



# Mellanox Update @ STAC NYC 2019

High Performance Datacenter Networking



October 2019

 Lior Paster  
 [LiorP@mellanox.com](mailto:LiorP@mellanox.com)  
+1-646-320-2772

# Mellanox is the Leading Supplier of Standard, Open, Intelligent End-to-End Interconnect Solutions for the datacenter

10/25/40/50/100 and soon 200/400 GbE Ethernet

EDR (100Gbps) and HDR (200Gbps) InfiniBand



# Company Overview

**\$1.09B**  
2018 Revenue

**\$265M**  
2018 Cash From Operations

**\$439M**  
2018 Ending Cash & Investments

**~2,800**  
Employees worldwide



Company HQ:

- Israel (Yokneam)
- USA (California)
- Worldwide Offices

# Mellanox acquisition by nVidia: Update

- We are not nVidia yet...
- Mellanox acquisition by nVidia is planned to be finalized this quarter
  - Pending Chinese regulatory approvals
- Combined company to lead the market in
  - High performance datacenter networking and processing of data
  - Best of breed AI and Deep Learning solutions
  - Efficient data movement for compute & storage

# Mellanox in Numbers





# 1

Shipping 1 Million  
Ethernet NIC's every  
Quarter



# 20

Mellanox Celebrating  
20 Years in business

