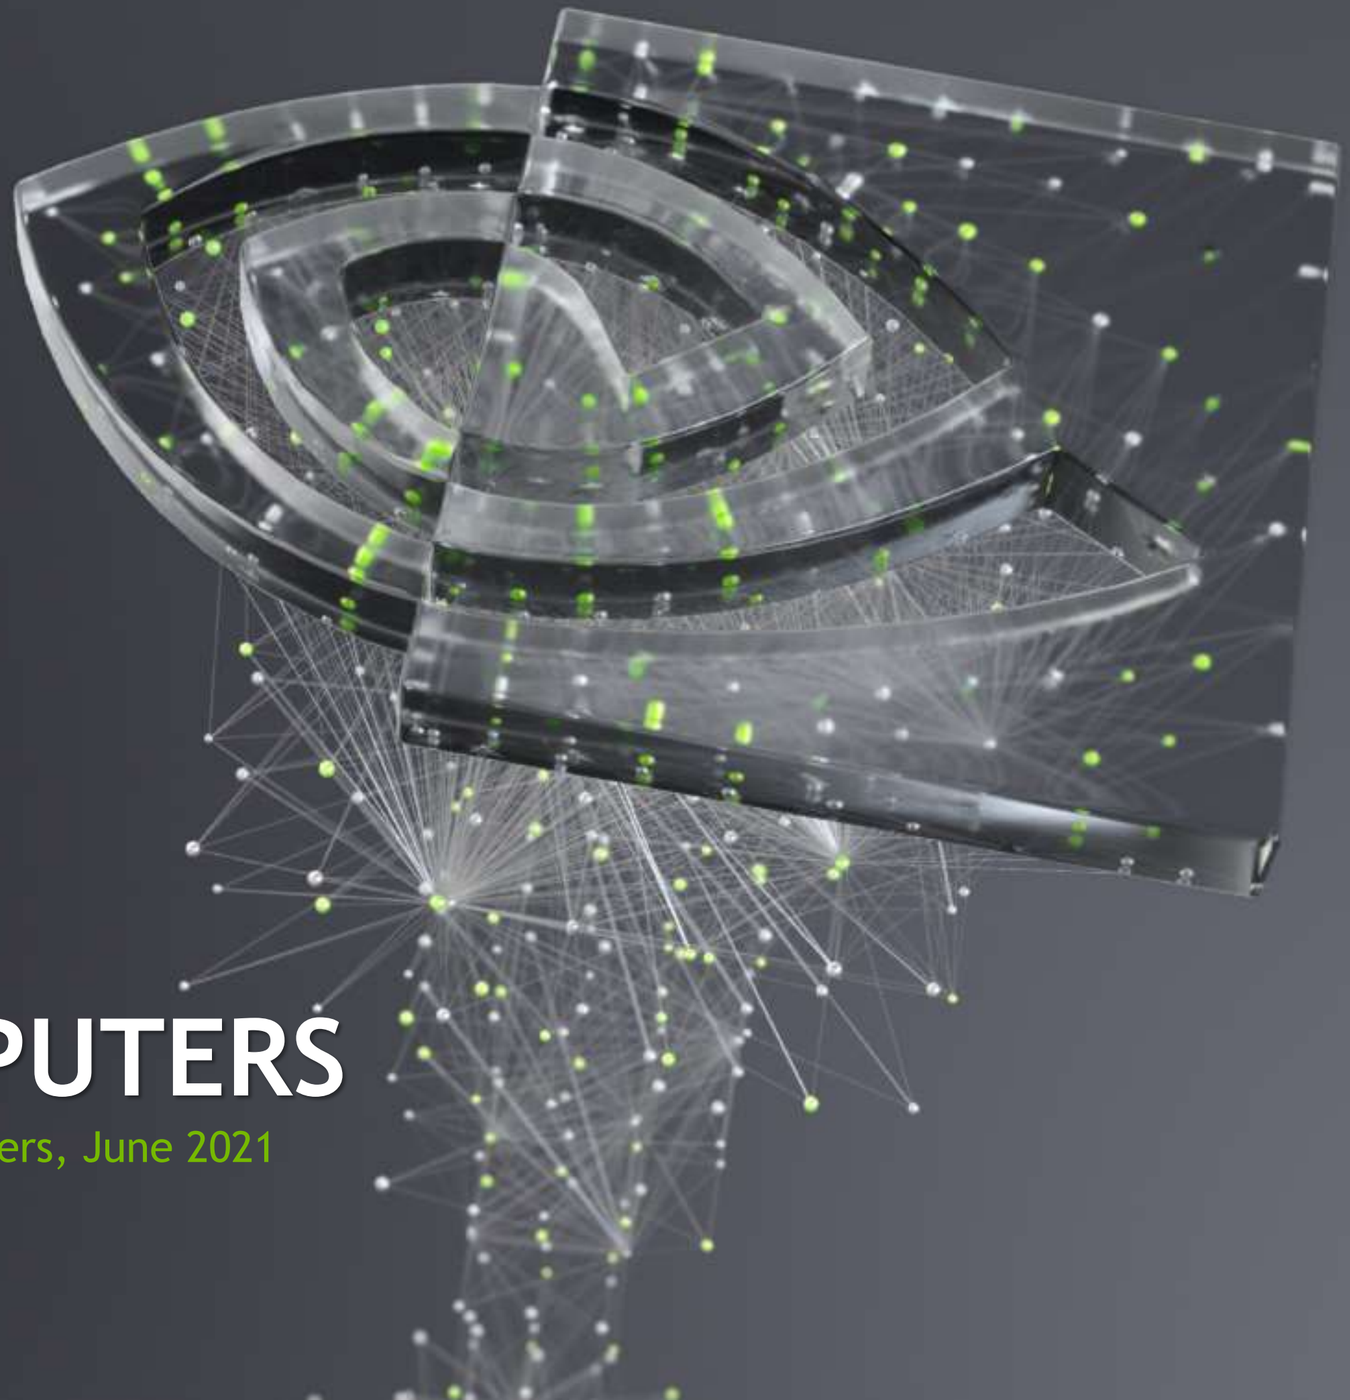

ACCELERATING THE TOP500 SUPERCOMPUTERS

The Interconnect View on the TOP500 Supercomputers, June 2021

NVIDIA INFINIBAND AND ETHERNET ACCELERATE WORLD-LEADING SUPERCOMPUTERS ON THE JUNE'21 TOP500 LIST

NVIDIA GPU or Network (InfiniBand, Ethernet) accelerate 342 systems or 68% of overall TOP500 systems

