



Interconnect Your Future

Paving the Road to Exascale

June 2017

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

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


CPU-Centric (Onload)



Data-Centric (Offload)



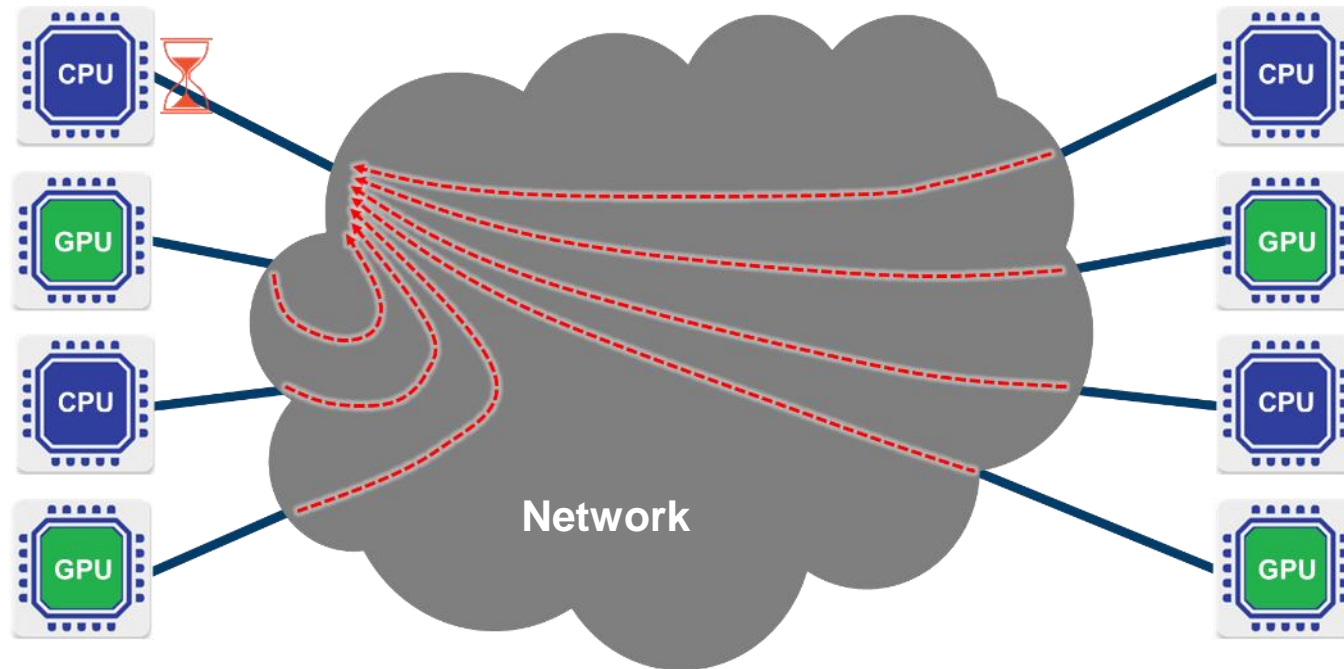
**Must Wait for the Data
Creates Performance Bottlenecks**



Analyze Data as it Moves!

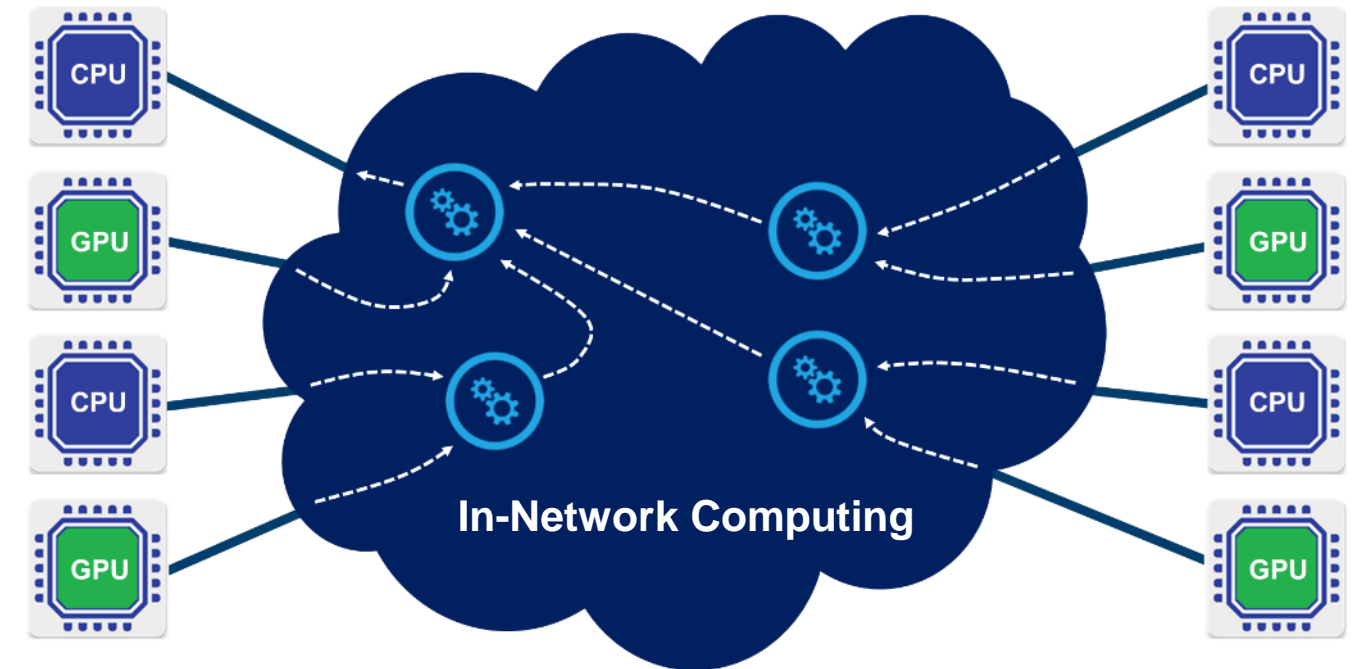
Faster Data Speeds and In-Network Computing Enable Higher Performance and Scale

CPU-Centric (Onload)



HPC / Machine Learning
Communications Latencies of 30-40us

Data-Centric (Offload)



