



## Paving the Road to Exascale

New Era of Performance through Co-Design






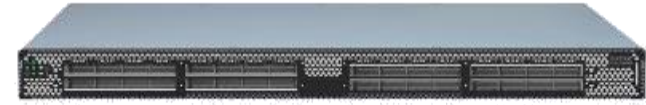




Oklahoma Supercomputing Symposium | 2015

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




# High-Performance Designed 100Gb/s Interconnect Solutions



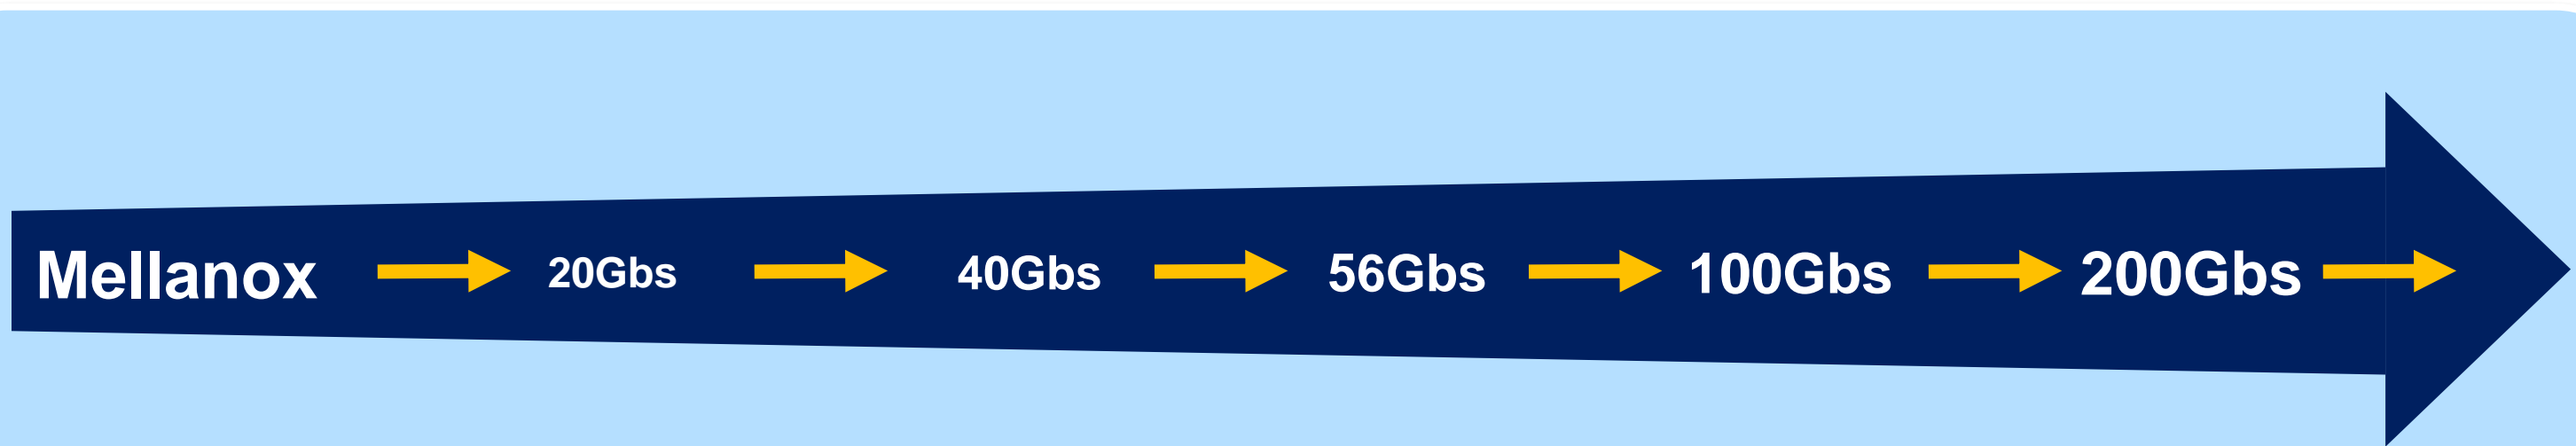
Adapters		<p><b>100Gb/s Adapter, 0.7us latency</b> <b>150 million messages per second</b> <b>(10 / 25 / 40 / 50 / 56 / 100Gb/s)</b></p>		
Switch		<p><b>36 EDR (100Gb/s) Ports, &lt;90ns Latency</b> <b>Throughput of 7.2Tb/s</b></p>		
Switch		<p><b>32 100GbE Ports, 64 25/50GbE Ports</b> <b>(10 / 25 / 40 / 50 / 100GbE)</b> <b>Throughput of 6.4Tb/s</b></p>		
Interconnect		 <b>Copper (Passive, Active)</b>	 <b>Optical Cables (VCSEL)</b>	 <b>Silicon Photonics</b>