InfiniBand accelerates seven of the top ten supercomputers in the world

NVIDIA BlueField DPU and HDR InfiniBand Networking accelerate the world's 1st academic cloud-native supercomputer at Cambridge University

NVIDIA InfiniBand and Ethernet networking solutions connect 318 systems or 64% of overall TOP500 platforms

InfiniBand accelerates 170 systems, 21% growth compared to June 2020 TOP500 list

InfiniBand accelerates #1, #2 supercomputers in the US, #1 in China, #1, #2 and #3 in Europe

NVIDIA 25 gigabit and faster Ethernet solutions connect 62% of total Ethernet systems

POWERING NEXT WAVE OF ACCELERATED AI SUPERCOMPUTING

CAMBRIDGE-1
Cloud-Native Supercomputing

TESLA AUTOMOTIVE
Industrial AI Supercomputer

RECURSION
Computational Biology

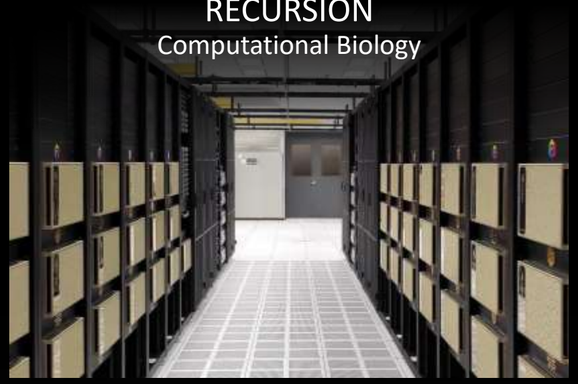

DGX SUPERPODS
Industrial, Cloud-Native AI

MICROSOFT AZURE
On-demand Supercomputing

HPC ADOPTS MIXED PRECISION
10+ Systems on HPL-AI

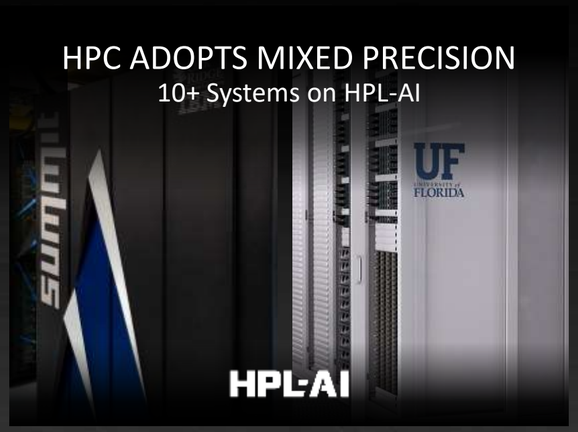

SUPERCLOUDS: CLOUD-NATIVE SUPERCOMPUTERS

Secured and Accelerated by NVIDIA BlueField DPU | Multi-Tenant Bare-Metal Performance