HPC / Machine Learning
Communications Latencies of 3-4us

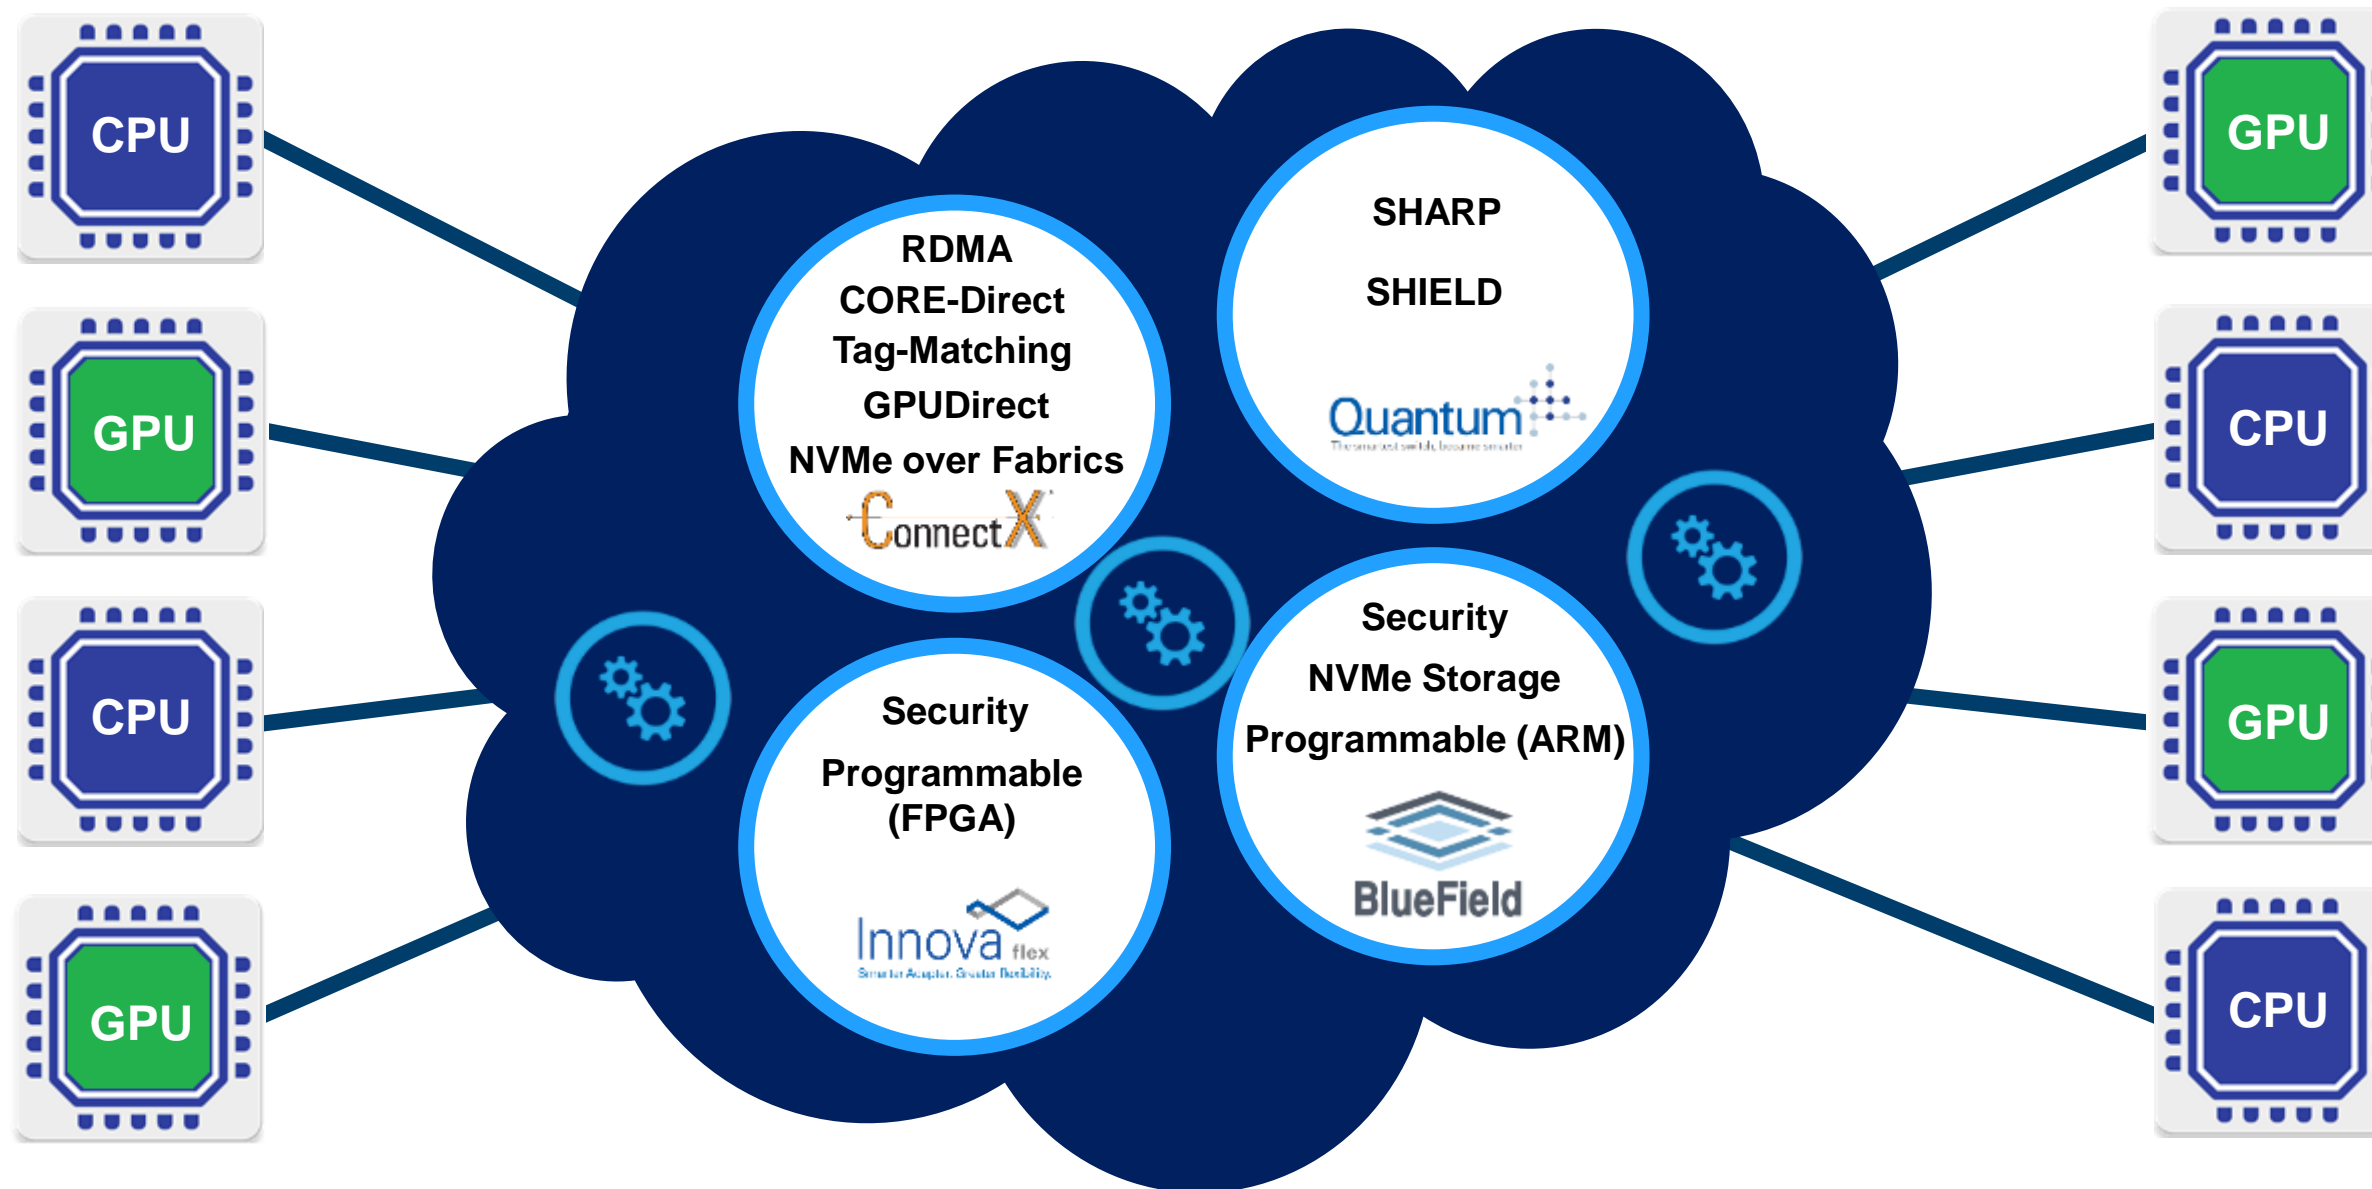
Intelligent Interconnect Paves the Road to Exascale Performance

