SDC 20
SNIA INDIA

Storage Developer Conference
December 4-5, 2020
BY Developers FOR Developers

Computational Storage
At the Edge and Beyond

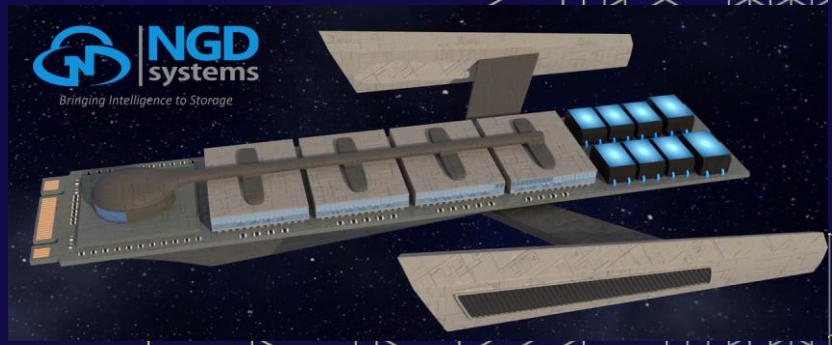

Scott Shadley
VP Marketing, NGD Systems
Director, Board of Directors, SNIA
Co-Chair, Computational Storage TWG, SNIA
Chair, Communications Steering Committee, SNIA

The Market Needs a New Way to Look at Storage.

Pain Points

Physical Space

Available Power

Scaling Mismatch

Bottleneck Shuffle

Traditional storage architectures are in **trouble**.

Scaling requirements are **not met** with existing solutions

One CPU to many storage devices **creates bottlenecks**

These bottlenecks exist, we currently just shift where they reside

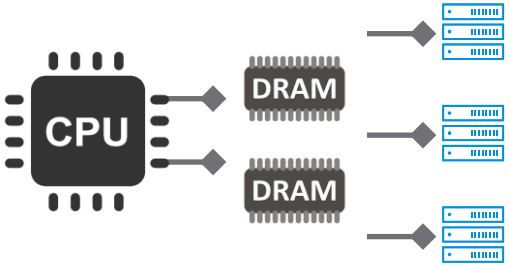

Technologies that '**compose**' these elements just move the bottleneck

A way to augment and support **without wholesale change** is needed

NVMe is Simply Not Enough.

More Lanes, More Traffic

- IDC predicts we will churn out 175 zettabytes of data in 2025
- **Storing** the data is easy
- **Working** on the Data will be hard
- **Solve** this with Computational Storage
- **DPUs** alone just *SHIFT* the bottleneck

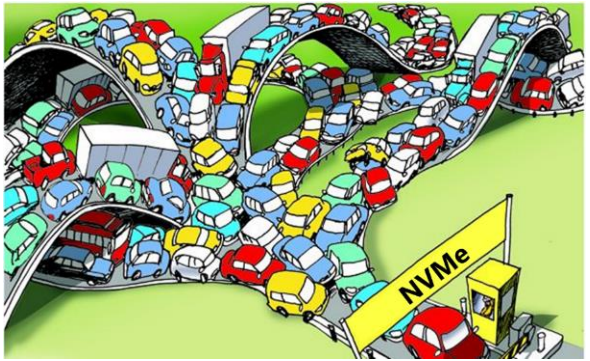

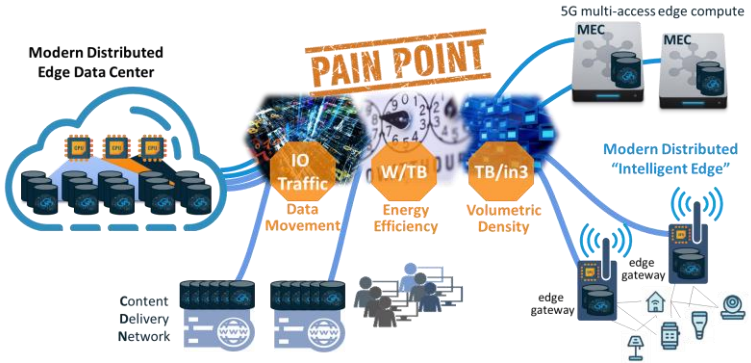

The Market Solution Using Computational Storage.

Value Add

- ✓ Distributed Processing
- ✓ Faster Results
- ✓ Lower Power
- ✓ Smaller Footprint

Scaling compute resources with storage provides access to results faster

Computational Storage resources 'offload' work from the overtasked CPU

Seamless architectures create new 'servers' with each storage device added

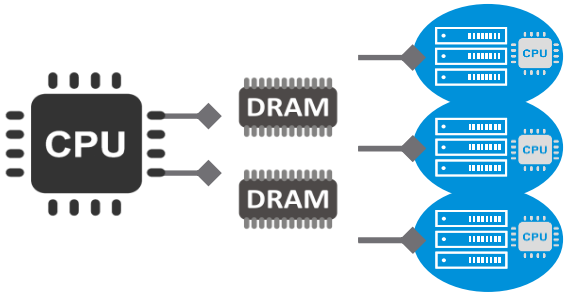

This **'Server in a Server'** Architecture provides value across many use cases

Additional CPU resources for the cost of Storage without added Rack Space

Computational Storage Needs are Now and the Future.

Pain Points - Before

- ✗ Physical Space
- ✗ Available Power
- ✗ Scaling Mismatch
- ✗ Bottleneck Shuffle

Computational Storage Benefits

- ✓ Smaller Footprint
- ✓ Lower Power
- ✓ Distributed Processing
- ✓ Faster Results

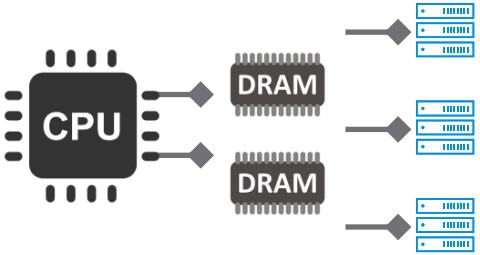

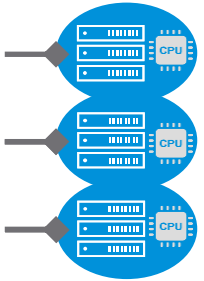

Innovative Computational Storage Uses.

SMART CITIES & AUTOMATION

EDGE GATEWAYS, MECs, SERVERS

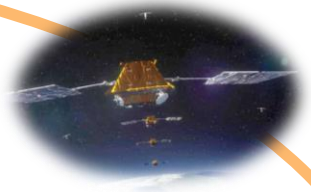

VMware Tanzu

TensorFlow

APACHE TRAFFIC CONTROL

Azure IoT Edge

hadoop

PyTorch

Parquet

bitfusion

Linux

docker

AWS IoT Greengrass Core

Keras

GREENPLUM DATABASE

mongo DB

Caffe2

DATACENTERS & CDN DELIVERY

AUTONOMOUS ANYTHING

Some Real-World Computational Storage Edge Use.

