SDC 20
SNIA EMEA
FEBRUARY 4-5, 2020
TEL AVIV, ISRAEL
STORAGE DEVELOPER CONFERENCE

Computational Storage
Architectural Discussion

Eli Tiomkin
NGD Systems

Today's Learning Opportunities (TLO).

- EDGE needs CSD in M.2, EDSFF
- Architecture Our Way – CSD with PCSS
- AI – ML – CSD – the Overlap
- Hadoop & DB – CSD – Growth, Scale

Data, Data, Data. But Don't Take Our Word For it.

3

What Are You Doing with Your Data Today?

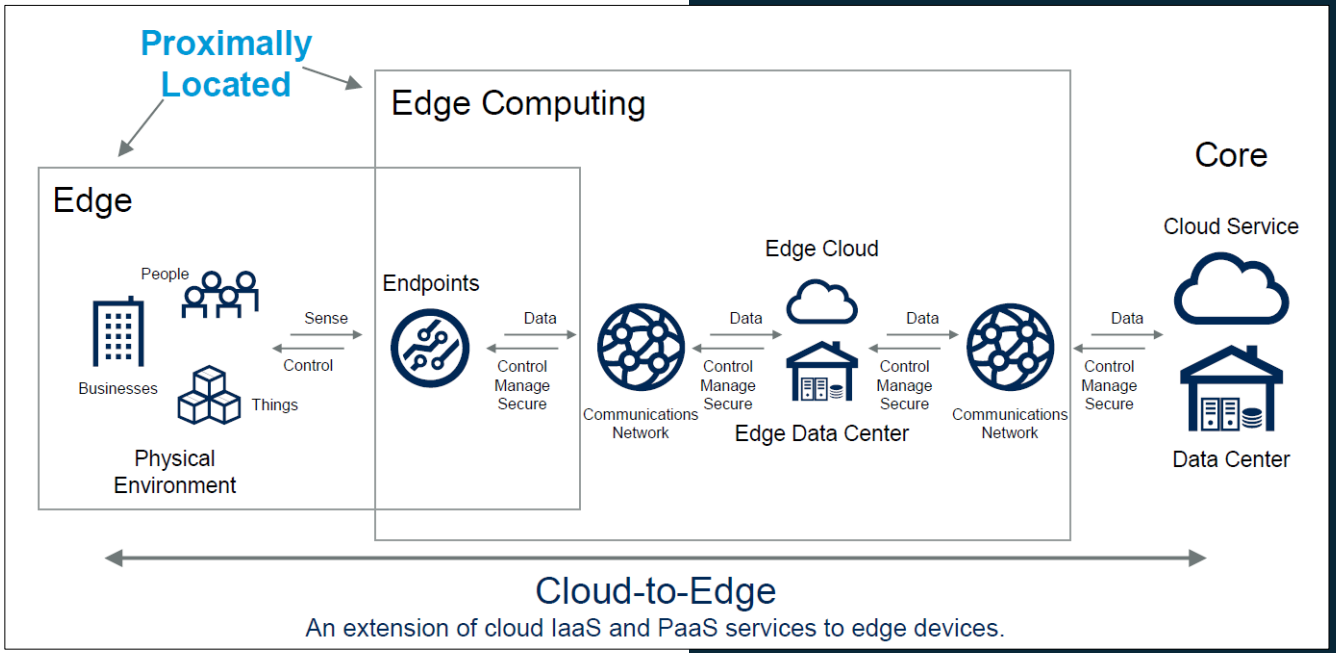

Source: Gartner - Bittman

It's No Longer Black and White.

What is Driving Our Data Analytics Issues?

Weeding through the Noise at the Edge

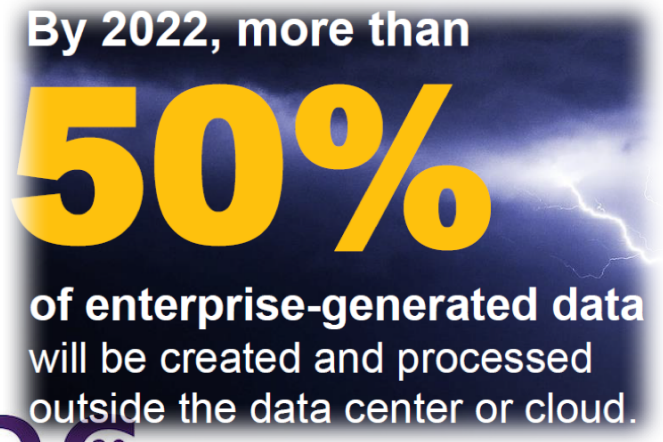

Source: Gartner - Bittman

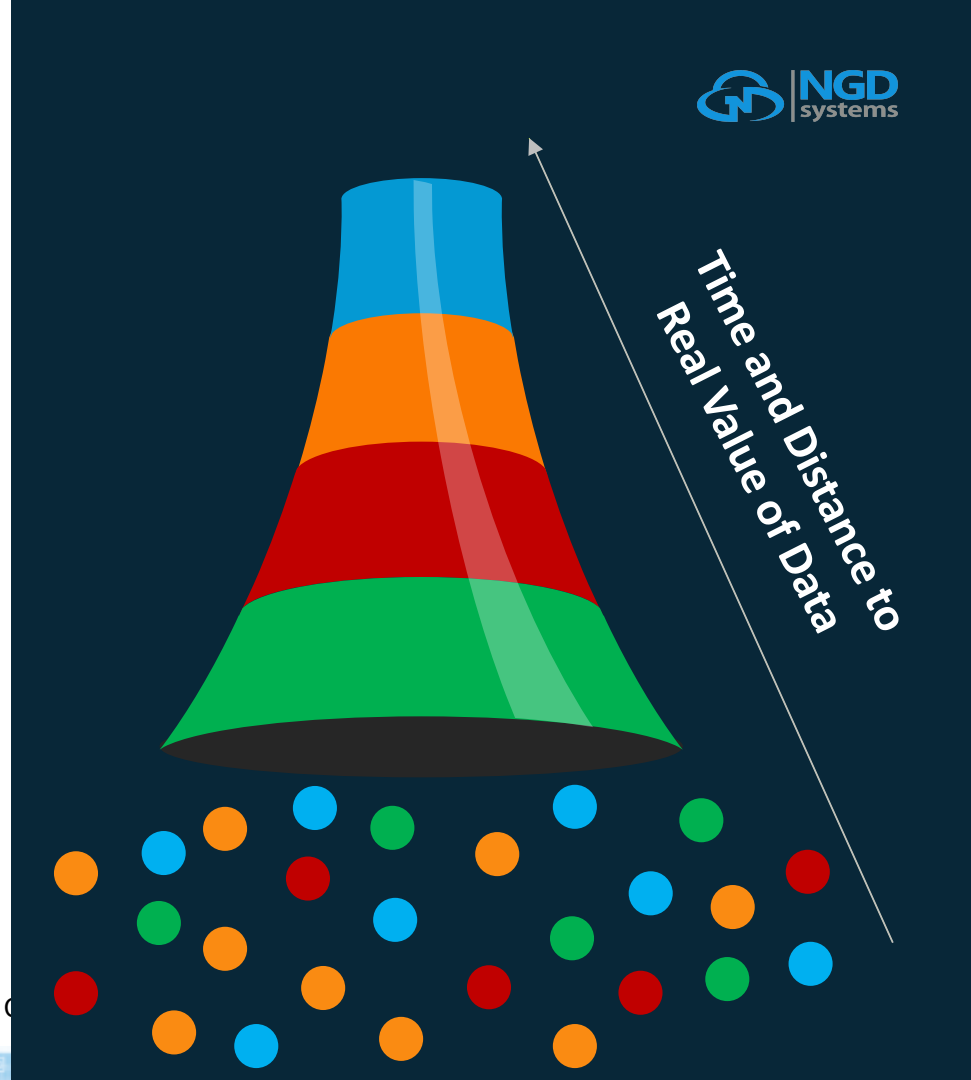

From Edge Sensors to Centralized Cloud

Computational Storage opportunities exist throughout the distributed compute environment

