High Performance Computing

Scot Schultz, Director HPC/Technical Computing

April 2015

Mellanox TECHNOLOGIES

Connect. Accelerate. Outperform.™

Interconnect Solutions Leadership – Software and Hardware

Comprehensive End-to-End Interconnect Software Products

Comprehensive End-to-End InfiniBand and Ethernet Hardware Products

ICs	Adapter Cards	Switches/Gateways	Metro / WAN	Cables/Modules
	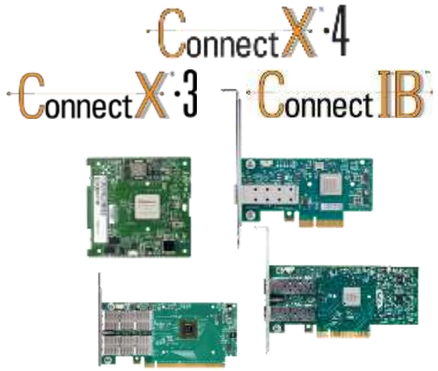			

Proud to Accelerate Future DOE Leadership Systems (“CORAL”)

U.S. DEPARTMENT OF ENERGY

OAK RIDGE National Laboratory

“Summit” System

Lawrence Livermore National Laboratory

“Sierra” System

5X – 10X Higher Application Performance versus Current Systems

Mellanox EDR 100Gb/s InfiniBand, IBM POWER CPUs, NVIDIA GPUs

Mellanox EDR 100G Solutions Selected by the DOE for 2017 Leadership Systems

Deliver Superior Performance and Scalability over Current / Future Competition

Mellanox Accelerated World-Leading HPC Systems

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

System Example: NASA Ames Research Center Pleiades

- 20K InfiniBand nodes
- Mellanox end-to-end scalable FDR and QDR InfiniBand
- Supports variety of scientific and engineering projects
  - Coupled atmosphere-ocean models
  - Future space vehicle design
  - Large-scale dark matter halos and galaxy evolution
- Leveraging InfiniBand backward and future compatibility

High-Resolution Climate Simulations

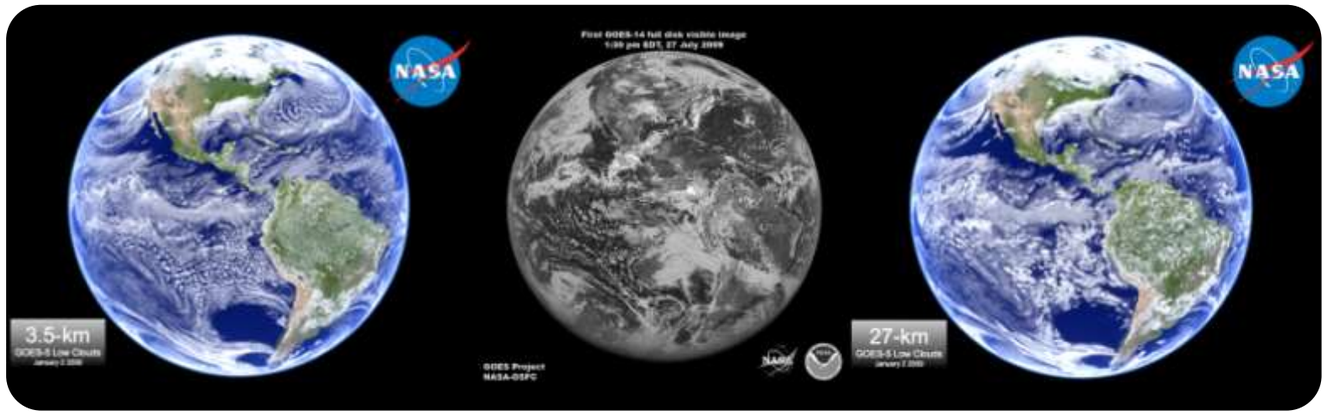

The Future is Here

Entering the Era of 100Gb/s

Adapters

ConnectX®-4

100Gb/s Adapter, 0.7us latency

150 million messages per second

(10 / 25 / 40 / 50 / 56 / 100Gb/s)

Switch

SwitchIB™

36 EDR (100Gb/s) Ports, <90ns Latency

Throughput of 7.2Tb/s

Interconnect

LinkX™

Copper (Passive, Active)

Optical Cables (VCSEL)

