THE WORLD'S MOST EFFICIENT SUPERCOMPUTER FOR AI AND DEEP LEARNING

NVIDIA DGX SATURNV

THE GREENEST PATH TO EXASCALE

EFFICIENCY

The NVIDIA® DGX SATURNV taps into the compute power of 1800 NVIDIA® DGX-1™ server nodes to drive new levels of deep learning and data science.

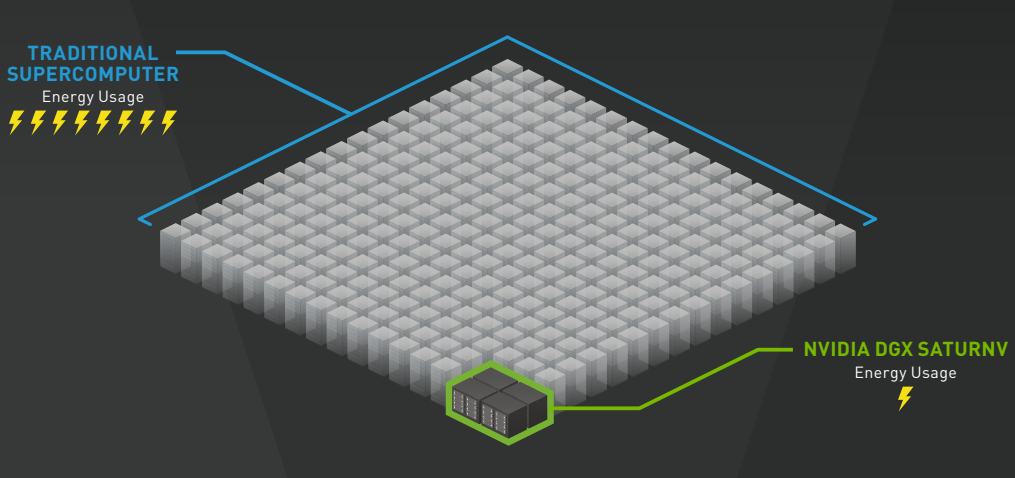

SATURNV is raising the bar for energy efficiency (making the Green500 with 15 GFLOPS per watt of FP64 efficiency) with a total expected computational capacity footprint of 660 petaFLOPS of AI horsepower.

The GREEN 500

NVIDIA is lighting a path for enterprises to build GPU-accelerated data centers of the future, breaking the barriers of Moore's law scaling, while offering more compute in less space than ever before.

FASTER NEURAL NETWORKS TRAINING FOR SMARTER RESULTS

INTELLIGENCE

The deeper the neural network, the more abstract concepts it can learn, the more intelligent it becomes. This can be the difference between a network identifying a square versus understanding that an image is a specific type of cancer cell.

TAP INTO THE WORLD'S FASTEST GPUs

POWER

Neural networks are the backbone of artificial intelligence, but training them takes an incredible amount of time and compute power. With GPUs, hundreds of networks can now be trained in parallel, accelerating solutions for some of the world's hardest problems through AI.