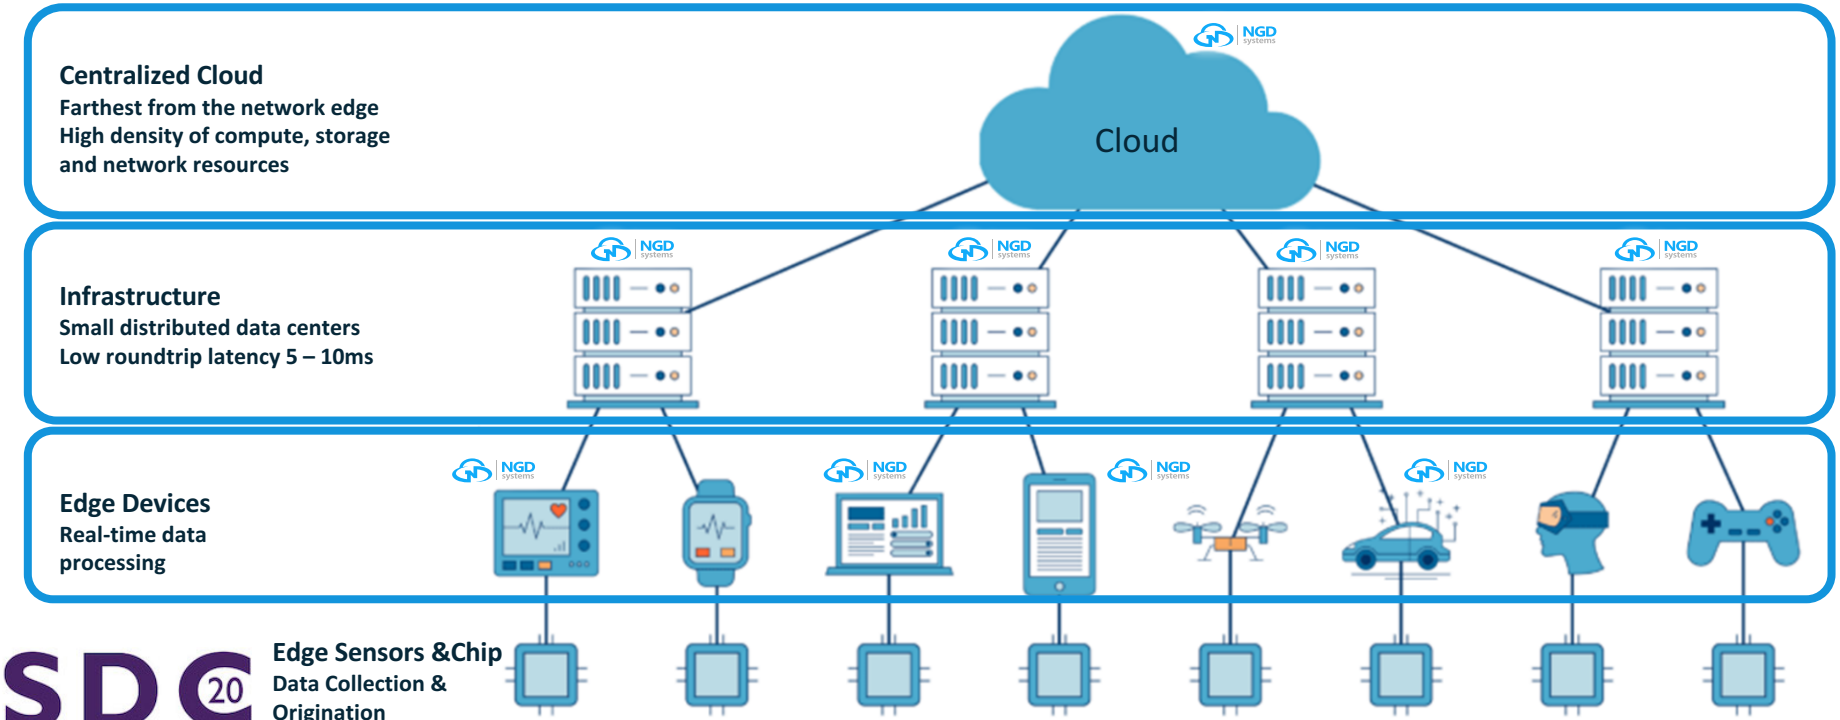

Innovative Computational Storage Uses.

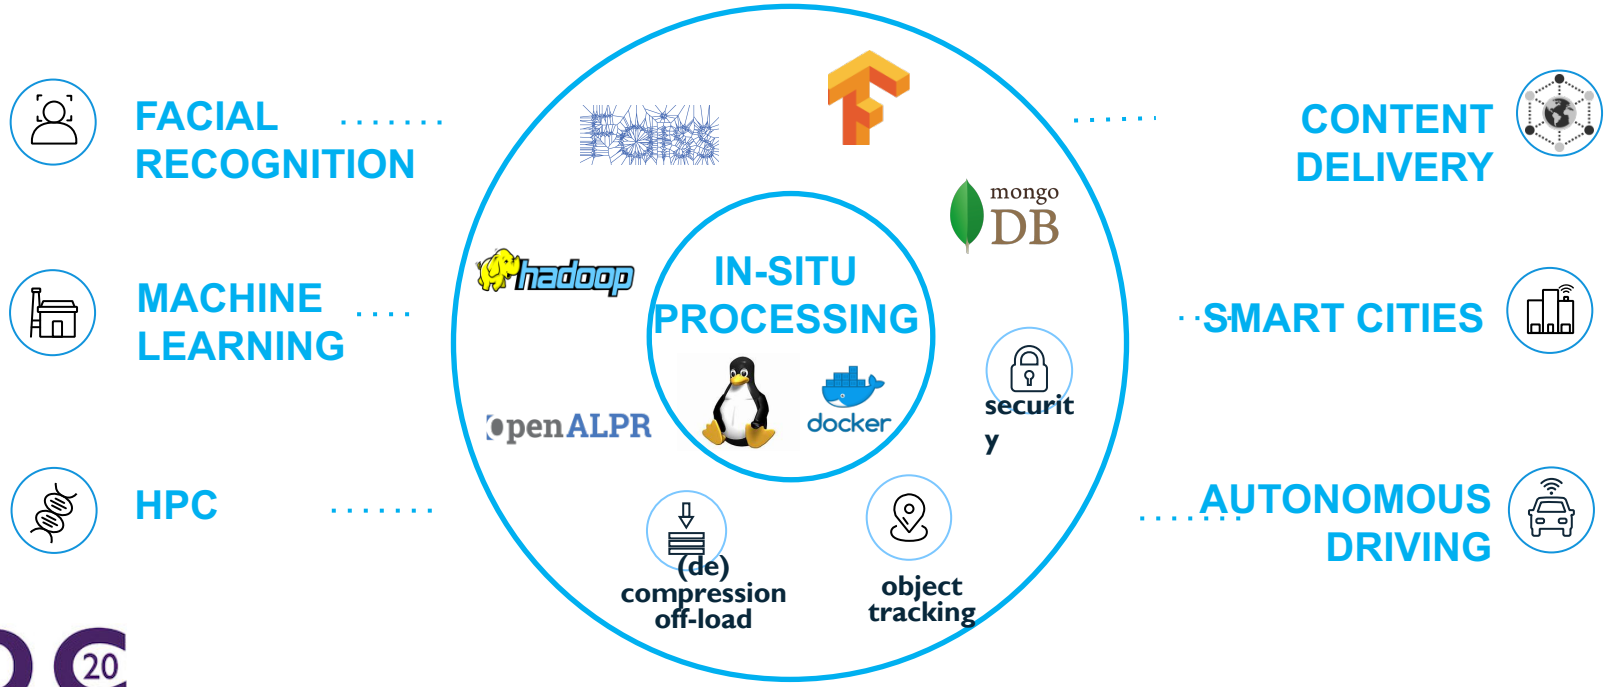

Highest Capacity, Lowest Power.