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



**Smart Interconnect to Unleash The Power of All Compute Architectures**

# Technology Roadmap – One-Generation Lead over the Competition



## Terascale

3<sup>rd</sup>



**TOP500 2003**  
Virginia Tech (Apple)

1<sup>st</sup>



**“Roadrunner”**  
Mellanox Connected

## Petascale



## Exascale

**OAK RIDGE**  
National Laboratory  
“Summit” System

**Lawrence Livermore**  
National Laboratory  
“Sierra” System

2000


2005

2010

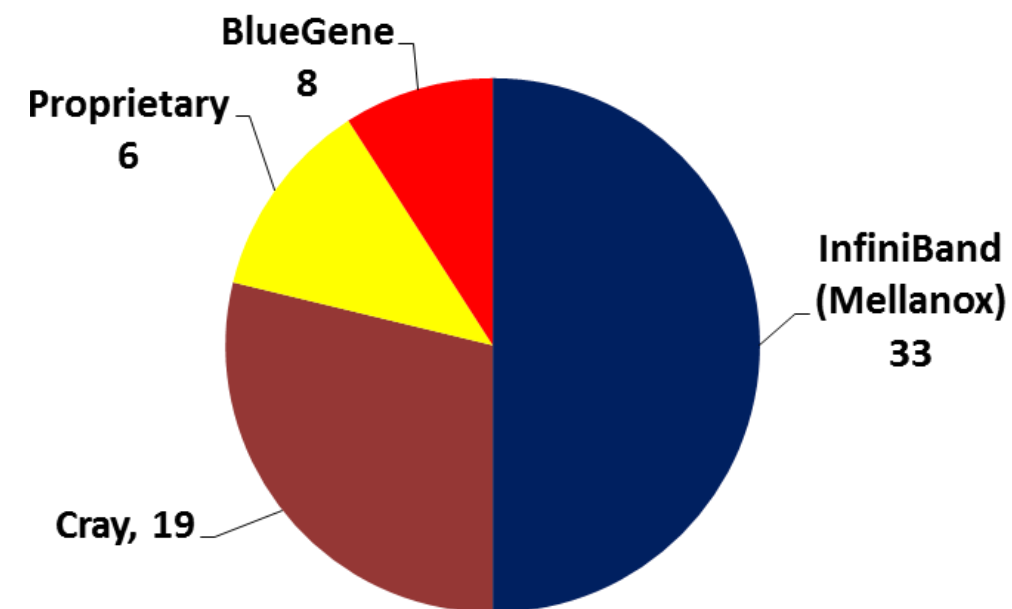
2015

2020

## TOP500 Interconnect Trends



## PetaFlops Systems on the TOP500 list

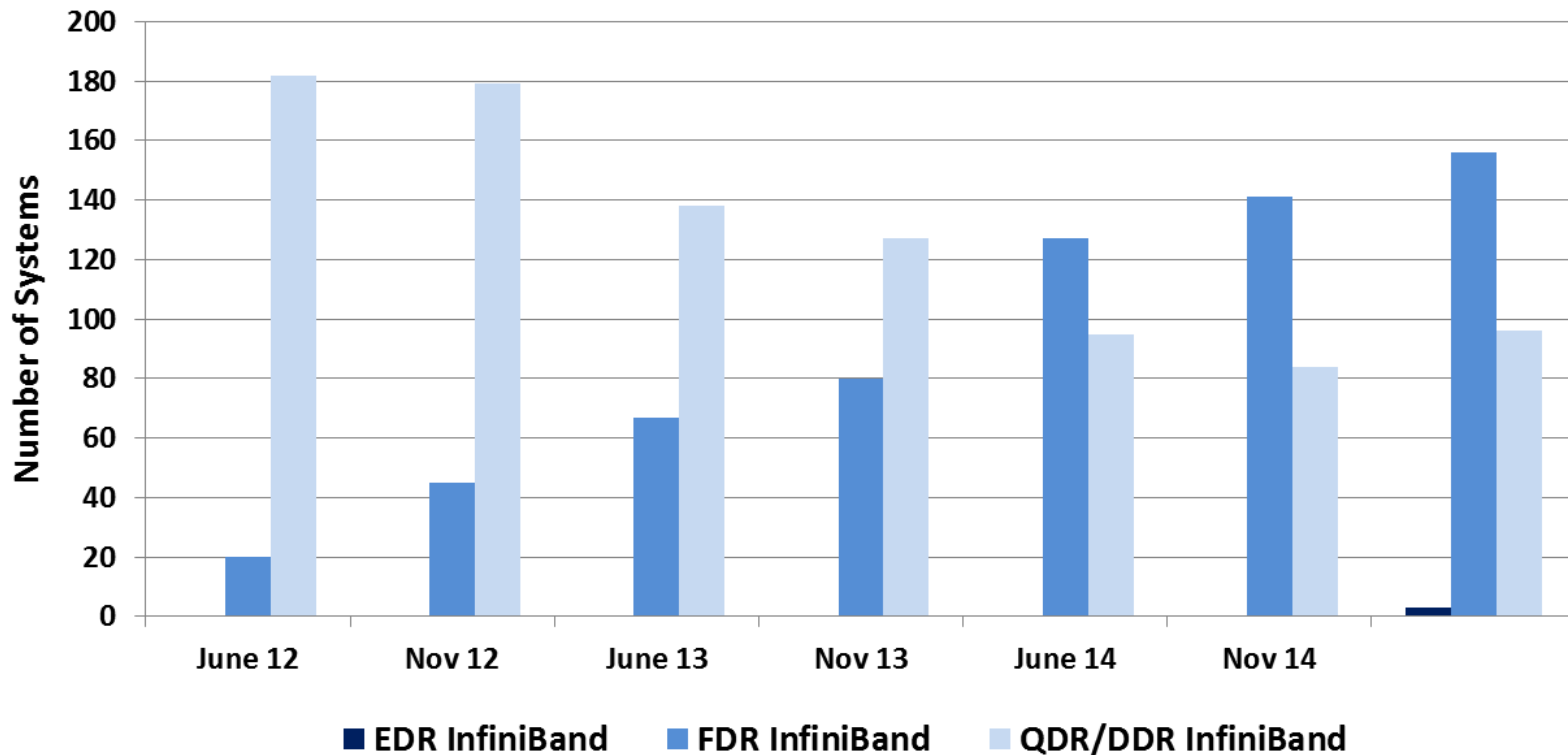


- InfiniBand is the de-facto interconnect solution for performance demanding applications