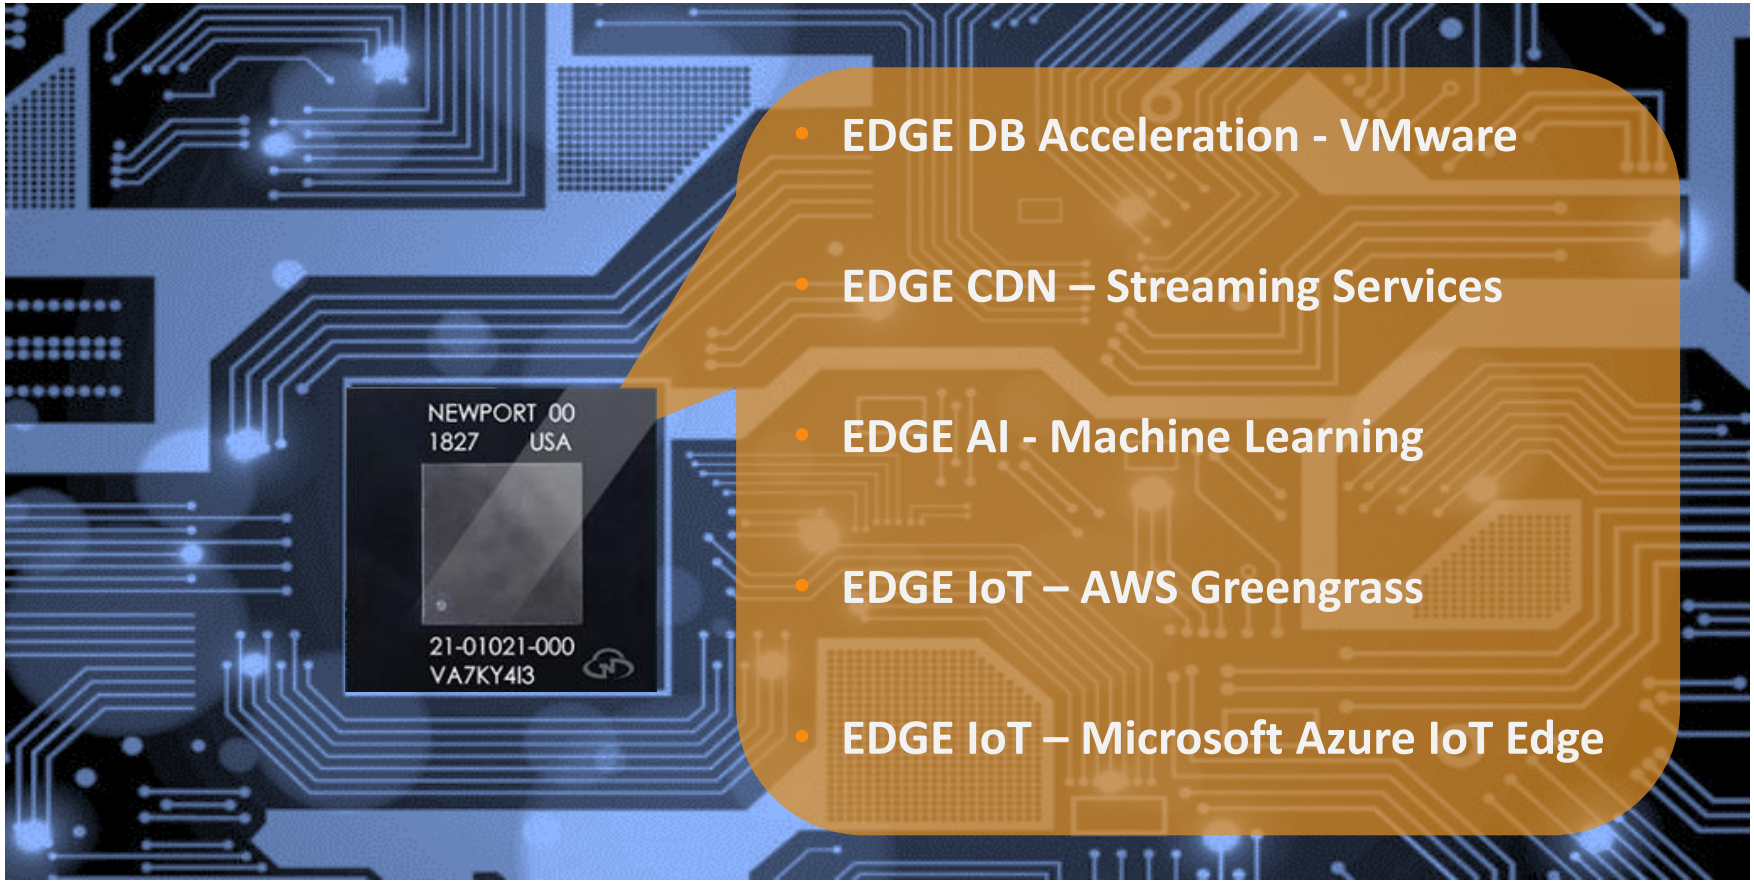

- EDGE DB Acceleration - VMware
- EDGE CDN – Streaming Services
- EDGE AI - Machine Learning
- EDGE IoT – AWS Greengrass
- EDGE IoT – Microsoft Azure IoT Edge

VMware Edge Analytics

- **Value**
  - Rack-scale system reduced to 6U
  - **70% Rackspace savings**
  - **50% overall system TCO (pending)**
- **Problem**
  - GreenplumDB Nodes are per CPU
  - Rack-scale needed for Nodes, not storage or processing
- **Solution**
  - Allocate Nodes per NGD Drive to reduce Server Count and need.
- **Implementation**
  - Migrate CPU Nodes to NGD Core
  - Resilience and Fail over still in place
  - Virtualized GPU across NGD Cores

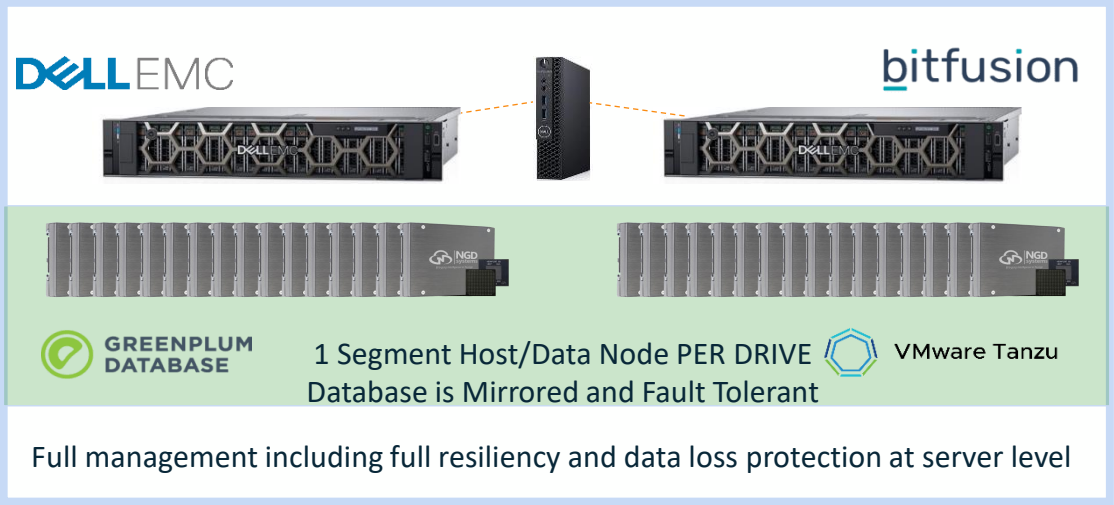

Integration of NGD Systems Devices to vSphere

LINK TO ONLINE DEMO – VMWorld 2020

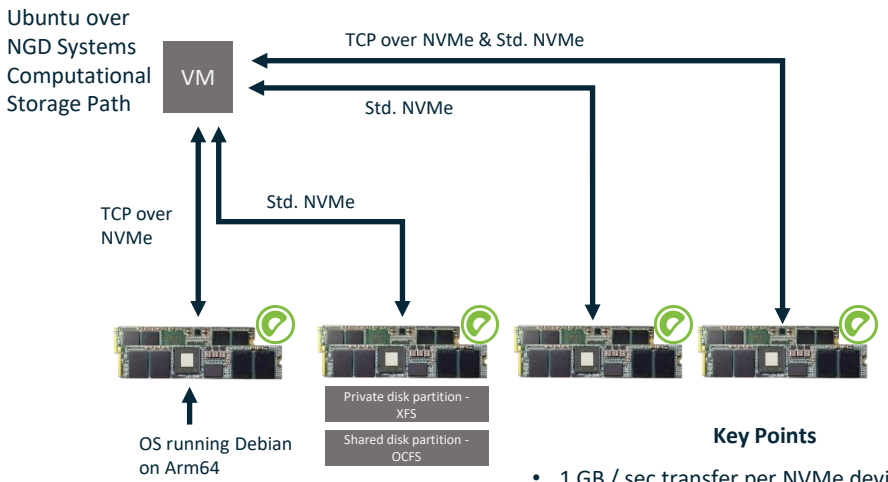

Key Points

- 1 GB / sec transfer per NVMe device
- 16TB capacity per device
- Simultaneous addressing as storage device & as remote compute node
- PCI passthrough allows native use by VMs
- Greenplum running on each node