Silicon Photonics

Enter the World of Scalable Performance – 100Gb/s Adapter

ConnectX-4: Highest Performance Adapter in the Market

InfiniBand: SDR / DDR / QDR / FDR / EDR

Ethernet: 10 / 25 / 40 / 50 / 56 / 100GbE

100Gb/s, <0.7us latency

150 million messages per second

OpenPOWER CAPI technology

CORE-Direct technology

GPUDirect RDMA

Dynamically Connected Transport (DCT)

Ethernet offloads (HDS, RSS, TSS, LRO, LSOv2)

Enter the World of Scalable Performance – 100Gb/s Switch

Switch-IB: Highest Performance Switch in the Market

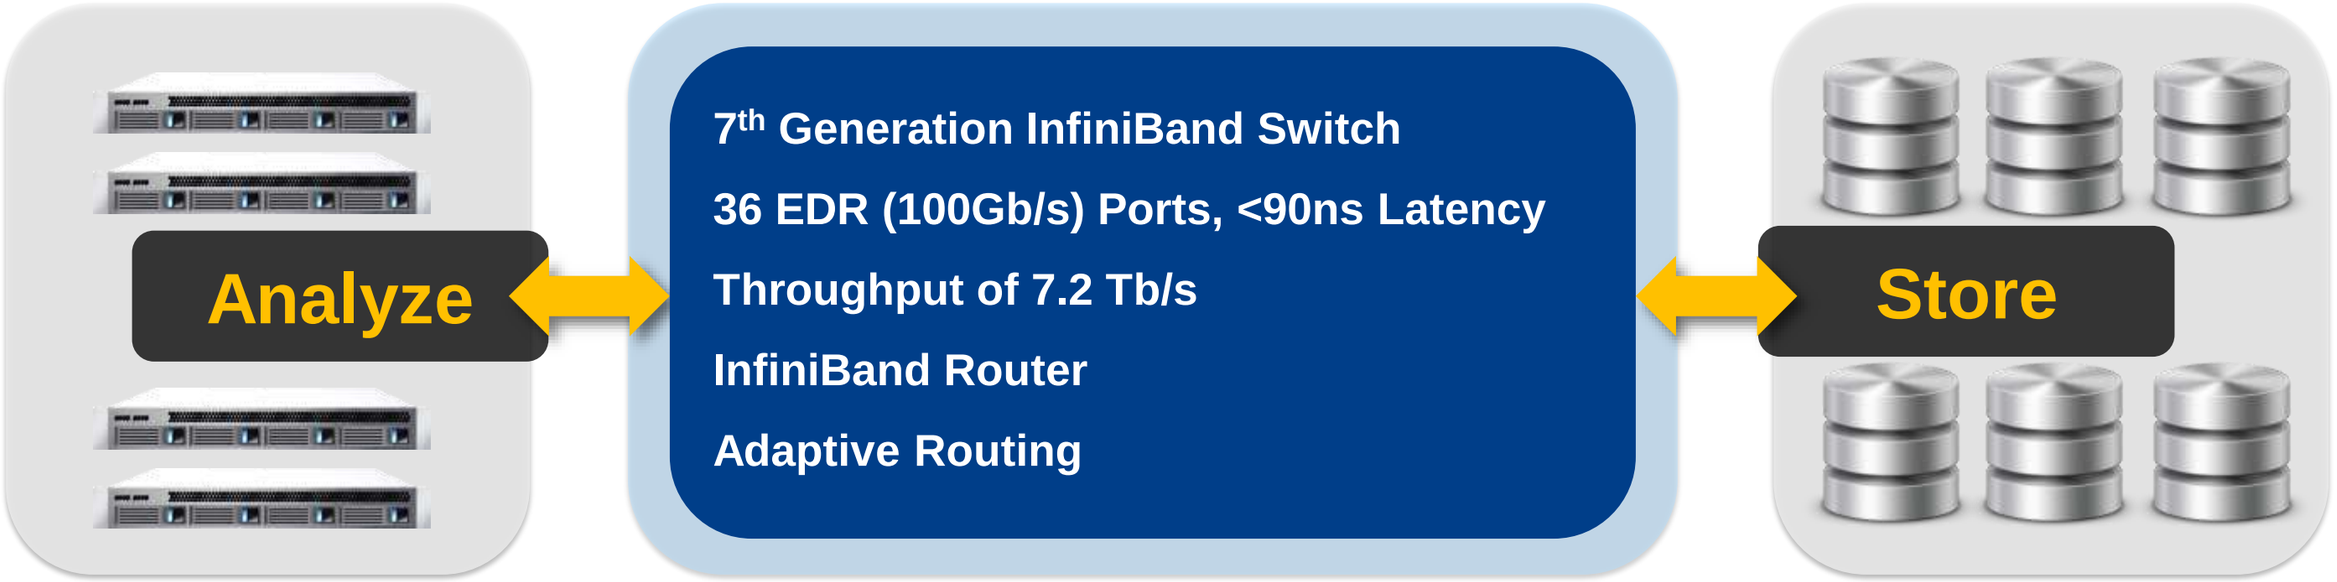

SwitchIB™

Mellanox QSFP 100Gb/s Cables