BARE-METAL PERFORMANCE

MULTI TENANCY

EDGE COMPUTING

University of Cambridge Wilkes-3

World’s first academic Cloud-Native Supercomputing

UNIVERSITY OF CAMBRIDGE

NVIDIA SELENE

Featuring NVIDIA DGX A100 640GB

#6 TOP500 | #1 MLPerf | #1 Industrial System

4,480 A100 GPUs

560 DGX A100 system

850 InfiniBand 200G HDR switches

14PB of high-performance storage

2.8 EFLOPS of AI peak performance

63 PFLOPS HPL @ 24GF/W

NVIDIA COMPUTER OR NETWORK ACCELERATE 68% TOP500 SYSTEMS

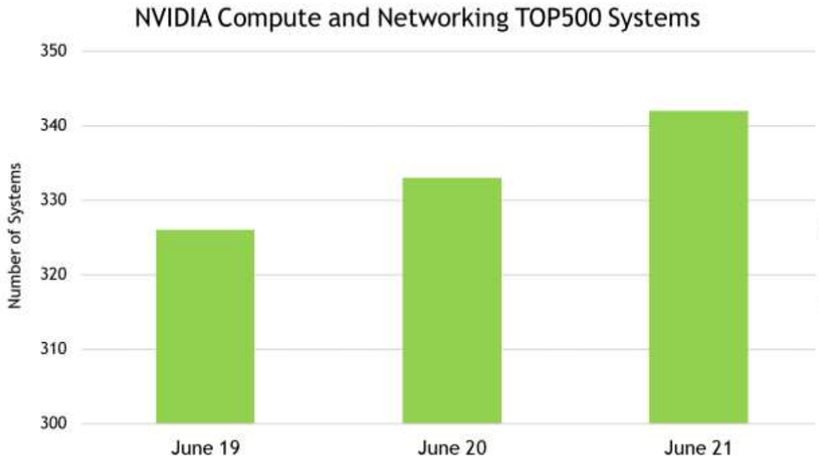

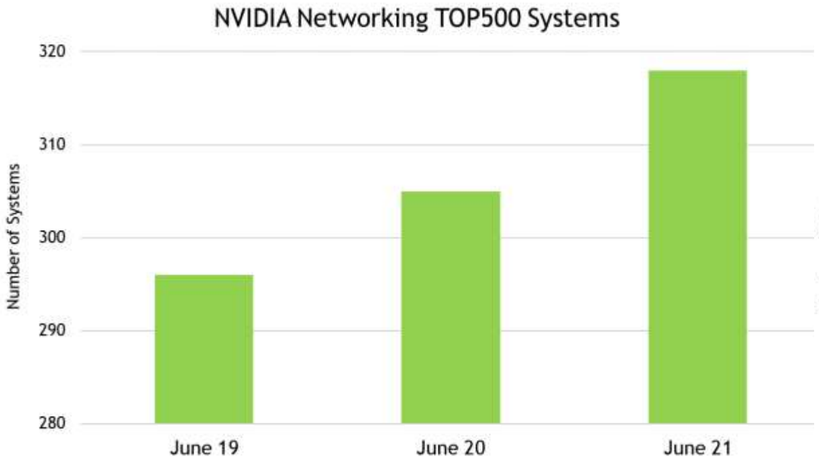

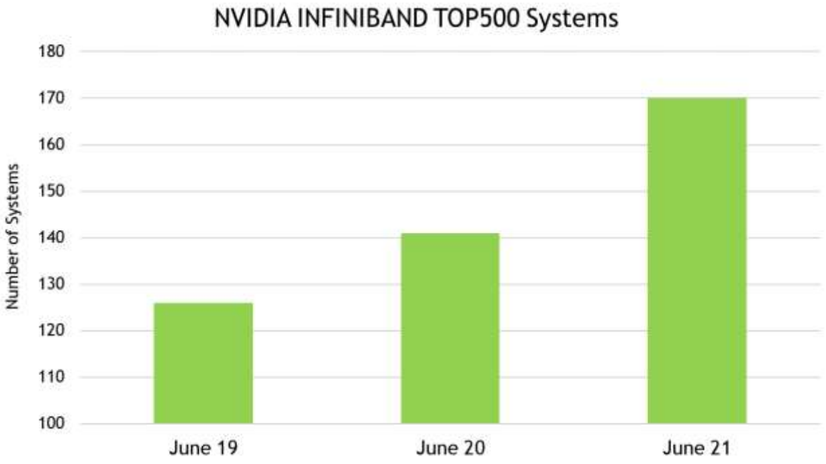

Seven of the world’s top 10 supercomputers use NVIDIA InfiniBand

NVIDIA compute or network accelerates 342 systems on the TOP500 list

NVIDIA InfiniBand and Ethernet networking accelerate 318 systems on the TOP500 list

TOP500 INTERCONNECT TRENDS

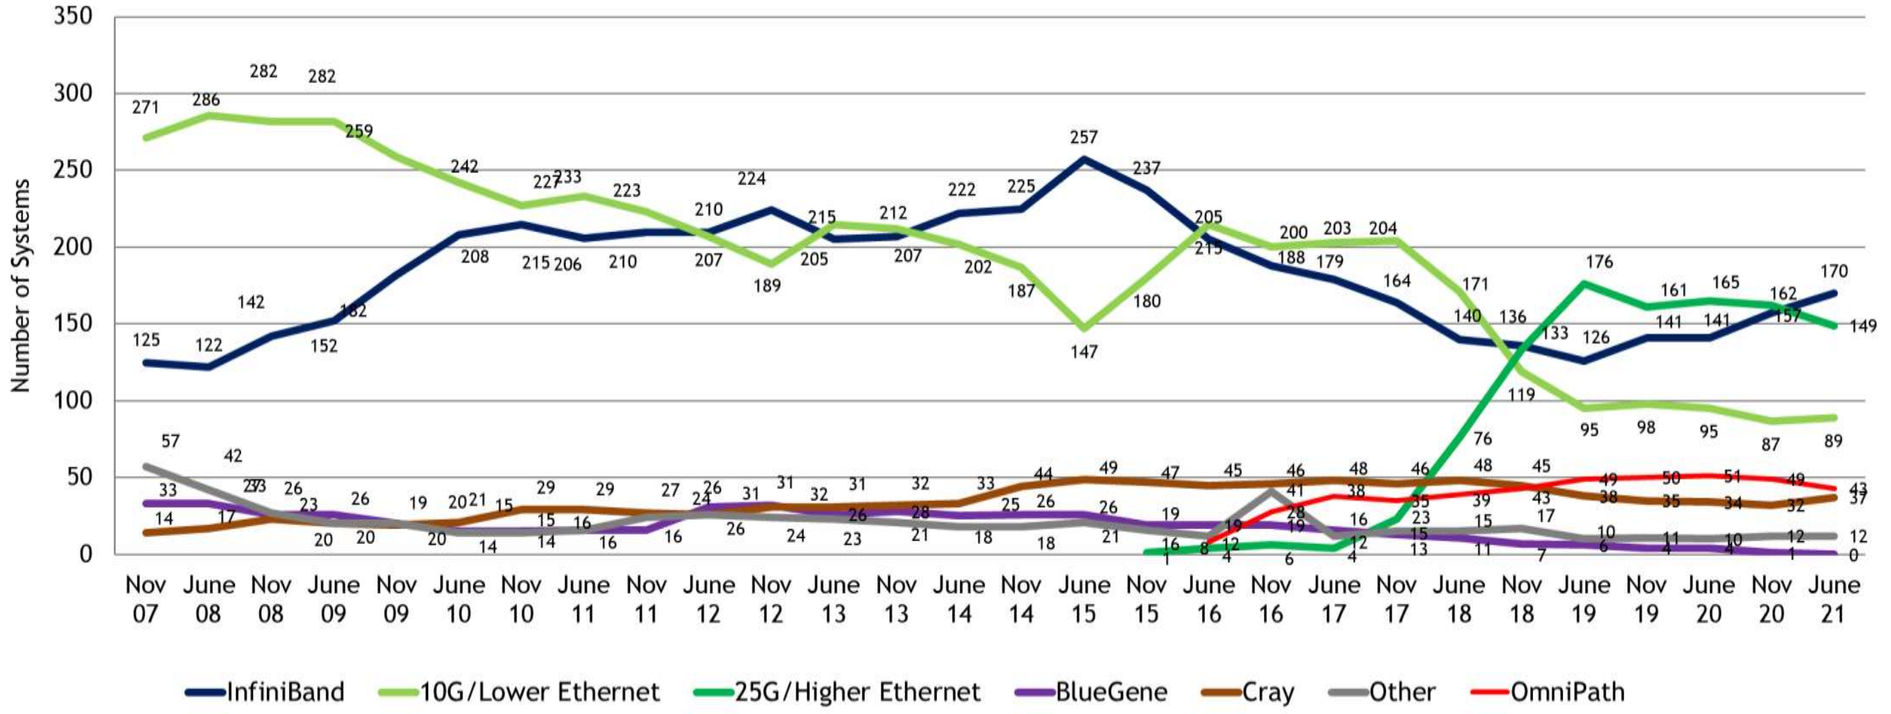

The TOP500 list includes both HPC and Cloud / Hyperscale platforms

InfiniBand continues its leadership as the most used interconnect solution for HPC platforms

