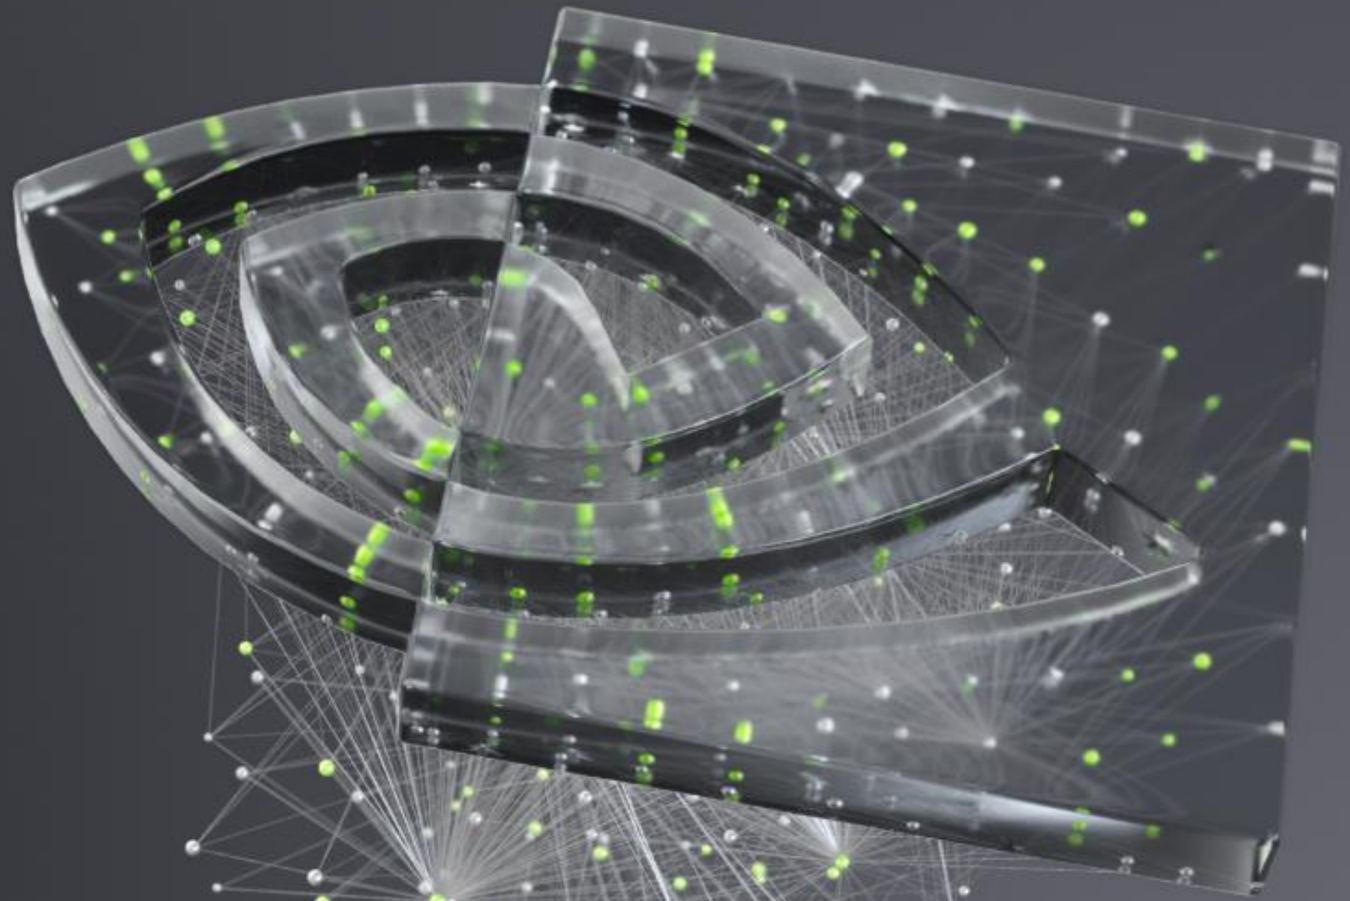

优化架构设计充分发挥算力极致性能

DGX POD架构分析以及为客户带来的益处

September 2021

AI IS EVERYWHERE

A Massive Opportunity

Exploding Model Complexity

Parameters of large models can no longer fit on the main Memory of even the largest GPU or multi-GPU server

Impractical Training Times

Mathematical operations required to train these large models requires months or years on sizable systems

Distributed Training Limitations

Parallelism, where the model is split across multiple GPUs, breaks down for large models due to increased inter-node communications

Full Stack solution

Allows Data Monsters to focus on application development

Most Complex Models

Training GPT-3 175B: 1 month on a 140-node DGX SuperPOD vs 355 years on one V100 GPU

Fastest Time to Solution

Very high GPU utilization achieved from parallelization and DOCA minimizes the training time

Unmatched Parallelism

DGX POD multi-rail high performance InfiniBand network optimized for SHARP v2 in-network acceleration

Full Solutions

Move from Ingredients & Recipes To fully Validated & Bundled Solutions

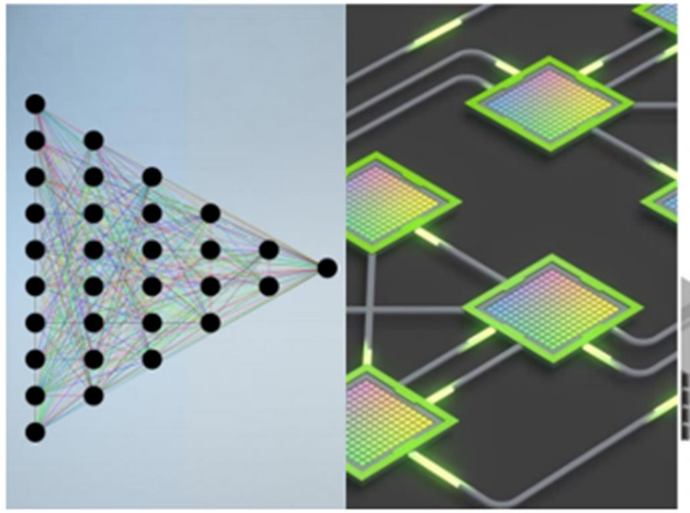

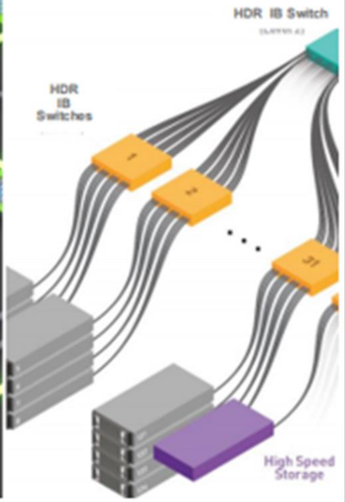

THE CHALLENGE OF AI INFRASTRUCTURE

Enterprise AI requires time, expertise and the right approach to architecture

MAAS

Provisioning Software Stack

slurm workload manager

Schedulers

Prometheus Grafana ZABBIX pagerduty docker

Operations

User Apps

Deploy Production AI applications | Deploy Kubernetes on K8S nodes

Deploy DL/ML development tools | Bootstrap all nodes

Deploy management services | Provision all node(s)

Deploy Slurm on Slurm nodes | Prepare provisioning node

DGX Components

Ethernet

IB

Networking

Block Storage

Shared Filesystems

Parallel Filesystems

Object Storage

Storage

Configuration Management | Datacenter Monitoring

Driver, Firmware, Security Fixes | Troubleshooting / Diagnostics

Data Access | Job Scheduling, Container Mgmt

Resource Management | Software Upgrades

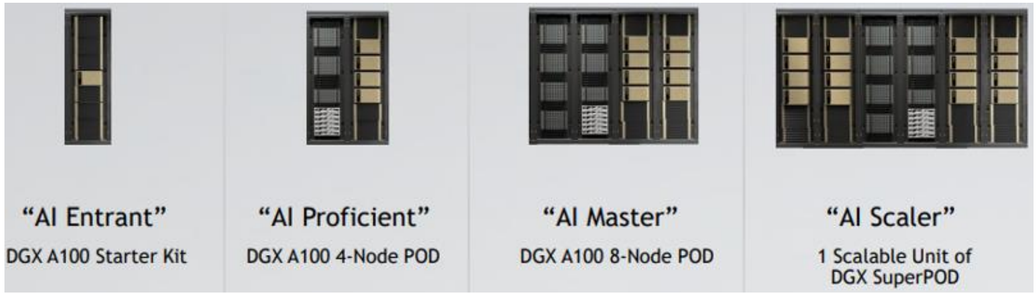

Reference Designs

Simple Deployment | **Simple Management** | **Predictable Performance**

NVIDIA

DGX POD

Datacenter is a new computer

- NVIDIA DGX POD is a reference architecture that incorporates best practices for compute, networking, storage, power, cooling, and more, in an integrated AI infrastructure design built on NVIDIA DGX.
- Delivered as fully integrated, ready-to-deploy offerings through our NVIDIA Partner Network, these solutions make your data center AI deployments simpler and faster for IT that accelerates AI instead of stalling it
- Purpose built and workflow optimized

AI Stack Optimizations from NGC are developed, tested and proven on DGX before any other system

