SDC 19

September 23-26, 2019
Santa Clara, CA

Hardware Security for Performance Based TCP Attached Storage

Scott Schweitzer
Solarflare a **Xilinx** Company

30-Seconds of Background

- 17 years with IBM, first 11 at Research

- 3 years with NEC as Itanium Product Mgr

- 8 years with Myricom as HPC/HFT Sales/Mktg
- 6 years with Solarflare as Sales/Mktg

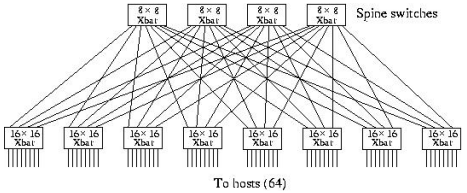

- Now Xilinx outbound DCG Mktg

DATA
CENTER GROUP

Software Based Acceleration

Our Increasing Appetite for Data

- Since the start of this century, digital storage has grown exponentially
- "Big Data" apps now consume petabytes of data at once
  - Data Mining/Warehousing
  - Real-Time Analytics
  - Artificial Intelligence
  - Business Intelligence

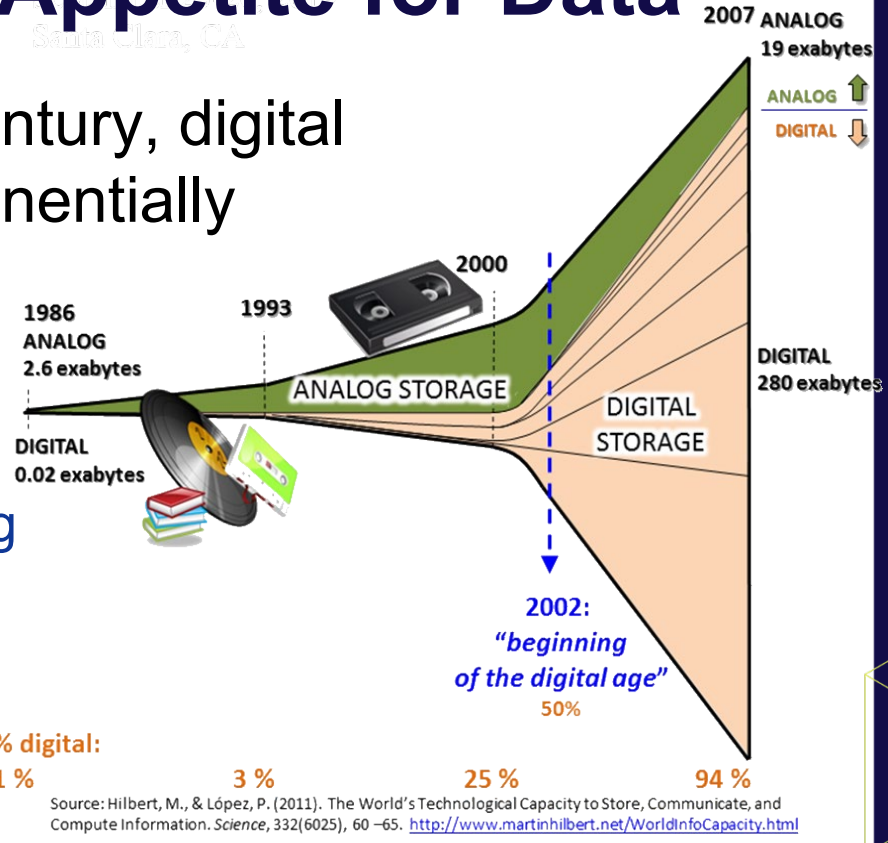

Source: Hilbert, M., & López, P. (2011). The World's Technological Capacity to Store, Communicate, and Compute Information. *Science*, 332(6025), 60 –65. http://www.martinhilbert.net/WorldInfoCapacity.html

Kernel I/O *Slows Down* Data

- Big Data ***requires*** fast I/O solutions
- Normal I/O solutions run all I/O traffic through the OS Kernel
  - Memory copies
  - Context switch – 1400ns Intel E5520*
  - Lag and increased CPU usage
- While great for managing "uncontrolled" environments, this approach impacts performance of cloud environments