1. VM Directpath IO for NVMe devices
    1. up to 15 into a VM
2. TCP connected jump box allows addressing of devices from network.
3. Two partitions
    1. one shared w/OCFS
    2. one dedicated to Greenplum

Computational Storage	Parallel Database with integrated Analytics	vSphere & Bitfusion
Embed compute with storage, offloading main server, improving performance on smaller systems by reducing data transfer to main system and enabling on-chip intelligence	Query across NVMe devices in parallel, making effective use of computational storage. Embedded analytics allowing analytics free of resources on the main system. Seamless replication of data to backup host.	Ability to offer Edge resiliency with vSAN, HA, FT. GPU acceleration for computational storage w/ Bitfusion. Effective use of limited host resources.

Content Traffic Control

- **Value**
  - Impressive Performance impact
  - **50% faster step performance**
  - **10% overall system improvement**
- **Problem**
  - Focus on Time to First Frame (TtFf)
  - Complex System to be pulled apart to find single step for impact
- **Solution**
  - Identified a Single Storage Instance, allocated Storage and processing to NGD Computational Storage
- **Implementation**
  - Migrated IP Geo Database from Network-attached HDD Storage to On-Router Local NGD Computational Storage Drives

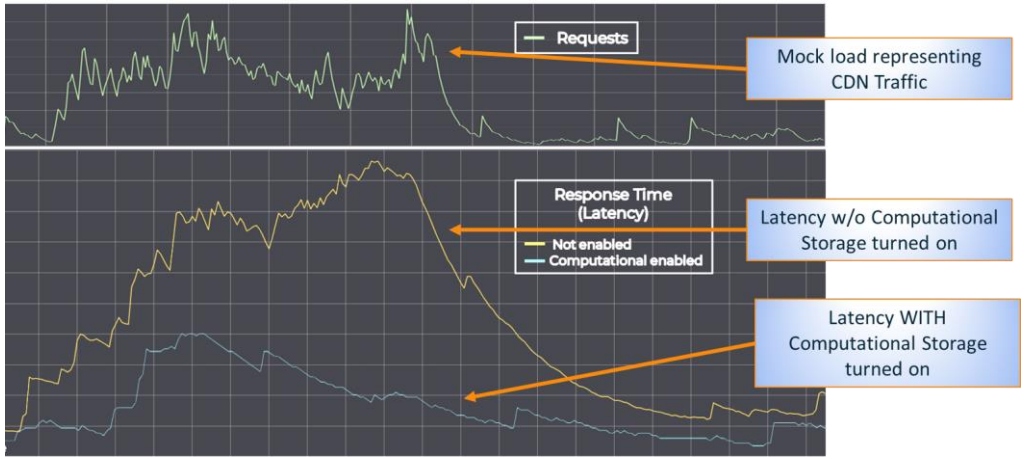

CDN Traffic Routing Offload – The Setup BEFORE.

CDN Clients

Step 1

Step 3

Step 4

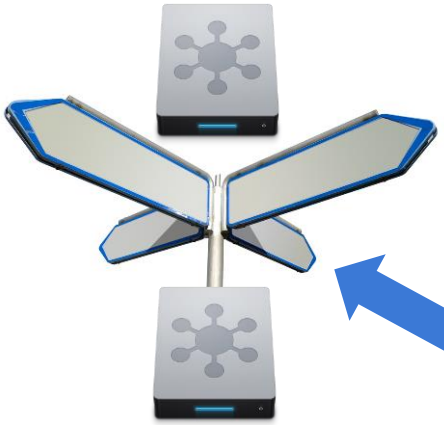

Traffic Router

Step 2 – Network-based Latency Issues are in this process step

Edge – Mid – ORIGIN Servers

IP Geo Database

Steps 1-4 Before Computational Storage

CDN Traffic Routing Offload – The Setup AFTER.

CDN Clients

Step 1

Step 3

Step 4

Edge – Mid – ORIGIN Servers

Step 2 – Localized Traffic increases efficiency

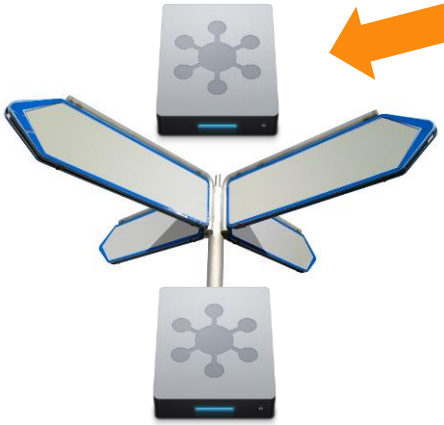

Traffic Router

Add Computational Drives to host IP Geo Database

By Removing this step, >50% of this step is saved, netting >10% OVERALL CDN Efficiency

Steps 1-4 with Computational Storage

AWS Greengrass IoT

- **Value**
  Certified for use with all AWS IoT Use Cases for Storage
- **Problem**
  Greengrass is a powerful Edge tool that is not serviced well for analytics locally
- **Solution**
  By migrating Greengrass and Lambda into NGD Drive, streamlined results are available
- **Implementation**
  Using an NGD Drive, Lambda functions usually run on embedded CPU are offloaded to Drive to accelerate system performance

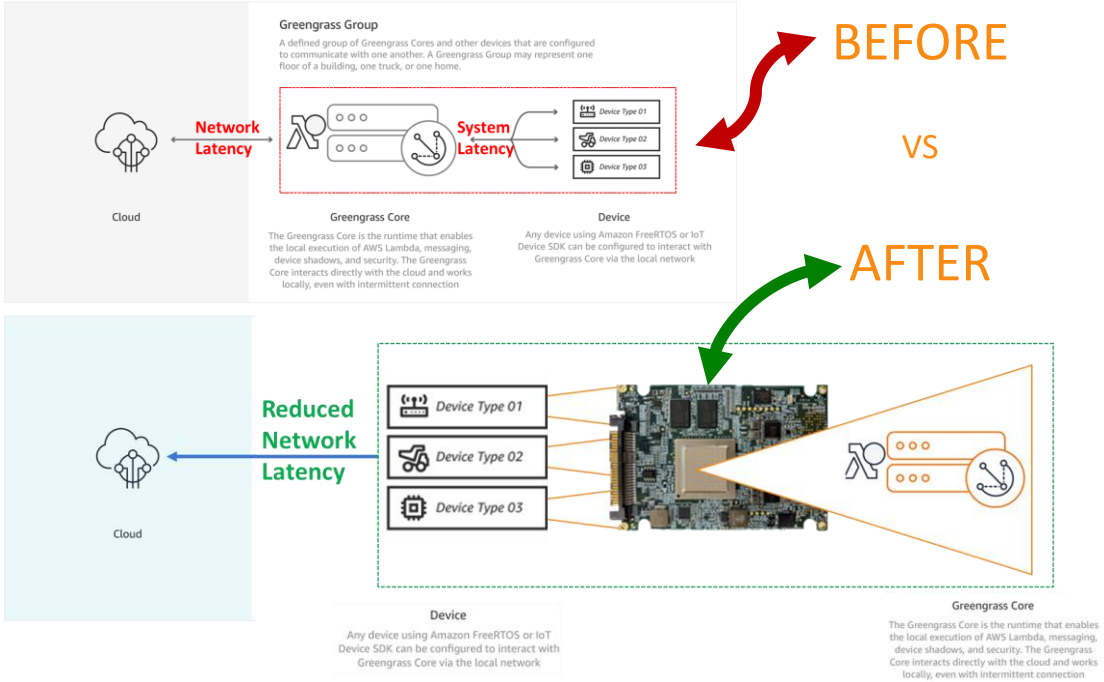

Microsoft Azure IoT Edge

LINK TO ONLINE VIDEO

- **Value**
  - Remove Local CPU requirements
  - Create Small Scale Platforms
  - Lower Power, Cost
- **Problem**
  - IoT Platforms require compute but cannot afford CPU/GPU/DPU in many small systems
- **Solution**
  - Allocate IoT Edge OS in NGD Drive to reduce CPU requirements
- **Implementation**
  - Run IoT Edge OS and link to Azure through NGD Drive, NOT host CPU
  - Allows for use of RaspberryPi vs Intel Xeon

