
KIOXIA
WEKA
NetApp
G2M
RESEARCH
The Explosion in Imagery from Radiometry, Cryo-EM, and Other Imaging Technologies: Can Storage Keep Pace?
Multi-Vendor Webinar
Tuesday Nov 9, 2021

G2M Research Introduction and Ground Rules

Mike Heumann
Managing Partner, G2M Research

Webinar Agenda

9:00-9:05	Ground Rules and Webinar Topic Introduction (G2M Research)
9:06-9:35	Sponsoring Vendor presentations on topic (10 minute each)
9:36-9:43	Panel Discussion Question #1
9:44-9:44	Audience Survey #1
9:45-9:52	Panel Discussion Question #2
9:53-9:53	Audience Survey #2
9:54-10:00	Panel Discussion Question #3
10:01-10:08	Audience Q&A (8 minutes)
10:09-10:10	Wrap-Up

The Explosion of Imaging Technologies

- The number of imaging technologies in the medical and bio/pharmaceutical world continues to grow in scope:
  - Cryo-Electron Microscopy
  - Computerized Axial Tomography (CAT)
  - Magnetic Resonance Imaging (MRI)
  - Positron Emission Tomography (PET)
  - Single-Photon Emission Computed Tomography (SPECT)
- The storage for these various imaging sources in AI/ML learning instances often take petabytes/tens of petabytes

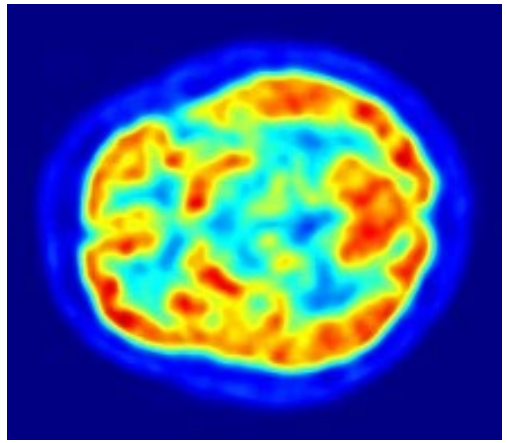

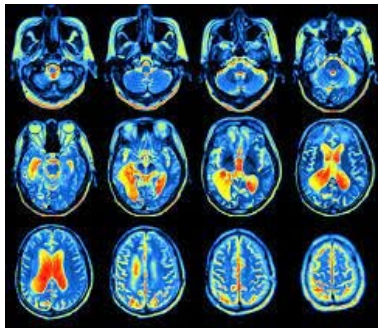

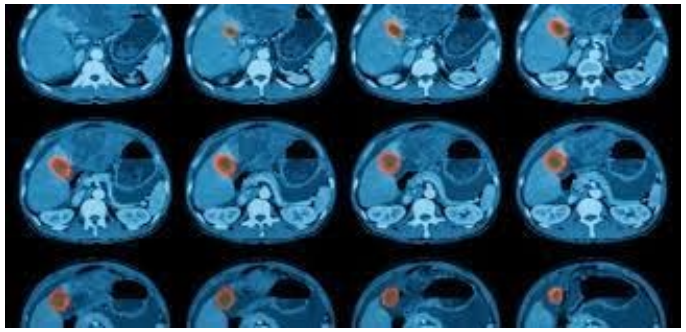

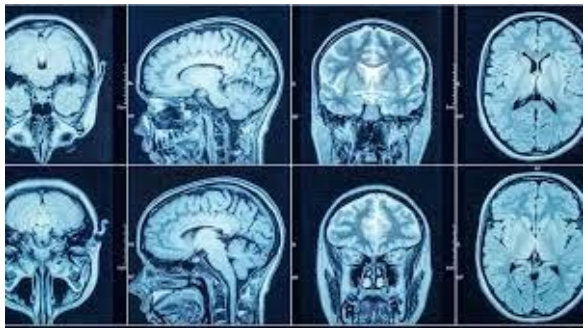

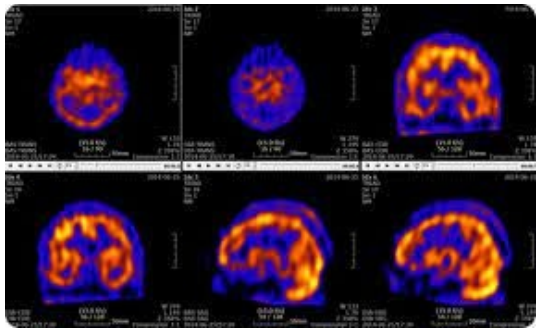

Storage Architectures for Imaging

- As the size of storage pools for medical/bio/pharma imaging continue to grow, building responsive storage architectures becomes more critical
- Chief storage requirements include:
  - Performance (bandwidth, latency, latency consistency)
  - Multi-tiering (including on-premises and cloud storage pools)
  - Data management (what data is where?)

