




2022 OFA Virtual Workshop

CLOUD NATIVE SUPERCOMPUTING


Gilad Shainer

NVIDIA

DIVERSITY OF APPLICATIONS REQUIRES ARCHITECTURAL FLEXIBILITY




DIVERSITY OF APPLICATIONS REQUIRES ARCHITECTURAL FLEXIBILITY



NVIDIA QUANTUM-2

400G NDR InfiniBand Cloud-Native Supercomputing




Multi-Tenant
Bare-Metal Secure



Performance
Isolation



Programmability




In-Network
Computing



Precision
Timing


NVIDIA QUANTUM-2

400G NDR InfiniBand Cloud-Native Supercomputing




QUANTUM-2 SWITCH

Optimized Multi-Tenant In-Network Computing
64-Ports of 400 Gbps or 128-Ports of 200 Gbps
3X Higher Switching Throughput
6.5X Higher Scalability >1M Nodes with DF+
32X More AI Acceleration Engines



CONNECTX-7 INFINIBAND

400 Gbps Crypto Accelerations
4X In-Network Computing Performance
2X RDMA / GPUDirect Throughput
2X message rate
PCIe Gen5



BLUEFIELD-3 INFINIBAND

16 Arm 64-Bit Cores
16 Core / 256 Threads Datapath Accelerator
400 Gbps Crypto Accelerations
4X In-Network Computing Performance
2X RDMA / GPUDirect Throughput
PCIe Gen5

IN-NETWORK COMPUTING ACCELERATED SUPERCOMPUTING

Software-Defined, Hardware-Accelerated, InfiniBand Network

Networking Services

End-to-End	High Throughput	Extremely Low Latency	High Message Rate
	RDMA	GPUDirect RDMA	GPUDirect Storage
	Adaptive Routing	Congestion Control	Smart Topologies



In-Network Computing Services

Adapter/DPU	All-to-All	MPI Tag Matching	Data Reductions (SHARP)	Switch
	Programmable Datapath Accelerator	Data processing units (Arm cores)	Self Healing Network	
End-to-End	Data security / tenant isolation			End-to-End



MULTI-TENANT SUPERCOMPUTING CLOUD

Molecular Dynamics (LAMMPS) Example



HPC ON SUPERCOMPUTING



HPC ON THE CLOUD




HPC ON CLOUD-NATIVE SUPERCOMPUTING



PERFORMANCE ISOLATION - MICROSOFT AZURE

Quantum InfiniBand Congestion Control

Number of Iterations



HPC ON SUPERCOMPUTING

HPC ON THE CLOUD

HPC ON CLOUD-NATIVE SUPERCOMPUTING

CLOUD-NATIVE SUPERCOMPUTING

Bare-metal Secured Infrastructure

Higher Application Performance

From the Edge to the Main Data Center




BARE-METAL
PERFORMANCE




MULTI
TENANCY



EDGE
COMPUTING



CLOUD-NATIVE SUPERCOMPUTING INFRASTRUCTURE



MULTI-TENANT ISOLATION


Zero-Trust Architecture

Secured Network Infrastructure and Configuration

Computational Storage

Tenant Service Level Agreement (SLA)

32K Concurrent Isolated Users on Single Subnet



HIGHER APPLICATION PERFORMANCE

DPU-Accelerated HPC Communications


Collective Offloads

Active Messages




Smart MPI Progression

Data Compression

User-defined Algorithms




NON-BLOCKING MPI PERFORMANCE



HIGHER APPLICATION PERFORMANCE

100% Communication - Computation Overlap



Courtesy of: Ohio State University MVAPICH Team and X-ScaleSolutions




32 servers, Dual Socket Intel® Xeon® 16-core CPUs E5-2697A V4 @ 2.60 GHz (32 processes per node), NVIDIA BlueField-2 HDR100 DPUs and ConnectX-6 HDR100 adapters, NVIDIA HDR Quantum Switch QM7800 40-Port 200Gb/s HDR InfiniBand, 256GB DDR4 2400MHz RDIMMs memory and 1TB 7.2K RPM SATA 2.5" hard drive per node.



HIGHER APPLICATION PERFORMANCE

Higher App Performance, MPI Collectives Offload



Courtesy of: Ohio State University MVAPICH Team and X-ScaleSolutions



MVAPICH




X-ScaleSolutions

32 servers, Dual Socket Intel® Xeon® 16-core CPUs E5-2697A V4 @ 2.60 GHz (32 processes per node), NVIDIA BlueField-2 HDR100 DPUs and ConnectX-6 HDR100 adapters, NVIDIA HDR Quantum Switch QM7800 40-Port 200Gb/s HDR InfiniBand, 256GB DDR4 2400MHz RDIMMs memory and 1TB 7.2K RPM SATA 2.5" hard drive per node.




HIGHER APPLICATION PERFORMANCE

Georgia Tech & Sandia National Lab - minMD




16 servers, Dual Socket Intel® Xeon® 16-core CPUs E5-2697A V4 @ 2.60 GHz (32 processes per node), NVIDIA BlueField-2 HDR100 DPUs and ConnectX-6 HDR100 adapters, NVIDIA HDR Quantum Switch QM7800 40-Port 200Gb/s HDR InfiniBand, 256GB DDR4 2400MHz RDIMMs memory and 1TB 7.2K RPM SATA 2.5" hard drive per node.



MORPHEUS AI CYBERSECURITY FRAMEWORK

UFM Cyber-AI with Morpheus for Higher Security



NVIDIA DOCA

Enabling Broad DPU Partner Ecosystem

Software Application Framework for BlueField DPUs

DOCA is for DPUs What CUDA is for GPUs

Protects Developer Investment for Future DPUs

Certified Reference Applications, APIs & Partner Solutions

Rich Partner Ecosystem Across Industries and Workloads

<https://developer.nvidia.com/networking/doca>