NVIDIA connects most of 25G and faster Cloud and Hyperscale Ethernet systems

TOP500 NVIDIA INTERCONNECT TREND

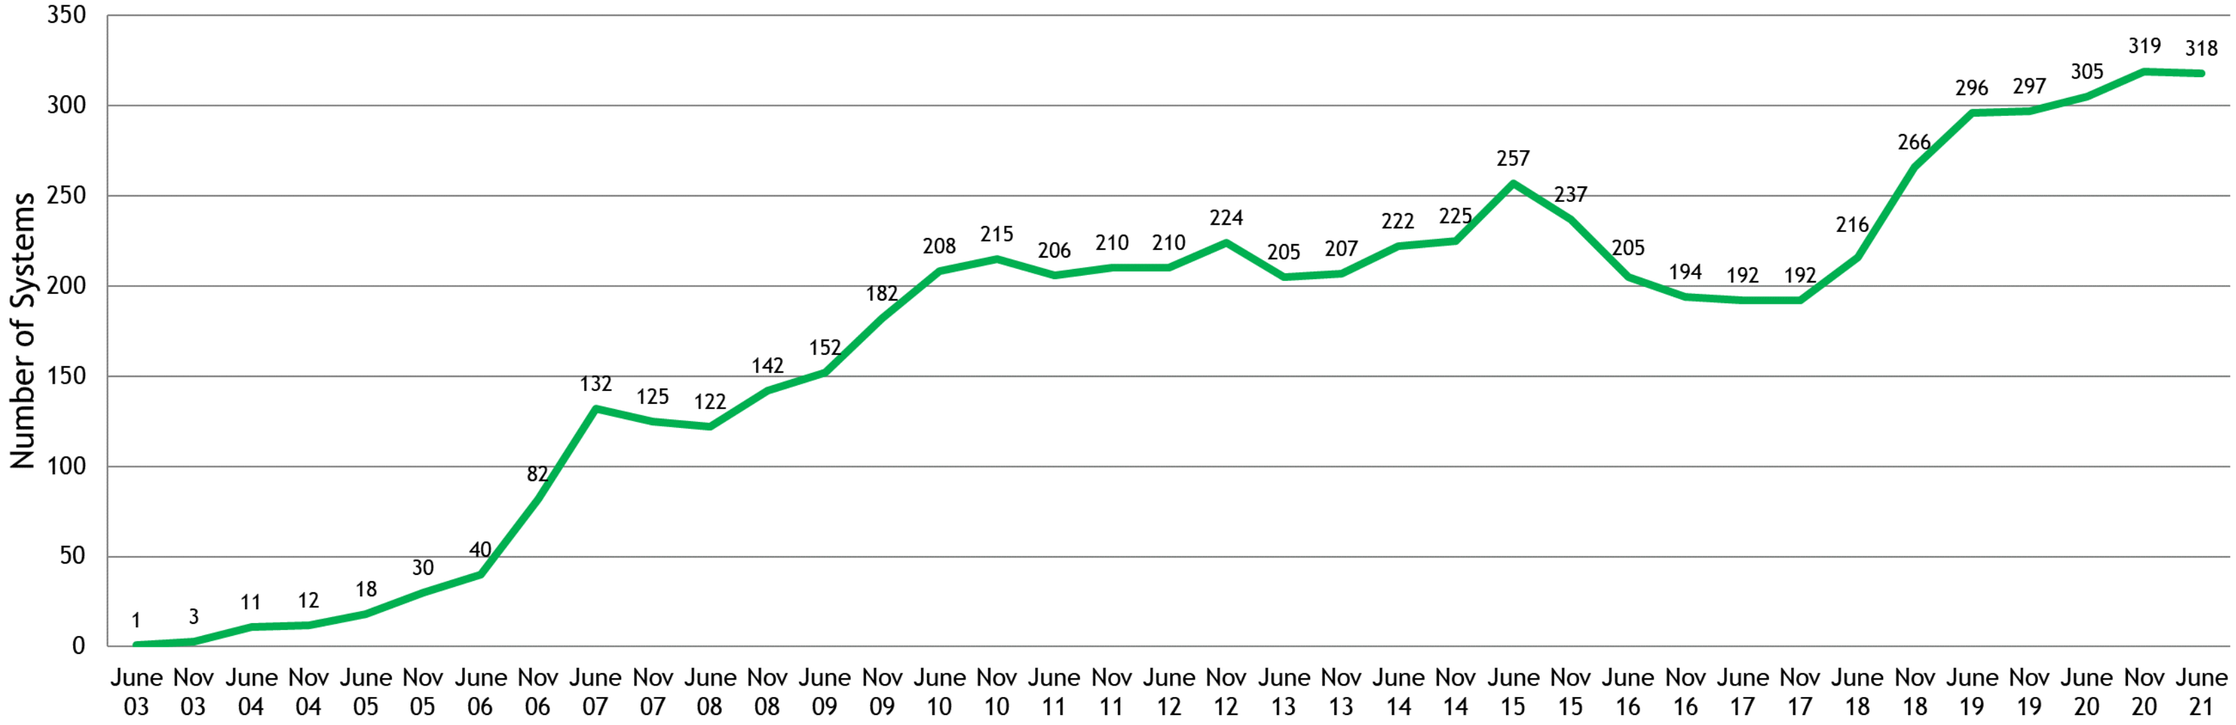

NVIDIA InfiniBand and Ethernet Interconnect Solution Accelerate Majority of TOP500 Platforms