Complete Solution of 100Gb/s Copper and Fiber Cables

Copper Cables

VCSEL AOCs

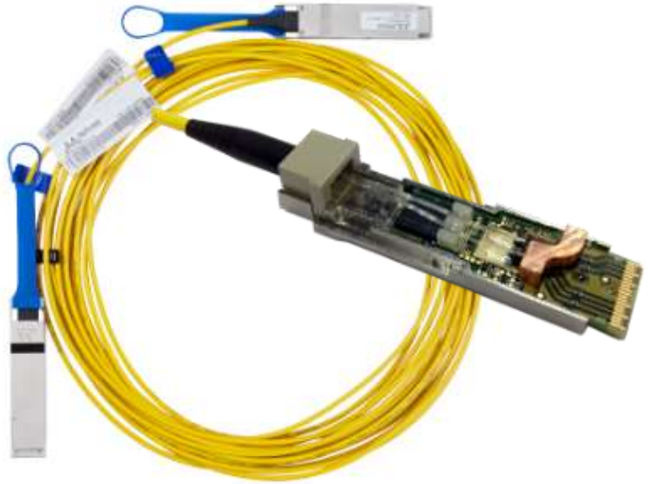

Silicon Photonics AOCs

LinkX™

Making 100Gb/s Deployments as Easy as 10Gb/s

Mellanox HPC-X™ Components and Accelerators

HPC-X™

Mellanox ScalableMPI - Message Passing Interface based on Open MPI
Mellanox ScalableSHMEM - One-sided Communications Library
Mellanox ScaleableUPC – Berkeley UPC parallel programming language library
Mellanox MXM – Messaging Accelerator optimized for underlying hardware
Mellanox FCA – Fabric Collectives Accelerator supporting MPI-3
Profiling Tools …IPM, Benchmarking Tools….and MORE…