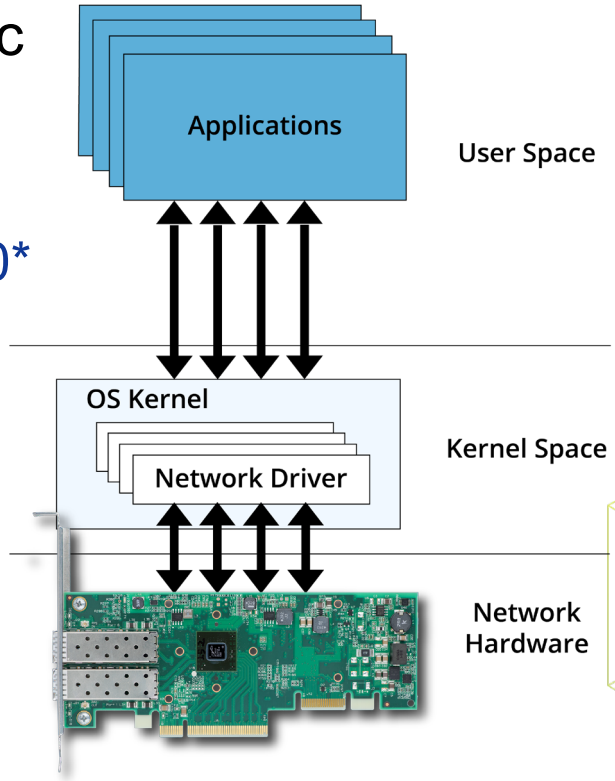

First There Was User Space Acceleration

- By bypassing the kernel, user space I/O solutions overcome this issue
  - No context switching
  - No memory copies
- User space I/O increases bandwidth while decreasing CPU utilization
  - Improved CapEx and OpEx
  - Better solution scalability

