Doing Large Scale AI Faster IBM SPECTRUM STORAGE FOR A and NVIDIA DGX

DGX REFERENCE ARCHITECTURE SOLUTION





Jacci Cenci, Technical Marketing **Engineer, NVIDIA Systems Engineering**

THE TRUE TCO OF AN AI PLATFORM



1. Designing and Building an AI Compute Platform - <u>from Scratch</u>

AI TRAINING REQUIRES FULL STACK INNOVATION



Workload: Time (in minutes) to train ResNet-50 network.

STACKING DGX

Aggregating Resources - Scaling Out



- Preconditioned to Scale
- Preconditioned for effective Multi-Node -MultiGPU scaling
- Preconditioned to aggregate resources
- Horovod allowed NVIDIA to scale from one to eight GPU's
- Optimizing model training for self-driving sensing and perception technologies
- Leading to faster, safer systems.



SCALING WITH HOROVOD

One Process per GPU - One Data pipeline per GPU



- At NVIDIA, Horovod training jobs are run on <u>DGX</u> <u>SATURNV</u> cluster. From there, it runs in Docker containers (hosted on <u>NGC</u>)
- To train their self-driving systems for example, NVIDIA uses TensorFlow images that come with Horovod preinstalled on them alongside CUDA, CuDNN, and NCCL.
- With Horovod, researchers experience a scaling factor greater than seven times on an eight GPU system, with hundreds of multi-GPU jobs launched per day per perception model (e.g., lane detector, road signs, etc.).
- The process of launching jobs and finding optimized parameters from NVIDIA's AI training and inference infrastructure.





IBM SPECTRUM STORAGE FOR AI WITH NVIDIA DGX

The Engine to Power Your AI Data Pipeline

HARDWARE

- NVIDIA DGX-1 | up to 9x DGX-1 Systems
- IBM Spectrum Scale NVMe Appliance | 40GB/s per node, 120GB/s in 6RU | 300TB per node
- NETWORK: Mellanox SB7700 Switch | 2x EDR IB with RDMA

SOFTWARE

- NVIDIA DGX SOFTWARE STACK | NVIDIA Optimized Frameworks
- **IBM:** High performance, low latency, parallel filesystem
- IBM: Extensible and composable



IBM SPECTRUM STORAGE AND NVIDIA DGX ARCHITECTURE



IBM SPECTRUM STORAGE and NVIDIA DGX servers

A scalable software defined infrastructure

Performance:

- NVME throughput of 120GB/s in a rack.
- Over 40GB/s sustained random reads per 2U.

Flexible to Grow as Needed:

- Up to 9 DGX-1 servers in one rack (72 GPUs)
- Scalability: 1 server to 8 Exabytes and a Yottabyte of files

Extensible for the AI Data Pipeline:

• Support for any tiered storage based on policies. Send data to low cost tiers like object, tape, or cloud

Multi-protocol support: Scale, NFS, POSIX, CIFS, Object, Block, OpenStack etc.

IBM SPECTRUM STORAGE FOR AI WITH NVIDIA DGX: SCALABLE REFERENCE ARCHITECTURES

Scaling with NVIDIA DGX-1

- Start with a single IBM Spectrum Scale NVMe and a single DGX-1
- Grow capacity in a cost-effective, modular approach
- Each config delivers balanced performance, capacity and scale
- IBM Spectrum Scale NVME all-flash appliance is power efficient to allow maximum flexibility when designing rack space and addressing power requirements











9:3 Configuration

🕺 NVIDIA

IBM STORAGE WITH NVIDIA DGX: FULLY-OPTIMIZED AND QUALIFIED

Performance at Scale

For multiple DGX-1 servers, IBM Spectrum Scale on NVMe architecture demonstrates linear scale up to full saturation of all DGX-1 server GPUs

The multi-DGX server image processing rates shown demonstrate scalability for Inception-v4, ResNet-152, VGG-16, Inception-v3, ResNet-50, GoogLeNet and AlexNet models



