



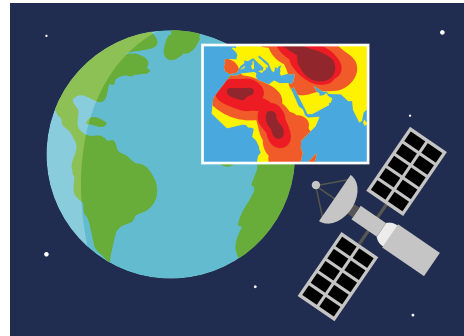
THREE REASONS TO DEPLOY NVIDIA TESLA P100 IN YOUR DATA CENTER

From scientific discovery to artificial intelligence, HPC is an important pillar that fuels the progress of humanity. Modern HPC data centers are currently solving some of the greatest challenges facing the world today.  With traditional CPUs no longer delivering the performance gains they used to, the path forward for HPC data centers is GPU-accelerated computing.


NVIDIA® Tesla® is the leading platform for accelerated computing and powers some of the largest data centers in the world—delivering significantly higher throughput while saving money. NVIDIA Tesla P100 powered by NVIDIA® Pascal™ architecture is the computational engine for scientific computing and artificial intelligence. Here are three powerful reasons to deploy NVIDIA Tesla P100 GPUs to your data center.


Reason 1: Be Prepared for the AI Revolution

The AI revolution is here, and every data center should be equipped for it.  AI is the engine behind consumer services we use every day, like web searches and video recommendations. In HPC, AI is enabling new ways to solve complex scientific challenges in bioinformatics, drug discovery, and high-energy physics.



AI is Monitoring Earth's Vitals

NASA is using AI to measure the effects of carbon and greenhouse gas emissions on the planet. 

NVIDIA Tesla P100 is the computational engine driving the AI revolution and enabling HPC breakthroughs. For example, researchers at New York's Icahn School of Medicine at Mount Sinai are using deep learning to analyze over 100,000 patient health records to predict patients likely to develop serious illnesses and provide treatment up to one year before traditional diagnoses. 




AI Predicts And Prevents Disease

Mount Sinai is using deep learning to give doctors a life-saving edge by identifying high-risk patients before diseases are diagnosed.

"Advances in Artificial Intelligence (AI) technology have opened up new markets and new opportunities for progress in critical areas such as health, education, energy, and the environment."


-Executive Office of the President National Science Technology Council

Reason 2: Top Applications are GPU-Accelerated

Over 450 HPC applications are already GPU-optimized in a wide range of areas including quantum chemistry, molecular dynamics, climate and weather, and more. 


In fact, an independent study by Intersect360 Research shows that 70% of the most popular HPC applications, including 10 of top 10 have built-in support for GPUs.

With most popular HPC applications and all deep learning frameworks GPU-accelerated, every HPC customer would see the majority of their data center workload benefit from GPU-accelerated computing.

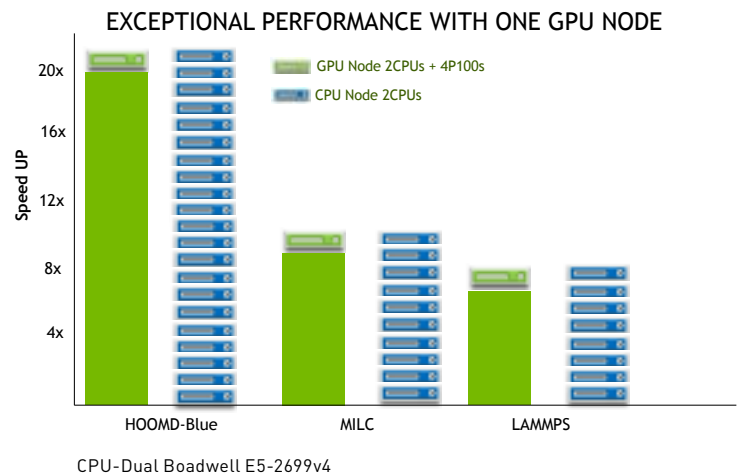


Reason 3: Boost Data Center Productivity & Throughput

Data center managers all face the same challenge: how to meet the demand for computing resources that often exceed available cycles in the system.

The NVIDIA Tesla P100 dramatically boosts the throughput of your data center with fewer nodes, completing more jobs and improving data center efficiency. 

A single server node with P100 GPUs can replace up to 20 CPU nodes. For example, for MILC, a single node with four P100's will do the work of 10 dual socket CPU nodes while for HOOMD Blue a single P100 node can replace 21 CPU nodes. With fewer overheads on networking and cables, strong nodes provide high application throughput at substantially reduced costs.



Find systems powered by Tesla GPUs at www.NVIDIA.com/where-to-buy-tesla.