In-Network Computing to Enable Data-Centric Data Center



In-Network Computing Key for Highest Return on Investment

In-Network Computing to Enable Data-Centric Data Centers



In-Network Computing Key for Highest Return on Investment

InfiniBand
Just Got
Smarter

In-Network Computing

SHARP



10X Performance Acceleration

Critical for High Performance Computing and Machine Learning Applications

InfiniBand
Just Got
Smarter


Self-Healing Technology



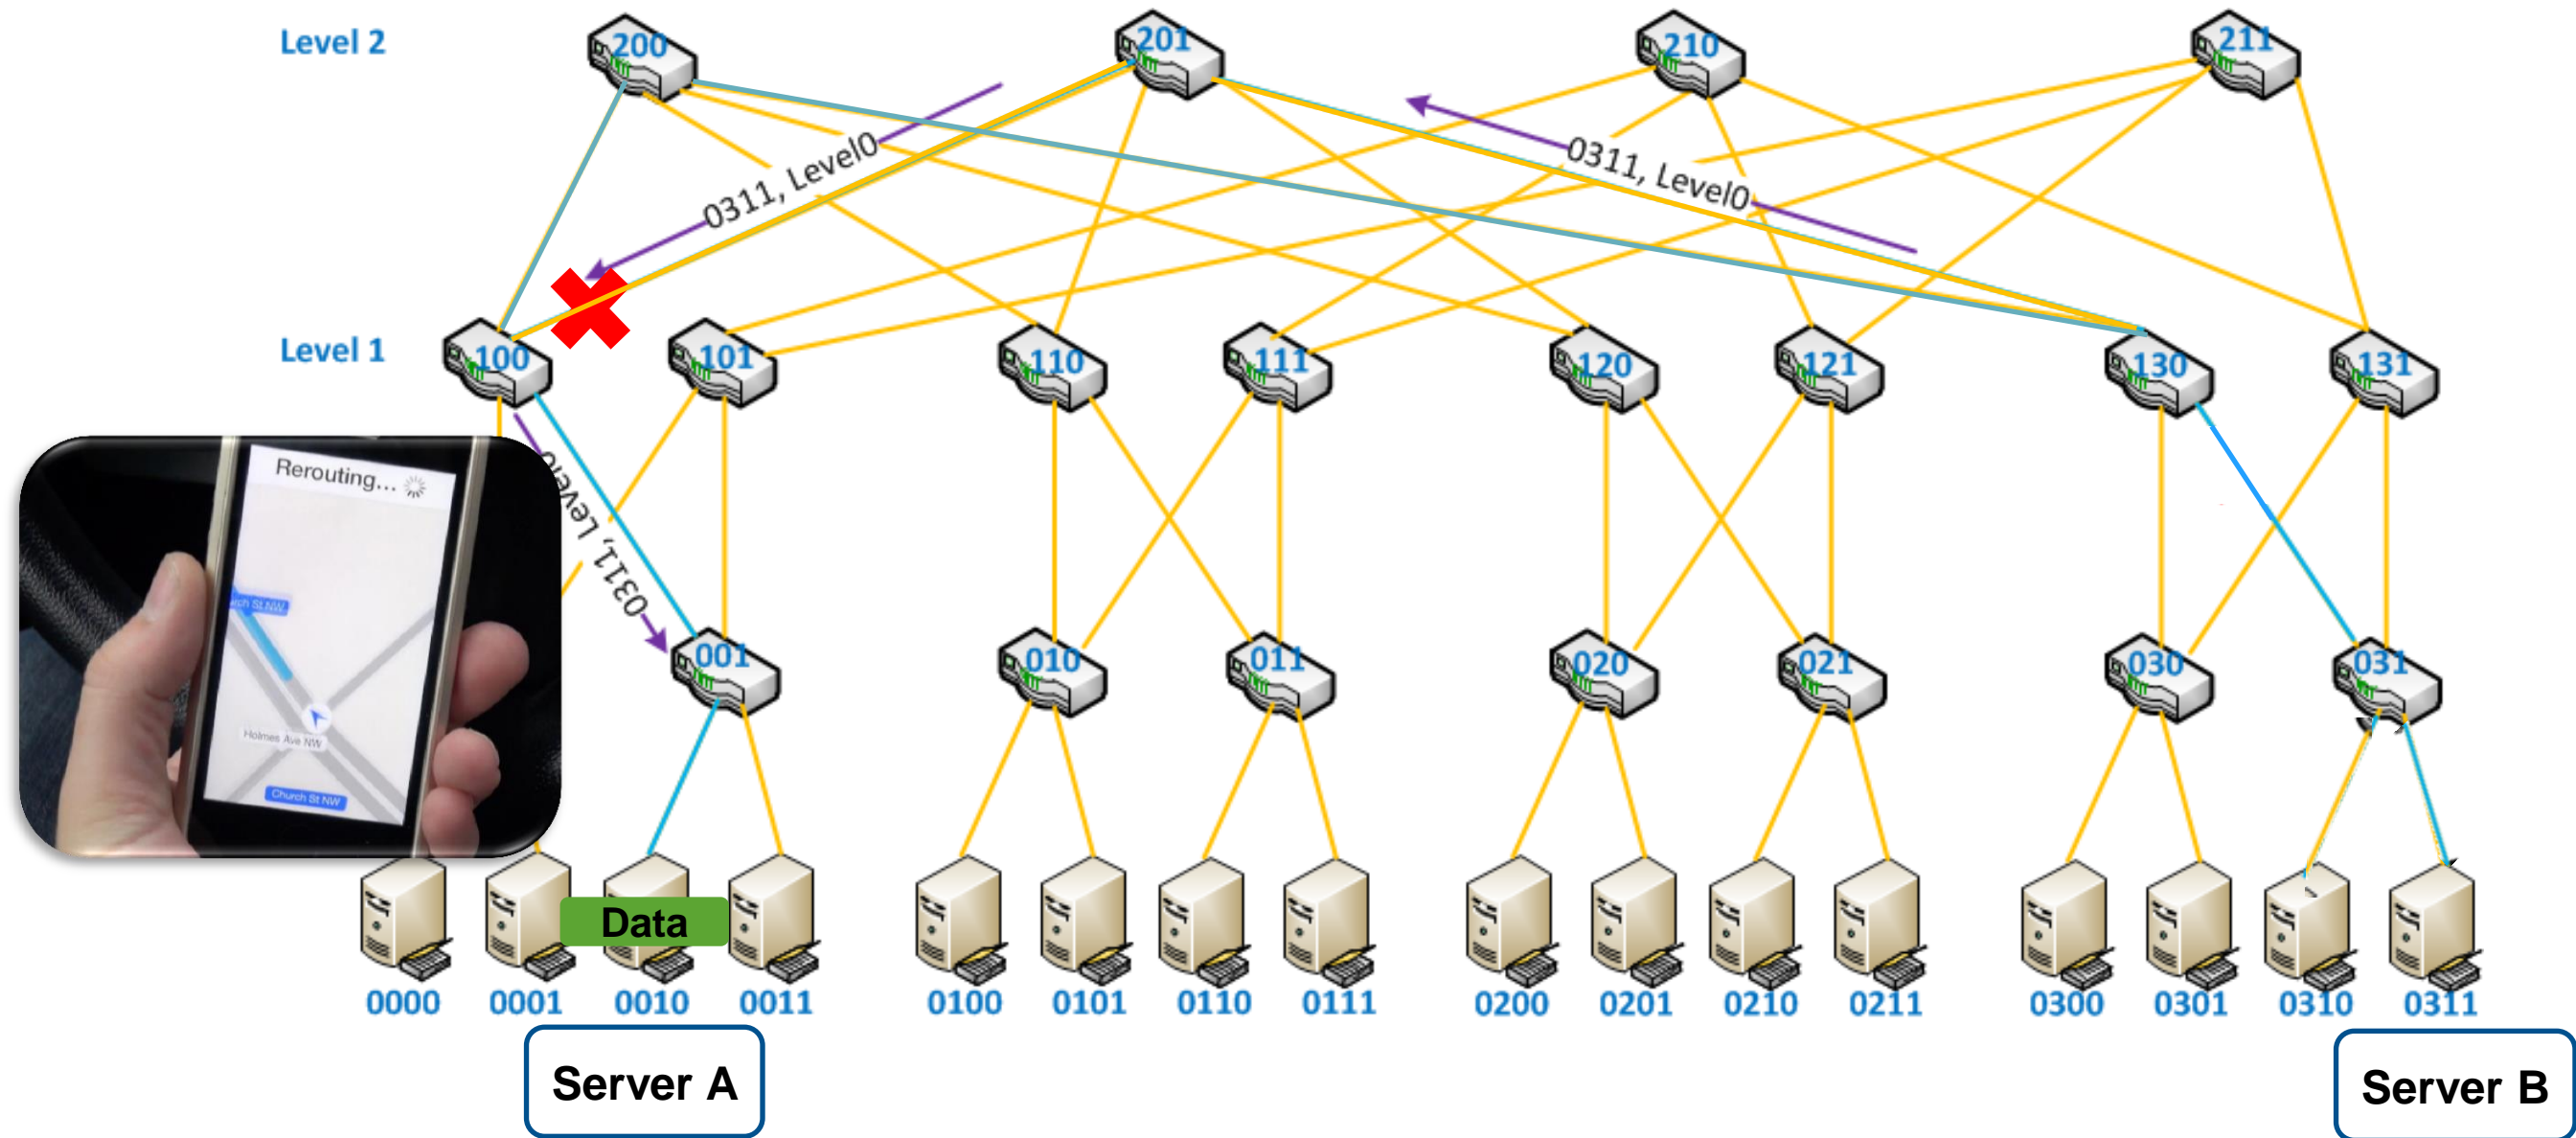
5000X Faster Network Recovery

Enable Unbreakable Data Centers

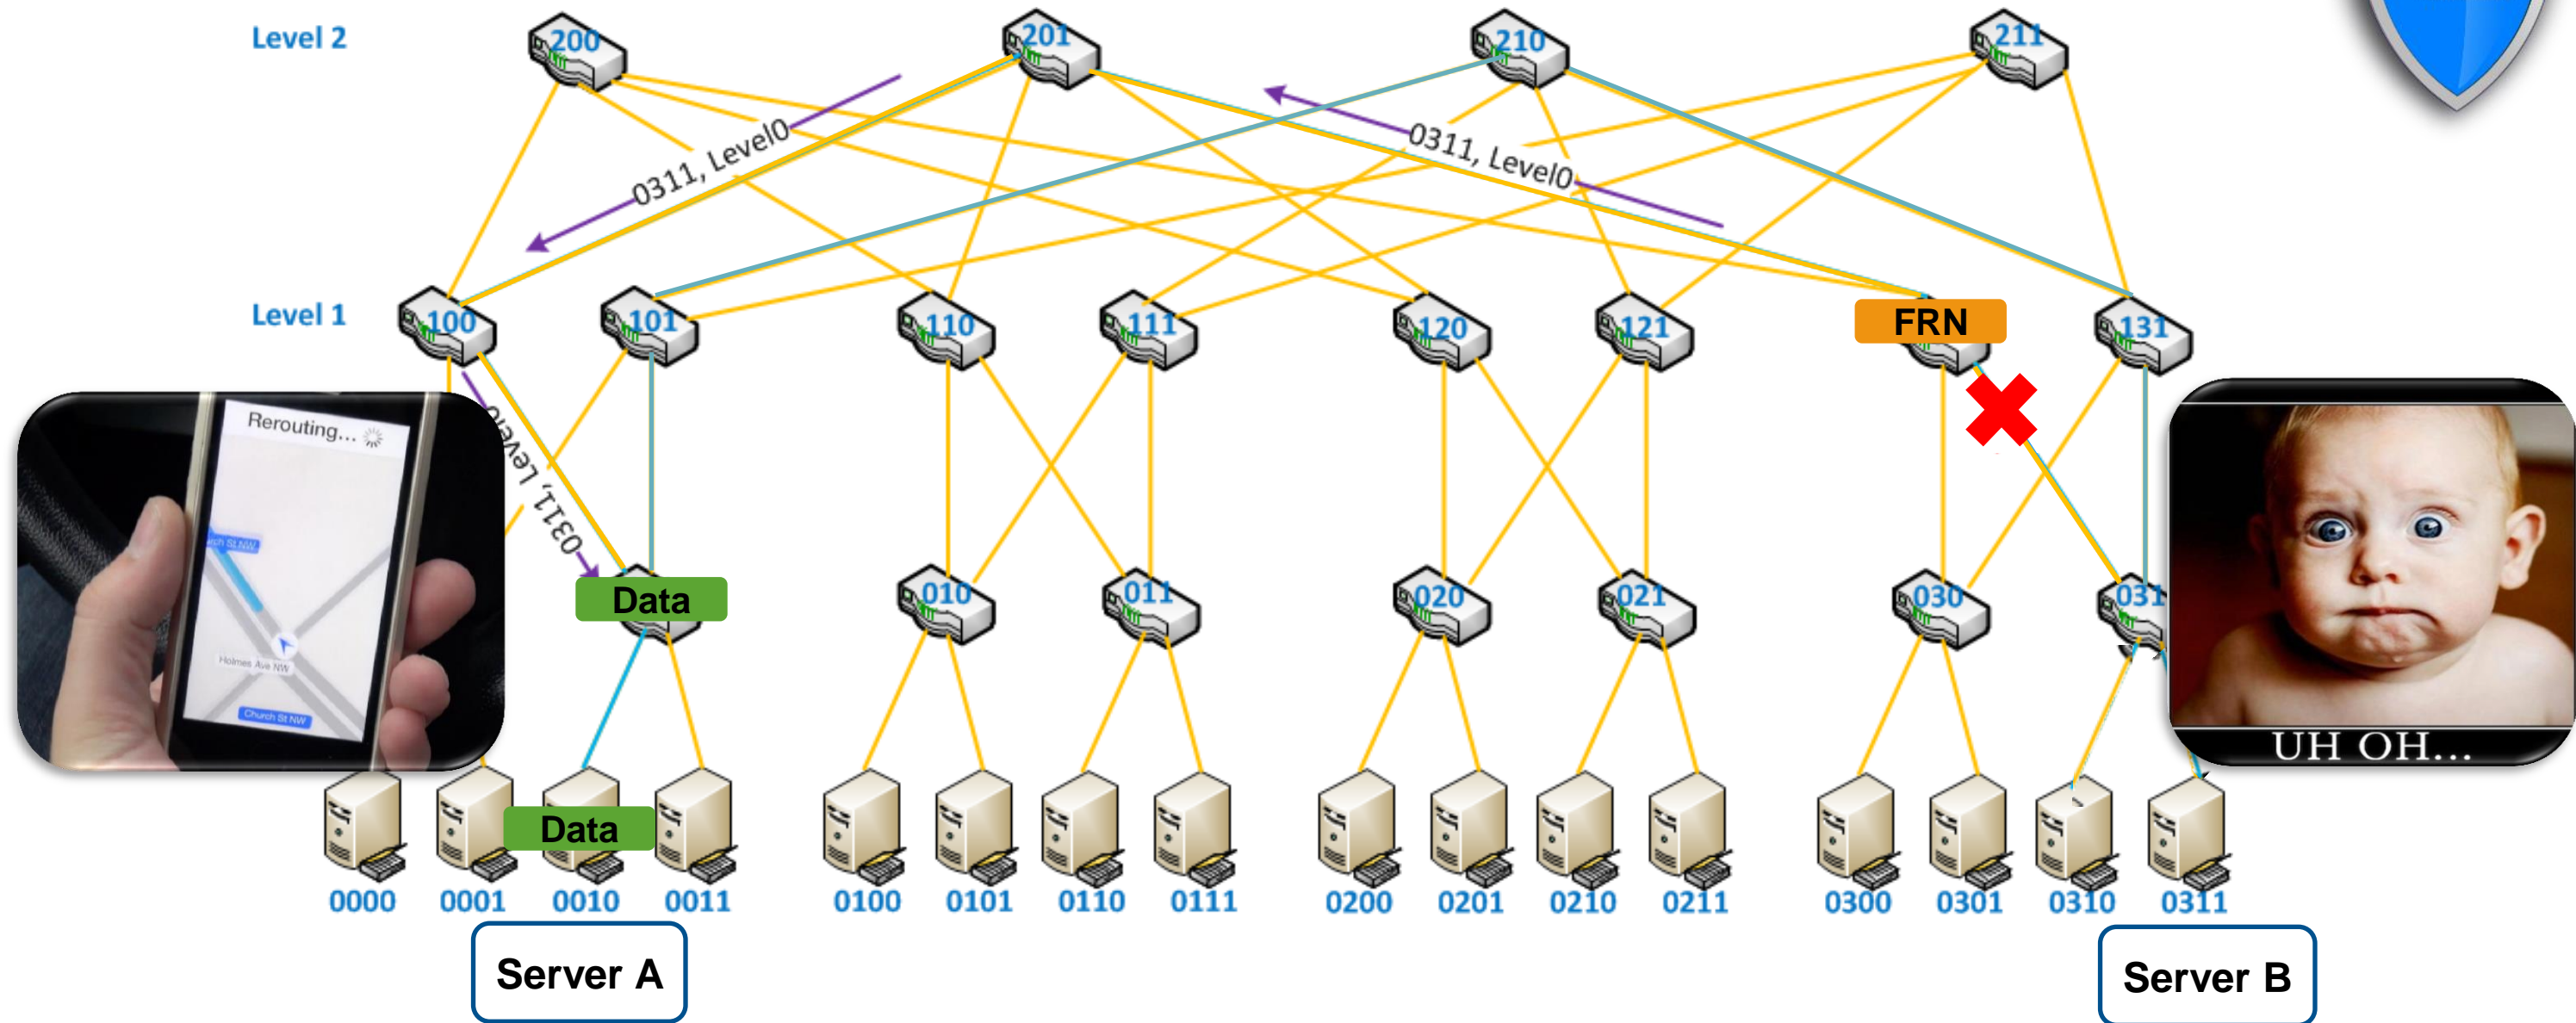
Consider a Flow From A to B