OS Optimizations enable you to keep up with the latest commercial drivers while ensuring stability

Architecture Optimizations across GPU, CPU, RAM, IO, networking and storage deliver the best performance despite similar specs from alternatives

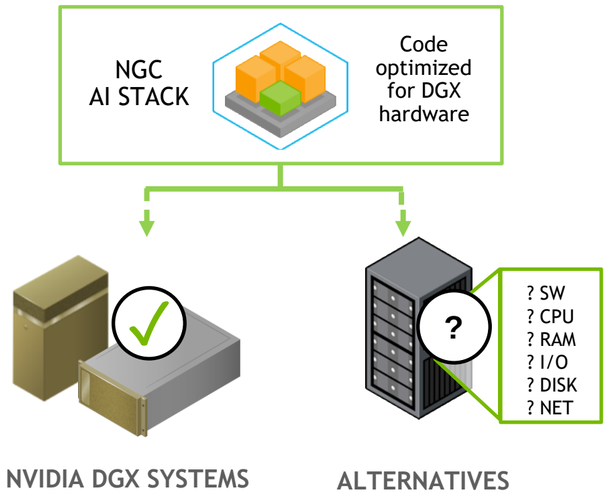

DGX POD VS DGX SUPERPOD

DGX POD

- Storage partner brand
- Multiple storage partners
- Starts with 2/4/8-node configurations
- InfiniBand/RoCE compute, Ethernet or InfiniBand storage
- Software – DeepOps, Customer
- Typical customer- using it as AIaaS majority of work is limited to single node with some work scaling across multiple nodes

DGX SuperPOD

- NVIDIA brand
- Top Storage venders, incl. NetApp as approved storage
- Scalable Unit (20 DGX systems) based design
- Separate InfiniBand fabric for storage and compute
- Base Command Manager software
- Typical customer – center of excellence, large multi-node requirements

Argonne NATIONAL LABORATORY UF UNIVERSITY of FLORIDA WASP

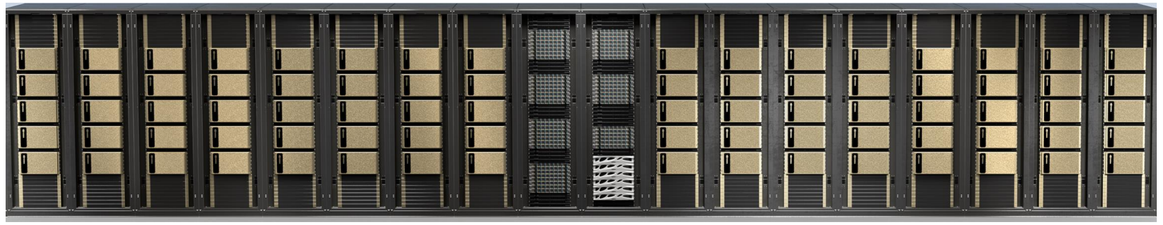

RIGHT SIZING

For reference only

Model Size	Number of GPUS
1.7B	32
3.6B	64
7.5B	128
18B	256
39B	512
76B	1024
145B	1536
310B	1920
530B	2520
1T	3072