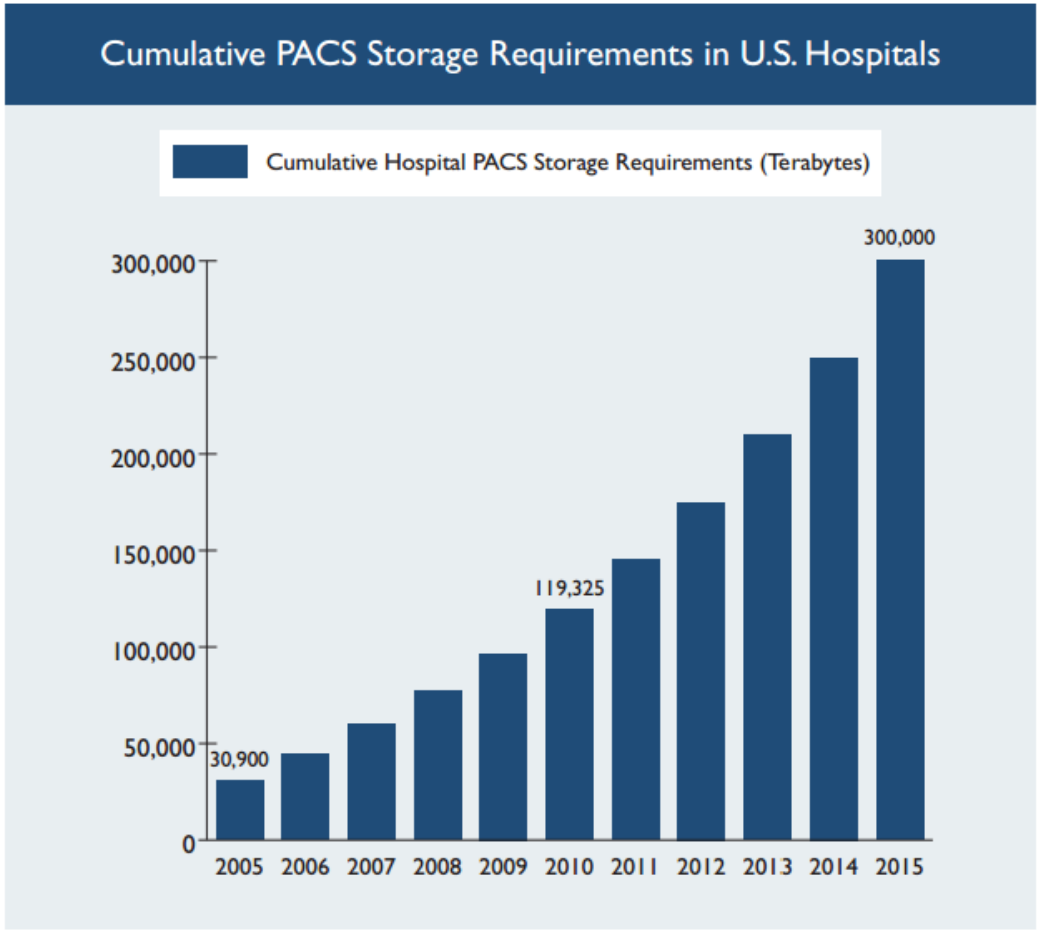

Source: Frost & Sullivan

Panelists

KIOXIA

Matt Hallberg
Sr Product Marketing Manager
www.kioxia.com

WEKA

Shimon Ben-David
Field Chief Technical Officer
www.weka.io

NetApp

Esteban Rubens
Principal, Healthcare AI Practice
www.netapp.com

G2M RESEARCH

Mike Heumann
Principal Analyst
www.g2minc.com

KIOXIA

Matt Hallberg

Senior Product Marketing Manager

www.kioxia.com

Storage Use Case: High Performance Nano-Scale Brain 3D Visualization

High performance SSDs enable interactive 3D and 8K visualization of brain images

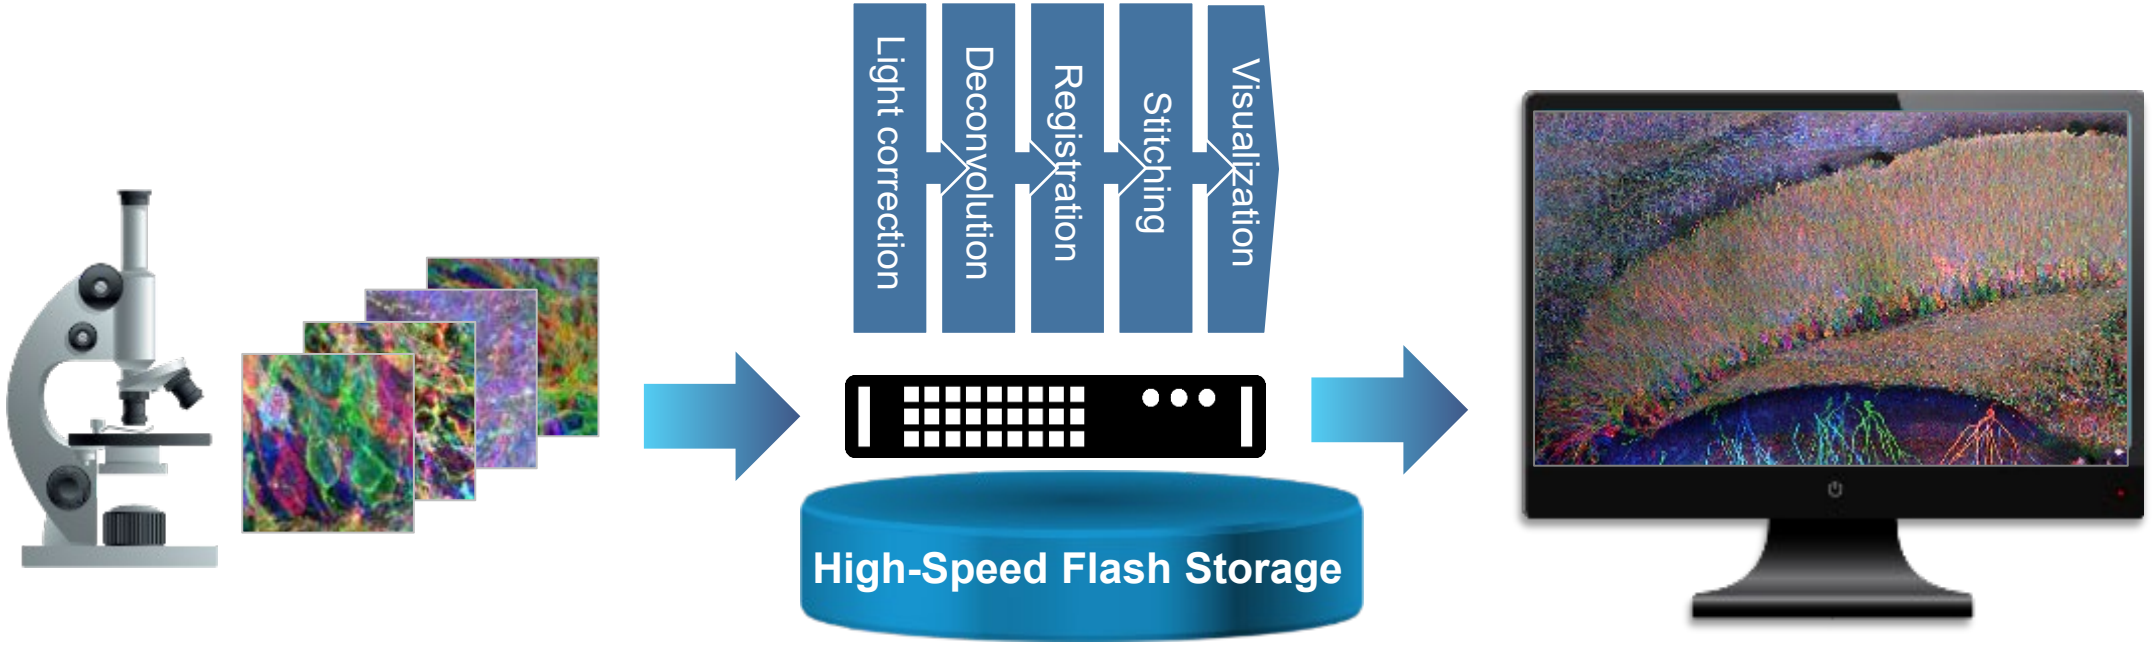

Image capture using Expansion Microscopy

Flash-based high-speed data processing system

Real-time 3D exploration of whole brain images

*** Microscopy images courtesy of Shoh Asano, Edward Boyden, MIT Media Lab**

Additional Information: https://www.media.mit.edu/articles/mapping-the-brain-at-high-resolution/

8K/3D High Capacity Interactive Visualization of Nano-scale Brain Neuron Demo

8K Interactive Visualization of Brain

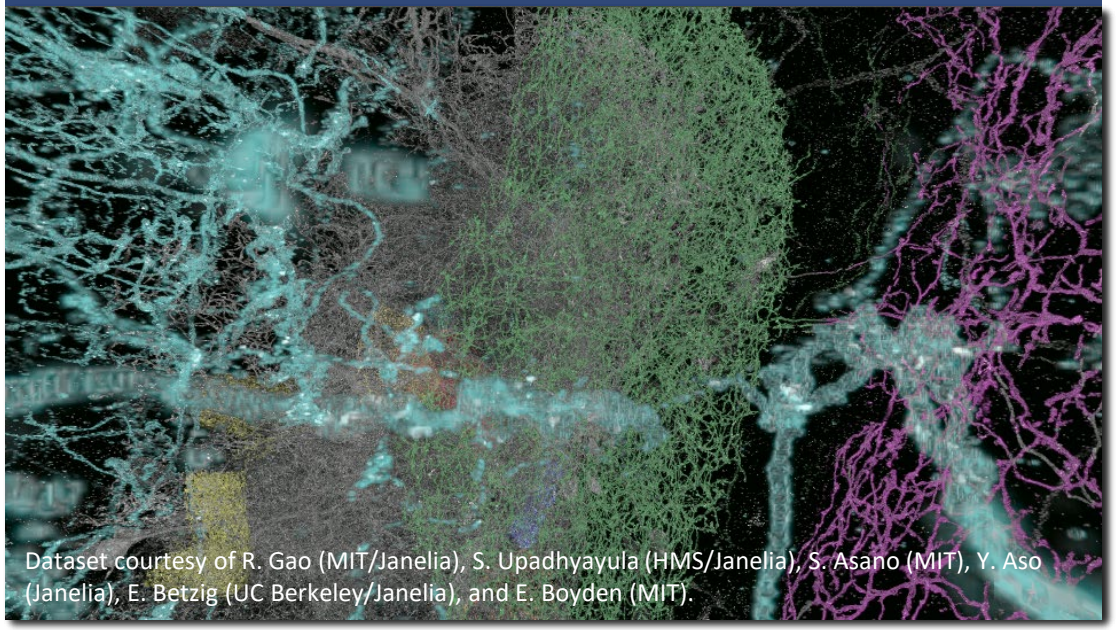

Dataset courtesy of R. Gao (MIT/Janelia), S. Upadhyayula (HMS/Janelia), S. Asano (MIT), Y. Aso (Janelia), E. Betzig (UC Berkeley/Janelia), and E. Boyden (MIT).

