




# Interconnect Your Future

Paving the Road to Exascale

June 2018

# HPC and AI Needs the Most Intelligent Interconnect



SmartNIC



System on a Chip

Higher

Data Speeds

Faster

Data Processing

Better

Data Security



Adapters



Switches



Cables &  
Transceivers



# Mellanox Accelerates Leading HPC and AI Systems

**OAK RIDGE**  
National Laboratory



'Summit' CORAL System  
World's Fastest HPC / AI System  
9.2K InfiniBand Nodes



国家超级计算无锡中心  
National Supercomputing Center in Wuxi



Wuxi Supercomputing Center  
Fastest Supercomputer in China  
41K InfiniBand Nodes



Lawrence Livermore  
National Laboratory



'Sierra' CORAL System  
#2 USA Supercomputer  
8.6K InfiniBand Nodes





# Mellanox Accelerates Leading HPC and AI Systems



UNIVERSITY OF  
**TORONTO**  
SciNet



Fastest Supercomputer in Canada  
Dragonfly+ Topology  
1.5K InfiniBand Nodes



**AIST**  
NATIONAL INSTITUTE OF  
ADVANCED INDUSTRIAL SCIENCE  
AND TECHNOLOGY (AIST)




Fastest HPC / AI System in Japan  
1.1K InfiniBand Nodes



Sandia  
National  
Laboratories



**NNSA**  
National Nuclear Security Administration




'Astra' Arm-Based Supercomputer  
NNSA Vanguard Program  
2.6K InfiniBand Nodes

# Data Centric Architecture to Overcome Latency Bottlenecks


Intelligent Interconnect Paves the Road to Exascale Performance

CPU-Centric (Onload)