GPT models ranging from
1 billion to 1 trillion parameters

~ 4 to 8 DGX A100 POD

~ 1 to 2 SUs DGX A100 SuperPOD

~ 3 to 7 SUs DGX A100 SuperPOD

~ 10 to 20 SUs DGX A100 SuperPOD

Efficient Large-Scale Language Model
Training on GPU Cluster

5 REASONS TO USE DGX POD

Tremendous Value Driving customers to use DGX based solutions

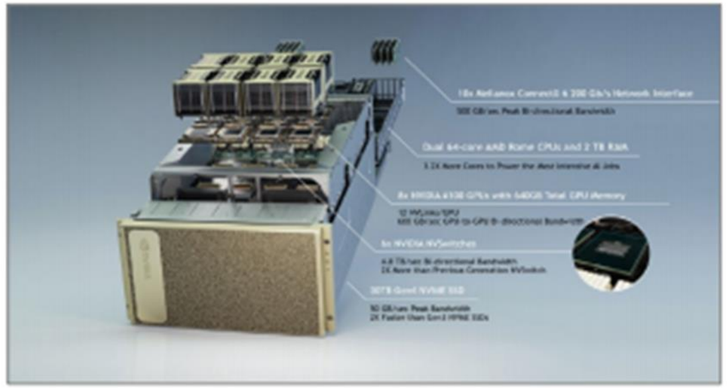

Hardware Optimized for AI

Instead of A Re-Purpose Community Server

QA Hardened Software that Just Works

And Gets Better As Its What NVIDIA Uses

Solutions Speed Deployment

Tested & Scalable Solutions From AI Experts

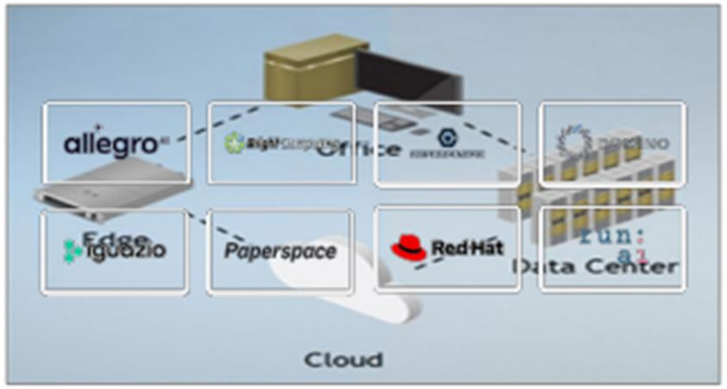

More than a Box

Platform That's Ready to Go

Full-Stack AI Expertise in One Place

The Only Way To Access End-To-End NVIDIA

THE BIG PICTURE

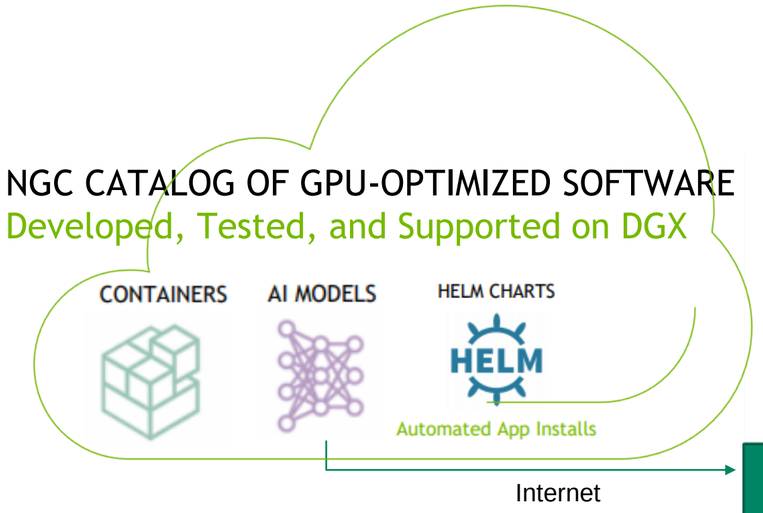

Internet

Considerations to Install in Air-Gapped Clusters

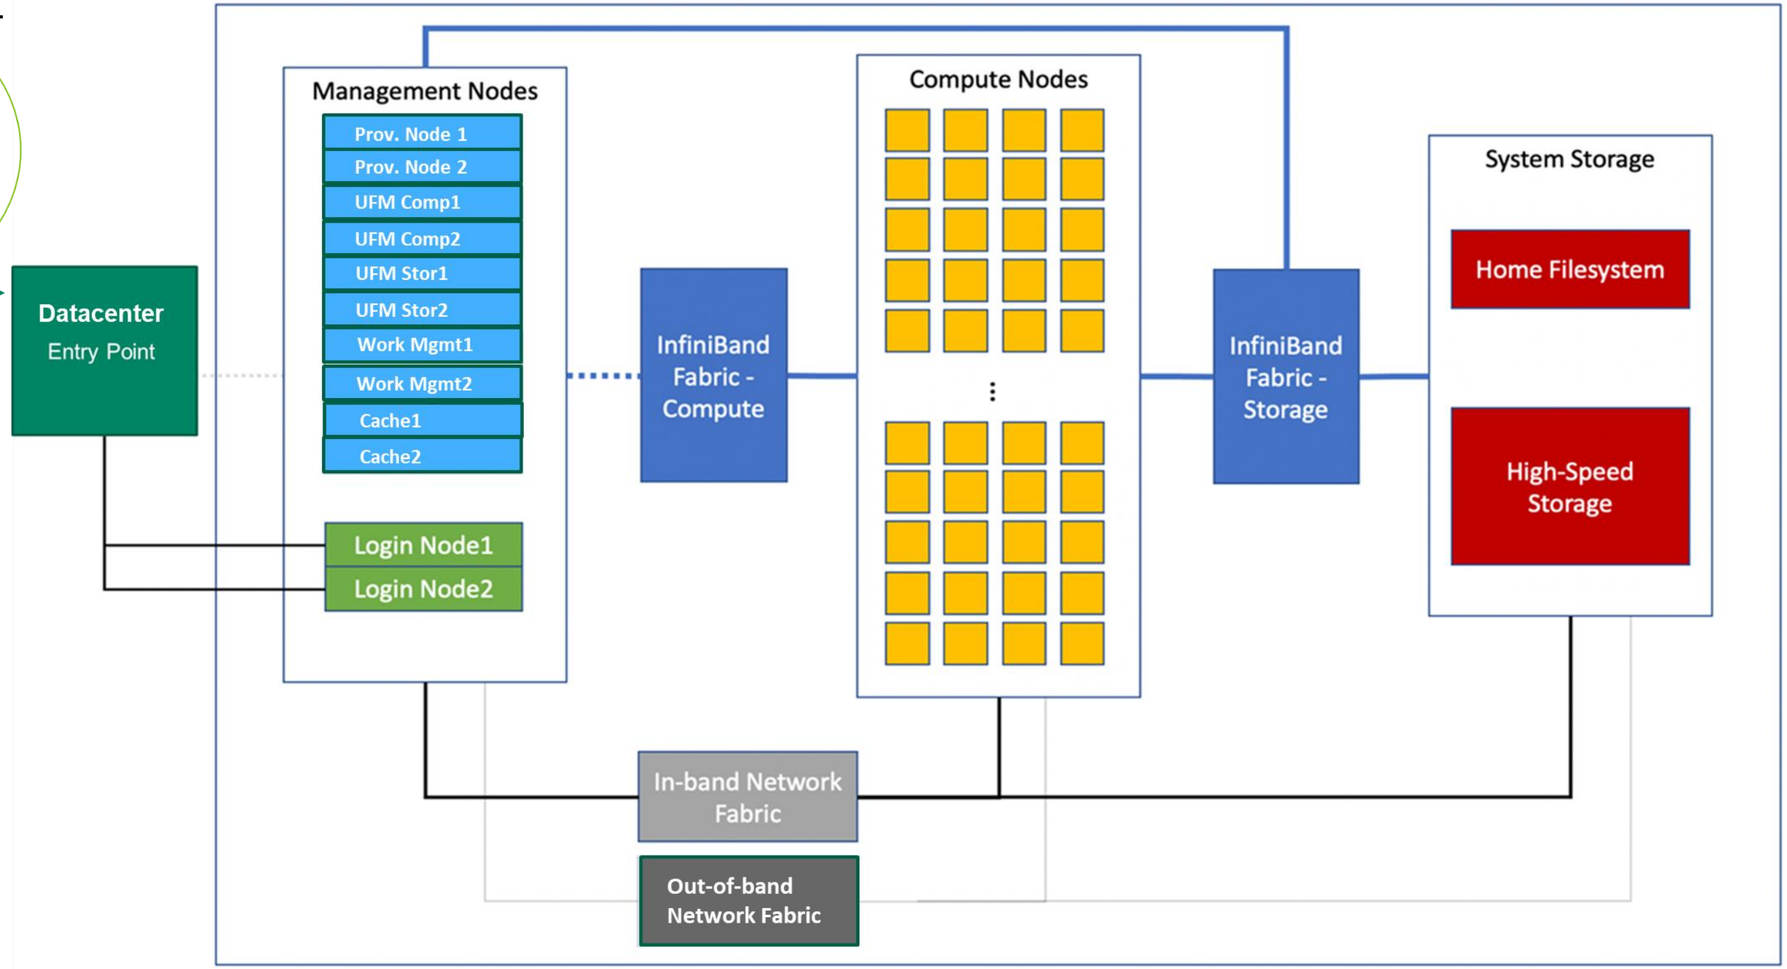

SOLUTION LOGICAL DESIGN

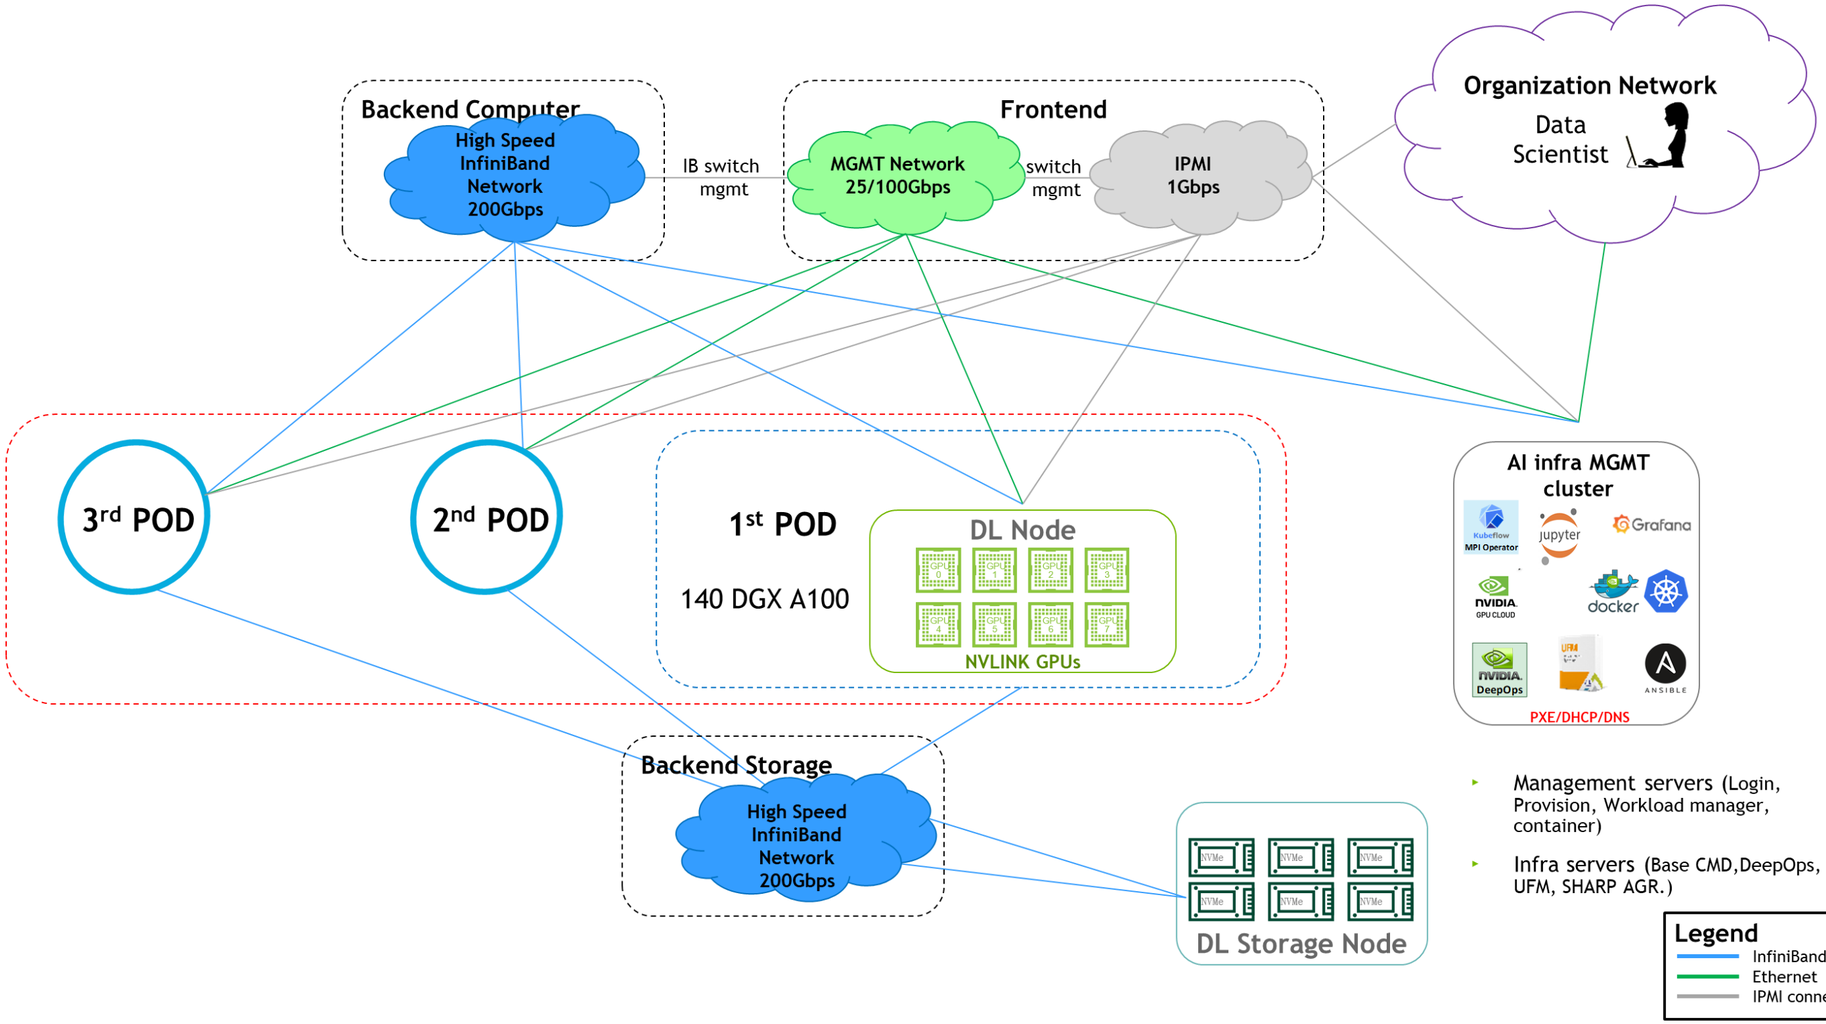

The DGX SuperPOD is a fully-tested system to support all AI and HPC at scale. We don't provide knobs to allow customer to tune architecture for specific workloads. It is a capability-class supercomputer and supports all workloads. -ARB

