



# Smart Networks and Services workshop

Benny Koren  
Mellanox Technologies


Brussels 4 July, 2019




© 2019 Mellanox Technologies 1




## Mellanox Technologies




SmartNIC




System on a Chip



Adapters




Switches




Cables & Transceivers

**Higher** Data Speeds  
**Faster** Data Processing  
**Better** Data Security



SmartNIC: Smart Network Interface Card

© 2019 Mellanox Technologies 2



## Cloud and Compute next generation disruptions

- Cloud Infrastructure Computing
- Security
- Distributed AI workload with network based acceleration
- Switch optical co-packaging

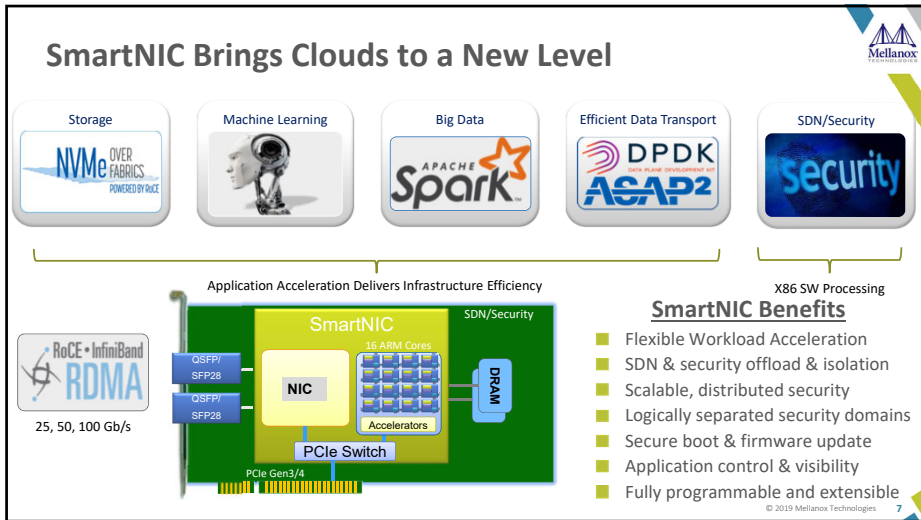
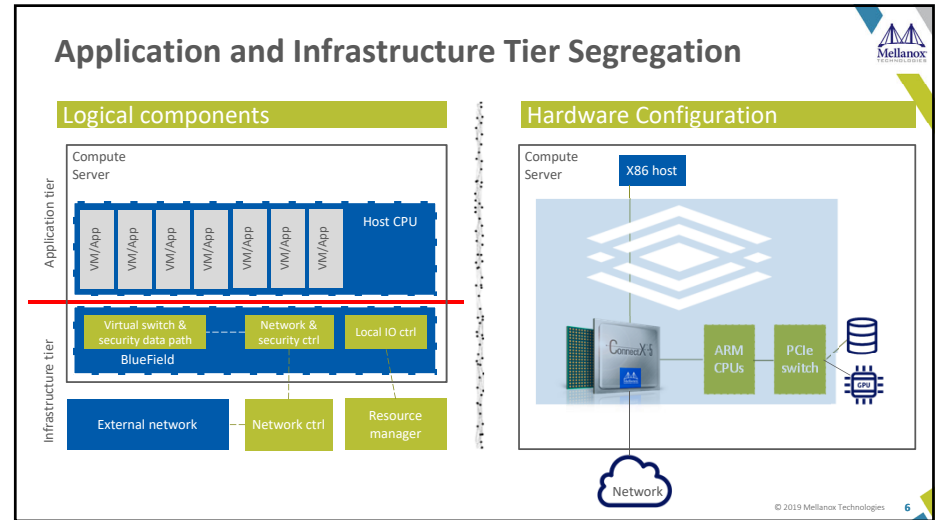
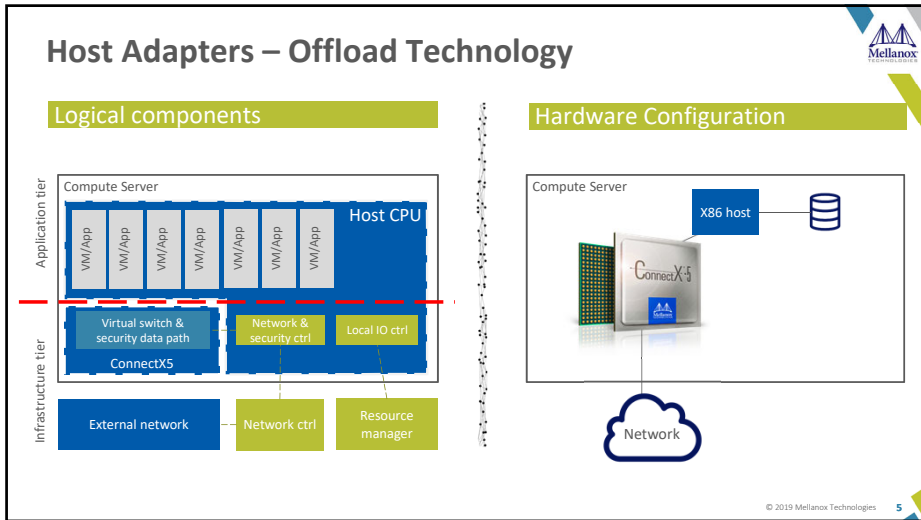
© 2019 Mellanox Technologies 3