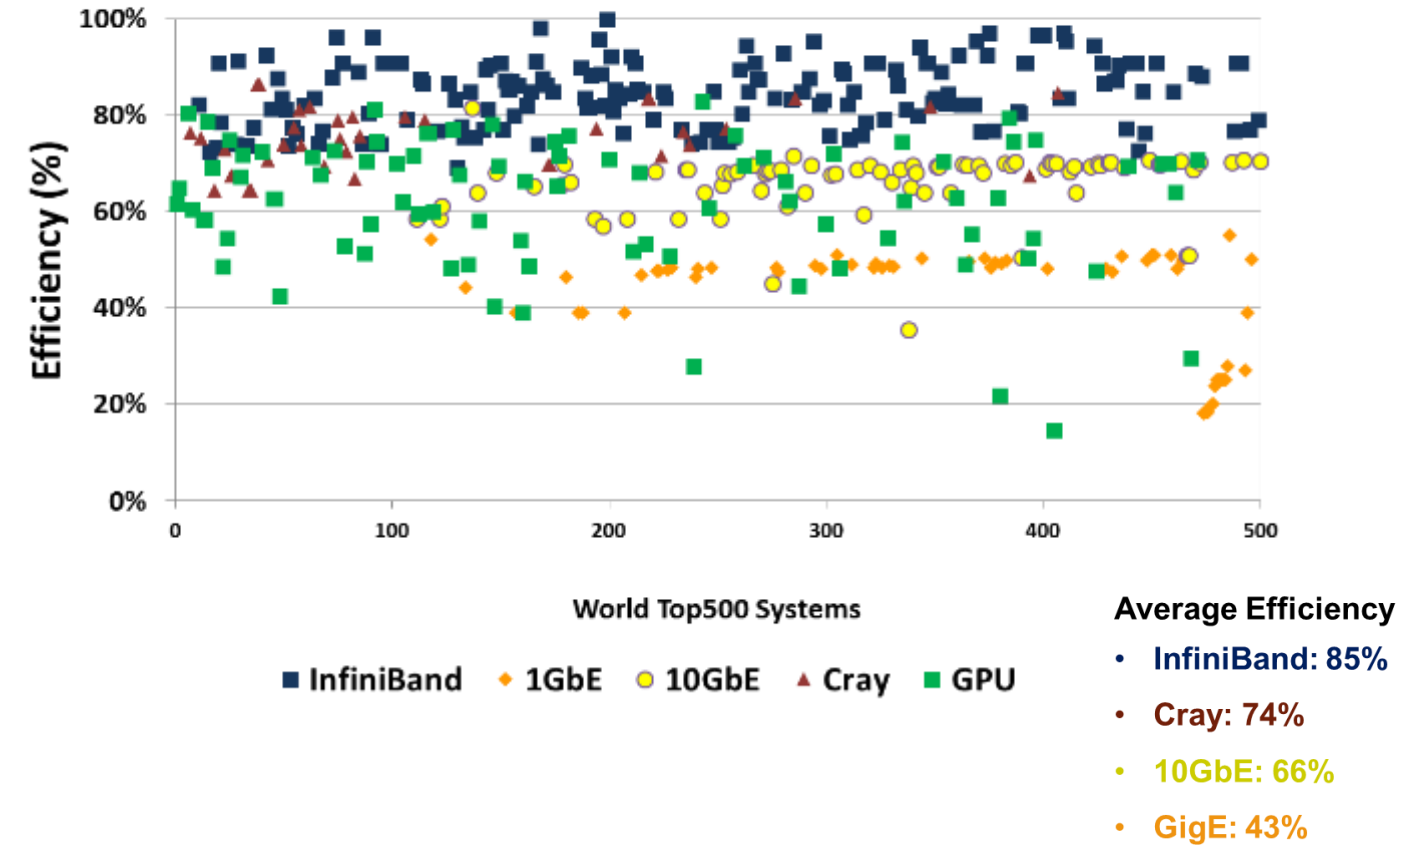
# TOP500 InfiniBand Accelerated Systems



## InfiniBand Accelerated TOP500 Systems



## World Leading Compute Systems Efficiency Comparison



- Number of Mellanox FDR InfiniBand systems grew 23% from June'14 to July'15
- EDR InfiniBand entered the list with 3 systems




# Mellanox Accelerated World-Leading HPC Systems




Connecting Half of the World's Petascale Systems (examples)

# The Road to Exascale Computing





**Discrete Level**



**System Level**



# Exascale will be Enabled via Co-Design Architecture



Software – Hardware



Hardware – Hardware (e.g. GPU-Direct)