THE IMPACT OF IBM STORAGE + NVIDIA DGX ON TIMELINE



2. Deploying an Integrated, Full-Stack AI Solution using DGX Systems

THE IMPACT OF IBM STORAGE + NVIDIA DGX ON TIMELINE



2. Deploying an Integrated, Full-Stack AI Solution using DGX Systems

IBM & NVIDIA REFERENCE ARCHITECTURE

Validated design for deploying DGX at-scale with IRM

Download at https://bit.ly/2GcYbg0

Learn more about DGX RA Solutions at: <u>https://bit.ly/20pXYeC</u>



Figure 12: Solution architecture diagram

IBM Spectrum Storage for AI with NVIDIA DGX[®] – Reference Architecture of Infrastructure Solutions for AI Workloads

The engine to power your AI data pipeline

Artificial intelligence (AD) – including deep learning (DU) and machine learning (MU) – is the engine rapidly powering innovation across industries from healthcare to autonomous vehicles and agriculture. By 2020, IBM⁴ projects that the world's volume of digital data will exceed 44 zettalytes ¹ Organisations that recognize the value of their data for decisions and actions are turning to DL systems that can rapidly ingest, accurately interpret, and quickly provide key data insights from the volumes of new data generated nov and in the future.

Enterprises are increasing investment in AL research and innovation, with related patents growing more than 30% and academic papers by 13% during the last decade.² Arguably, only one kind of enterprise will survive and thrive in the future – the data-crivee enterprise.

Highly performant and scalable DL systems must excel at data ingestion, data preparation, data training, and verification, and deliver inferences and classifications while handling the growing demands of DL in the organi-

From Concept to World-Record Setting Supercomputer in Three Weeks

The CIRCE Supercomputer and The MLPerf Benchmark Suite

- The MLPerf benchmark suite can be used as a proxy to characterize the performance of GPU-based computer systems.
- Developed with support from over 40 commercial and academic organizations including NVIDIA, Google, Microsoft, Facebook, Intel, AMD, Stanford, and Harvard.
- MLPerf is the first benchmark covering such a broad range of AI algorithms and is representative of many of the workloads in use across industry, research and government.

Features

Component	Technology	Description
Compute Nodes	36 NVIDIA DGX-2H servers	•576 Tesla V100 SXM3 GPUs
		•18.4 TB of HBM2 memory
		 72 petaFLOPS via Tensor Cores
Compute Network	Mellanox InfiniBand CS7510	 324 EDR/100 Gbps ports
	Director Class switch	 Eight connections per node
Storage Network	Mellanox InfiniBand CS7520	 216 EDR/100 Gbps ports
	Director Class switch	 Two connections per node
Management Networks	Arista 7060CX2 and 7010	Each DGX-2H server has ethernet
	ethernet switches	connections to both switches
High Speed Storage	IBM Elastic Storage Server (ESS),	 276 TB raw storage
	GS4S model	 40 GB/s read performance
Management Software	DGX POD management software	
	NVIDIA GPU Cloud (NGC)	NGC provides the best
		performance for all DL
User Runtime		frameworks
Environment	Slurm	Slurm is used for the
		orchestration and scheduling of multi-cell and multi-cell
		training jobs

Optimized System for Multi-node DL and HPC

- Consists of 15 Racks with 36 DGX-2H servers
 - Compute: 12 racks of DGX-2H servers
 - Utility: one rack for the management servers and ethernet networking InfiniBand: one rack for the two Director Class InfiniBand switches Storage: one rack for the IBM ESS nodes
- Total of 576 V100 GPUs
- Built using the NVIDIA DGX POD reference architecture and was configured to be a scalable and balanced system providing maximum performance.
- The DGX-2H server is an enhanced version of DGX-2 server which features faster Tesla V100 GPUs running at 450 watts per GPU and higher-frequency CPUs.
- Utilizes 16 Tesla V100 GPUs and NVIDIA NVSwitch[™] technology to be the world's first 2.1 PetaFLOPS server.

