Unigroup Presentation

Mellanox TECHNOLOGIES

Connect. Accelerate. Outperform.™

Mellanox Overview

Ticker: MLNX NASDAQ®

- Leading provider of high-throughput, low-latency server and storage interconnect
  - Ethernet and InfiniBand up to 100Gb/s
  - Reduces application wait-time for data
  - Dramatically increases ROI on data center infrastructure

- Company headquarters:
  - Yokneam, Israel; Sunnyvale, California
  - ~1,950 employees worldwide

- Solid financial position
  - FY'14 revenue of $463.6M
  - Cash + investments @ 6/30/15 = $467.2M

- Mellanox NYC Team (@ One Liberty Plaza)
  - Asaf Wachtel – Sr. Director, Business Development
  - Lior Paster – Director, Business Development
  - Eitan Rabin – Director, Advanced Software Solutions
  - Dor Juravski – Senior Systems Engineer

Agenda

- Technology / Market Update
  - Ethernet State of the Market Update
  - InfiniBand Introduction & State of the Market Update
  - Ethernet vs InfiniBand Comparisson
  - Offload Technologies
    - RDMA / ROCE
    - OS bypass
    - Offloads in Virtualized Environments
  - Use case examples
- Mellanox Company / Portfolio Update
- Advanced Topics / Q&A

Interconnect Speed Evolution

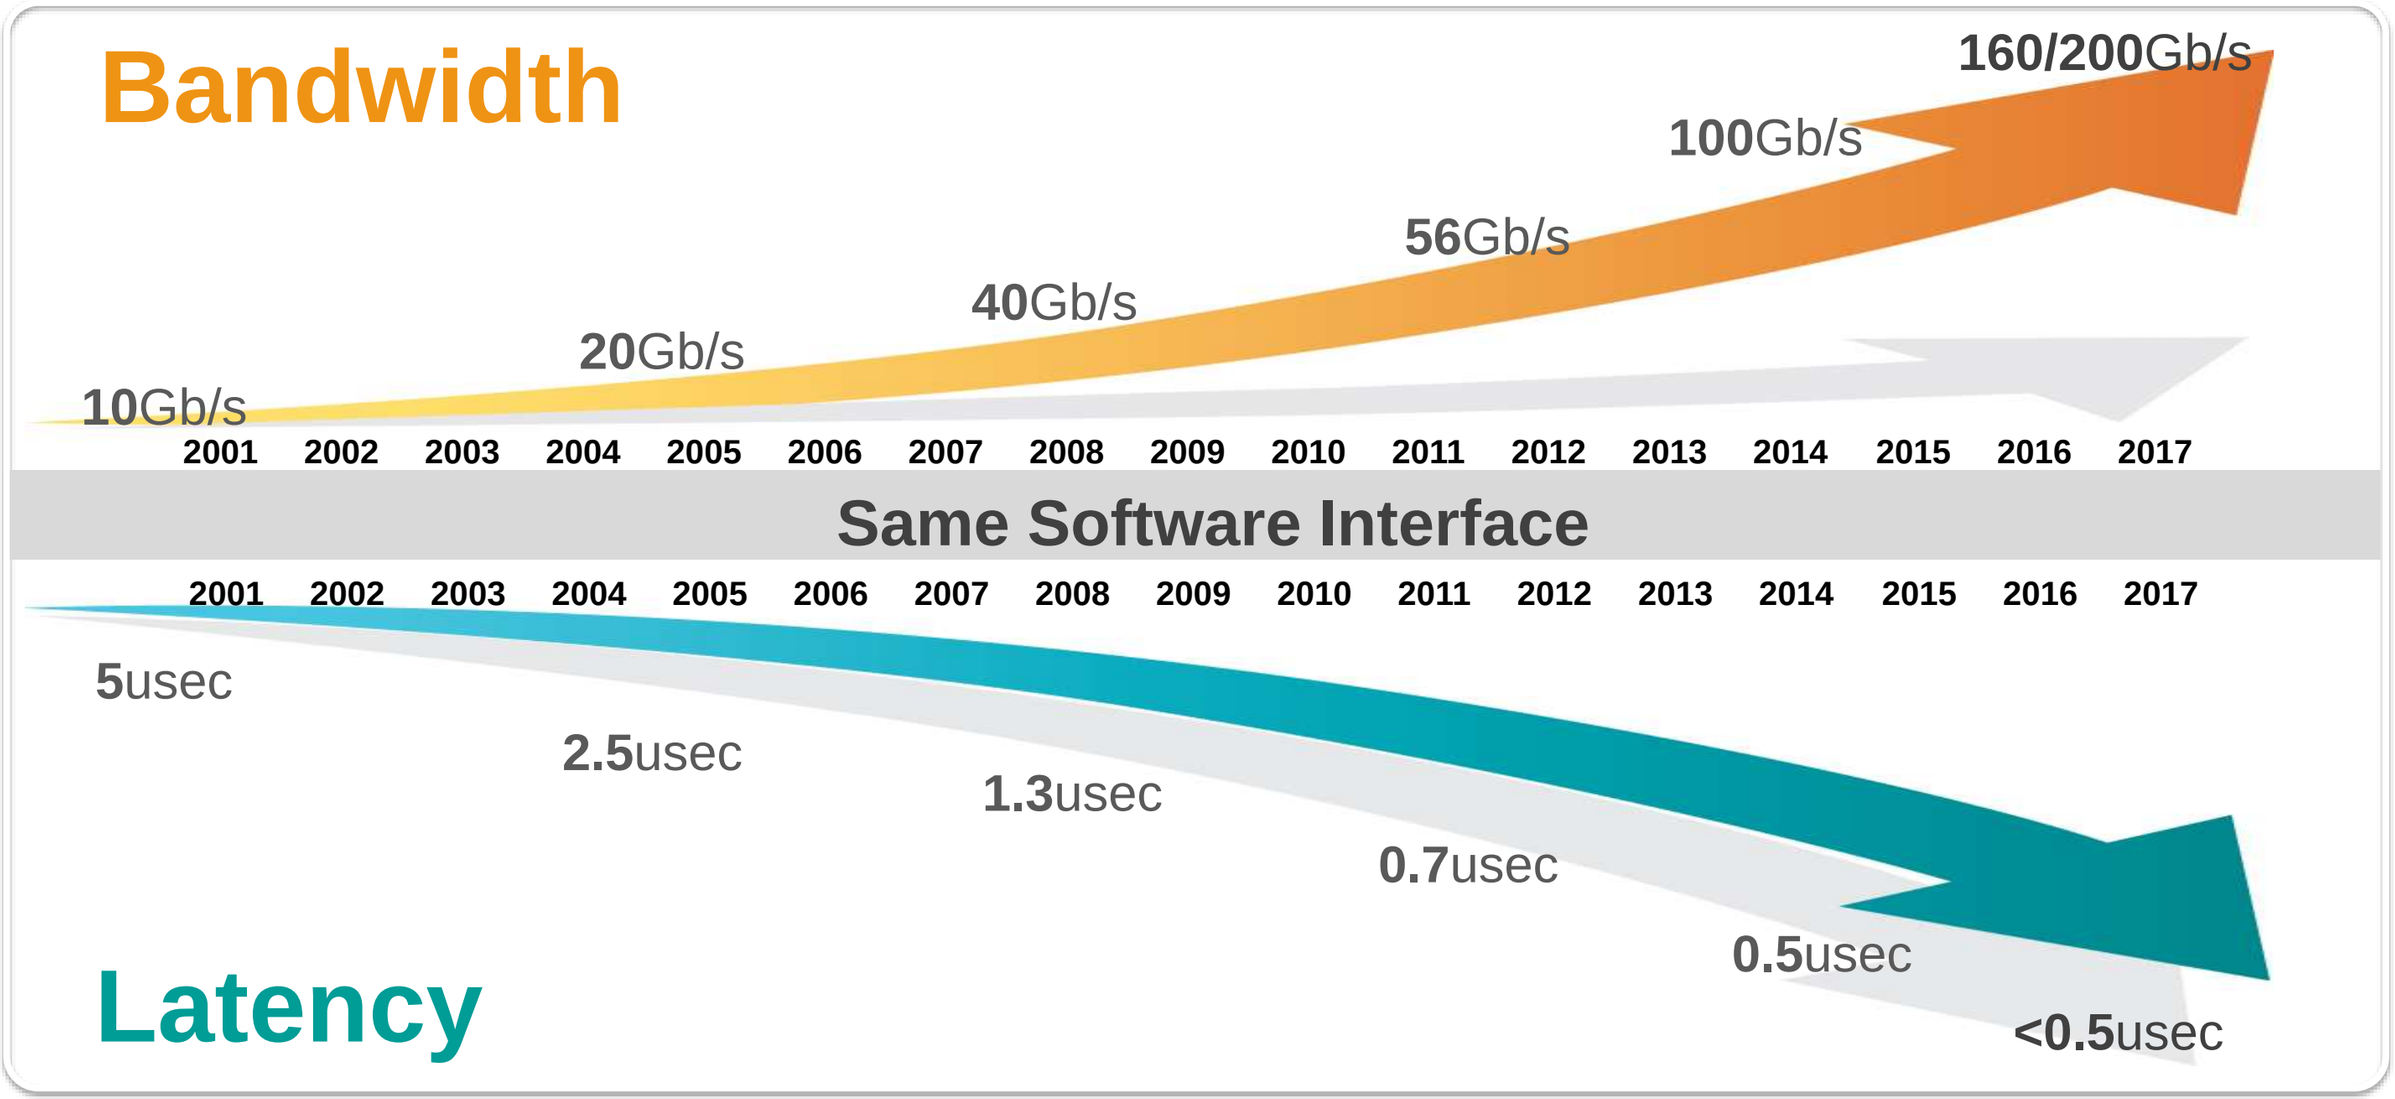

Data Center Trends driving High Performance Networking

- Multi-core, Multi-processor architectures
- “Cloudification” – Seamless access to resources anywhere, anytime
- “Big Data” – Exponential growth in Volume, Variety and Velocity
- Storage Media – Flash / NVMe (HDD no longer bottleneck)

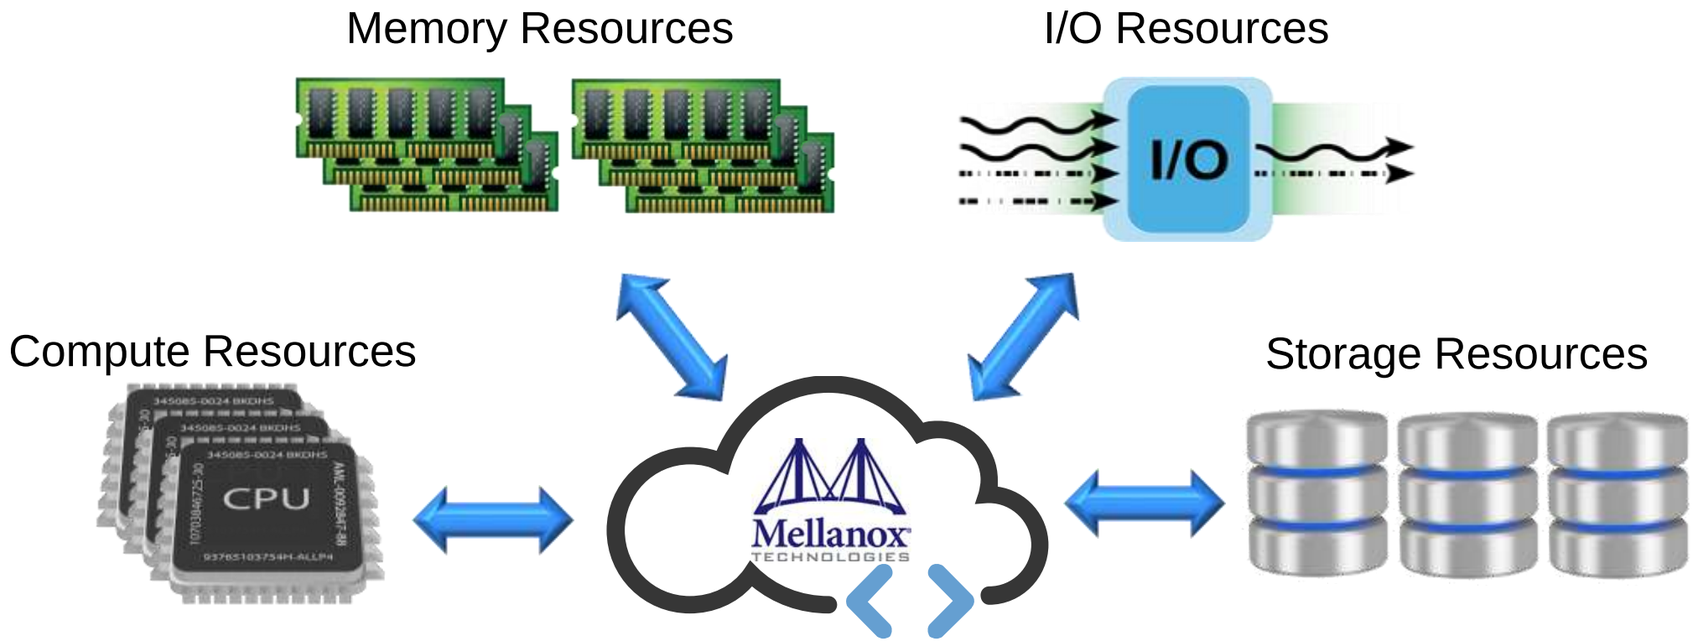

Ethernet State of the Market

Ethernet: Typical Components

Automation & Monitoring

Management software

Switches

SN2700

SN2100

SN2410

SX1710 / SX1036

SX1012

SX1024 / SX1400

SX1016

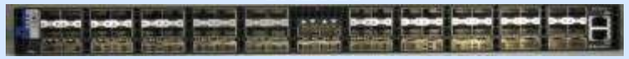

NICs

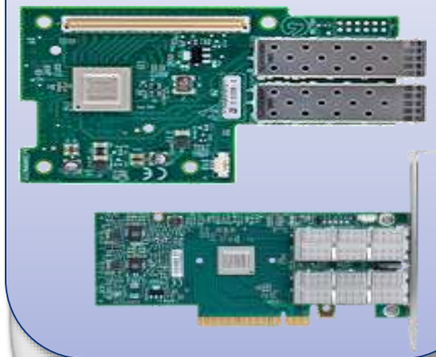

Cables

Cables

NICs

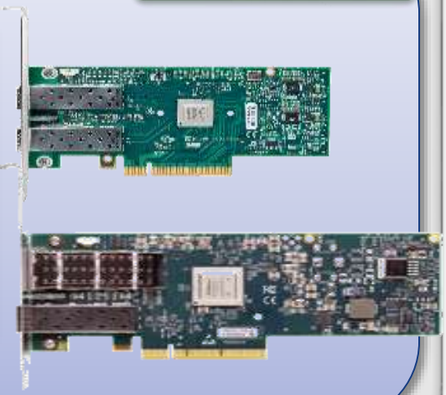

25/50/100GbE – The Future is Here!

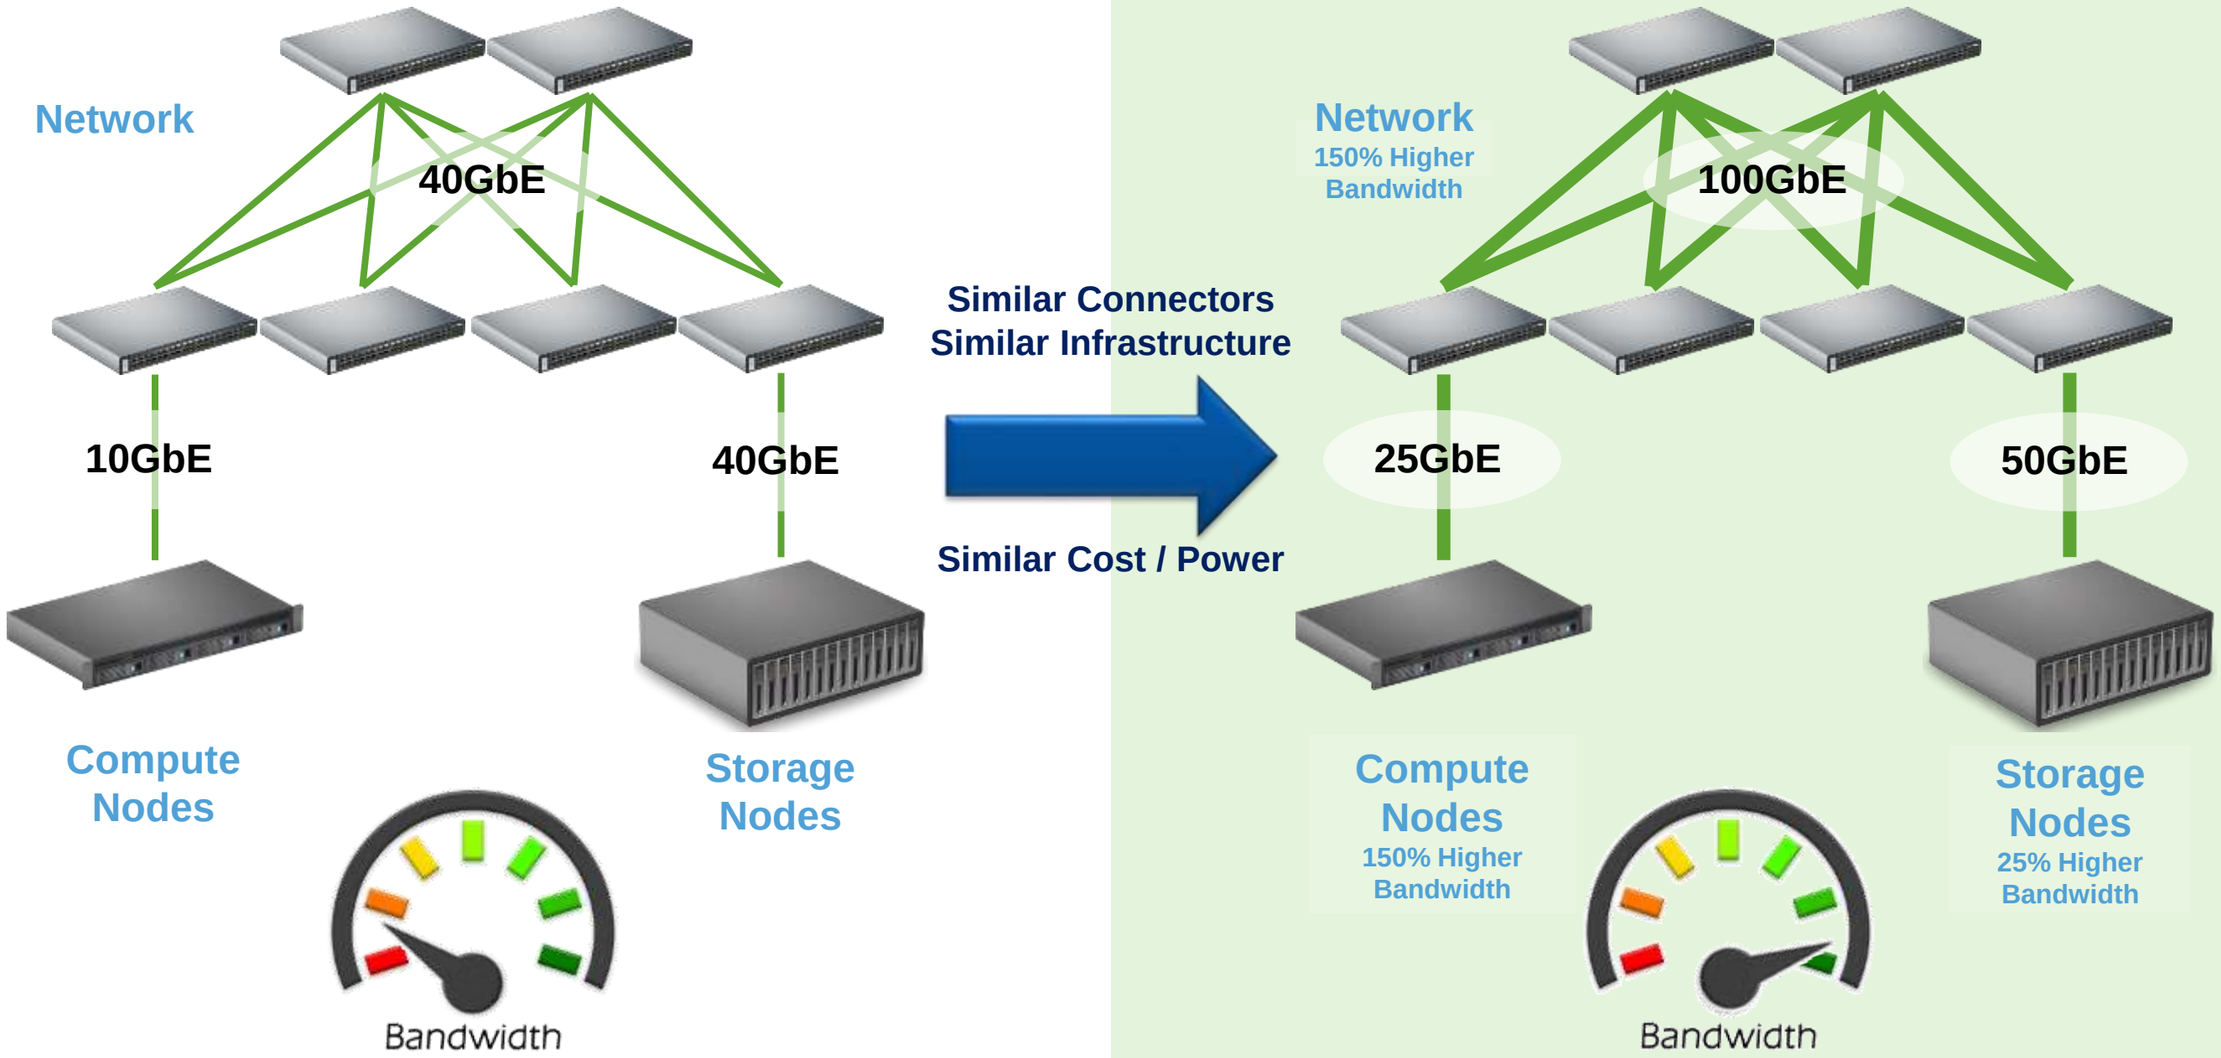

Ethernet Aggregation Switch Network

- Scalable cloud, web 2.0, and Enterprise datacenters require large aggregation Ethernet switch network
  - Connecting racks for 10Gb/E based servers and storage
- Legacy solutions utilize large modular switches
  - Expensive (both CAPEX and OPEX)
  - Rapid obsolescence
  - Market share of modular switches is on the decline
- Users prefer to use of Top-of-Rack switches to form Virtual Modular Switch solutions
  - A more cost effective and flexible solution
  - Allows over-subscription to be handled on a rack basis