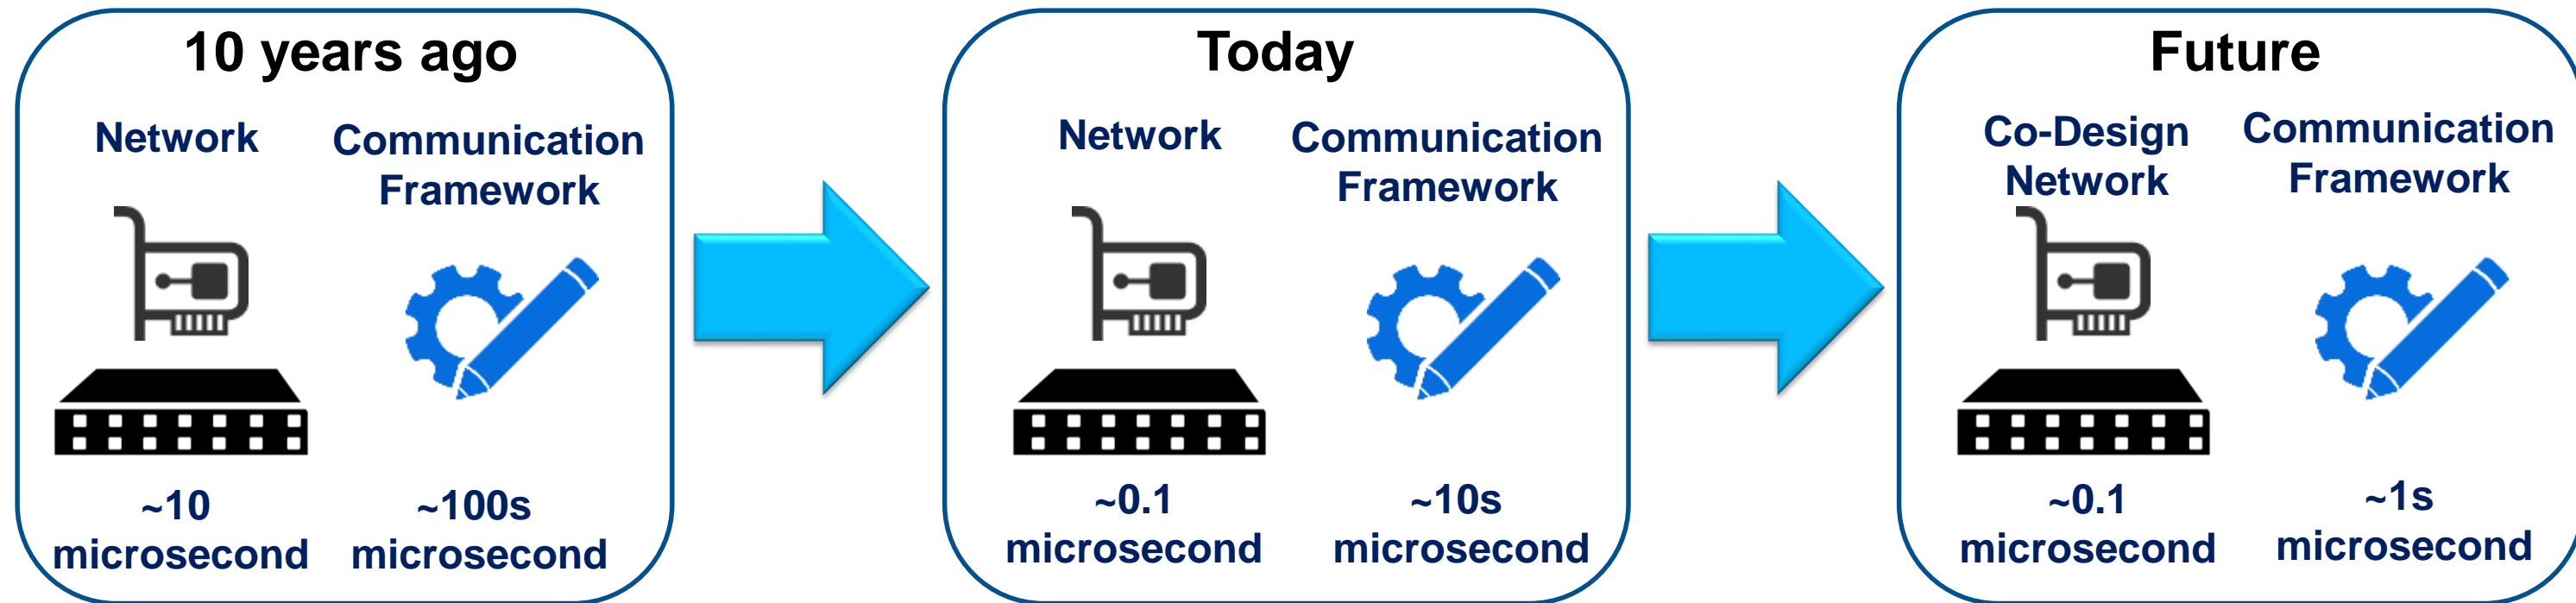
Software – Software (e.g. OpenUCX)



Industry – Users – Academia



**Standard, Open Source, Eco-System**  
**Programmable, Configurable, Innovative**



- Today: Network devices are in 100ns latency today
- Challenge: How to enable the next order of magnitude performance improvement?
- Solution: Co-Design - mapping the communication frameworks on all active devices
- Result: reduce HPC communication frameworks latency by an order of magnitude