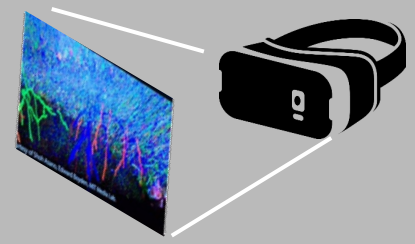

HMD Interactive Visualization

Storage Performance Requirement

Current: ~1PB, ~10GB/s

Future: >EB, >TB/s

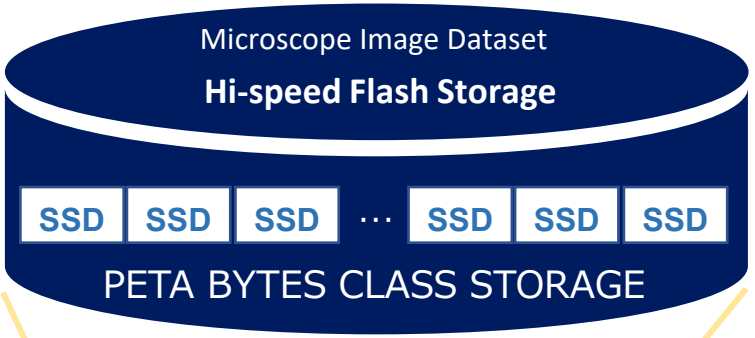

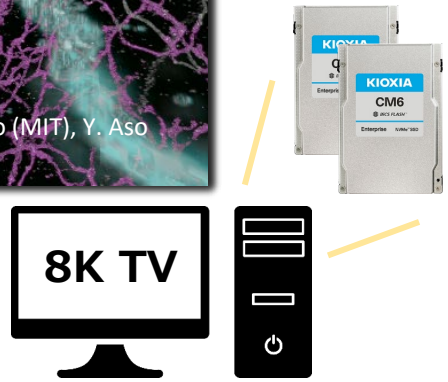

NVMe-oF™ over Ethernet

External JBOF

KUMOSCALE™

Real-time Image Analysis | **Interactive Visualization** | **Real-time Data Matching** | **Real-time Database Query**

NVMe-oF is a trademark of NVM Express, Inc.

Interactive 3D Visualization of Terabyte-sized Nano-scale Brain Image in 8K Resolution

Overview

- 8K (7,680 x 4,320 Pix) visualization that accelerates neuroscience research
  - High-capacity, high-performance SSDs (storage) enable processing of terabyte-class 3D microscope images
- Through high resolution and hi-speed rendering, approximately 1 trillion voxels of 3D image datasets can be interactively visualized

➔ Accelerates the research of nano scale and macro scale structures of neurons

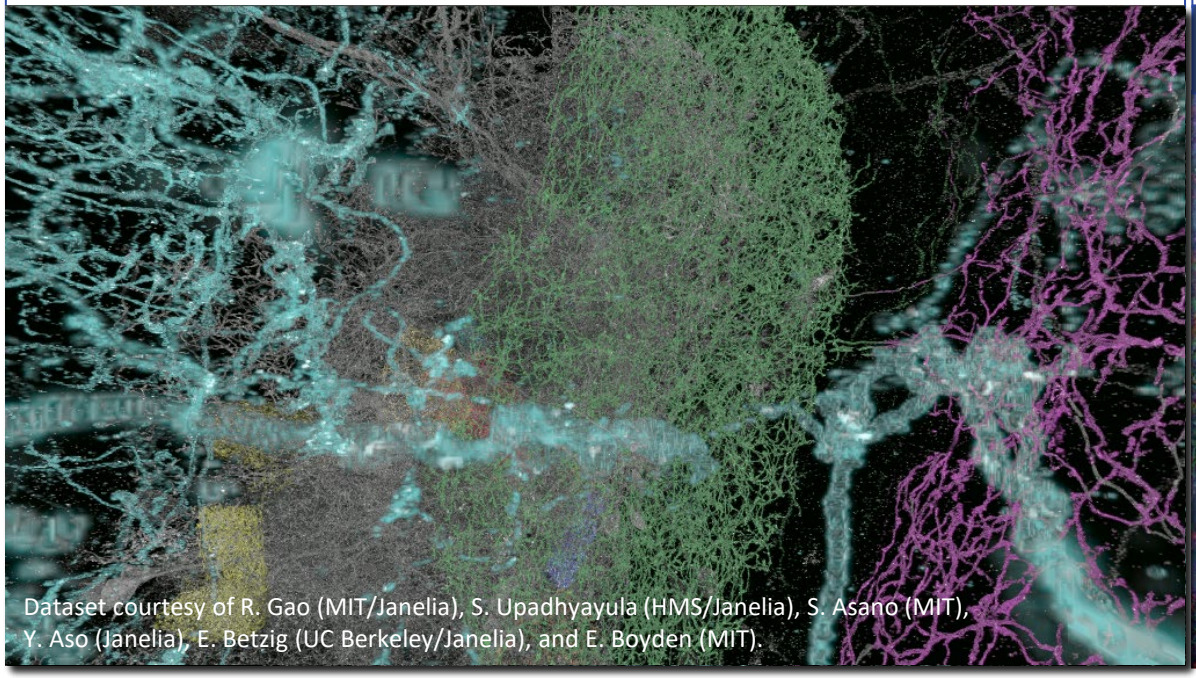

Dataset courtesy of R. Gao (MIT/Janelia), S. Upadhyayula (HMS/Janelia), S. Asano (MIT), Y. Aso (Janelia), E. Betzig (UC Berkeley/Janelia), and E. Boyden (MIT).

Why 8K?

Accommodates to various spatial scales within the brain

- Neurons link in milimeter scale: (>10^3 µm)
- Small processes are nano scale structure: (<10-1µm)

8K (Super Hi-Vision) enable visualization in scales of 4-digit order

- 8K ≒ 10^4

Example:

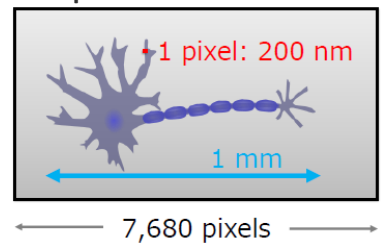

Rendering

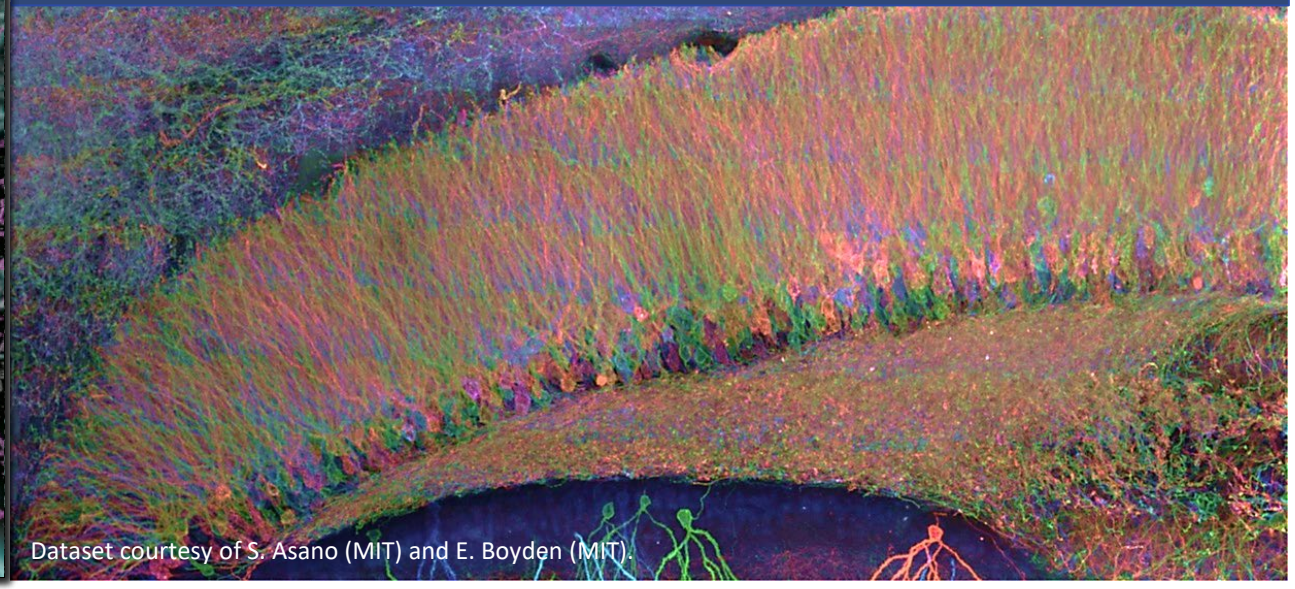

Dataset courtesy of S. Asano (MIT) and E. Boyden (MIT).

KIOXIA

Y. Bando, K. Hiwada, M. Kanaya, T. Ito, S. Asano, M. Bove, and E. Boyden.
Interactive 3D Visualization of Terabyte-sized Nanoscale Brain Images at 8K Resolution, Neuroscience 2017 Abstract 531.11, 2017.

Why Local Storage Matters

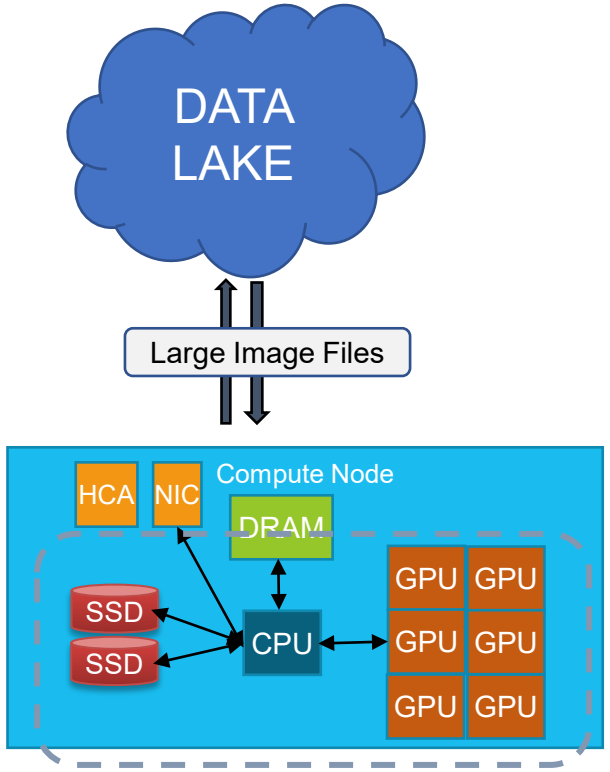

- The training phase of machine learning is the most-resource intensive set of operations
  - Datasets are growing at a fast pace, MRIs can reach up to TBs, and training sets can be composed of thousands of images
  - Whether you are running on RAM or on local storage, the local storage needs to be able to handle reads and writes in blazing fashion and with little impact to overall latency
  - Moving the data in, moving the data out, and checkpointing all need to be completed quickly to minimize the idleness of the GPUs
- PCIe® 4.0 SSDs' noticeable benefits with file copying and other I/O tasks versus PCIe 3.0 SSDs
  - For sequential workloads
    - Up to 7000MB/s on reads
    - Up to 4200MB/s on writes
  - For random workloads
    - 1M+ IOPS on random reads
    - 70K+ on random writes
  - PCIe 4.0 SSDs are also able to take advantage of 3D flash's higher densities, allowing for up to 30TB in a 2.5" SSD

Why Remote Storage Matters

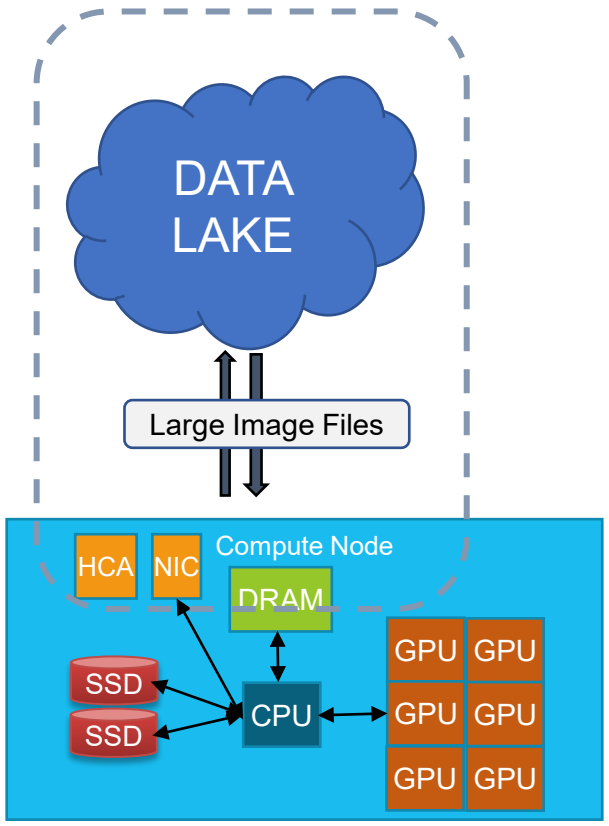

- Networking speeds, technologies, and topologies have greatly advanced over the past few years
  - 200GbE NICs
  - RoCEv2 / RMDA over Ethernet
  - NVMe® over Fabrics (NVMe-oF™) deployments
  - Etc.
- Data sets are comprised of thousands of files all of which need to be sent and received to and from the compute node to minimize downtime
- The storage behind the NIC(s) should be optimized for sequential performance to send data for processing and receive processed data. The faster the offload can occur, the faster the local storage can send the data to the GPUs
  - NVMe-oF deployments show their strength here by improving performance and reducing latency
  - Using NVMe SSDs for staging “warm” data vs cold data (cheaper storage) optimizes the spend on remote storage
  - GPUDirect® technology can also take advantage of the remote NVMe SSDs

QCT / AMD / Broadcom / KIOXIA PCIe® 4.0 SSD Demo @ Microsoft Ignite 2019

G2M RESEARCH

23.4 GiB/s

PCIe 3.0

52.1 GiB/s

PCIe 4.0

Benchmark: Storage I/O for Virtualized App

- 100% sequential 129KiB reads
- More than double the performance over PCIe 3.0 configuration
- Each drive delivers over 7 GiB/s!