INTERCONNECTS - TOP10 SUPERCOMPUTERS

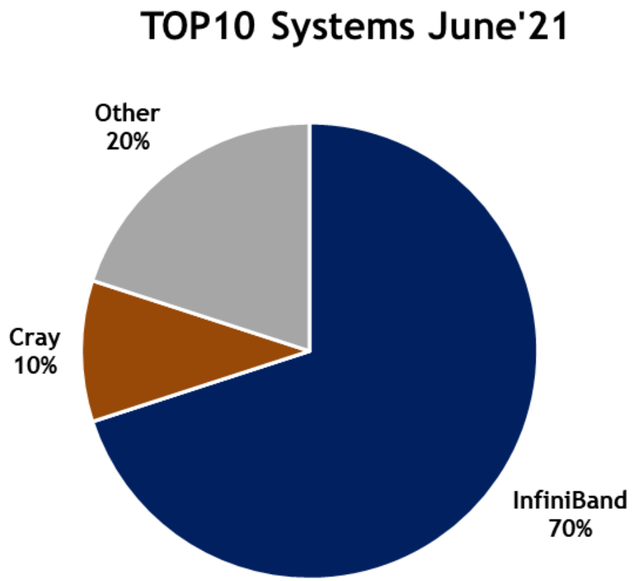

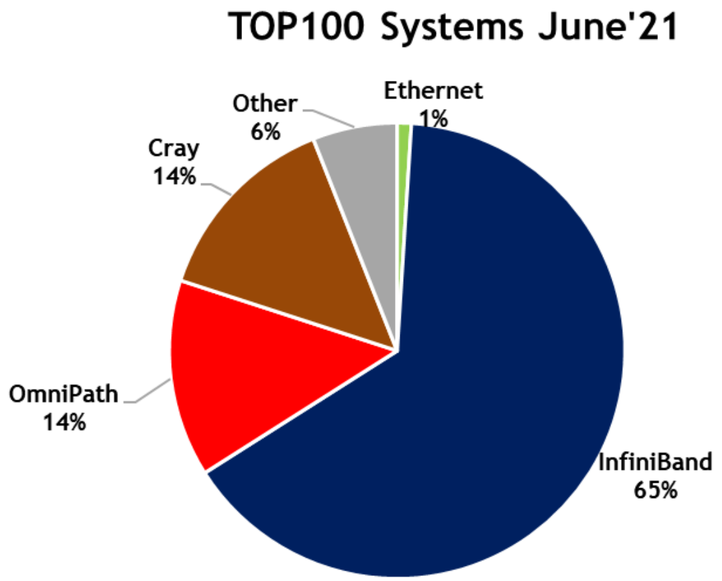

InfiniBand accelerates 7 of Top 10 supercomputers

InfiniBand accelerates 65 of top 100 supercomputers, 23% increase year-over-year (June’20 to June’21)

INTERCONNECTS – TOP100, 200, 300, 400, 500 DISTRIBUTION

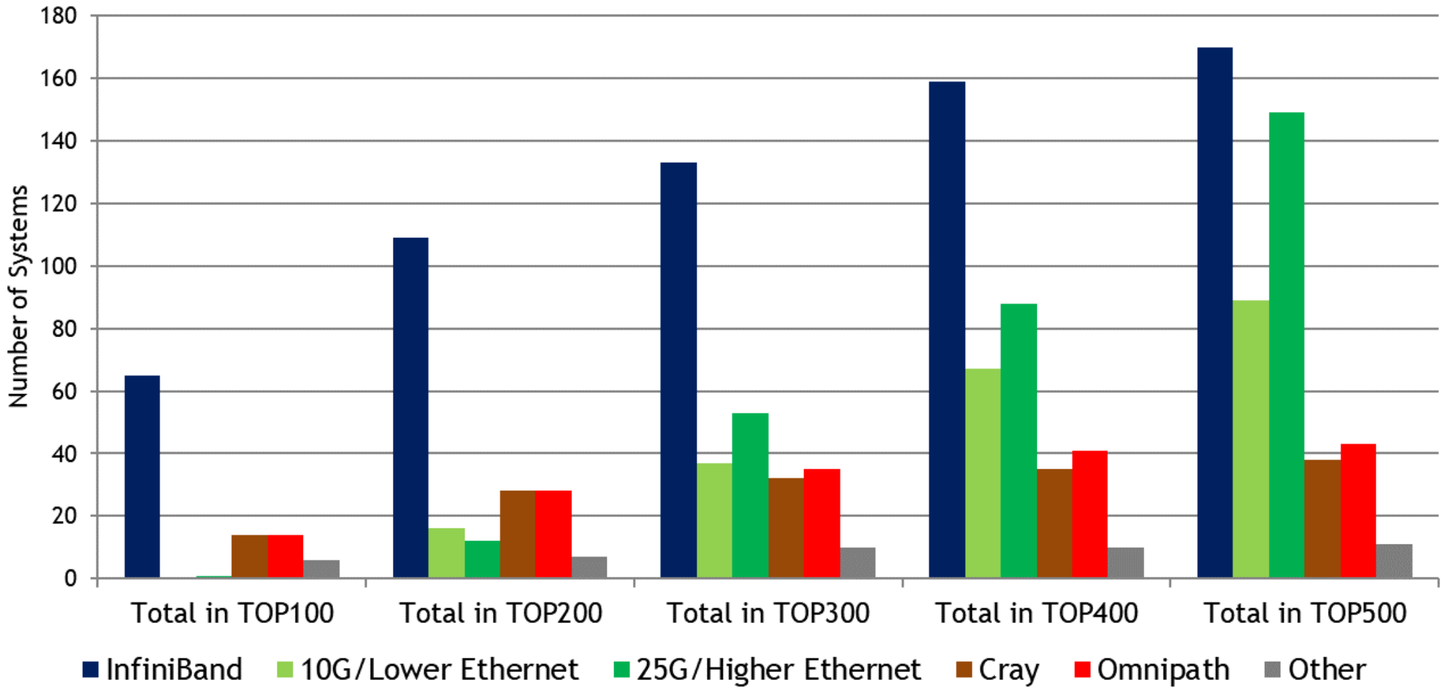

InfiniBand is the most used high-speed interconnect of top supercomputers

Superior performance, scalability, efficiency and return-on-investment

NVIDIA NDR 400G INFINIBAND: NEXT-GENERATION INFINIBAND ARCHITECTURE

ConnectX-7 Adapter

NDR 400G InfiniBand
In-Network Computing

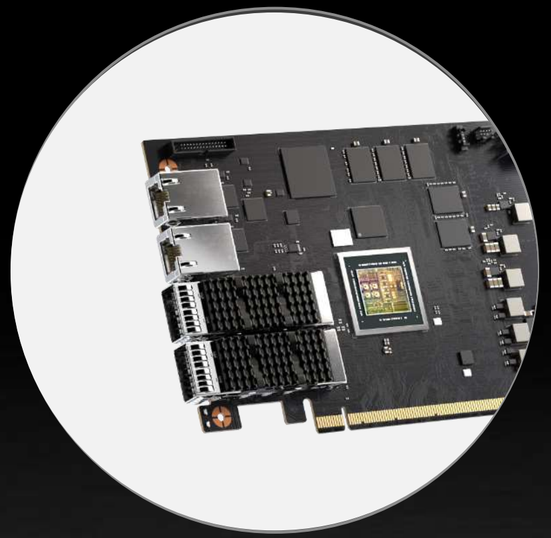

BlueField-3 DPU

NDR 400G InfiniBand with Arm Cores
AI Application Accelerators
Programmable Datapath
In-Network Computing

