



Interconnect Your Future
Enabling the Best Datacenter Return on Investment

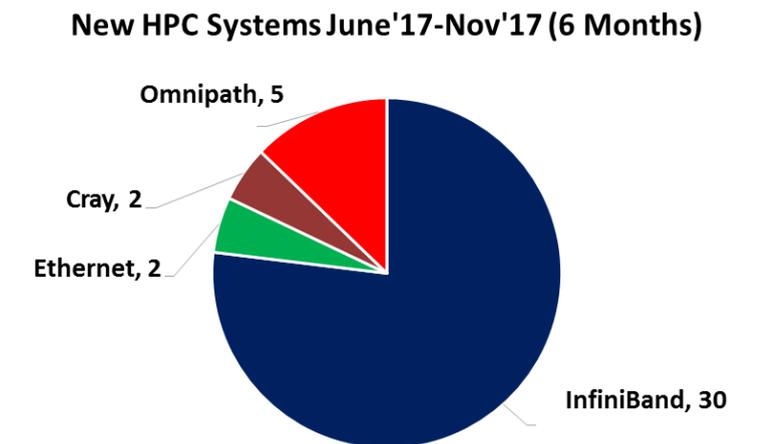
TOP500 Supercomputers, November 2017

 **Mellanox**
TECHNOLOGIES
Connect. Accelerate. Outperform.™

InfiniBand Accelerates Majority of New Systems on TOP500



- InfiniBand connects 77% of new HPC systems (June'17 – Nov'17)
- InfiniBand connects 6 times more new systems versus proprietary interconnects (June'17 - Nov'17)
- InfiniBand connects 15 times more new systems versus Ethernet (June'17 - Nov'17)
- Mellanox accelerates the fastest supercomputer in the world
- InfiniBand connects 2 of the top 5 supercomputers (#1, #4)
- Mellanox connects 39% of overall TOP500 systems (192 systems, InfiniBand and Ethernet)
- InfiniBand connects 33% of the total TOP500 systems (164 systems)
- InfiniBand connects 60% of the HPC TOP500 systems
- 25G Ethernet first appearance on the Nov'17 TOP500 list (China Hyperscale company) - 19 systems, all are Mellanox connected



InfiniBand is the Interconnect of Choice for HPC Infrastructures
Enabling Machine Learning, High-Performance, Web 2.0, Cloud, Storage, Big Data Applications

National Supercomputing Center in Wuxi, China #1 on the TOP500 List



- 93 Petaflop performance, 3X higher versus #2 on the TOP500
- 41K nodes, 10 million cores, 256 cores per CPU
- Mellanox adapter and switch solutions



* Source: "Report on the Sunway TaihuLight System", Jack Dongarra (University of Tennessee), June 20, 2016 (Tech Report UT-EECS-16-742)

National Supercomputing Center in Wuxi, China

#1 on the TOP500 List



Japan Agency for Marine-Earth Science and Technology

#4 on the TOP500 List, 19 Petaflops



facebook.

Facebook AI Supercomputer

#35 on the TOP500 List



nvidia

NVIDIA AI Supercomputer

#36 on the TOP500 List



- EDR InfiniBand In-Network Computing technology key for scalable Deep Learning systems
- RDMA accelerates Deep Learning performance by 2X, becomes de-facto solution for AI

Mellanox Accelerates Next Generation HPC and AI Systems



Summit CORAL System



Sierra CORAL System



Fastest Supercomputer in Japan



Fastest Supercomputer in Canada

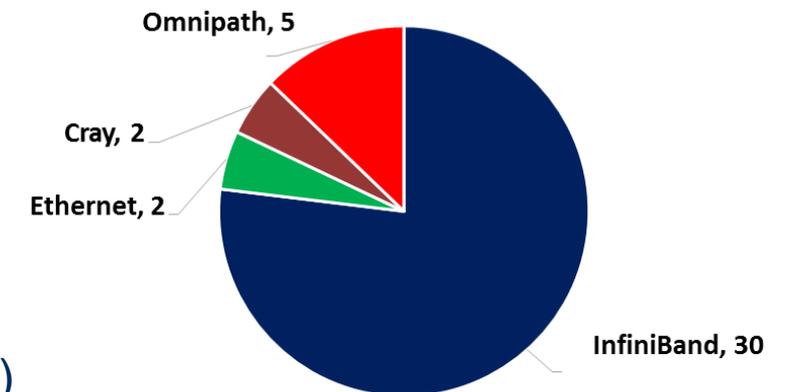
InfiniBand Dragonfly+ Topology



Highest Performance with InfiniBand

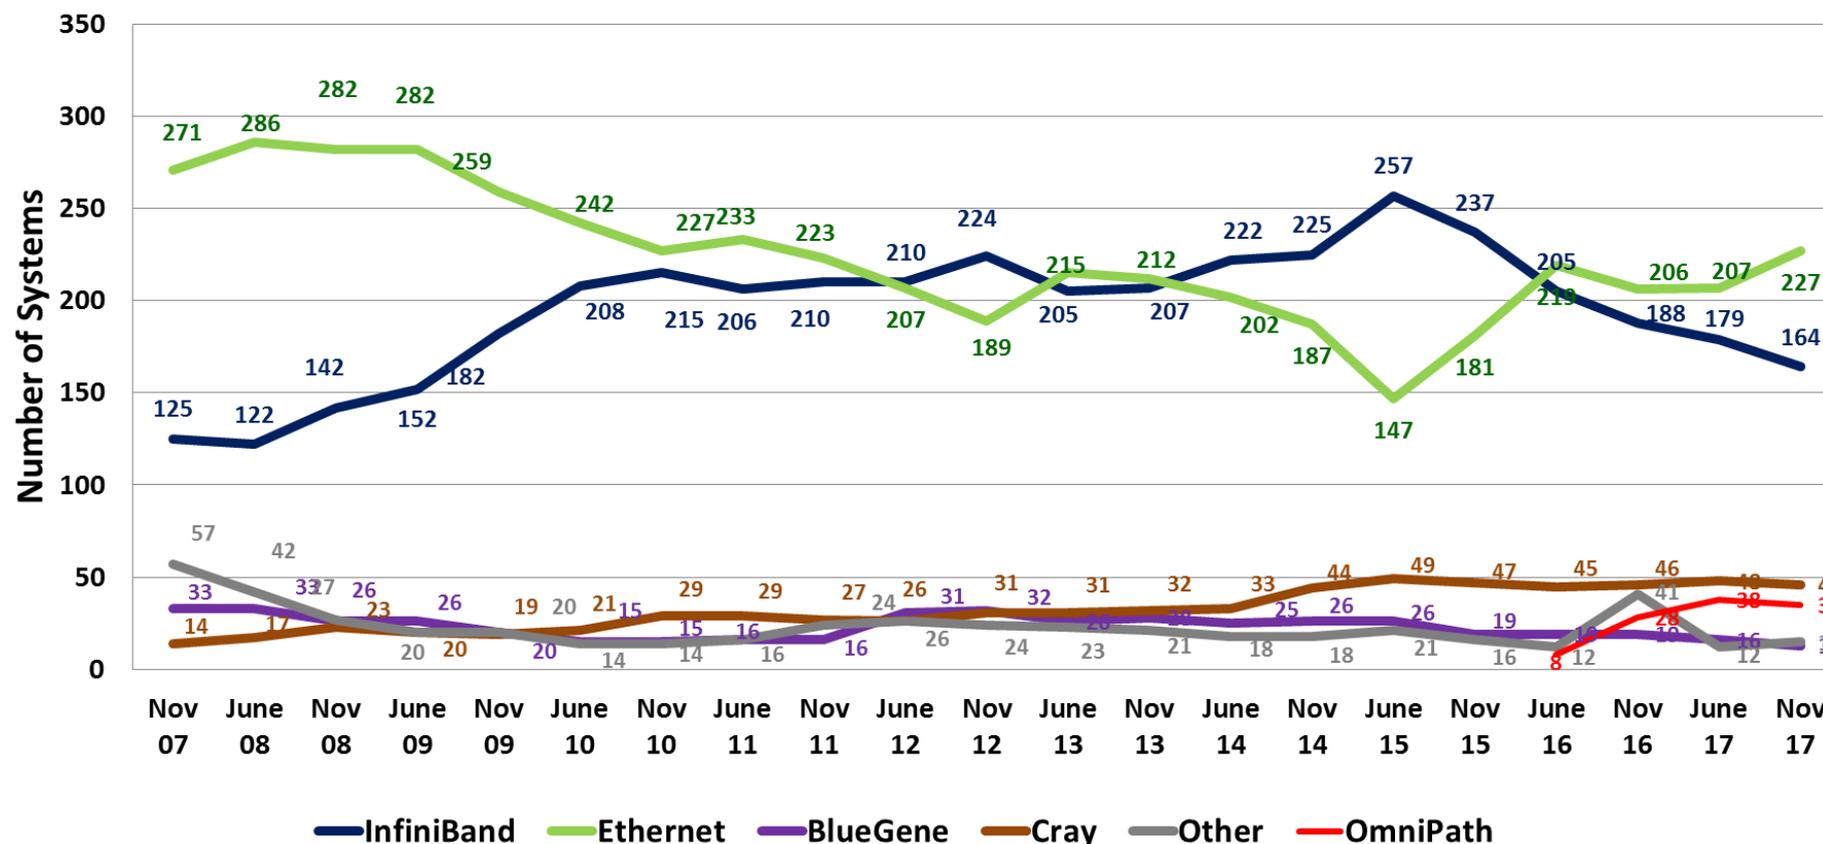
- Mellanox accelerates the fastest supercomputer on the list
- InfiniBand connects 2 of top 5 systems - #1 (China) and #4 (Japan)
- InfiniBand connects 6 times more new HPC systems versus OmniPath (June'17 - Nov'17)
- InfiniBand connects 15 times more new HPC systems versus Ethernet (June'17 - Nov'17)
- InfiniBand connects 77% of new HPC systems (June'17 – Nov'17)
- Mellanox connects 39 percent of overall TOP500 systems (192 systems, InfiniBand and Ethernet)
- InfiniBand connects 33 percent of the total TOP500 systems (164 systems)
- InfiniBand connects 60 percent of the HPC TOP500 systems
- 25G Ethernet first appearance on the Nov'17 TOP500 list (China Hyperscale company) - 19 systems, all are Mellanox connected
- Mellanox connects all of 25G, 40G and 100G Ethernet systems
- InfiniBand is the most used high-speed Interconnect on the TOP500
- InfiniBand is the preferred interconnect for Artificial Intelligence and Deep Learning systems
- Mellanox enable highest ROI for Machine Learning, High-Performance, Cloud, Storage, Big Data and more applications

