



NVIDIA DGX-2 THE WORLD'S MOST POWERFUL DEEP LEARNING SYSTEM FOR THE MOST COMPLEX AI CHALLENGES

The Challenge of Scaling to Meet the Demands of Modern AI and Deep Learning

Deep neural networks are rapidly growing in size and complexity, in response to the most pressing challenges in business and research. The computational capacity needed to support today's modern AI workloads has outpaced traditional data center architectures. Modern techniques that exploit use of model parallelism are colliding with the limits of inter-GPU bandwidth, as developers build increasingly large accelerated computing clusters and push the limits of data center scale. A new approach is needed—one that delivers almost limitless AI computing scale in order to break through the barriers to achieving faster insights.

Performance to Train the Previously Impossible

Increasingly complex AI demands unprecedented levels of compute. NVIDIA® DGX-2™ is the world's first 2 petaFLOPS system, packing the power of 16 of the world's most advanced GPUs and accelerating the newest deep learning model types that were previously untrainable. With groundbreaking GPU scale, you can train models 4X bigger on a single node. In comparison with legacy x86 architectures, DGX-2's ability to train ResNet-50 would require the equivalent of 300 servers with dual Intel Xeon Gold CPUs, costing over \$2.7 million dollars.

NVIDIA NVSwitch—A Revolutionary AI Network Fabric

Leading-edge research demands the freedom to leverage model parallelism and requires never-before-seen levels of inter-GPU bandwidth. NVIDIA has created NVSwitch to address this need. Like the evolution from dial-up to ultra-high speed broadband, NVSwitch delivers a networking fabric for the future, today. With DGX-2, model complexity and size are no longer constrained by the limits of traditional architectures. Embrace model-parallel training with a networking fabric that delivers 2.4TB/s of bisection bandwidth for a 24X increase over prior generations. This new interconnect "superhighway" enables limitless possibilities for model types that can reap the power of distributed training across 16 GPUs at once.

SYSTEM SPECIFICATIONS

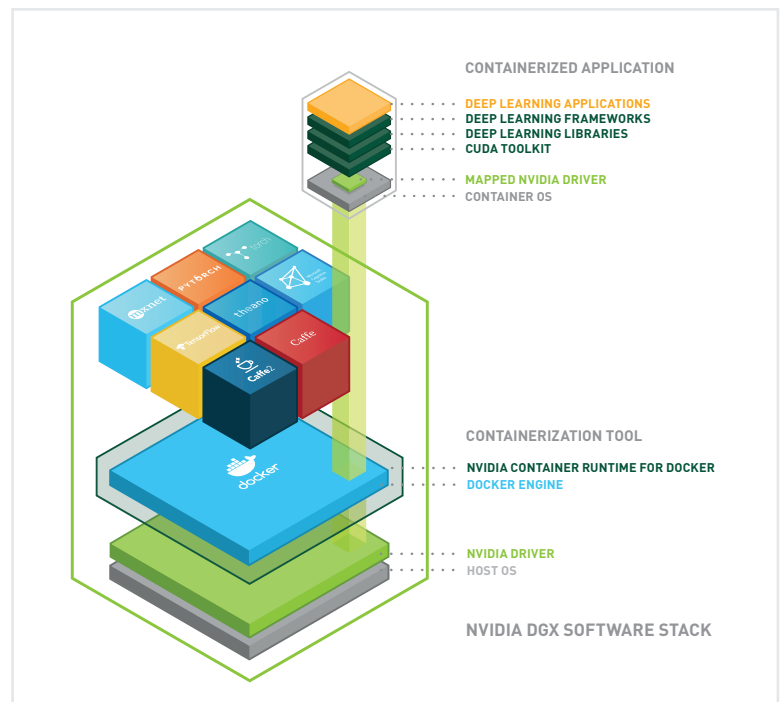
GPUs	16X NVIDIA® Tesla® V100
GPU Memory	512GB total
Performance	2 petaFLOPS
NVIDIA CUDA® Cores	81920
NVIDIA Tensor Cores	10240
NVSwitches	12
Maximum Power Usage	10kW
CPU	Dual Intel Xeon Platinum 8168, 2.7 GHz, 24-cores
System Memory	1.5TB
Network	8X 100Gb/sec Infiniband/100GigE Dual 10/25/40/50/100GbE
Storage	OS: 2X 960GB NVME SSDs Internal Storage: 30TB (8X 3.84TB) NVME SSDs
Software	Ubuntu Linux OS Red Hat Enterprise Linux OS (See software stack for details)
System Weight	360 lbs (163.29 kgs)
Packaged System Weight	400lbs (181.44kgs)
System Dimensions	Height: 17.3 in (440.0 mm) Width: 19.0 in (482.3 mm) Length: 31.3 in (795.4 mm) - No Front Bezel 32.8 in (834.0 mm) - With Front Bezel
Operating Temperature Range	5°C to 35°C (41°F to 95°F)

AI Scale on a Whole New Level


Modern enterprises need to rapidly deploy AI power in response to business imperatives and scale-out AI, without scaling-up cost or complexity. We've built DGX-2 and powered it with DGX software that enables accelerated deployment and simplified operations—at scale. DGX-2 delivers a ready-to-go solution that offers the fastest path to scaling-up AI, along with virtualization support, to enable you to build your own private enterprise grade AI cloud. Customers can choose to use either pre-installed Ubuntu Linux Host OS, popular among developers, or install Red Hat Enterprise Linux Host OS if IT teams prefer to integrate DGX-2 with their existing Red Hat data center deployment. Now businesses can harness unrestricted AI power in a solution that scales effortlessly with a fraction of the networking infrastructure needed to bind accelerated computing resources together. With an accelerated deployment model, and an architecture purpose-built for ease of scale, your team can spend more time driving insights and less time building infrastructure.

Access to AI Expertise

With DGX-2, you benefit from NVIDIA's AI expertise, enterprise-grade support, extensive training, and field-proven capabilities that can jump-start your work for faster insights. Our dedicated team is ready to get you started with prescriptive guidance, design expertise, and access to our fully-optimized DGX software stack. You get an IT-proven solution, backed by enterprise-grade support, and a team of experts who can help ensure your mission-critical AI applications stay up and running.



NVIDIA DGX-2 Delivers 195X Faster Deep Learning Training



Workload: ResNet-50, BS=256, 90 epochs to solution | CPU: dual Xeon Platinum 8180 | DGX-1 GPU: 8X Tesla V100 32 GB | DGX-2 GPU: 16X Tesla V100 32 GB

For more information, visit www.nvidia.com/DGX-2