Quantum-2 Switch

64-ports NDR 400G InfiniBand
128-ports 200G NDR200
In-Network Computing

Cable

Copper Cables
Active Copper Cables
Optical Transceivers

ANNOUNCING NVIDIA NDR 400G INFINIBAND SYSTEMS

In-Network Computing Accelerated Network for Cloud-Native Supercomputing at Any Scale

2x
Data Throughput
400 Gigabits per Second

32x
More AI Acceleration
NVIDIA SHARP In-Network Computing

6.5x
Higher Scalability
>1M nodes with 3 hops

5x
Switch System Capacity
>1.6 Petabit per Second

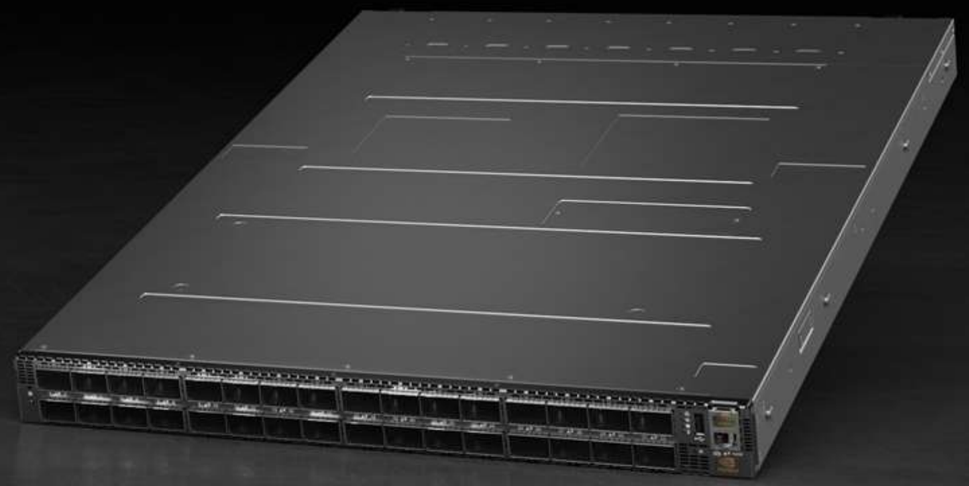

ANNOUNCING NVIDIA NDR 400G INFINIBAND SYSTEMS

In-Network Computing Accelerated Network for Cloud-Native Supercomputing at Any Scale