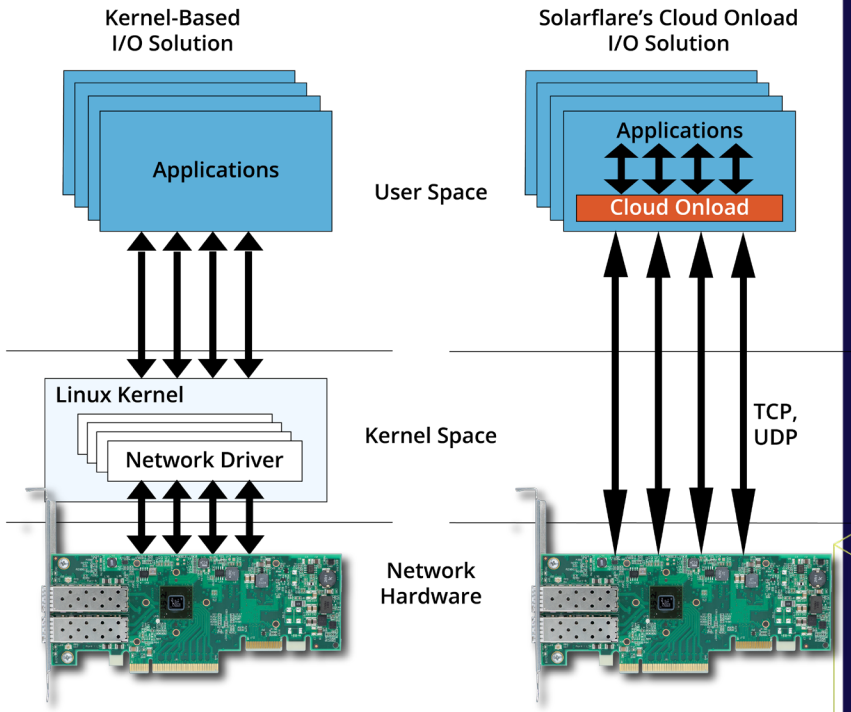

User Space Cloud App Acceleration

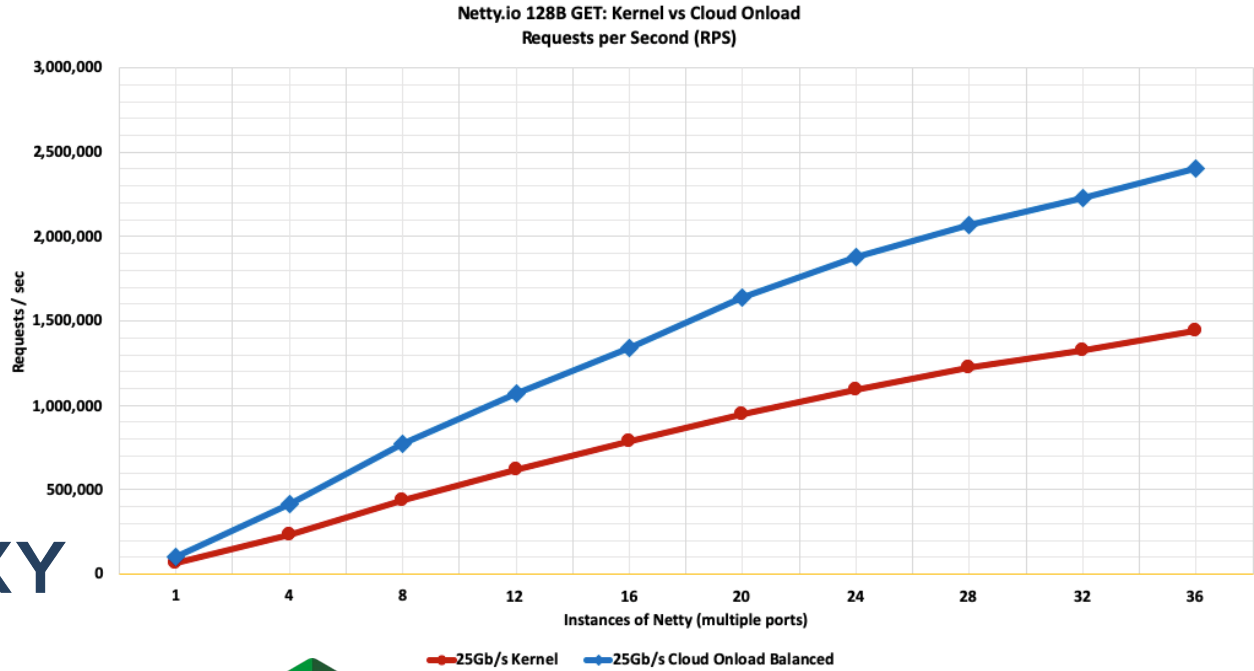

User Space Performance is Proven

- Pioneered user-space acceleration for the most demanding markets:
  - FinTech
  - Managed Service Providers
- Onload in the cloud provides performance gains for:
  - IMDBs
  - Software Load Balancers
  - HTTP and Reverse Proxy Servers
- Primarily software based application acceleration