Industry leading W/TB

Industry's Only 16-Channel M.2

Industry's Largest Capacity U.2

8TB/8W M.2

32TB/12W

16TB/12W E1.S

Form Factor	Capacity (TB)	MAX Power (W)
M.2 22110	Up to 8	8
EDSFF E1.S	Up to 12	12
EDSFF E1.L	Up to 32	12
U.2 15mm	Up to 32	12
AiC FHTQL	Up to 64	15

Deve

Today's Learning Opportunities (TLO).

- EDGE needs CSD in Compact Form Factor
- Architecture One Way – CSD with PCSS
- AI – ML – CSD – the Overlap
- Hadoop & DB – CSD – Growth, Scale

Complete Solution and Disruptive Technology.

1st FULLY INTEGRATED COMPUTATIONAL STORAGE SOLUTION

SoC Controller
Industry's First 14nm

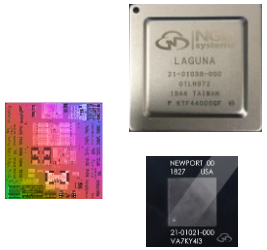

Management

- Modular firmware
- Efficient algorithm
- Flash characterization

Optimized Hardware

"In-Situ Processing" Computational Storage Stack

- Full fledged on drive OS
- Light virtualization
- Quad-core 64-bit application processor
- Hardware acceleration

Linux

docker

ARM

SDC 20 SNIA EMEA

The Scalable, ASIC-based Computational Storage Drive.

An enterprise class device capable of processing workloads in storage at the source

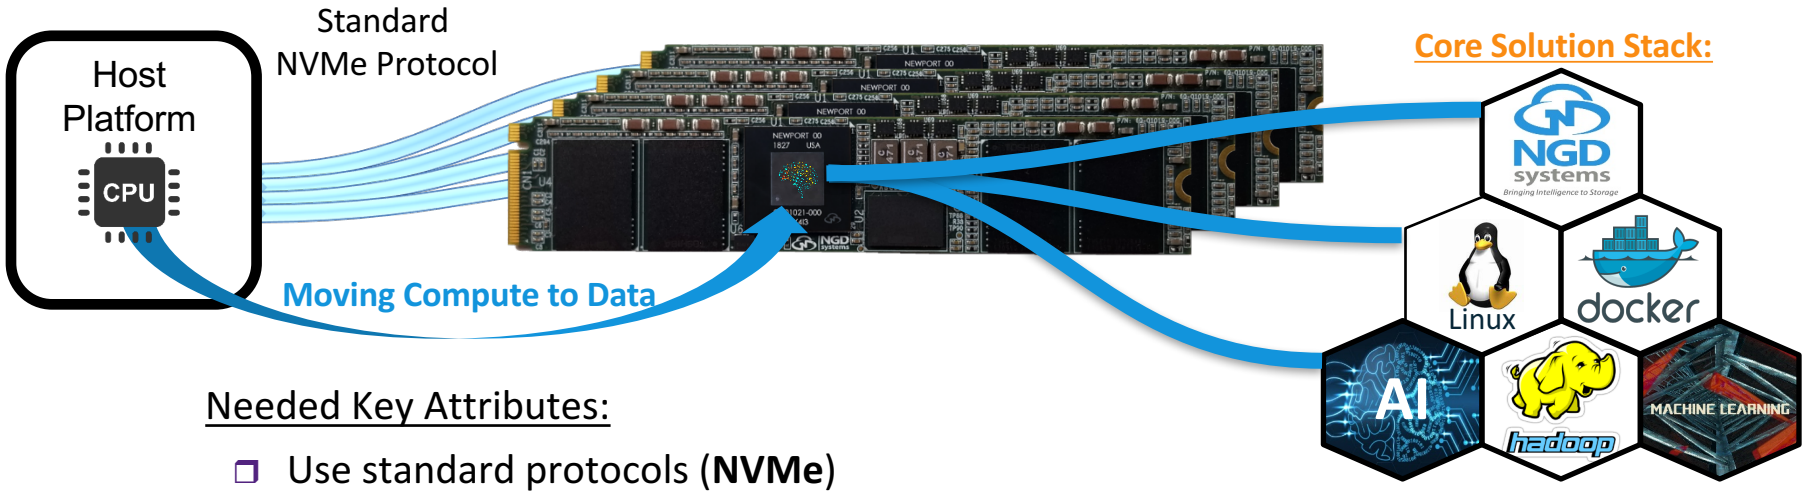

Needed Key Attributes:

- Use standard protocols (**NVMe**)
- Minimize data movement (Faster Response, Lower **W/TB**)
- Improve (**TB/in³**) with maximize (Customer **TCO**)

SDC 20 SNIA EMEA

The Data Lives on Storage.

Why Not Work on it There?