Microsoft Windows Server 2019, Azure Stack HCI 2 node cluster

- Build 177763.775 Data Center Edition
- Hyper-V and Storage Spaces Direct
- Two-way mirror volumes
- VM Fleet with diskspd 2.0.21a
- FIO

QCT D43K-1U Server(2)

- 2 x AMD Epyc 7742
- 512GiB DRAM
- 4x KIOXIA CM5 PCIe 3.0, 7.68T SSD
- 4x KIOXIA CM6 PCIe 4.0, 7.68T SSD
- Broadcom NetXtreme™-P2100G 200Gbps NIC

There's more to come: Ethernet Bunch of Flash (EBOF) over NVMe-oF™

G2M RESEARCH

Construction

- (6) x 200 Gbps high speed network connectivity
- (24) bays @ 830K / drive, 20M IOPs per fully loaded system

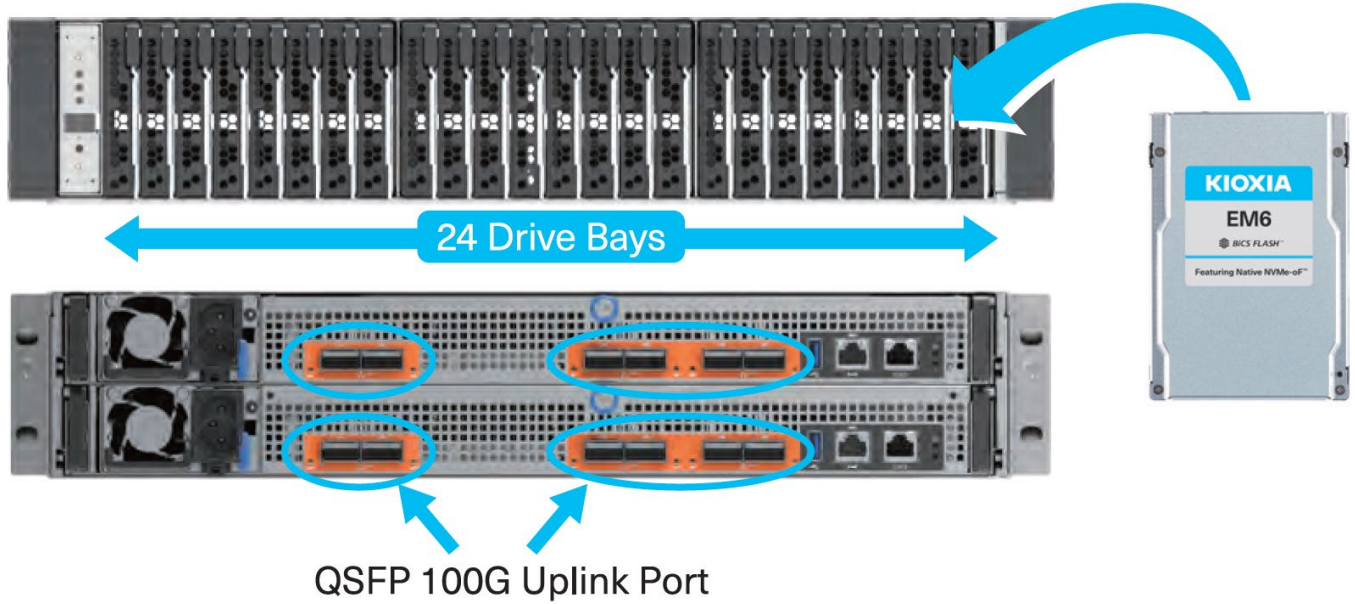

Application

- Data Centers
  - Just add EBOF(s) to existing RoCEv2
- High Performance Computing
  - Burst buffer storage
  - DRAM offload to storage for IO phase after computation
  - Lustre®
- GPU Direct for AI/ML
  - GPU Direct bypasses CPU/DRAM by writing directly through local storage or via NIC / remote node
- Expansion Storage shelf
  - NVMe-oF via RDMA over Converged Ethernet (RoCEv2) is increasing in popularity as an alternative to Infiniband with minimal performance tradeoffs

Lustre - The Lustre® trademark is jointly owned by OpenSFS and EOFS

KIOXIA CM6 Series Enterprise NVMe® SSDs

- **Enterprise PCIe® 4.0, NVMe 1.4 SSDs**
- **Form factors: 2.5-inch, 15mm Z-height**
- **Proprietary KIOXIA architecture: controller, firmware and BiCS FLASH™ 96-layer 3D TLC memory**
- **Dual-port design for high availability applications**
- **6th generation flash die failure recovery**
- **High performance with lower power consumption**
- **Power loss protection (PLP) and end-to-end data protection**
- **Suited for 24x7 enterprise workloads**
- **Data security options: SIE, SED, FIPS 140-2 Certified**
- **Six power mode settings**
- **Available now**

			CM6 (Mixed-Use)					CM6 (Read-Intensive)					
Endurance		DWPD	3					1					
User Capacity*		GB	800	1600	3200	6400	12800	960	1920	3840	7680	15360	30720
Sequential Read	128KB(QD32)	MB/s	6900	6900	6900	6900	6900	6900	6900	6900	6900	6900	6850
Sequential Write	128KB(QD32)	MB/s	1400	2800	4200	4000	4000	1400	2800	4200	4000	4000	4000
Random Read	4KB(QD256)	KIOPS	800	1300	1400	1400	1400	800	1200	1400	1300	1400	900
Random Write	4KB(QD32)	KIOPS	100	215	350	325	330	50	100	170	170	170	70

* Definition of capacity - KIOXIA Corporation defines a kilobyte (KB) as 1,000 bytes, a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1Gbit = 230 bits = 1,073,741,824 bits, 1GB = 230 bytes = 1,073,741,824 bytes and 1TB = 240 bytes = 1,099,511,627,776 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

Note: Specifications are subject to change

KIOXIA CD6 Series Data Center NVMe® SSDs

- **Data Center PCIe® 4.0, NVMe 1.4 SSDs**
- **Form factors: 2.5-inch, 15mm Z-height**
- **Proprietary KIOXIA architecture: controller, firmware and BiCS FLASH™ 96-layer 3D TLC memory**
- **Single-port design, optimized for data center class workloads**
- **6th generation flash die failure recovery**
- **Designed for high density storage deployments**
- **Power loss protection (PLP) and end-to-end data correction**
- **Data security options: SIE, SED, FIPS 140-2 Certified**
- **Five power mode settings**
- **Available Now**

			CD6 (Mixed-Use)					CD6 (Read-Intensive)				
Endurance		DWPD	3					1				
User Capacity*		GB	800	1600	3200	6400	12800	960	1920	3840	7680	15360
Sequential Read	128KB(QD32)	MB/s	5800	5800	6200	6200	5500	5800	5800	6200	6200	5500
Sequential Write	128KB(QD32)	MB/s	1300	1150	2350	4000	4000	1300	1150	2350	4000	4000
Random Read	4KB(QD256)	KIOPS	700	700	1000	1000	750	700	700	1000	1000	750
Random Write	4KB(QD32)	KIOPS	90	85	160	250	110	30	30	60	85	30

* Definition of capacity - KIOXIA Corporation defines a kilobyte (KB) as 1,000 bytes, a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1Gbit = 230 bits = 1,073,741,824 bits, 1GB = 230 bytes = 1,073,741,824 bytes and 1TB = 240 bytes = 1,099,511,627,776 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

