THREE REASONS TO DEPLOY NVIDIA TESLA V100 IN YOUR DATA CENTER

From scientific discovery to artificial intelligence, high performance computing (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.
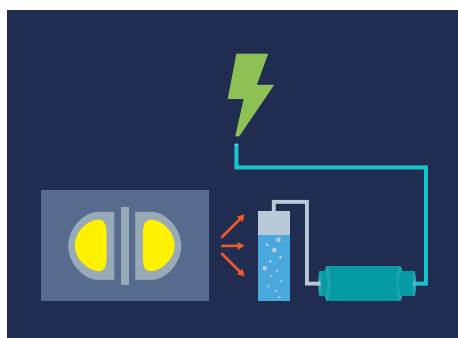
NVIDIA® Tesla® is the leading platform for accelerated computing and powers some of the largest data centers in the world—delivering significantly higher throughput while saving money.

NVIDIA Tesla 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 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. 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 domains like bioinformatics, drug discovery, and high-energy physics.

AI IS HELPING RESEARCHERS CREATE CLEAN ENERGY USING FUSION REACTORS

Princeton University is using AI to predict and respond to disruptions in a tokamak reactor (a magnetic fusion reactor device). 🐦

NVIDIA Tesla V100 is the computational engine driving the AI revolution and enabling HPC breakthroughs. For example, researchers at the University of Florida University and University of North Carolina leveraged GPU deep learning to develop ANAKIN-ME (ANI) to reproduce molecular energy surfaces at extremely high (DFT) accuracy and 1-10/millionths of the cost of current computational methods. 🐦

AI SPEEDS DRUG DEVELOPMENT

Today's CPU-based methods can process as many as two million candidates per year, but researchers at the Universities of Florida and North Carolina have developed a way to use GPU-based deep learning to accomplish the same work in under two minutes.

"A fast and accurate description of molecular energetics can lead to breakthroughs in many fields including drug discovery and materials science. ANI may start the next revolution in computational chemistry,"
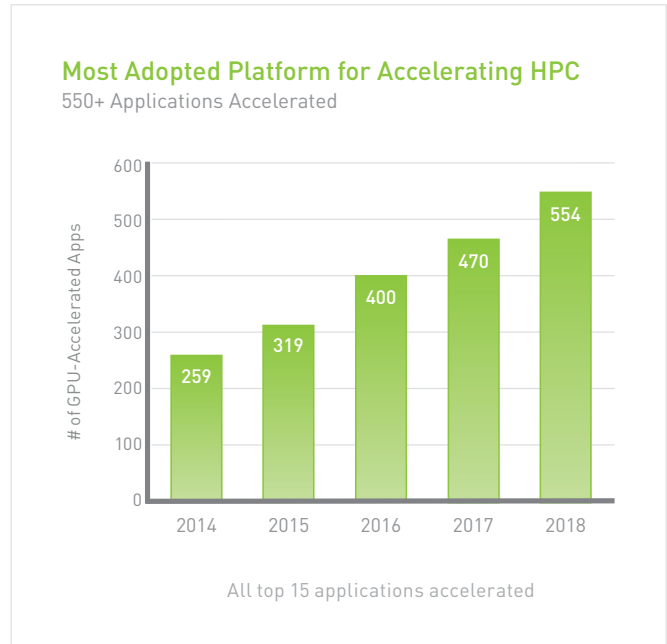***Justin S. Smith a Chemistry Graduate student at the University of Florida.***

Reason 2: Top Applications are GPU-Accelerated

Over 550 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 the top 10 have built-in support for GPUs.

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

Most Adopted Platform for Accelerating HPC
550+ Applications Accelerated

Year	# of GPU-Accelerated Apps
2014	259
2015	319
2016	400
2017	470
2018	554

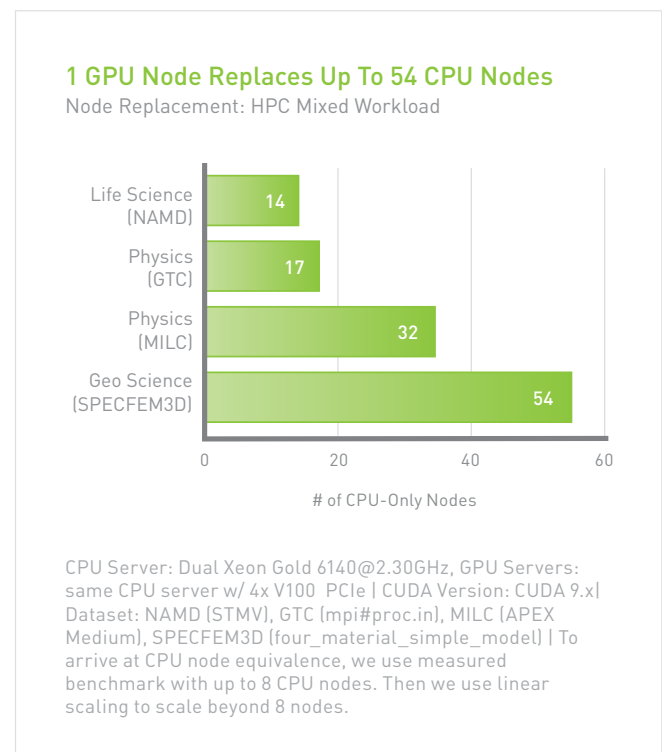All top 15 applications accelerated

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.

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

A single server node with V100 GPUs can replace over 60 CPU nodes. For example, for SPECFEM3D, a single node with four V100's will do the work of 53 dual-socket CPU nodes while for NAMD a single V100 node can replace 13 CPU nodes. With lower networking, power, and rack space overheads, accelerated nodes provide higher application throughput at substantially reduced costs.

1 GPU Node Replaces Up To 54 CPU Nodes
Node Replacement: HPC Mixed Workload

Application	# of CPU-Only Nodes
Life Science (NAMD)	14
Physics (GTC)	17
Physics (MILC)	32
Geo Science (SPECFEM3D)	54

CPU Server: Dual Xeon Gold 6140@2.30GHz, GPU Servers: same CPU server w/ 4x V100 PCIe | CUDA Version: CUDA 9.x| Dataset: NAMD (STMV), GTC (mpi#proc.in), MILC (APEX Medium), SPECFEM3D (four_material_simple_model) | To arrive at CPU node equivalence, we use measured benchmark with up to 8 CPU nodes. Then we use linear scaling to scale beyond 8 nodes.

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

NVIDIA.