New HPC Systems June'17-Nov'17 (6 Months)



Paving The Road to Exascale Performance

TOP500 Interconnect Trends

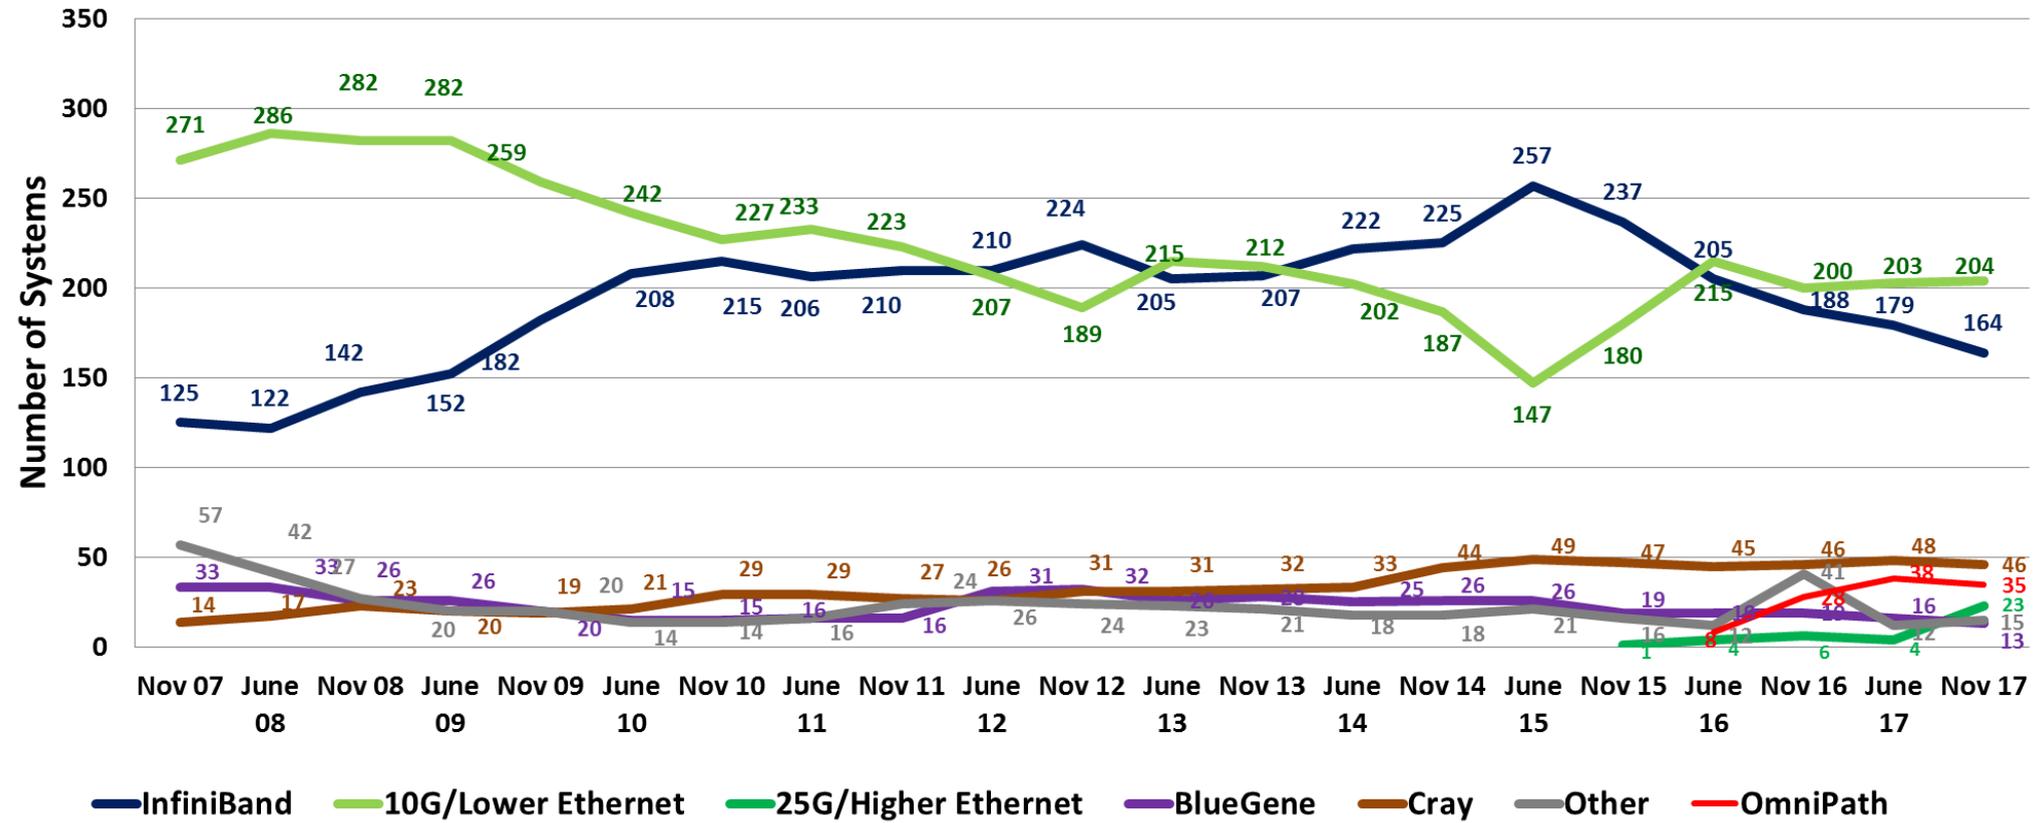


The TOP500 List has Evolved to Include Both HPC and Cloud / Web2.0 Hyperscale Platforms
For the HPC Platforms, InfiniBand Continues it's Leadership as the Most Used Interconnect Solution
Mellanox Connects all 25, 40 and 100G Ethernet Systems

TOP500 Interconnect Trends (Ethernet Speed Details)



TOP500 Interconnect Trends

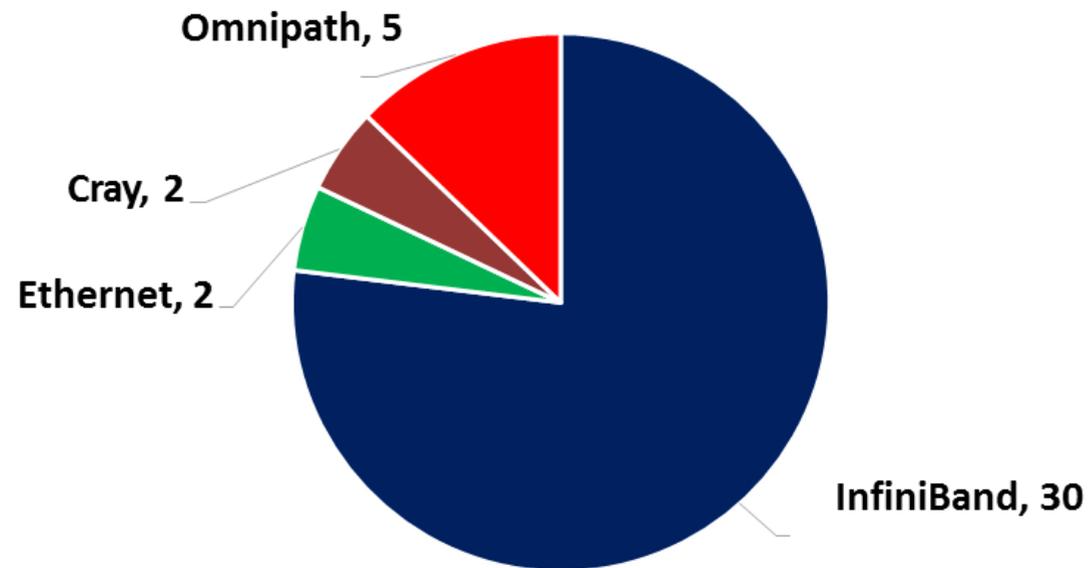


The TOP500 List has Evolved to Include Both HPC and Cloud / Web2.0 Hyperscale Platforms
For the HPC Platforms, InfiniBand Continues it's Leadership as the Most Used Interconnect Solution
Mellanox Connects all 25, 40 and 100G Ethernet Systems

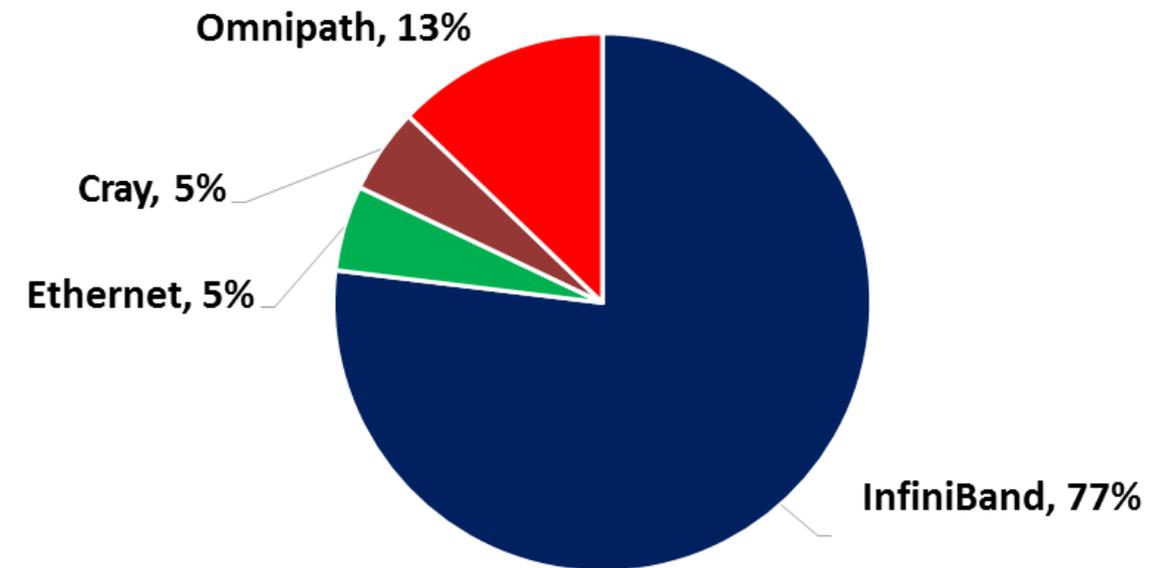
New TOP500 HPC Systems June'17 – Nov'17 (6 Months Period)



New HPC Systems June'17-Nov'17 (6 Months)



New HPC Systems June'17-Nov'17 (6 Months)



