



Interconnect Your Future with Mellanox

Paving the Road to Exascale

Forum Teratec - June 2017

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

Company Headquarters

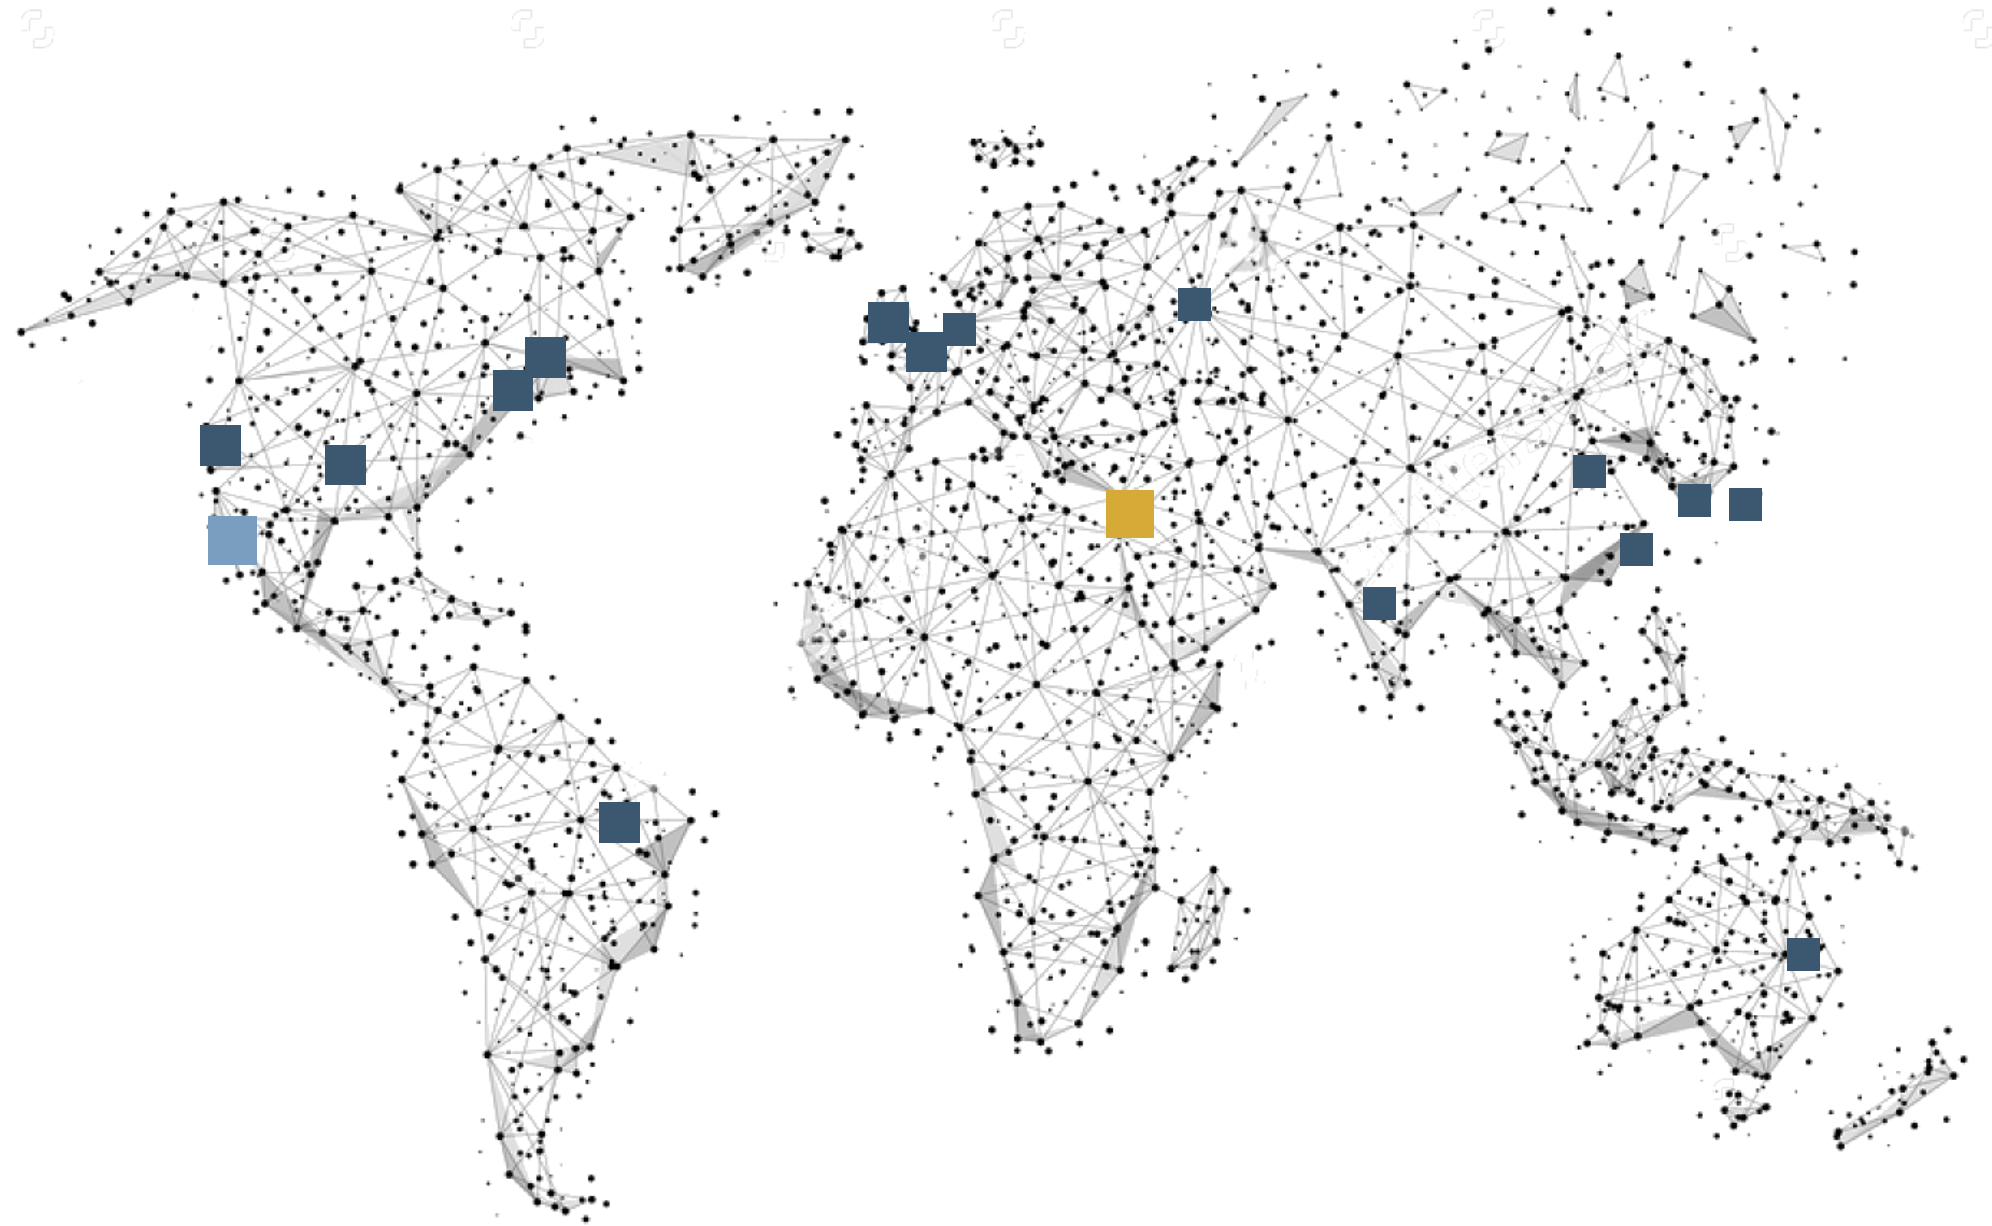
- Yokneam, Israel
- Sunnyvale, California
- Worldwide Offices