DESIGNING FOR PERFORMANCE

Key technologies and architecture deep dive

- Contention free communication and maximum inter-GPU reduction bandwidth
- Modular design, Start Small, Scale Huge
- Separate network for Compute vs Storage
- SharpV2 offload and adaptive routing support
- Data center ready and instant production

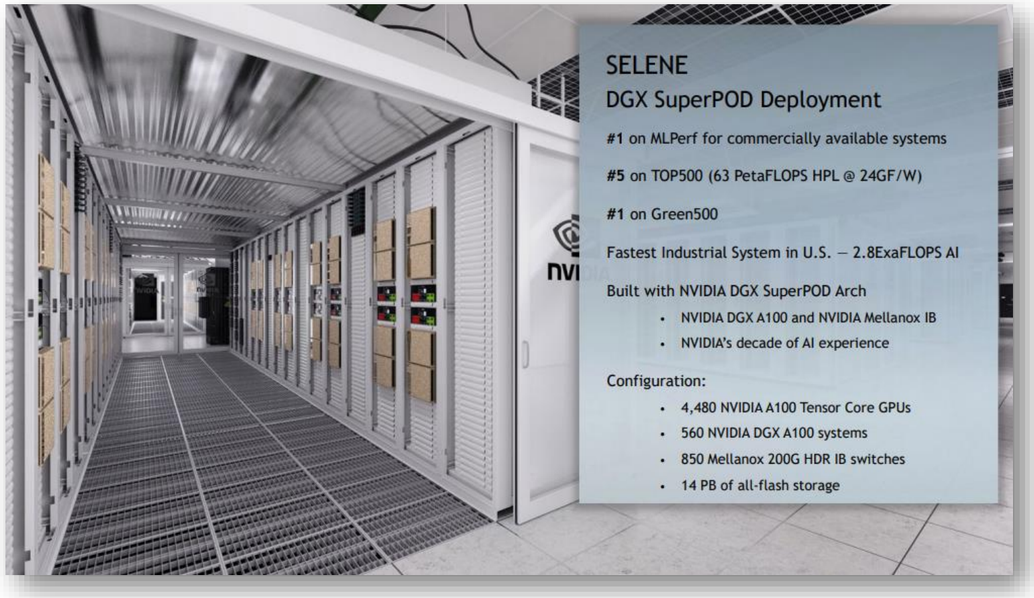

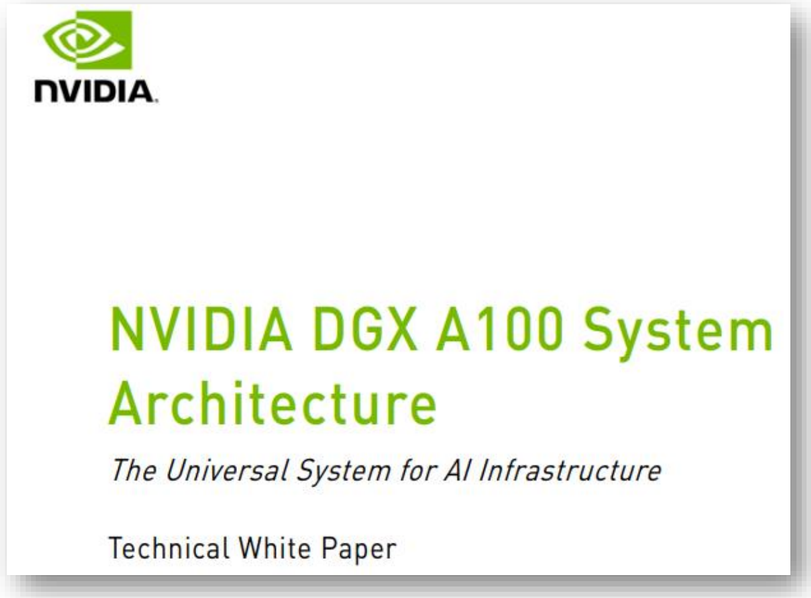

RAIL-OPTIMIZED TOPOLOGY DESIGN

Maximum inter-GPUs reduction bandwidth

Rail-optimized topology design properties:

- Keep rail affinity connectivity.
- Maximize the number of servers that are reachable on each switch hop.
- Reduce NVIDIA SHARP blast radius by enabling a way to run with less rails.
- Maximize reduction capabilities for multi-rail configurations when framework can utilize it.

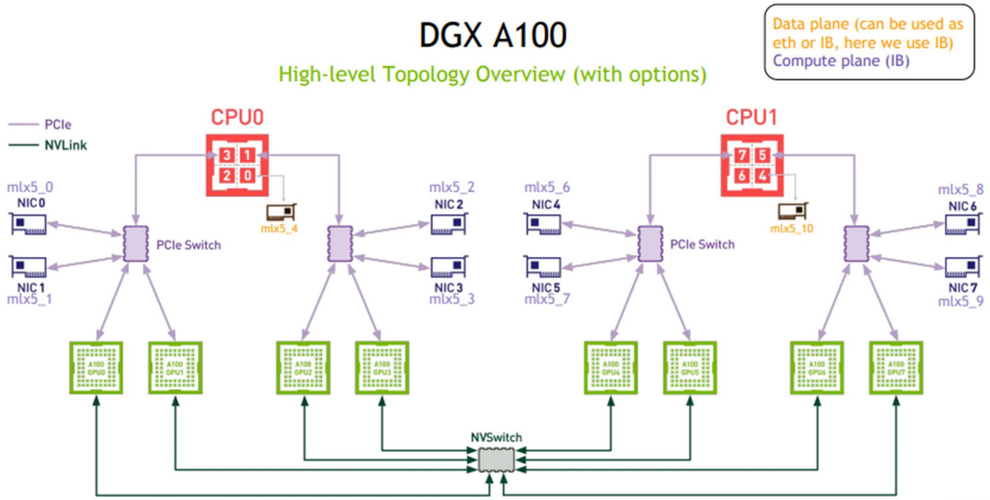

DGX A100 node layout – 8 InfiniBand rails and intra server GPU interconnection based on NVLink

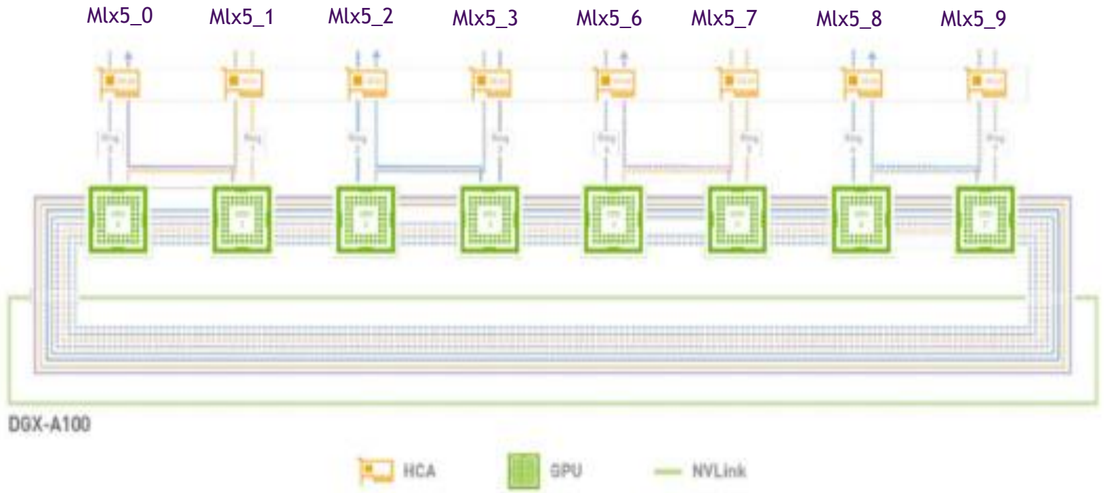

Conversion of intra-node reduction into multi-rail inter-node reduction Vector into 8 pieces

SHARP ACCELERATION

Scalable Hierarchical Aggregation and Reduction Protocol

- **Accelerating HPC applications**
  - Scalable High Performance Collective Offload
  - Collectives - Barrier, All-Reduce
  - Functions - Sum, Min, Max, Min-loc, max-loc, OR, XOR, AND
  - Integer and Floating-Point, 16 / 32 / 64 bit
  - Significantly reduce MPI collective runtime
  - Increase CPU availability and efficiency
  - Enable communication and computation overlap

- **Accelerating Machine Learning applications**
  - Prevents the many-to-one Traffic Pattern
  - Improves performance for AI/ML applications that use streaming
  - High bandwidth reduction for large messages
  - NCCL, GPUDirect RDMA
  - Network Operator & GPU Operator automate GPU Direct RDMA on Kubernetes

LEVERAGE FULL CAPABILITY OF SHARP ACCELERATION

- Hardware offload
- Reduces the amount of data running on the network
- Reduces operation latency
- Scalability