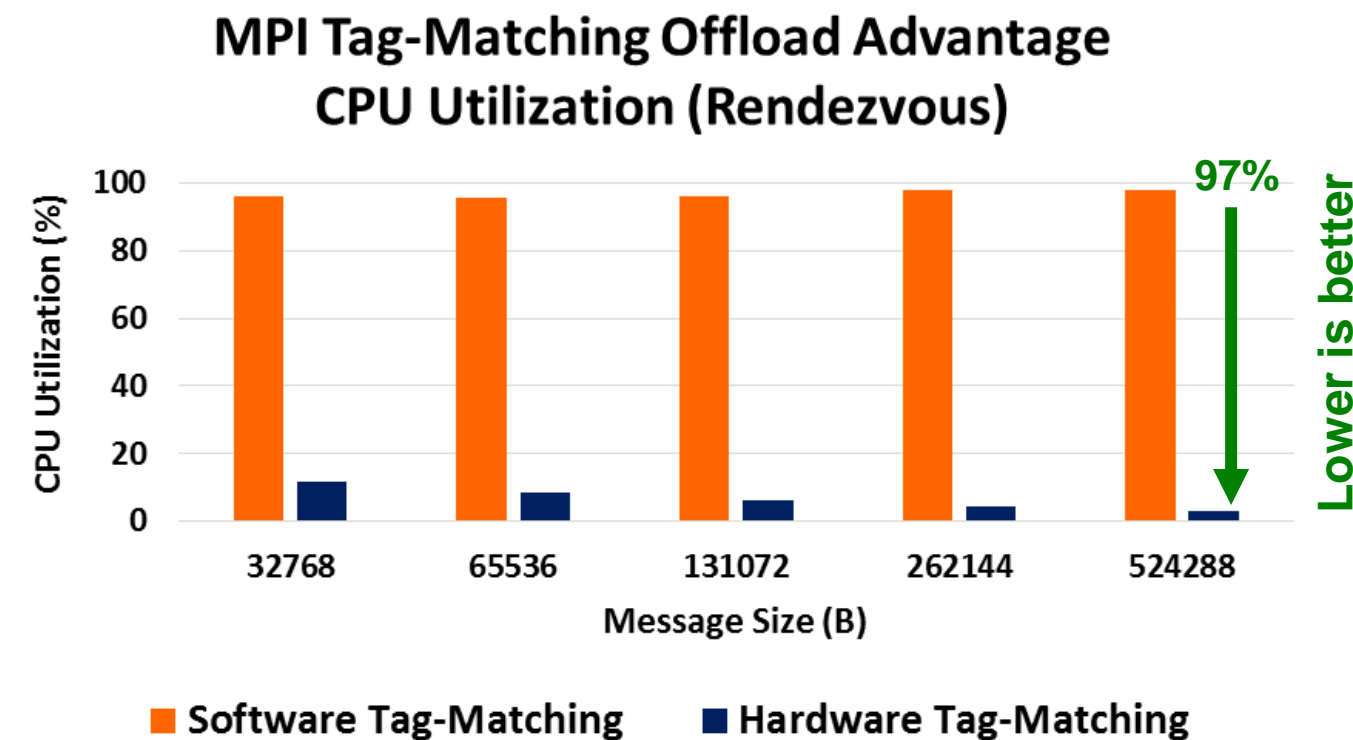
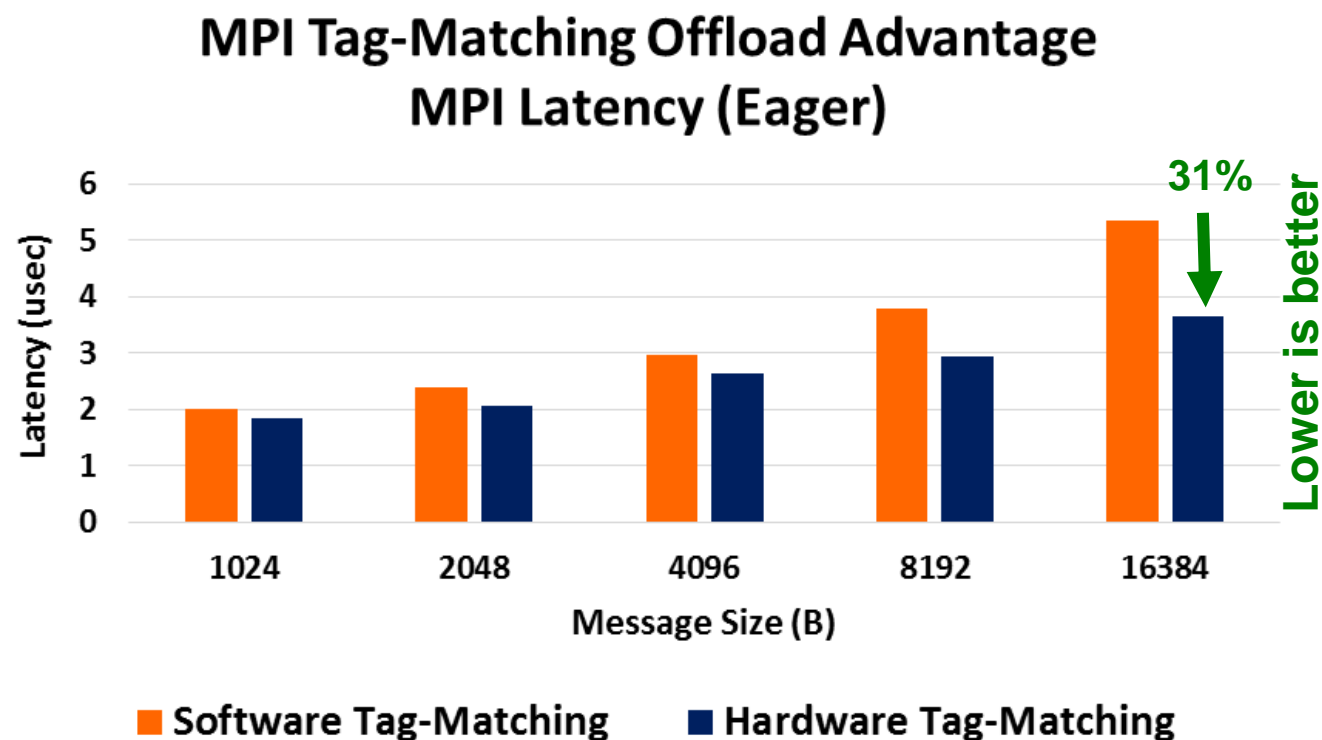
The Simple Case: Local Fix



The Remote Case: Using FRN's (Fault Recovery Notifications)




MPI Tag-Matching Offload Advantages




- 31% lower latency and 97% lower CPU utilization for MPI operations
- Performance comparisons based on ConnectX-5