15%
Faster Deep Learning
Recommendations

17%
Faster Natural
Language Processing

15%
Faster Computational
Fluid Dynamics Simulations

60%
Lower
Power Consumption

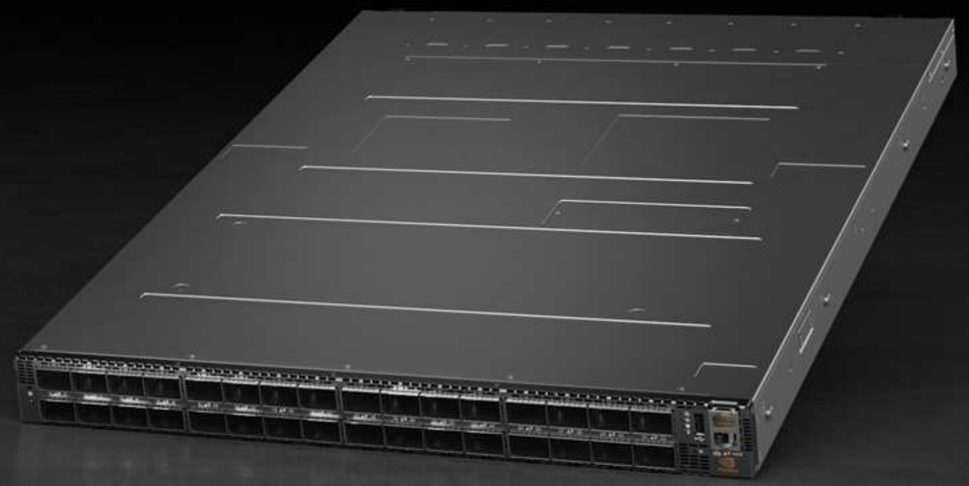

ANNOUNCING NVIDIA NDR 400G INFINIBAND SYSTEMS

In-Network Computing Accelerates Cloud-Native Supercomputing at Any Scale

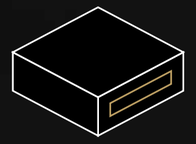

64 NDR Ports
128 NDR200 Ports

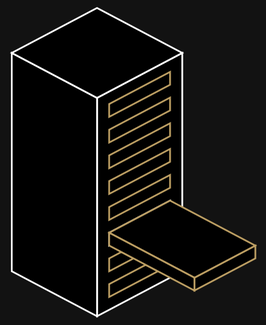

512 NDR Ports
1024 NDR200 Ports

