





THREE REASONS TO DEPLOY NVIDIA TESLA V100 IN YOUR RESEARCH CENTER


From scientific discovery to artificial intelligence, high performance computing (HPC) is an important pillar that fuels the progress of humanity. Modern research 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 is GPU-accelerated computing.

NVIDIA® Tesla® is the leading platform for accelerated computing and powers some of the largest research centers in the world—delivering significantly higher throughput while saving money. NVIDIA Tesla V100  powered by NVIDIA Volta architecture is the computational engine for scientific computing and artificial intelligence. Here are three powerful reasons to deploy NVIDIA Tesla V100 GPUs in your research center.

Reason 1: HPC + AI=Unprecedented Opportunities

The AI revolution is here, and every research center should be equipped for it.  AI extends traditional HPC by allowing researchers to analyze large volumes of data for rapid insights where simulation alone cannot fully predict the real world, in areas like medical imaging, bioinformatics, drug discovery, and high-energy physics.

Tesla V100 is the computational engine driving the AI revolution and enabling HPC breakthroughs. For example, researchers at Mayo Clinic are using deep learning for radiogenomics, using the rich information contained in multi-modal images such as multi-parametric magnetic resonance (MR) images to determine genomic properties of tumors with high accuracy. 




AI IS HELPING DETECT GROWTH PROBLEMS IN CHILDREN

To help detect growth-related problems in children, MGH & BWH Center for Clinical Data Science has developed an automated bone-age analyzer that delivers results that are 99% accurate in seconds versus days.




AI PREDICTS AND PREVENTS DISEASE

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

“What we are able to do now with the technology and deep learning was almost unimaginable just five years ago.”


Bradley J. Erickson – Director, Radiology Informatics Lab, Mayo Clinic

Reason 2: Top Applications are GPU-Accelerated

Over 500 HPC applications are already GPU-optimized in a wide range of areas including quantum chemistry, molecular dynamics, bioinformatics, and more. 

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

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




Reason 3: Boost Research Center Productivity & Throughput

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

The NVIDIA Tesla V100 dramatically boosts the throughput of your institution's data center with fewer nodes, completing more jobs and improving data center efficiency.

A single server node with V100 GPUs can replace up to 50 CPU nodes. For example, for HOOMD-BLUE, a single node with four V100's will do the work of 43 dual socket CPU nodes while for MILC, a single V100 node can replace 14 CPU nodes. With lower networking, power, and rack space overheads, accelerated nodes provide higher application throughput at substantially reduced costs.



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