New TOP500 HPC Systems Added to TOP500 List Between June'17 to Nov'17

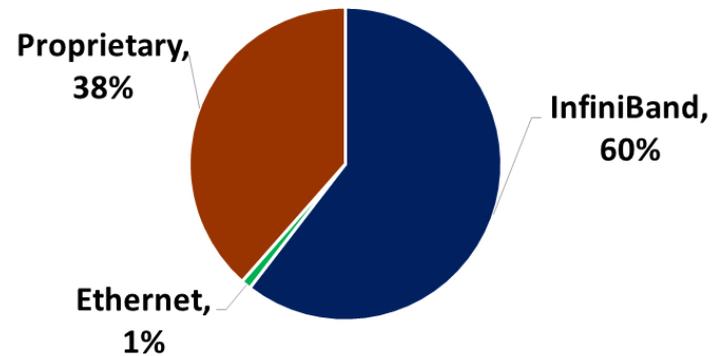
InfiniBand Accelerates 6X New Systems Compare to OmniPath

InfiniBand Accelerates 15X New Systems Compare to Ethernet

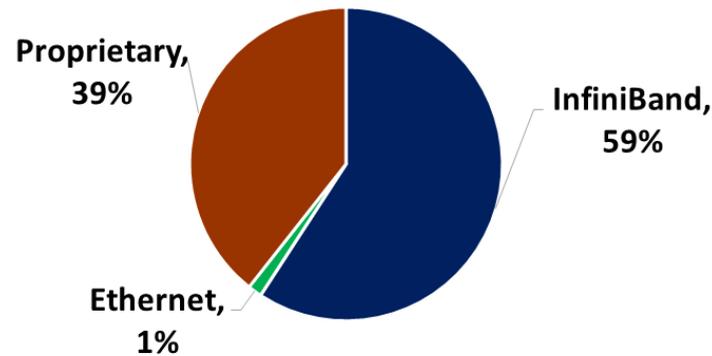
TOP500 HPC Systems TOP500 Interconnect Trends



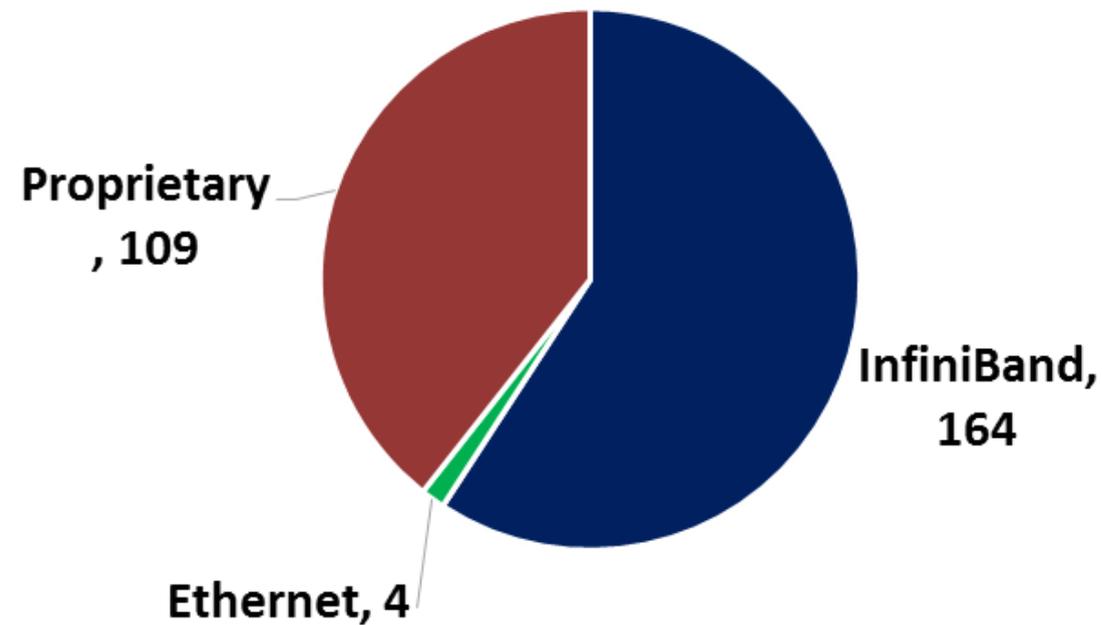
TOP500 June'17 - % of HPC Systems



TOP500 Nov'17 - % of HPC Systems

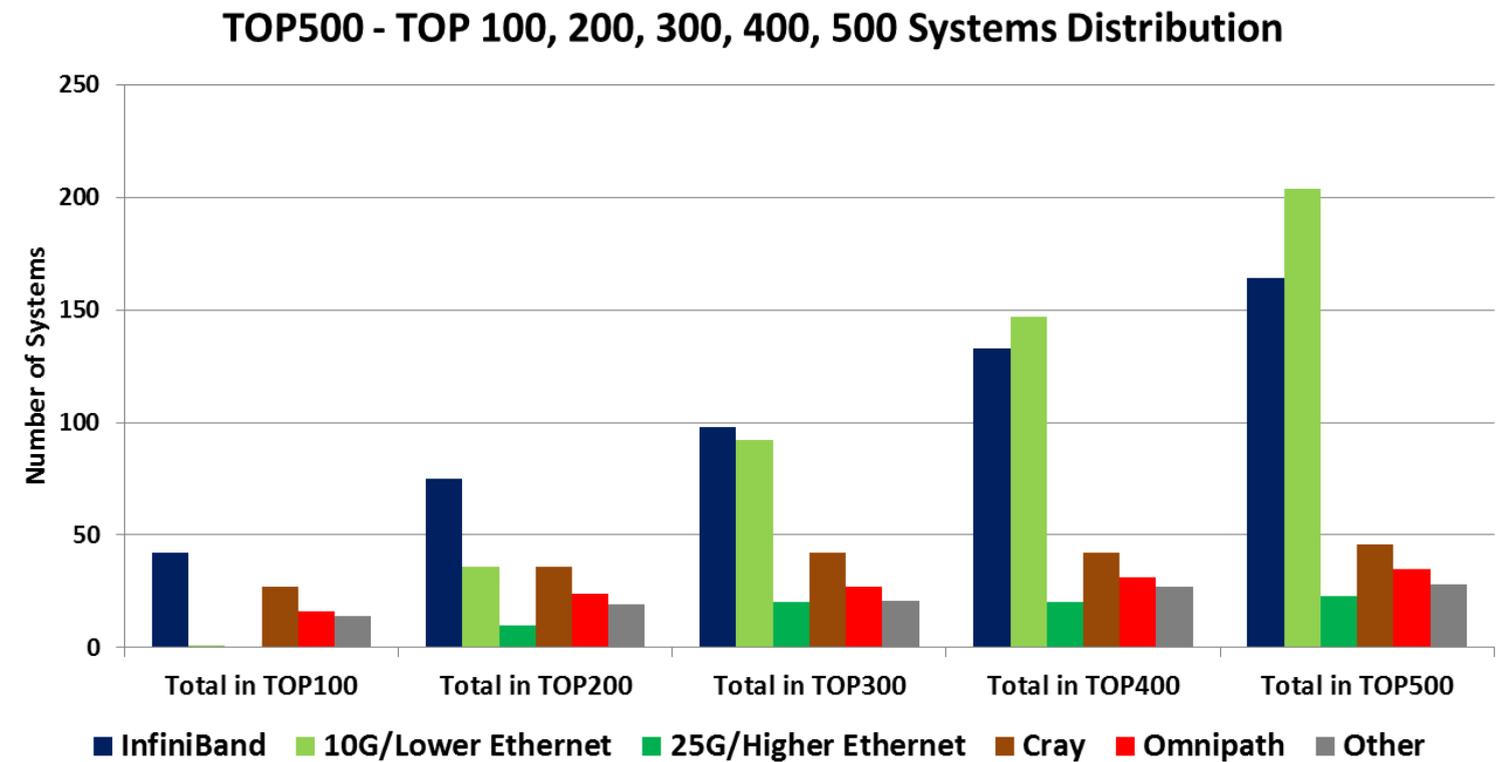
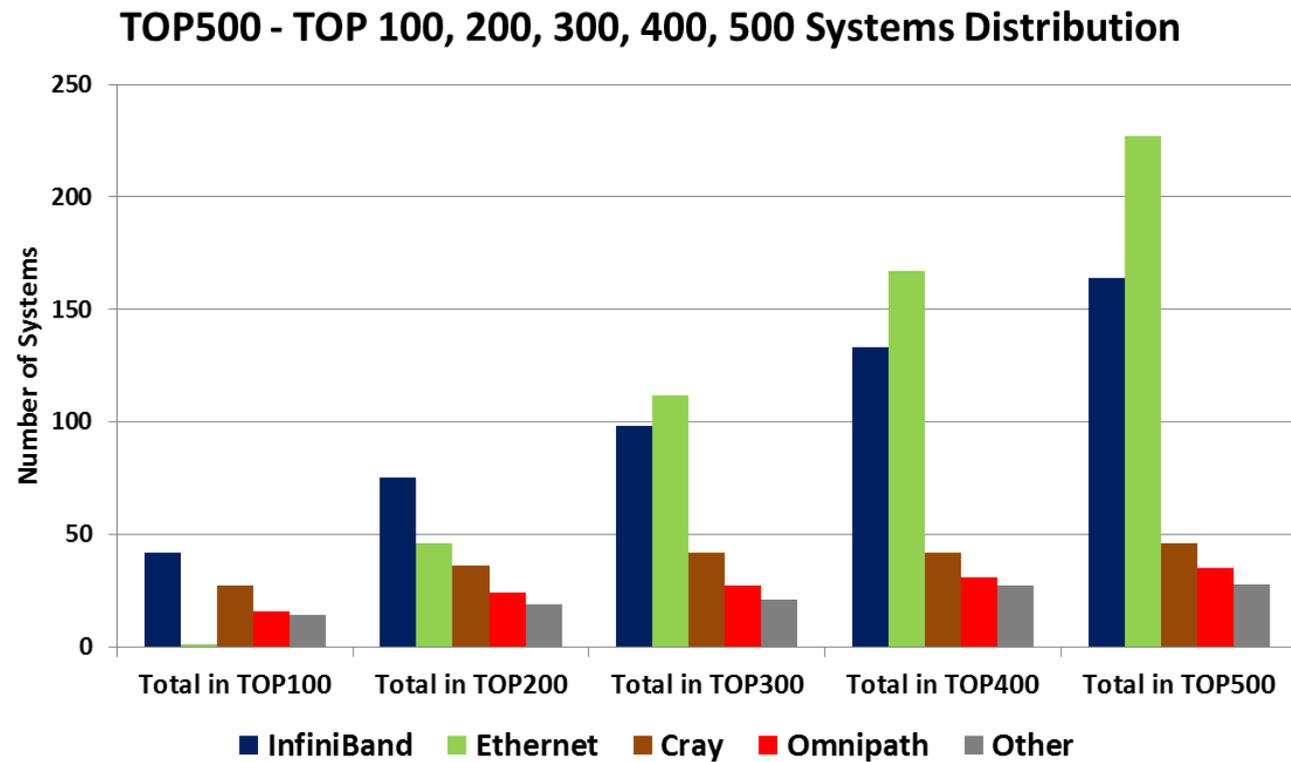