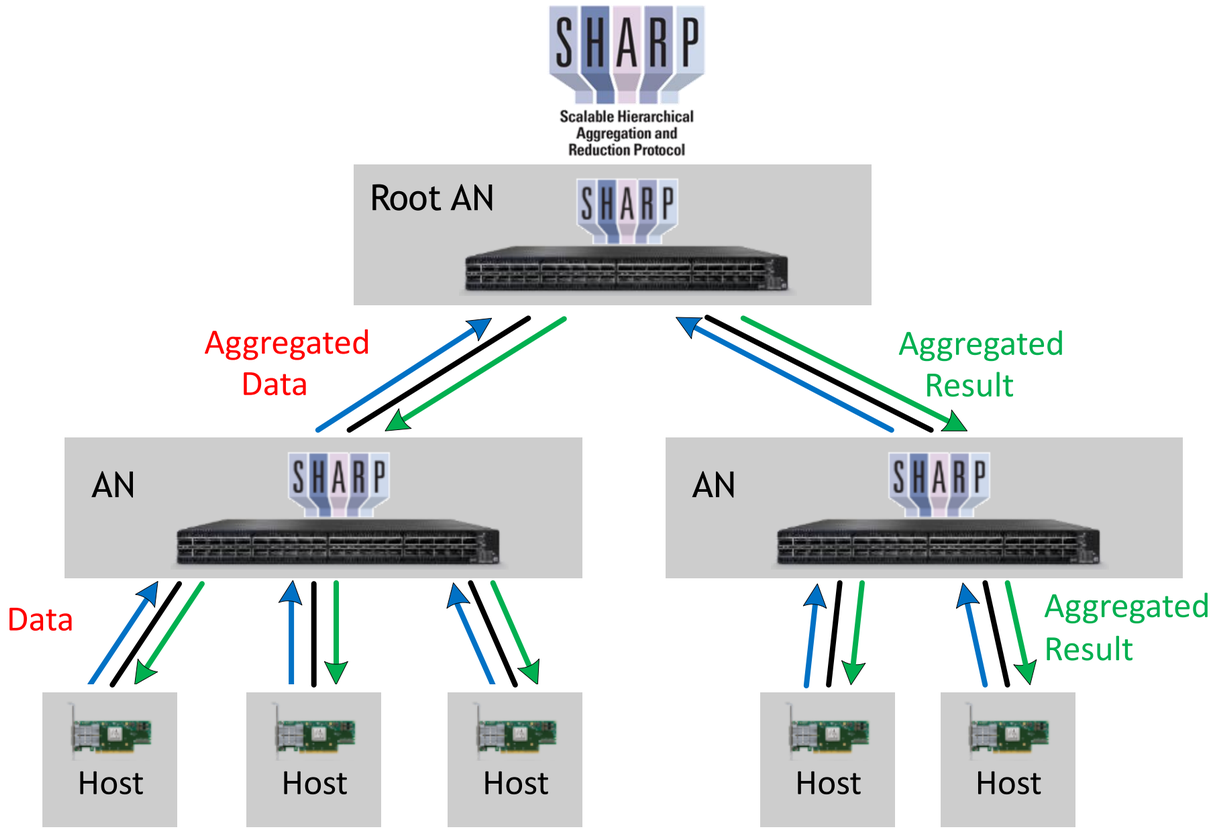

BASIC ALGORITHMS COMPARISON

The following table provides a complexity comparison (number of stages over N end nodes and network usage for a vector with S Bytes) between NVIDIA SHARP and between trees, ring, and reduce-scatter all gather operations.

Complexity Comparison

Algorithm	Sent data per end node	Received data per end node	Complexity
SHARP	S	S	O(1)
Double binary trees	~2S	~2S	Log_n(N)
Ring	~2S	~2S	O(N)
Reduce scatter all gather	~2S	~2S	Log_n(N)

O(1) represents a function that always takes the same take regardless of input size.
O(n) represents the complexity of a function that increases linearly and in direct proportion to the number of inputs.
O(log n) represents a function whose complexity increases logarithmically as the input size increases.

NCCL SHARP INTEGRATION

- Collective network Plugin
- Replace Inter-node tree with SHARP Tree
- Keeps Intra-node ring
- Aggregation in network switch
- Streaming from GPU memory with GPU Direct RDMA
- 2x BW compared to NCCL-TREE