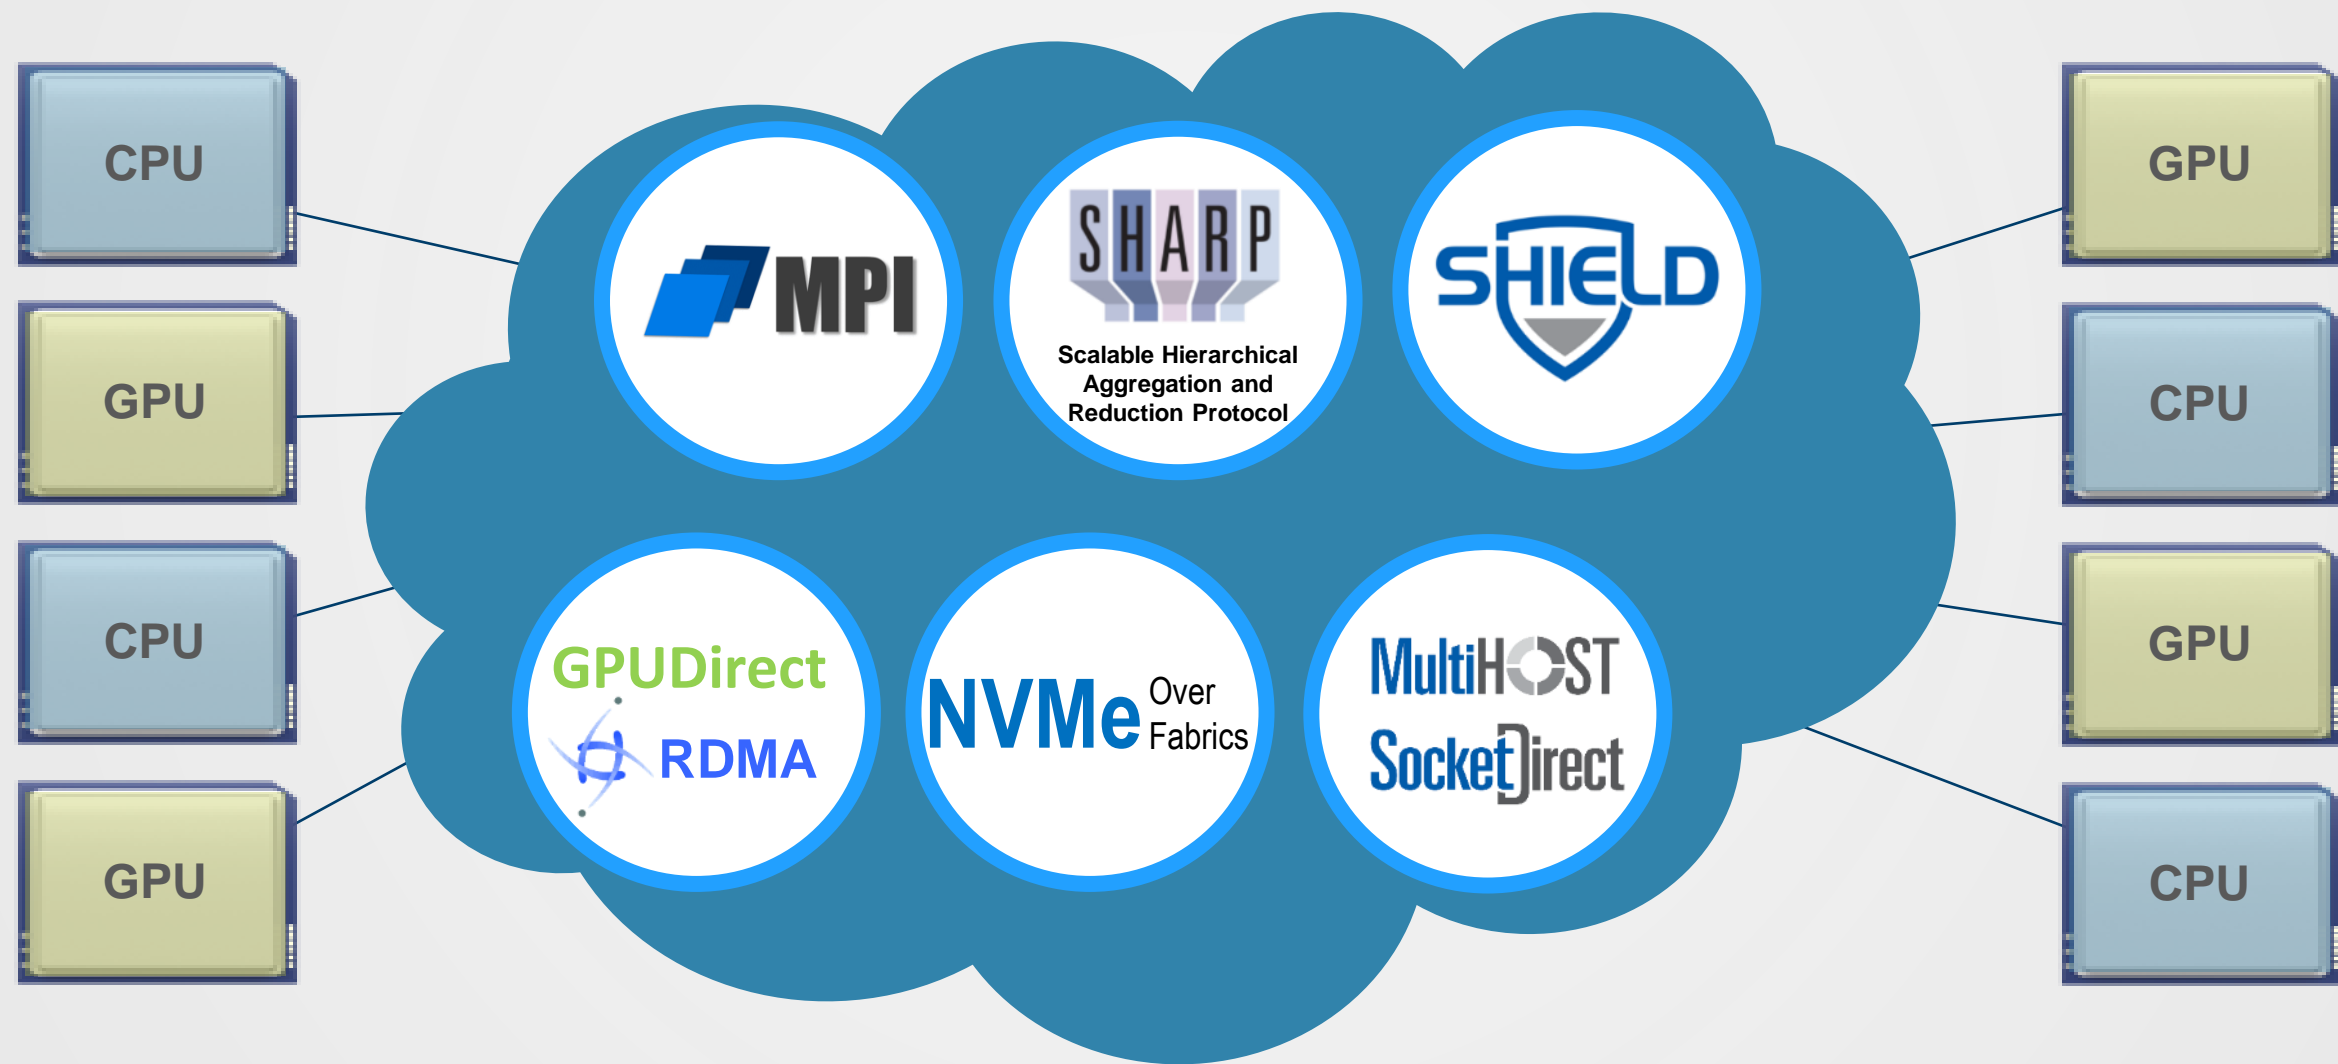
Communications Latencies  
of 30-40us

Data-Centric (Offload)




Communications Latencies  
of 3-4us

# In-Network Computing to Enable Data-Centric Data Centers




# SHARP AllReduce Performance Advantages (128 Nodes)

Allreduce Latency



Allreduce Latency



SHARP enables 75% Reduction in Latency  
Providing Scalable Flat Latency



# Unified Communication - X Framework

WEB:

[www.openucx.org](http://www.openucx.org)

<https://github.com/openucx/ucx>

Mailing List:

<https://elist.ornl.gov/mailman/listinfo/ucx-group>

[ucx-group@elist.ornl.gov](mailto:ucx-group@elist.ornl.gov)



# Highest-Performance 200Gb/s Interconnect Solutions

Adapters

ConnectX<sup>®</sup> 6

200Gb/s Adapter, 0.6us latency  
200 million messages per second  
(10 / 25 / 40 / 50 / 56 / 100 / 200Gb/s)



Switch

Quantum  
The smartest switch, became smarter

40 HDR (200Gb/s) InfiniBand Ports  
80 HDR100 InfiniBand Ports  
Throughput of 16Tb/s, <90ns Latency



Switch

Spectrum<sup>™</sup> 2

16 400GbE, 32 200GbE, 128 25/50GbE Ports  