Note: Specifications are subject to change

WEKA

Shimon Ben-David

Chief Technology Officer

www.weka.io

Genomics Environment with Legacy Solutions

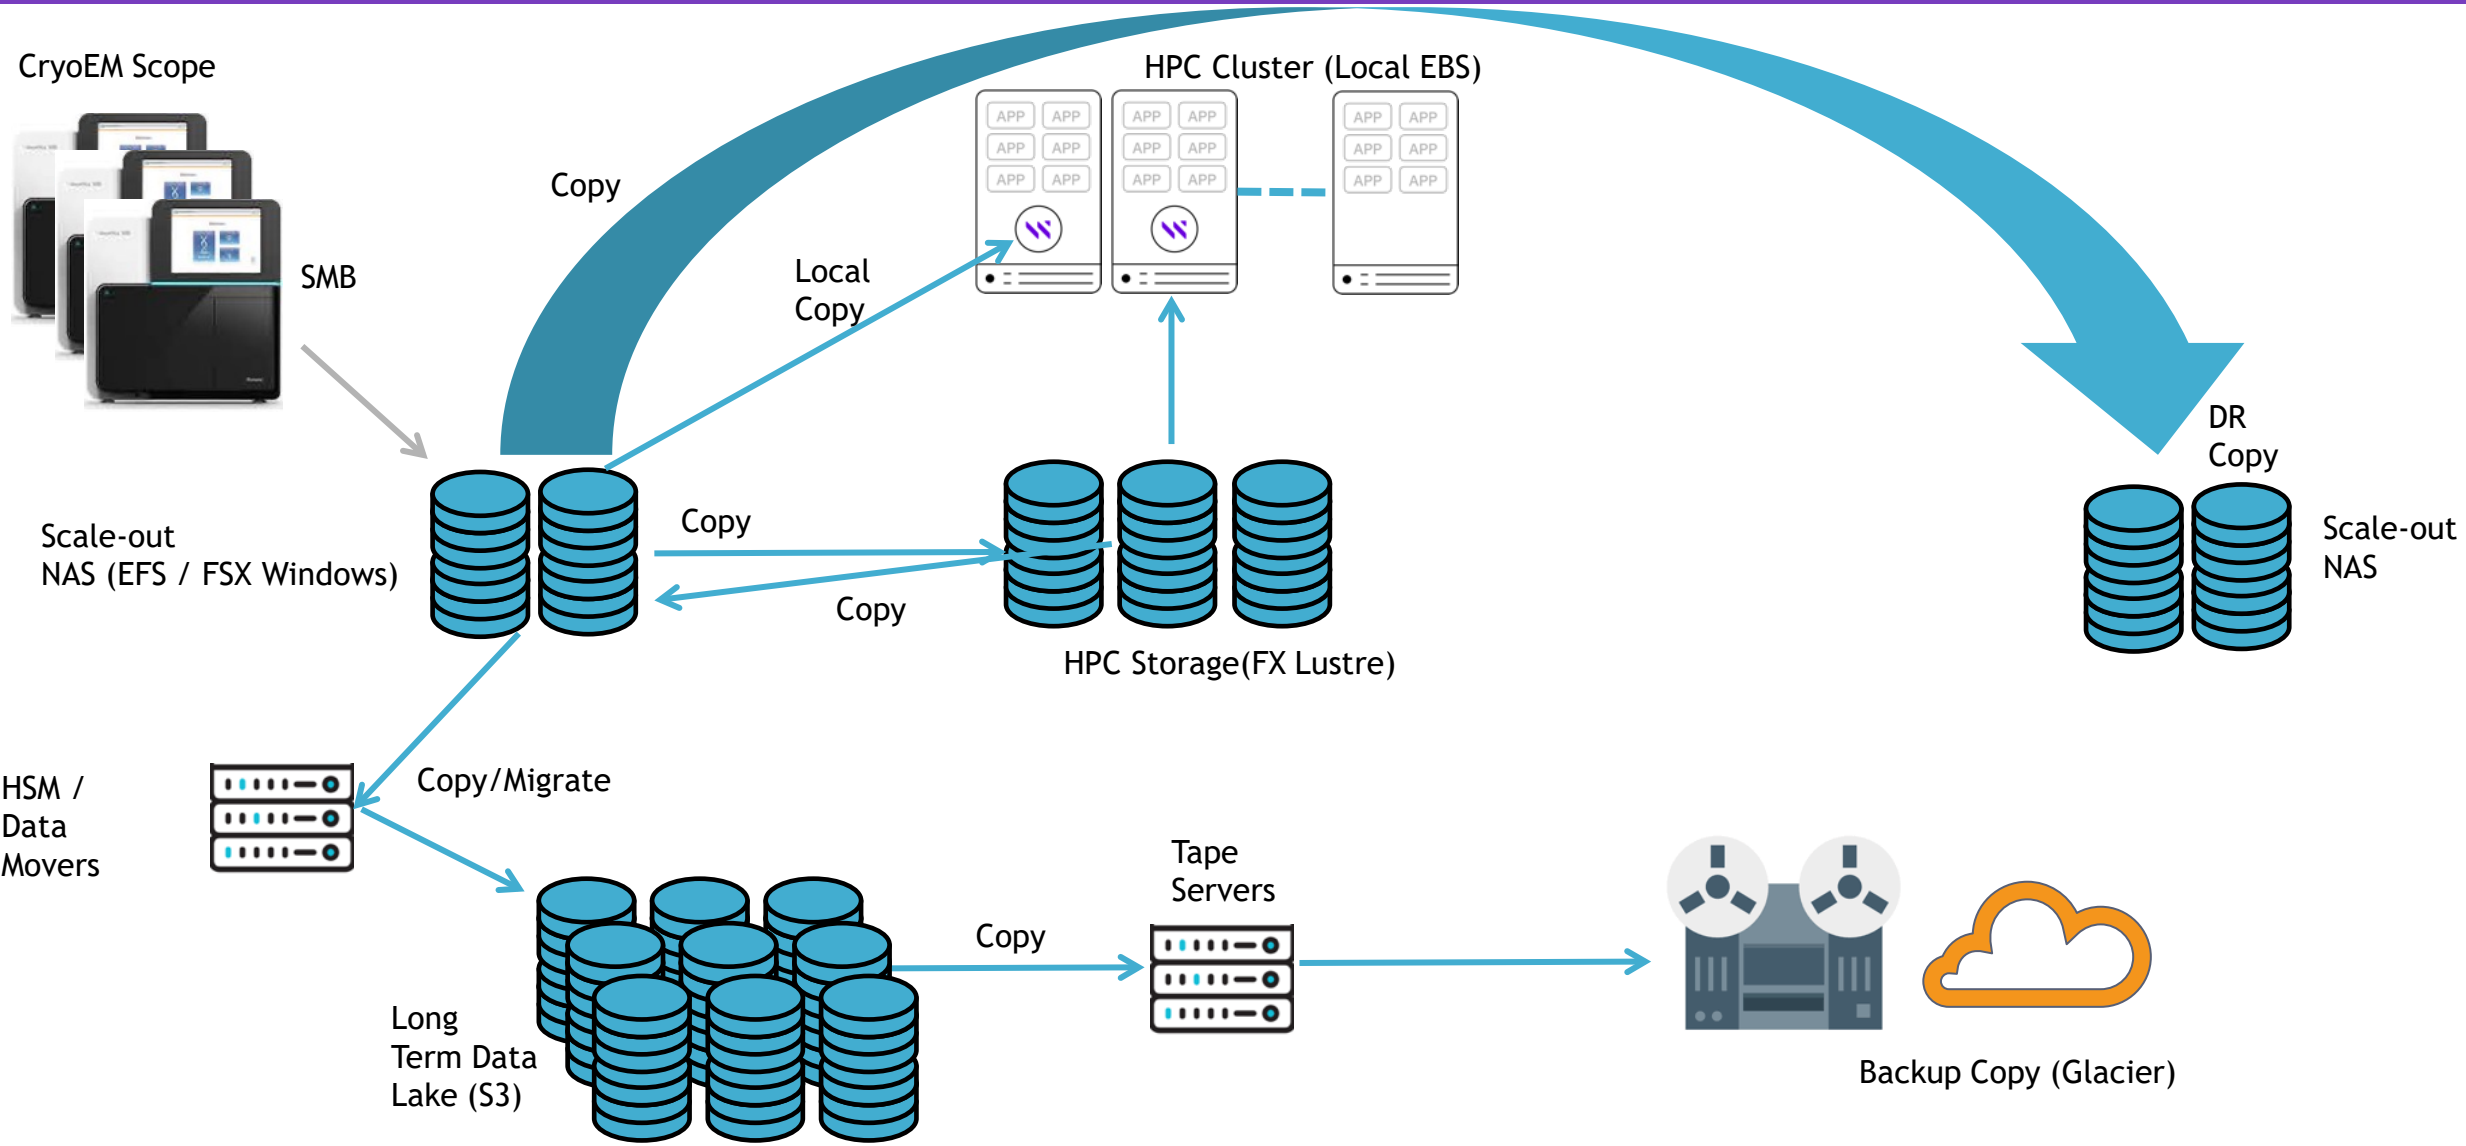

Solution: Zero-Copy Architecture

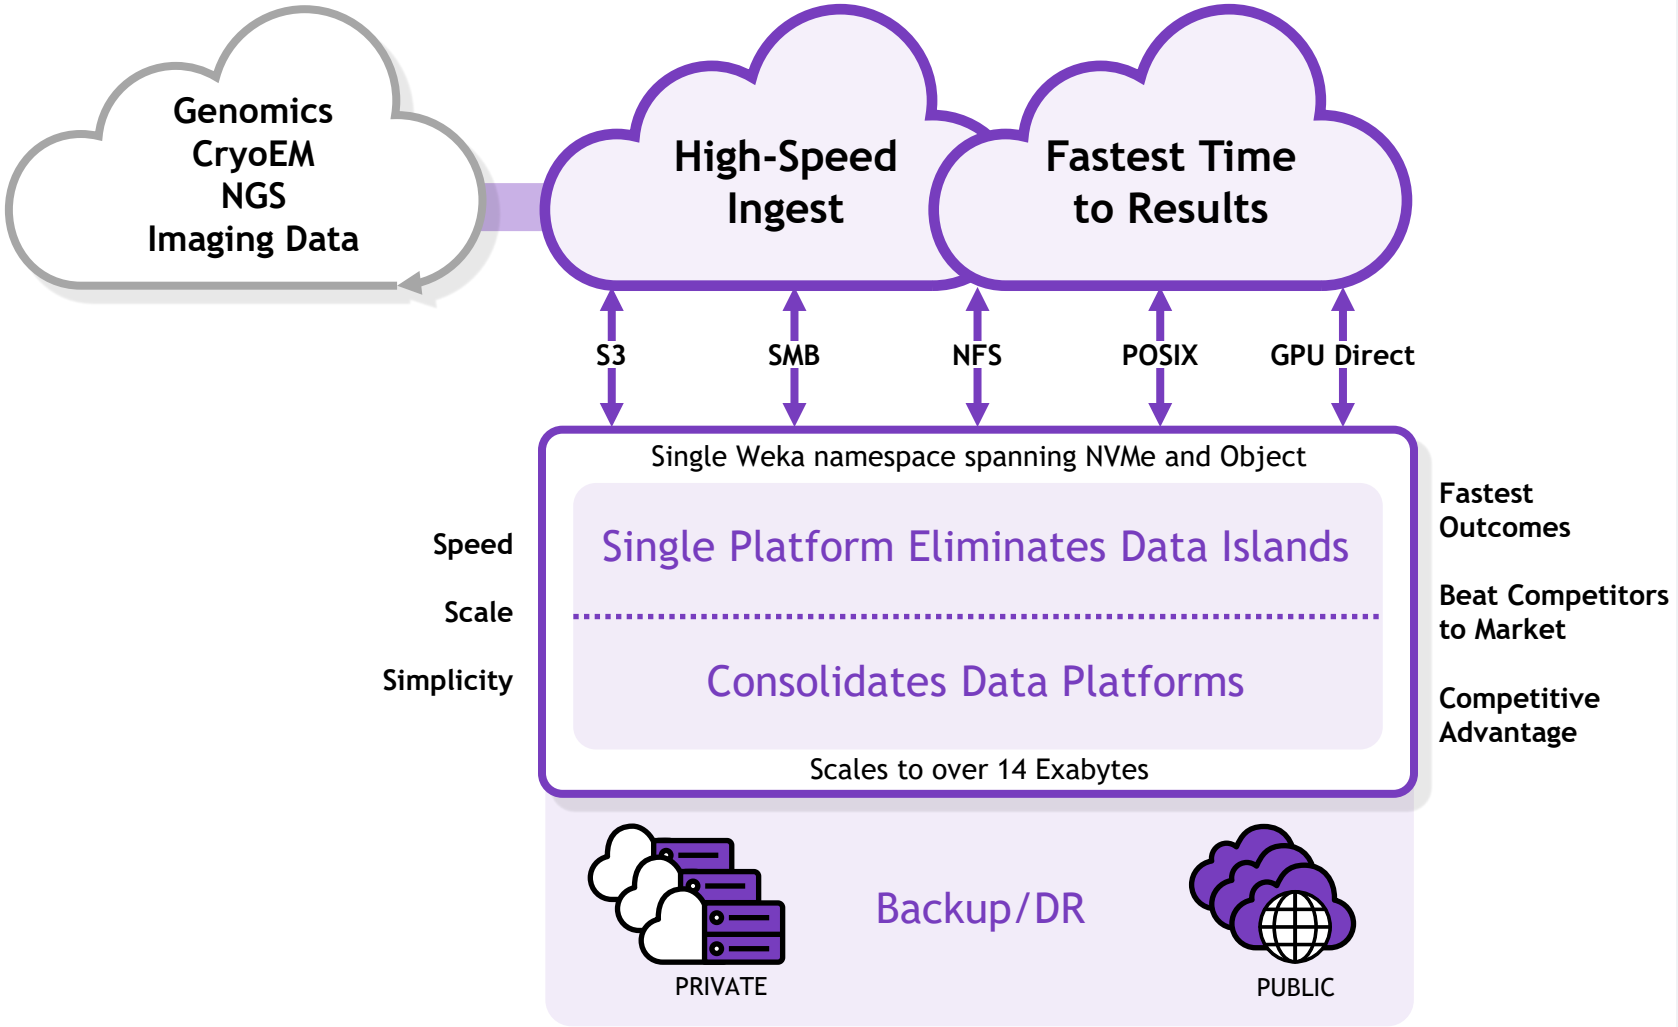

After Weka

Set It & Forget It

Simply write data to Weka and we transparently place data to the appropriate location

Limitless Platform - Any Application

Accelerate applications which encompasses small file, large file, and meta data intensive workflows

Flexible Storage Consumption

Any way you want.