Mellanox In-Network Computing Technology Deliver Highest Performance

Static Routing Performance




Adaptive Routing Performance



95% Network Utilization with Adaptive Routing

GPU-GPU Internode MPI Latency

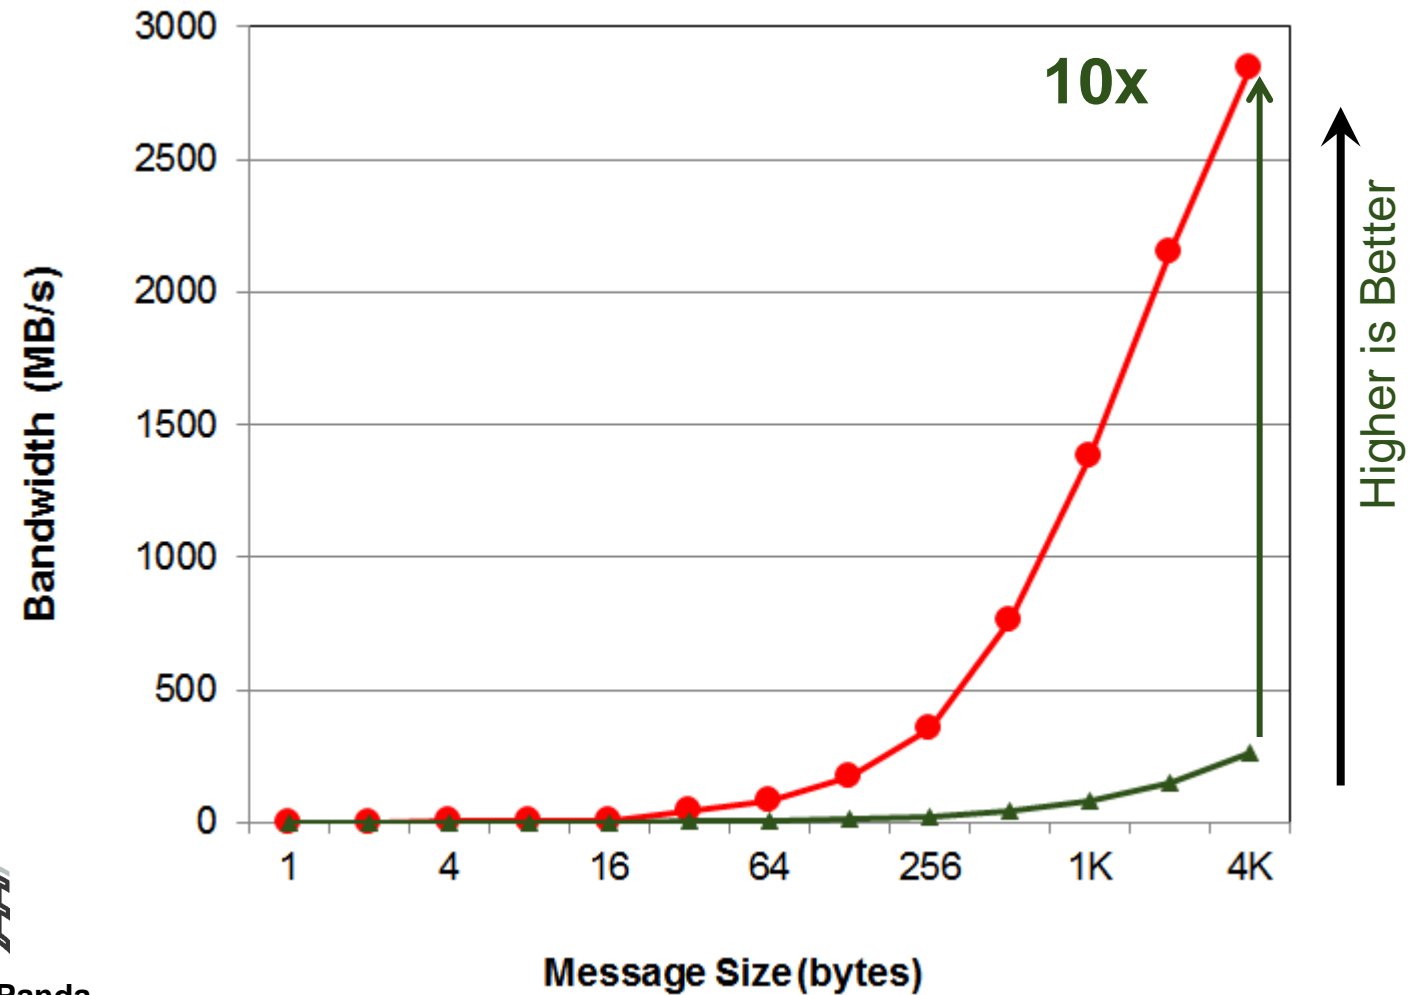


Lower is Better



Source: Prof. DK Panda

GPU-GPU Internode MPI Bandwidth




Higher is Better

88% Lower Latency