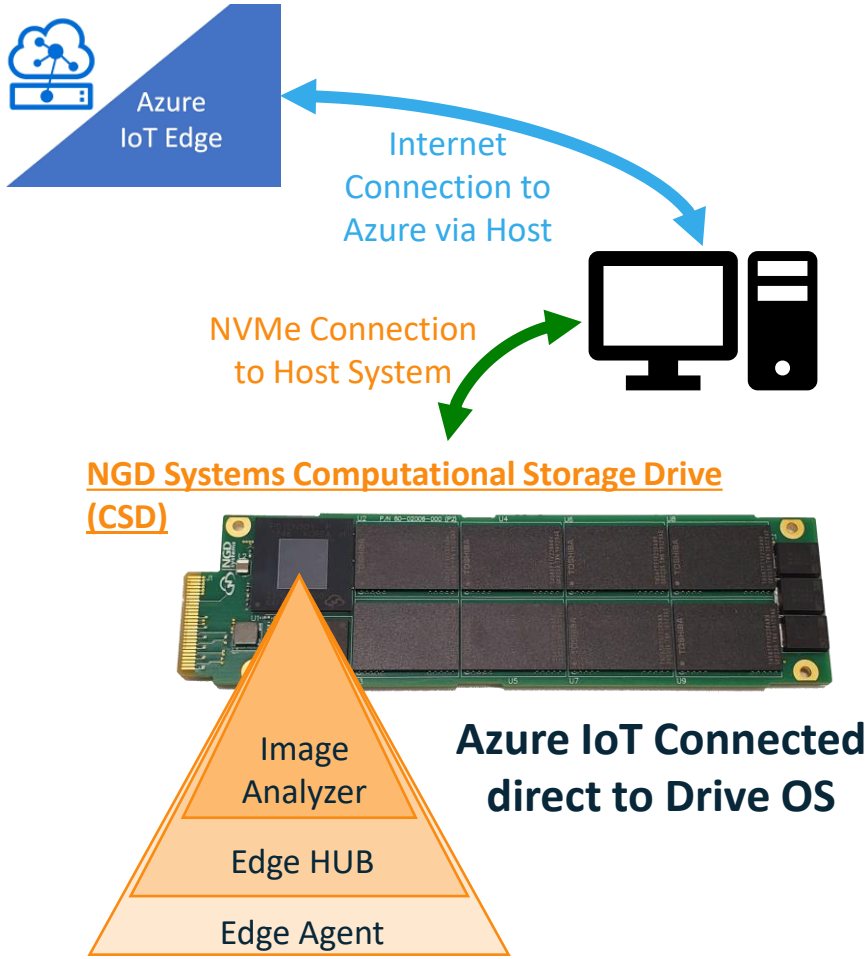

Microsoft Azure IoT – Running ON Drive.

HOST System Linux

Used for simply saving data

ON DRIVE Linux

Used for ALL Azure Interaction

NGD's Computational Storage device is powered by Azure Iot Edge

https://www.youtube.com/watch?v=D7Ab8zIi3kw

Demonstration shows use of Computer Vision Cloud-Based Application executed using only the NGD Systems Newport Drive OS, No Host Interaction with Cloud Required

Machine Learning At Scale – Not Just One Way.

- Four neural networks Evaluated
  - **MobilenetV2**
  - NASNet
  - **SqueezeNet**
  - InceptionV3Quad-core
- Tested with 24 CSDs
  - 32TB capacity each
- Training data stored on CSDs
- Using an AIC 2U-FB201-LX server
  - Intel® Xeon® Silver 4108 CPU
  - 32GB DRAM

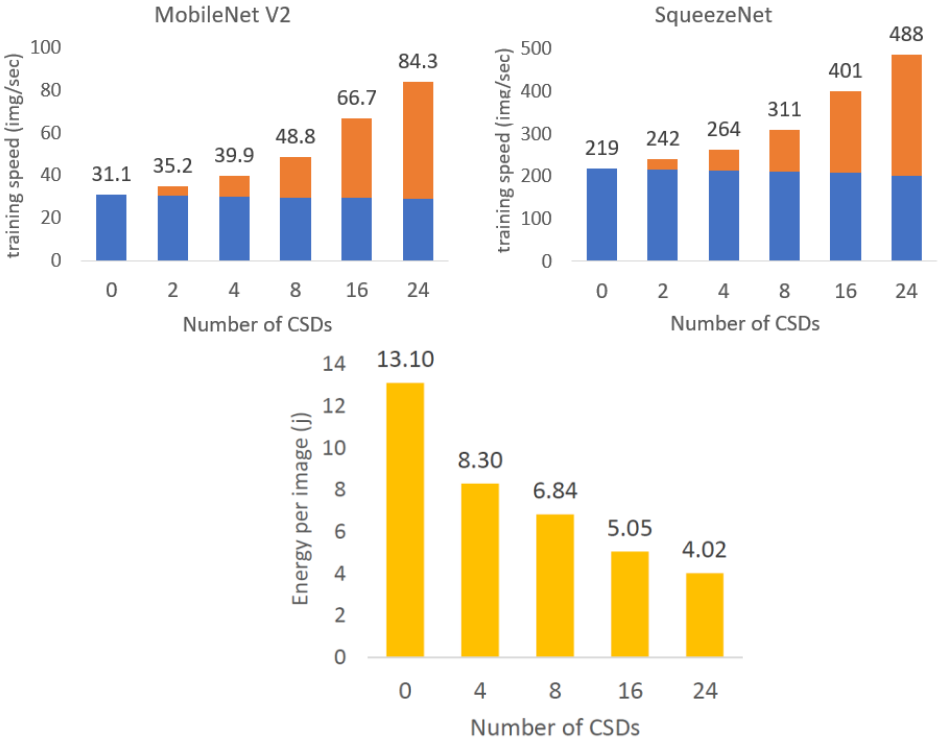

CSD Executed, Power Saved, GPU-like.