TOP500 Nov'17 - Number of HPC Systems



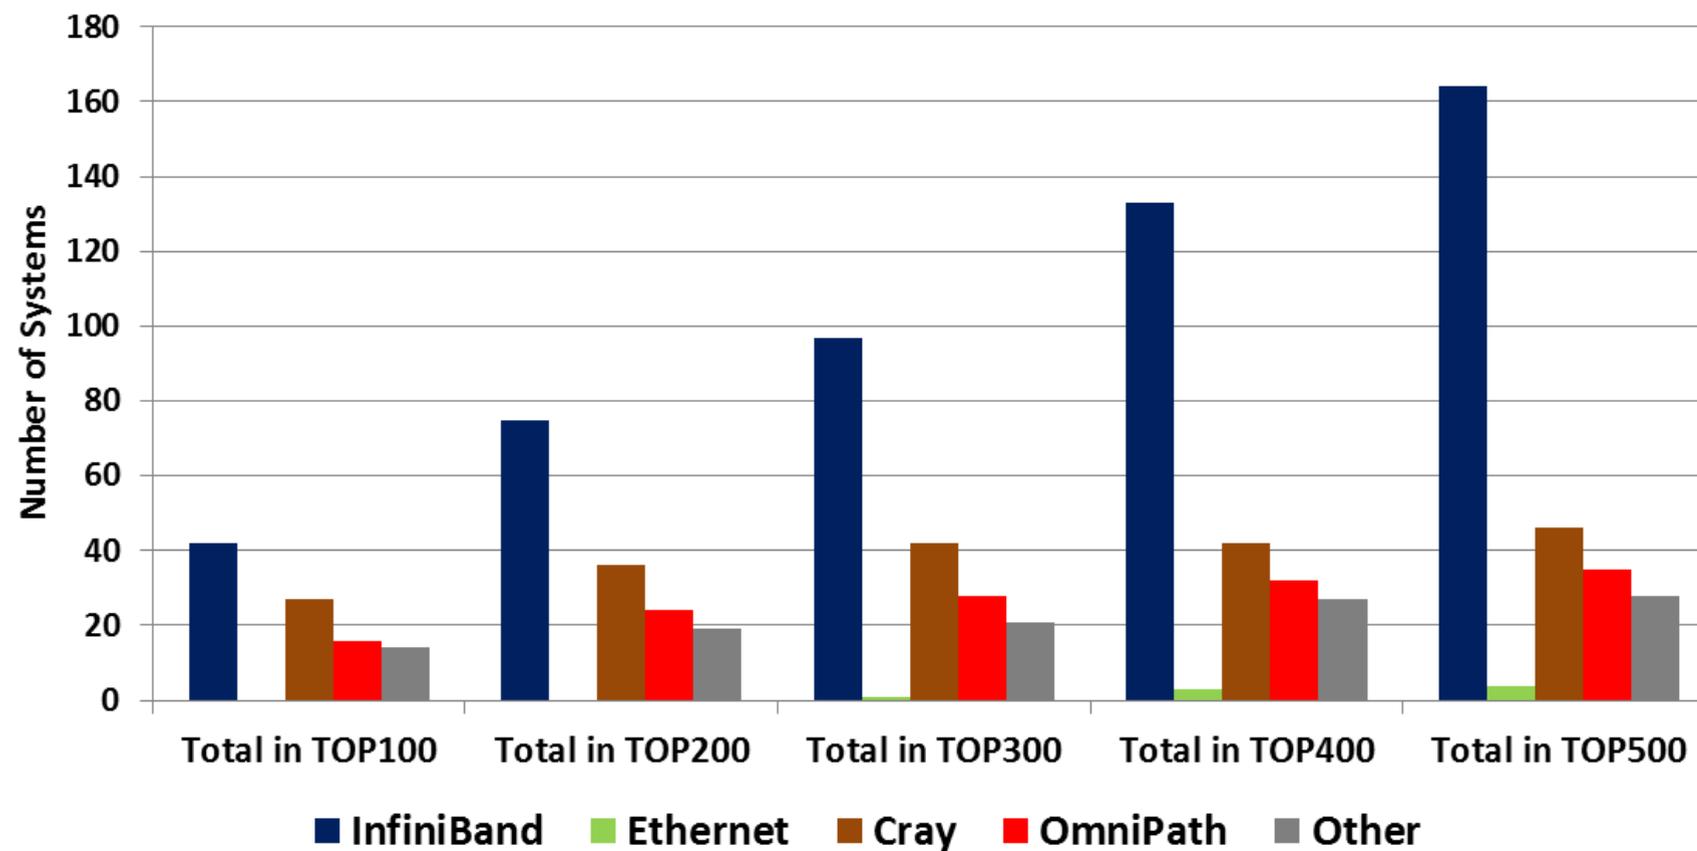
TOP500 List Includes HPC and Non-HPC Systems

InfiniBand is the Interconnect of Choice for HPC Systems



InfiniBand is The Most Used High-Speed Interconnect of the TOP500
Superior Performance, Scalability, Efficiency and Return-On-Investment

TOP500 - TOP 100, 200, 300, 400, 500 Systems Distribution HPC Systems Only



InfiniBand is The Most Used Interconnect For HPC Systems

Superior Performance, Scalability, Efficiency and Return-On-Investment

Maximum Efficiency and Return on Investment



- Mellanox smart interconnect solutions enable In-Network Computing and CPU-Offloading
- Critical with CPU accelerators and higher scale deployments
- Ensures highest system efficiency and overall return on investment

InfiniBand

System: NASA, System Efficiency: 84%

System: NCAR, System Efficiency: 90%

System: US Army, System Efficiency: 94%

Omni-Path

System: CINECA, System Efficiency: 57%

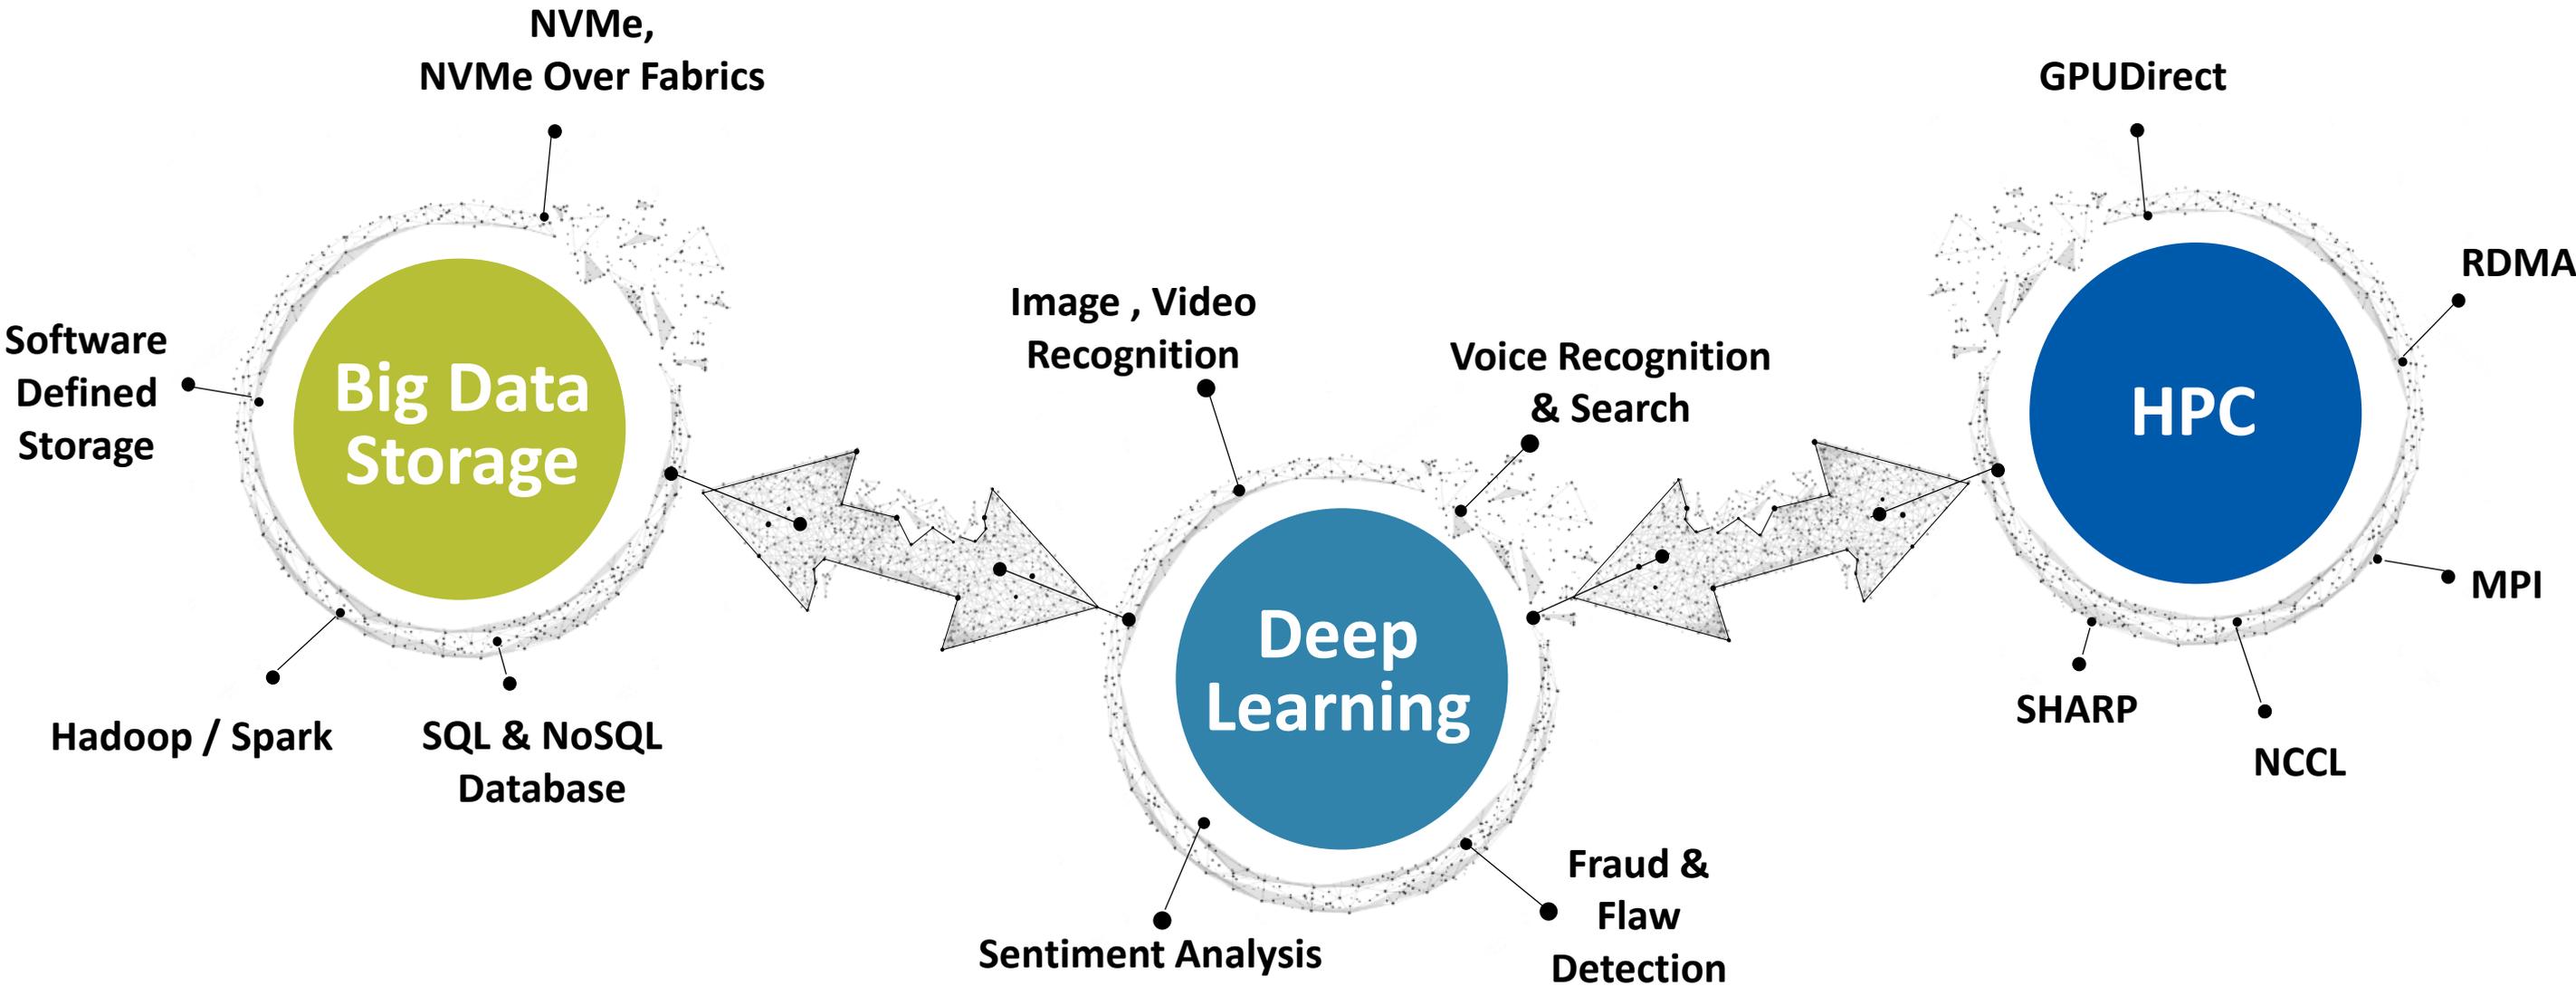
System: Barcelona, System Efficiency: 62%