SSDs Have Changed Storage

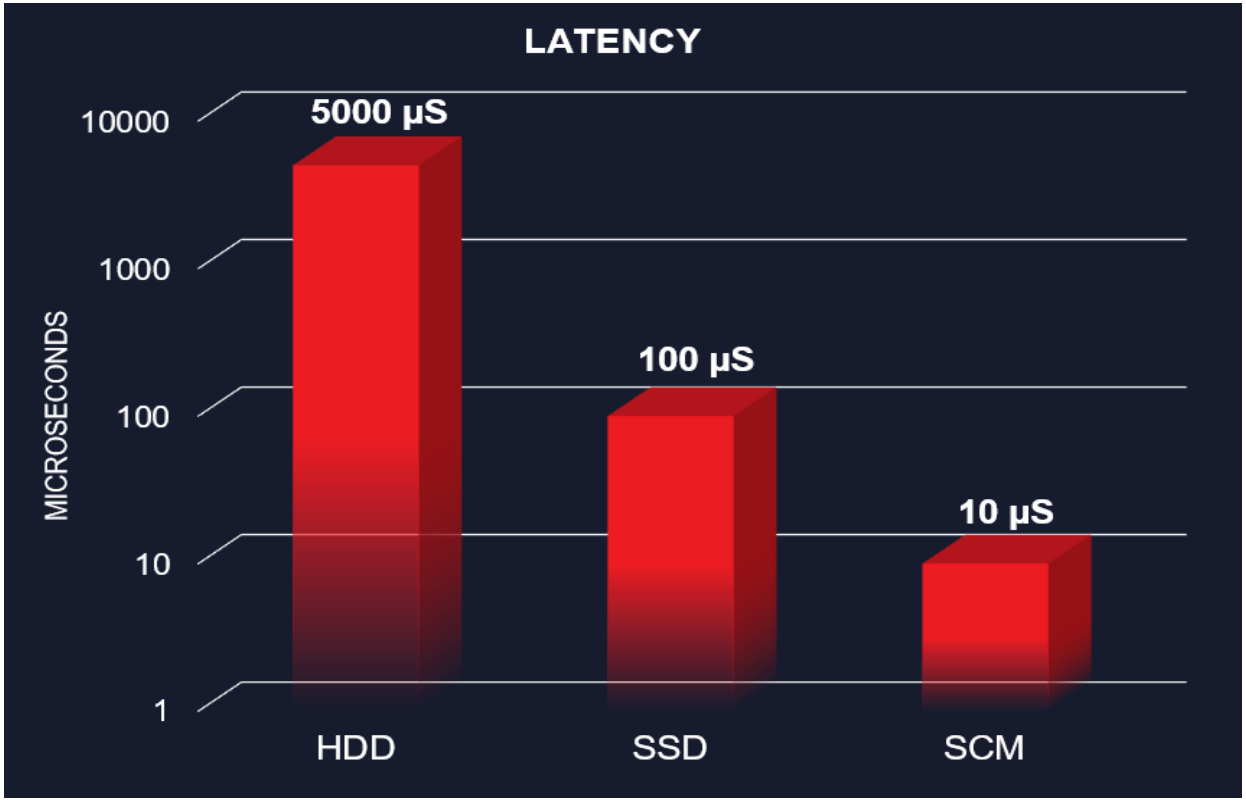

NVMe-oF Options – Protocol Stacks

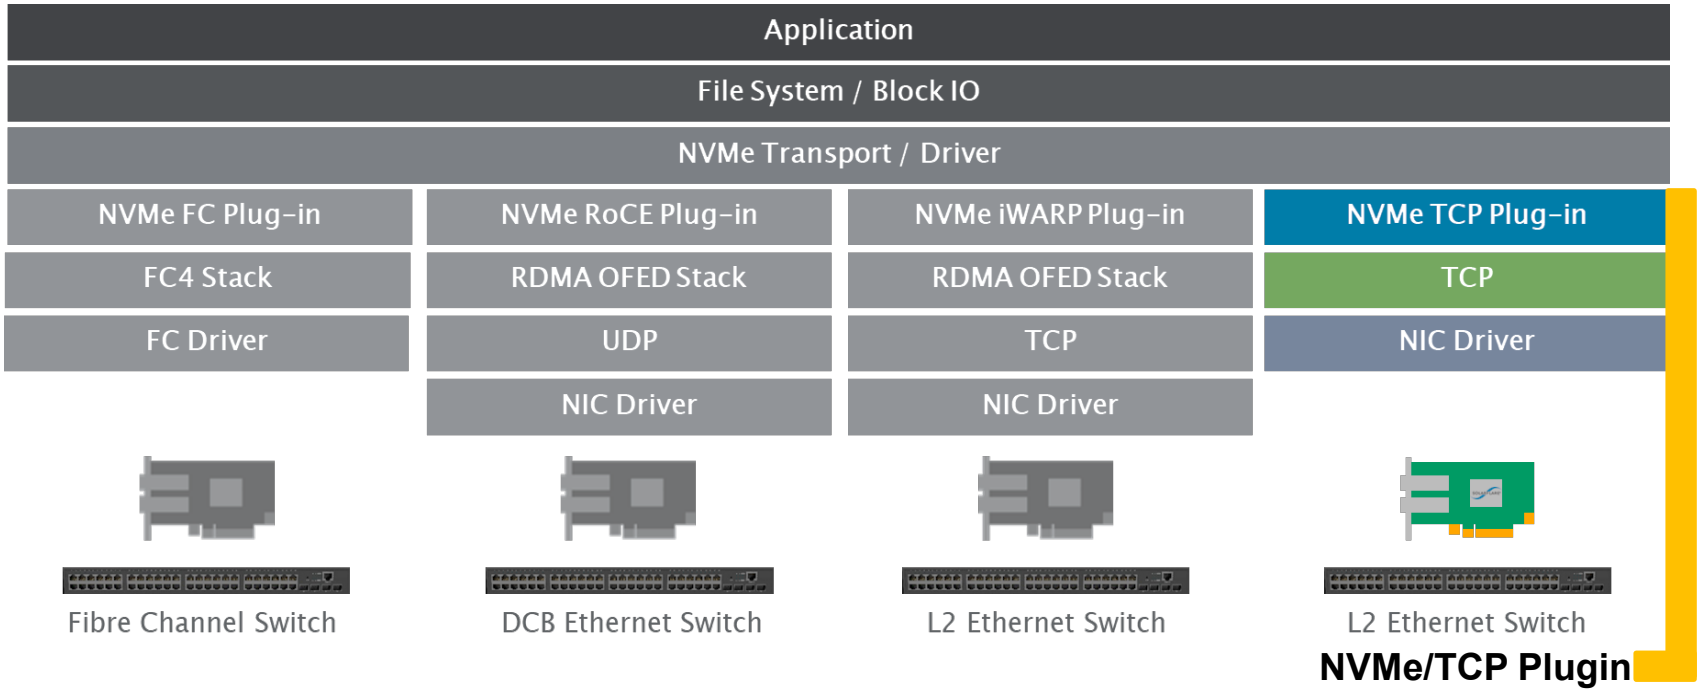

NVMe/TCP Testing on SuperMicro Server

10/25/50/100G
TCP

Initiator "App" Server
Red Hat RHEL Linux 7.4
SAMSUNG 1750 SSD

Ethernet Network

VPN

10/25/50/100G
TCP

SAMSUNG
NVMe
Flash
Storage

Target Storage Server
Red Hat RHEL Linux 7.4
SAMSUNG 1750 SSD

Tested by Solarflare

High Throughput | Low Latency

FIO Test Results	TCP
1 Thread 4K READ Throughput	35MB/s
32 Thread 4KB READ Throughput	1010 MB/s
Random 4K READ Latency	99 µs