Definitive AI Support at the Edge

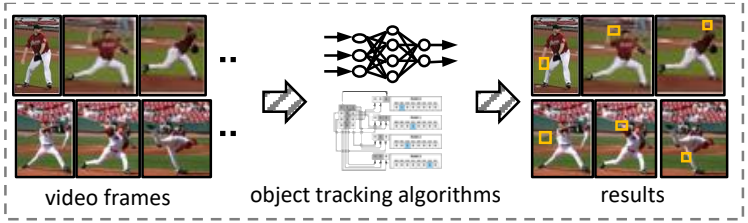

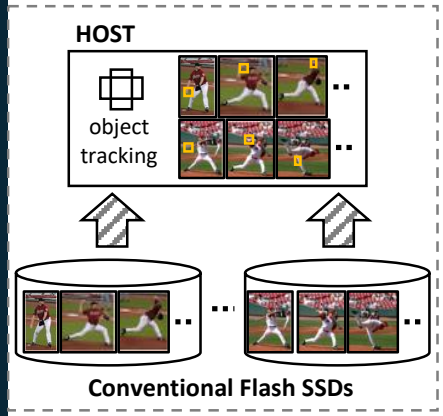

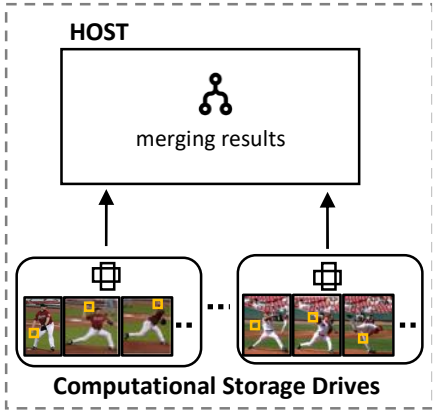

- Scalable solution
- Energy efficiency gains (~10x)

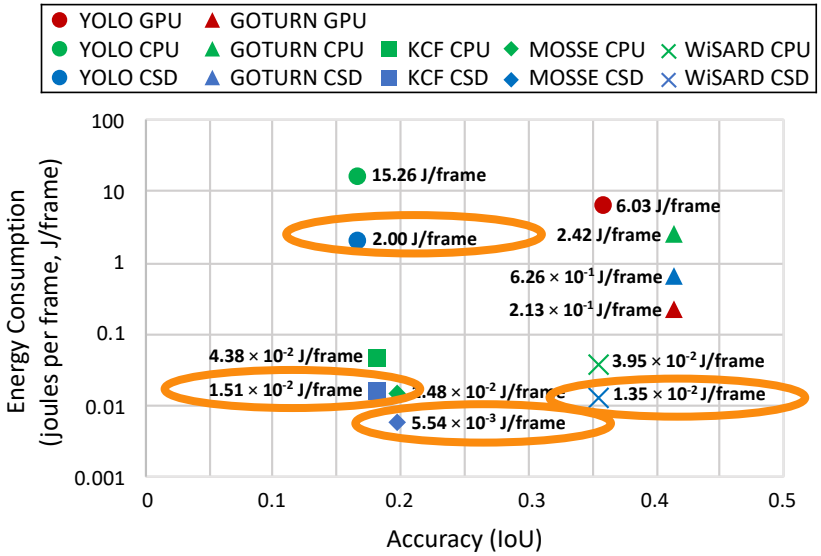

CSD Results Are equal in accuracy with Less Power

Computational Storage, Beyond the Edge.

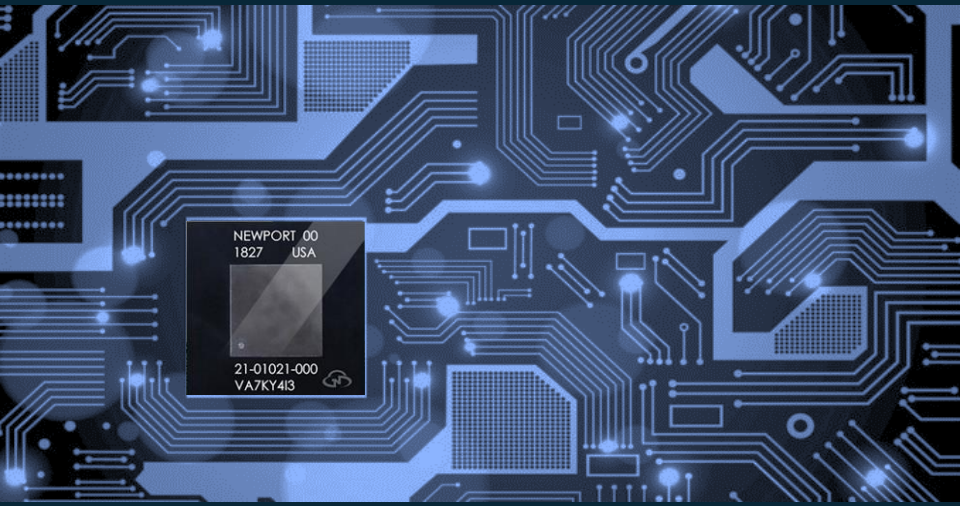

- eDiscovery – OCR Anything
- Compression Acceleration – GZIP
- AI Inference in the Datacenter – FAISS
- Data Search – Elasticsearch
- Distribute Processing – MongoDB
- Even more On our SITE!!

eDiscovery – OCR Anything!

DATA STORAGE SCIENCE

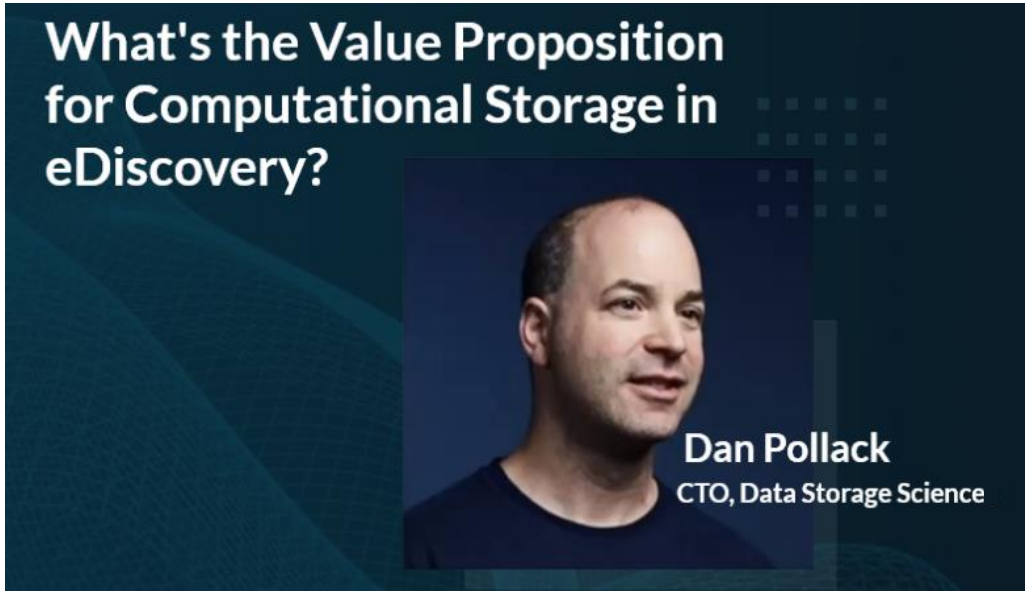

Gzip Performance With Computational Storage.

