



## Interconnect Your Future

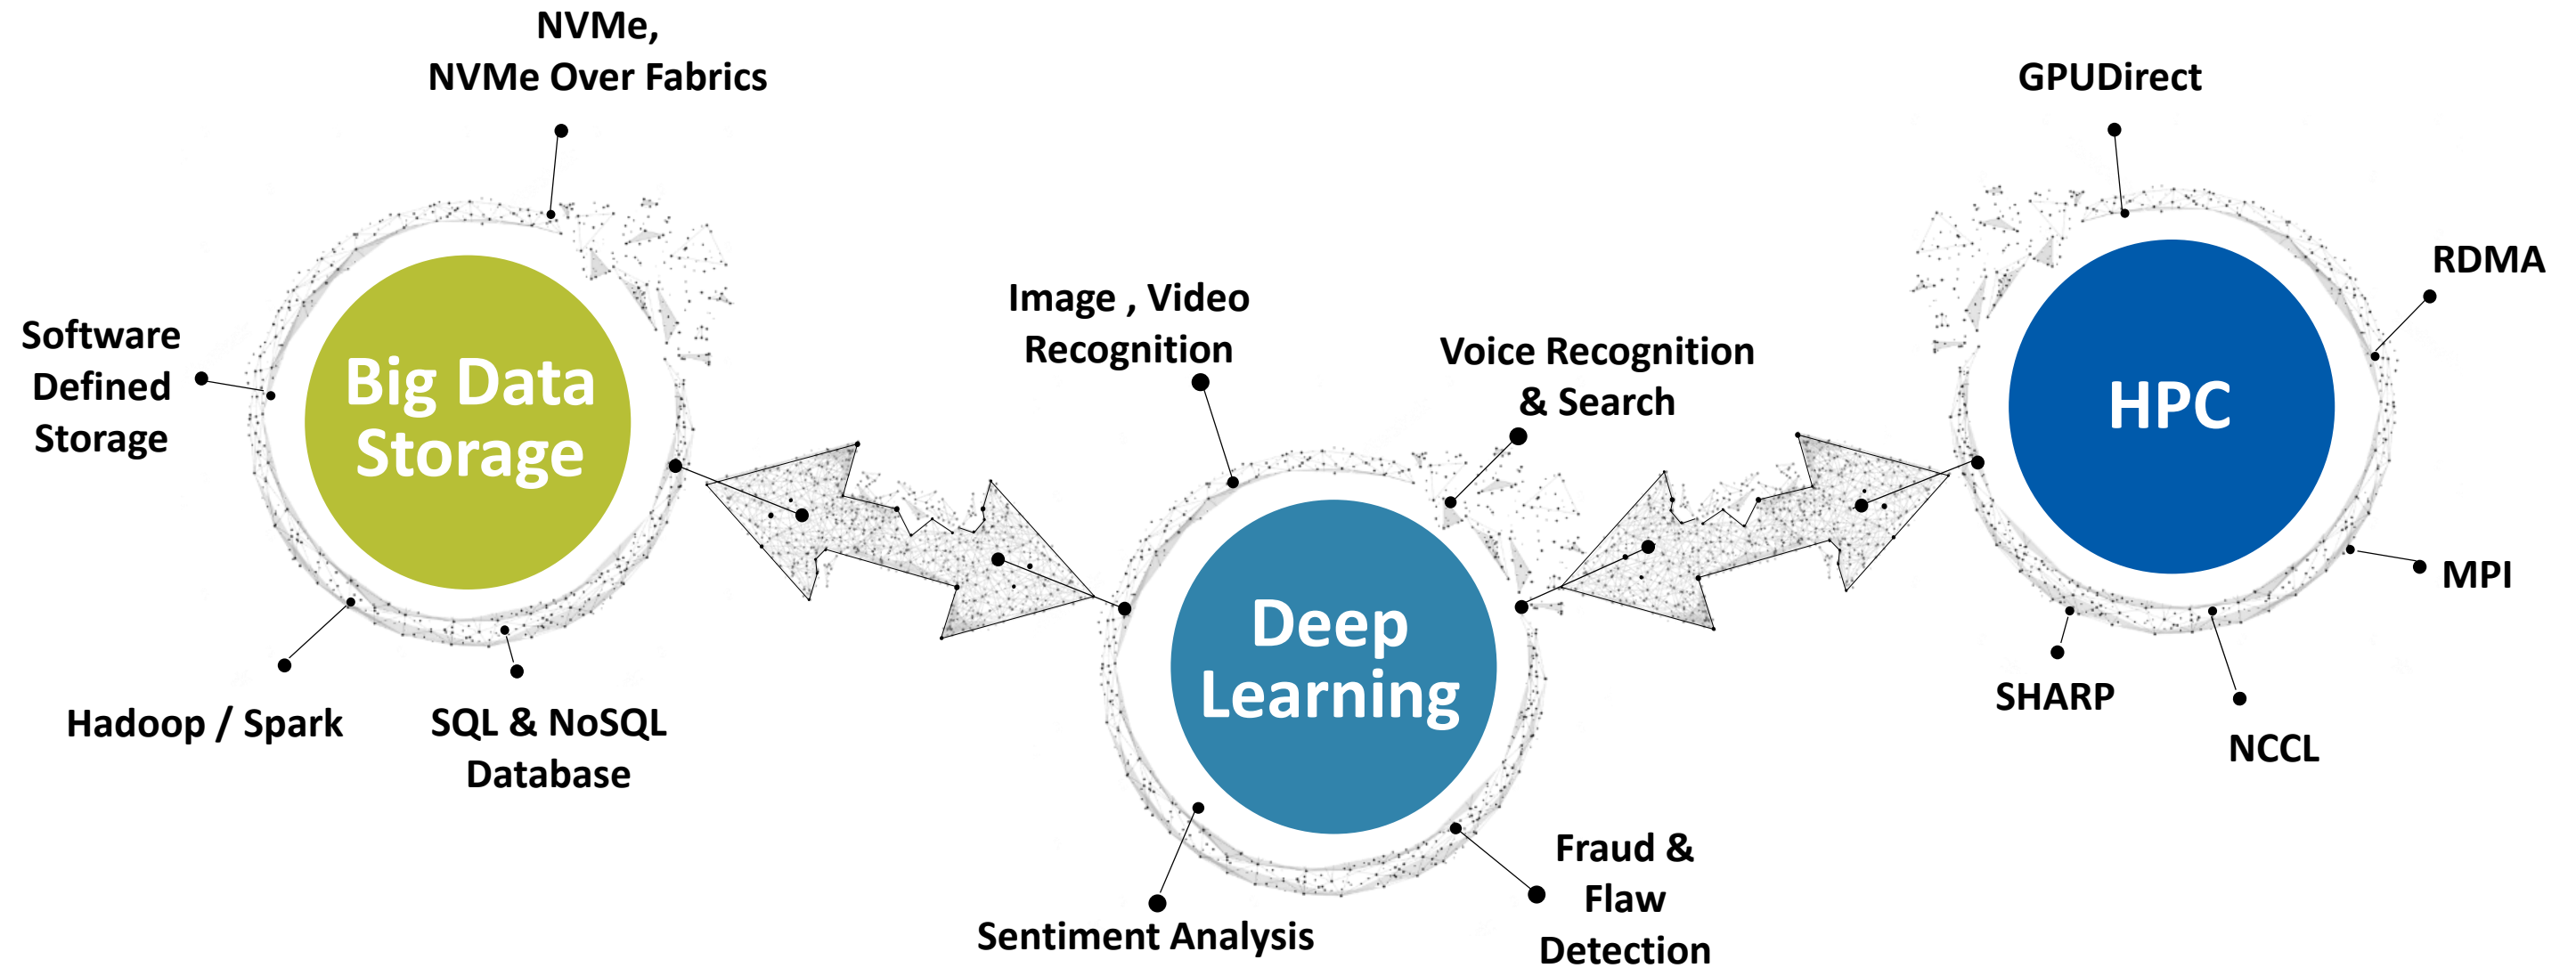
Paving the Road to Exascale

Supercomputing, November 2017


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



# Same Interconnect Technology Enables a Variety of Applications



# Artificial Intelligence Needs the Most Intelligent Interconnect



SmartNIC



System on a Chip

**Higher**

**Data Speeds**

**Faster**

**Data Processing**

**Better**

**Data Security**



Adapters



Switches




Cables &  
Transceivers




# The Need for Intelligent and Faster Interconnect


## CPU-Centric (Onload)



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



## Data-Centric (Offload)



**Analyze Data as it Moves!**

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

# 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

## Unleashing the Power of all Compute Architectures



**Highest Performance & Scalability for X86, Power, GPU, ARM & FPGA-based Compute and Storage Platforms**



# HPE Apollo 70 with Mellanox ConnectX Adapters



ConnectX<sup>®</sup> 4 Lx



## Best Ethernet Connectivity

On Board ConnectX<sup>®</sup>-4 Lx - 10G/25GE Ethernet  
Smart offloading - Virtual Overlay Network, OVS, RDMA

ConnectX<sup>®</sup> 5




## Best HPC and AI Connectivity

100Gb/s InfiniBand and Ethernet  
MPI Accelerations engines (Tag Matching, Collectives, Rendezvous protocol)  
GPUDirect RDMA, Adaptive Routing, Atomic Operations enhancements  
Burst Buffer Offload, NVMe over Fabric

Choice, Flexibility  
Highest Performance

HPE Apollo 70



Hewlett Packard  
Enterprise

# Highest-Performance 100Gb/s Interconnect Solutions

Adapters

ConnectX<sup>®</sup> 5

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



Switch

SwitchIB<sup>™</sup> 2

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



Switch

Spectrum<sup>™</sup>

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



Interconnect

LinkX<sup>™</sup>

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



VCSELs, Silicon Photonics and Copper

Software


HPC-X<sup>™</sup>

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





## HPCG Performance on ARM (40x40x40)



# 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



Interconnect

LinkX<sup>™</sup>

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



VCSELs, Silicon Photonics and Copper

Software

HPC-X<sup>™</sup>

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





# Thank You