## Cloud Infrastructure Computing



© 2019 Mellanox Technologies 4



## Security Challenges in the Data Center

- Traditional perimeter based security model is broken
- Need for privacy drives adoption of cryptography to protect data
- New regulation requirements (e.g. GDPR)
- Zero Trust in the data center!
- Software based security and cryptography is very CPU intensive



## Innovative Security Approach

### Zero Utilization: Transparent encryption accelerations

- Protection of Data-in-Motion and Data-at-Rest
- Host's CPU is fully offloaded from encryption functions
- SmartNIC enables fully Isolated control plane and key management



### Zero Trust: Security controls are built around applications

- Bare-metal, VM and Container complete visibility and line-rate mitigation
- Security infrastructure runs on SmartNIC in an isolated environment
- Resilient SDN micro-segmentation solutions




### Zero Touch: Adaptive, automated Security at the edge of the network

- SmartNIC Implements security applications independent of the host
- Agnostic Solution for different operational environments (bare-metal, legacy, on-prem)
- Enables massive infrastructure scalability



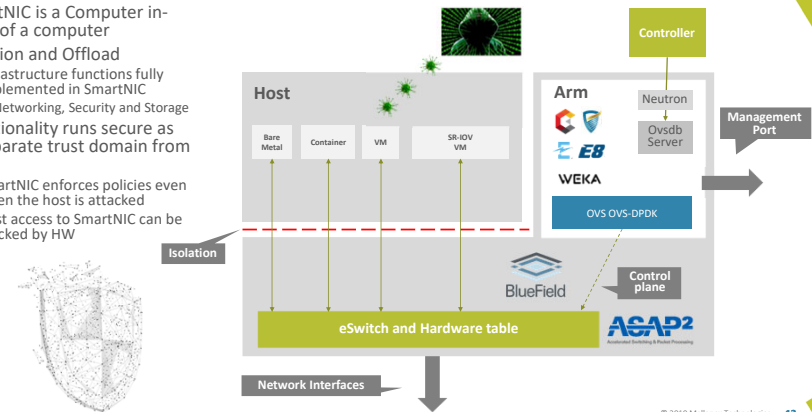
## Storage Security Overview

- Threats to stored information within the data center
  - Theft of storage media
  - Sniffing sensitive data traversing the network
  - Unauthorized access to data from other VMs or Users
- Security functions
  - Data-at-Rest encryption (AES-XTS)
  - Key Management
  - Authentication/Authorization services
- In-line Block-Level Data Encryption
  - Encryption and decryption of data to/from storage
  - Applicable for NVMe-oF / iSER / NFS
  - Protection between users sharing the same storage resource (different keys)
  - Transparent to users, minimal performance impact




## Functional Isolation with SmartNIC

- SmartNIC is a Computer in-front of a computer
- Isolation and Offload
  - Infrastructure functions fully implemented in SmartNIC
    - Networking, Security and Storage
- Functionality runs secure as in separate trust domain from host
  - SmartNIC enforces policies even when the host is attacked
  - Host access to SmartNIC can be blocked by HW