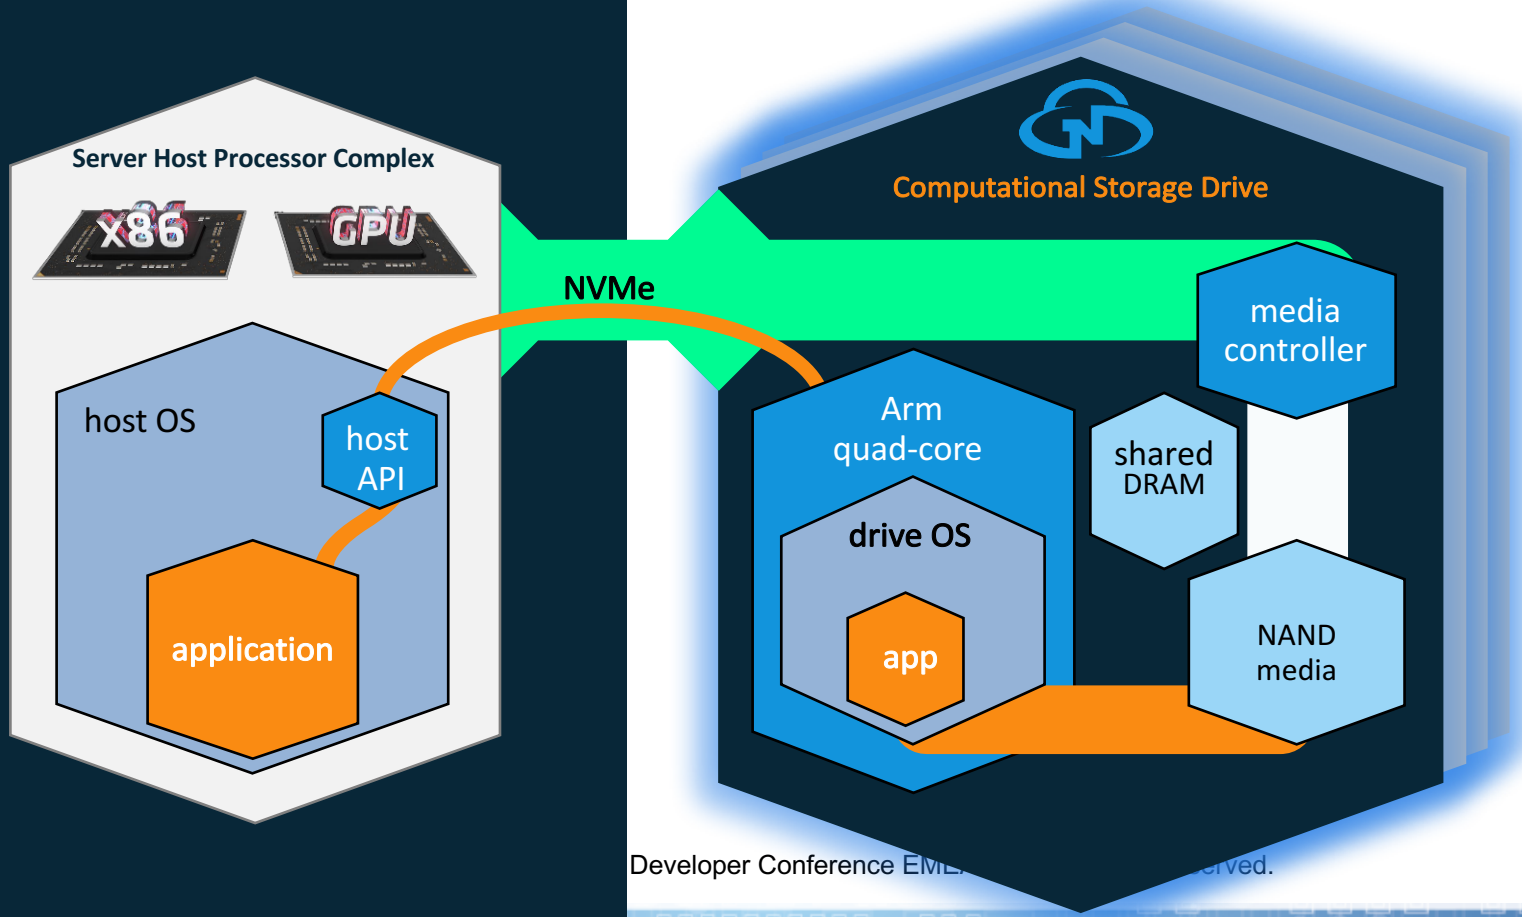

Today's Learning Opportunities (TLO).

- EDGE needs CSD in Compact Form Factor
- Architecture One Way – CSD with PCSS
- AI – ML – CSD – the Overlap
- Hadoop & DB – CSD – Growth, Scale

Azure IoT Edge Implementation

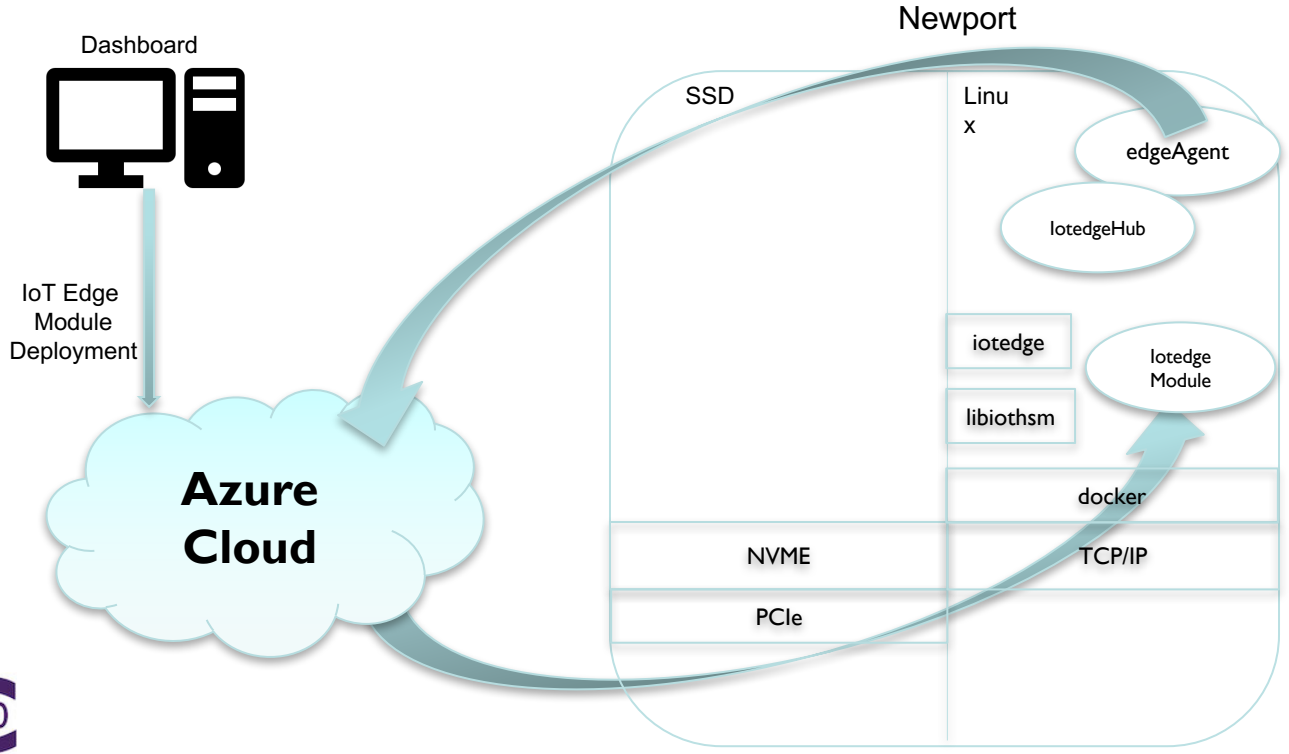

Using Computational Storage Drives for ML.

MobileNet Object Classification

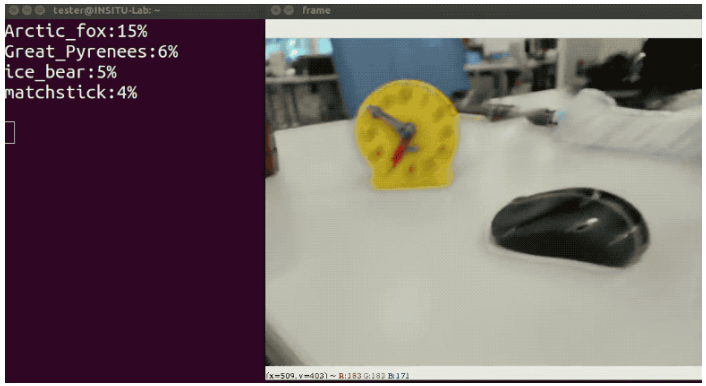

TensorFlow & Keras API

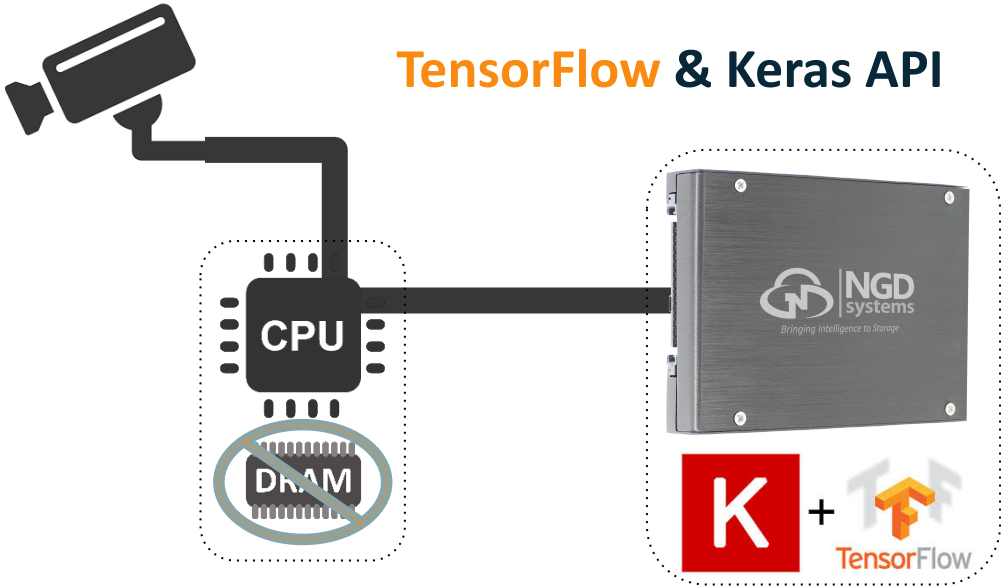

No Host Interaction Required

The Leader in Computational Storage Solutions

Computational Storage Edge Deployment View.