Mellanox Delivers Highest Application Performance

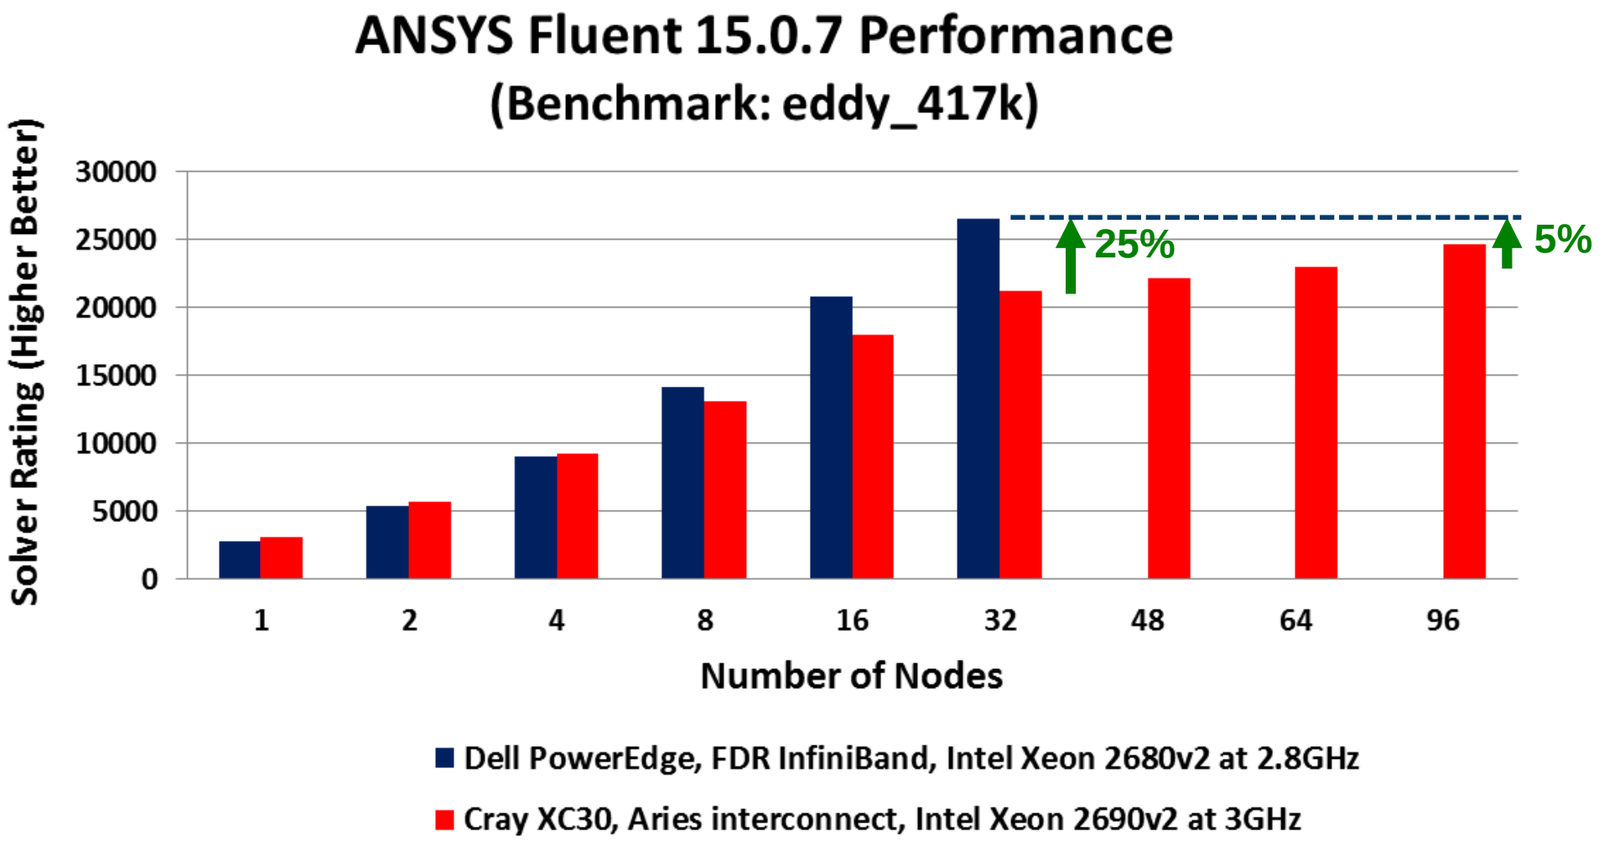

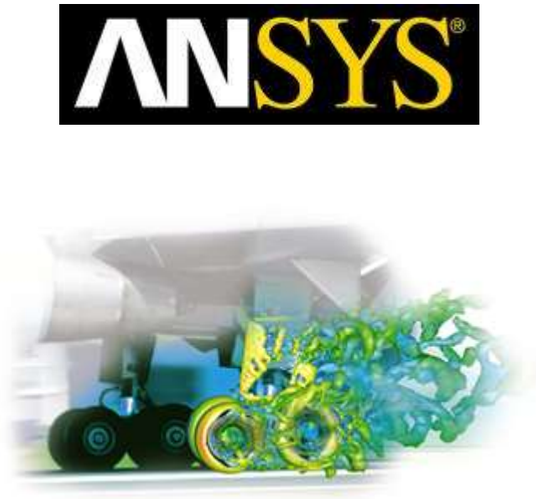

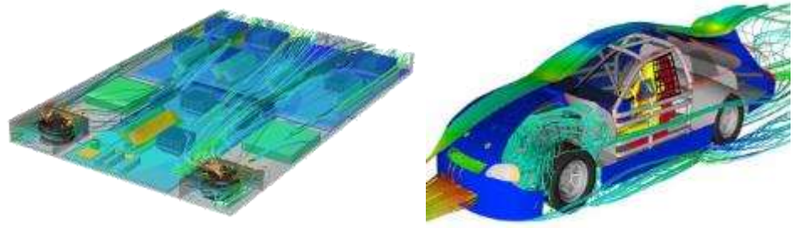

HPC-X™ Delivers World Record Performance!