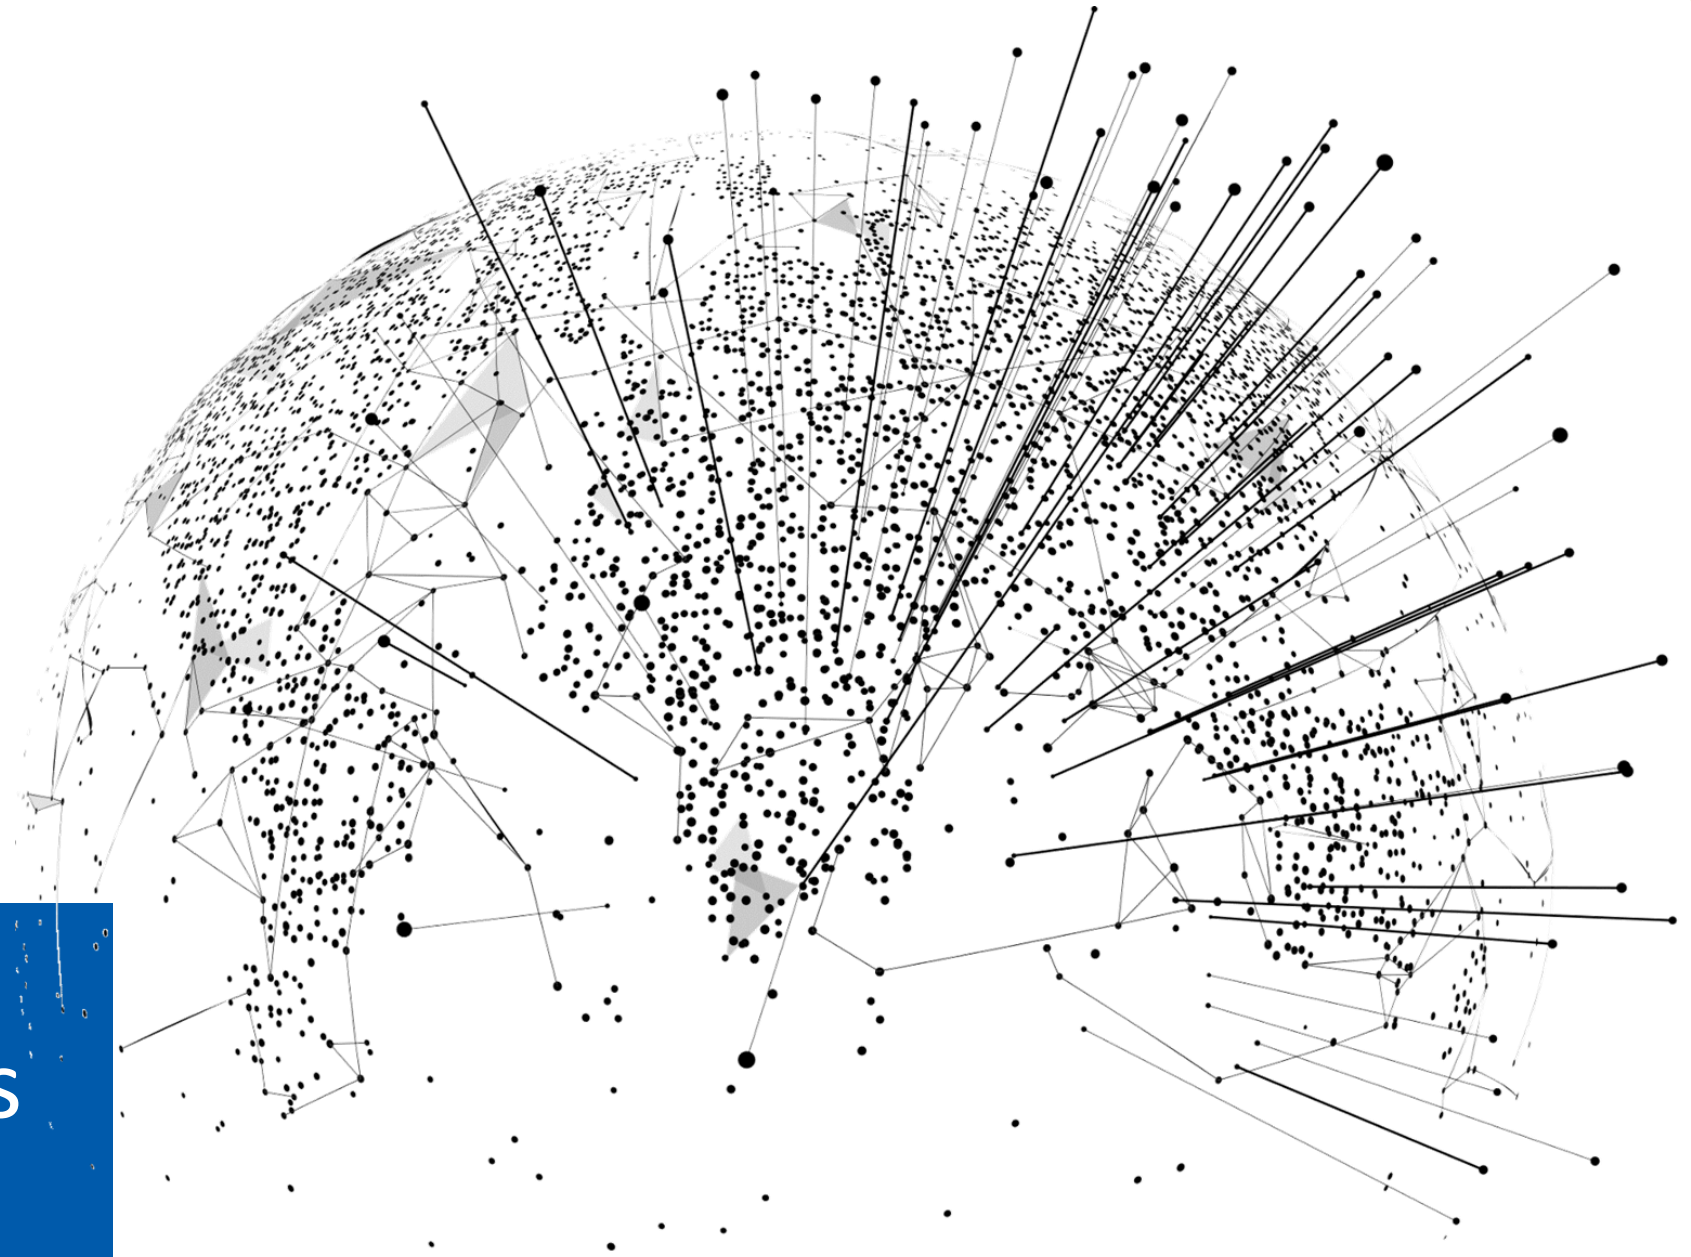


Ticker: MLNX



# 2

NYC-based Stock Exchanges  
use Mellanox





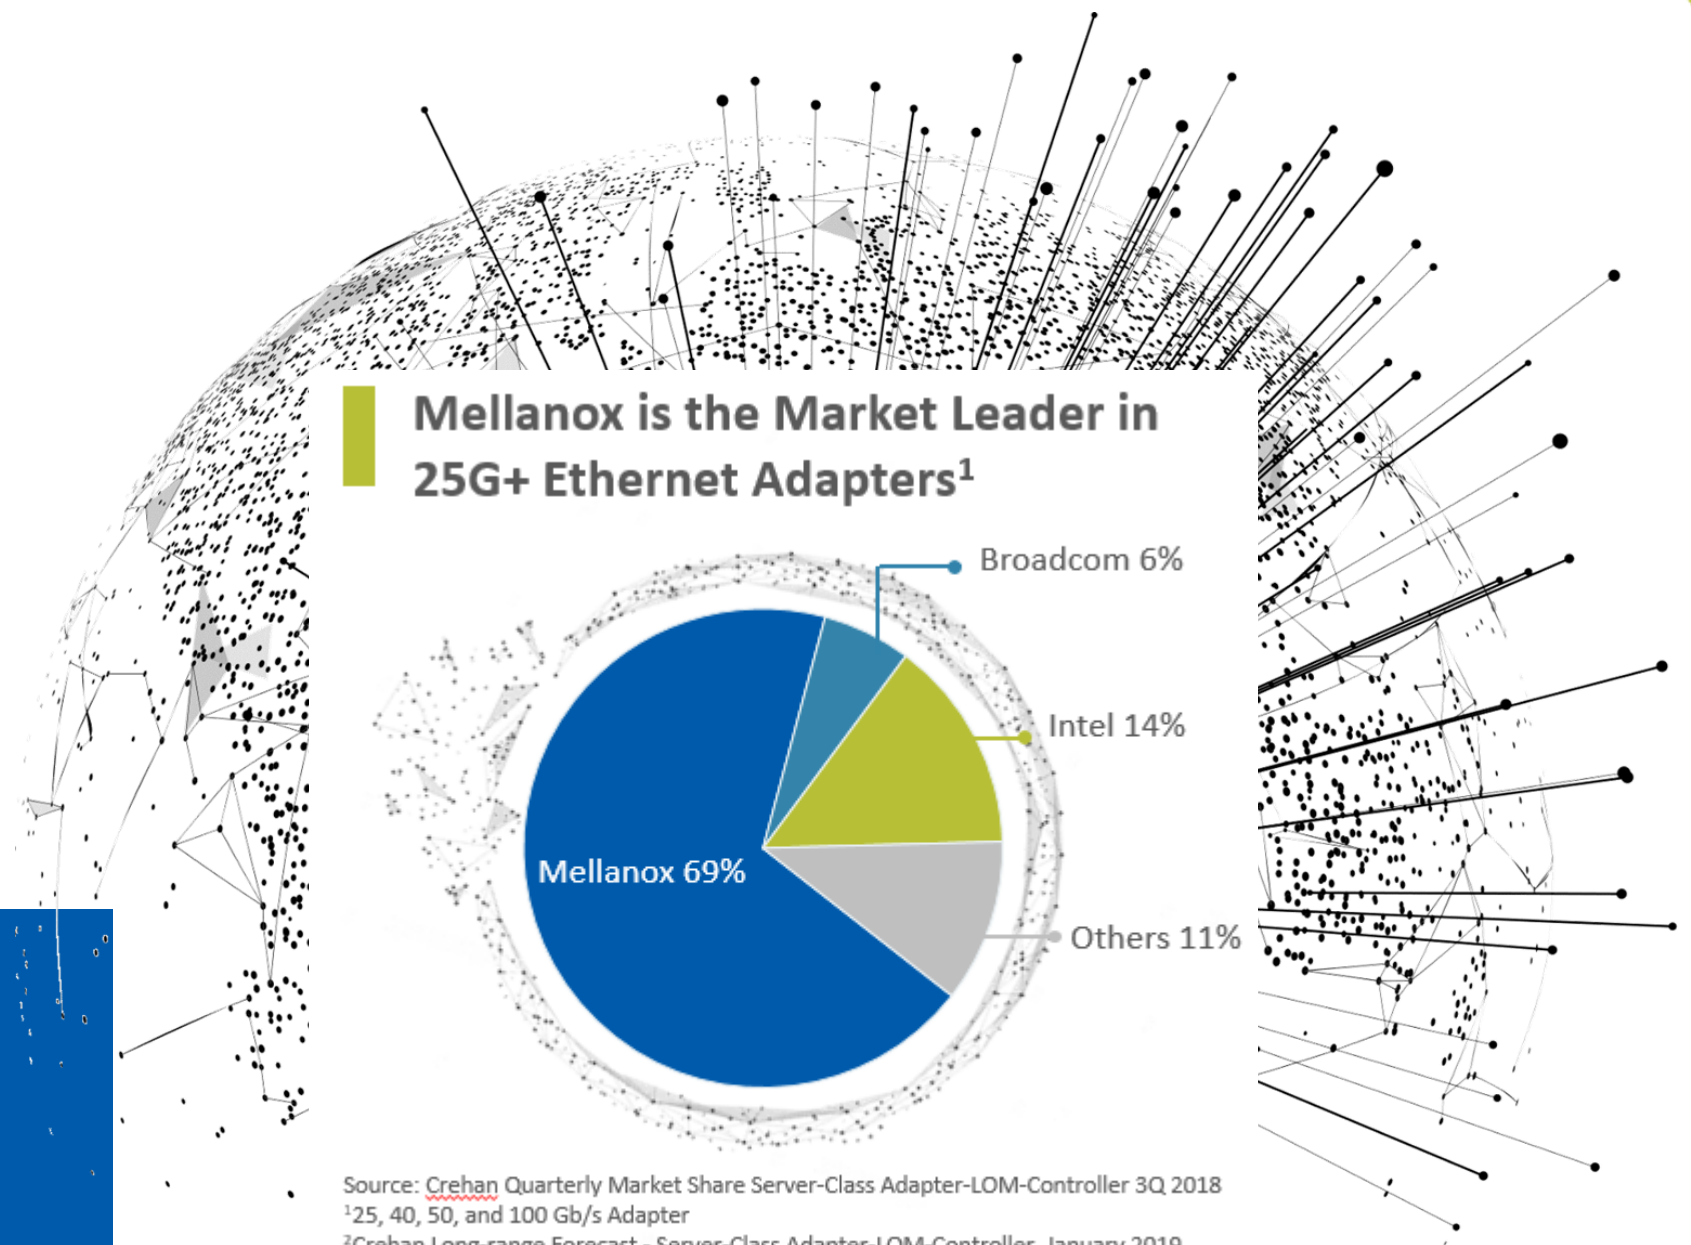
# 3

NYC-based Cloud Providers  
use Mellanox



# 69

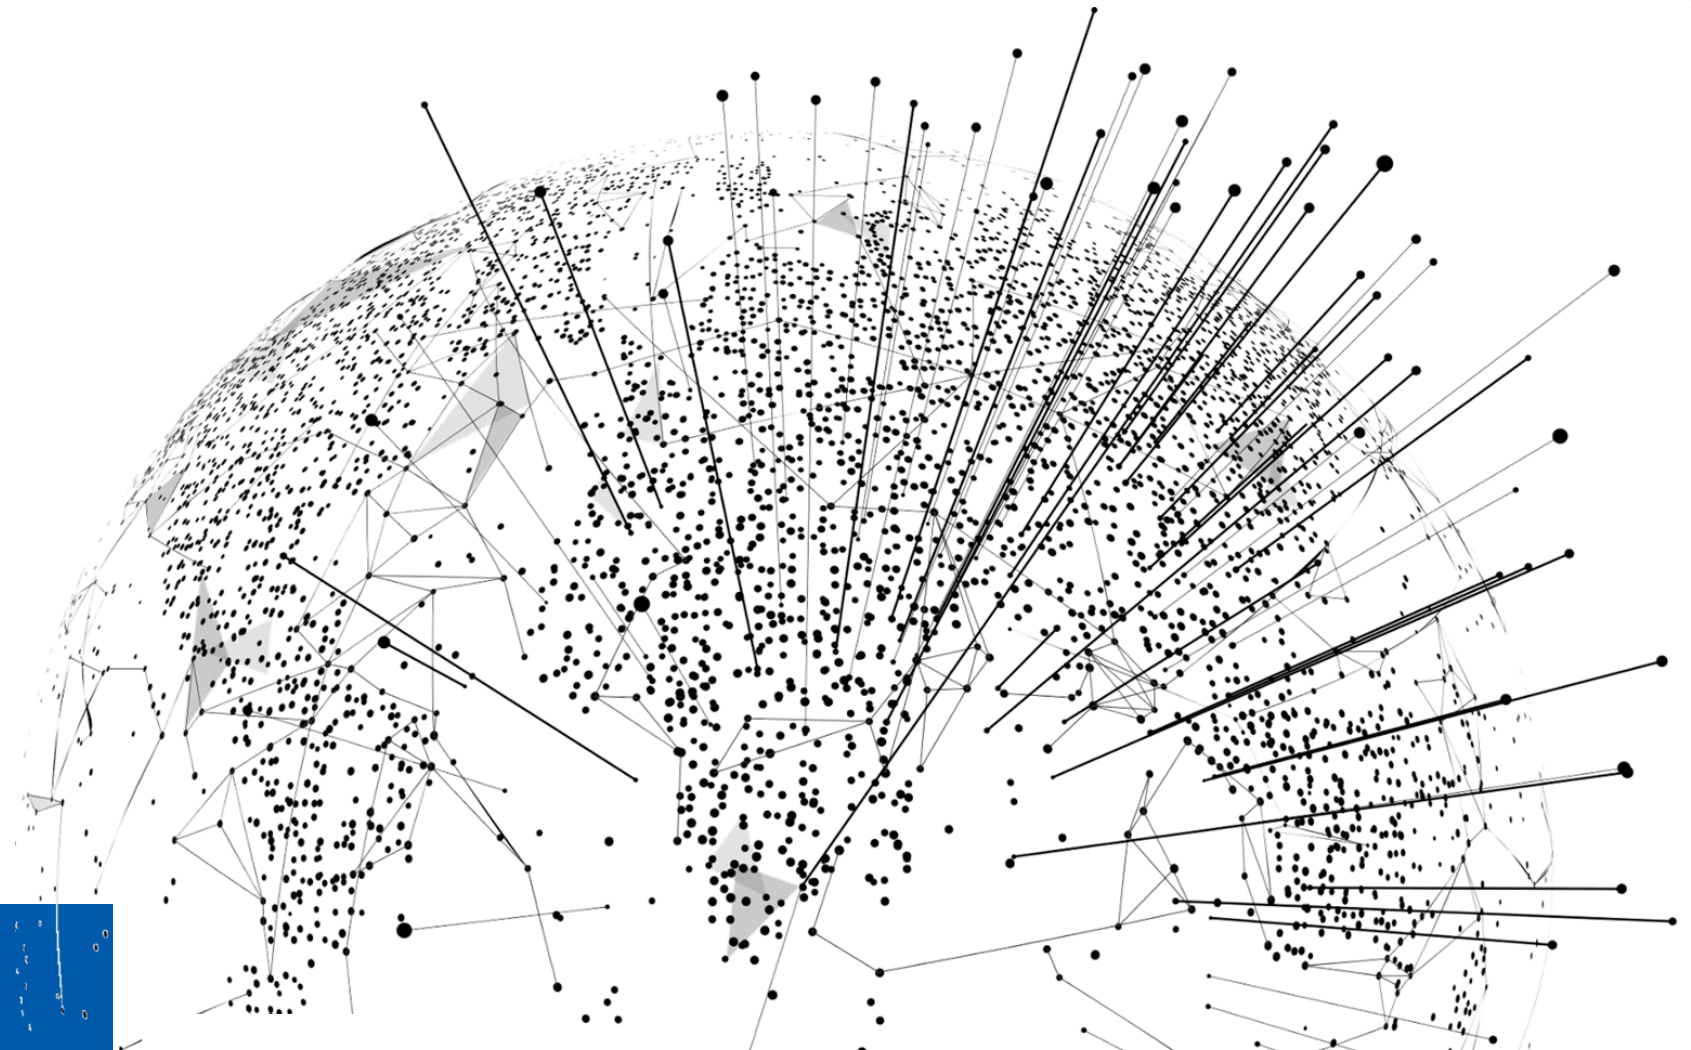
69% of NIC's sold in USA (25GbE and above) are Mellanox