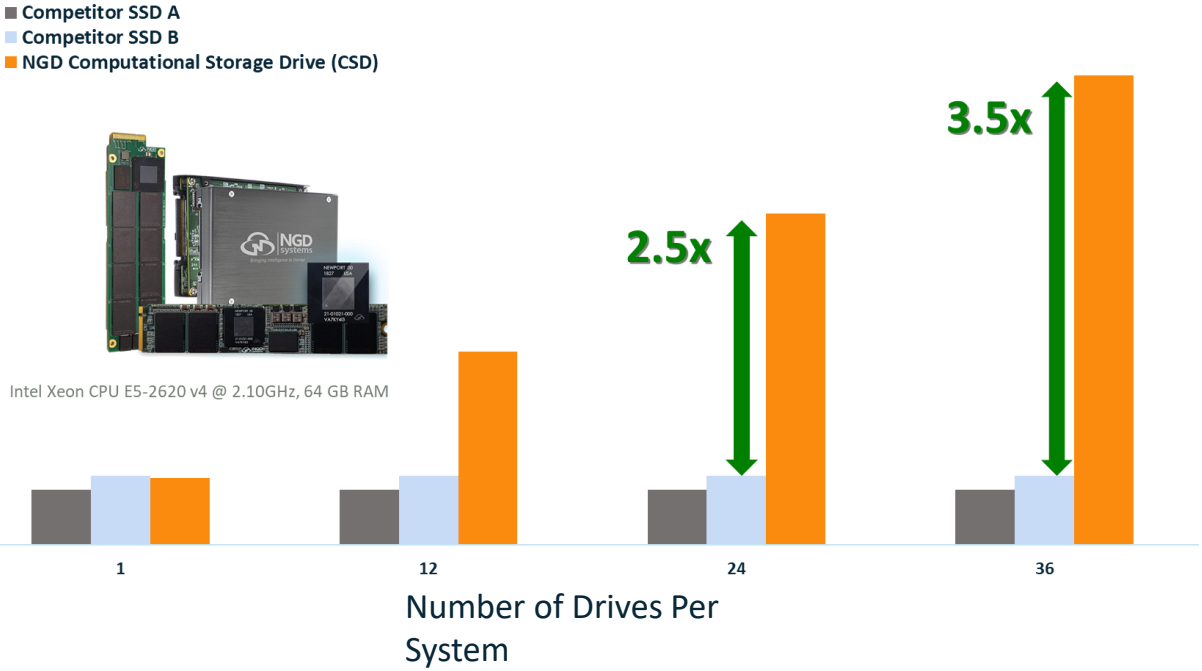

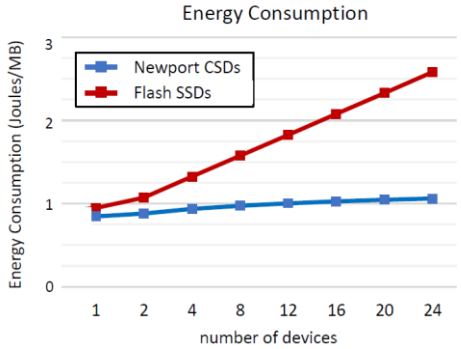

Nearest Neighbor Search

- **Value**
  - **Load Time Reduced > 95%**
  - **Search Time Reduced > 60%**
  - **Power Savings of > 60%**
- **Problem**
  - Databases growing at exponential rates
  - Load and Search time key blocks in getting results
  - 10M → 1 Billion → 1 Trillion
- **Solution**
  - Load Time Reductions due to NGD Drive Running Offload of Application code
- **Implementation**
  - Server level implementation of scaled drive count and modified code to maximize value to customer

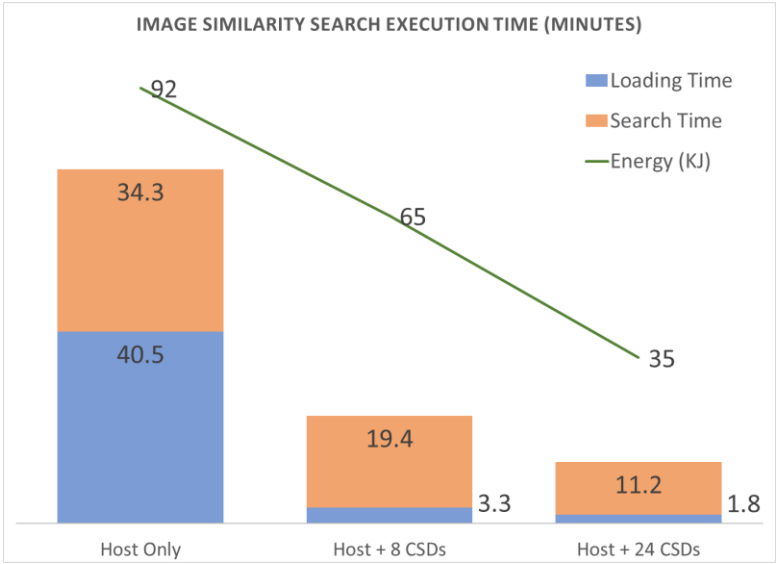

Search Efficiently on Computational Storage Devices.

elasticsearch

- **Customer choice of implementation**
  - Computational Storage Drives (CSDs) Only
    - No CPU/Memory utilized for tasks
  - Hybrid Solution of Host and CSDs
    - Maximize performance and power

- **Hybrid Model improves net TCO**
  - Easy to deploy
  - Improved performance
  - Reduced Memory footprint

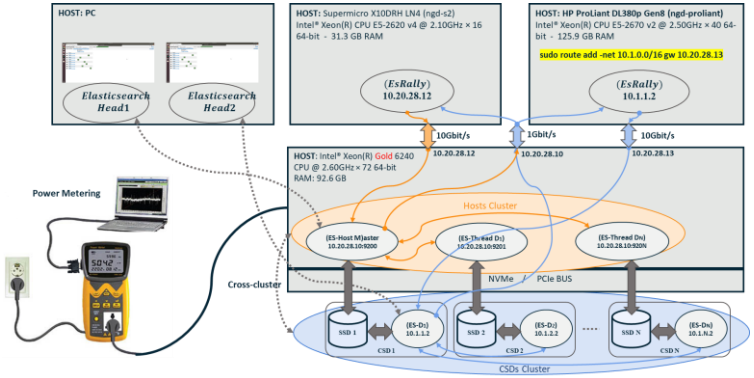

Intel® Xeon(R) Gold 6240 CPU @ 2.60GHz × 72 & 10Gbit/s Network

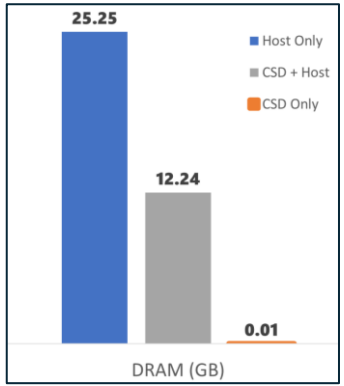

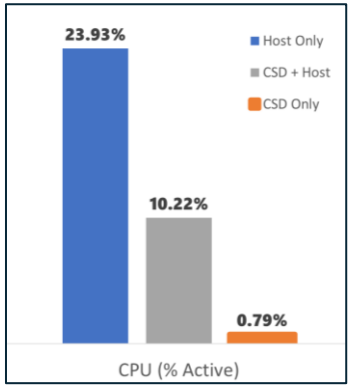

Database Management is About Cost and Performance

- **ROI quickly realized due to dramatic savings on Cap-Ex and Op-EX when using Computational Storage Drives (CSDs)**
  - **CapEx Reduced by 40%**
  - **OpEx Reduced by 30%**
- **Replication and Sharding come at a High Cost in Historical Configurations**
  - Many servers, multiple Hosts, Wasted Storage
- **Applying Computational Storage Drives to act as MongoDB Nodes and Shards at the Drive**
  - Each drive can act as a Host, Scaling Nodes as you add Capacity

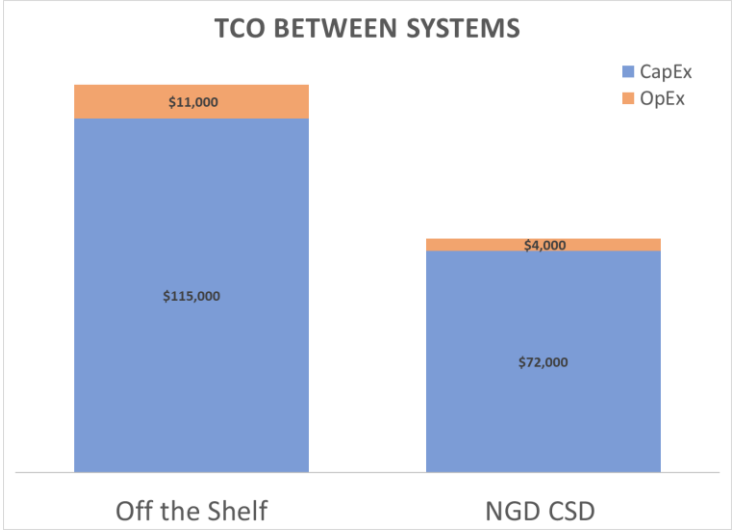

Computational NVMe Products at a Glance.