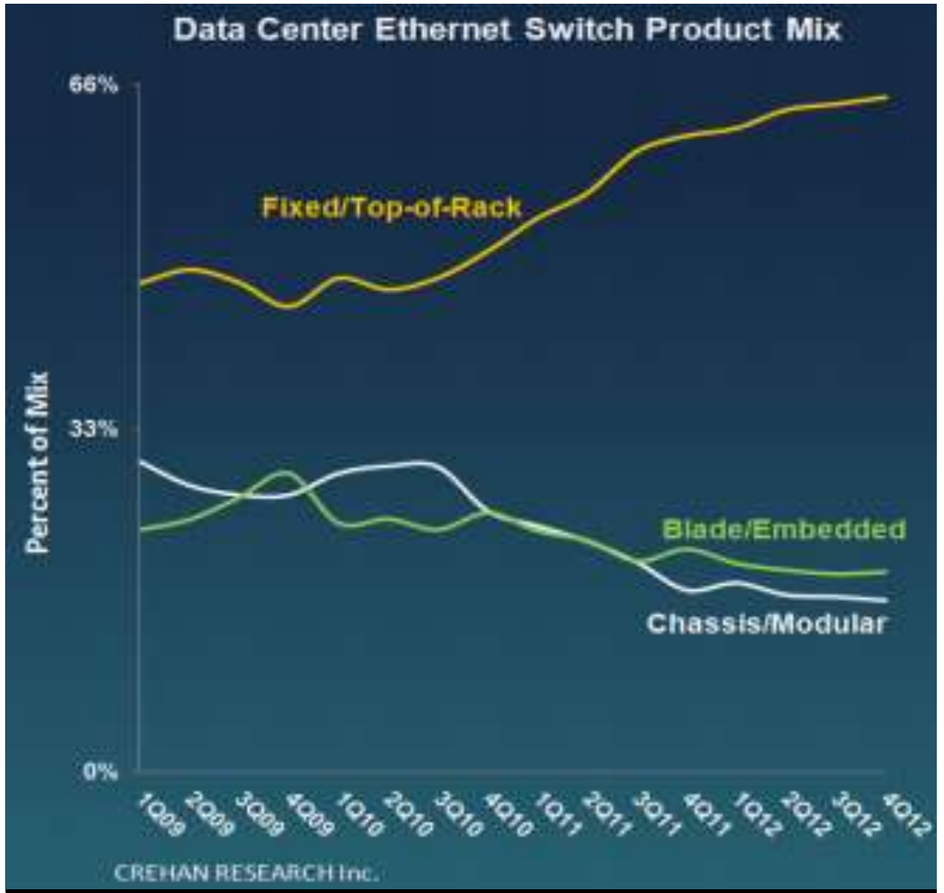

"Data center Ethernet Fixed/Top-of-Rack switches continue to reach record levels, as this form-factor not only sees increasing server access deployment, but also data center aggregation deployments."

Open Ethernet: Separation Of Hardware And Software

- Open Ethernet
  - The alternative to traditional closed Ethernet switches
- Delivers the freedom of choice
  - Any software can run on any hardware
  - Independent change of hardware or software
- Switch as a server experience
  - Home grown solutions
  - Open source packages
  - Merchant SW packages

WE ❤ CHOICE

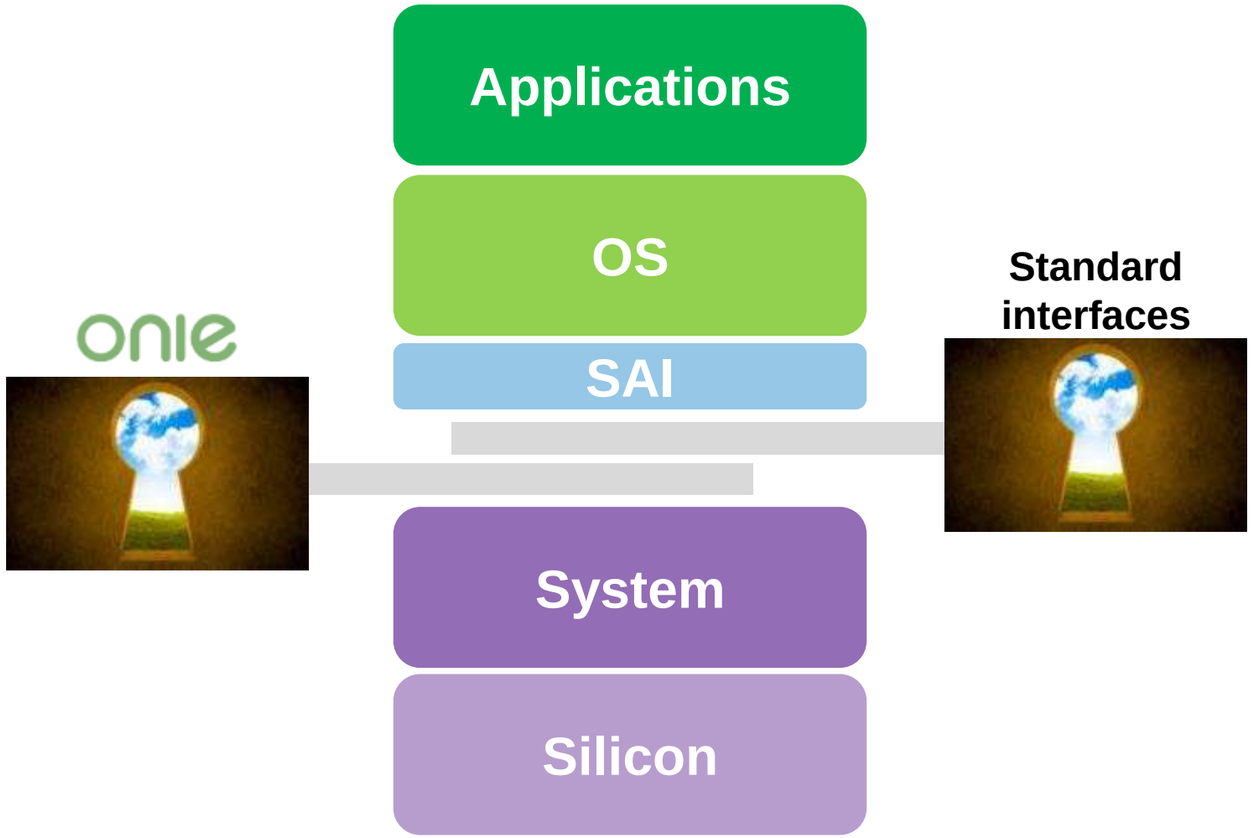

InfiniBand Intro &
State of the Market

History Lesson:

From Mainframes to Grid Computing

The Old Way	HPC Clusters
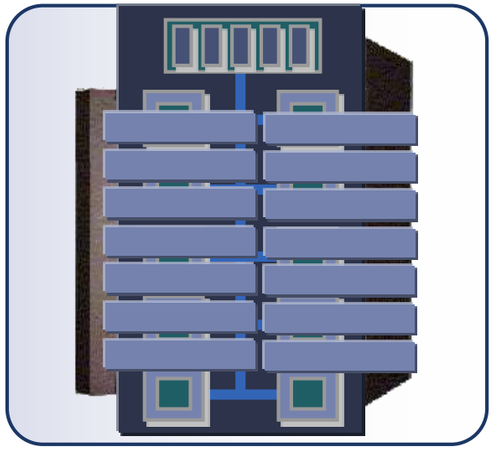	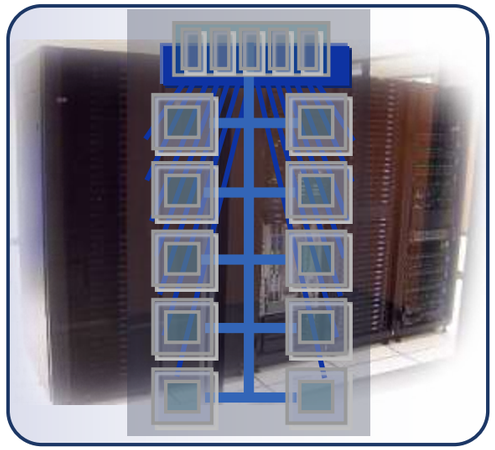
Expensive ✘	Commodity Based ✓
Low Utilization ✘	Optimized Utilization ✓
Limited Scalability ✘	Linear Scalability ✓
Simple to Manage ✓	Simple to Manage ✓

History Lesson:

InfiniBand - First Standard High Speed Interconnect

- Highest Speed Standard Interconnect
  - up to 100Gb/s
- Lowest Latency
  - <200ns latency per switch hop
  - <1us application to application latency
- Hardware Based Protocol Offload and RDMA
  - Enabling low CPU%, Zero copy and OS Bypass
- Linear Scalability
  - Multipathing and Clos topologies
- Lossless fabric
- InfiniBand drivers already in the Linux kernel
  - Offered by Red Hat and SuSE/Novell
- Interoperability driven eco-system
  - OpenFabrics
  - Interoperability events twice a year
- Supports standard TCP/IP transport

InfiniBand Trade Association (IBTA)

- Founded in 1999
- Actively markets and promotes InfiniBand from an industry perspective through public relations engagements, developer conferences and workshops
- Resources:
  - InfiniBand software is developed under OpenFabrics Open Source Alliance
    - http://www.openfabrics.org/index.html
  - InfiniBand standard is developed by the InfiniBand Trade Association
    - http://www.infinibandta.org/home
- Steering Committee Members:

The InfiniBand Architecture

- Host Channel Adapter (HCA)
  - Device that terminates an IB link and executes transport-level functions and support the verbs interface
- Switch
  - A device that routes packets from one link to another of the same IB Subnet
- Subnet Manager
  - Central management of topology and health
- Router - Optional
  - A device that transports packets between different IBA subnets
- Gateway - Optional
  - Bridges InfiniBand to Ethernet or Fibre Channel

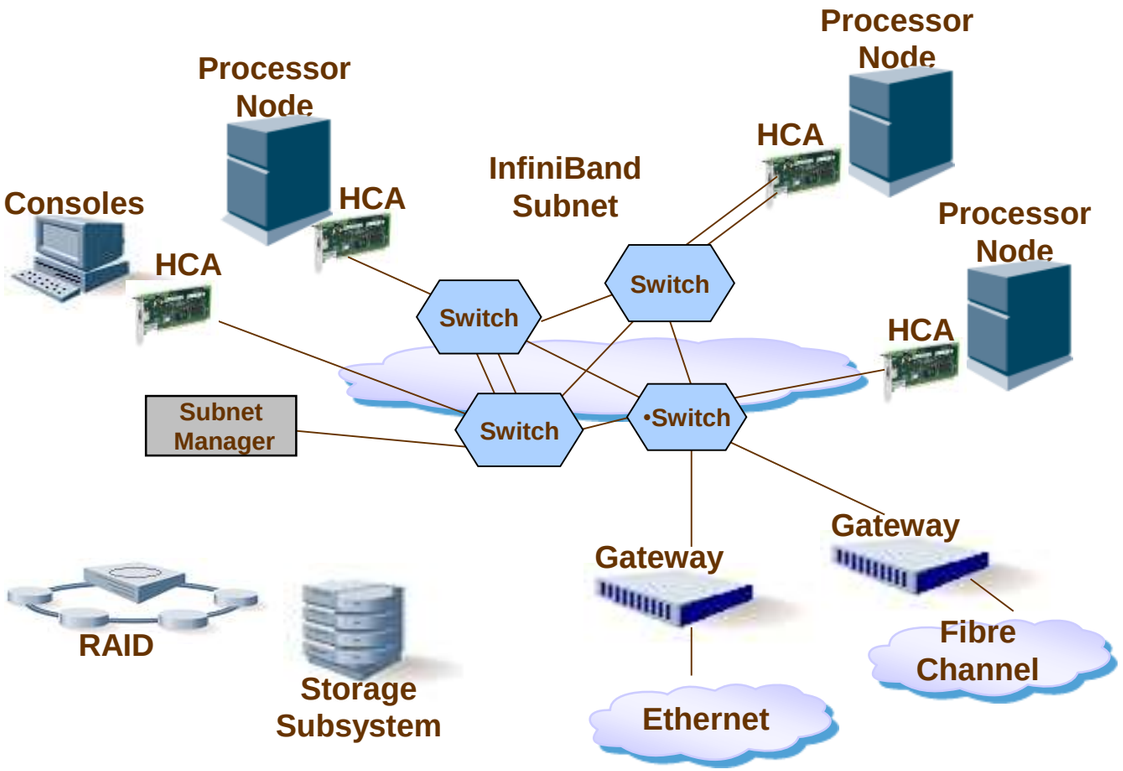

InfiniBand Roadmap

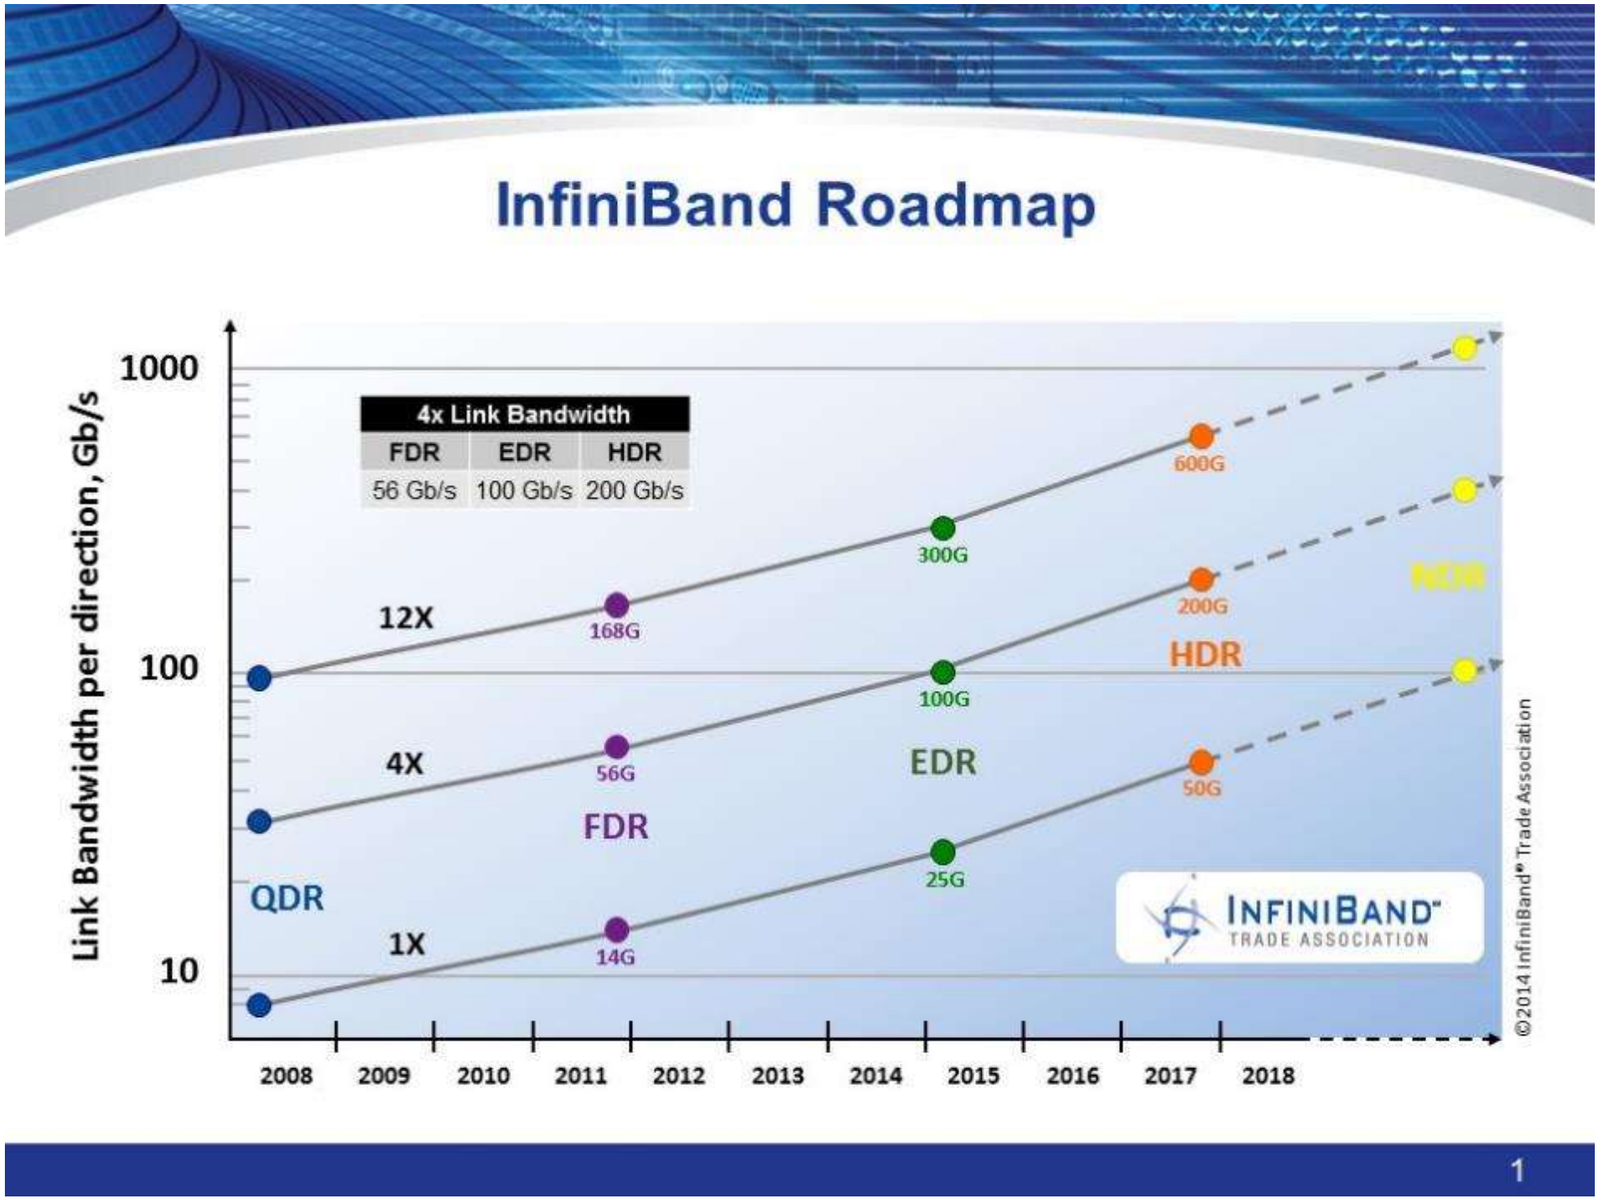

SDR - Single Data Rate
DDR - Double Data Rate
QDR - Quad Data Rate
FDR - Fourteen Data Rate
EDR - Enhanced Data Rate
HDR - High Data Rate
NDR - Next Data Rate

Long Haul InfiniBand - MetroX

Long Haul RDMA for Campus/Metro

Fully Optical Layer Management

40Gb/s link speed using DWDM

Optical Switch High Availability

Ethernet or InfiniBand – Where are they today?