Source: Crehan Quarterly Market Share Server-Class Adapter-LOM-Controller 3Q 2018  
<sup>1</sup>25, 40, 50, and 100 Gb/s Adapter  
<sup>2</sup>Crehan Long-range Forecast - Server-Class Adapter-LOM-Controller, January 2019


# 5

Top Machine Learning  
Frameworks support  
RDMA/RoCE



Caffe

 Microsoft

 PaddlePaddle

  
TensorFlow

 NVIDIA



# 3

Among TOP 3 contributor to  
Linux Kernel

Source: <https://lwn.net/Articles/798505/>





# 6

# DX

A complex network visualization consisting of numerous black dots (nodes) connected by thin black lines (edges). The nodes are distributed across the right side of the slide, with a higher density in the center and right, creating a sense of depth and connectivity. The overall shape is roughly spherical or hemispherical.

Delivers Line-rate IPSEC and  
TLS encrypt/decrypt  
Acceleration in ASIC

# NIC's



# Mellanox ConnectX-5 Ex



- High Performance & commodity NIC “in one”
  - HFT acceleration (kernel bypass, verbs)
  - Cloud acceleration (Vxlan, OVS, DPDK acceleration in hardware)
  - Storage acceleration (NVMeF RoCE, NVMeoTCP, storage encryption and erasure coding offload in hardware)
  
- What do you get with ConnectX-5 ?
  1. VMA Kernel Bypass: Low Latency, Open Source, use in many places in production
  2. RoCE/RDMA/Verbs for lower latency market data and messaging
  3. High Performance Capture (P2D)
  4. PTP (TimeKeeper, LinuxPTP)
  
- EASY to use: Linux, VMware, Windows Inbox Drivers
- 1-line to run VMA with pre-tuned profiles...
- “Upstream First” methodology = bug-free Linux Inbox Drivers, rigorous community validation

