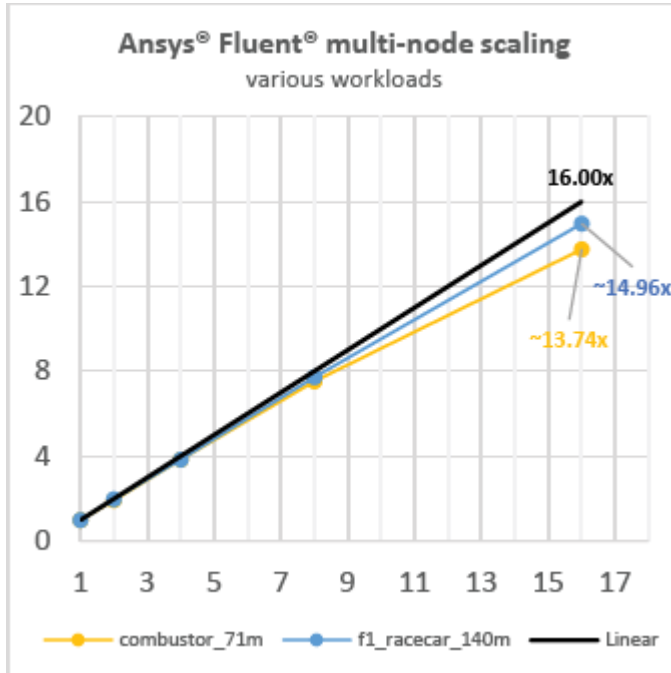


Figure 1: hpc7a.96xlarge scaling performance at 1, 2, 4, 8, and 16 nodes (combustor\_71m and f1\_racecar\_140m benchmarks)

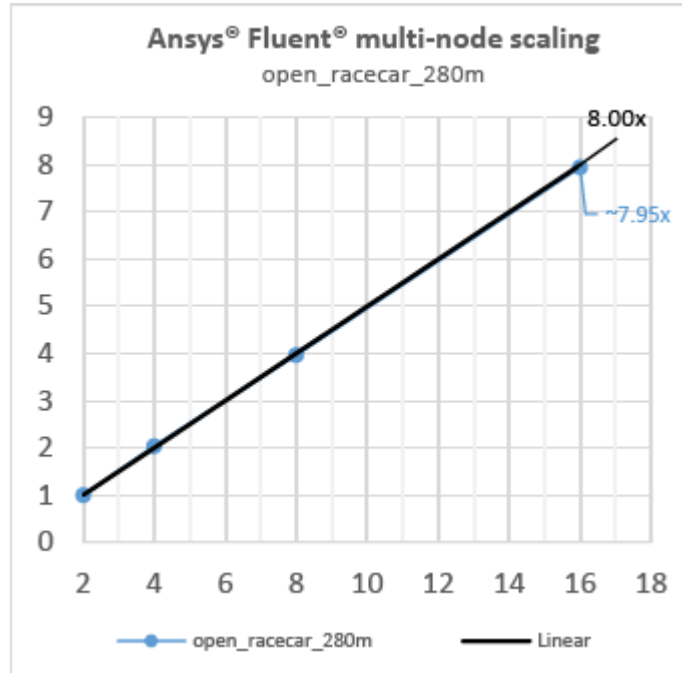


Figure 2: hpc7a.96xlarge scaling performance at 2, 4, 8, and 16 nodes (open\_racecar\_280m benchmark)

## 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 high 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.<sup>4</sup>

## ANSYS FLUENT

Ansys Fluent is a general-purpose computational fluid dynamics (CFD) and multi-physics tool that helps empower you to go further and faster as you optimize your product’s performance. Fluent contains the broad physical modeling capabilities needed to model flow, turbulence, heat transfer, and reactions for industrial applications—ranging from air flow over an aircraft wing to combustion in a furnace, from bubble columns to oil platforms, from blood flow to semiconductor manufacturing, and from clean room design to wastewater treatment plants. Fluent covers a broad reach, including special models with capabilities to model in-cylinder combustion, aero-acoustics, turbo machinery, and multiphase systems.<sup>5</sup>

## SYSTEM CONFIGURATION

	HPC6A INSTANCE <sup>1</sup>	HPC7A INSTANCE <sup>1</sup>
Instance Type	hpc6a.48xlarge	hpc7a.96xlarge
# of Instances	1	1, 2, 4, 8, and 16
Testing Timeframe	July and August, 2023	July and August, 2023
Solver	Ansys Fluent 2022 R1	Ansys Fluent 2023 R1
MPI	IntelMPI 2021.3	IntelMPI 2021.9
OS	Amazon Linux <sup>®</sup> 2	Amazon Linux <sup>®</sup> 2

Table 1: Instance and software configurations

## TEST METHODOLOGY

Ansys provides a standard set of benchmarks across different models that represent typical usage and cover a range of different model sizes. The benchmarks were run for the following models and sizes for both Hpc7a and Hpc6a instances: External Flow over Aircraft Wing 14M (aircraft\_wing\_14m), Vehicle Exhaust Model 33M (exhaust\_system\_33m), Flow through Combustor 71M (combustor\_71m), External Flow over F1 Racecar 140M (f1\_racecar\_280M), and External Flow over Open Racecar 280M (open\_racecar\_280m). Fluent reports Solver Rating as the standard metric for evaluating the performance of a run. Each benchmark was run three times for each data point and the average of the 3 Solver Ratings 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 the system under test (Hpc7a instance type) to the reference system (Hpc6a instance type). Tables 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:

- [AMD EPYC™ 9004 Series Processors](#)
- [AMD EPYC™ Products](#)
- [AMD EPYC™ Tuning Guides](#)

## REFERENCES

1. 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.
2. Information obtained from <https://aws.amazon.com/ec2/instance-types/hpc7a/>.\*
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 <https://www.amd.com/en/technologies/infinity-guard>. GD-183
5. Please see <https://www.Ansys.com/products/fluids/Ansys-fluent>\* for detailed information about Ansys Fluent.

## AUTHORS

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

## RELATED LINKS

- [Amazon EC2 Hpc7a Instances\\*](#)
- [Elastic Fabric Adapter\\*](#)
- [Amazon FSx for Lustre\\*](#)
- [AWS ParallelCluster\\*](#)
- [AWS Nitro System\\*](#)
- [Ansys\\*](#), [Ansys and AMD\\*](#), and [Ansys HPC\\*](#). Please contact Ansys [here\\*](#)
- [AMD EPYC™ Processors](#) and the [AMD Documentation Hub](#)

*\*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, low-latency inter-node communications. These instances offer efficient scaling and are ideal for compute-intensive, latency-sensitive HPC workloads.

### AMD EPYC 9004 FOR HPC

4th Gen AMD EPYC processors deliver blazing per-core performance thanks to fast CPU frequencies, low 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.

### ANSYS

Ansys offers a broad portfolio of engineering simulation software that helps customer solve complex design challenges, rapidly innovate and easily validate design ideas, and predict the performance of future products.

### ANSYS FLUENT

Collaboration between AMD and Ansys offers high performance and scalability for Computational Fluid Dynamics (CFD) workloads. Customers across many industries can benefit from the technical partnership between AMD and Ansys.

## 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. Ansys, CFX, Fluent, LS-DYNA, Mechanical, RedHawk, and any and all Ansys, Inc. brand, product, service and feature names, logos and slogans are registered trademarks or trademarks of Ansys, Inc. or its subsidiaries in the United States or other countries under license. 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.