- Large breadth of **SSD** solutions and capacity options
- Leading **W/TB** Energy Efficiency
- Industry's only 16-Channel M.2
- Largest capacity NVMe U.2

Form Factor	Availability	Raw Capacity TLC (TB)	QLC Capacity (early 2021)
M.2 22110	NOW	up to 8	>12
U.2 15mm	NOW	up to 32	>60
EDSFF E1.S	NOW	up to 12	>16
E3.S	Pending Spec	up to 32	>60
LFF 3.5 Drive	Request	up to 64	>100

M.2

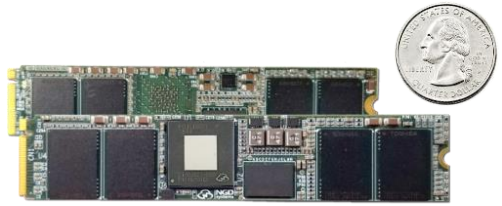

U.2

EDSFF

The Data Lives on Storage.

Why Not Work on it There?

One Host Many Drives

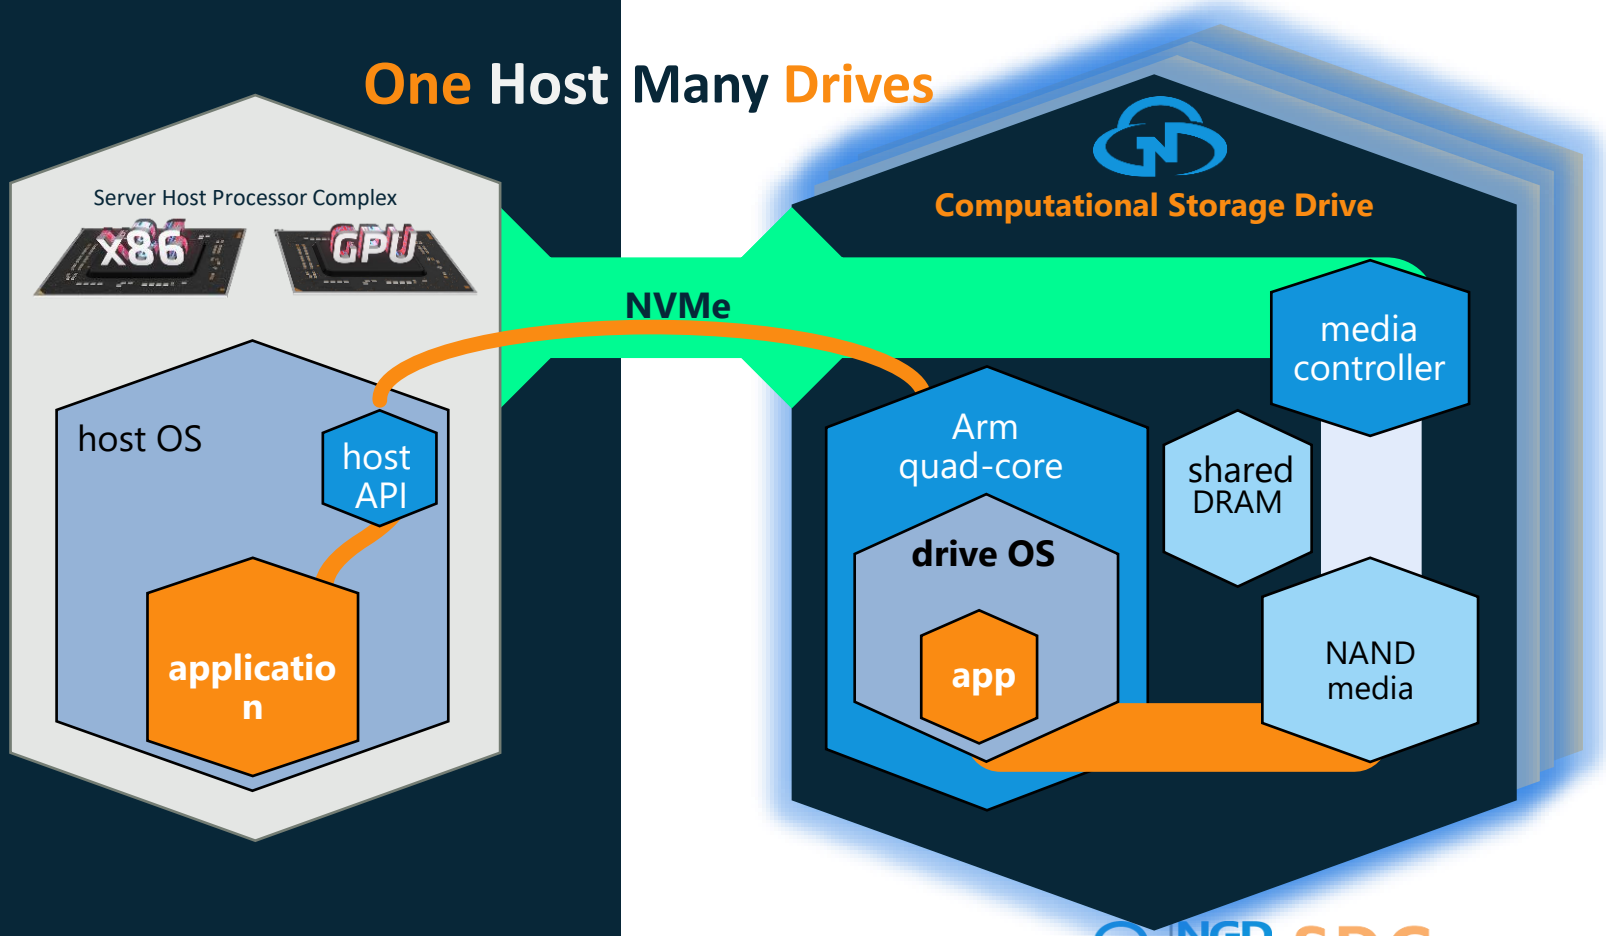

Delivering a Wholistic NVMe Storage Solution.

SINGLE ASIC Design

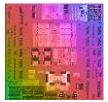

Complete Management

- Modular firmware
- Efficient algorithm
- Flash characterization
- Flash agnostic SLC/MLC/TLC/QLC

Optimized Hardware

In-Situ Processing Stack For Computational Storage

- Full fledged on drive OS — Linux
- Light virtualization — docker
- Quad-core 64-bit application processor — ARM
- Hardware acceleration (i.e. AES 256 XTS/GCM)

Use as SSD or Activate CSD (Computational Storage Drive)

Seamless Programming Model

Market Leadership Driven by Customer Feedback.

Computational Storage

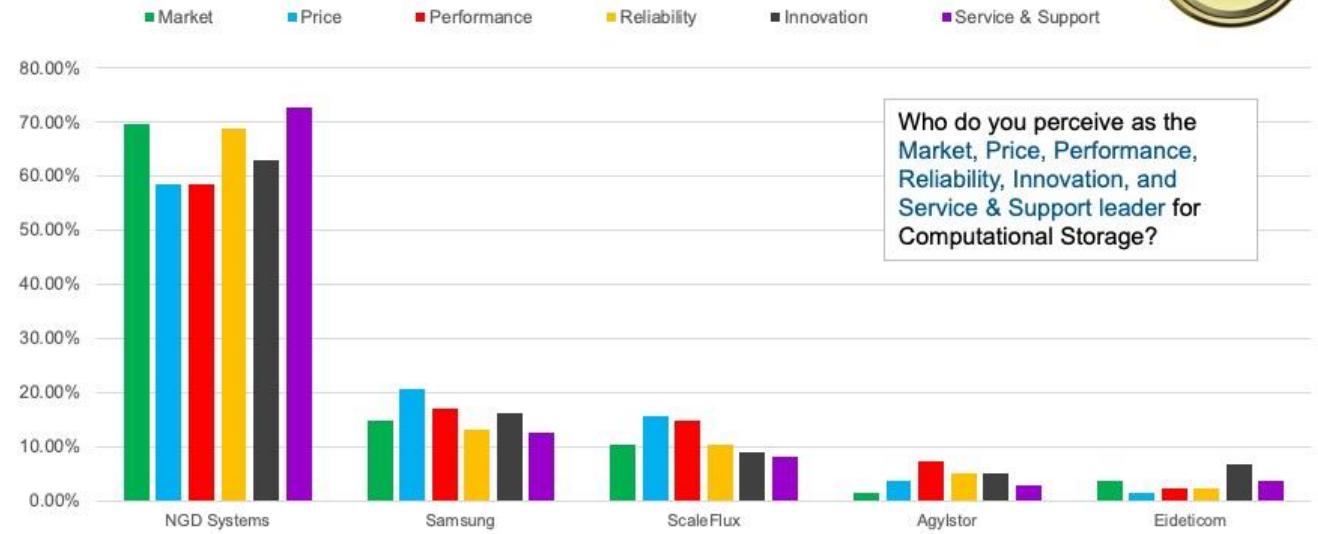

2020/21 IT Brand Leader Survey 12

Computational Storage is Not New, NGD Got it Right.