# In Network Computing

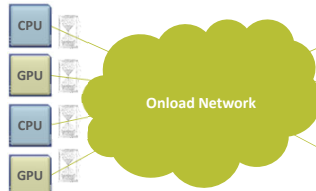


© 2019 Mellanox Technologies 13

## The Need for Intelligent and Faster Interconnect

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

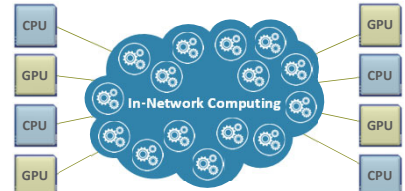
**CPU-Centric (Onload)**



Onload Network



Must Wait for the Data  
Creates Performance Bottlenecks

**Data-Centric (Offload)**



In-Network Computing

Analyze Data as it Moves!  
Higher Performance and Scale

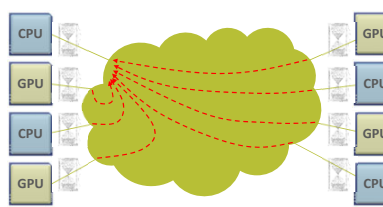



© 2019 Mellanox Technologies 14

## Data Centric Architecture to Overcome Latency Bottlenecks

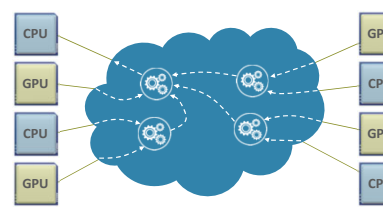
Intelligent Interconnect Paves the Road to Exascale Performance

**CPU-Centric (Onload)**





Communications Latencies of 30-40us

**Data-Centric (Offload)**








Communications Latencies of 3-4us

© 2019 Mellanox Technologies 15



## Accelerating All Levels of HPC / AI Frameworks

<b>Application</b>	<ul style="list-style-type: none"> <li>Data Analysis</li> <li>Real Time</li> <li>Deep Learning</li> </ul> 
<b>Communication</b>	<ul style="list-style-type: none"> <li>Mellanox SHARP In-Network Computing</li> <li>MPI Tag Matching</li> <li>MPI Rendezvous</li> <li>Software Defined Virtual Devices</li> </ul> 
<b>Network</b>	<ul style="list-style-type: none"> <li>Network Transport Offload</li> <li>RDMA and GPU-Direct RDMA</li> <li>SHIELD (Self-Healing Network)</li> <li>Enhanced Adaptive Routing and Congestion Control</li> </ul> 
<b>Connectivity</b>	<ul style="list-style-type: none"> <li>Multi-Host Technology</li> <li>Socket-Direct Technology</li> <li>Enhanced Topologies</li> </ul> 



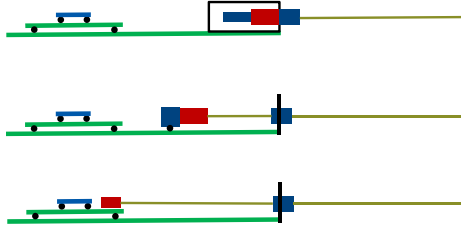
© 2019 Mellanox Technologies 16

# Optical co packaging




© 2019 Mellanox Technologies 17

## Where is the Optics ?



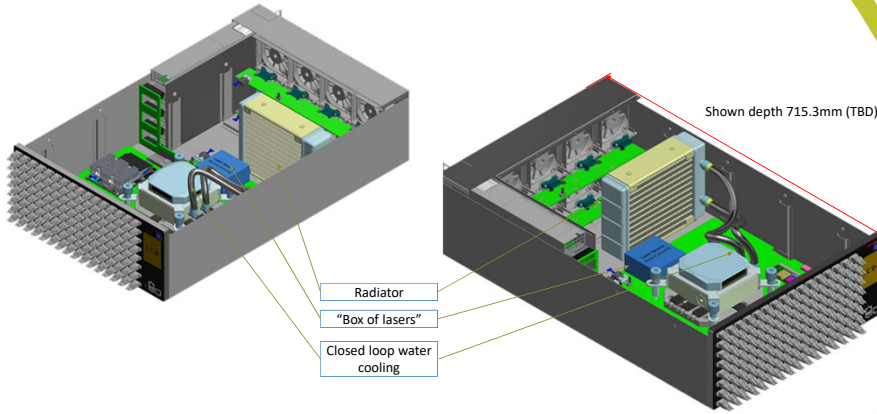
- Pluggable Optics
- On Board Optics ( COBO )
- Co-Packaged optics

Co-packed solution enables lower system power ( 30%-45% )




© 2019 Mellanox Technologies 18

## System overview



Shown depth 715.3mm (TBD)

- Radiator
- "Box of lasers"
- Closed loop water cooling



© 2019 Mellanox Technologies 19