Higher Performance With 1/3 of The Compute Infrastructure

HPC-X Performance Advantage

58% Performance Advantage!

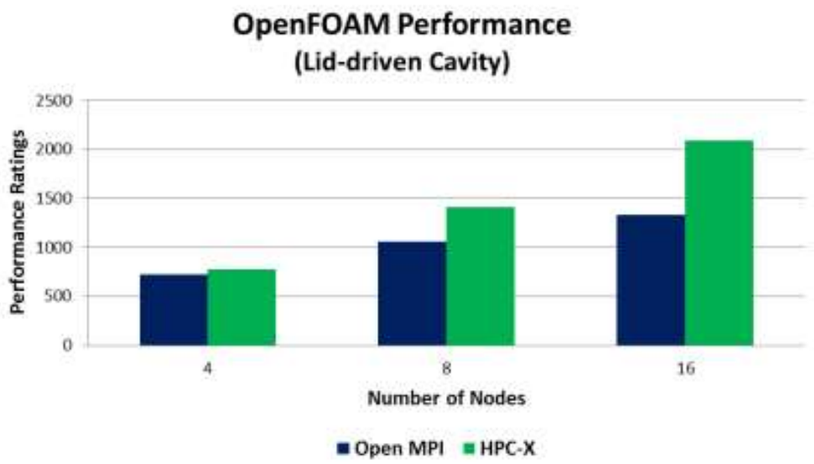

21% Performance Advantage!

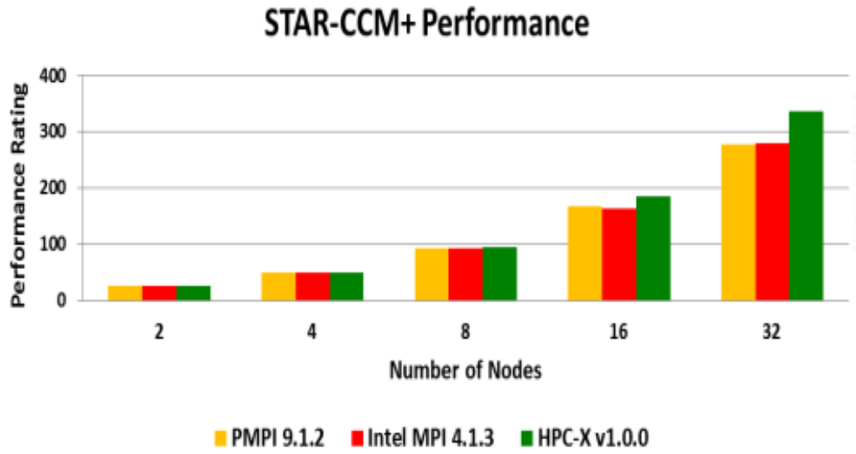

20% Performance Advantage!

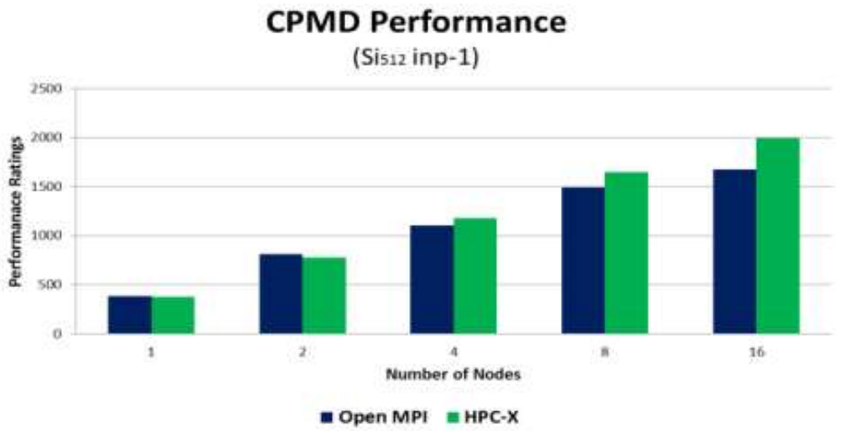

OpenFOAM

CD-adapco

CPMD

Enabling Highest Applications Scalability and Performance

HPC-X Performance Advantage – LS-DYNA

17% Performance Advantage!