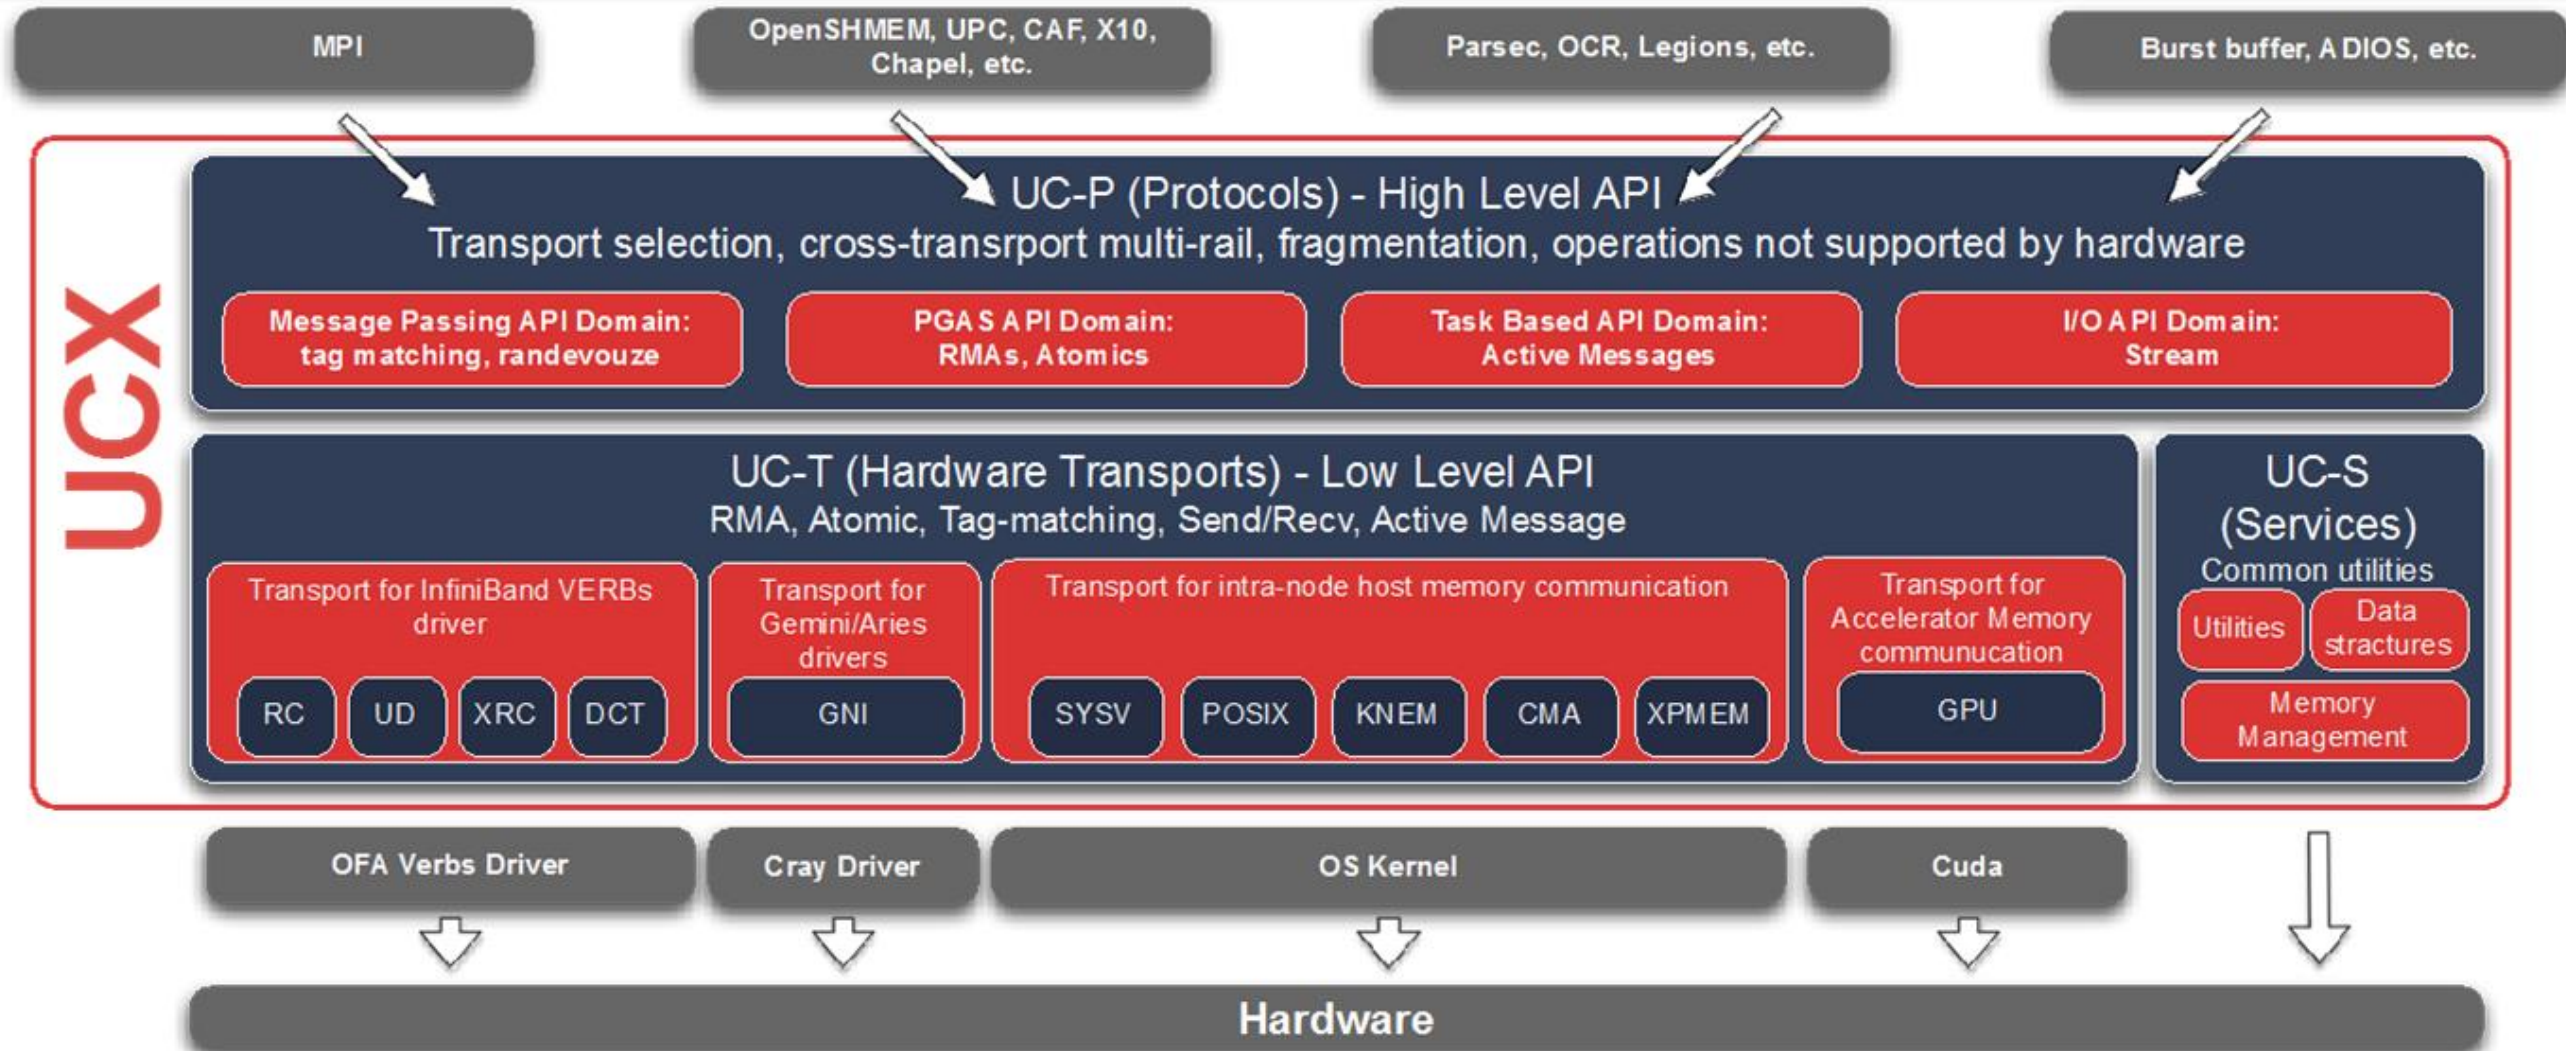
**Co-Design Architecture Paves the Road to Exascale Performance**