# Mellanox VMA Kernel Bypass


## KERNEL BYPASS

Reduce Kernel overhead with direct network adapter access

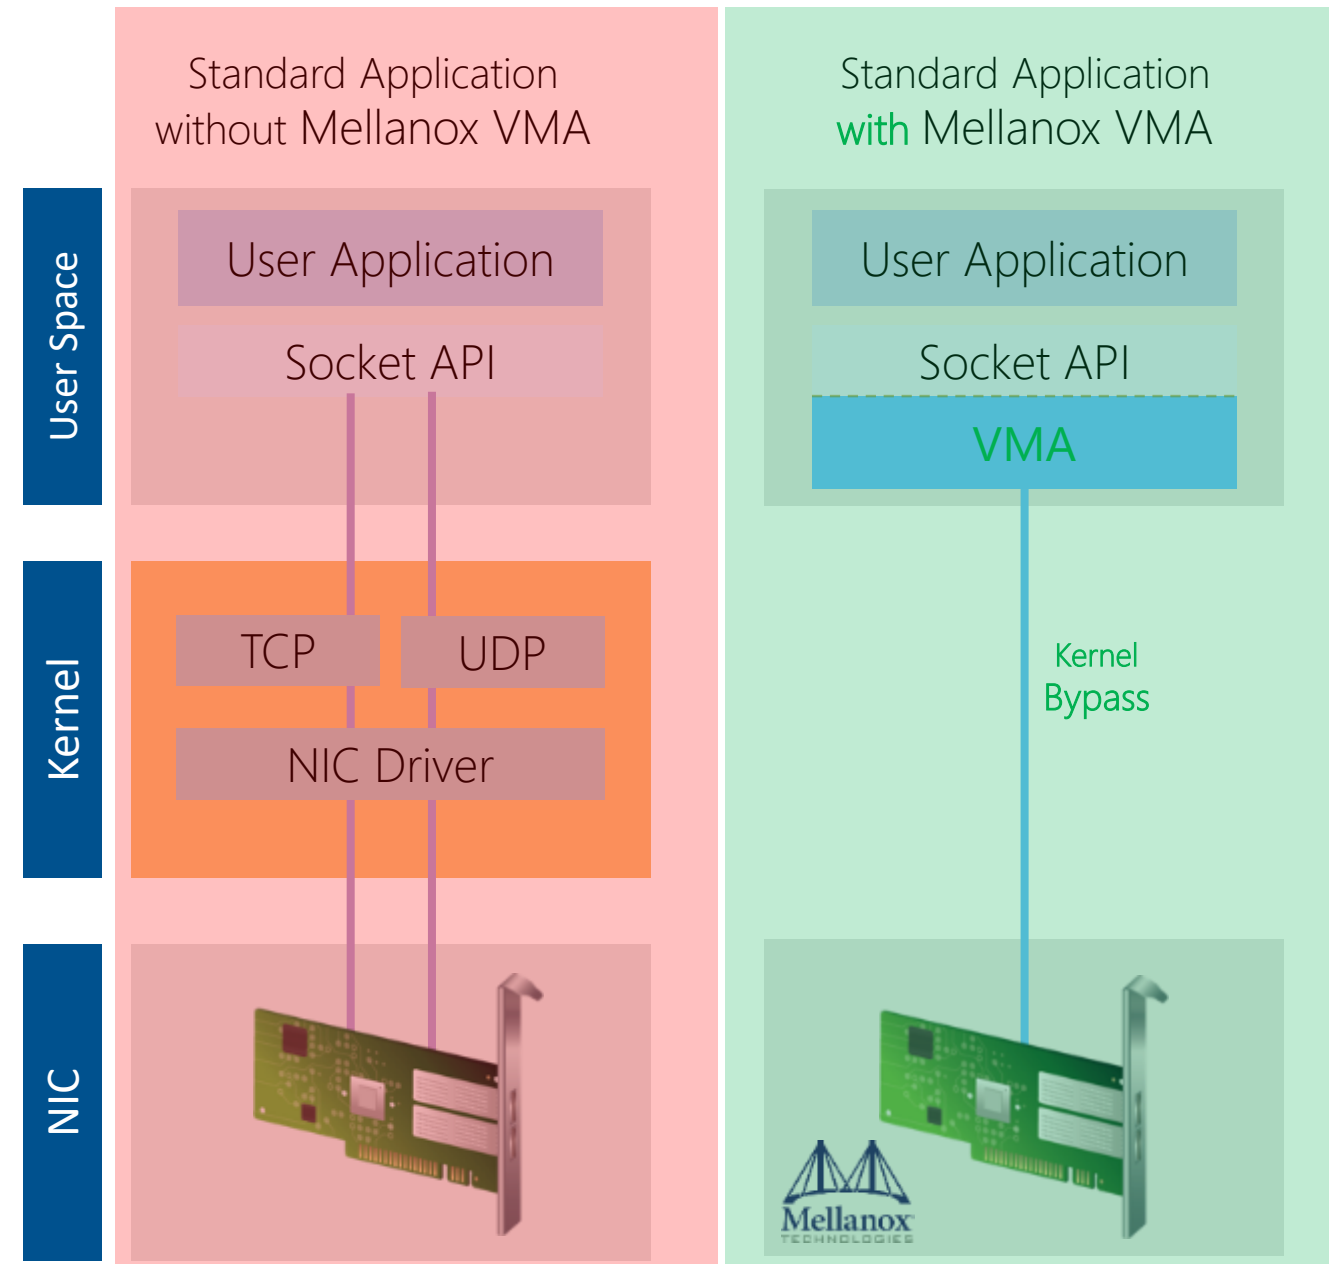


## SINGLE SIDED

no application changes –  
Standard sockets TCP, UDP  
(Unicast, Multicast)



[https://www.mellanox.com/related-docs/applications/SB\\_HighFreq\\_Trading.pdf](https://www.mellanox.com/related-docs/applications/SB_HighFreq_Trading.pdf)





# BlueField SmartNIC



Wouldn't it be nice to have ConnectX-5 goodies plus **compute cores** on a NIC?

Lots of use cases come to mind...

## What's in it for me?

- Host Isolation → full feature, low latency NIC, running Linux, as a true firewall on every host
  - Enterprise class agent-less isolation & micro segmentation with Guardicore
- **NVMe SNAP: simple NVMe storage disaggregation** for Linux, Windows, VMware

## Call to action

- Check out Guardicore, our security partner
- Talk to us, and Pavilion and Pure on storage
- Test BlueField today – ASK ME HOW

# Ethernet Switch



# Mellanox Spectrum 100GbE Ethernet Switch

- Spectrum performance is proved in the [Tolly Report](#)
- Smart, Fully Shared Buffer, see [Tolly Report](#)
