



Leverage HPC on AWS, powered by NVIDIA GPUs, to solve your most complex workload challenges



High Performance Computing (HPC) continues to fuel advances in various industries, from healthcare with drug discovery, to energy with more efficient operations, to transportation with autonomous driving.

Running HPC on premises has historically been the primary way of processing complex workloads. However, nowadays, organizations recognize the advantages of the cloud over finite on-premise systems. Due to its bursting capabilities, many HPC workloads can benefit from the cloud's scalable computing capacity.

Additionally, investing significant resources in HPC systems on premises can affect future organizational agility, as typical hardware refresh cycles lock teams into aging systems with reduced performance.

Meanwhile, instant access to the latest technologies are changing the market perception of HPC in the cloud. A recent study¹ showed that almost 50% of HPC users are altering on-premises deployments due to the cloud.

It doesn't come as a surprise that HPC cloud computing is projected to grow by 17.6% per year² over the next five years compared to only 7-8% growth for on-premises HPC servers.

To help organizations solve complex computational problems faster - AWS brings a portfolio of HPC services supported by highly performant NVIDIA GPUs. Over the past decade, NVIDIA accelerated computing has provided 1,000x performance increases³, thanks to advances in GPU architectures and the full stack of software built on top of them.

Running HPC workloads on AWS means organizations can access the agility and flexibility of AWS for on-demand GPU-accelerated infrastructure to address their most demanding workloads while keeping their costs in check.

Advancing complex workloads with AWS and NVIDIA

Performance acceleration delivered by AWS in collaboration with NVIDIA can benefit workloads across industries.



Energy and Weather

Pursue renewable energy sources, build smarter, more resilient grid operations, expedite energy exploration and production, and help ensure safe conditions for workers and communities.



Financial Services

Satisfy fast-changing client expectations, reduce time to insight, discover strategies, spotlight fraud, remain secure, and drive business growth.



Life Sciences and Healthcare

Accelerate structure-based drug design. Build and run data applications that are retrospective, real-time or predictive to accelerate the journey from genomic data to genomic insight.



Design and Engineering

Model, develop, and run product simulations sooner. Execute simulations at scale and in parallel to reduce costs and achieve faster time to results.

How GPU-accelerated HPC power can benefit your organization

Combining AWS global reach and a portfolio of HPC services with the most powerful NVIDIA GPUs and GPU-optimized software delivers the scale and performance needed for today's most demanding HPC workloads.



Scale flexibly

Flexible configurations and virtually unlimited infrastructure allow you to scale up or down as your workload demands. Access to AWS HPC services, powered by NVIDIA GPUs, grants the freedom of running your largest HPC applications anywhere.

Customer example:

A quantum chemistry researcher, Jean-Philip Piquemal, used NVIDIA GPUs on AWS, coupled with an HPC setup to power high-resolution molecular dynamics simulations.





Reduce costs

Leverage AWS' pay-as-you-go infrastructure, procuring only the capacity you need for the duration you need. On-demand GPU-powered [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) instances provide flexibility to start and stop workloads according to user demand and budget. You can also save up to 90% off on-demand prices with [AWS Spot pricing](#) for workloads that can tolerate interruptions.

Customer example:

Volkswagen Group Research estimates that it could save up to 70 percent of its existing hardware cost by using HPC on AWS powered by NVIDIA GPUs.



Optimize with leading-edge technologies

Predictive maintenance with digital twins powered by NVIDIA Omniverse on AWS can help reduce the frequency of maintenance checks, minimize downtime, and save costs.

Customer example:

Siemens Energy, a leading supplier of power plant technology, built a Digital Twin to support predictive maintenance of their power plants.



Leverage Artificial Intelligence and Machine Learning

With access to AWS high-performance computing, data storage, and machine learning tools, engineers and researchers can analyze large volumes of data, think outside the box and derive insights faster.

Customer example:

Lyft's autonomous vehicle (AV) division runs millions of simulations each year to improve the performance and safety of self-driving vehicles.



Shorten time to results

Running your most computationally-intensive simulations in parallel can support testing efforts, increase overall productivity, improve the product, and shorten the time to insight. Amazon EC2 instances, powered by NVIDIA and available on demand, give organizations the freedom to test different hypotheses in a shorter time frame and quickly change course during research.

Customer example:

FLSmidth accelerated mission-critical simulations, reducing the need for on-premise resources by almost four times.





Get started with HPC on AWS, powered by NVIDIA ›

¹ https://hyperionresearch.com/wp-content/uploads/2022/05/Hyperion-Research_ISC22-Market-Update_Cloud-Update.pdf

² https://hyperionresearch.com/wp-content/uploads/2022/06/Hyperion-Research_ISC22-Market-Update_May-30-2022_Combined.pdf

³ <https://blogs.nvidia.com/blog/2021/11/09/ai-millionfold-leaps/>