NVMe/TCP Benchmarking Read Latency

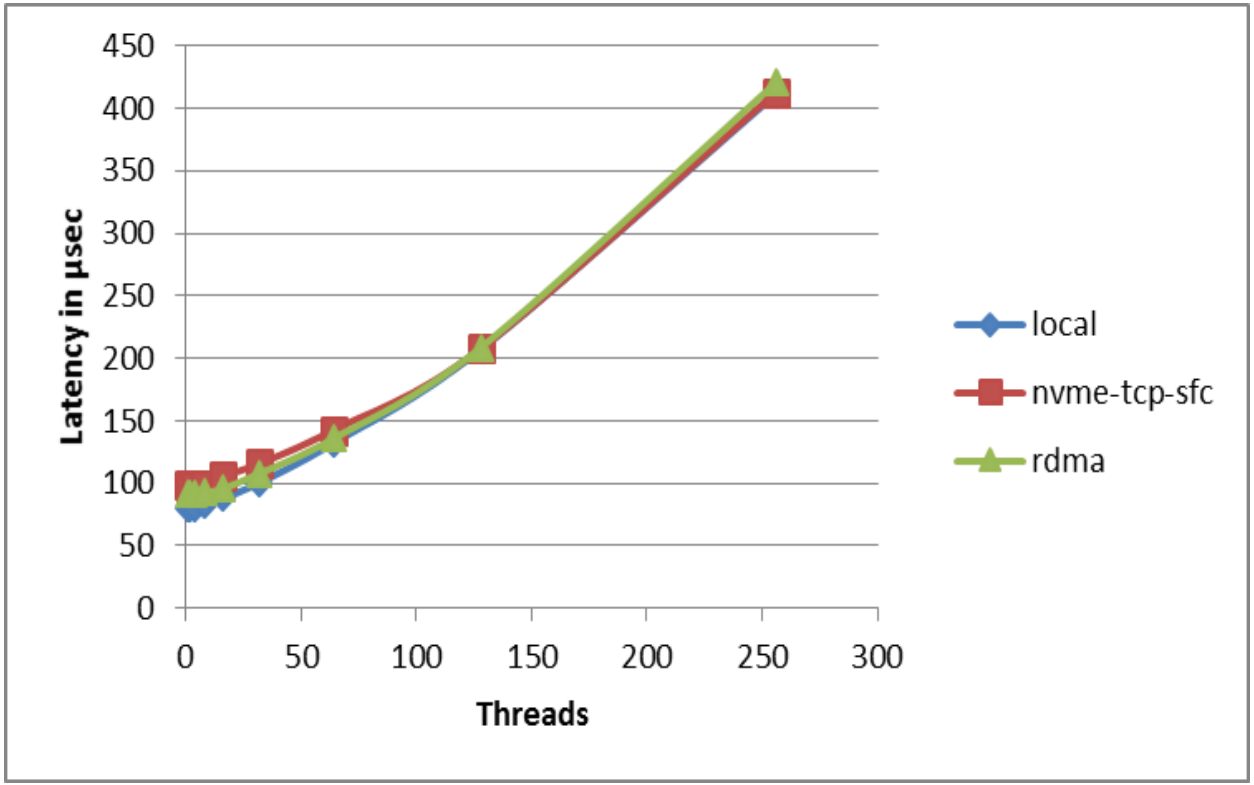

Hardware Based Acceleration

FPGAs are a Critical Acceleration Component

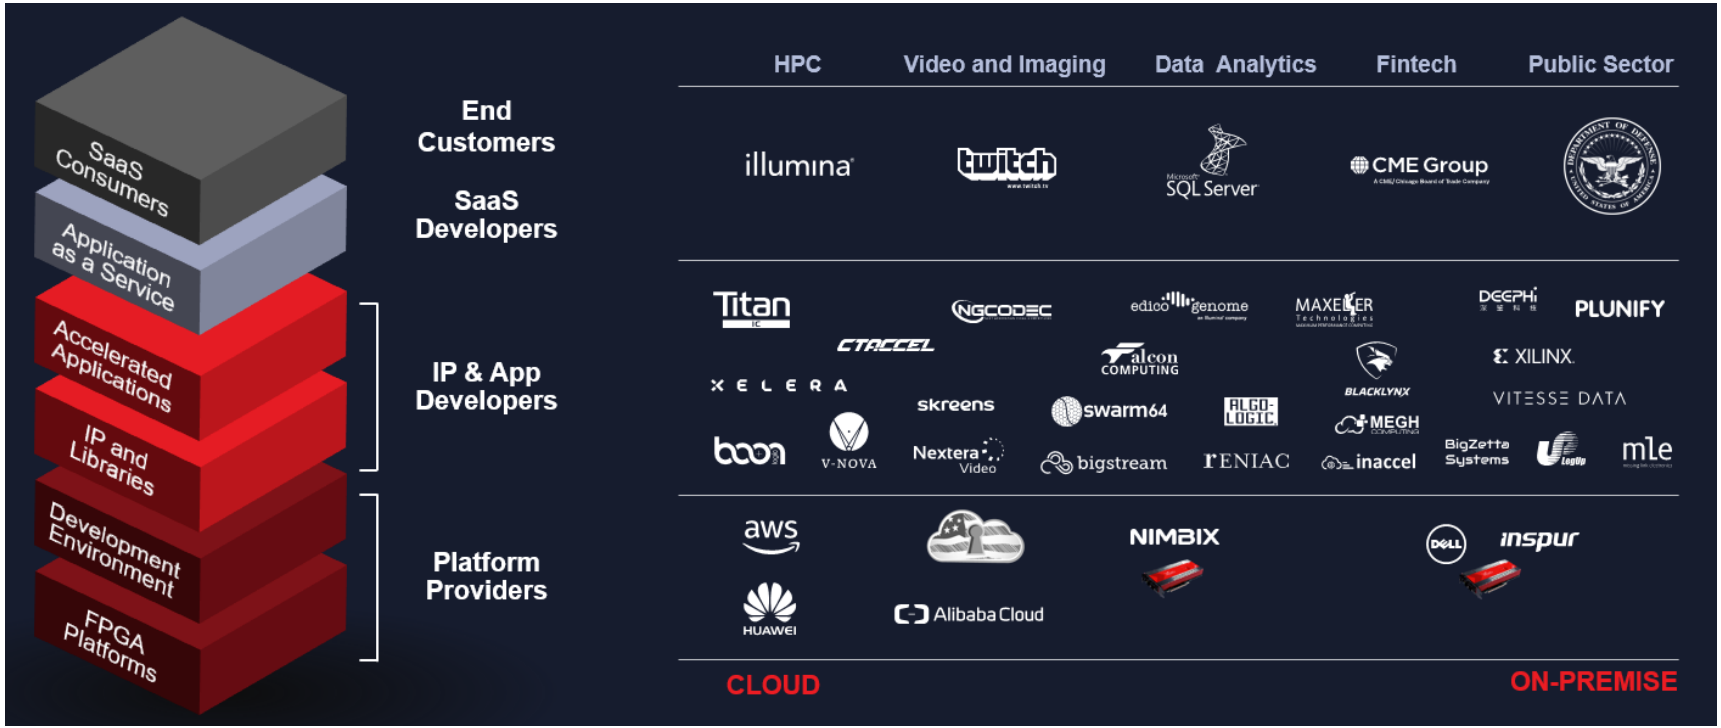

Applications Accelerated with FPGAs

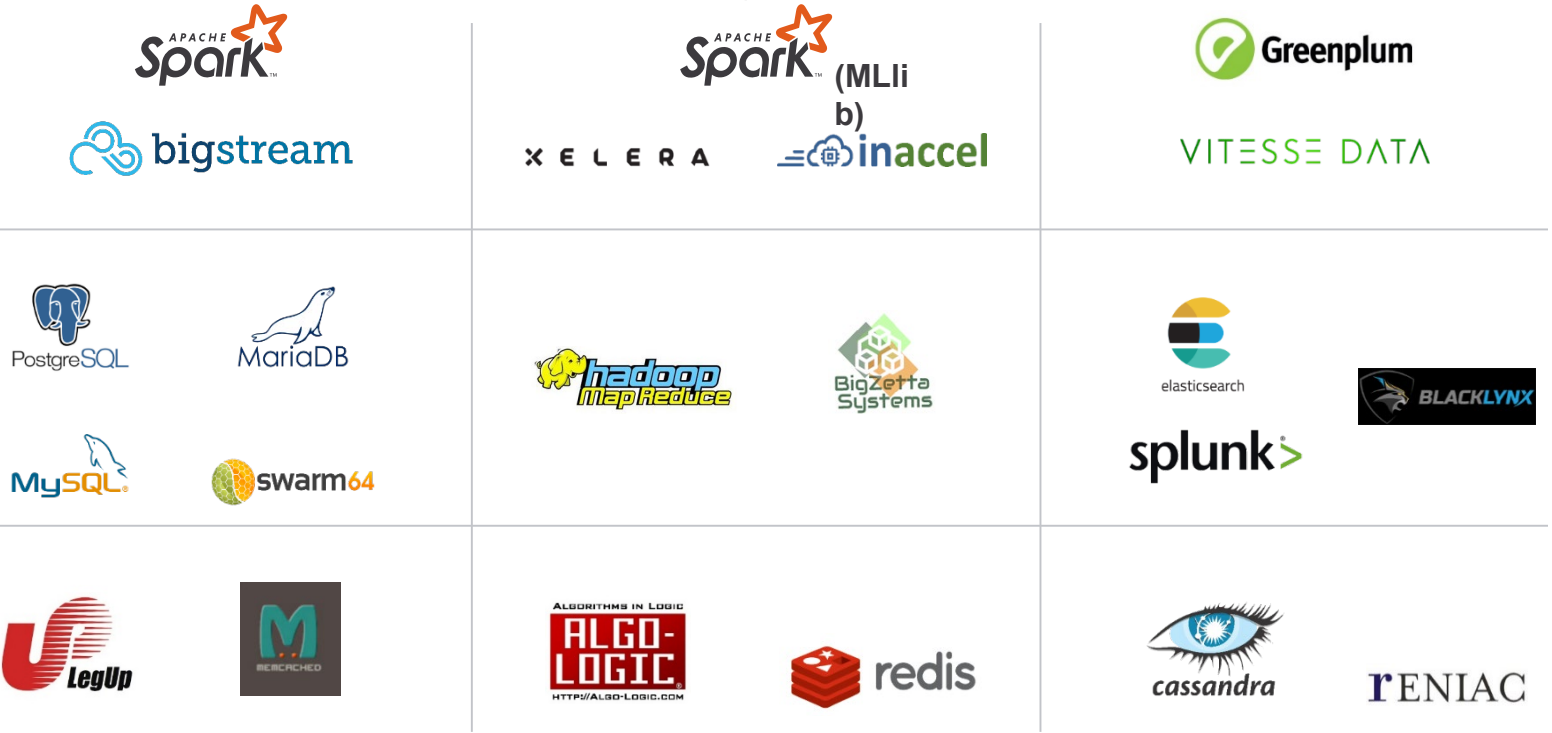

Accelerated Apache Spark

ETL, Streaming Analytics, Spark SQL, Kafka, Hive

Zero code change

Intelligent, automatic computation across CPUs and Alveo

2x- 30x acceleration

Average 4x-5x on Alveo

Hardware Based Security for Network Storage

Four Keys to Hardware Based Security

- Tamper Resistant NIC platform
- Agentless Enforcement & Telemetry
- Secure External Control Plane
- Centralized Orchestration

Tamper Resistant NIC Platform

- Firmware must be digitally signed & checked