ACM Transactions on Storage - Special Section on Computational Storage

	Processing Resources	Programming Model	Media Access	Communication to Host
Research Prototypes				
• Do et al [15]	• 32-bit embedded processor (< 200MHz)	• Firmware (C code)	• Raw access	• SAS
• Kang et al [41]	• 32-bit embedded processor (< 200MHz) (shared with firmware)	• Firmware (C code)	• Raw access	• SATA
• Seshadri et al [67]	• 32-bit embedded processor (< 200MHz)	• Primitive OS (C code)	• Raw access (DRAM emulator)	• PCIe (RPC protocol)
• Jun et al [40]	• FPGA	• Bluespec [53]	• Raw access	• PCIe (RPC protocol)
• Gu et al [28]	• 32-bit processor (< 800MHz) • Pattern matcher HW	• Flow-based model (C/C++)	• Raw access	• PCIe
Commercial Products				
• Samsung [49]	• FPGA	• Bare metal	• Raw access	• PCIe
• ScaleFlux [66]	• FPGA	• Bare metal	• Raw access	• PCIe
• Eideticom [19]	• FPGA	• Bare metal	• Raw access	• PCIe
• NGD [74]	• 64-bit processor (> 1.5GHz) • FPGA	• Full-fledged OS (Linux)	• Clustered file system • Disk file system • Raw access	• PCIe (RPC protocol)

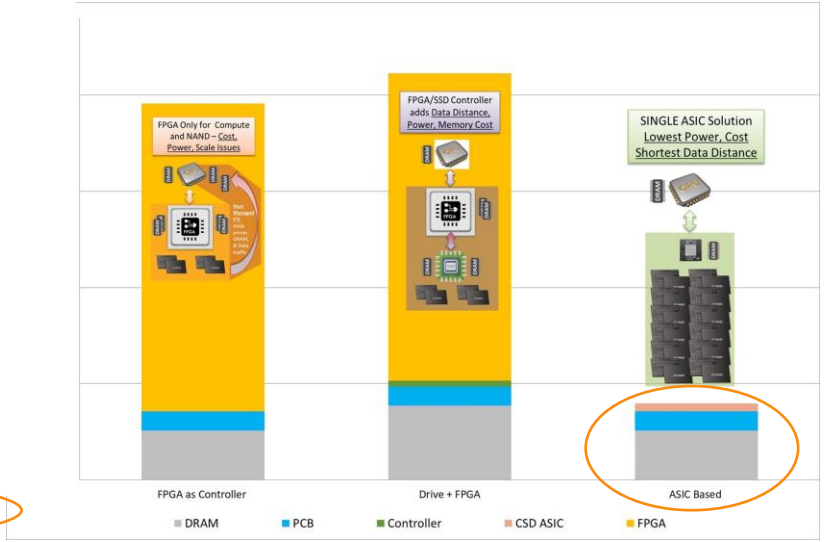

Computational Storage Recognized as Transformational!

Hype Cycle for Storage and Data Protection Technologies, 2020

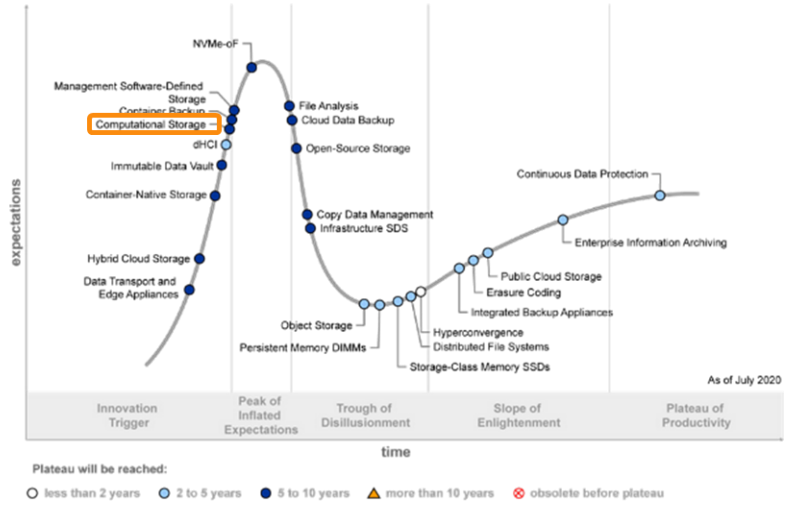

Source: Gartner
ID: 441602

Hype Cycle for Compute Infrastructure, 2020

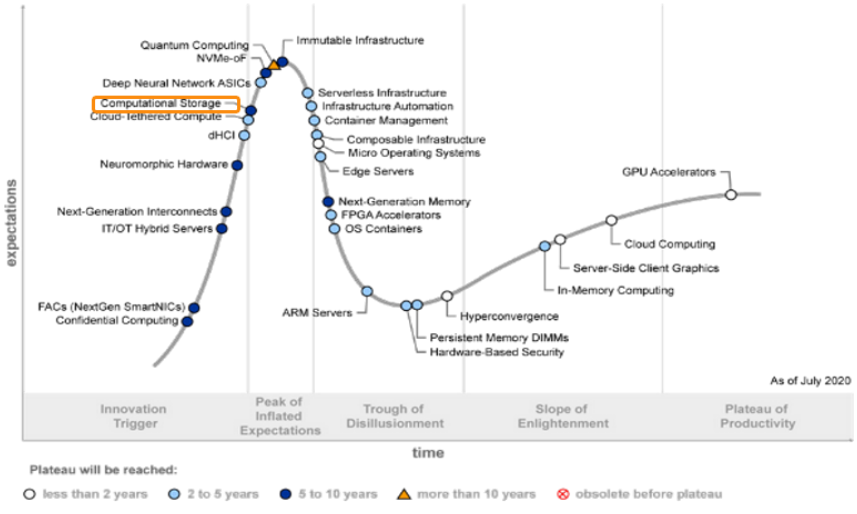

Source: Gartner
ID: 448100

Gartner COOL VENDOR

"Business Impact Areas (per Hype Cycle):
There is both a cost and time factor involved in shuffling terabytes of data around. CS can provide **material performance benefits** to data-intensive applications, especially in edge computing. Combined with its low power footprint, **CS increases the performance-per-watt ratio**, therein **decreasing power consumption costs for applications at the edge**." – Jeff Vogel, Julia Palmer – Gartner Research

Innovation at Heart.
Customer Solutions Delivered.

- **Computational Storage**
- **Data Management**
- **Flash Management**
- **High-Capacity Management**
- **Power Management**

Broad & Defensible Patent Portfolio
32 Granted, **8** Pending

NSF / SBIR Grant awardee 2016, 2017, 2018, 2019, 2020

How Can We Help?

Please take a moment to rate this session.

Your feedback matters to us.