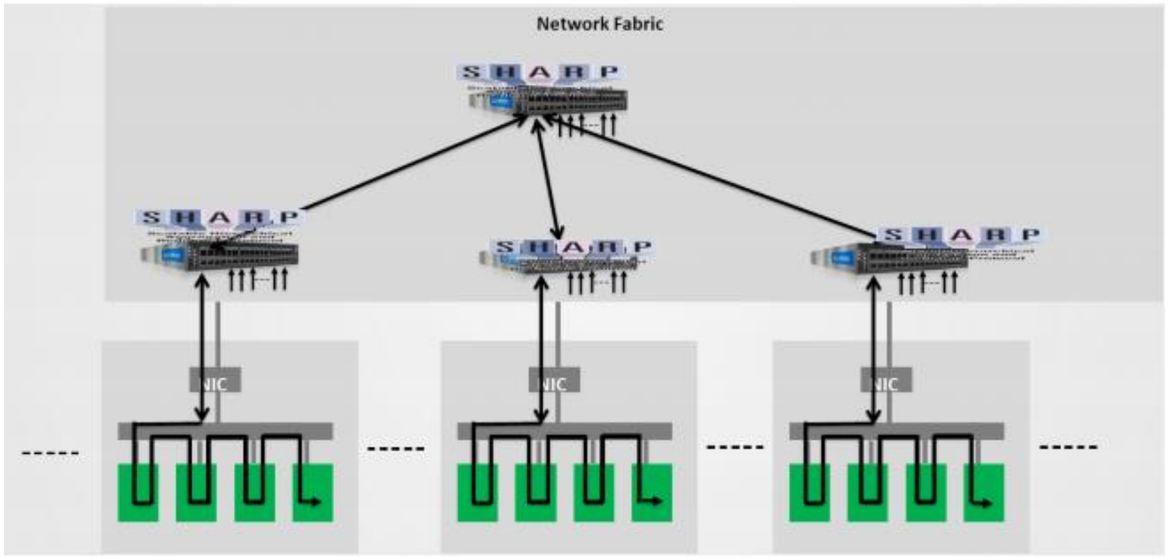

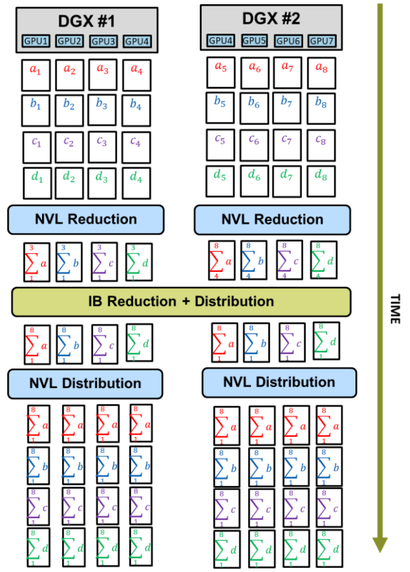

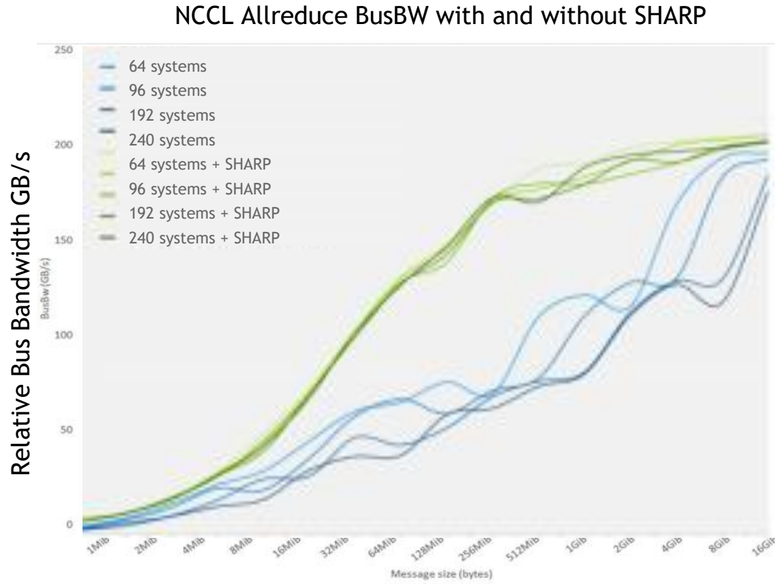

SHARP Enables 2X Higher Data Throughput for NCCL

NGC GPU-OPTIMIZED SOFTWARE HUB

Simplifying DL, ML, and HPC Workflows

Containers
DL, ML, HPC
tuned, tested and optimized

Collections
use-case based curated content in one easy-to-use package

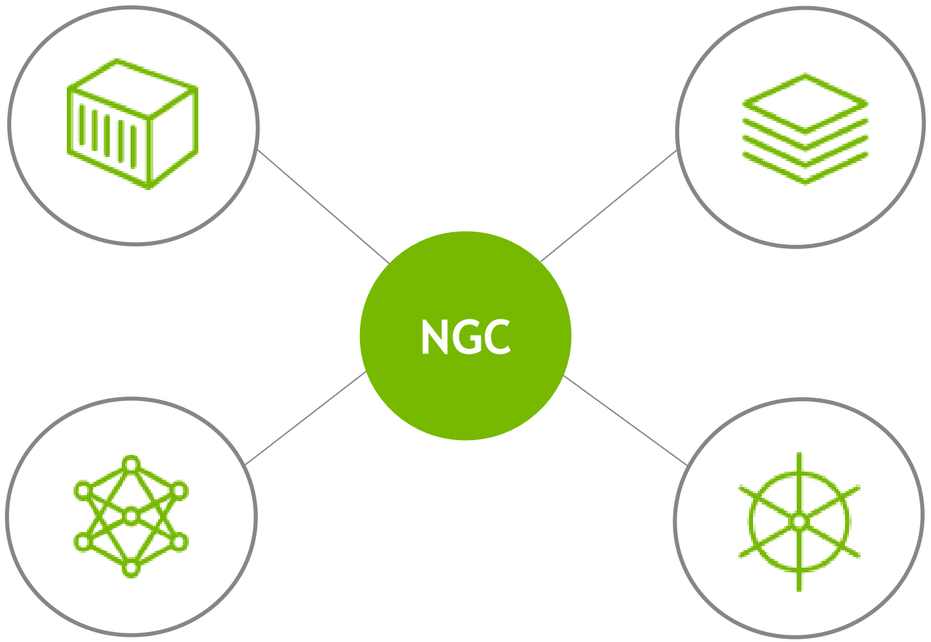

Models and Resources
pre-trained models for a
wide range of common AI tasks

Helm Charts
automate software deployment on Kubernetes clusters

TensorFlow | PyTorch | more

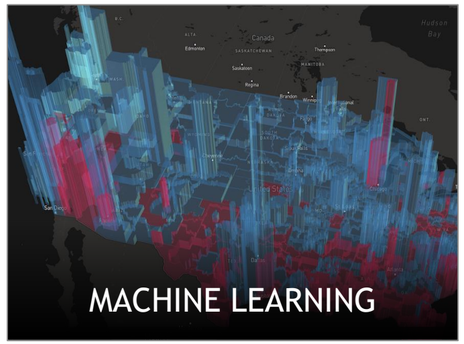

RAPIDS | H2O | more

NAMD | GROMACS | more

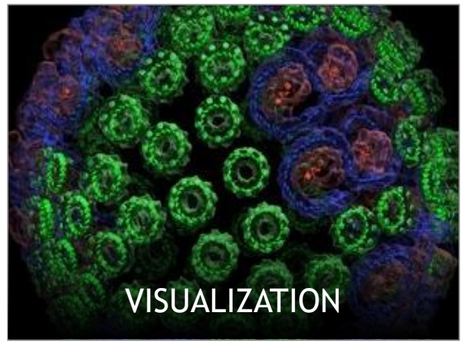

ParaView | IndeX | more

NGC CONTAINERS ACCELERATING WORKFLOWS

Easily Deploy Latest Software, Anywhere

Optimized for Performance

Monthly DL container releases offer latest features and superior performance on NVIDIA GPUs

Scalable Performance

Supports multi-GPU and multi-node systems for scale-out and scale-up environments

Designed for Enterprise and HPC Environments

Supports Docker and Singularity runtimes

Run Anywhere

NVIDIA DGX, PCs, workstations, and OEM servers
On bare metal or in virtual machines
From Core to the Edge
On-Prem to Hybrid to Cloud

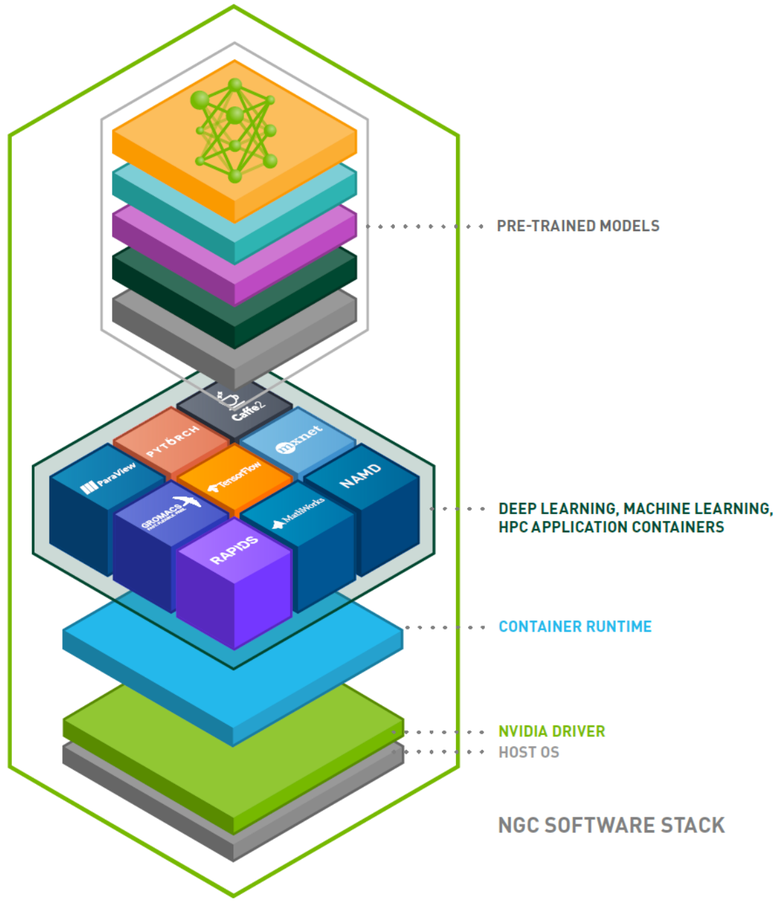

BASE COMMAND MANAGER

Simplify AI Infrastructure Management for Supercomputers

- DGX SuperPOD with Base Command Manager
- DGX SuperPOD moves from a hardware BOM to turnkey solution with software & services
- Predictable performance
- Simplified Infrastructure management
- Scheduling, resource utilization, analytics, etc.

Dashboard / Analytics

Infrastructure Monitoring

Resource Allocation

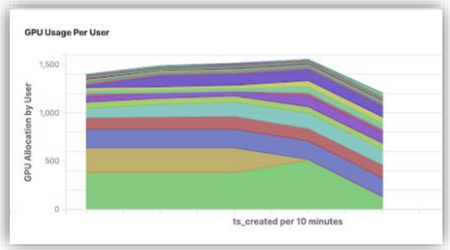

DGX SuperPOD deployment

Infrastructure Management

Deployment | Provisioning | Monitoring | Logging & Alerting | Scheduling

NVIDIA

NVIDIA SELENE

Now Featuring NVIDIA DGX A100 640GB

#6 Top500 | #1 MLPerf | #1 Industrial System

4,480 A100 GPUs

560 DGX A100 systems

850 Mellanox 200G HDR switches

14 PB of high-performance storage

2.8 EFLOPS of AI peak performance

63 PFLOPS HPL @ 24GF/W

CAMBRIDGE-1

Accelerating U.K. healthcare research

- Cambridge-1 ranked No. 41 on the TOP500 list of world's fastest computers.
- Delivers more than 400 petaflops of AI performance and eight petaflops of Linpack performance.
- Based on the NVIDIA DGX SuperPOD reference architecture, the system packs 80 NVIDIA DGX A100 systems
- AI communication frameworks and libraries are bandwidth-sensitive, and they play a critical role in determining application performance
- Using DPU to offload libraries reduces operating system jitter, while dramatically increasing application performance, which is key to enabling a cloud-native supercomputer architecture