Rack Layout



Network Topology



STORAGE SYSTEM REQUIREMENTS

Provide a Hierarchical Storage System

- Minimizes time to stage data to local storage
- Allows for training of DL models that require DL models which require peak IO performance up to 10 GB/s and data sizes which exceed local storage cache.
- Provides a large, cost-effective, storage area for long-term data storage for data that are not in active use



IBM SPECTRUM SCALE CONFIGURATION Specific Deployment

- 276 TB total flash storage
- Four storage disk shelves and two ESS nodes
- One management server
- Total of 14 IB ports allocated to the management server and ESS nodes for client data and inter cluster traffic
- Total of one 100 Gbps ports allocated on the management server.
- Total of six 1 Gbps or 10 Gbps ethernet ports allocated for BMC/XCAT. Note that this is an allocation on internal switch for the appliance.

DGX-2H STORAGE AND CACHING Hierarchy

Storage Hierarchy Level	Technology	Total Capacity	Read Performance
RAM	DDR4	1.5 TB per node	> 100 GB/s
Internal Storage	NVMe	30 TB per node	> 15 GB/s
High-Speed Storage	IBM SPECTRUM SCALE using SSD	276 TB	> 40 GB/s aggregate > 15 GB/s per node
Long Term Storage	NFS using SSD	> 1 PB	10 GB/s aggregate 1 GB/s per node
Home	NFS using NVMe	30 TB	N/A

Supporting companies





First Industry Benchmark for Measuring AI Performance



A broad ML benchmark suite for measuring performance of ML software frameworks, ML hardware accelerators, and ML cloud platforms.

https://mlperf.org/

NVIDIA Sets New Records for AI Performance

NVIDIA Deep Learning Platform Continues to be the "Gold Standard" for Al

NVIDIA Sets New Records in AI Performance:

- Captured top spots in MLPerf, ran models of all kinds of complexity, in the industry's first comprehensive Al benchmark.
- Tensor Core GPUs are the fastest and combined with CUDA the most versatile platform for AI.
- We also have the most affordable platform, available everywhere from desktop to cloud services

NVIDIA's Record-Setting Platform Available Everywhere to All Al Developers:

- Software innovations used to achieve this industry-leading performance is available free of charge via the NGC container registry.
- NGC containers are available for all key AI frameworks and can be used anywhere, on desktops, workstations, servers and all leading cloud services

State-of-the-art AI Computing Requires Full Stack Innovation:

• NVIDIA's industry-leading results on MLPerf benchmarks are a testament to years of relentless innovation in architecture, systems, and a full software stack.

MLPerf Usages & Workloads

Diverse Use-Cases Towards a Full Performance Picture

Usage	Network	Data-Set
Image Classification	ResNet-50 v1.5	ImageNet
Object Detection (Heavy Weight)	Mask R-CNN	сосо
Object Detection (Light Weight)	Single-Shot Detector	сосо
Translation (recurrent)	NMT	WMT English-German
Translation (non-recurrent)	Transformer	WMT English-German
Recommendation	Neural Collaborative Filtering	MovieLens 20M
Reinforcement Learning	Mini-Go	

MLPERF RESULTS - SINGLE NODE

Results are Time to Complete Model Training



MLPERF RESULTS - AT SCALE

Results are Time to Complete Model Training



Test Platform: For Image Classification and Translation (non-recurrent), DGX-1V Cluster. For Object Detection (Heavy Weight) and Object Detection (Light Weight), Translation (recurrent) DGX-2H Cluster. Each DGX-1V, Dual-Socket Xeon E5- 2698 V4, 512GB system RAM, 8 x 16 GB Tesla V100 SXM-2 GPUs. Each DGX-2H, Dual-Socket Xeon Platinum 8174, 1.5TB system RAM, 16 x 32 GB Tesla V100 SXM-3 GPUs connected via NVSwitch

THANK YOU!