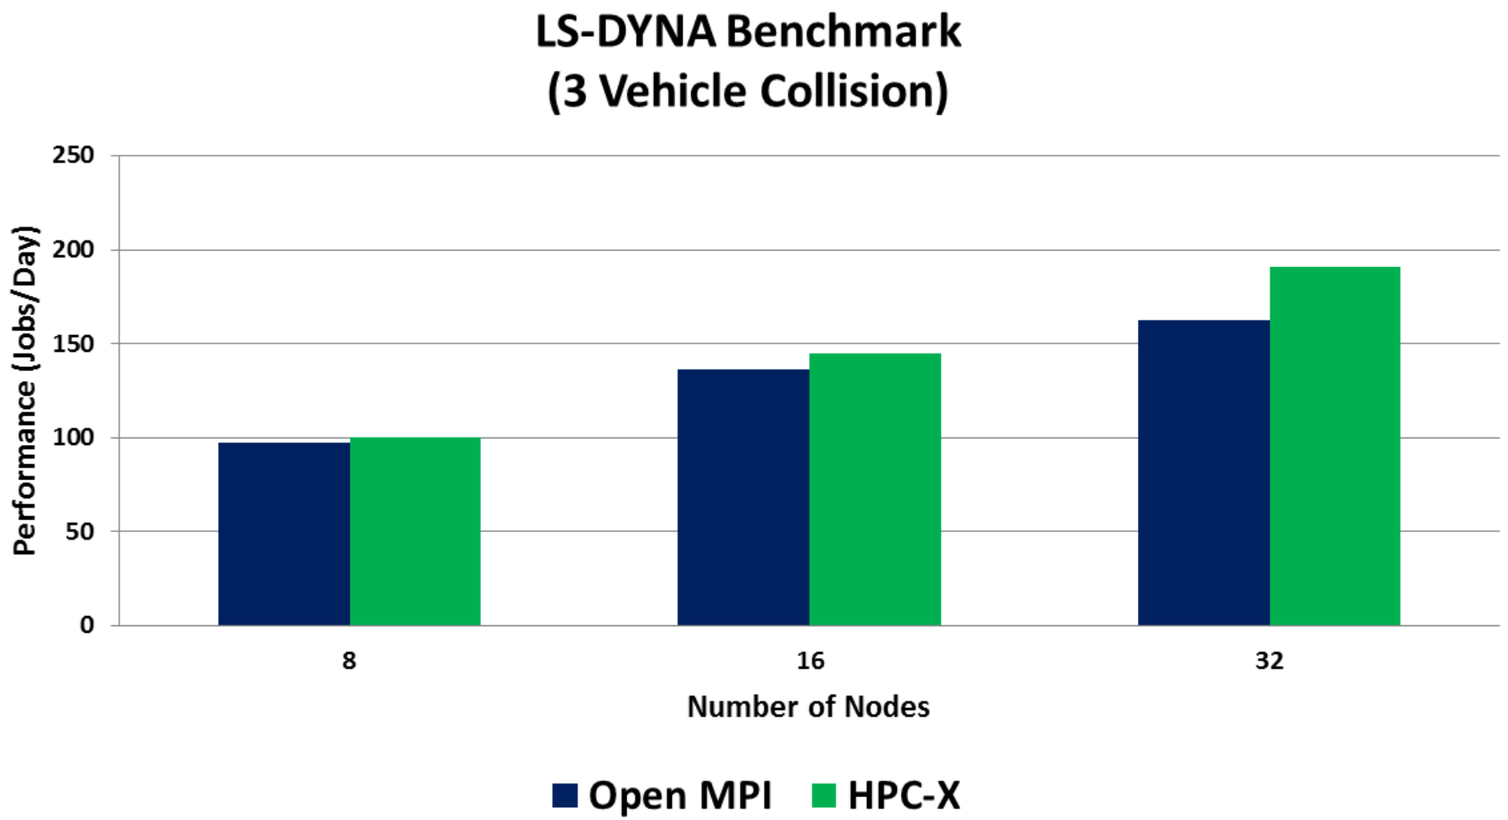

HPC-X™

LSTC
Livermore Software Technology Corp.

HPC-X™ Performance for Commercial FEA

Mellanox Interconnect Advantages

- 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 software provides leading performance for MPI, OpenSHMEM/PGAS and UPC
- Superiors applications offloads: RDMA, Collectives, scalable transport
- Backward and future compatible

Speed-Up Your Present, Protect Your Future
Paving The Road to Exascale Computing Together

Thank You

Mellanox TECHNOLOGIES

Connect. Accelerate. Outperform.™