System: TACC, System Efficiency: 53%

1.7X Higher System Efficiency!

43% System Resources not Utilized!

Mellanox Interconnect Technology Enables a Variety of Applications



Mellanox In-Network Computing Delivers Highest Performance



In-Network Computing



10X

Performance Acceleration

Critical for HPC and Machine Learning Applications



In-Network Computing



35X

Performance Acceleration

Delivers Highest Application Performance



Self-Healing Technology



5000X

Faster Network Recovery

Unbreakable Data Centers

GPUDirect™ RDMA

GPU Acceleration Technology



10X

Performance Acceleration

Critical for HPC and Machine Learning Applications



RDMA Supercharges Leading AI Frameworks



Mellanox Supercharges Leading AI Companies



60%
Higher ROI

50%
Lower CapEx
& OpEx

30-100% Higher Return on Investment

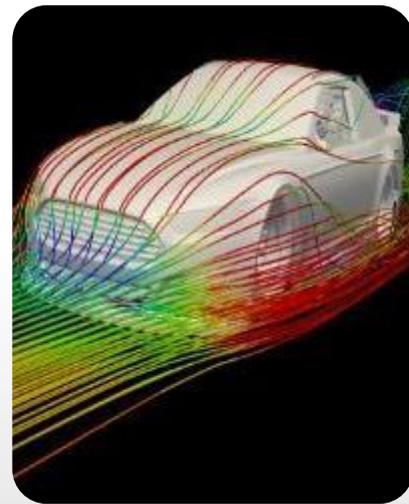
Up to **50%** Saving on Capital and Operation Expenses

Highest Applications Performance, Scalability and Productivity



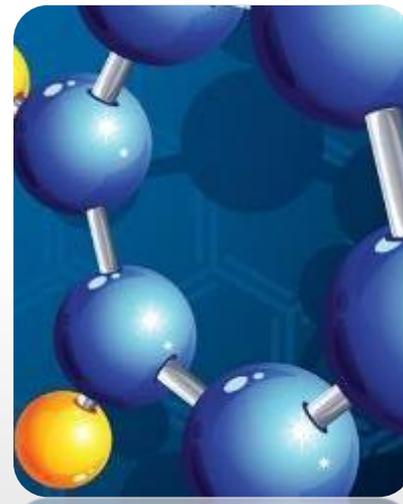
Weather

1.3X Better



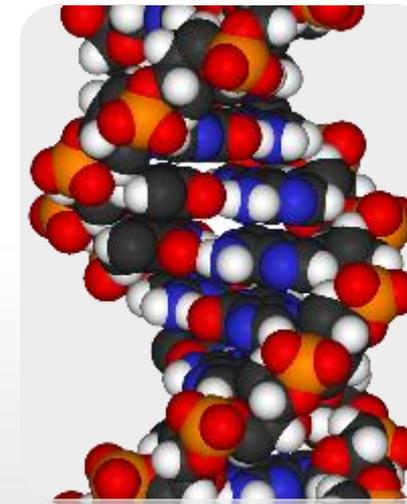
Automotive

2X Better



Chemistry

1.4X Better



**Molecular
Dynamics**

1.7X Better



Genomics

1.3X Better

NUMBERS
NEVER
LIE



InfiniBand Always Wins:

**Best Performance.
Best R.O.I.**



	OmniPath	InfiniBand
System Cost	\$491,280 With Free Network	\$278,238

\$213,042 **Cost Savings**

43% **Lower TCO**

Based on 32-node Scalable Unit Performance



NUMBERS
NEVER
—LIE—



InfiniBand Always Wins:



**Best Performance.
Best R.O.I.**

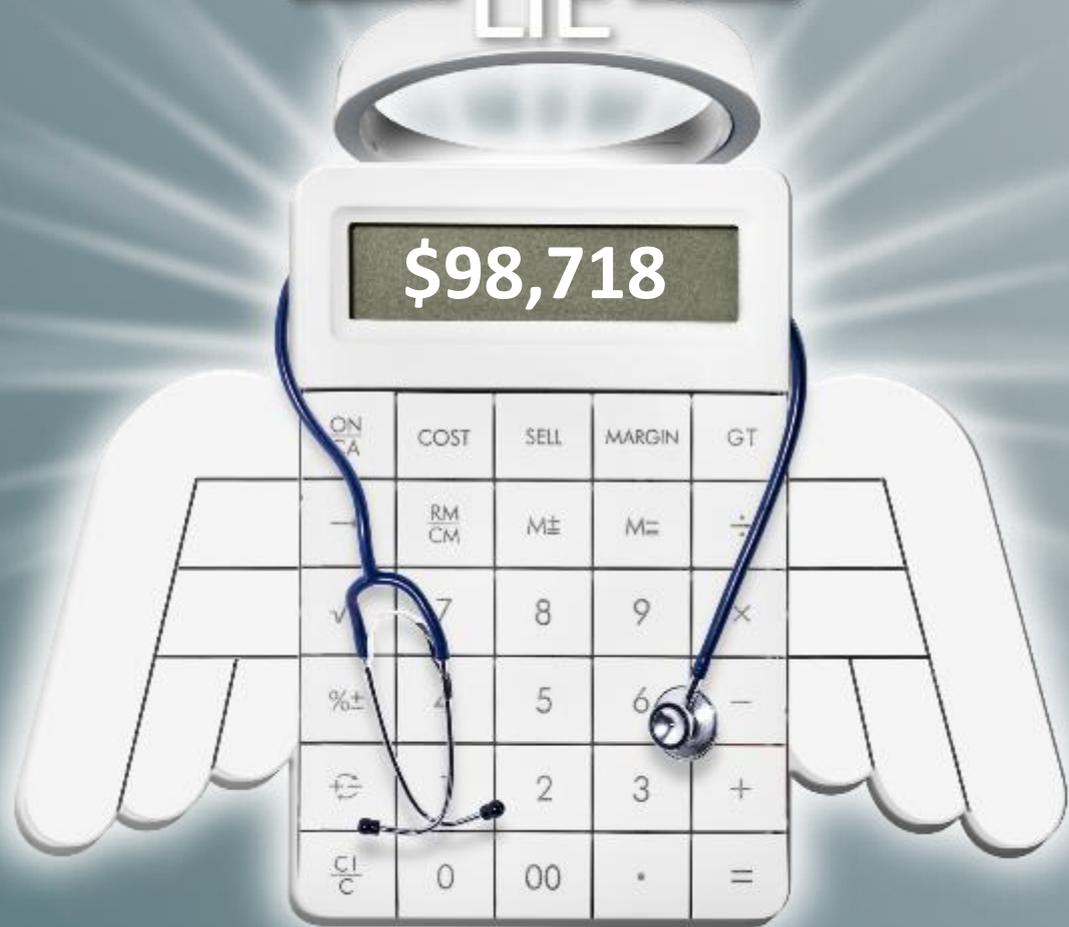


	OmniPath	InfiniBand
System Cost	\$982,560 With Free Network	\$540,870

\$441,690 **Cost Savings** | **45%** **Lower TCO**

Based on 64-node Scalable Unit Performance

NUMBERS
NEVER
— LIE —



InfiniBand Always Wins:

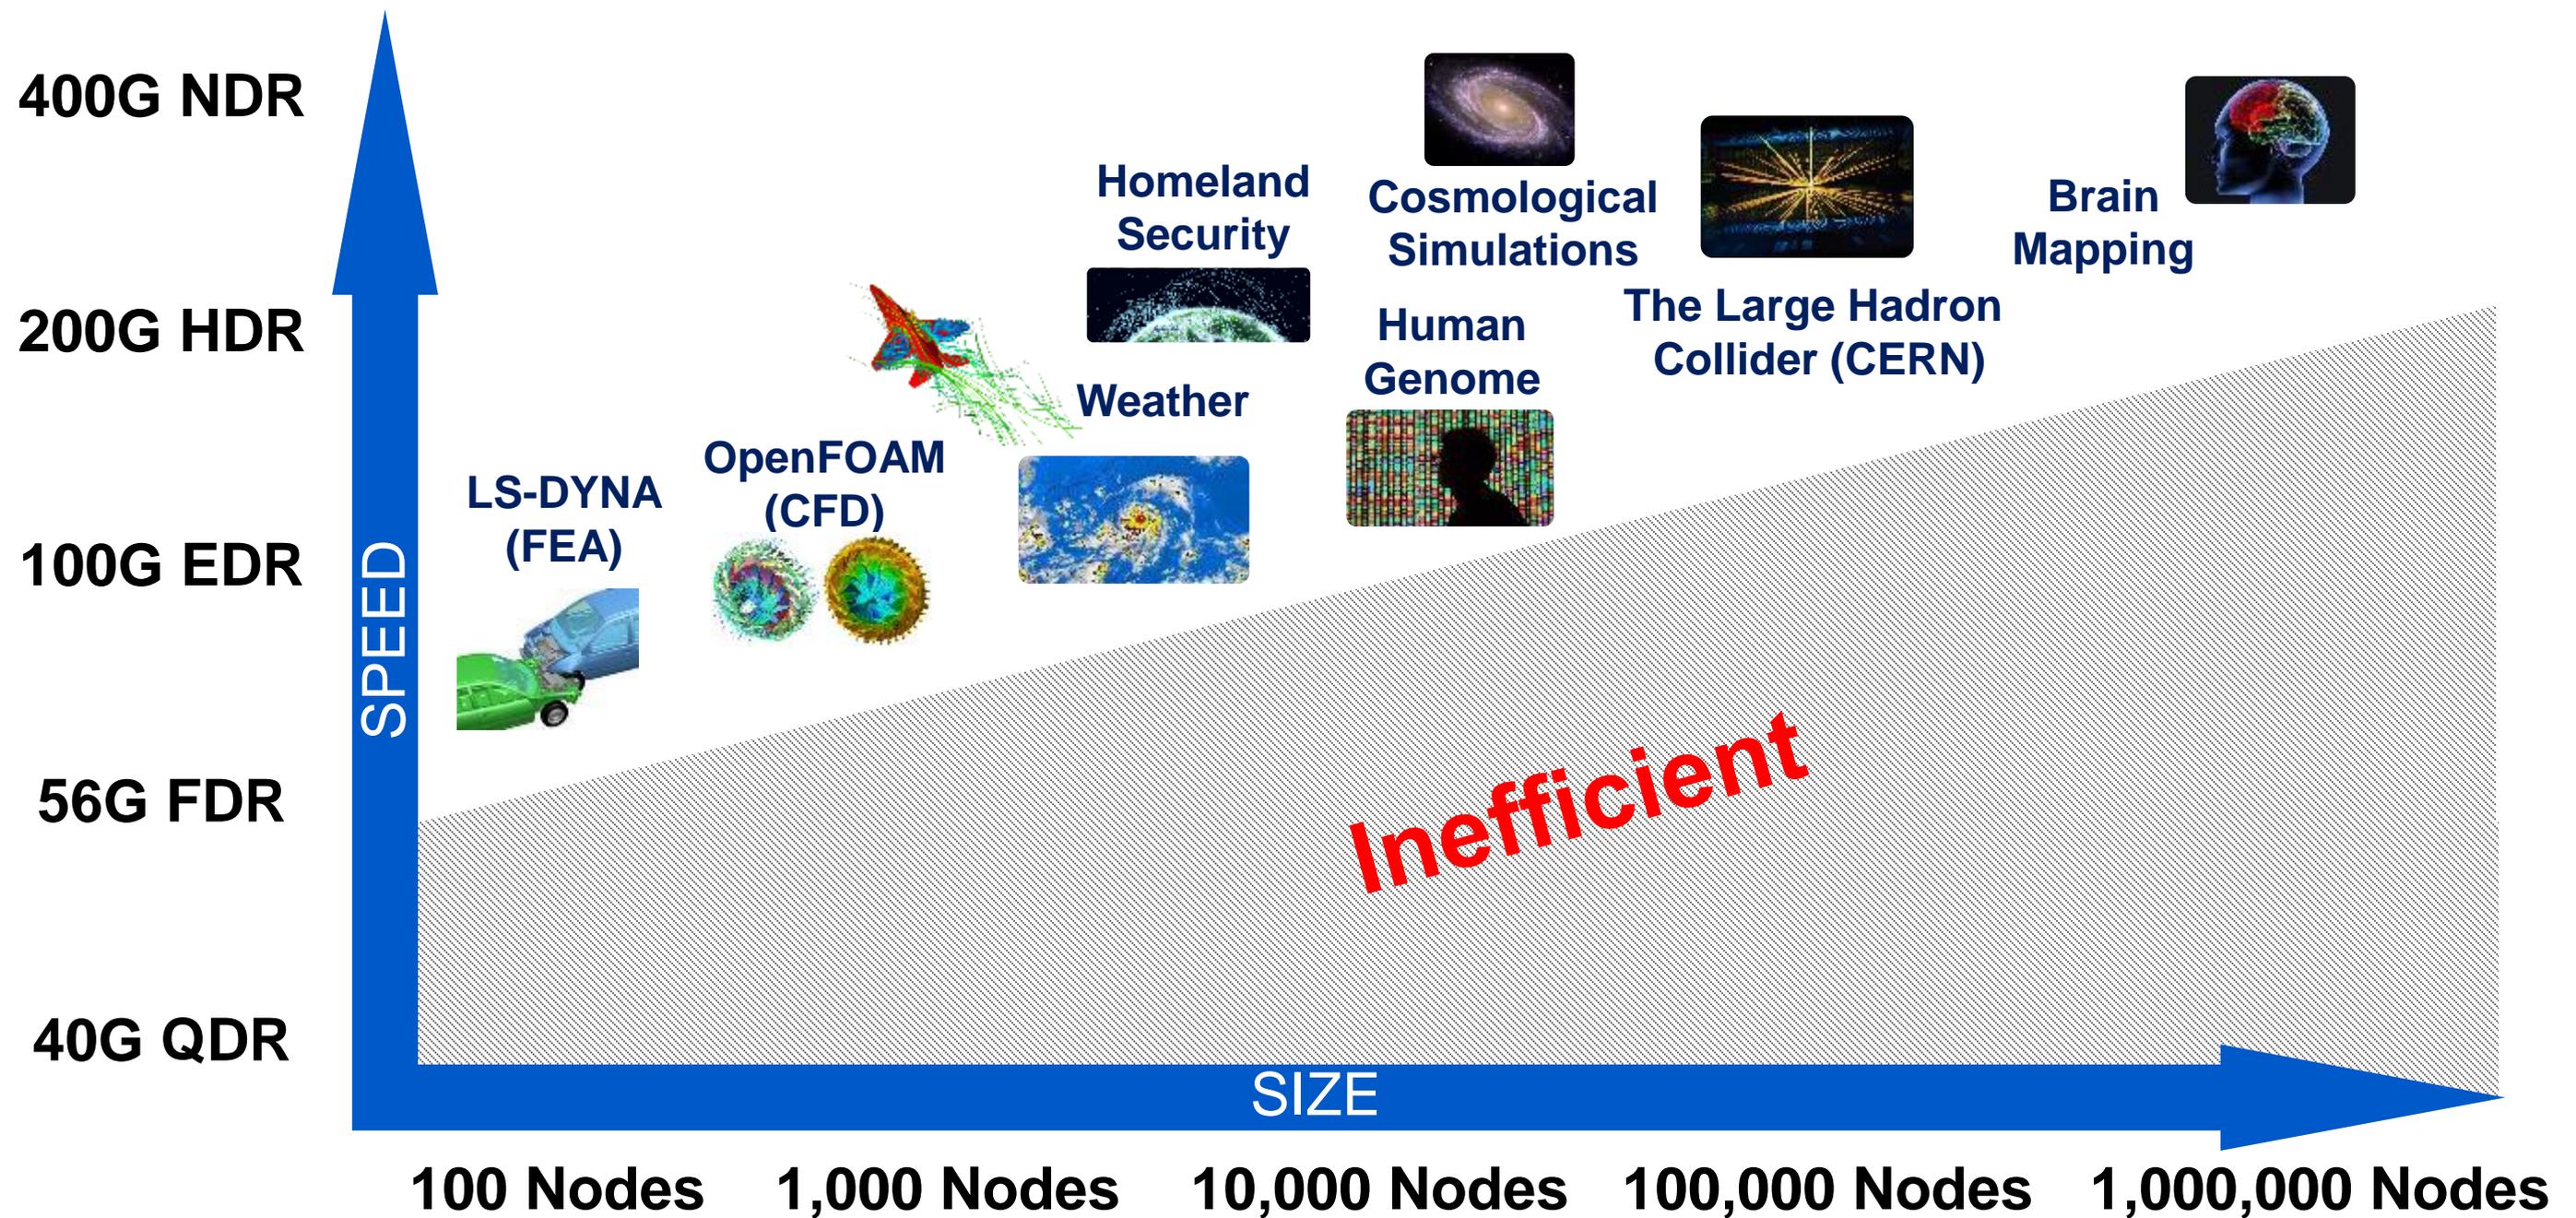
**Best Performance.
Best R.O.I.**



	OmniPath	InfiniBand
System Cost	\$245,640 With Free Network	\$146,922
\$98,718	Cost Savings	40% Lower TCO