Ethernet has advanced significantly over the past Decade

- Layer 3 Switching vs. Routing
- Cut-thru Switching vs. Store and Forward
- Higher Bandwidth / Lower Latency
- Lossless Ethernet (DCB)
- RDMA over Converged Ethernet (RoCE)

However in Certain Aspects InfiniBand Continues to Lead

- Centralized Management
- Linear Forwarding Table
- Ability to scale to thousands of nodes at L2
- L2 Multipathing
- Simplified HW based Multicast
- Still Highest Bandwidth / Lowest Latency
- Faster moving ecosystem

Networking Offloads

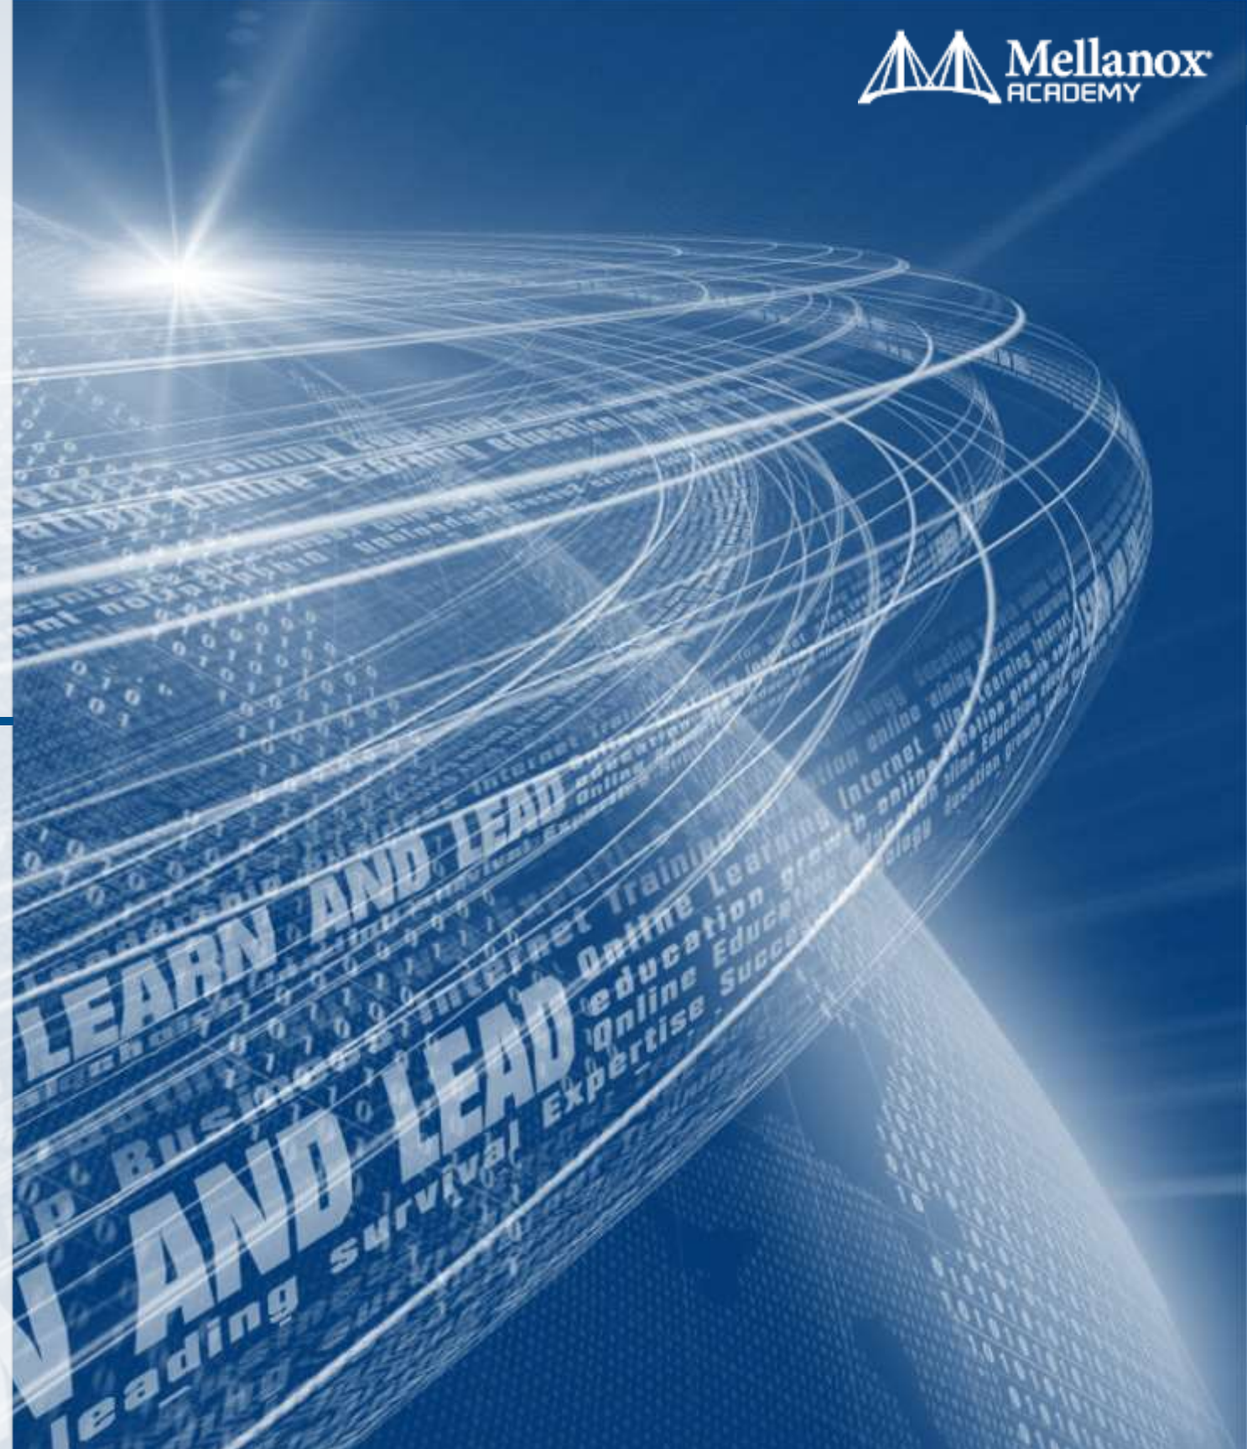

Offloading Your CPU

- Technology that enables data transfers without involving the CPU
- The CPU is free to perform other tasks
- Kernel is bypassed:
  - Latency is reduced
  - Throughput is increased
- Reliable Transports
- RDMA

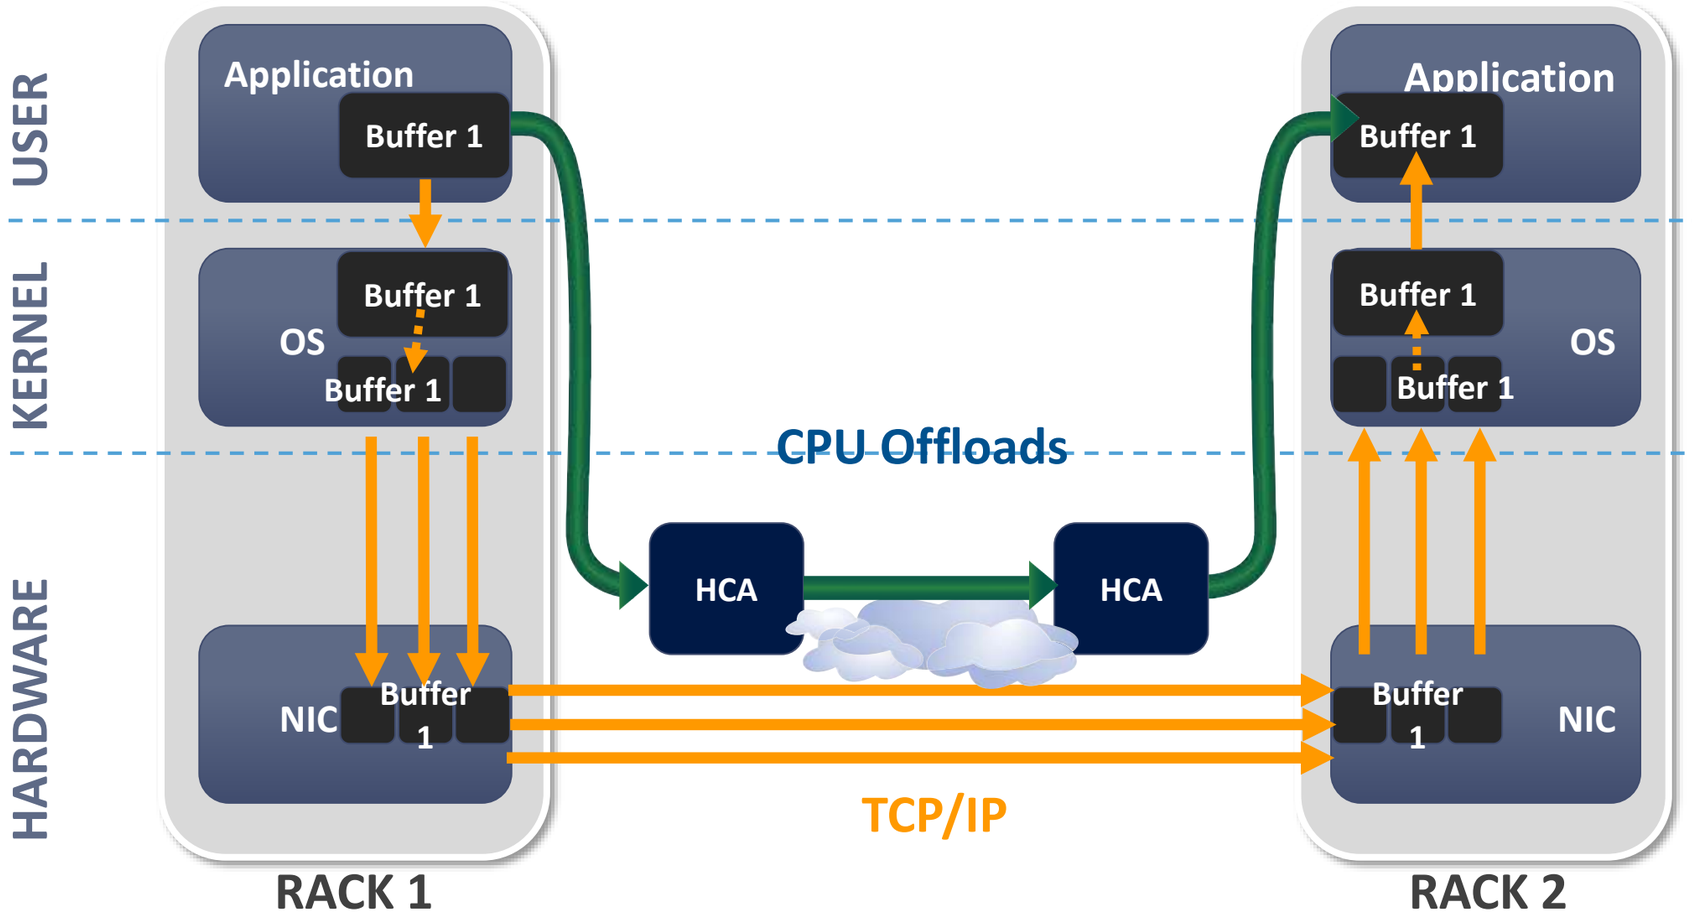

Offloads data transfer methods

There are two predominant methods for offloading data transfer:

- Channel semantics:

The receive side application allocates a buffer in its virtual memory space registering it in the receive queue.

The send side application sends the data through its send queue without having visibility of the receive side buffers.

  - SEND / RECEIVE

- Memory semantics:

The receive side application registers a buffer in its virtual memory space giving access to that buffer to a remote peer which can than target it with:

  - RDMA READ
  - RDMA WRITE

Hardware Offload Transport Responsibilities

- Segmentation/Reassembly:
  - The network adapter on the **send** side segments a message into multiple packets
  - The network adapter on the **receive** side reassemble packets into a message
- Reliability
  - Connection oriented vs. Datagram

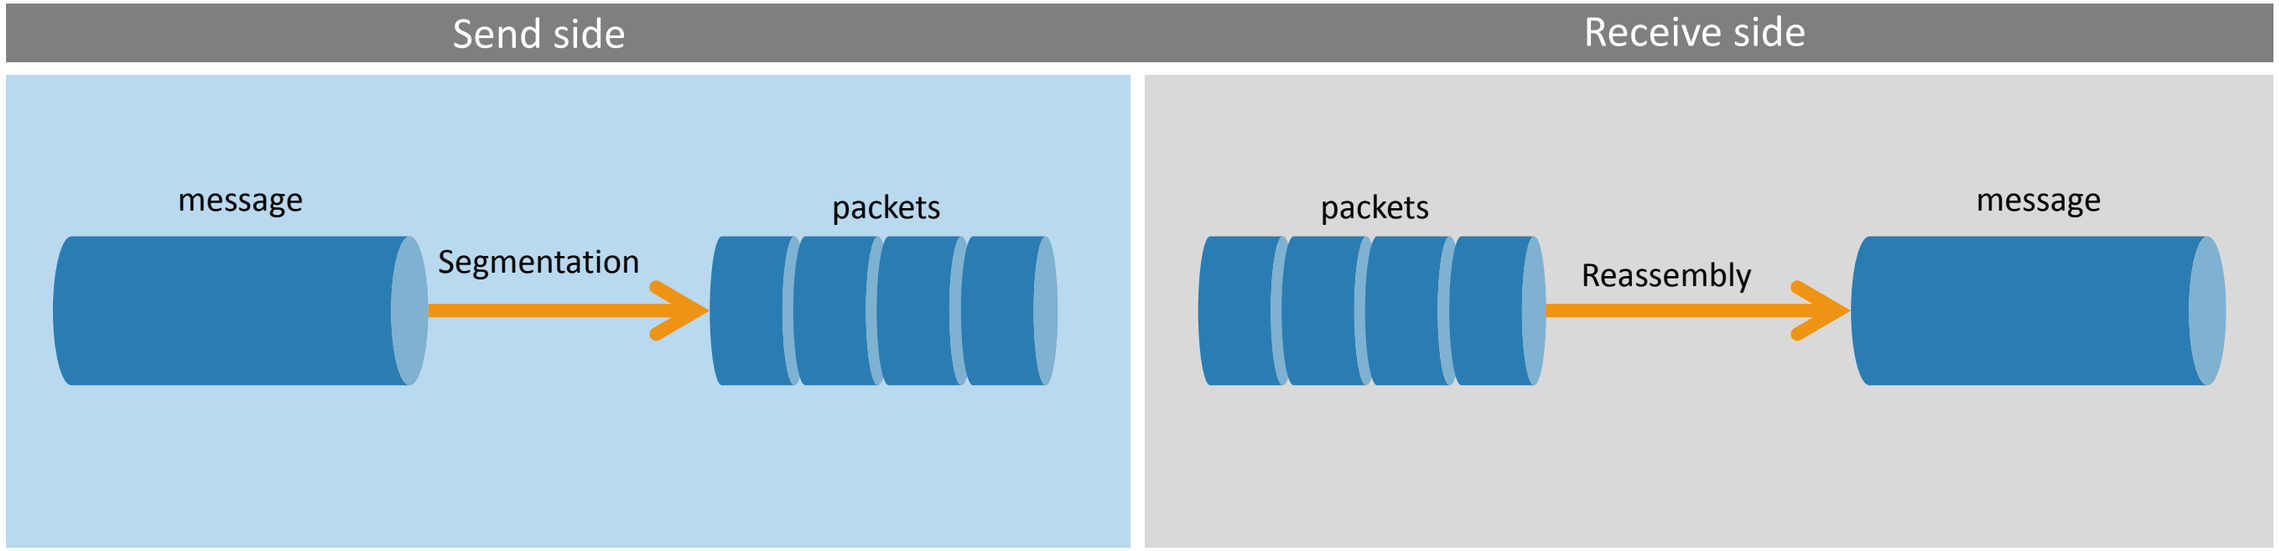

Adapter Offload - Application Services

- Provides end-to-end communication services for applications
- End to end “virtual channel”, connecting two applications that exist in entirely separate address spaces, should be created
- The endpoints of the channel are called Queue Pairs (QPs), each QP represents one end of a channel
- Queue Pairs are the structure by which applications access hardware based transport services

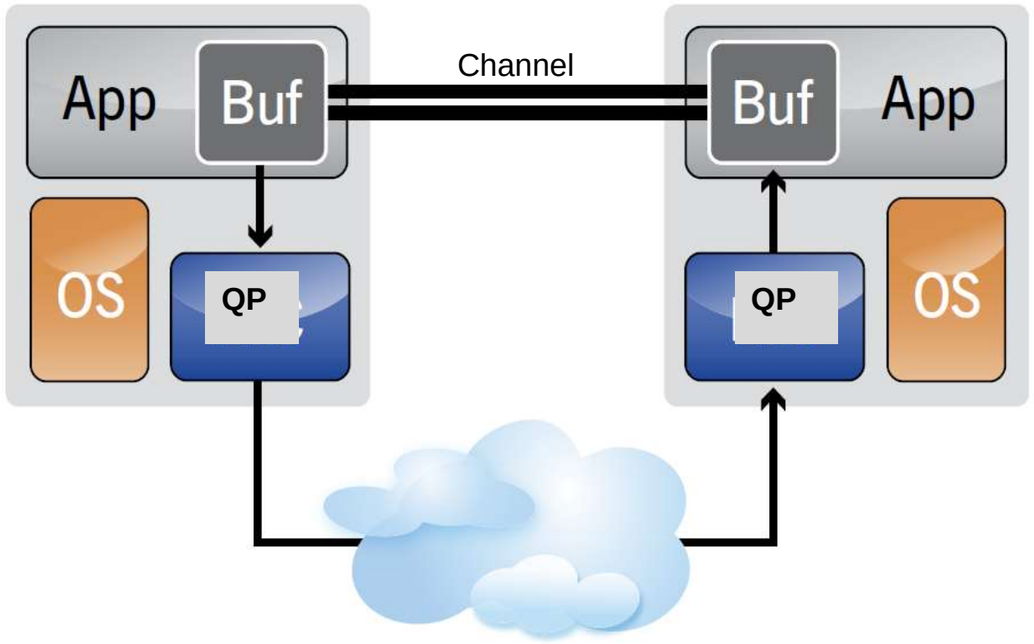

Queue Pairs (QP)

- A QP contains two work queues; a Send Queue and a Receive Queue
- Send and receive queues are created as a pair
- A Queue Pair is identified by its **Queue Pair Number**
- Applications have direct access to the QPs
- If an application requires more than one connection, more QPs are created

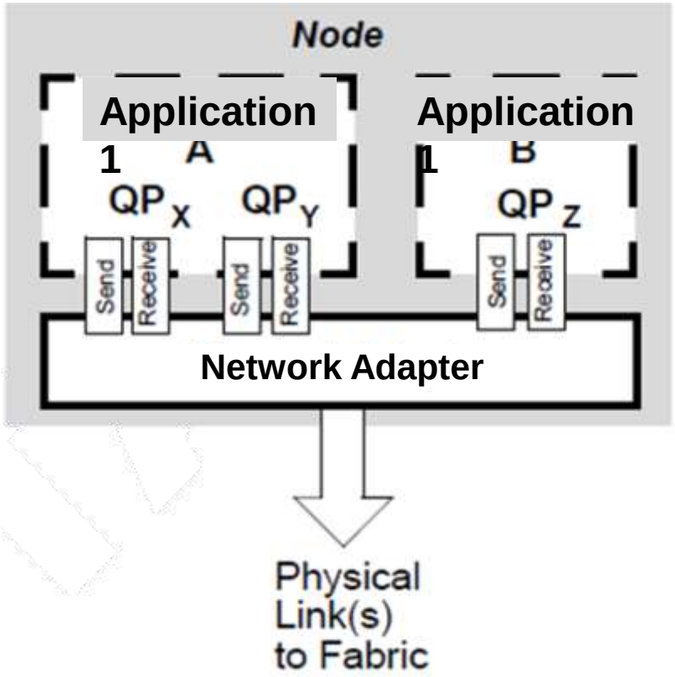

Work Flow

- An application posts a Work Request (WR) to a work queue
- A WR placed in the work queue is called Work Queue Element (WQE)
- WQE is an instruction to transmit a message on the channel
- When the network adapter completes a WQE a Completion Queue Element (CQE) is placed on a completion queue
- CQE notifies the application about the status of completed WR

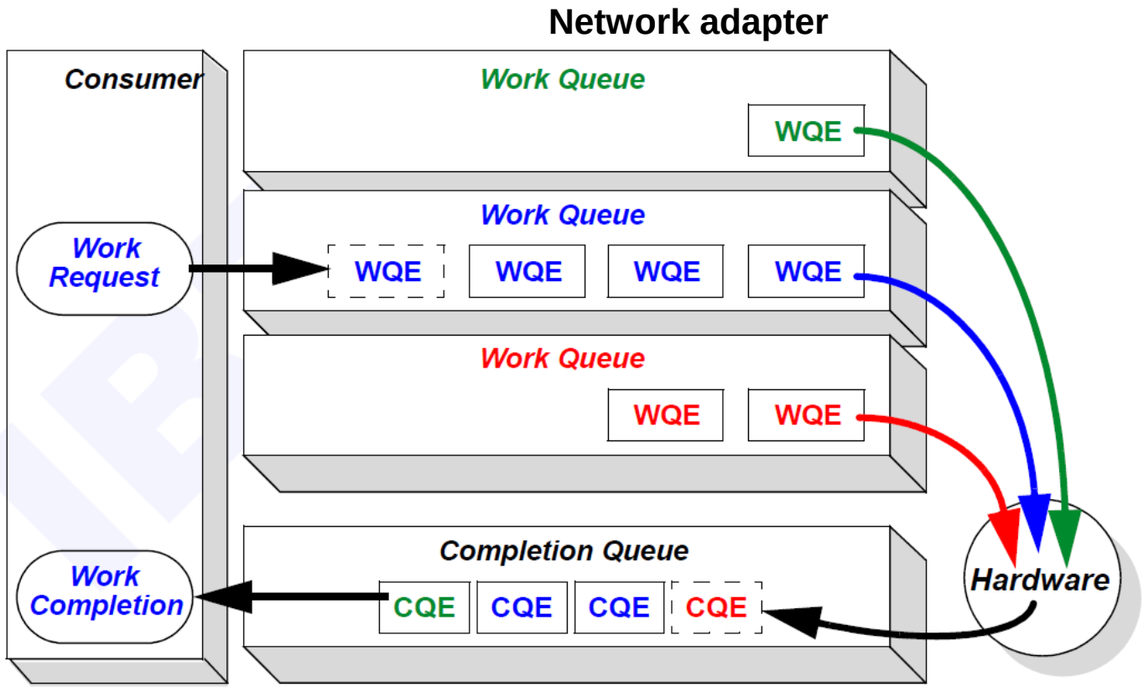

Send / Receive Operation Example

3. The network adapter executes the send WQE:

- gathers data from the memory region according to gather entries
- sends data on the wire
- generates a CQE