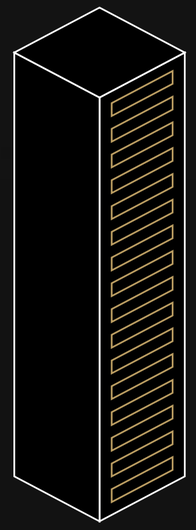

1024 NDR Ports
2048 NDR200 Ports

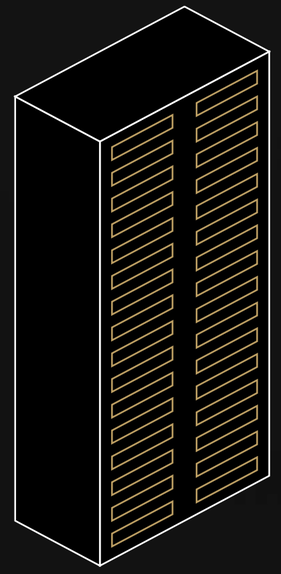

2048 NDR Ports
4096 NDR200 Ports

FROM SUPERCOMPUTERS TO SUPERCLOUDS: CLOUD-NATIVE SUPERCOMPUTERS

DOCA ENABLING GROWING PARTNER ECOSYSTEM

BLUEFIELD-3 NEXT GENERATION 400G DATA CENTER INFRA PROCESSOR

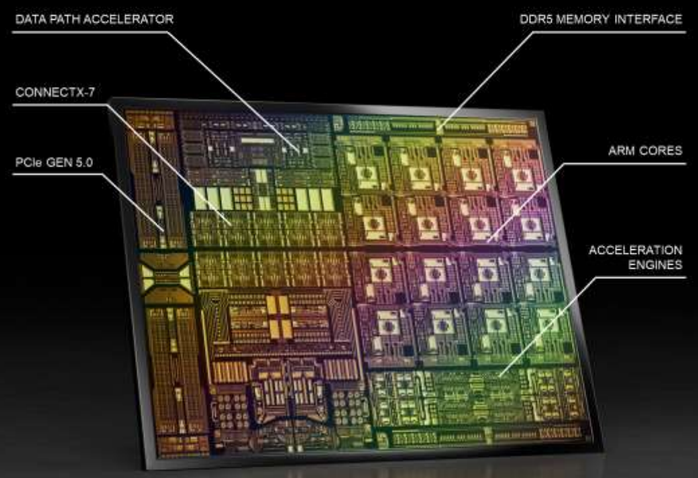

NVIDIA NDR 400G INFINIBAND IN-NETWORK COMPUTING INTERCONNECT

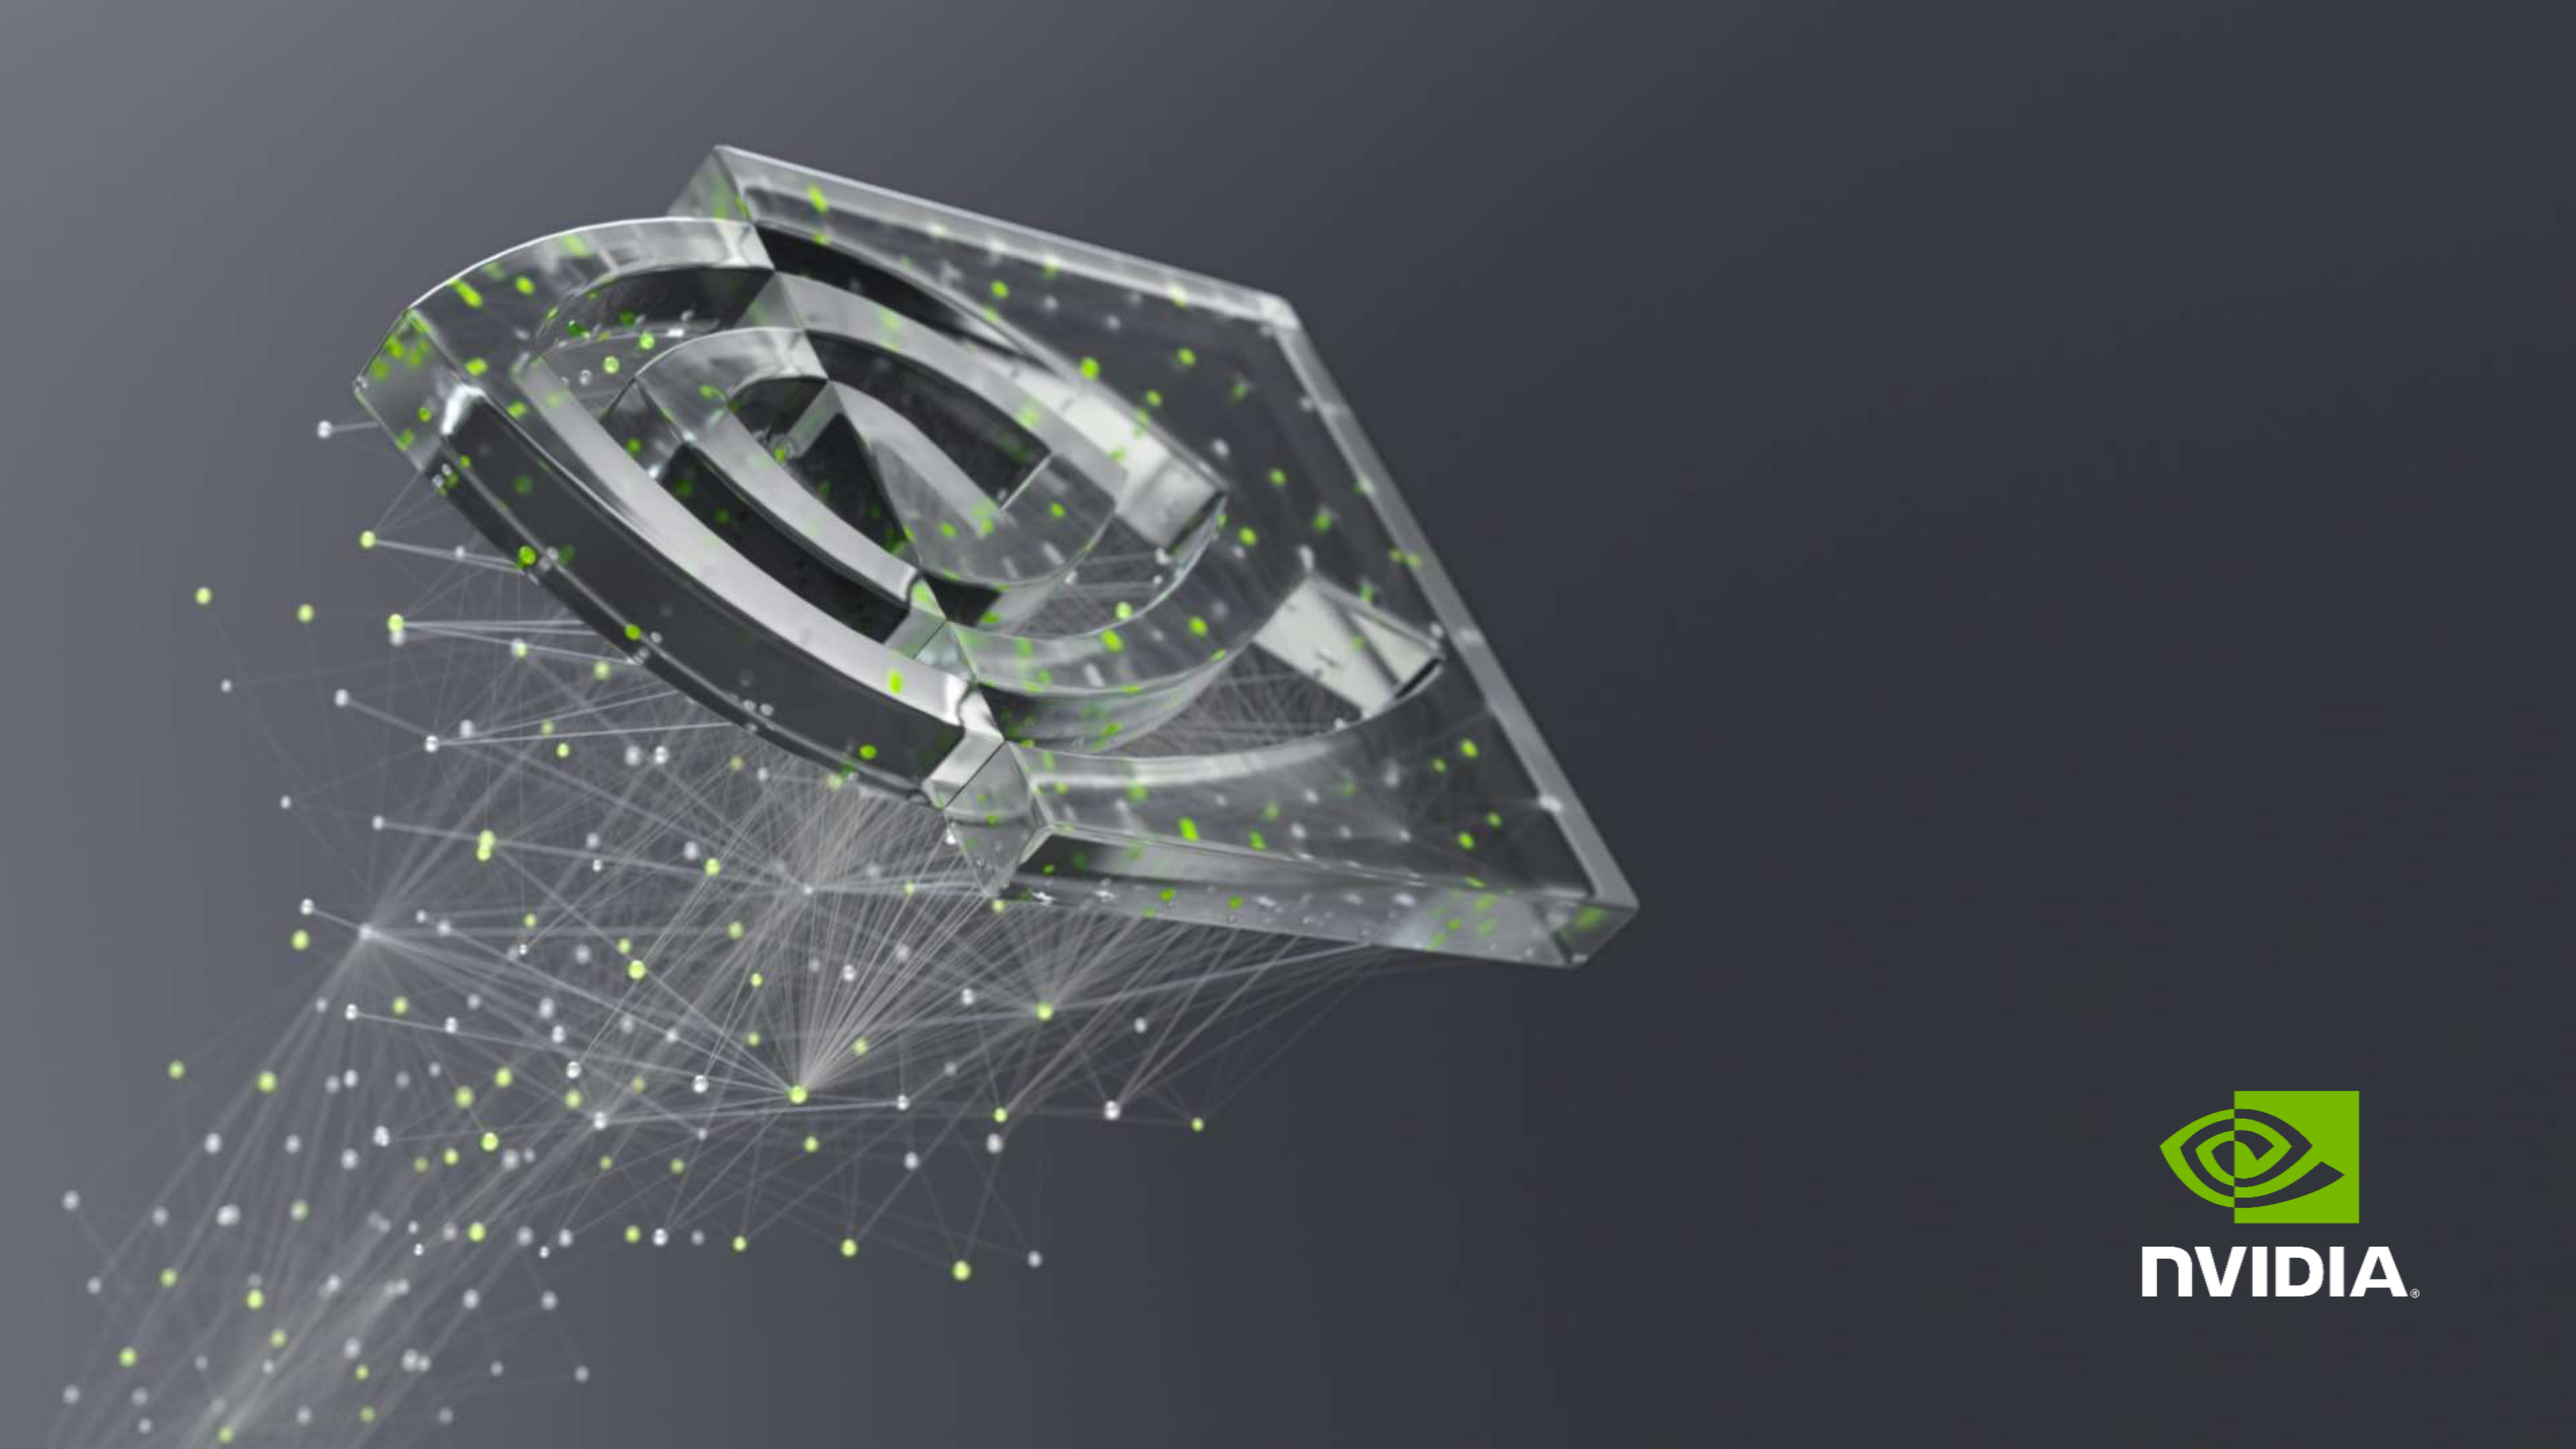
NVIDIA