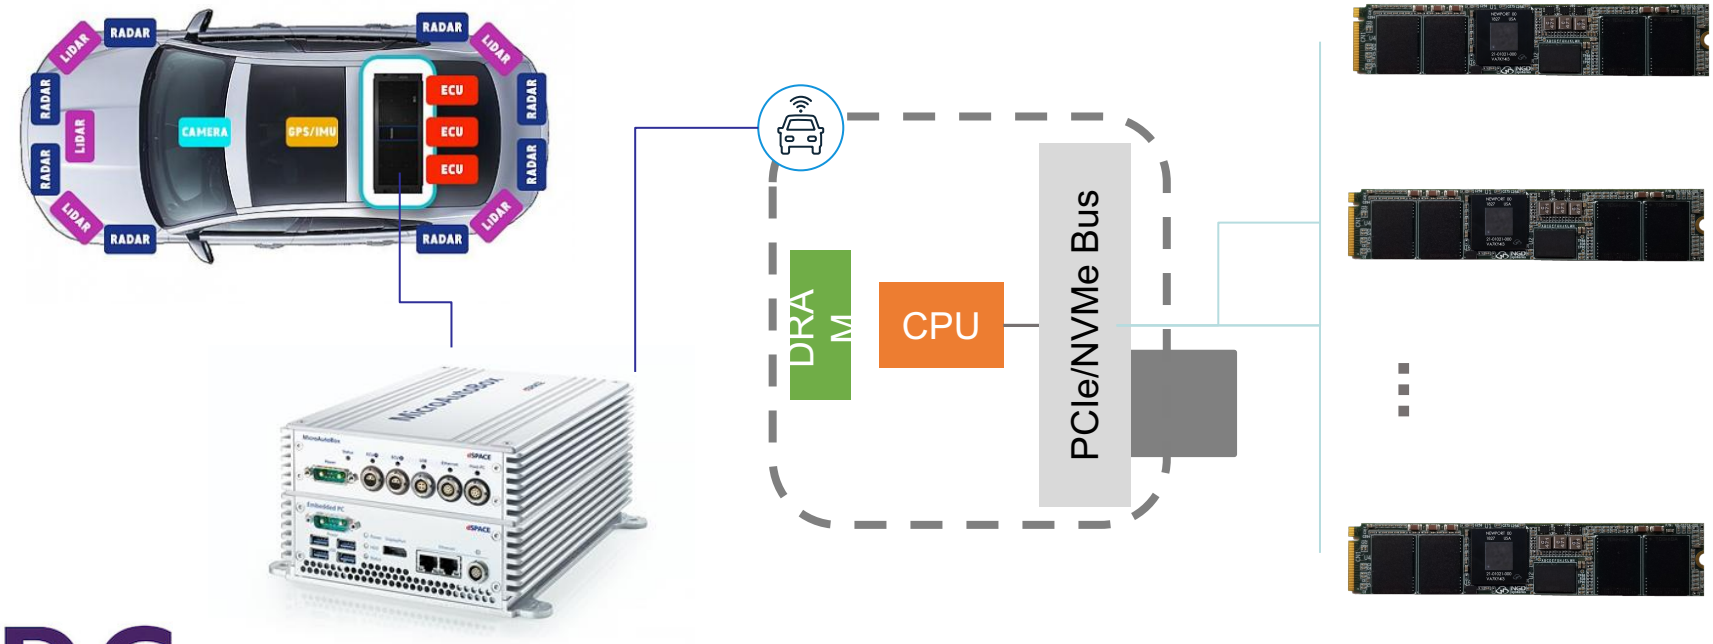

Computational Storage Server Deployment View.