4. When the message arrives to the receive side, network adapter executes the receive WQE:

- scatters data to the memory region according to scatter entries
- generates a CQE

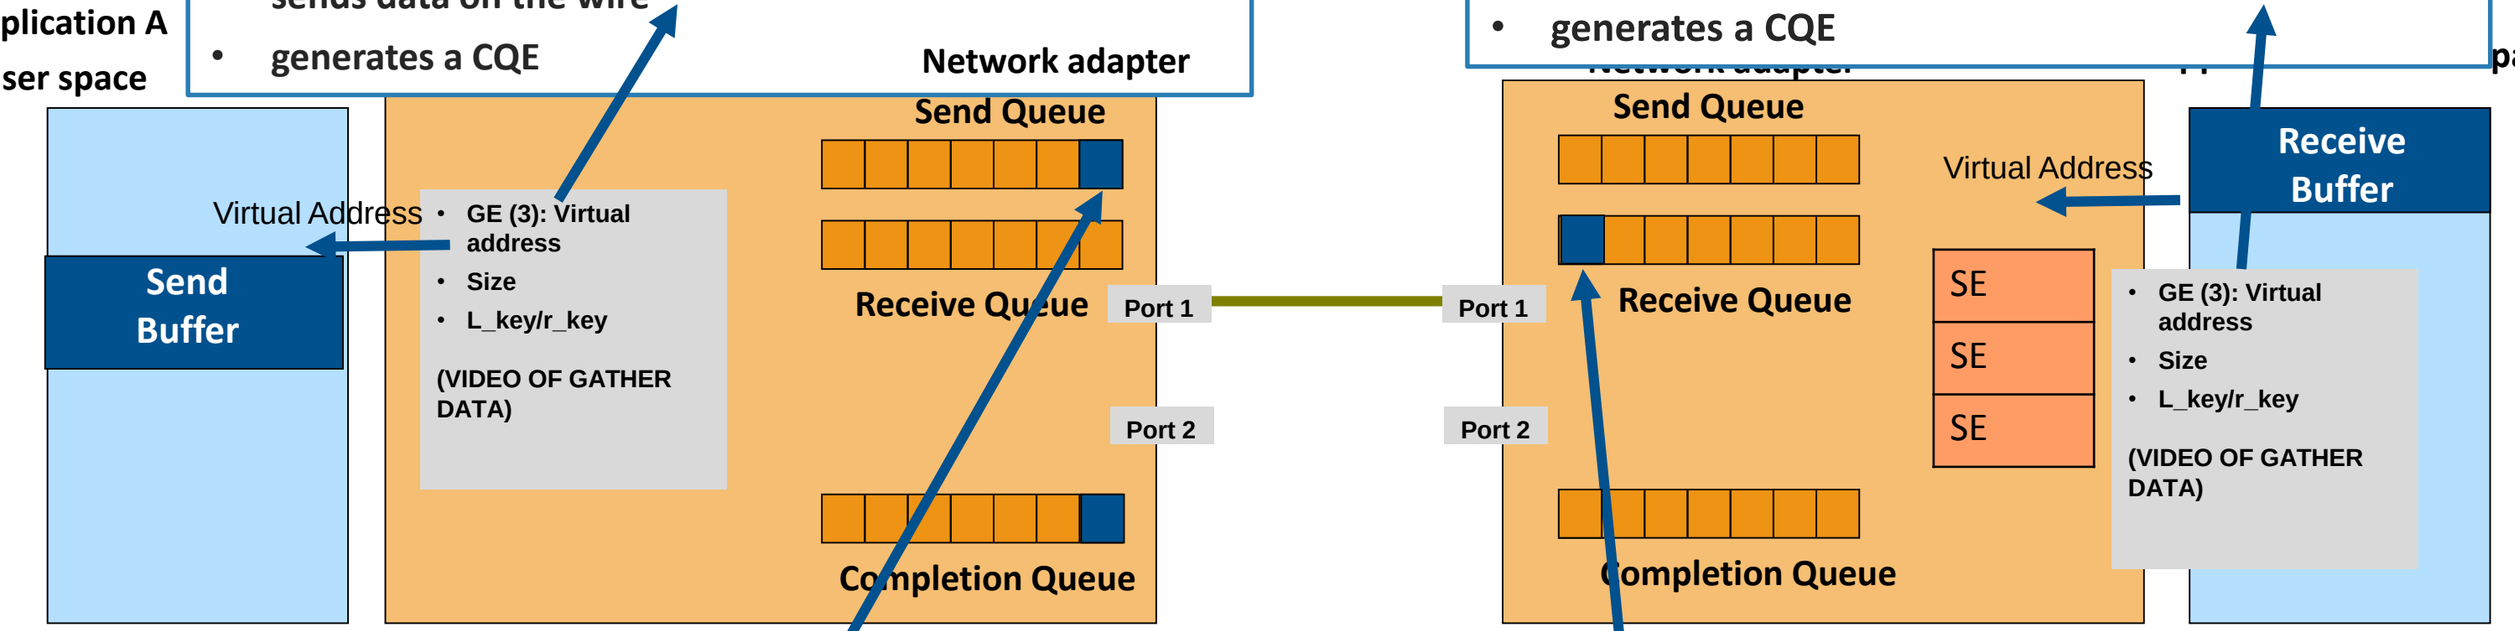

2. The send application places a Work Request in the Send Queue. The send work request contains a pointer to the

1. The receive application posts a Work Request in the Receive Queue. The receive work request contains a pointer to the virtual address.

Note that for Unreliable transport, CQE is generated after data is sent on wire.
For reliable transport, , CQE is generated after acknowledgment is received from remote side

Remote Direct Memory Access (RDMA)

Remote Direct Memory Access (RDMA)

- RDMA supports zero-copy networking by enabling the network adapter to transfer data directly to or from application memory:
  - RDMA Read
  - RDMA Write
- RDMA eliminates the need to copy data between application memory and the data buffers in the operating system

RDMA Write Operation Example

4. The adapter gathers data from the memory region according to gather entries, sends data on the wire, and generates a CQE.

3. When the RDMA Write arrives to the responder, it signal the requester that it is ready to receive.

5. When the data arrives to the responder, the network adapter checks the address and r_key, and writes data directly to the application memory.

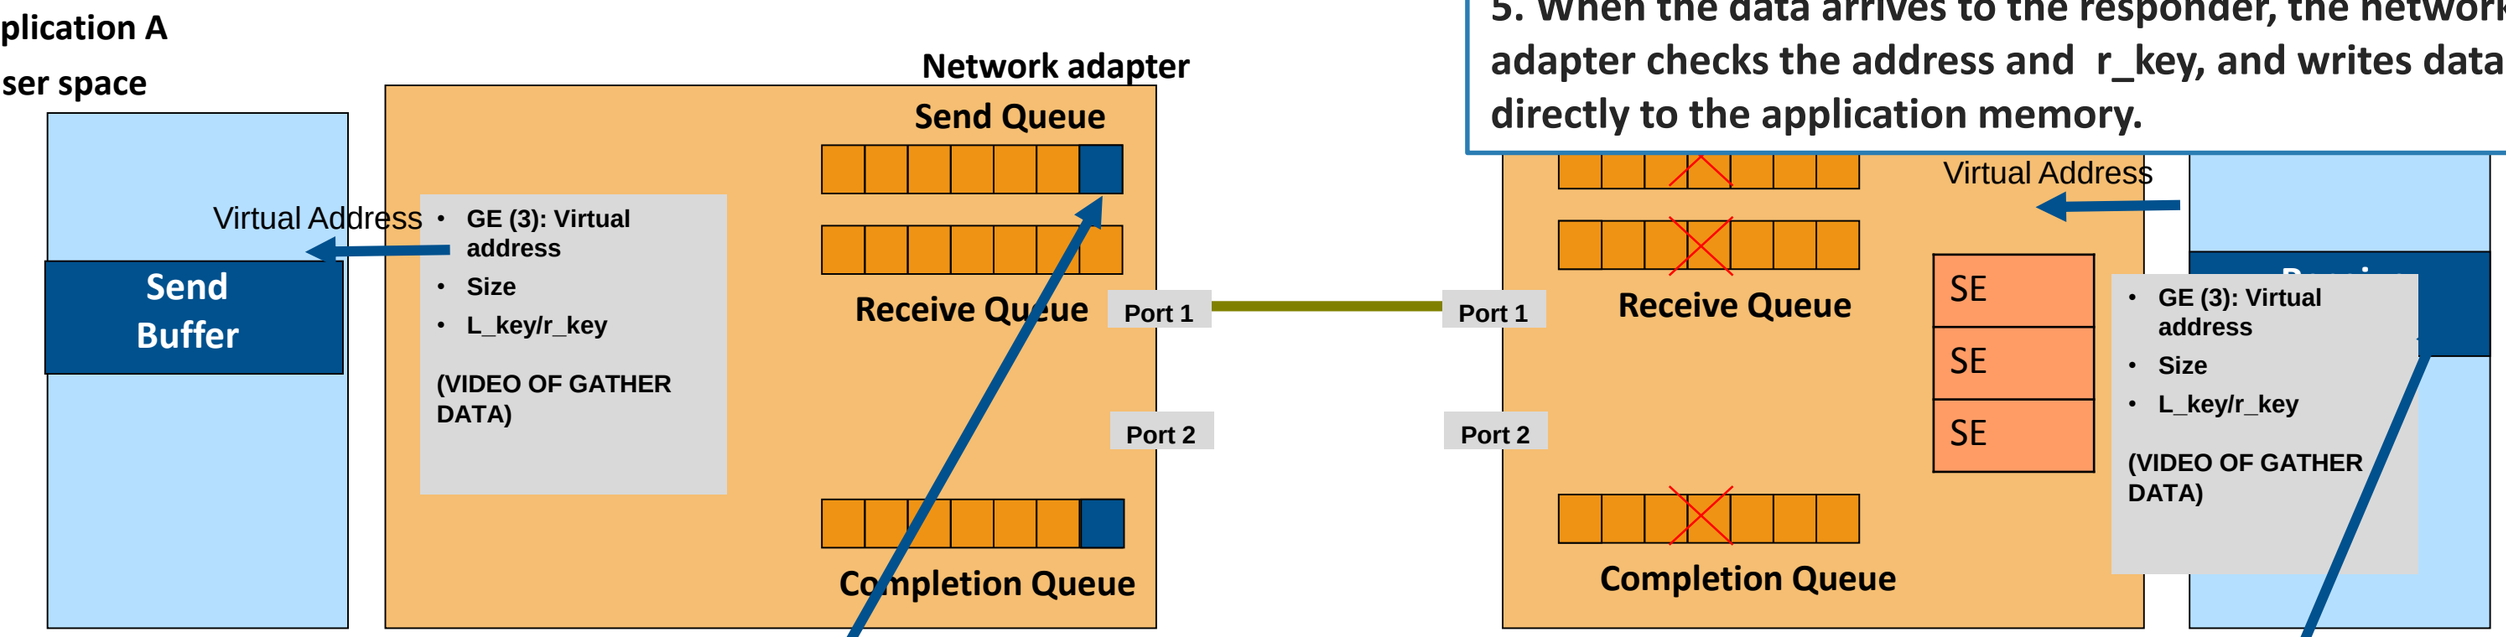

2. The requester places a RDMA write request in the Send Queue. The adapter executes RDMA Write WQE and sends to the responder R_key and virtual address.

1. Application B performs memory registration and passes address and r_key to application A.

RDMA Read Operation Example

3. When the RDMA read arrives to the responder, it signals the requester that it is ready to send. The network adapter checks the address and l_key, read data directly from the application memory and sends data on the wire. No Send WQE is required.

4. When the message arrives to the requester, the network adapter executes receive WQE. The adapter scatter data to the memory region according to scatter entries and generates a CQE.

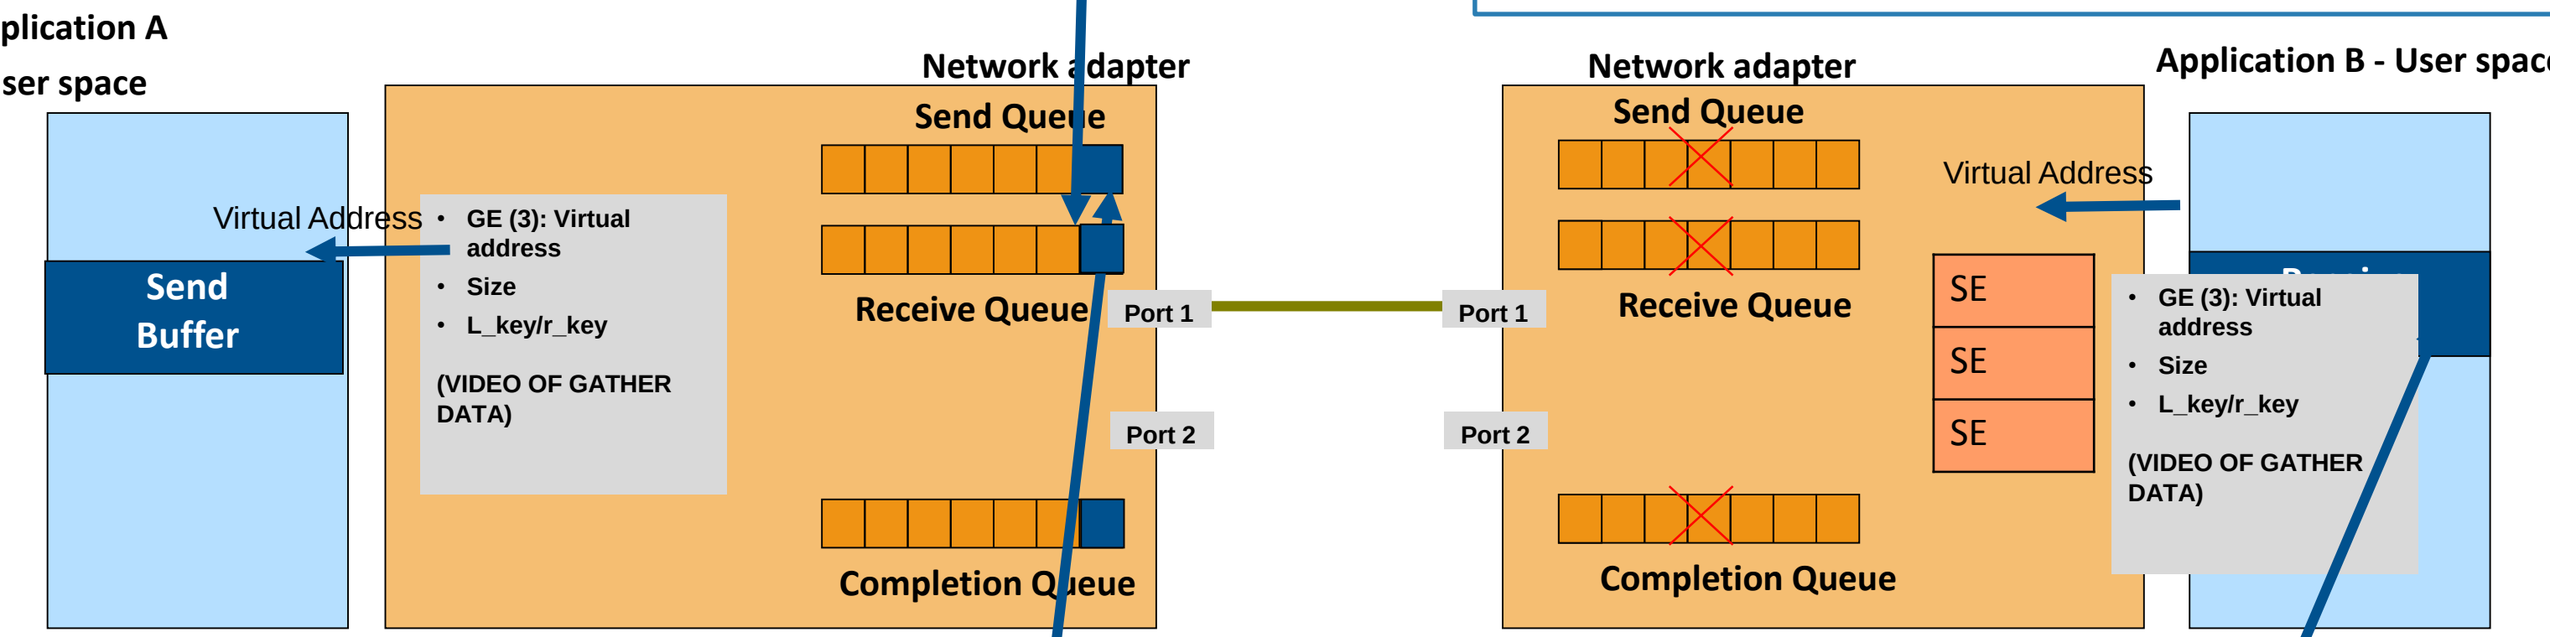

2. The requester places a RDMA read work request in the Send Queue. The adapter executes RDMA read WQE and sends to the responder l_key and virtual address.

The application posts a Work Request in the Receive Queue.

1. Application B performs memory registration and passes the virtual address and l_key to application A.

Reliable Connection (RC)

- A Queue Pair can send and receive messages to/from only one other QP
- Reliable transport:
  - The receiver sends Acknowledgment if packets arrive in order
  - The send QP maintains a timer to catch packets that did not arrive or for which an ACK was lost
  - Packet Sequence Number (PSN) – used by the receiver to identify lost packets
  - Retransmission reduces performance or may break a connection
- RC connection is very similar to a TCP connection

- Usages:
  - Implemented mainly with RDMA operations
  - Good for MPI applications, Storage, and multicast application for FSI customers

Unreliable Connection (UC)

- A Queue Pair can send and receive messages to/from only one other QP
- Unreliable transport:
  - Messages with errors are not retransmitted by the transport
  - Error handling must be provided by a higher level protocol

- Usages:
  - Good for streaming data

Unreliable Datagram (UD/Raw ETH)

- A Queue Pair may send and receive single-packet messages to/from any other QPs
- Message cannot be divided into packets - Message size equal to MTU
- Unreliable transport:
  - Ordering and delivery are not guaranteed
  - Delivered packets may be dropped by the receiver

*A UD connection is very similar to a UDP connection

- Usages:
  - Good for streaming data, MPI applications

ROCE (RDMA over Converged Ethernet)

- IBTA Collaboration on RoCE
  - Steering Committee: Cray, Emulex, HP, IBM, Intel, Mellanox, Microsoft, Oracle,
  - RoCE specification first released in 2010
  - Most widely deployed Ethernet RDMA standard

- Standardization paves way for multi-vendor interoperable solutions RoCEv2 Specification