SUPERPOD TIMELINE

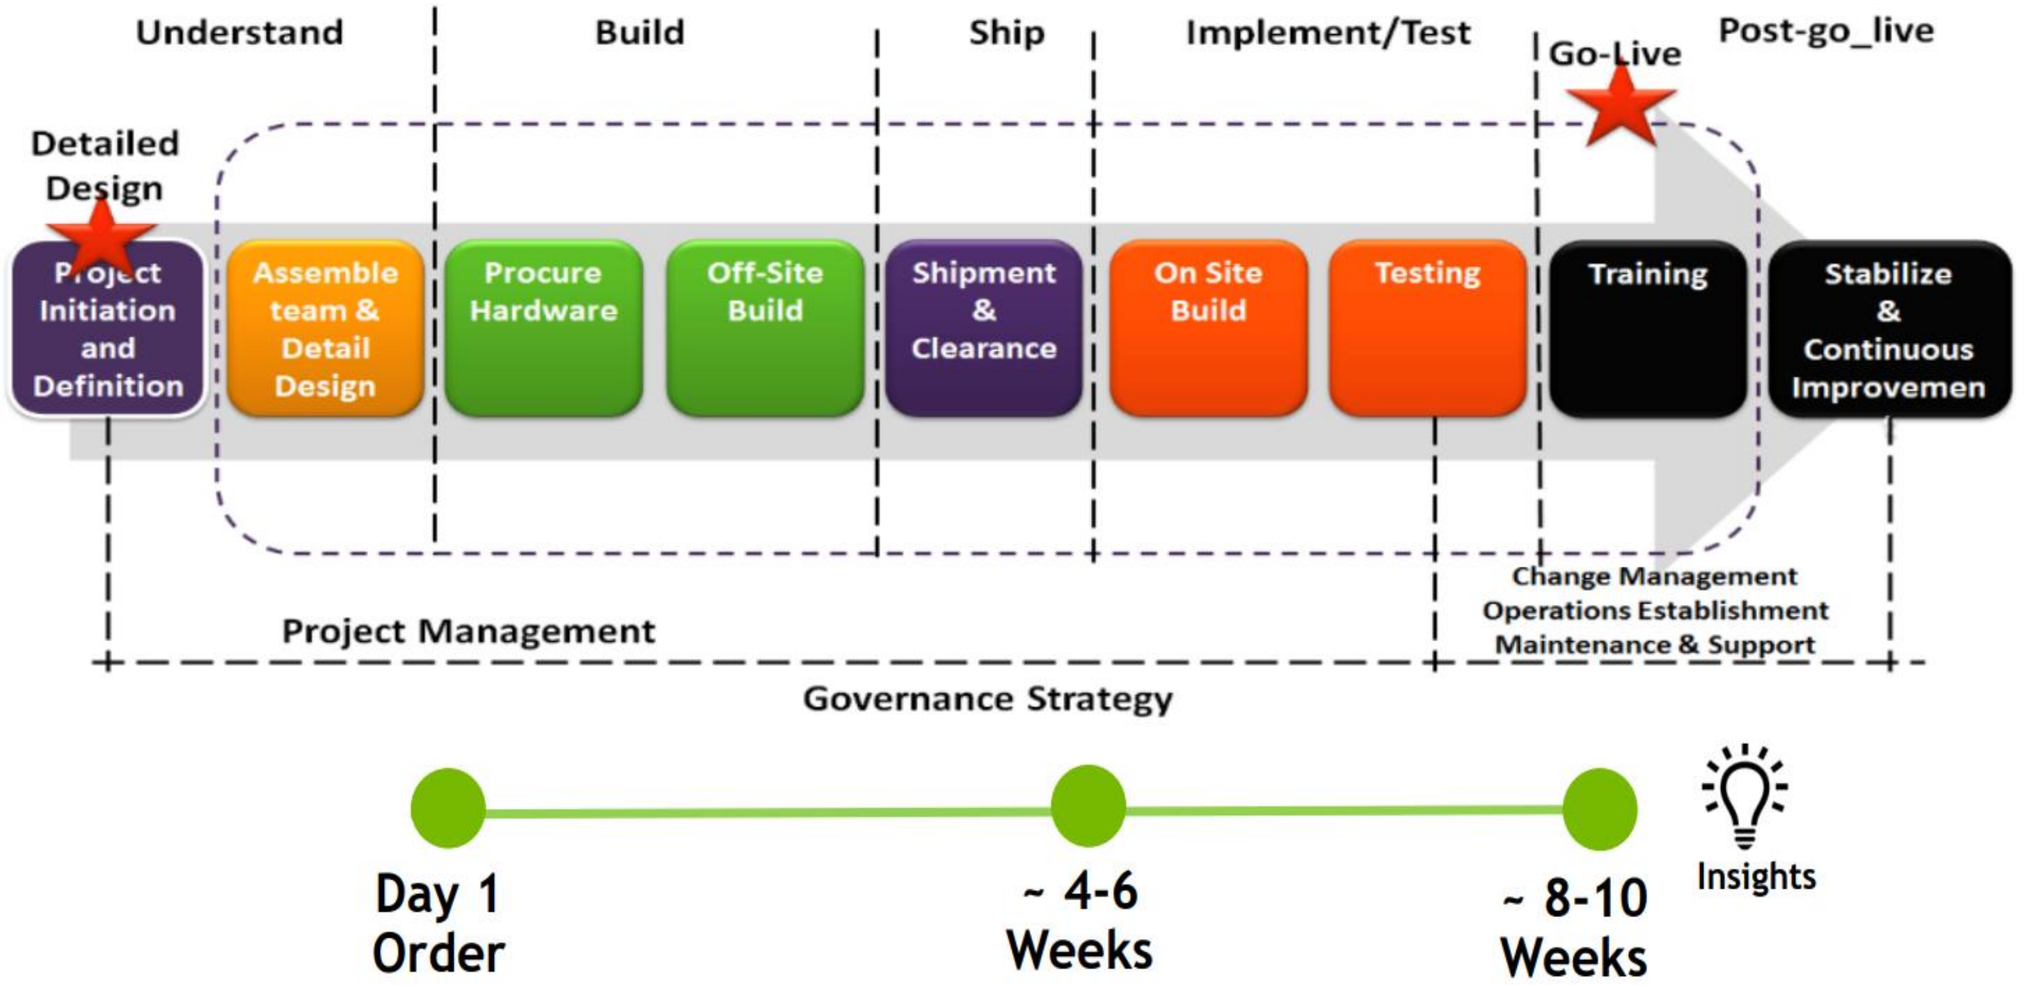

KEY TAKE AWAY- USE US

THE MOST EFFICIENT WAY TO DEPLOY AI

Engage With DGX POD NOW

The Only AI System Designed & Built By NVIDIA

The Art Of The Possible