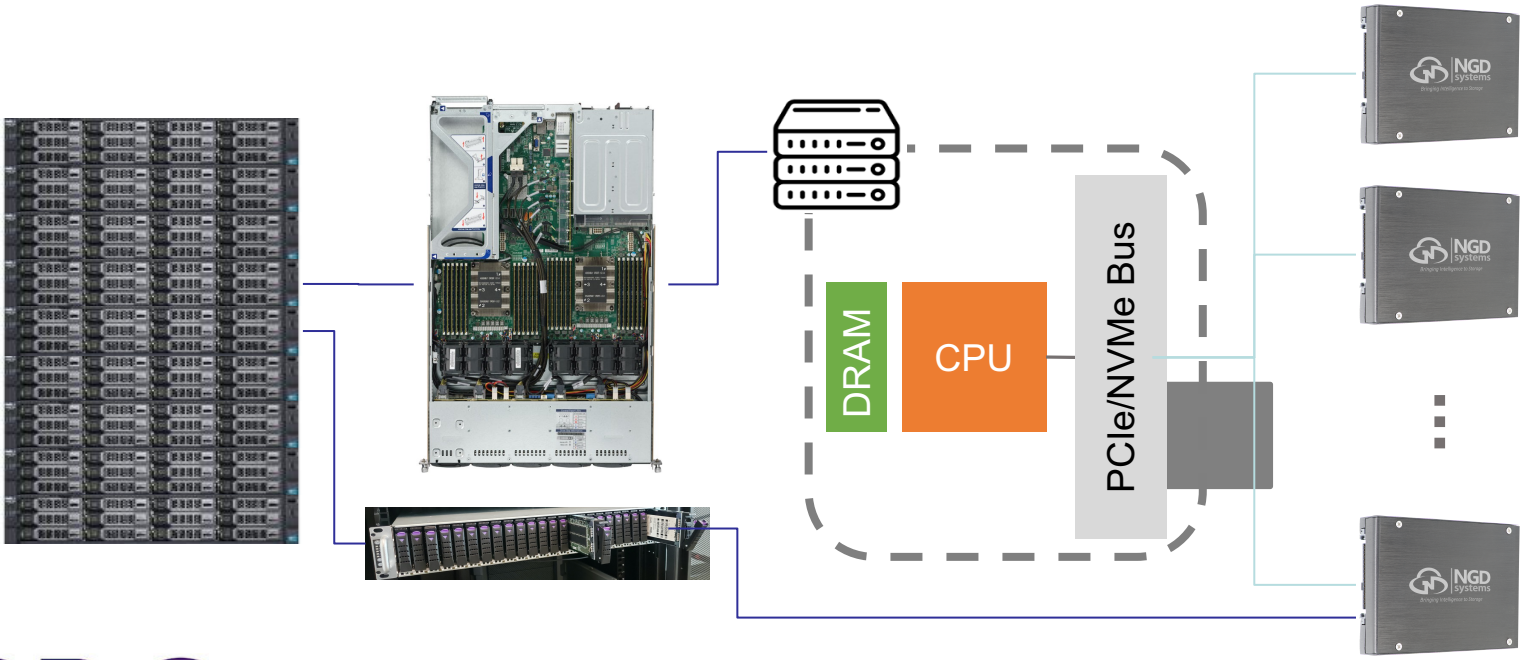

Amplifying TCO for Hadoop

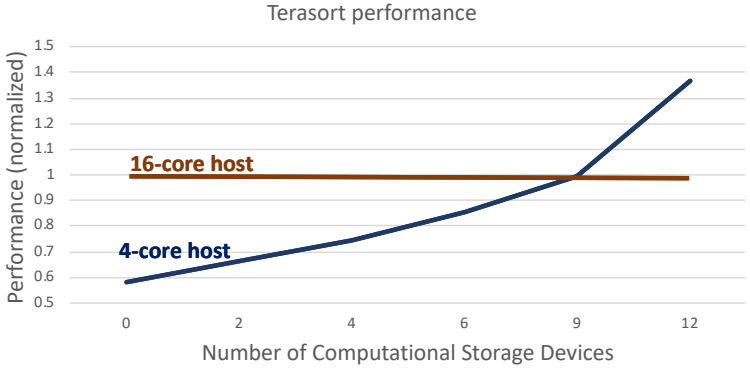

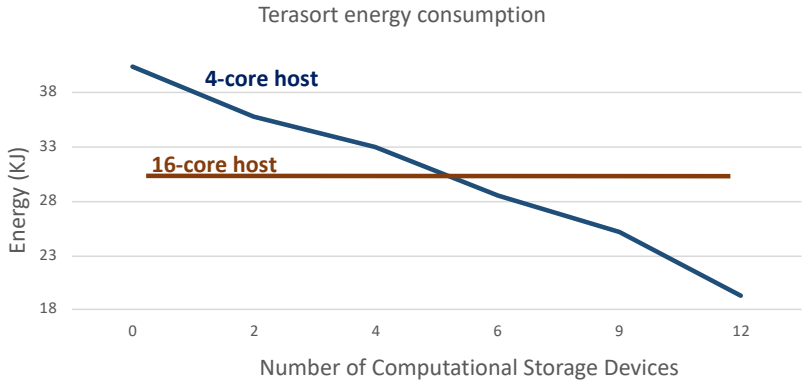

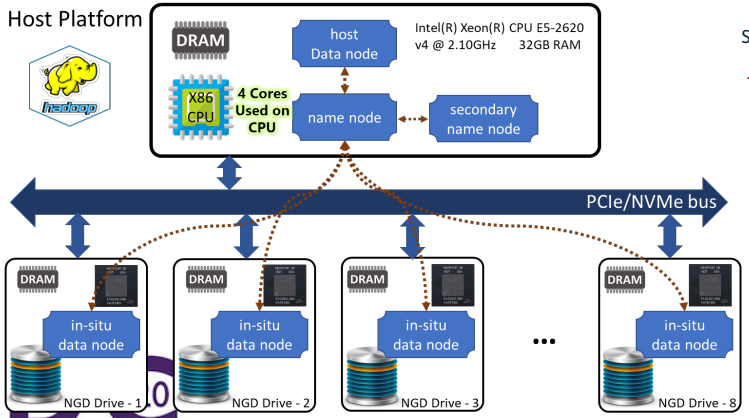

Datanode Config:
Single E5-2620v4, 32GB DRAM, 12*8 TB SAS HDD

18U Total Density in 18U = 864TB

Datanode Config:
Single E5-2620v4, 32GB DRAM, 36*8TB NVMe

3U Total Density in 3U = 864TB

432 Additional Drive Cores

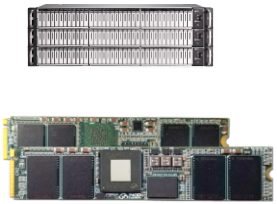

@ Scale
Saves Power!
Saves Space!
Saves Time!

SDC SNIA EMEA

Augmenting system performance

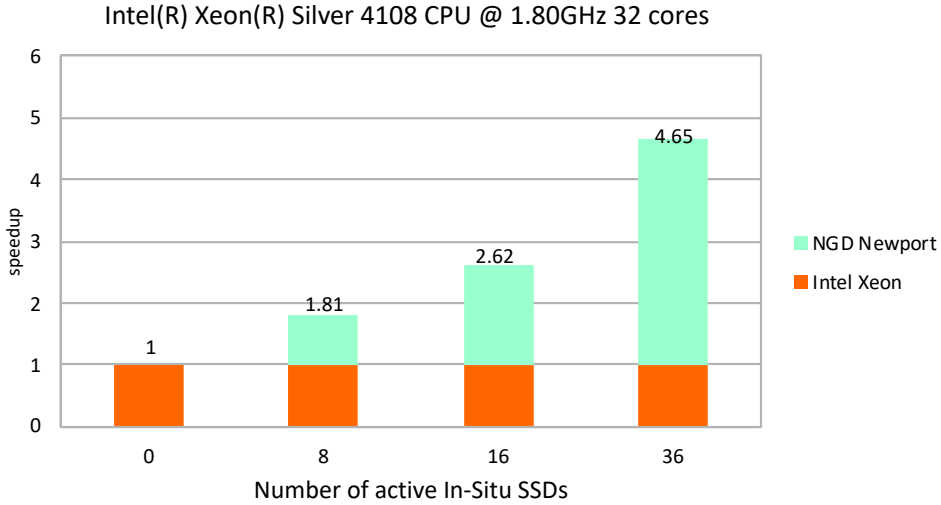

TensorFlow's MobileNetV2 (inferences/sec)

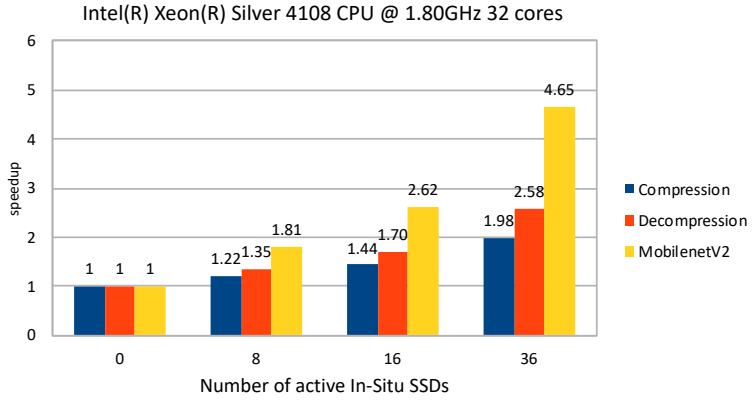

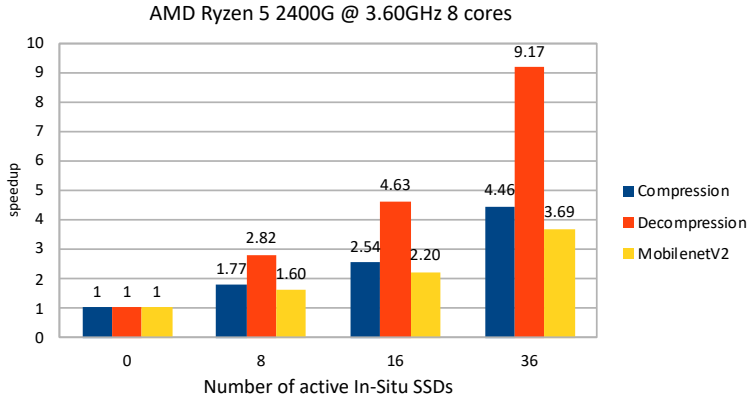

Compression (Gzip) Using Computational Storage

- Intel Xeon CPU E5-2620 v4 @ 2.10GHz, 64 GB RAM (4 channels)
- SSDs Tested:
  - NGD's Newport NVMe 8TB (Gen3 - 4 lanes)
  - Micron PRO 9300 NVMe 3.8TB (Gen3 - 4 lanes)
  - Intel DC P3600 NVMe 450GB (Gen3 - 4 lanes)