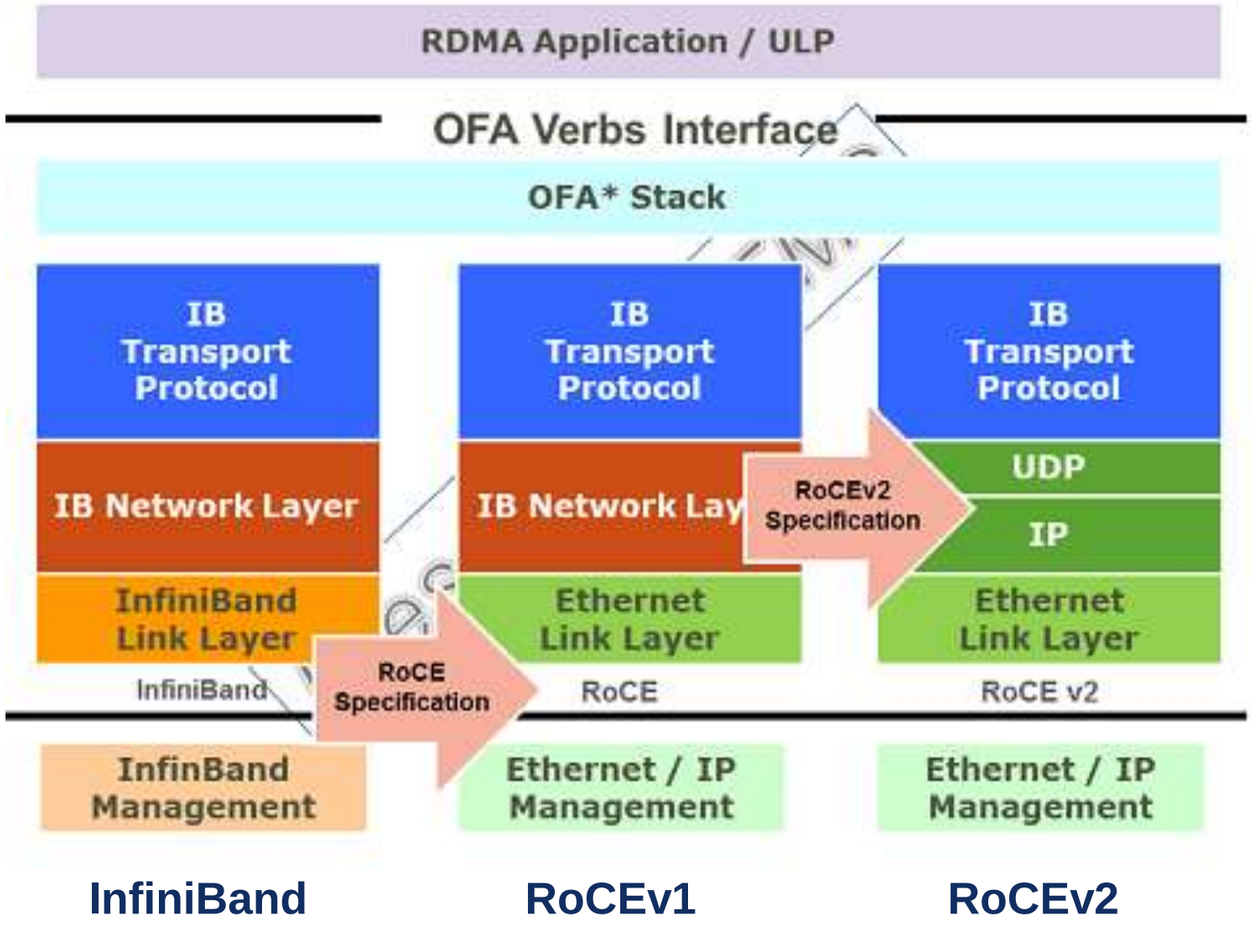

Messaging Acceleration (AKA Kernel bypass)

- Key Features:
  - Dynamically-linked user-space Linux library for accelerating messaging and streaming traffic.
  - A network traffic off-loader, transparently enhances performance of socket applications over Verbs API
  - Provides a unified acceleration platform for all standard communications methods such as TCP, UDP (MC and Unicast)
  - Active polling + kernel bypass eliminate cpu involvement and remove interrupts and context switches
  - A BSD-Socket compliant dynamically linked library
- Main Benefits:
  - Lowest latency, jitter and CPU usage
  - Highest throughput and packet rates
  - No application changes required
  - Single-sided implementation
  - Scales well to 100's of nodes, 1000's of subs

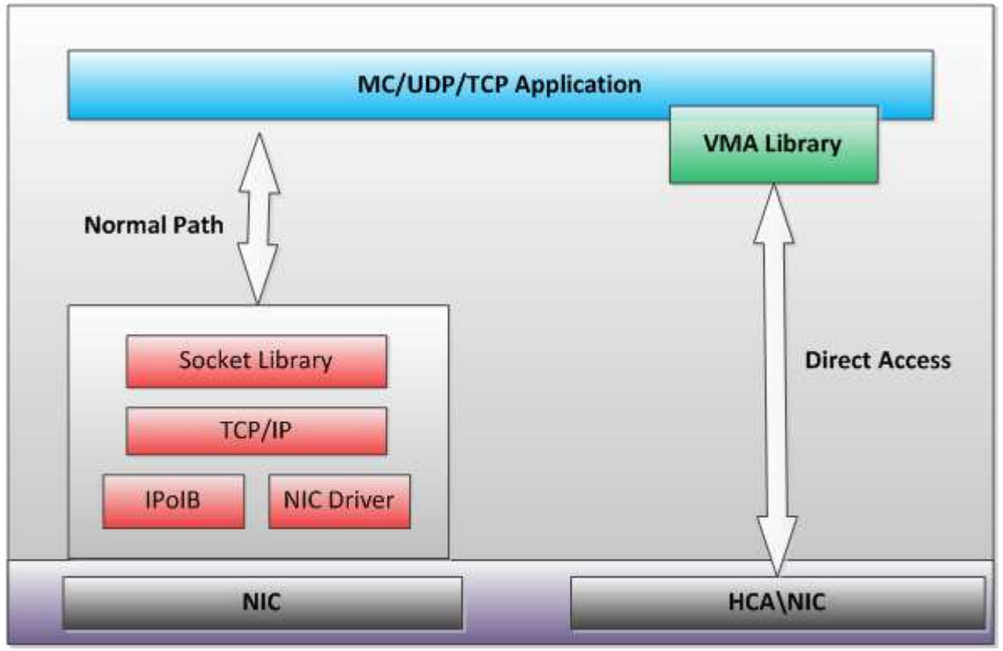

Microsoft Windows Server 2012 R2 Storage Solutions

- SMB Direct
  - SMB 3.0 over RDMA
  - New class of enterprise file storage
- Better performance, lower CPU overhead
- Fibre Channel replacement at a lower cost and higher performance
- Leverages Windows Server 2012 R2 Mellanox inbox drivers
  - InfiniBand & RoCE
- Accelerates Microsoft Hyper-V and SQL Server
- **No need to change the Application**

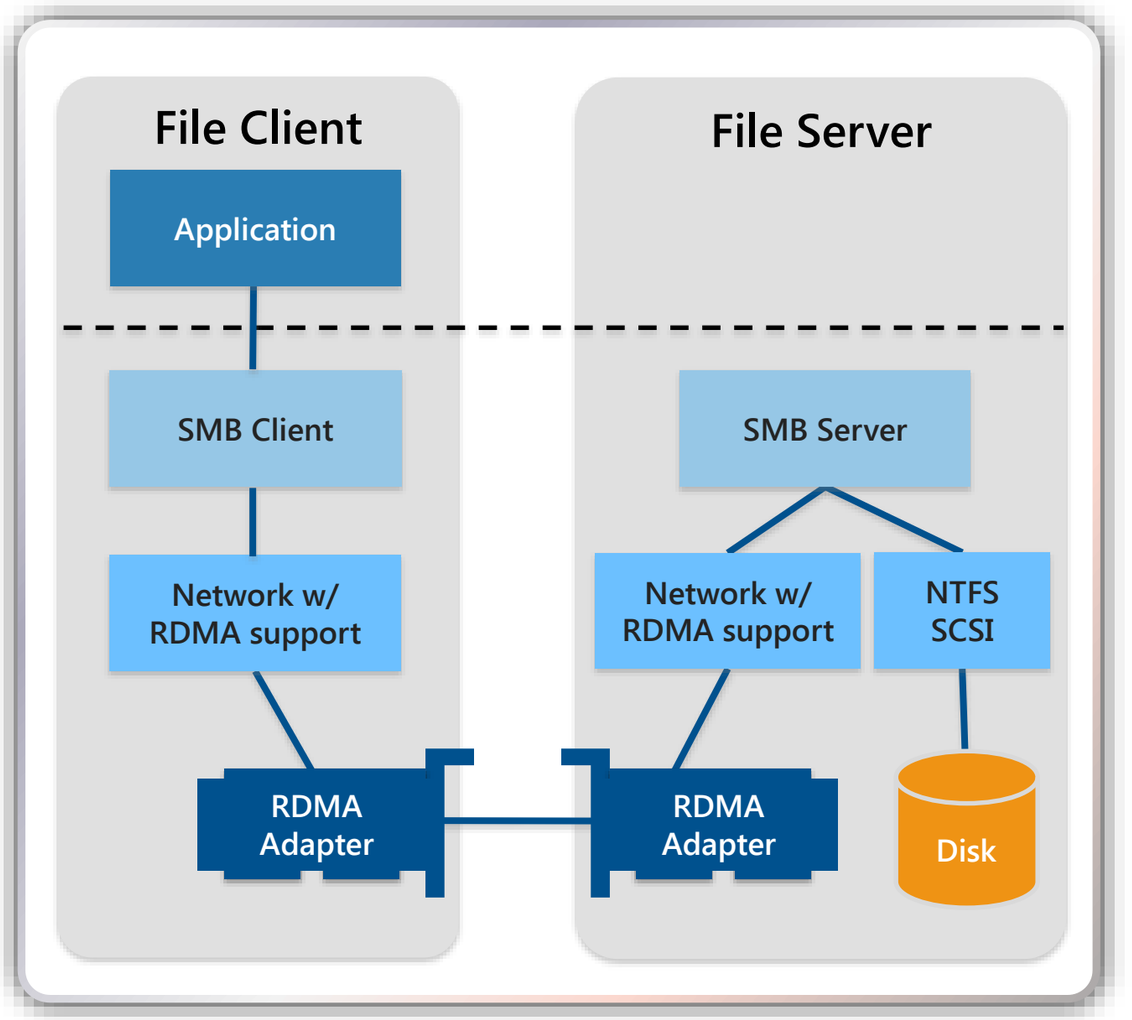

Offloads for Virtualization / Cloud

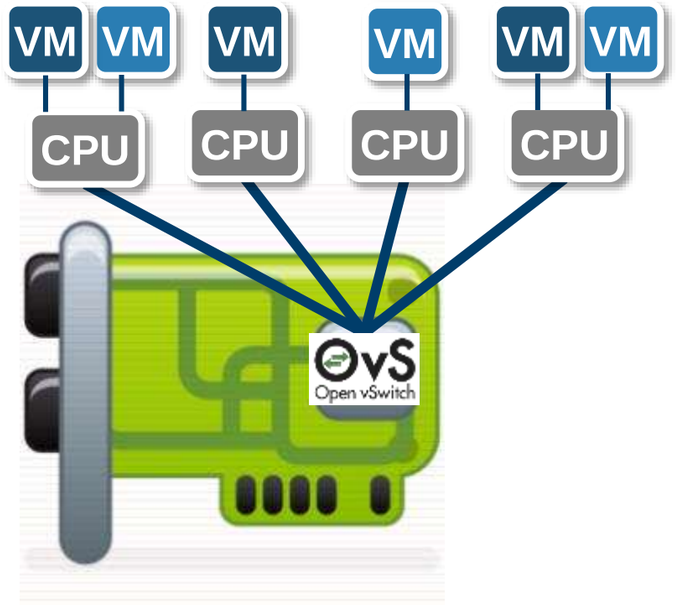

Embedded Switch
Hardware OVS Switch

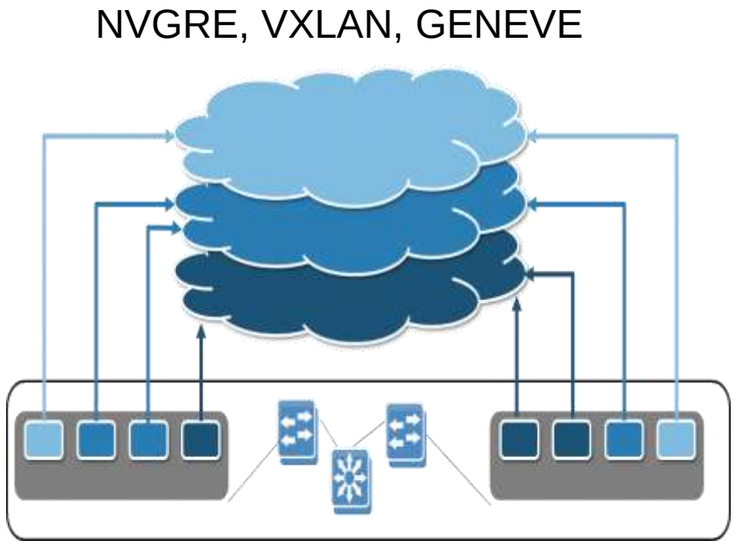

Virtual Overlay Network Acceleration

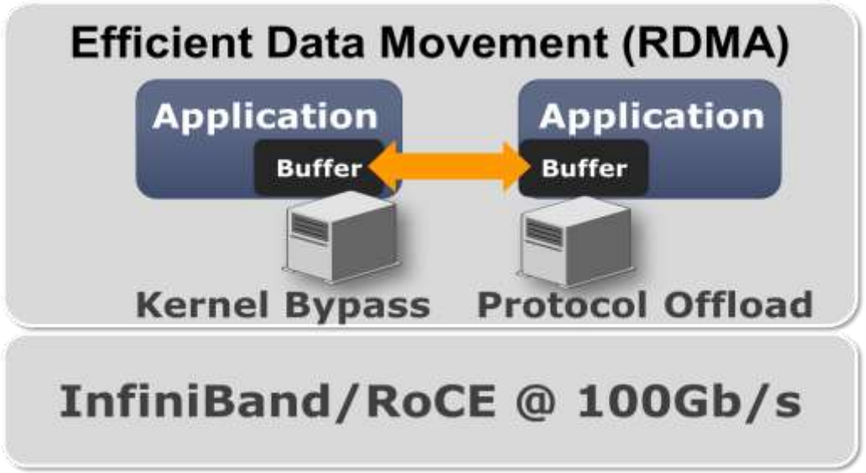

Efficient Data Movement With RDMA

- Efficient Data Movement
  - Multi-Host & eSwitch: Embedded hardware OVS switch – Advance Flow Steering Engine
  - Virtual network acceleration (VXLAN, NVGRE, GENEVE)
  - RDMA – Efficient Data Exchange - Low Latency, Low CPU Overhead

Use Cases of High Speed Networking

Use Case Example: RoCE/RDMA @ Microsoft Azure Cloud

"To make storage cheaper we use lots more network! How do we make Azure Storage scale? RoCE (RDMA over Converged Ethernet) enabled at 40GbE for Windows Azure Storage, achieving massive COGS savings"

Microsoft

Keynote
Albert Greenberg, Microsoft
SDN Azure Infrastructure

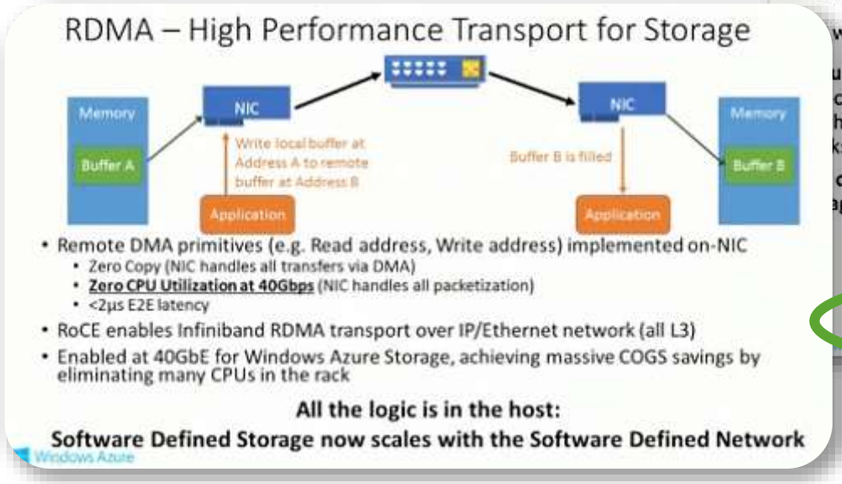

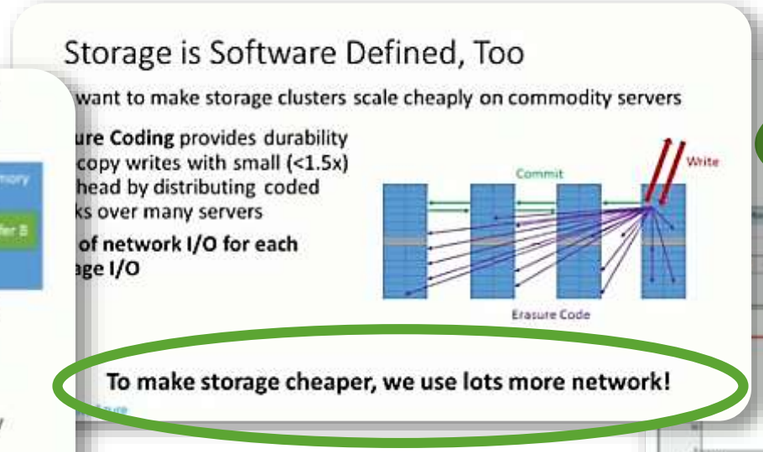

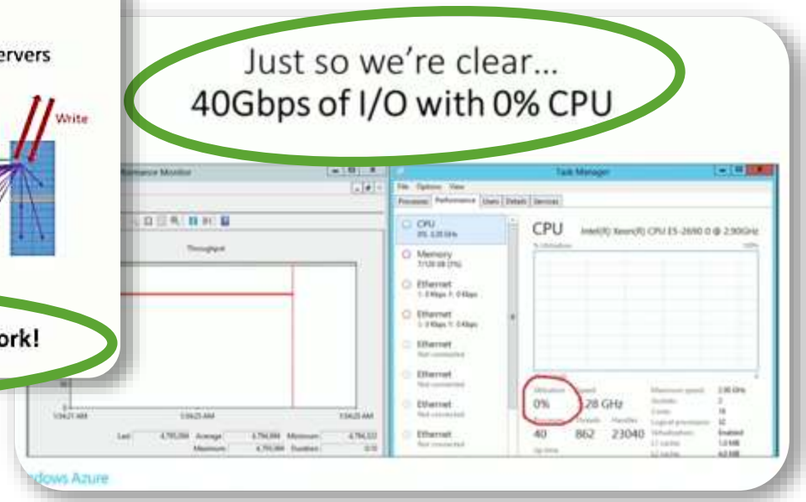

Microsoft Keynote at Open Networking Summit 2014 on RDMA

Use Case Example:
IB/RDMA in Enterprise Application Appliances

Microsoft IBM ORACLE TERADATA UNISYS

Exadata Database Machine

Exalogic Elastic Cloud

Oracle Database Appliance

Exalytics In-Memory Machine

Big Data Appliance

SPARC SuperCluster T4-4

RDMA is Pervasive in Enterprise Application Appliances

Use Case Example:
Powering the World's Fastest Exchange

SIX Swiss Exchange

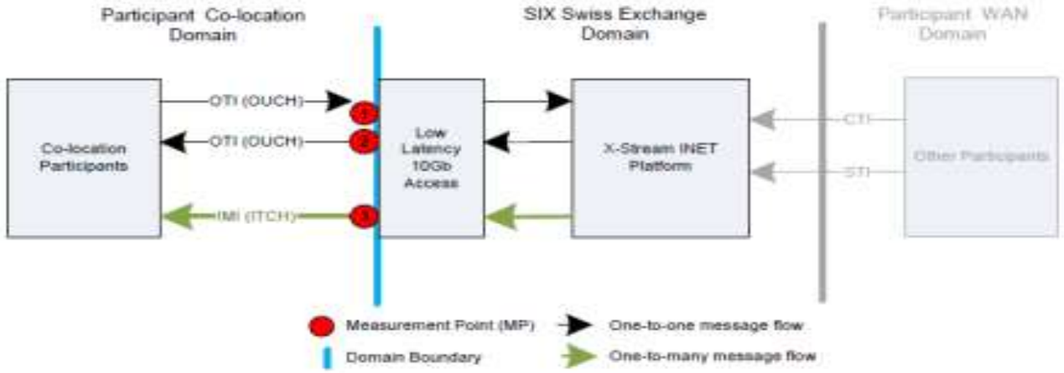

"Client focus combined with state-of-the-art technology are essential to attracting additional liquidity.

Right now, we are setting a new standard for trading, risk management and the future growth of SIX Swiss Exchange."