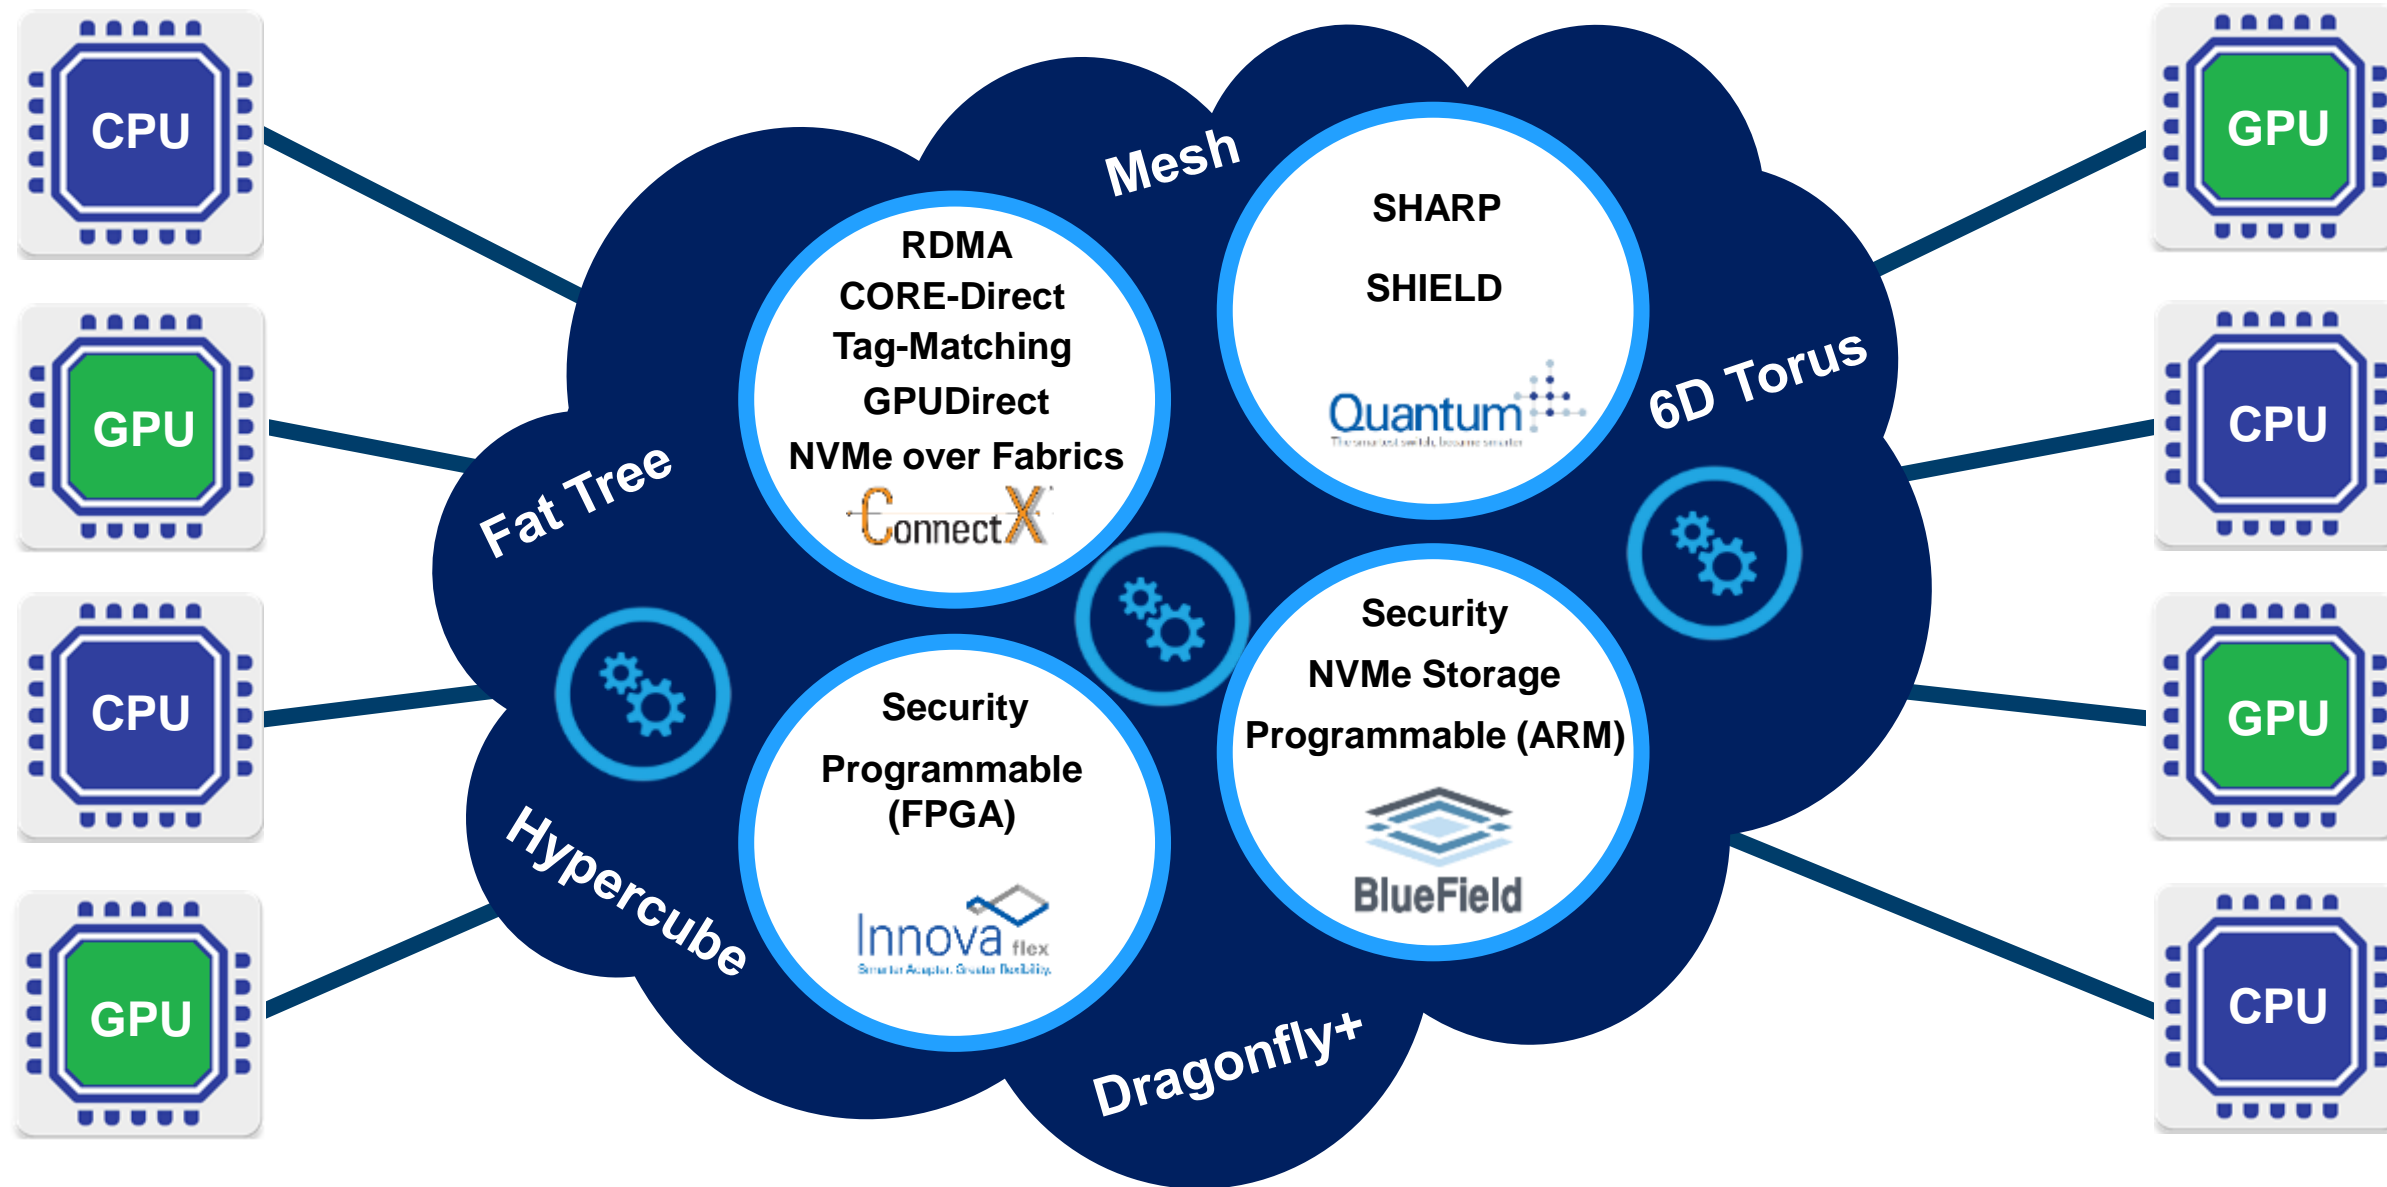
10X Increase in Throughput

- Machine Learning Software from Baidu
 - Usage: word prediction, translation, image processing
- RDMA (GPUDirect) speeds training
 - Lowers latency, increases throughput
 - More cores for training
 - Even better results with optimized RDMA











~2X Acceleration for Paddle Training with RDMA