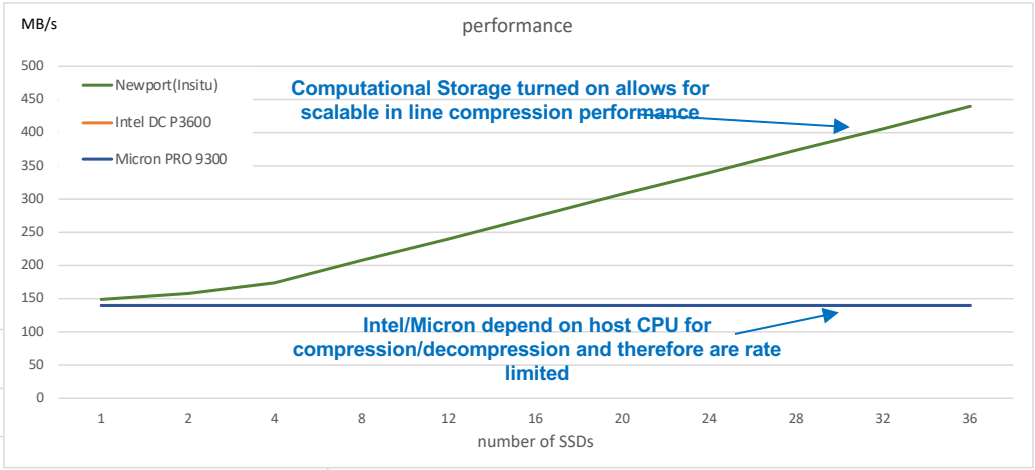

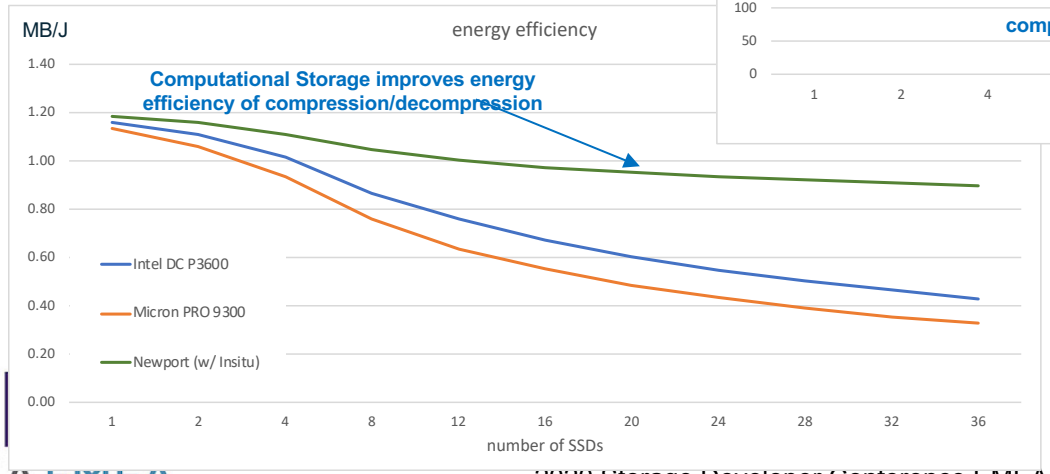

Using MongoDB within Computational Storage.

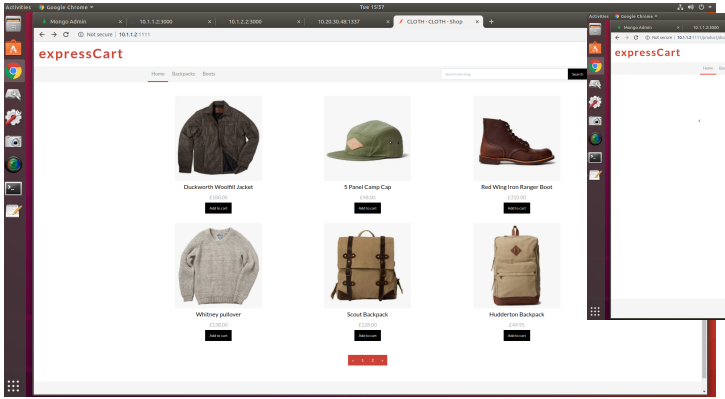

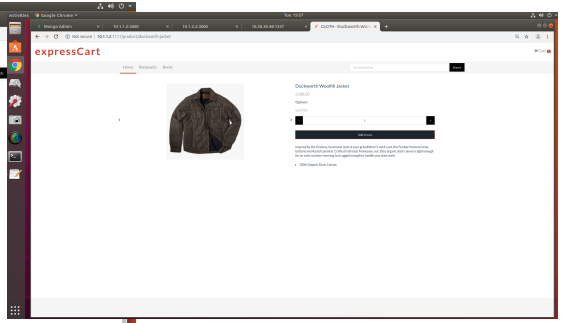

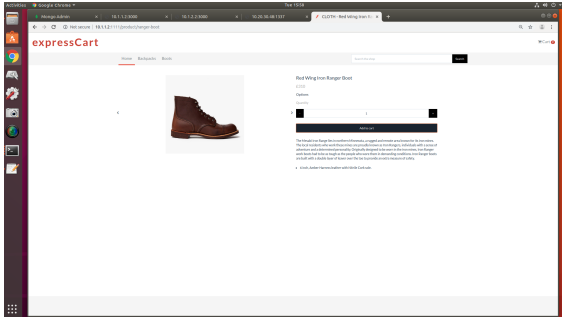

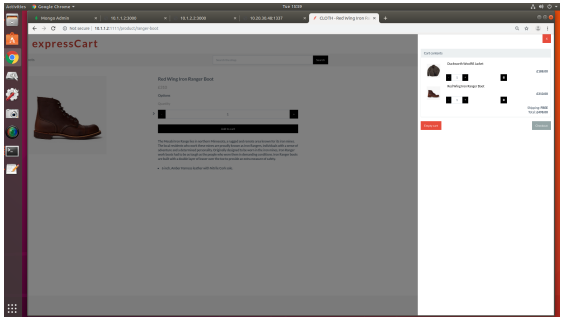

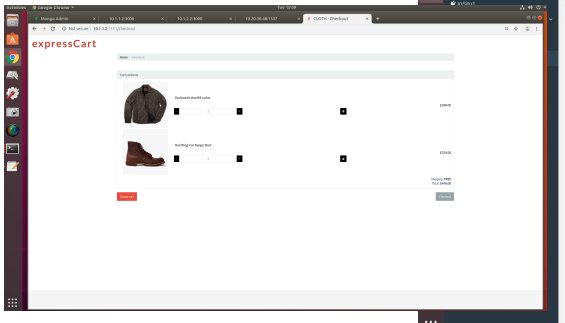

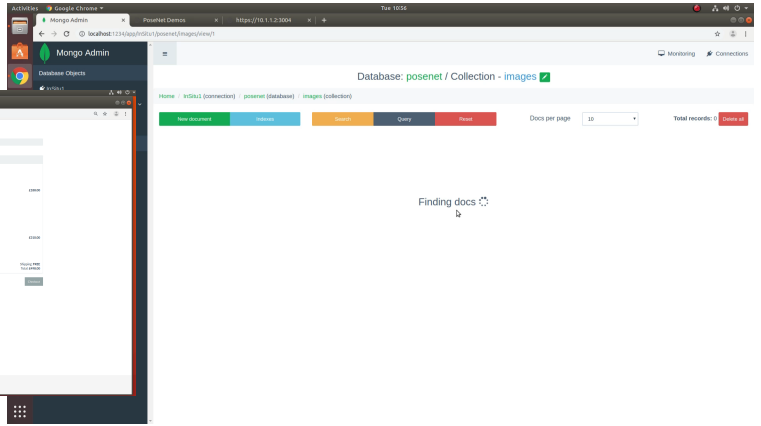

Scalable Computational Storage.

