Solution Brief

Driving Scientific Discovery

Application efficiency and performance for every supercomputing center.

Courtesy of Lawrence Berkeley National Lab.

Introduction

Supercomputing is driving breakthroughs in research domains such as **climate change**, **renewable energy**, **drug discovery**, and more. Large supercomputing centers support thousands of applications and researchers around the world. The research processes vast amounts of data and performs complex numerical simulations. The world's top supercomputers together utilize 16 billion annual CPU hours, highlighting the importance of these tools for advancing scientific discovery.

Applications

As simulations grow in complexity and size, they require increasingly large computational resources. The National Energy Research Scientific Computing Center (NERSC) Exascale Science Applications Program (NESAP) addresses this challenge by developing simulations, data analytics, and machine learning/deep learning (ML/DL) solutions. NESAP improves scalability and ensures access to computational resources—critical for researchers working with larger models and parameter space searches—to enable breakthrough scientific discoveries.

Use Cases

Scientists at the Lawrence Berkeley and Los Alamos National Laboratories are utilizing applications like MILC, BerkeleyGW, and the Exascale Atomistic Capability for Accuracy, Length, and Time (EXAALT) project to simulate and study the behavior of subatomic particles and molecules. These simulations help researchers understand the building blocks of matter and develop molecular dynamics simulation platforms that balance the costs of running simulations with the quality of results. Perlmutter, one of NERSC's supercomputers, is designed to leverage scientific AI models like DeepCAM, which is used to detect hurricanes and atmospheric rivers in climate data.

Accelerate, Reclaim, and Save

Running a selection of four key HPC and ML NESAP applications on a four-way **NVIDIA A100 Tensor Core** GPU-accelerated server, NERSC is seeing a 12X performance speedup over a dual-socket x86 server. At performance parity, a GPU-accelerated cluster consumes 588MWh less energy per month compared to a CPU-only system. If a researcher ran the same workload at the same speed on a four-way NVIDIA A100 cloud instance for a month, they would save $4 million compared to a CPU-only instance.

To learn more, visit: **nvidia.com/sustainable-computing**

Key Points

- Thousands of high-performance computing and AI applications
- 16 billion CPU hours consumed annually
- 12X performance speedup
- $4.2 million 588 megawatt hours (MWh) saved monthly

3X Lower Cost and 5X Less Energy

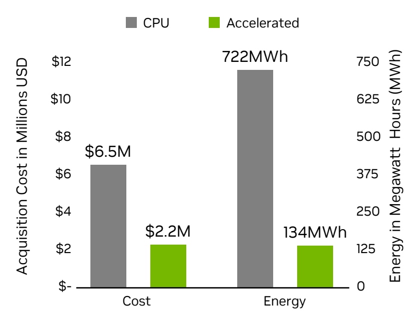

Based on the geometric mean of application speedups and energy consumption vs. dual AMD 7763 | benchmark applications | BerkeleyGW | DeepCAM | EXAALT | MILC and monthly prices on Microsoft Azure for instances standard_NC96ADS_A100_v4 and standard_D96ads_v5

> "Researchers need energy-efficient, performant ways for their simulation and AI applications to support their work. At NERSC, we've adopted accelerated computing to support research across a broad variety of science applications."

Nick Wright,
Chief Architect, NERSC

© 2023 NVIDIA Corporation and affiliates. All rights reserved. NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation and affiliates in the U.S. and other countries. Other company and product names may be trademarks of the respective owners with which they are associated. 2746602. DEC23