The Generation of In-Network Computing – 10X Higher Performance

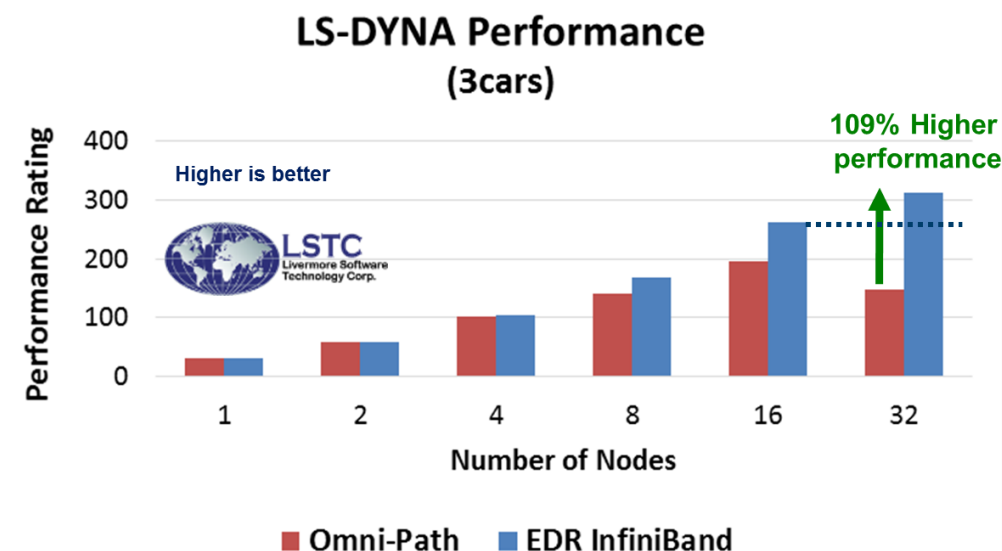

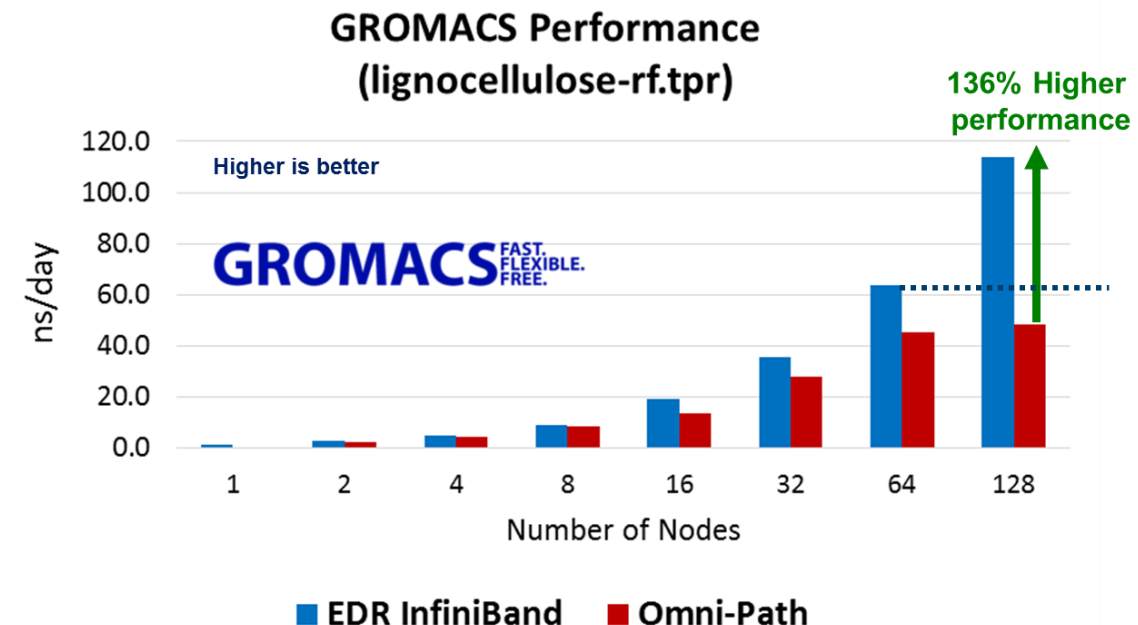



In-Network Computing Key for Highest Return on Investment

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

Applications Performance Comparison - Examples



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



Thank You