A New Storage Paradigm is Here

- **The “New Cloud” needs the Distributed Edge**
  - There is no longer just a ‘central’ storage location
- **Edge data growth challenges HW platforms**
  - Innovative form factors and high capacity for the Edge
- **In-Situ Processing brings ML closer to data**
  - Exploit data locality and enable distributed processing

SDC 20
SNIA EMEA
FEBRUARY 4-5, 2020
TEL AVIV, ISRAEL
STORAGE DEVELOPER CONFERENCE

Computational Storage

eli.tiomkin@NGDSystems.com
www.NGDSystems.com

TW: @NGDSystems

ML Training with Traditional Approach.

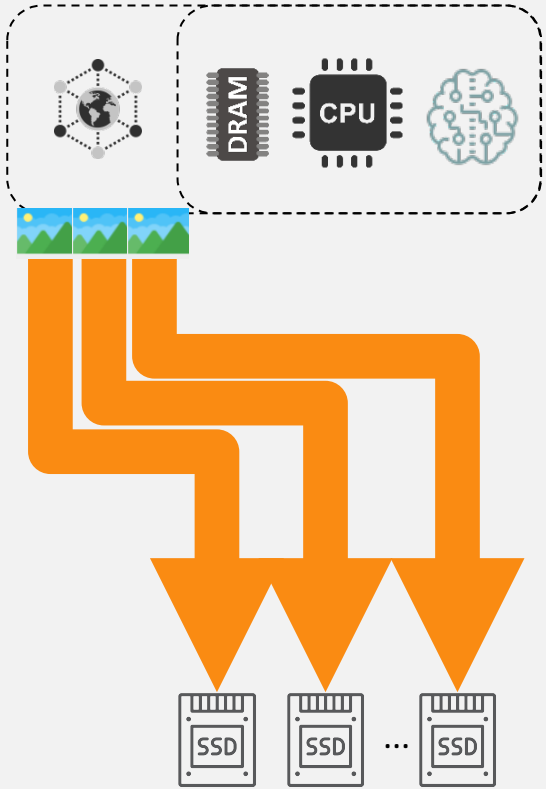

Load Data

ML Training with Computational Storage.

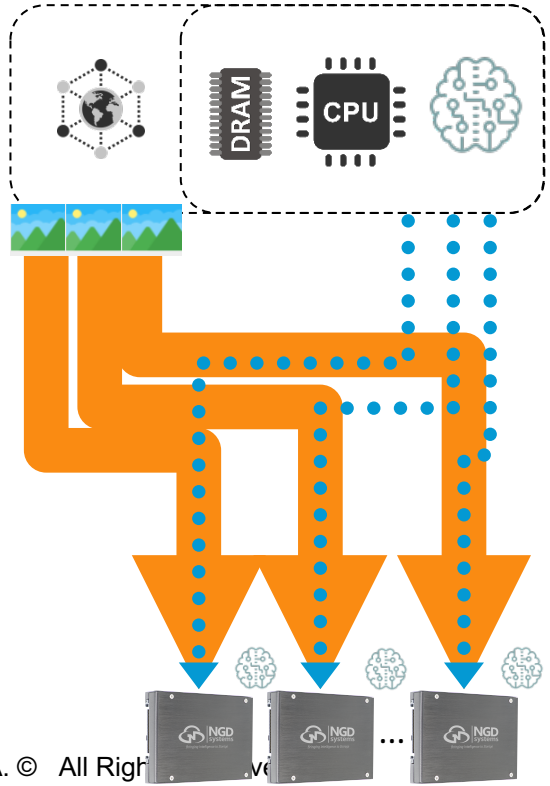

ML Training with Traditional Approach.

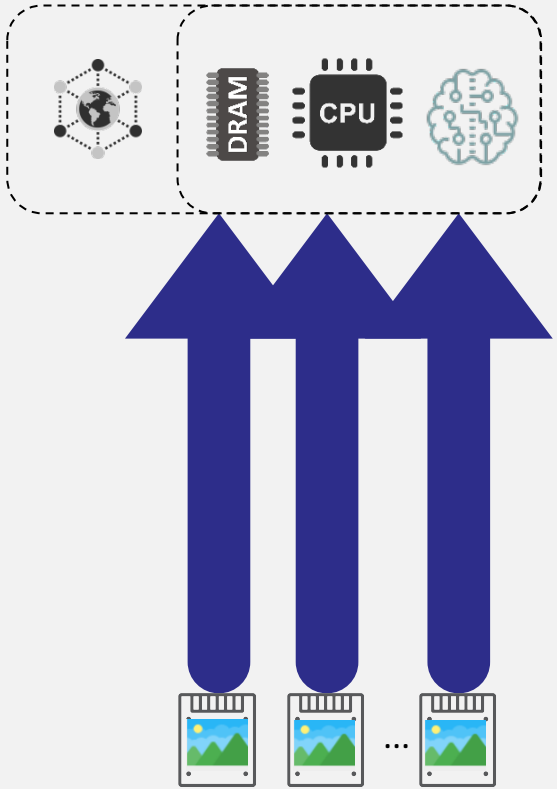

ML Training with Computational Storage.

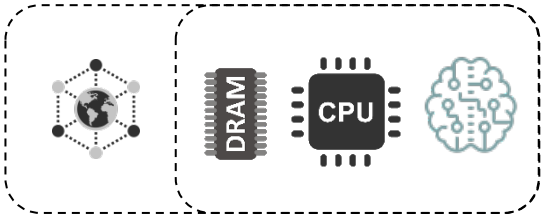

- No data movement
- No host CPU needed
- Distributed training

ML Training with Traditional Approach.

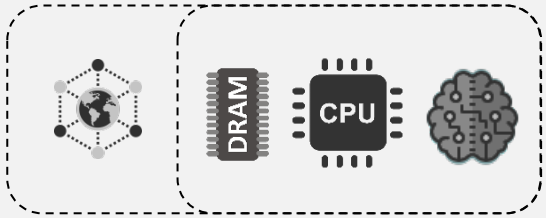

- Host CPU still needed
- No Parallelism

ML Training with Computational Storage.

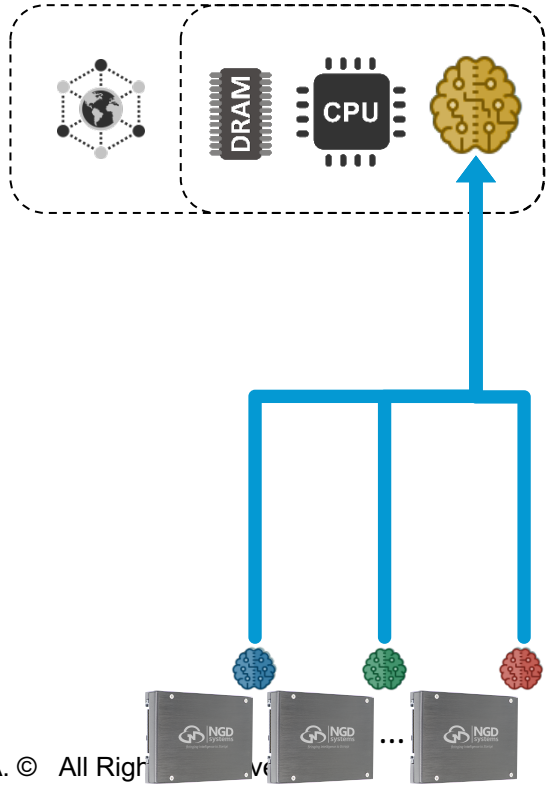

ML Training with Traditional Storage.

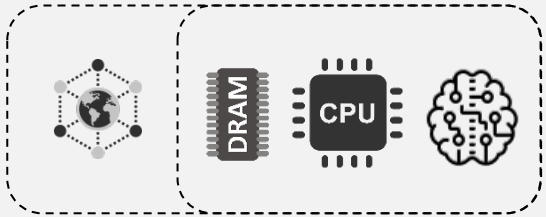

Repeat Steps

- Train
- Evaluate
- Update

ML Training with Computational Storage.

- Load Data
- Train
- Evaluate
- Update
- Deploy