Christian Katz, CEO SIX Swiss Exchange

http://www.six-swiss-exchange.com/download/participants/trading/x-stream_inet_performance_measurement_details.pdf

Record-breaking Door-to-Door Latency

Stock Exchanges Running InfiniBand

Use Case 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

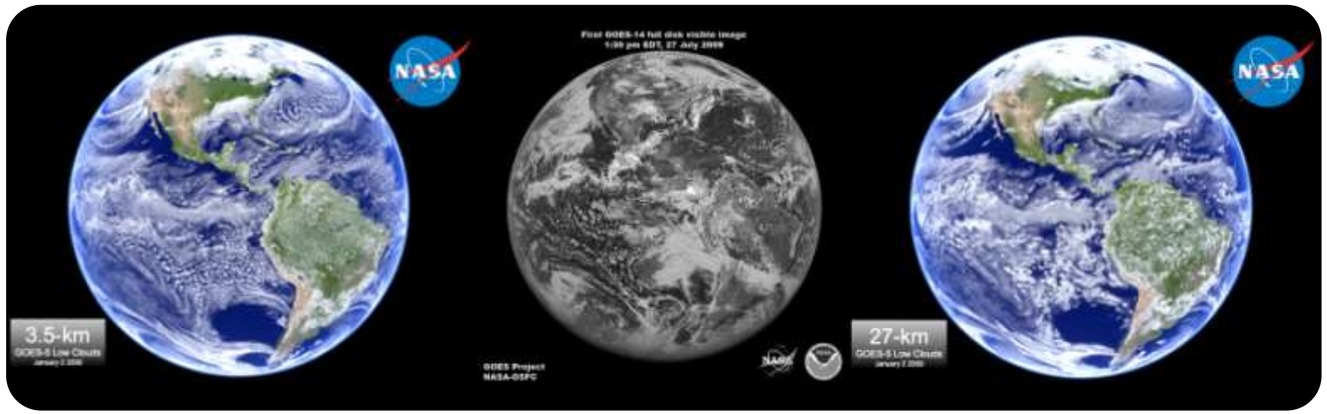

TOP500 Interconnect Trends

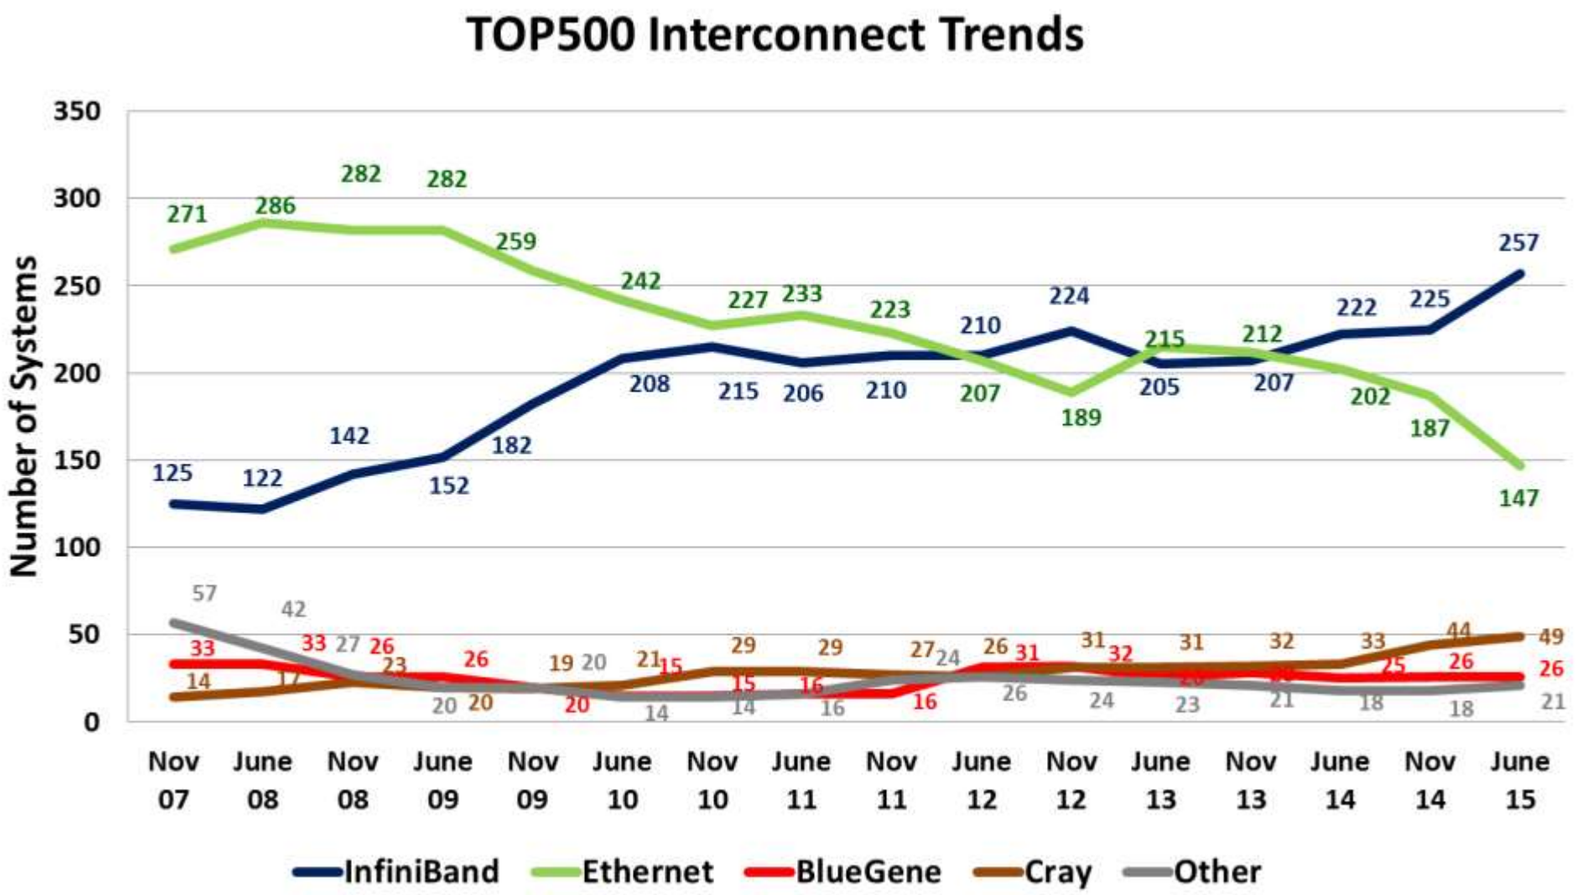

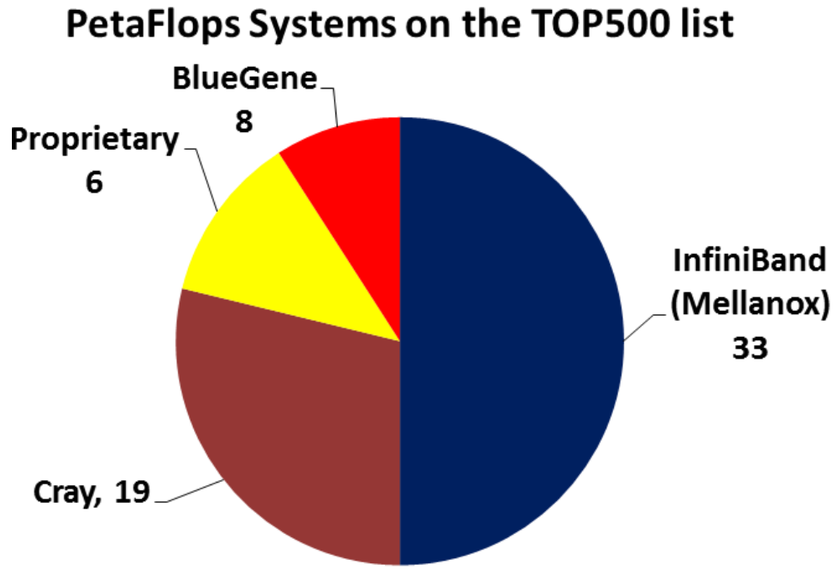

- InfiniBand is the de-facto interconnect solution for performance demanding applications

TOP500 InfiniBand Accelerated Systems

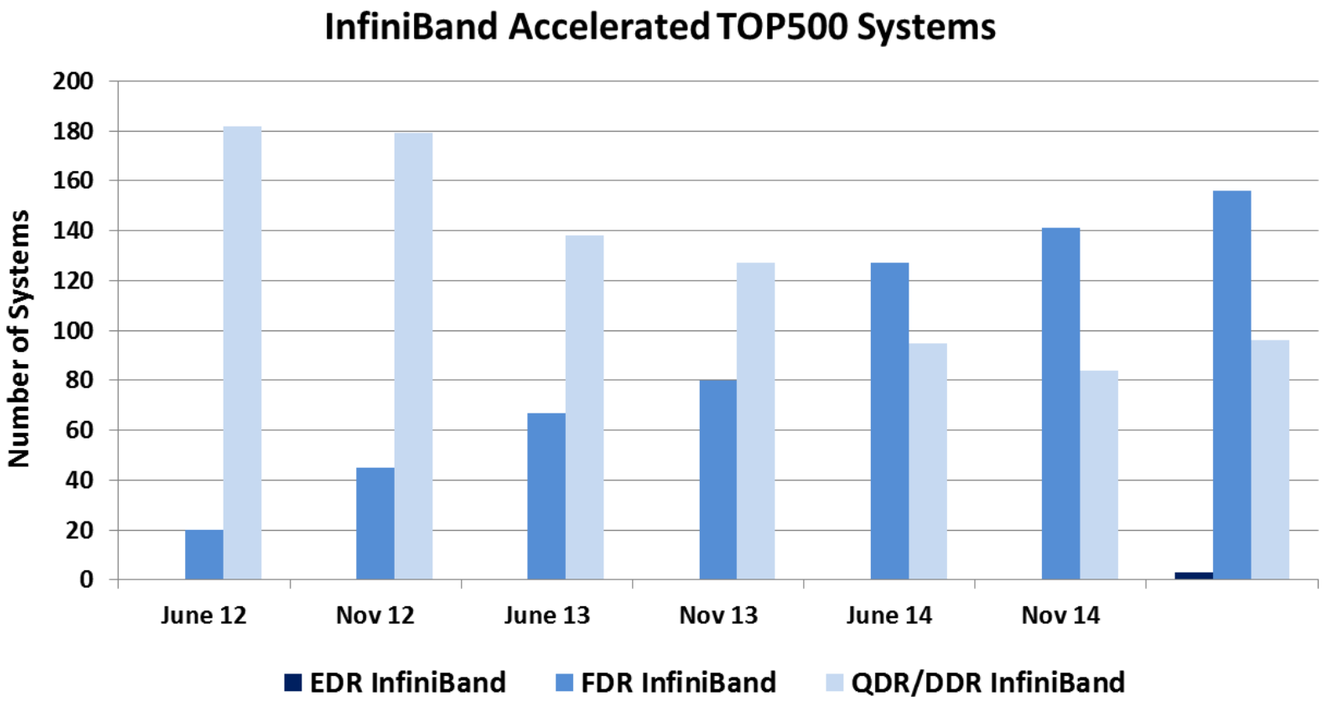

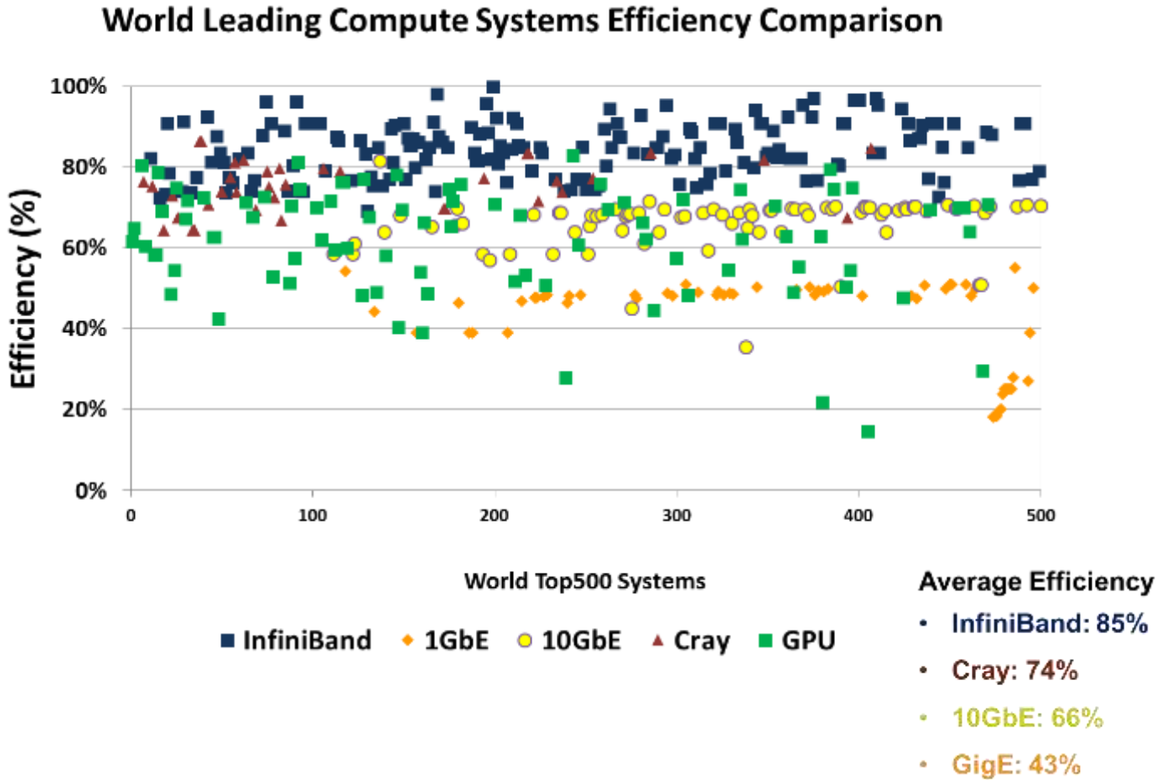

- Number of Mellanox FDR InfiniBand systems grew 23% from June'14 to June'15
- EDR InfiniBand entered the list with 3 systems

Mellanox Product Overview

Leading Supplier of End-to-End Interconnect Solutions

Comprehensive End-to-End InfiniBand and Ethernet Portfolio

ICs	Adapter Cards	Switches/Gateways	Software and Services	Metro / WAN	Cables/Modules
	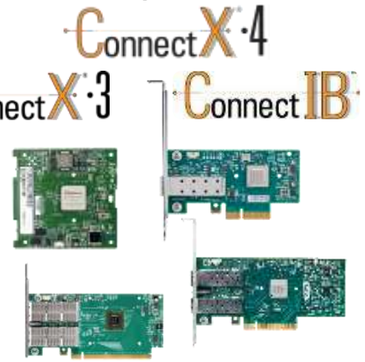				

At the Speeds of 10, 25, 40, 50, 56 and 100 Gigabit per Second

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

Switch

Spectrum™

32 100GbE Ports, 64 25/50GbE Ports

(10 / 25 / 40 / 50 / 100GbE)

Throughput of 6.4Tb/s

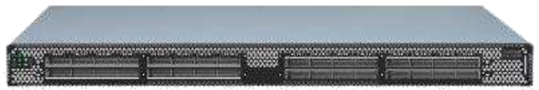

Interconnect

LinkX™

Copper (Passive, Active)

Optical Cables (VCSEL)

Silicon Photonics

Mellanox Adapters Portfolio

ConnectX-3 PRO

- 10GbE high runner
- Leading cloud offering
- On Ethernet, price parity with ConnectX-3

ConnectIB

- High Performance InfiniBand
- FDR switch ports connectivity
- Shares ConnectX-4 architecture

ConnectX-4

- Highest Performance, VPI
- Addresses all markets and needs
- Software features graduate roll out

ConnectX-4 Lx

- 10 / 25 / 40 / 50GbE focused
- Similar to ConnectX-4 feature-wise
- Availability: Q3’ 2015

Technical Comparison

Feature	ConnectX-3	ConnectX-3 Pro	ConnectX-4 Lx	ConnectX-4	Connect-IB
Max Speed	**FDR / 40/56GbE**	**FDR / 40/56GbE**	**10/40/25/50GbE /FDR**	**EDR / 100GbE**	**100Gb FDR X2 ports**
PCIe	Gen3 x8	Gen3 x8	Gen3 x8	Gen3 x16 x8	Gen3 x16 X8
Package	17x17mm	17x17mm	17X17mm	25x25mm	21x21mm
Message rate	33Mpps	36Mpps	60Mpps (40GbE)	150Mpps	137Mpps
Latency	0.7us	0.7us	0.71us	0.63us	0.7us
Max Power (2 ports)	6.1W	6.2W	8.7W	15.6W	10.6W

Main Feature Comparison

Feature	ConnectX-3	ConnectX-3 Pro	ConnectX-4 Lx	ConnectX-4	Connect-IB
Speed Rates	FDR/ 10/40/56GbE	FDR/ 10/40/56GbE	10/25/40GbE (1 port)/ 50GbE*	EDR/10/25/40/50/100GbE	FDR 100Gb X2 ports
Message Rate	33Mpps	36Mpps	60 Mpps (40GbE)	150Mpps	137Mpps
RDMA / RoCE	+	+	+	+	RDMA
CoreDirect	+	+	+	+	+
GPUDirect RDMA	+	+	+	+	+
Stateless Ethernet Offloads	+	+	+ LRO, LSOv2	+ LRO, LSOv2	
RSS (MAC, VLAN, 5 Tuple)	+	+	++	++	
Virtualization (SR-IOV)	+	+ 127 VFs	+ 64 PFs, 512 VFs	+ 64 PFs, 512 VFs	+
RoCE with Routing Capabilities		+	+	+	
Congestion Control QCN, ECN		+	+	+	
Overlay Networks Offload		+ VXLAN/NVGRE	+ VXLAN/NVGRE/GENEVE	+ VXLAN/NVGRE/GENEVE	
Overlay Networks Encap/Decap			+	+	
Offload of Tunneling Protocols			+	+	
Dynamically Connected Transport			+	+	+
Erasure Coding (RAID Offload)			+	+	
Multi Host			+	+	
T-10 / DIF Signature Handover				+	+
CAPI				+	

ConnectX-3 / Pro Adapters for Open Compute Project

- Open platform drives high volumes
- High performance OCP adapters
- Single and dual port 10GbE and 40GbE
- Servers from multiple ODMs
  - Quanta, Foxconn, Wiwynn, MiTAC
- OCP Server 2.0 compliant



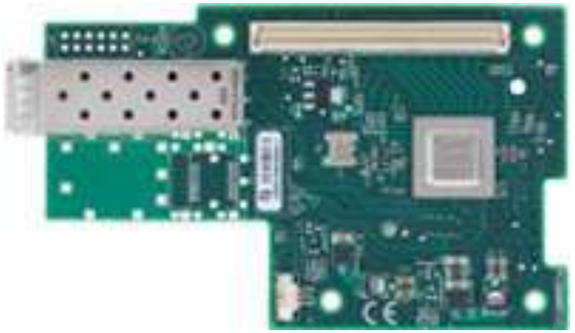


End-to-End Interconnect Solutions for All Platforms

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

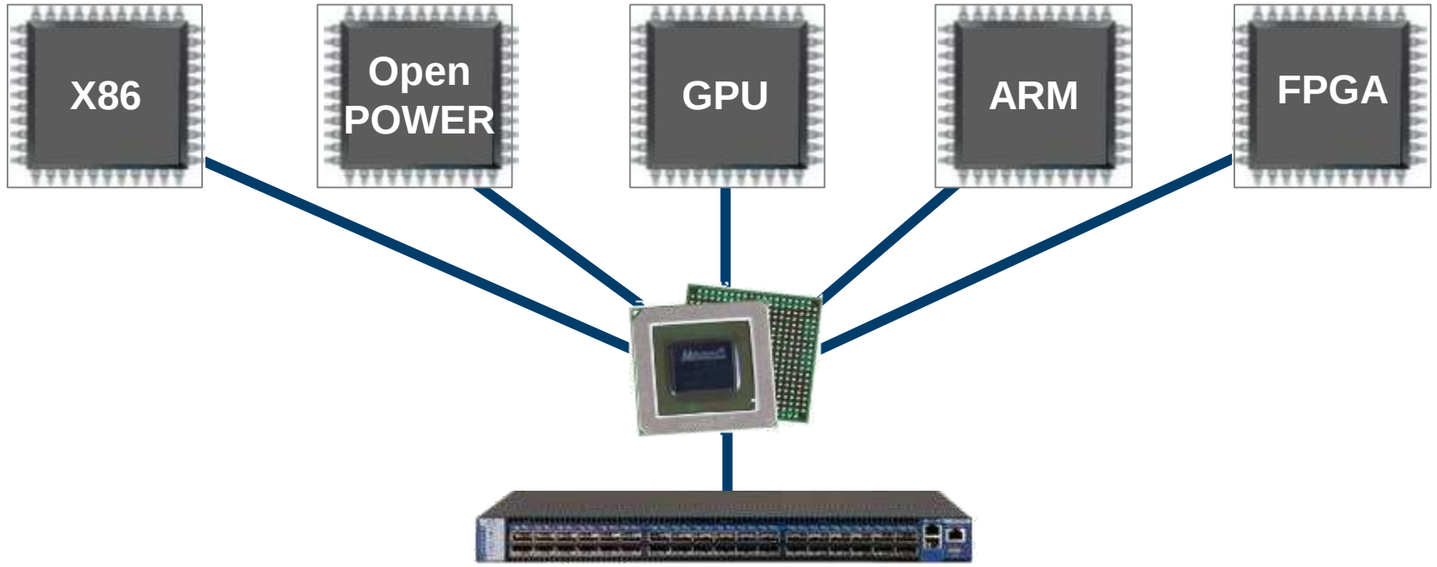

Smart Interconnect to Unleash The Power of All Compute Architectures

ConnectX-4 on Facebook OCP Multi-Host Platform (Yosemite)