Backup





THE FASTEST PATH TO AI SCALE ON A WHOLE NEW LEVEL

Today's business needs to scale-out AI, without scaling-up cost or complexity

- Powered by DGX software
- Accelerated AI-at-scale deployment and effortless operations
- Unrestricted model parallelism and faster time-to-solution



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

- First 2 PFLOPS System
- 16 V100 32GB GPUs Fully Interconnected
- NVSwitch: 2.4 TB/s bisection bandwidth
- 24X GPU-GPU Bandwidth
- 0.5 TB of Unified GPU Memory
- 10X Deep Learning Performance



NVIDIA DGX-2

Limitless AI for Exploration Without Boundaries

PERFORMANCE TO TRAIN THE PREVIOUSLY IMPOSSIBLE NVSWITCH: REVOLUTIONARY AI NETWORK FABRIC AI SCALE ON A WHOLE NEW LEVEL



2 PFLOPS AI SUPERCOMPUTER DELIVERING 10X PERFORMANCE 2.4 TB/sec GPU INTERCONNECT FABRIC SPEED OF SCALE WITHOUT COMPLEXITY

COMMON SOFTWARE STACK ACROSS DGX FAMILY



DESIGNED TO TRAIN THE PREVIOUSLY IMPOSSIBLE



10X PERFORMANCE GAIN IN LESS THAN A YEAR



NVIDIA CONFIDENTIAL. DO NOT DISTRIBUTE.



THE WORLD'S FIRST 16 GPU AI PLATFORM

- Revolutionary SXM3 GPU package design
- Innovative 2 GPU board interconnect
- 32GB HBM2 stacked memory per GPU





NVSWITCH: THE REVOLUTIONARY AI NETWORK FABRIC

- Inspired by leading edge research that demands unrestricted model parallelism
- Like the evolution from dial-up to broadband, NVSwitch delivers a networking fabric for the future, today
- Delivering 2.4 TB/s bisection bandwidth, equivalent to a PCIe bus with 1,200 lanes
- NVSwitches on DGX-2 = all of Netflix HD <45s

NVME SSD STORAGE

Rapidly ingest the largest datasets into cache

- Faster than SATA SSD, optimized for transferring huge datasets
- Dramatically larger user scratch space
- The protocol of choice for next-gen storage technologies
- 8 x 3.84TB NVMe in RAID0 (Data)
- 25.5 GB/sec Sequential Read bandwidth (vs. 2 GB/sec for 7TB of SAS SSDs on DGX-1)





LATEST GENERATION CPU AND 1.5TB SYSTEM MEMORY

Faster, more resilient, boot and storage management

- More system memory to handle larger DL and HPC applications
- 2 Intel Skylake Xeon Platinum 8168 -2.7GHz, 24 cores
- 24 x 64GB DIMM System Memory

THE ULTIMATE IN NETWORKING FLEXIBILITY

Grow your DL cluster effortlessly, using the connectivity you prefer

- Support for RDMA over Converged Ethernet (ROCE)
- 8 EDR Infiniband / 100 GigE
- 1600 Gb/sec Total Bi-directional Bandwidth with low-latency
- Also supports Ethernet mode: Dual 10/25 Gb/sec





FLEXIBILITY WITH VIRTUALIZATION

Enable your own private DL Training Cloud for your Enterprise

- KVM hypervisor for Ubuntu Linux
- Enable teams of developers to simultaneously access DGX-2
- Flexibly allocate GPU resources to each user and their experiments
- Full GPU's and NVSwitch access within VMs — either all GPU's or as few as 1

RELIABILITY, AVAILABILITY, AND SERVICEABILITY

As your most critical endeavors are realized on DGX-2, trust our purposebuilt architecture to keep you running

- Ability to lose a GPU and gracefully continue operations in degraded mode
- Reducing failure recovery time with new diagnostic capabilities
- Graceful recovery after faulty hardware replaced
- Full in-rack serviceability for many components

ENTERPRISE GRADE INFRASTRUCTURE FOR AI

If your AI platform is critical to your business, you need one designed with your most important endeavors in mind

- Enable round-the-clock AI operations thanks to architectural design features found only in DGX-2
- Get started fast, train faster, remain faster
- DGX-2 comes with same enterprise grade support included with every DGX system

NVIDIA DGX-2 LIMITLESS DEEP LEARNING FOR EXPLORATION WITHOUT BOUNDARIES

The World's Most Powerful Deep Learning System for the Most Complex Deep Learning Challenges

- Performance to Train the Previously Impossible
- Revolutionary AI Network Fabric
- Fastest Path to AI Scale
- Powered by NVIDIA GPU Cloud

For More Information: nvidia.com/dgx-2