~2,900


Employees worldwide



Ticker: MLNX





Approaching \$1 Billion in Annual Revenues



Historical Annual Revenue (\$M)


Source: Company financial statements

Unleashing the Power of all Compute Architectures




Highest Performance & Scalability for
Power, x86, GPU, ARM & FPGA-based Compute and Storage Platforms
10, 20, 25, 40, 50, 100 & 200Gb/s Speeds


CPU-Centric (Onload)



Data-Centric (Offload)



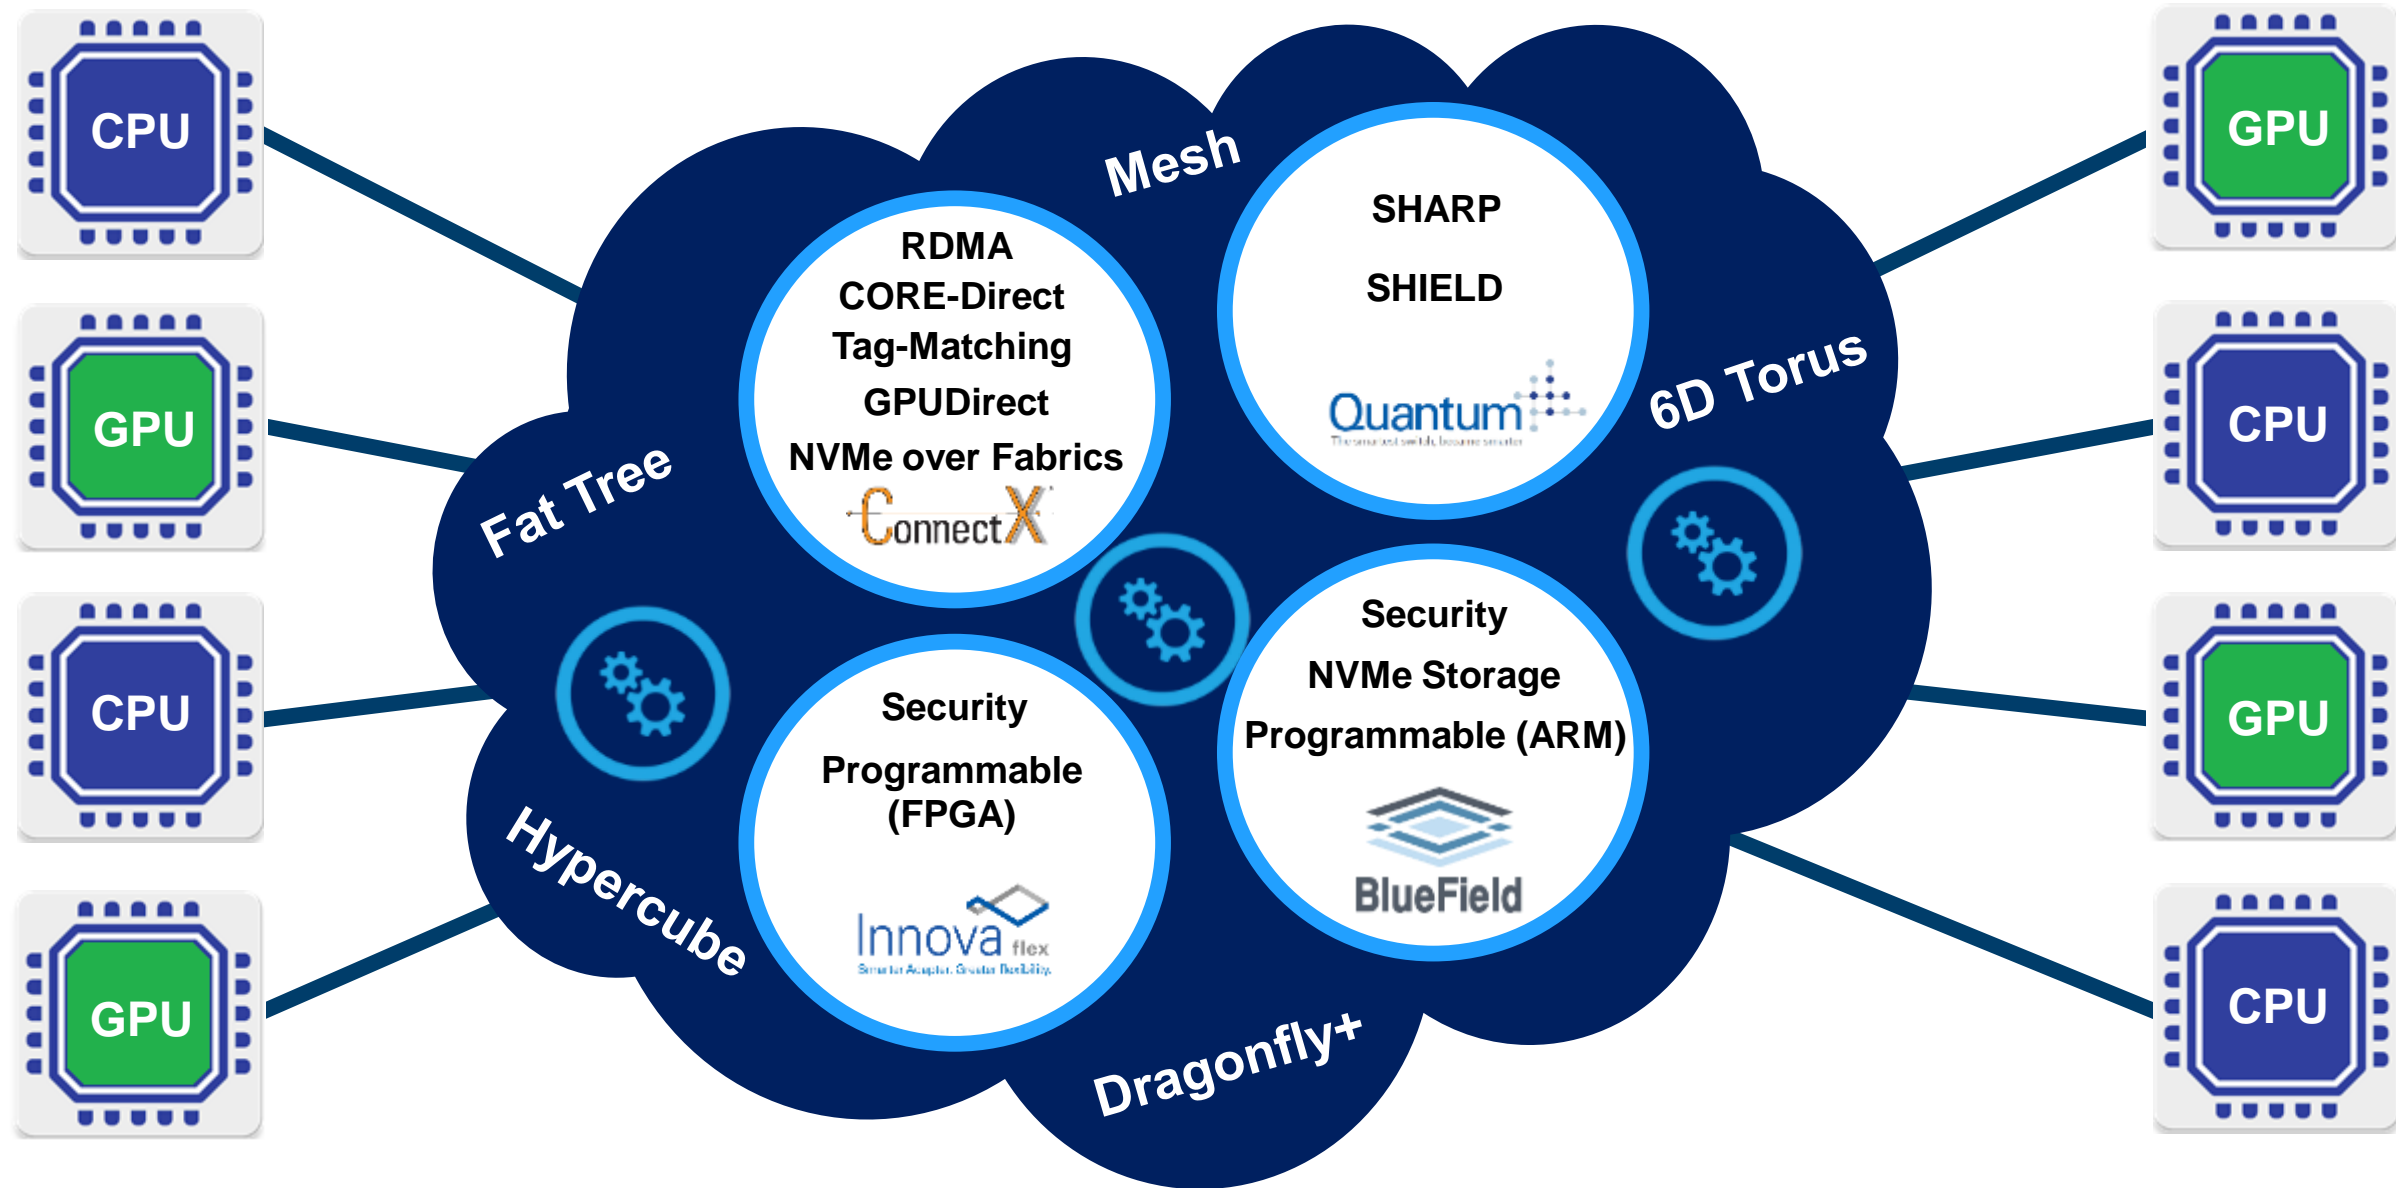
**Must Wait for the Data
Creates Performance Bottlenecks**



Analyze Data as it Moves!

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

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



In-Network Computing Key for Highest Return on Investment

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

Enabling the Future of Machine Learning Applications



HPC and Machine Learning Share Same Interconnect Needs

Enabling Most Efficient Machine Learning Platforms

■ OCP Big Sur Artificial Intelligence Platform



■ Real Time Fraud Detection



■ Machine Learning System with 400Gb/s



■ 18X Speedup For Image Recognition




■ 4X Speedup For Image Recognition



■ Data Analytics Image Recognition



■ World Record For Data Sort, 3X Faster



Allowing Machine learning to Perform Critical & Real Time Decisions



NVIDIA AI Supercomputer

#31 on the TOP500 List



Facebook AI Supercomputer

#32 on the TOP500 List



- EDR InfiniBand In-Network Computing technology key for scalable Deep Learning systems
- RDMA accelerates Deep Learning performance by 2X, becomes de-facto solution for AI

**Maximum Performance
and Highest
Return on Investment**



- **Higher** performance
- **Higher** efficiency
- **Higher** scalability
- **Higher** speed

**World's First
200G Switch**



**World's First
200G Adapter**



Thank You