ConnectX-4 Multi-Host Adapter

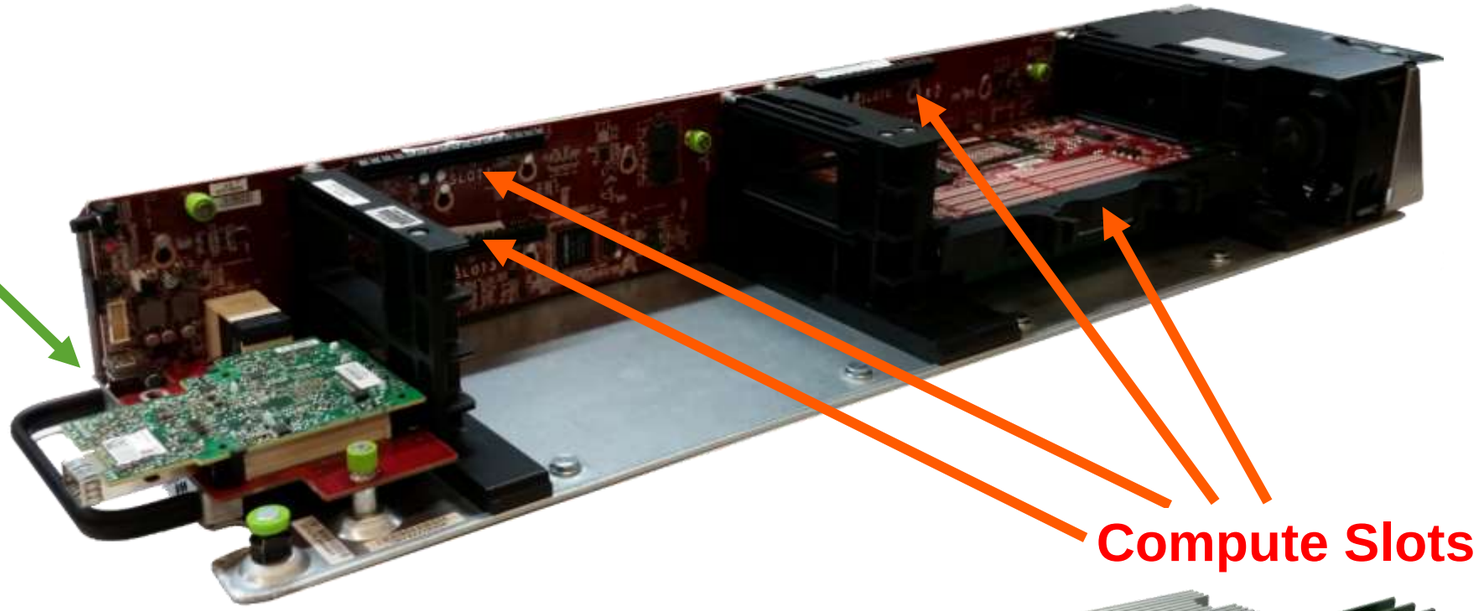

Compute Slots

The Next Generation Compute and Storage Rack Design

Industry’s Broadest Range of Cables and Transceivers

Passive Copper Cables

QSFP/SFP+: 0.5–7m

Hybrid Cables

Copper and fiber, QSFP/SFP+/MPO/LC

Colored Cables

QSFP/SFP+, Copper and Fiber

Short and Long Range Transceivers

Short Range QSFP/SFP+, 30-300m, MMF
Long range QSFP/SFP+, up to 10Km, SMF
Quad to Serial (QSA) Adapters, 40G QSFP to 10G SFP+

Active Optical Cables (AOC)

QSFP: 3-300m, MMF

ICs

TIA, VD, VOA, Mux, De-mux

Mellanox Ethernet Switches

Best ROI – Switch Silicon Example

- What matters in switch silicon?
  - Highest switching capacity
  - Lowest power
  - Lowest latency

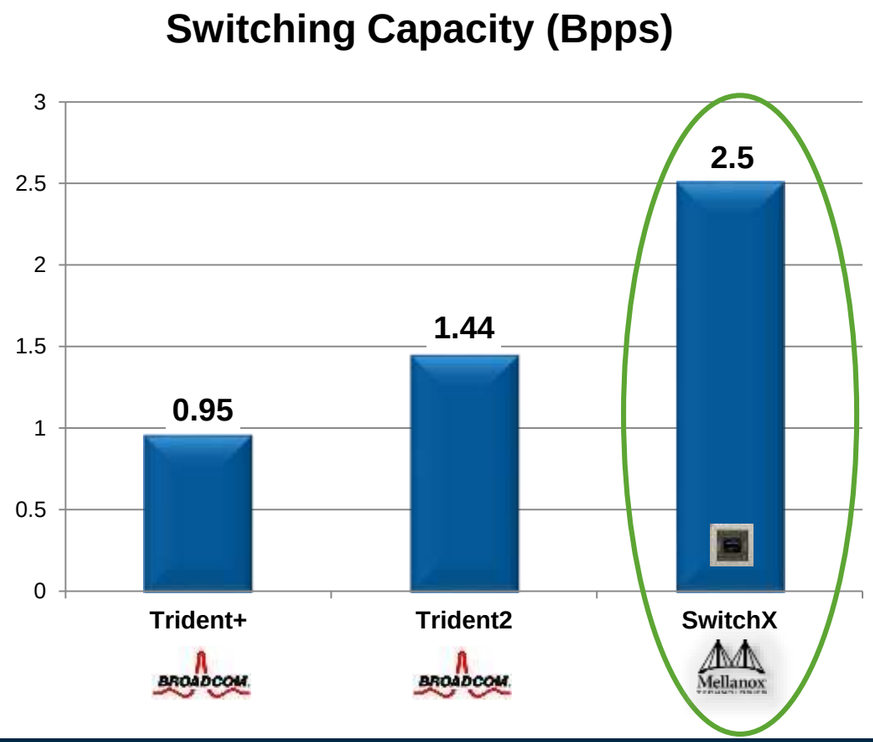

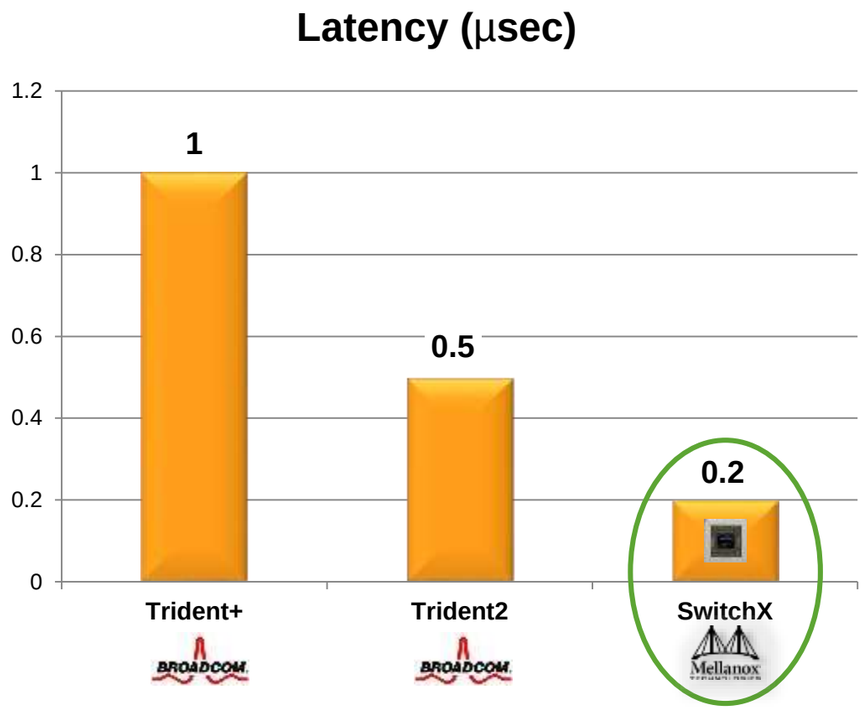

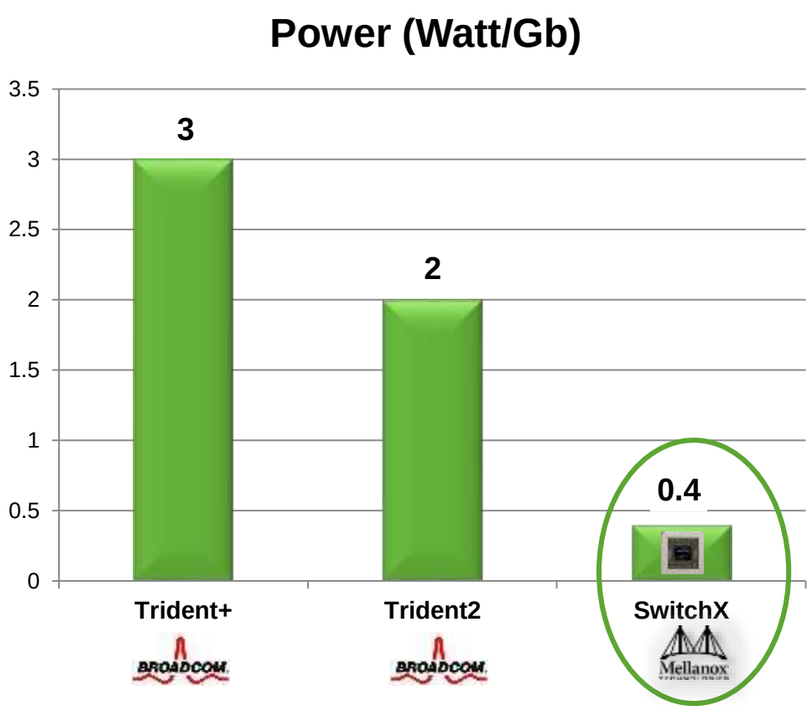

SwitchX-2 Ethernet Switch Portfolio

SX1710 (x86)/1036 (PPC)
The Ideal 40GbE ToR/Aggregation

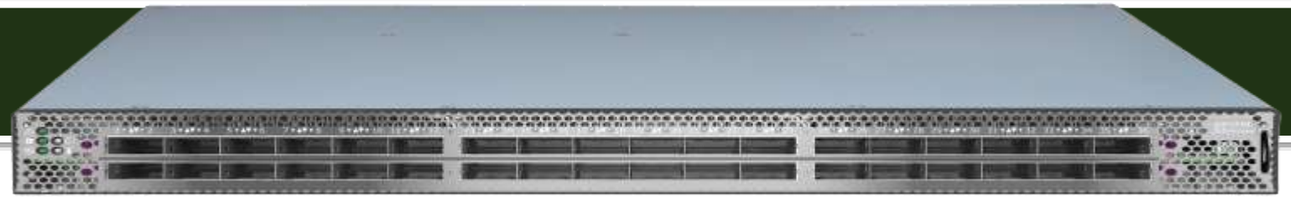

SX1024 (PPC)/1400 (x86)
Non-blocking 10GbE➔ 40GbE ToR

SX1012
Ideal storage/Database 10/40GbE Switch
Classic PoC switch

SX1016
Highest density 10GbE ToR

energy efficiency

SwitchX-2

SX1710 – 91W
SX1036 – 83W
SX1016 – 62W
SX1024 – 75W
SX1012 – 50W

220ns

Open Ethernet Spectrum Switch Portfolio

SN2700 – 32x100GbE (64x50GbE)
The Ideal 100GbE ToR / Aggregation

SN2410 – 8x100GbE + 48x25GbE
25GbE ➔ 100GbE ToR

SN2100 – 16x100GbE ports
Ideal storage/Database 25/100GbE Switch
Highest 25GbE Density Per 1RU

energy efficiency

Spectrum™

300ns

Zero Packet Loss

OPEN ETHERNET™

Mellanox Virtual Modular Switch® (VMS) Solution

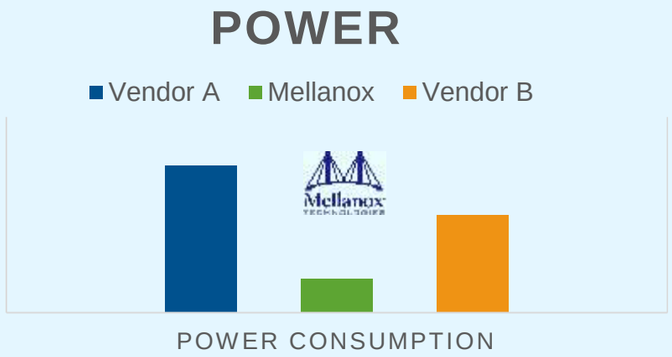

Cost Effective
- Pay as you grow

Standard L3 Scale-out
- ECMP over OSPF/BGP

Automation
- *Configured in minutes (VMS Wizard)*

Flexible
- *10/25/40/50/100GbE ports*

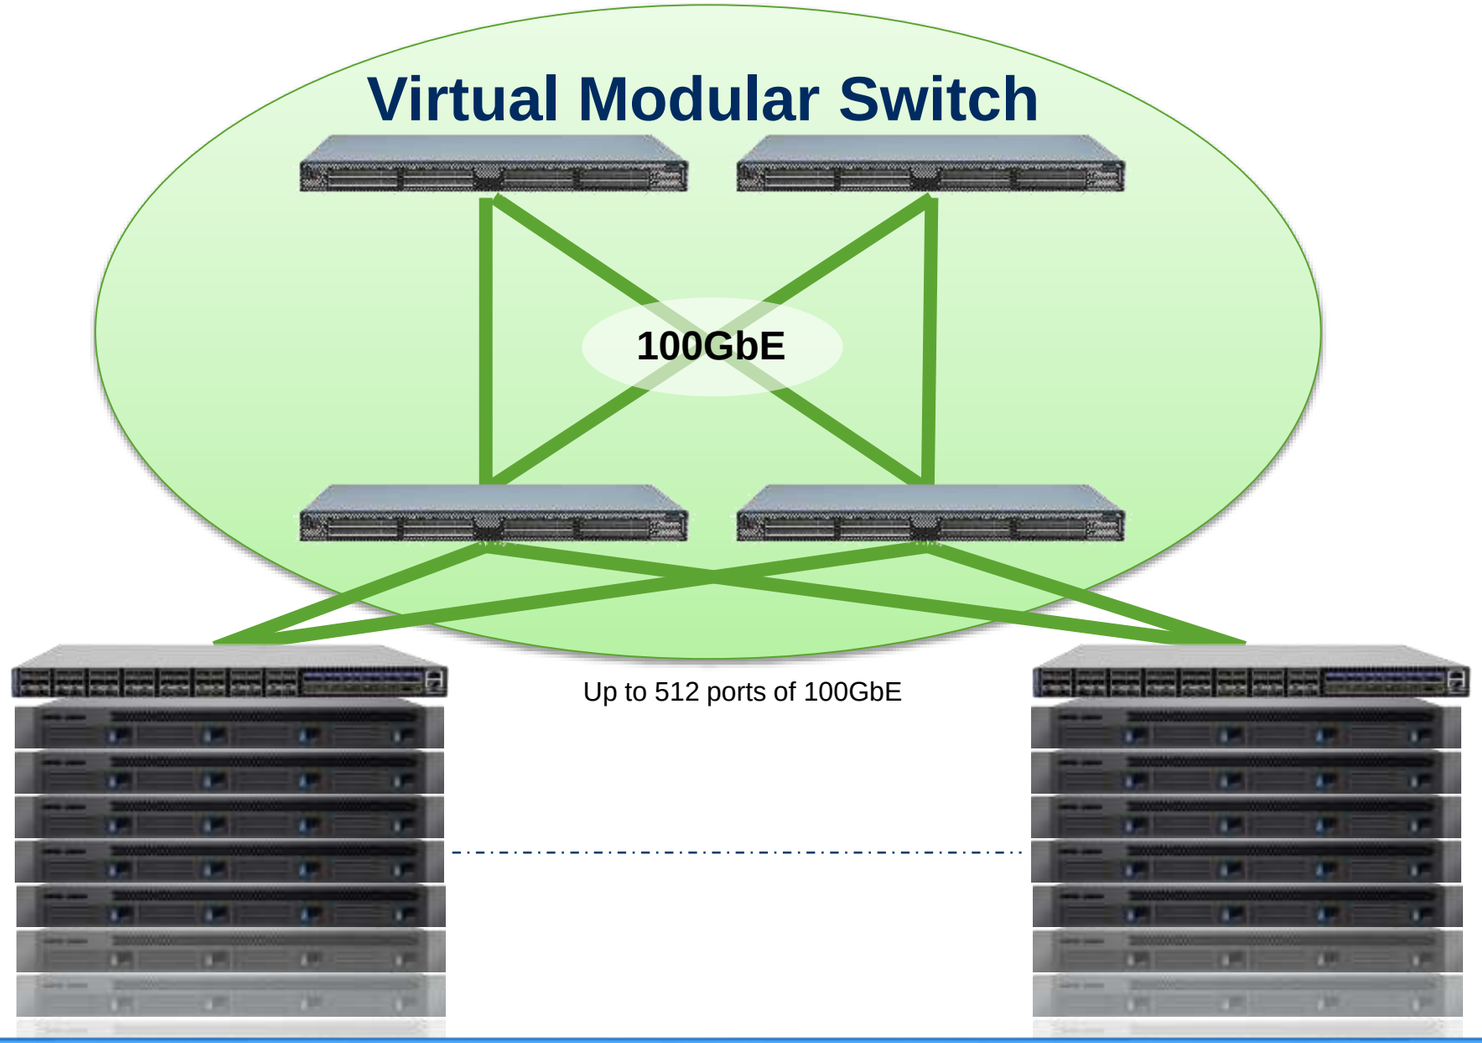

Cost Effective , Scalable , Resilient and Flexible Solution

Tolly Report – February 2015

TEST REPORT
Tolly.

#215111

February 2015

Commissioned by
Mellanox Technologies, Ltd.

Mellanox SwitchX-2 (SX1036) vs. Broadcom StrataXGS Trident II (Arista DCS-7050QX)

Performance Evaluation

Qualifying Data Center Ethernet Networks with RFC2544 at 40Gbps

EXECUTIVE SUMMARY

The demand for data center network performance continues to grow as multi-tenant, public/private clouds and enterprise workloads require that Ethernet switches deliver higher levels of reliability and guaranteed service level agreements. In this environment, unexpected packet loss is unacceptable. In the past Ethernet switches could easily "pass" RFC2544 with no packet loss and did not exhibit large variances in latency. Today, however, that is not the case with some vendors' high-speed switches.

Designing a switch ASIC which operates at 40GbE or higher rates is a different type of challenge and this may be the reason for this new phenomenon where switches fail to pass the very basic RFC2544 at L2 or L3. Today, with the extensive usage of text and short messages, Web2 and large clouds are seeing increasing portions of very small packets which changes the way the network operates.

Mellanox commissioned Tolly to benchmark the 40 Gigabit Ethernet performance of the Mellanox SwitchX-2 ASIC, implemented in the Mellanox SX1036 switch and compare that to the performance of the Broadcom StrataXGS Trident II ASIC, implemented in the Arista Networks DCS-7050QX switch. The Mellanox solution delivered 40GbE wire-speed layer 2 performance with zero frame loss at all frame sizes tested in tests of up to 36 ports. See Table 1.

THE BOTTOM LINE

The Mellanox SwitchX-2 ASIC delivers:

1. Zero-loss, wire-speed throughput at all frame sizes tested from 64- through 9212-byte jumbo frames compared to up to 20% loss and latency up to 97,980 ns for Arista Networks
2. Better latency than the Arista Networks DCS-7050QX at all frame sizes tested, up to 96% lower in one test
3. True cut-through switching, while the Arista Networks runs store & forward for 10GbE-10GbE traffic within the same rack for typical top-of-rack topologies

Predictable, High Performance Networks & No Packet Loss!