What is Possible, tell our story proudly

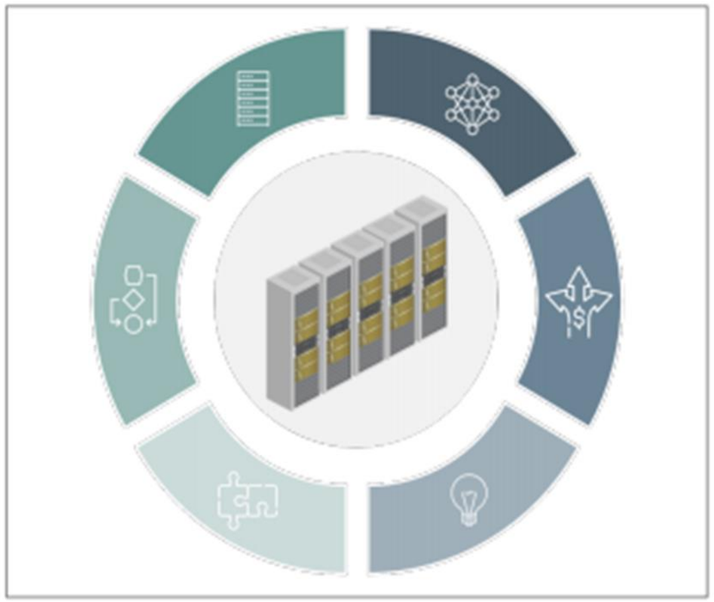

Use The DGX Team, Tools & Ecosystem

Made Life Simple Than Ever with DGX Based Full AI Solutions

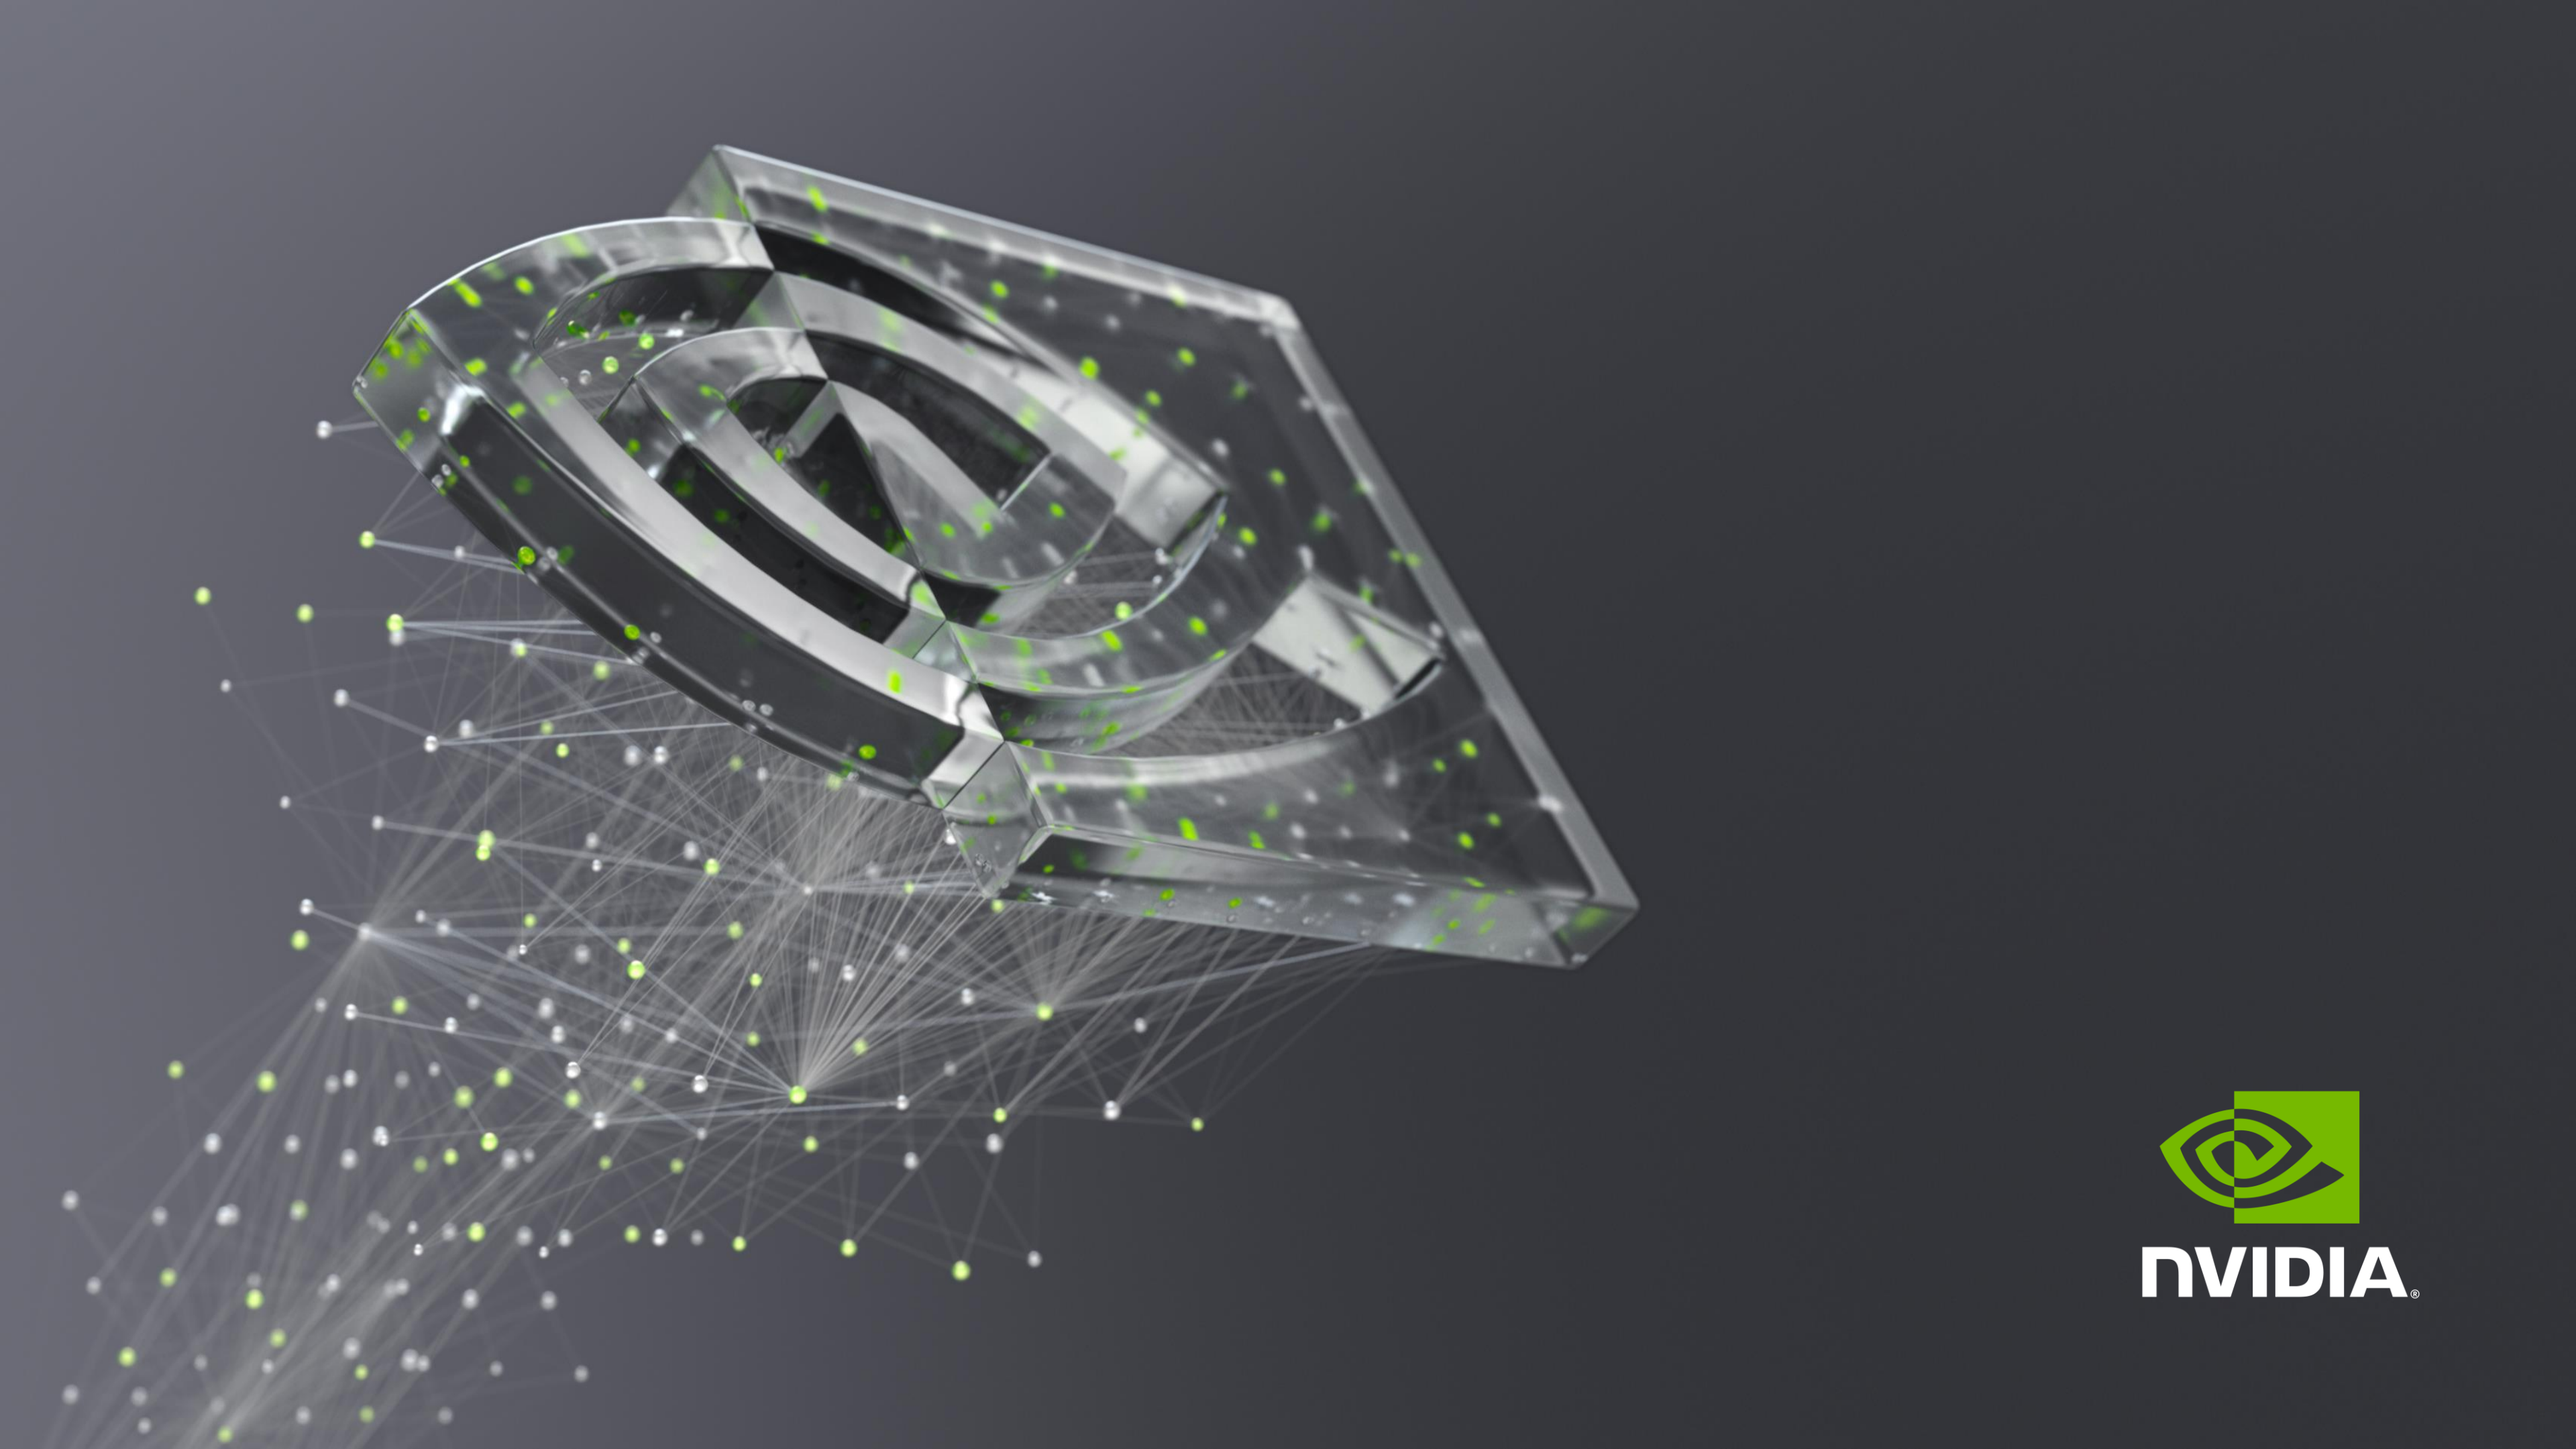
NVIDIA