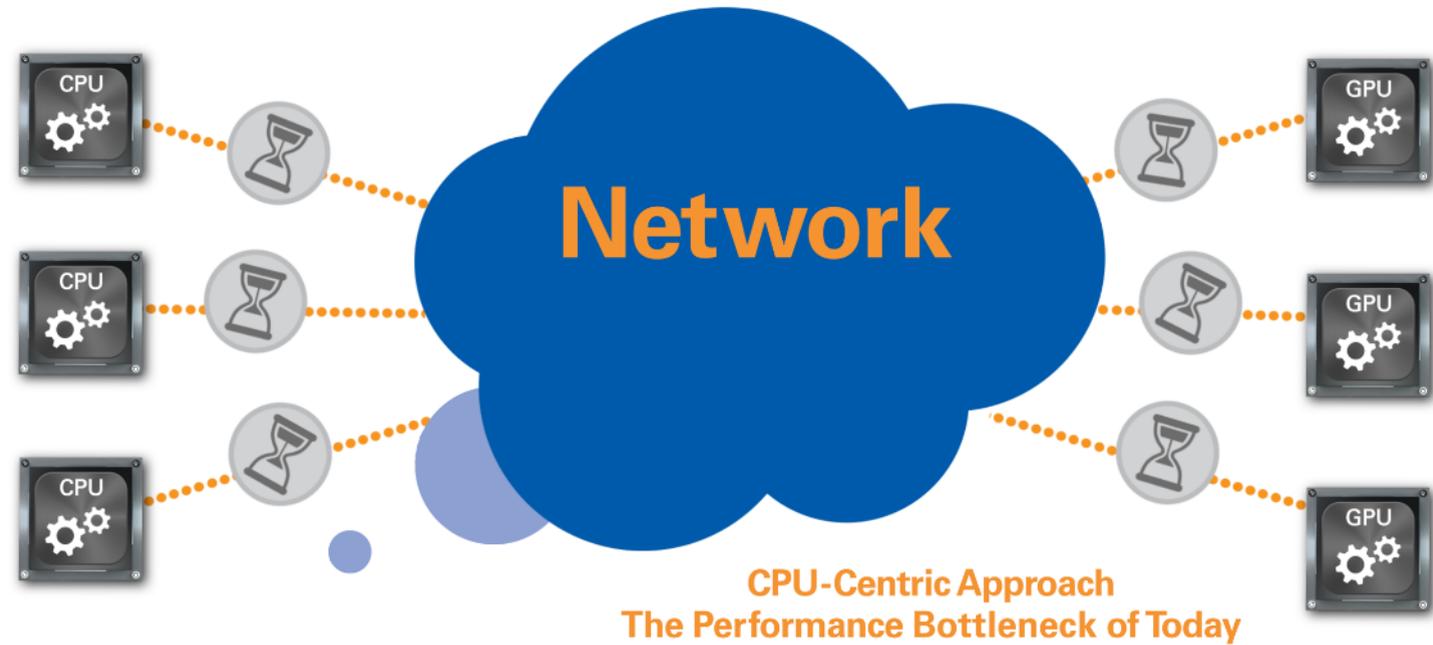
Based on 16-node Scalable Unit Performance

Interconnect Technology: The Need for Speed and Intelligence

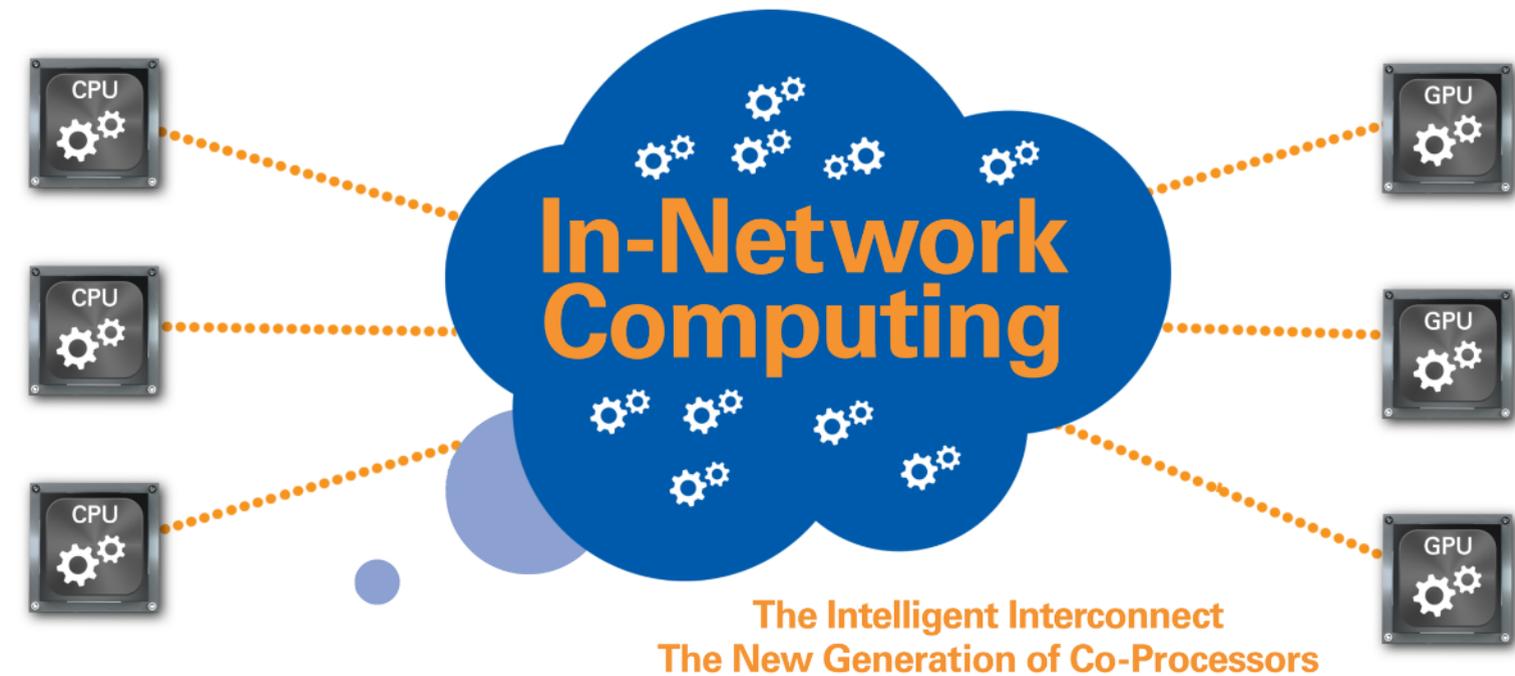


The Intelligent Interconnect Overcomes Performance Bottlenecks

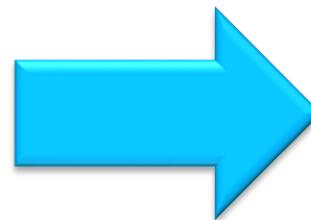
CPU-Centric (Onload)



Data-Centric (Offload)



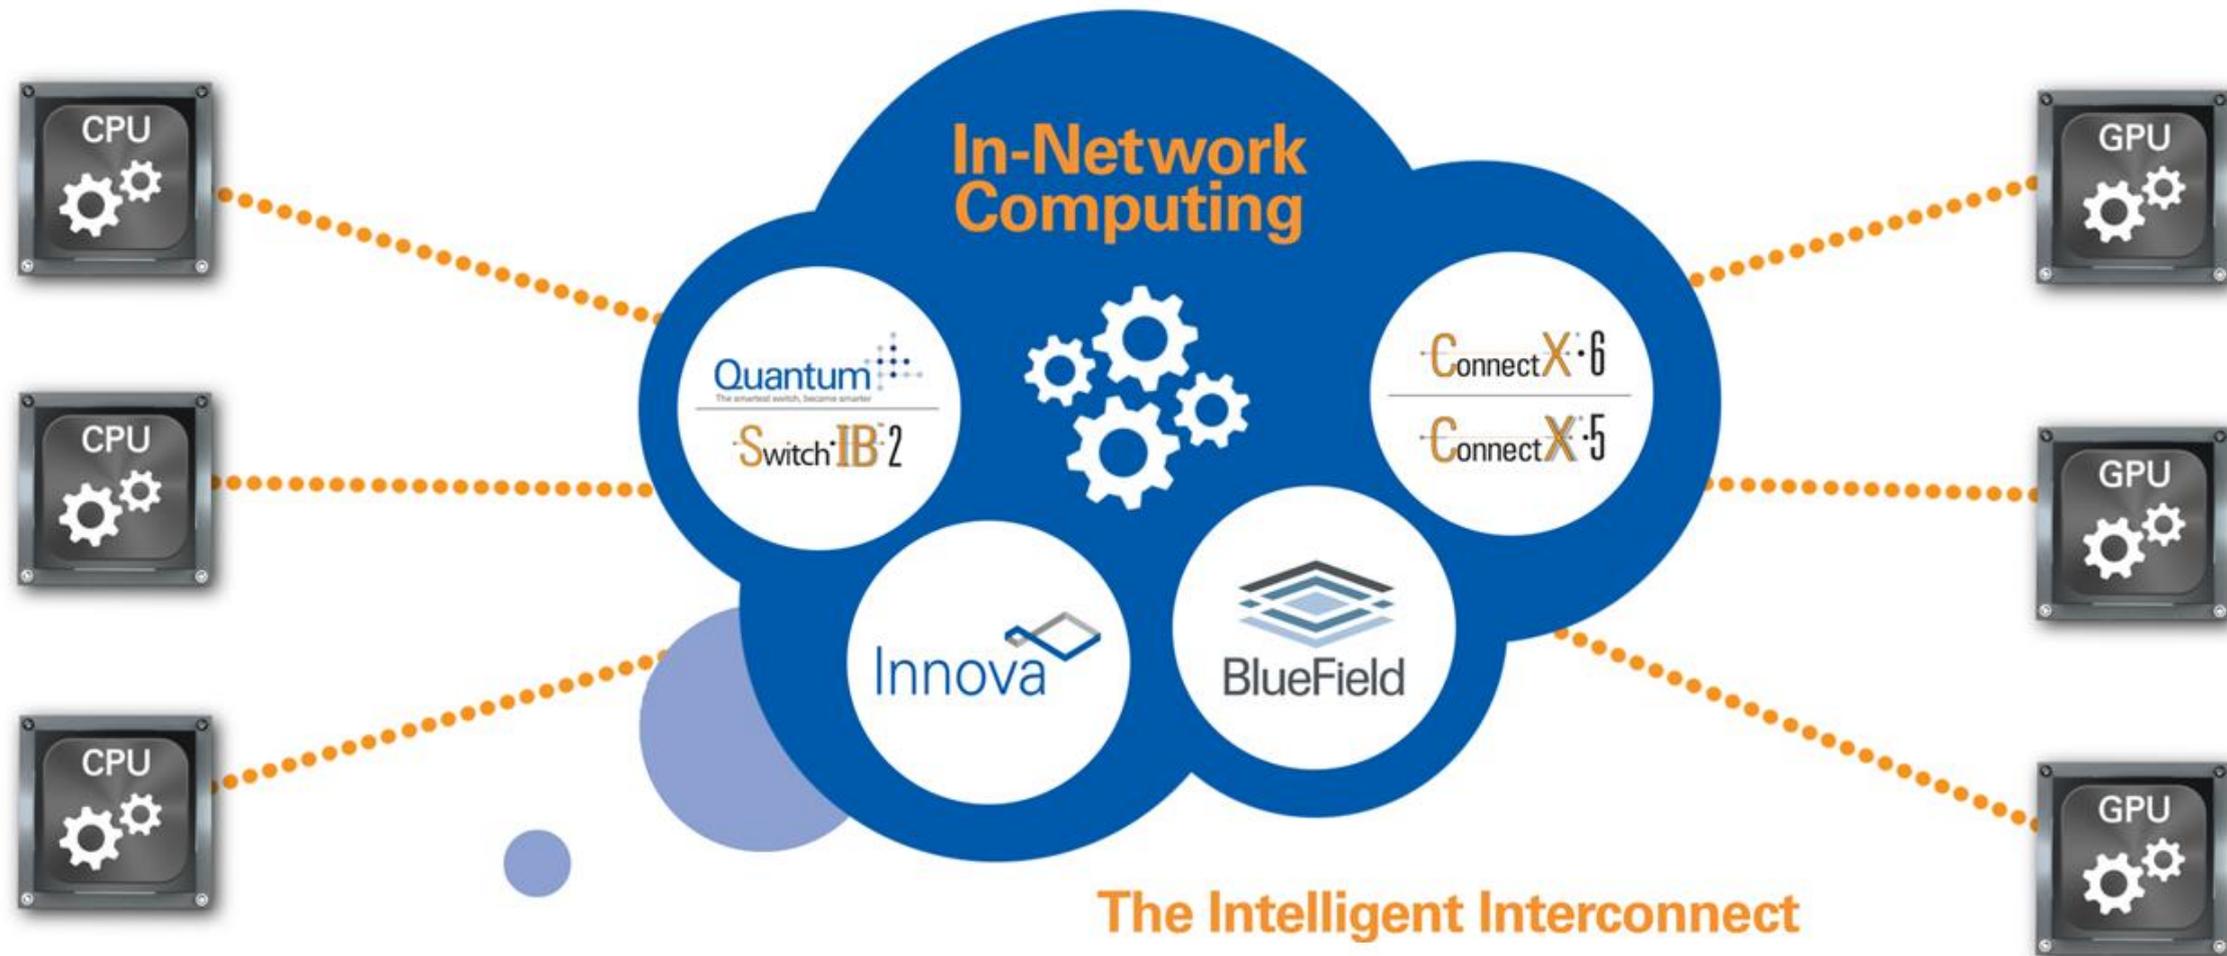
**Must Wait for the Data
Creates Performance Bottlenecks**