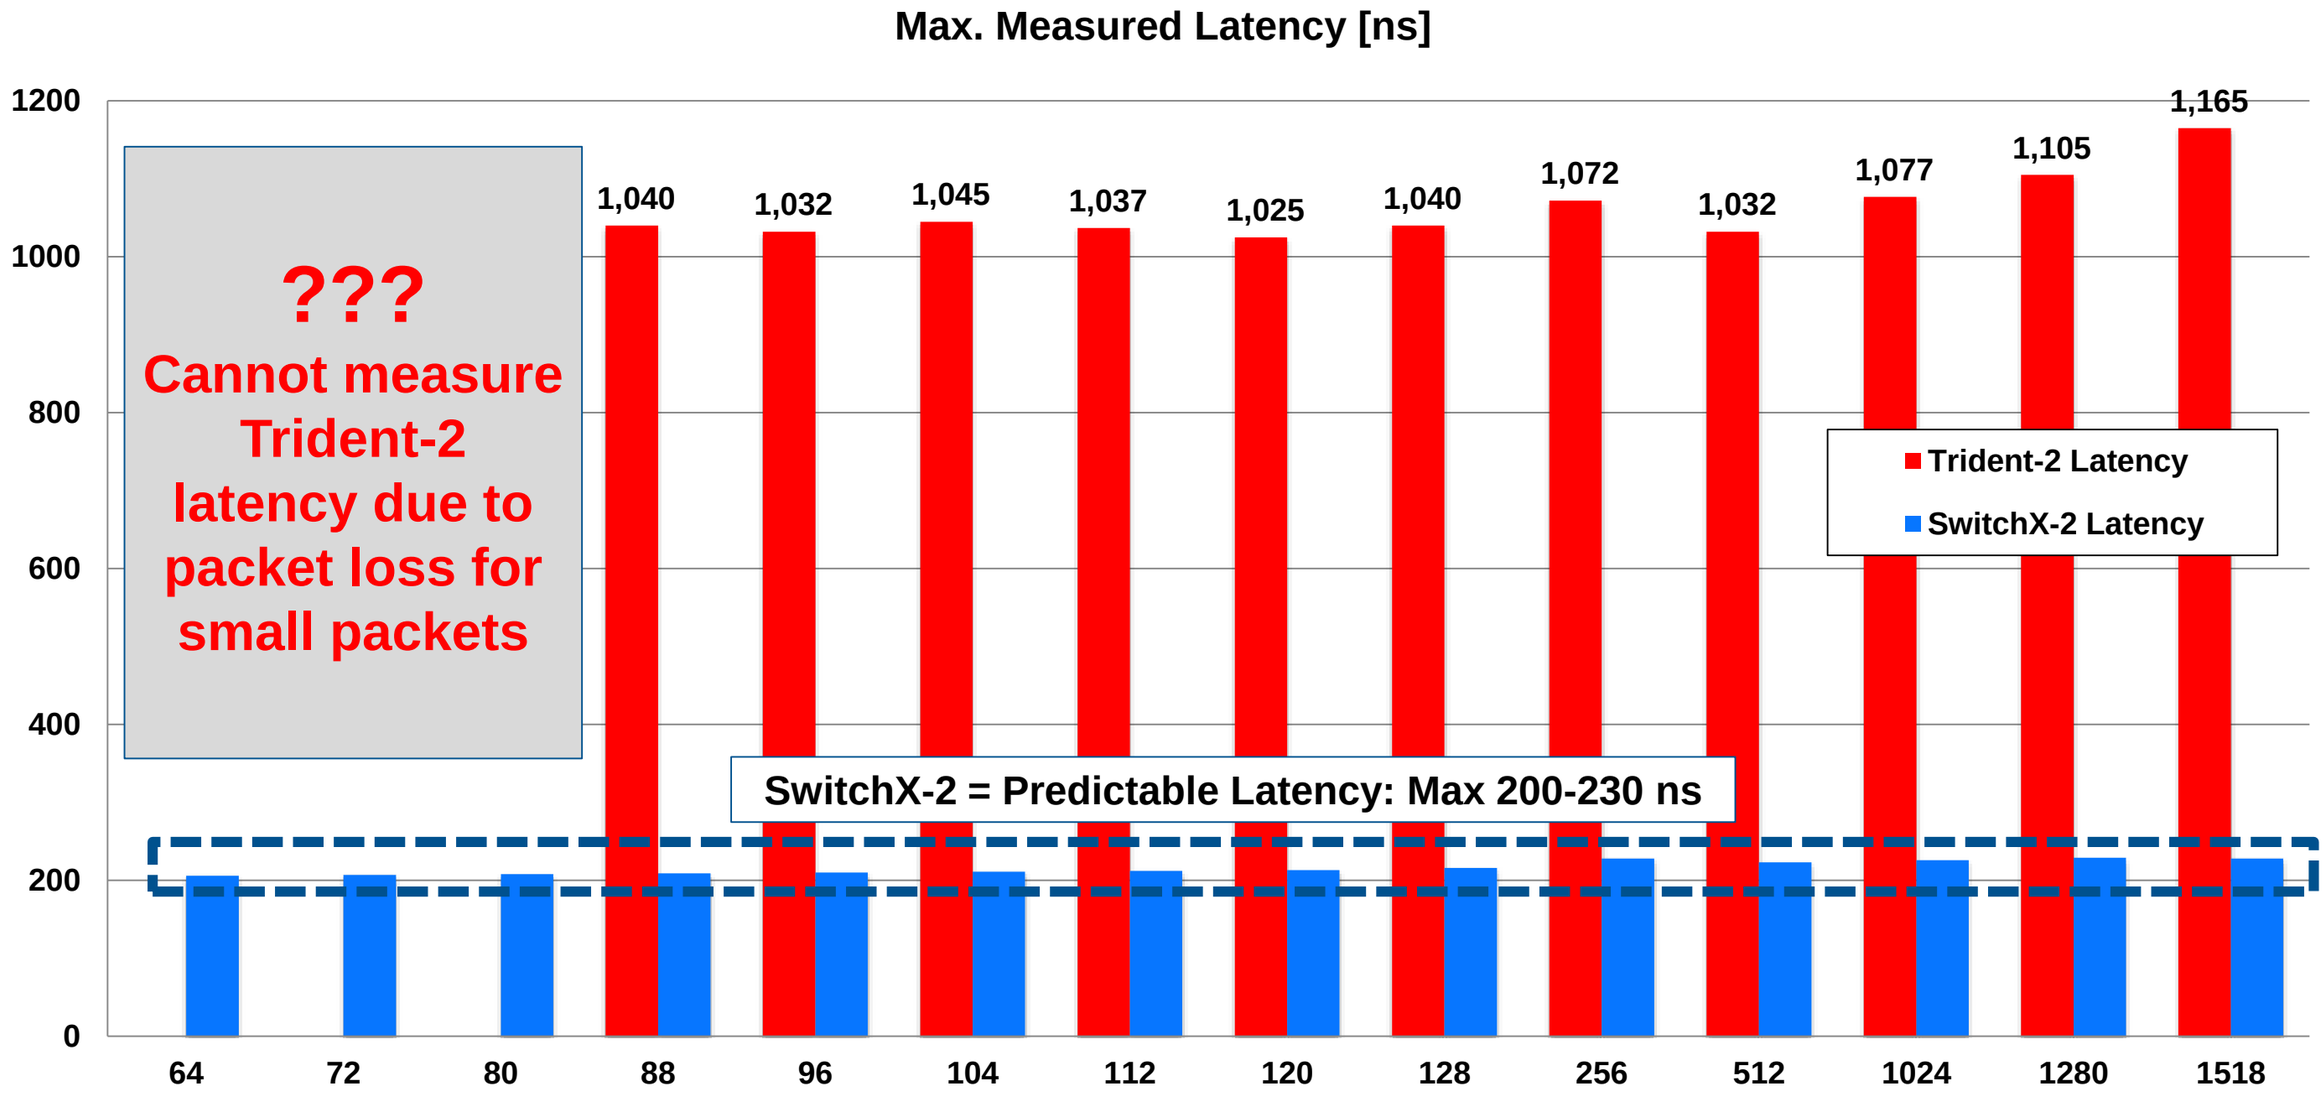

Open Ethernet Choice and Flexibility

Standard Interfaces

- ONIE
- SAI
- Linux

Any NOS/ Any application

- Open and closed source option
- Discrete components or full OS

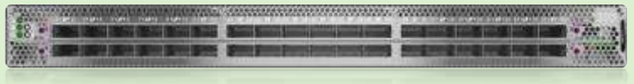

Turnkey / White Box Switch System

OPEN ETHERNET™

OCP activity

OPEN Compute Project

Forward Compatible

- 40GbE now, 100GbE next

SKA Project

- **World's Largest Radio Telescope**
- **64 x SX1012 in 64 Antenna stations connected to the data center via LR4 Optical Transceiver**
- **40 x SX1710 forming 384 ports VMS in the data center**
- **PIM-BiDir with load sharing and OSPF**

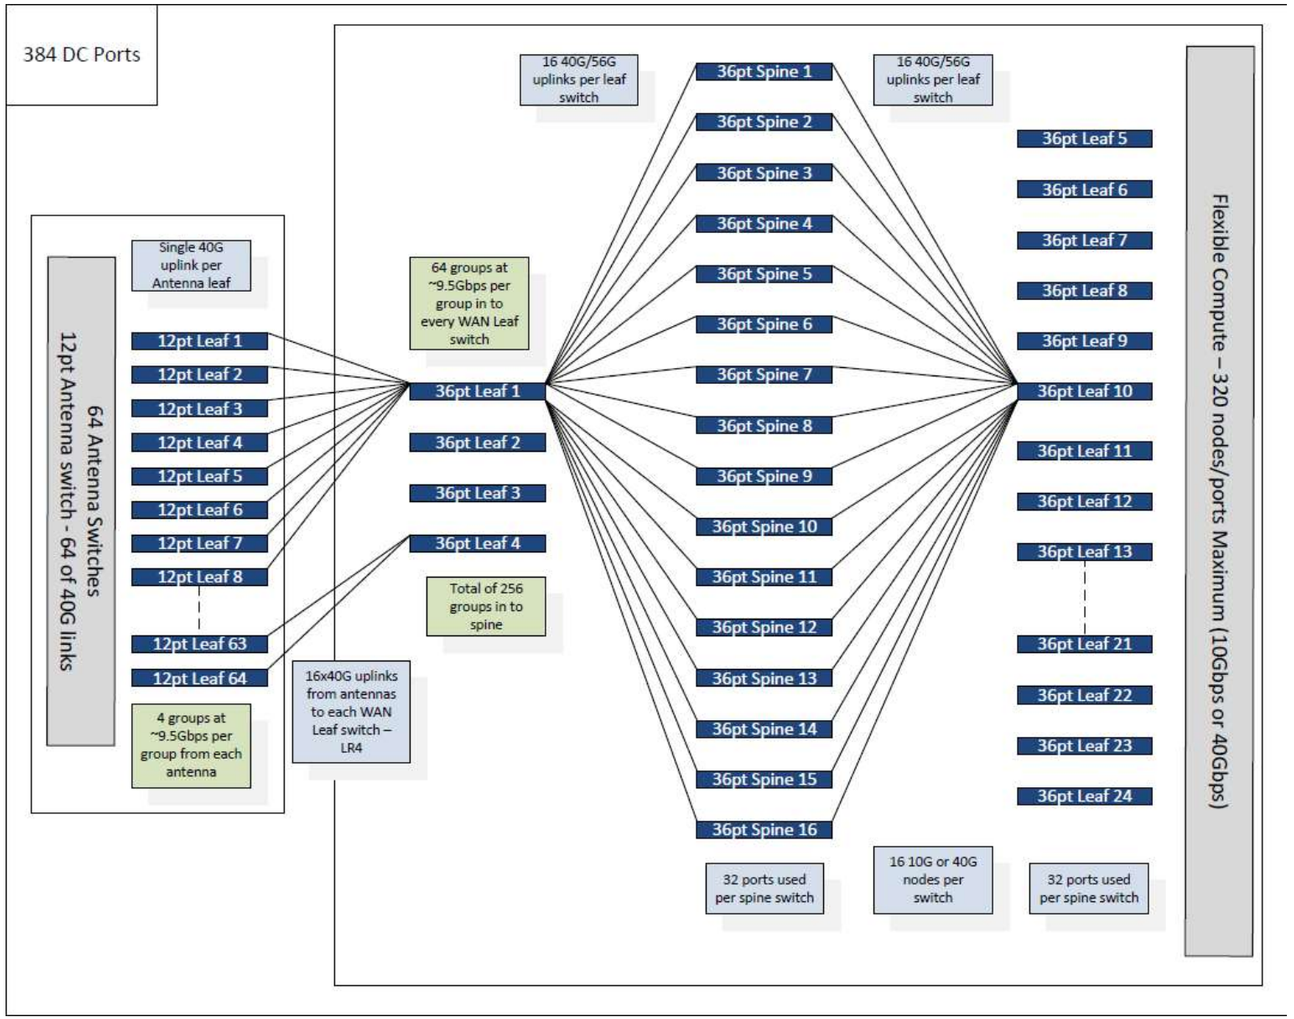

"The amount of data that will be generated by the MeerKAT antennas is equivalent to more than 5 million DVDs every day or about 64 DVDs per sec," said Francois Kapp, sub-system manager at SKA Africa. "This enormous amount of data requires high-bandwidth, low jitter networks. After a comprehensive review of the various products in the market, we concluded that Mellanox's SwitchX-2 Ethernet switches and optical cables will provide us with the best value enterprise-class Ethernet infrastructure that meets all the requirements for our systems."

http://www.mellanox.com/page/press_release_item?id=1521

Oil & Gas Company: Additional Tier Saves $ on the Core

- 1536 nodes cluster in two phases
- VMS based aggregation using VMS wizard
- Dramatic performance improvement

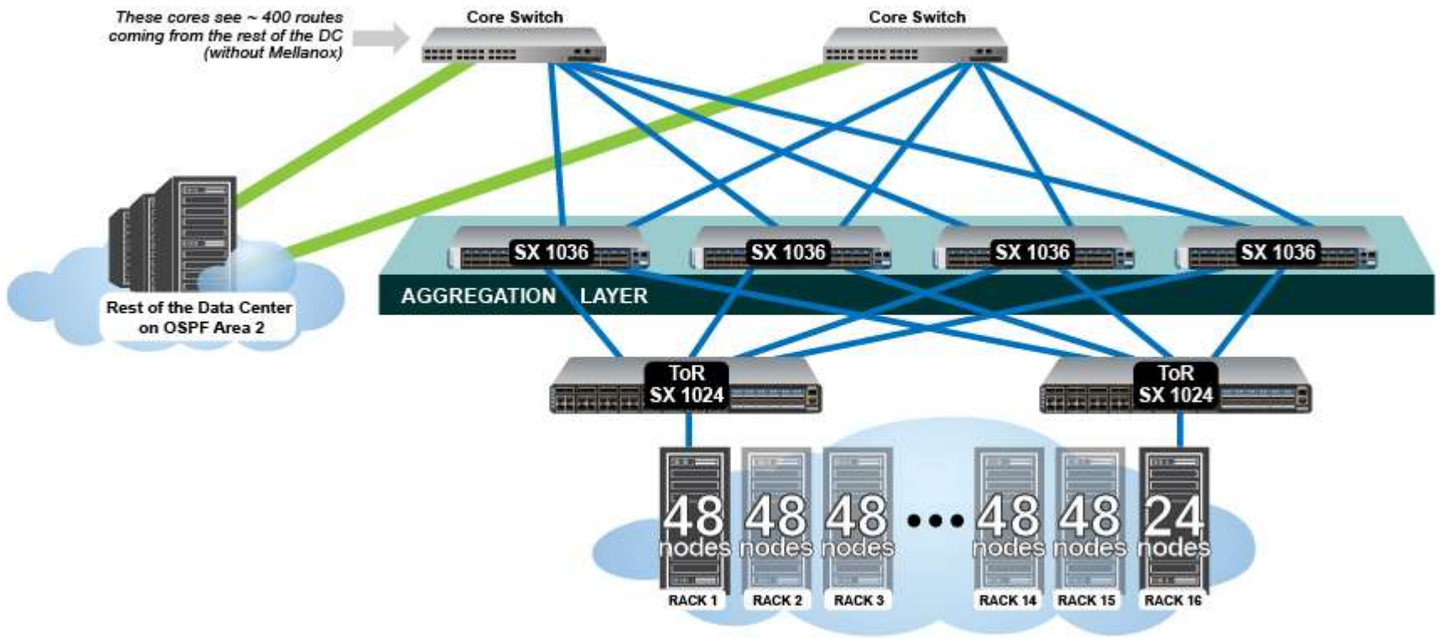

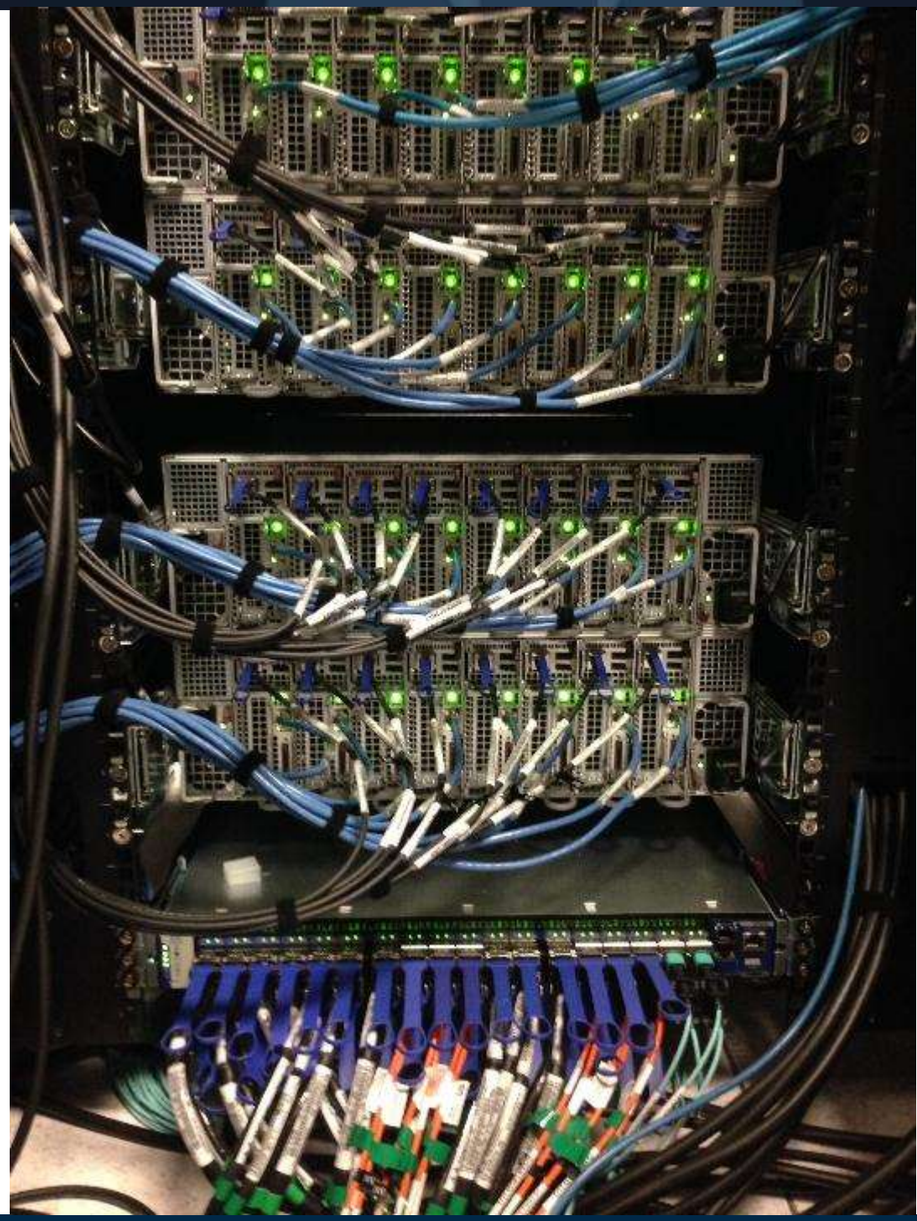

Advanced Topics / Q&A

Thank You

Faster Data Movement Enables Better Data Utilization

Keeping You One Generation Ahead

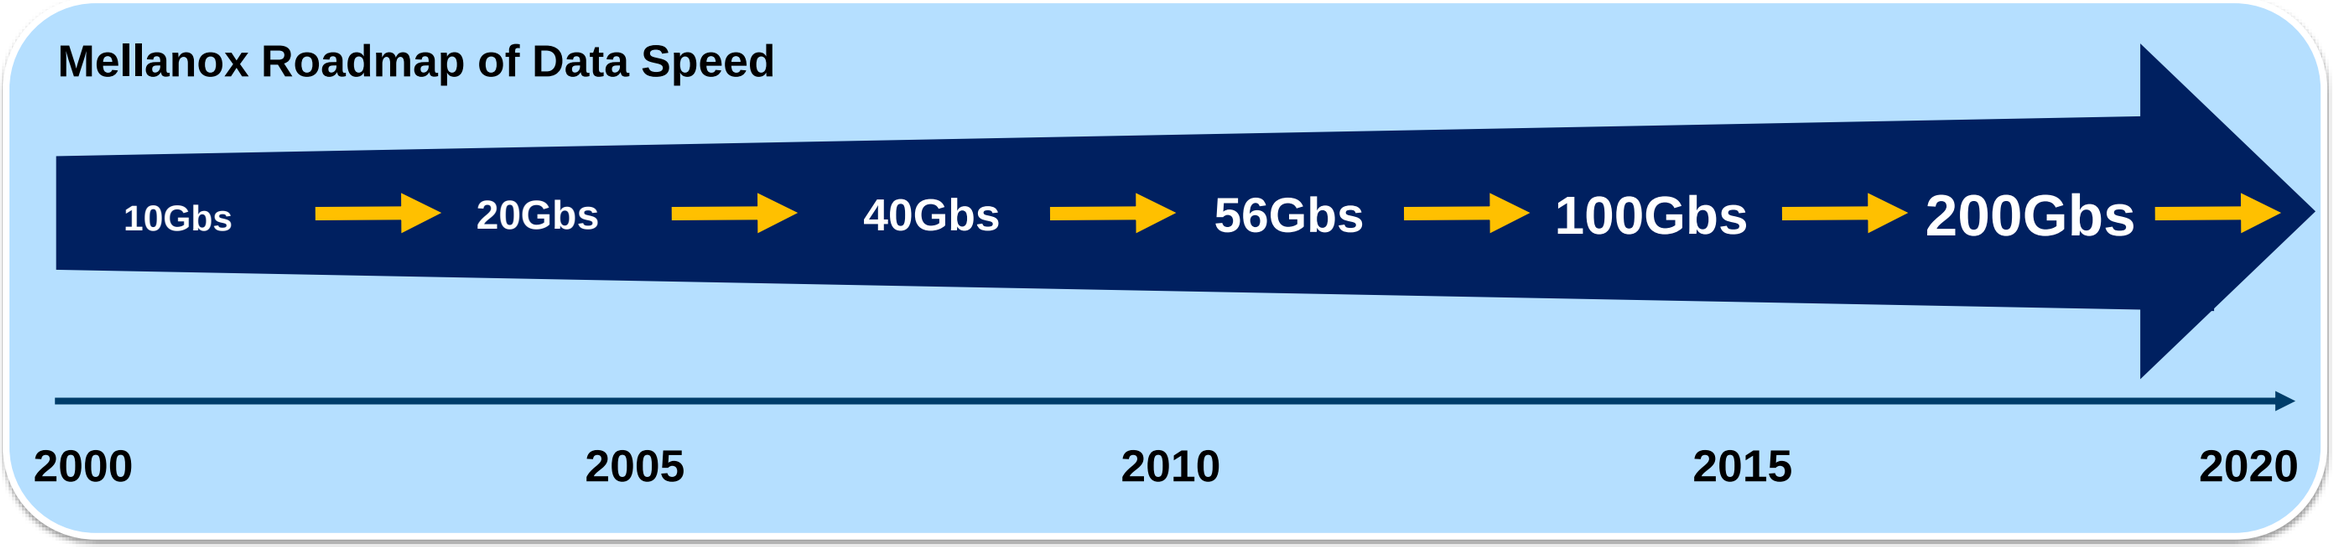

Enabling the Use of Data

Gbs – Gigabit per second