# The Next Generation HPC Software Framework

Collaborative Effort  
Industry, National Laboratories and Academia

## Applications



**Unified, Light-Weight, High-Performance Communication Framework**

## • HPC-X™

- Complete MPI, PGAS OpenSHMEM and UPC package
- Maximize application performance
- For commercial and open source applications
- Based on UCX (Unified Communication – X Framework)

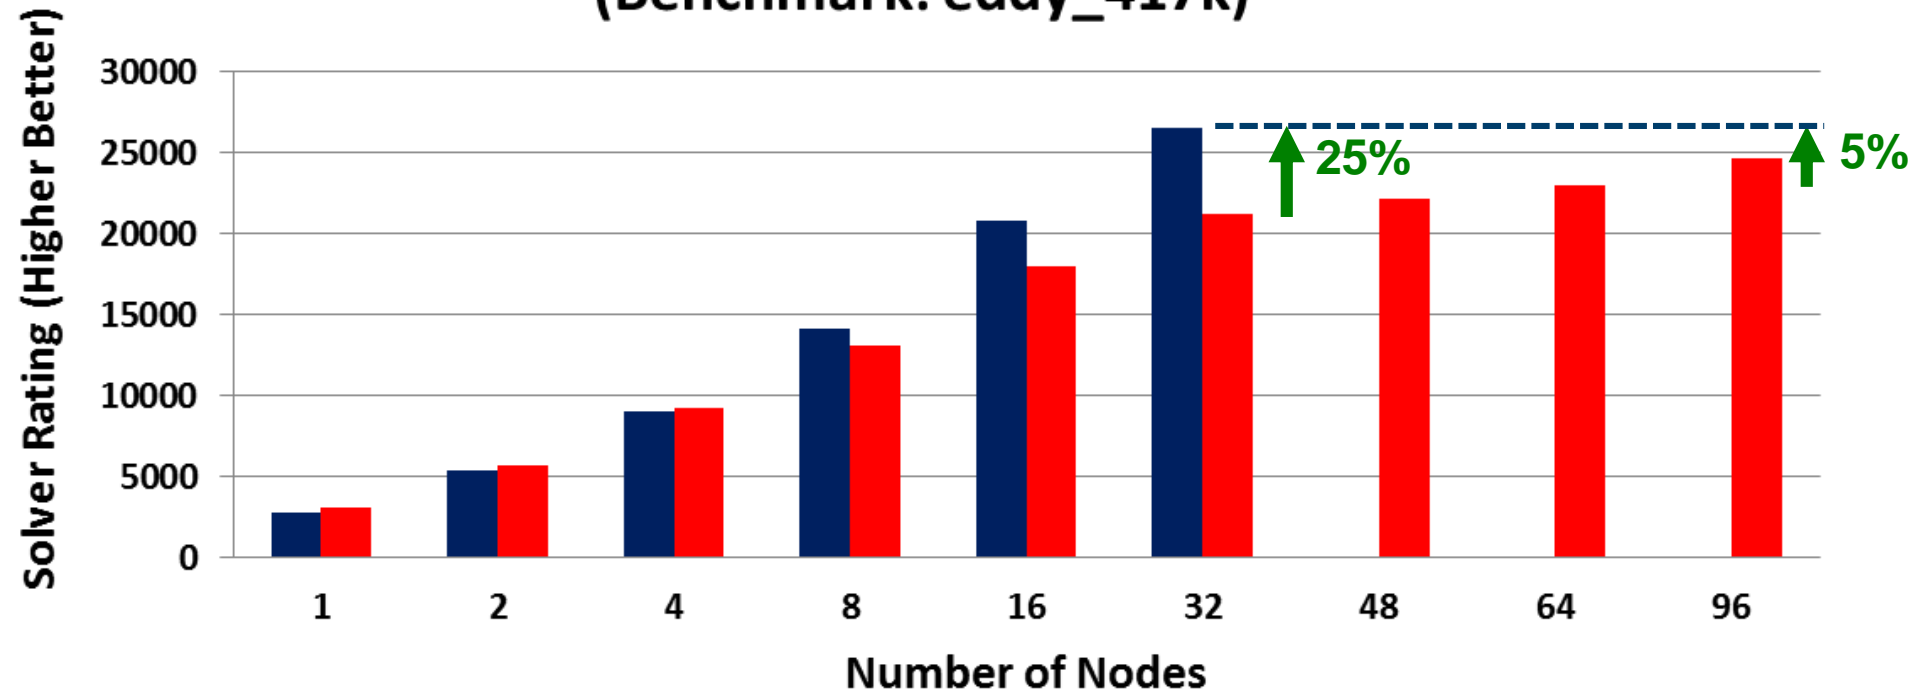





# Mellanox Delivers World Record Application Performance



## ANSYS Fluent 15.0.7 Performance (Benchmark: eddy\_417k)




- Dell PowerEdge, FDR InfiniBand, Intel Xeon 2680v2 at 2.8GHz
- Cray XC30, Aries interconnect, Intel Xeon 2690v2 at 3GHz




# The Performance Power of Co-Design


**OptiStruct Performance**  
(Engine\_Assy.fem)




■ FDR InfiniBand ■ EDR InfiniBand




**WRF Performance**  
(conus12km)




■ FDR InfiniBand ■ EDR InfiniBand



**GROMACS Performance**  
(d.dppc)



■ FDR InfiniBand ■ EDR InfiniBand



**40% Increase in Overall System Performance**

System: \$315K, InfiniBand FDR: \$35K (11%), InfiniBand EDR: \$52K (16%)

For 5% System Investment (\$17K), 40% Performance Increase! (\$126K value, 30% - \$95K)

- Mellanox solutions provide a proven, scalable and high performance end-to-end connectivity
- Flexible, support all compute architectures: x86, Power, ARM, GPU, FPGA etc.
- Standards-based (InfiniBand, Ethernet), supported by large eco-system
- Higher performance: 100Gb/s, 0.7usec latency, 150 million messages/sec
- HPC-X is leading HPC software based on UCX, which provides leading performance for MPI, OpenSHMEM/PGAS and UPC
- Exascale will be Enabled via Co-Design Architecture
- Only Mellanox provides superior applications offloads: RDMA, Collectives, scalable transport
- Backward and future compatible





Thank You