- In-ASIC Telemetry, Open API's to export to Cassandra, Grafana, Kibana..., for deep monitoring
- Network Virtualization (EVPN VXLAN) best-in-class scale
- Truly OPEN: supporting Cumulus, Mellanox OS, SONiC

# HFT-Specific Feature in Spectrum 1/2 Switch

- Spectrum(1) leads in performance with 300 nsec p2p latency @ 100Gig (Tolly Report)

## Spectrum-2 key new features:

- Hardware NAT at huge scale AND low latency (Q1 - Cumulus)
- 128 ports 10 and 25 gig on single ASIC
- Comprehensive L2 and L3 support, multicast
- Increased scale

*Not STAC Benchmarks*



# EDR and HDR Infiniband Update

- Half of Mellanox Business is Infiniband, growing FAST in AI/ML/DL
  
- EDR IB
  - **Lowest E2E latency, with 90 nanos in the switch and below 700 nanos in the NIC**
  - 100Gb/s rate
  - Flat network scales to 1000's of nodes
  - Self managing, self healing, “set up and go”
  - New telemetry and monitoring capabilities
  
- HDR IB
  - All the EDR goodies, plus:
  - **Support for 200Gb/s per port**
  - Supports 40x 200Gig or **80x 100Gig ports on a 1U switch**, single ASIC
  
- NDR is coming...

# Let Mellanox help you innovate

Lior Paster

[LiorP@mellanox.com](mailto:LiorP@mellanox.com)

+1-646-320-2772