1,800
NODES

15,000
V100 Tensor Core GPUs

1.8
exaFLOPS

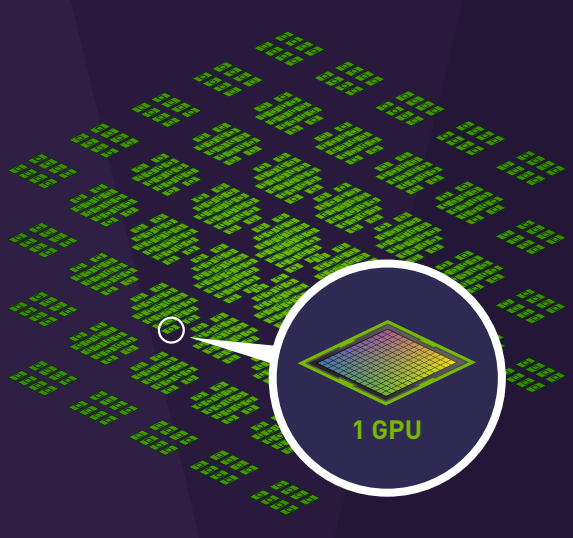

AI PERFORMANCE LEADERSHIP

AT SCALE

DGX SuperPOD

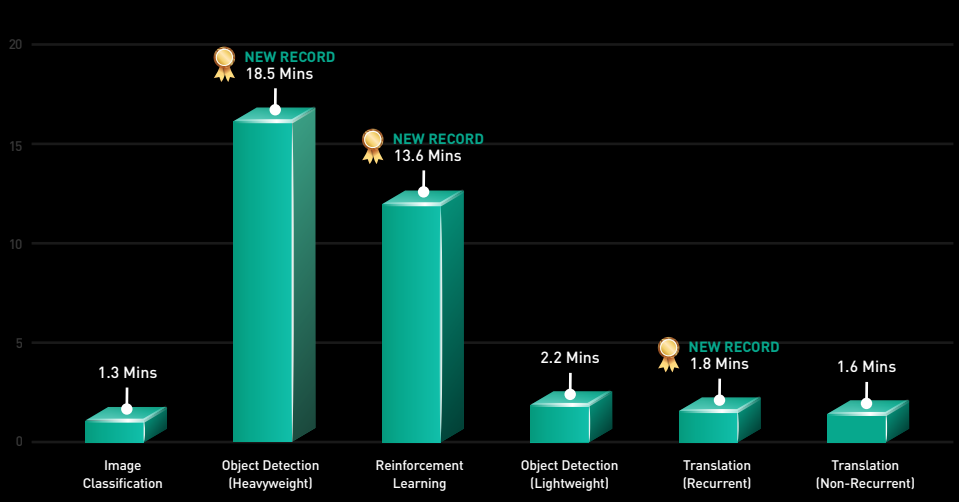

Test platform: NVIDIA DGX SuperPOD in different configurations. Each DGX-2H, Dual-Socket Xeon Platinum 8189, 384 GB system RAM, 16x 32 GB Tesla V100 GPUs connected via NVSwitch.
All results from MLPerf 0.6 training closed; retrieved from **www.mlperf.org** 10 July 2019. MLPerf name and logo are trademarks. See **www.mlperf.org** for more information.

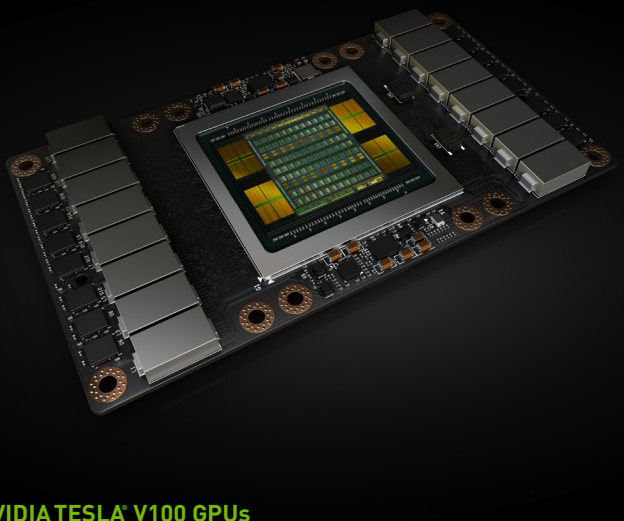

POWERED BY NVIDIA TESLA® V100 GPUs
Built on the latest NVIDIA Volta GPU architecture

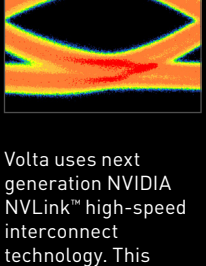

Volta uses next generation NVIDIA NVLink™ high-speed interconnect technology. This delivers 2X the throughput, compared to the previous generation of NVLink.

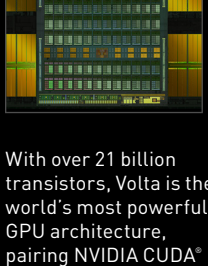

With over 21 billion transistors, Volta is the world's most powerful GPU architecture, pairing NVIDIA CUDA® and Tensor Cores, delivering the performance of an AI supercomputer in a GPU.

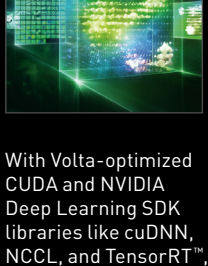

With Volta-optimized CUDA and NVIDIA Deep Learning SDK libraries like cuDNN, NCCL, and TensorRT™, the industry's top frameworks and applications can easily tap into the power of Volta.

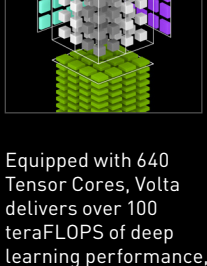

Equipped with 640 Tensor Cores, Volta delivers over 100 teraFLOPS of deep learning performance, a 5X increase compared to prior generation NVIDIA Pascal™ architecture.

TRANSFORMING INDUSTRIES WITH AI SUPERCOMPUTING

IMPACT

THE NATIONAL CANCER MOONSHOT INITIATIVE

NVIDIA is teaming with the National Cancer Institute, and U.S. Department of Energy to create an AI platform for accelerating cancer research.

NASA AMES GLOBAL CLIMATE CHANGE

A deep learning framework for Satellite Image Classification helps safeguard our planet by using satellite imagery and DeepSAT to measure the effects of carbon and greenhouse gases on crops, vegetation, and the urban landscape.

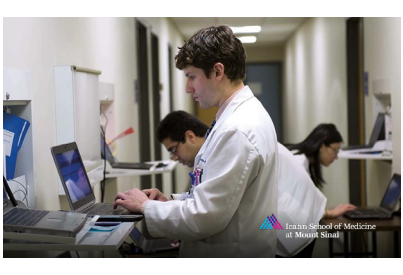

ICAHN SCHOOL OF MEDICINE

The school developed 'Deep Patient', a tool trained on thousands of patient records using GPU accelerators to identify high-risk patients.

BUILD YOUR OWN SATURNV WITH NVIDIA DGX-1 AND DGX-2

nvidia.com/dgx