- Log ALL attempts to load tampered firmware
- Be capable of terminating a secure connection without requiring an agent
- NO Root access to security functions
- Binding process to link NIC to central controller

Agentless Enforcement & Telemetry

- Agentless is means NO host based software is required
- Enforcement on all traffic, Rx and Tx
  - Five tuple, Regx or both depending on HW
  - Filtering & Alerting
  - Support for millions of rules and IPv6
  - Nominal latency impact
- Flow reporting to central controller
  - New flows
  - Packet and byte counts
  - Rates/flow if HW can support

Secure External Control Plane

- Pathway needs to exist so centralized management can be enabled

- On the NIC side this TCP management needs to be secure & fully autonomous

Centralized Orchestration

- Terminate secure control plane for all bound NICs under management
- SDN “Neutron” like controller with RestAPI
- Support for high level abstractions
- Cluster capable database to store all objects
- Double log everything, standard Linux system logs, and proprietary

Why Hardware Enforcement?

- Because software solutions are easily beat
- Hardware enforcement should be transparent to root users
- Adds another layer to the security onion
- Zero attackable local surface area
- Cloak server from all but service ports
  - Becoming a high performance Honeypot
- Extremely high performance ~200ns
  - Fortinet appliances by contrast 2500ns
- Compliance validation solution

Rise of the SmartNIC

All the above is possible and very probable with the next generation of SmartNICs