(10 / 25 / 40 / 50 / 100 / 200 GbE)  
Throughput of 6.4Tb/s



SoC

BlueField

System on Chip and SmartNIC  
Programmable adapter  
Smart Offloads



Interconnect

LinkX<sup>™</sup>

Transceivers  
Active Optical and Copper Cables  
(10 / 25 / 40 / 50 / 56 / 100 / 200Gb/s)



Software

HPC-X<sup>™</sup>

MPI, SHMEM/PGAS, UPC  
For Commercial and Open Source Applications  
Leverages Hardware Accelerations



# BlueField Product Line




System on Chip



BlueField™ Platform  
BF1100 & BF1200



SmartNIC  
Ethernet



SmartNIC VPI  
BF1660



Dual Port 100Gb/s  
Controller  
BF1600



Dual Port 100Gb/s  
Controller  
BF1700

## ■ Different SKUs

- # Cores
- Speeds
- Perf points

## ■ Different SKUs

- GPU and SSD
- 1U and 2U
- Up to 16 SSDs

## ■ Different SKUs

- 16/8/4 cores
- Dual 25Gb/s

## ■ VPI 2-ports 100G

- 100Gb/s / EDR
- 16 Cores

## ■ Different SKUs

- PCIe x16
- PCIe x32

## ■ Storage card

## Full Software Enablement

- SDK and Development Tools
- Networking , Security Features
- Full NVMe storage capability



# Thank You