Analyze Data as it Moves!

Faster Data Speeds and In-Network Computing Enable Higher Performance and Scale

In-Network Computing Enables Data-Centric Data Center



Faster Data Speeds and In-Network Computing Enable Higher Performance and Scale

Highest-Performance 100Gb/s Interconnect Solutions

Adapters

ConnectX[®] 5

100Gb/s Adapter, 0.6us latency
175-200 million messages per second
(10 / 25 / 40 / 50 / 56 / 100Gb/s)



Switch

SwitchIB[™] 2

36 EDR (100Gb/s) Ports, <90ns Latency
Throughput of 7.2Tb/s
7.02 Billion msg/sec (195M msg/sec/port)



Switch

Spectrum[™]

32 100GbE Ports, 64 25/50GbE Ports
(10 / 25 / 40 / 50 / 100GbE)
Throughput of 3.2Tb/s



Interconnect

LinkX[™]

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



VCSELs, Silicon Photonics and Copper

Software

HPC-X[™]

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



Highest-Performance 200Gb/s Interconnect Solutions

Adapters		200Gb/s Adapter, 0.6us latency 200 million messages per second (10 / 25 / 40 / 50 / 56 / 100 / 200Gb/s)	
Switch	 The smartest switch, became smarter	40 HDR (200Gb/s) InfiniBand Ports 80 HDR100 InfiniBand Ports Throughput of 16Tb/s, <90ns Latency	
Switch		16 400GbE, 32 200GbE, 128 25/50GbE Ports (10 / 25 / 40 / 50 / 100 / 200 GbE) Throughput of 6.4Tb/s	
Interconnect		Transceivers Active Optical and Copper Cables (10 / 25 / 40 / 50 / 56 / 100 / 200Gb/s)	 VCSELs, Silicon Photonics and Copper
Software		MPI, SHMEM/PGAS, UPC For Commercial and Open Source Applications Leverages Hardware Accelerations	



HDR InfiniBand the Most Scalable Switch

1.7X Better

80-Ports Top of Rack Switch

2.8X Better

3200 Nodes with Only 2-Tier Switch Network

4.6X Better

128K Nodes with Only 3-Tier Switch Network

400-Nodes



1.6X Switch Saving, 2X Cable Saving

InfiniBand: 15 ToR Switches; OmniPath: 24 ToR Switches

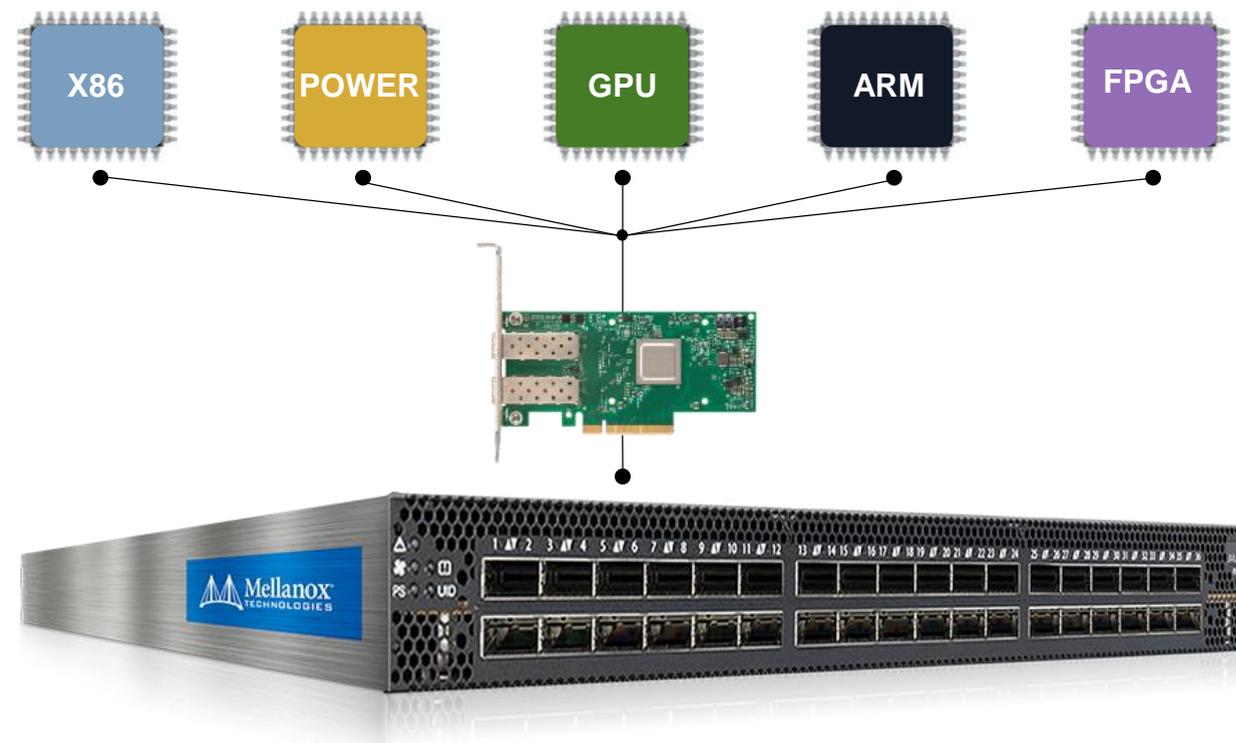
1600-Nodes



4X Cable Saving, 2X Power Saving

InfiniBand: 1 Modular Switch; OmniPath: 64 ToR + 2 Modular Switches

Highest Performance and Scalability for X86, Power, GPU, ARM and FPGA-based Compute and Storage Platforms 10, 20, 25, 40, 50, 56, 100 and 200Gb/s Speeds



Smart Interconnect to Unleash The Power of All Compute Architectures

InfiniBand The Smart Choice for HPC Platforms and Applications



“We chose a co-design approach supporting in the best possible manner our key applications. The **only interconnect that really could deliver** that was Mellanox’s InfiniBand.”



[Watch Video](#)

“In HPC, the processor should be going 100% of the time on a science question, not on a communications question. This is why the **offload capability of Mellanox’s network is critical.**”



[Watch Video](#)

“One of the big reasons we use InfiniBand and not an alternative is that we’ve got **backwards compatibility** with our existing solutions.”



[Watch Video](#)

“InfiniBand is the most advanced interconnect technology in the world, with dramatic communication overhead reduction that **fully unleashes cluster performance.**”



[Watch Video](#)

“We have users that move tens of terabytes of data and this needs to happen very, very rapidly. **InfiniBand is the way to do it.**”



[Watch Video](#)

“InfiniBand is the best that is required for our applications. It **enhances and unlocks the potential of the system.**”



[Watch Video](#)

TOP500 Mellanox Accelerated Supercomputers (Examples)



93
Petaflop

World's Fastest Supercomputer

41K nodes

Wuxi Supercomputing Center



19
Petaflop

Japan's Fastest Supercomputer

**Japan Agency for Marine-Earth
Science and Technology**



6
Petaflop

Enabling Space Exploration

20K nodes

NASA Ames Research Center

TOP500 Mellanox Accelerated Supercomputers (Examples)



Total Exploration Production



5.3
Petaflop

Enabling Energy Explorations



National Center for
Atmospheric Research



4.8
Petaflop

Enabling Weather Research



Exploration & Production ENI S.p.A.



3.2
Petaflop

Enabling Energy Explorations

- Deliver best return on investment
- Scalable, intelligent, flexible, high performance, end-to-end connectivity
- Standards-based (InfiniBand, Ethernet), supported by large eco-system
- Supports all compute architectures: x86, Power, ARM, GPU, FPGA etc.
- Offloading architecture: RDMA, application acceleration engines, etc.
- Flexible topologies: Fat Tree, Mesh, 3D Torus, Dragonfly+, etc.
- Converged I/O: compute, storage, management on single fabric
- Backward and future compatible

The Future Depends On Smart Interconnect



Thank You