Weka Data Platform

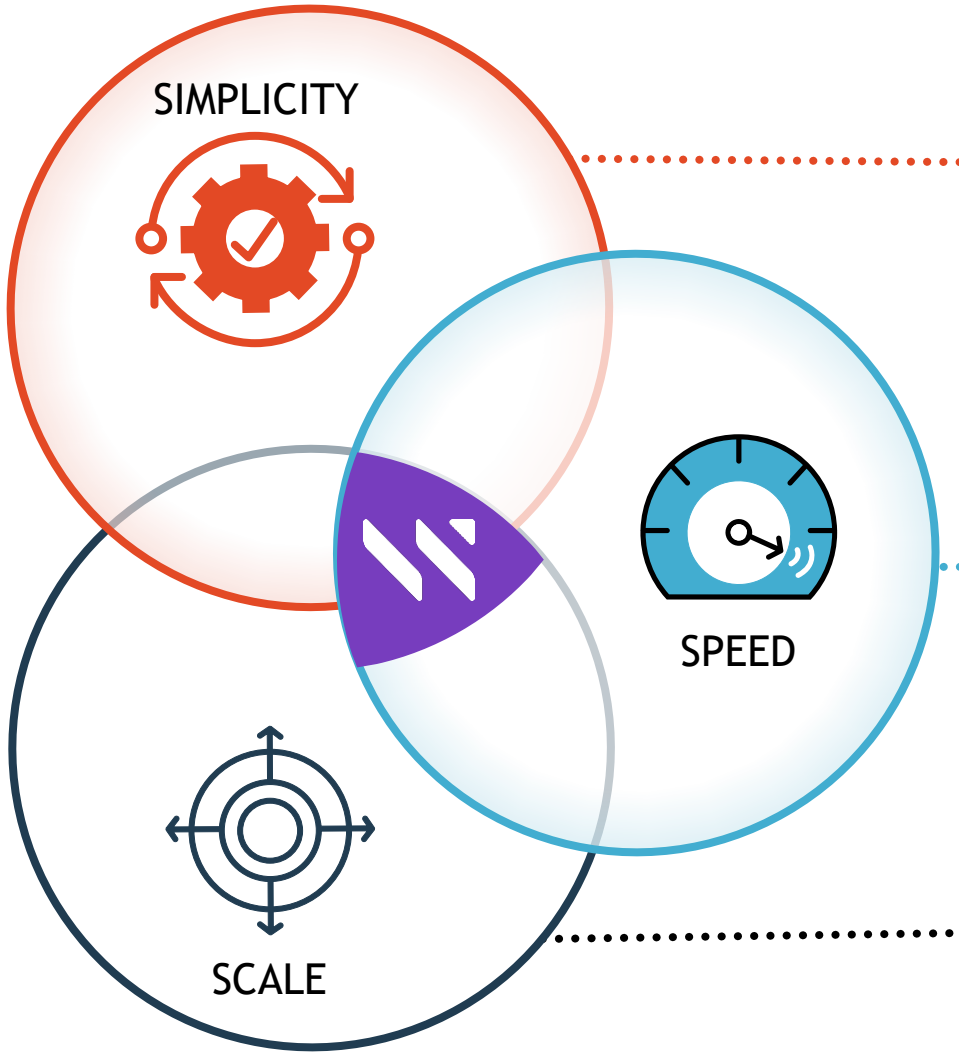

SIMPLICITY

- Unifies all your data
- Single pane of management
- Any platform- cloud or on-premises

SPEED

- Industry Leading Performance
- Resiliency without performance impact
- Amazing latency across diverse workloads

SCALE

- From Terabytes to Exabytes
- Data mobility to cloud
- Flash & disk for best economics

Weka Data Platform

Eliminate Data Silos

- Multi-protocol ready
  - SMB
  - NFS
  - S3/Object
  - GPUDirect Storage
  - CSI
- Diverse workloads
  - File and Object
  - Small and Large file workloads
  - IOPS/Throughput/Latency/Metadata

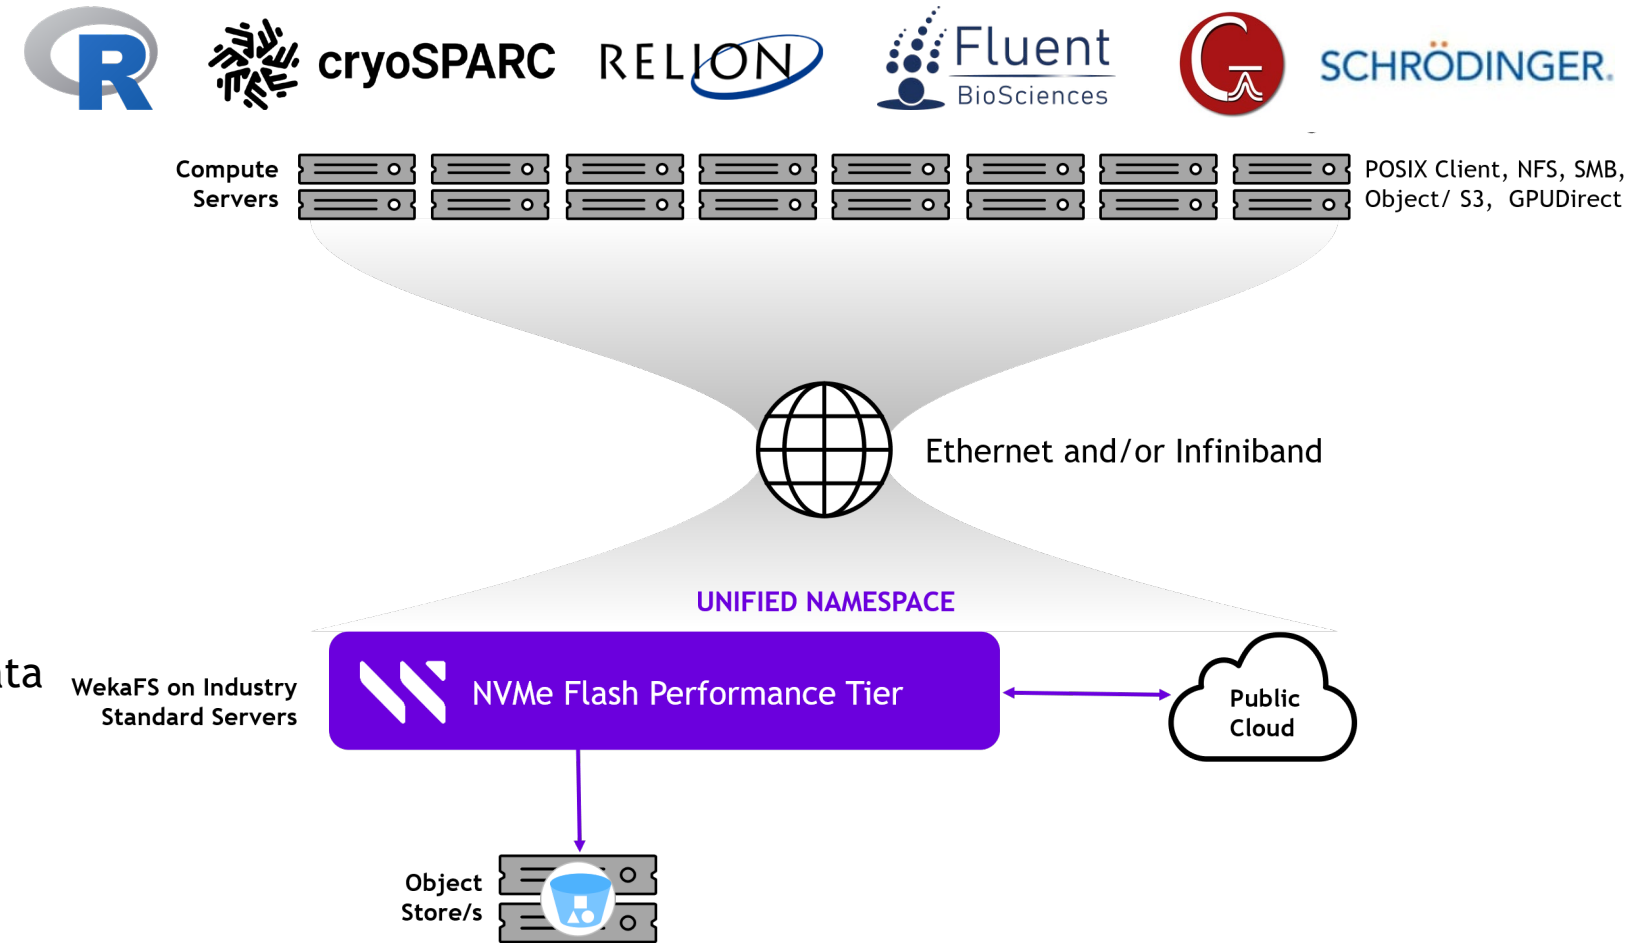

Cloud Ready - Data Mobility - Business Continuity

- Transparent backup, archive, data migration
- Snap-to-object allows burst to cloud
- Disaster recovery and business continuity baked into solution

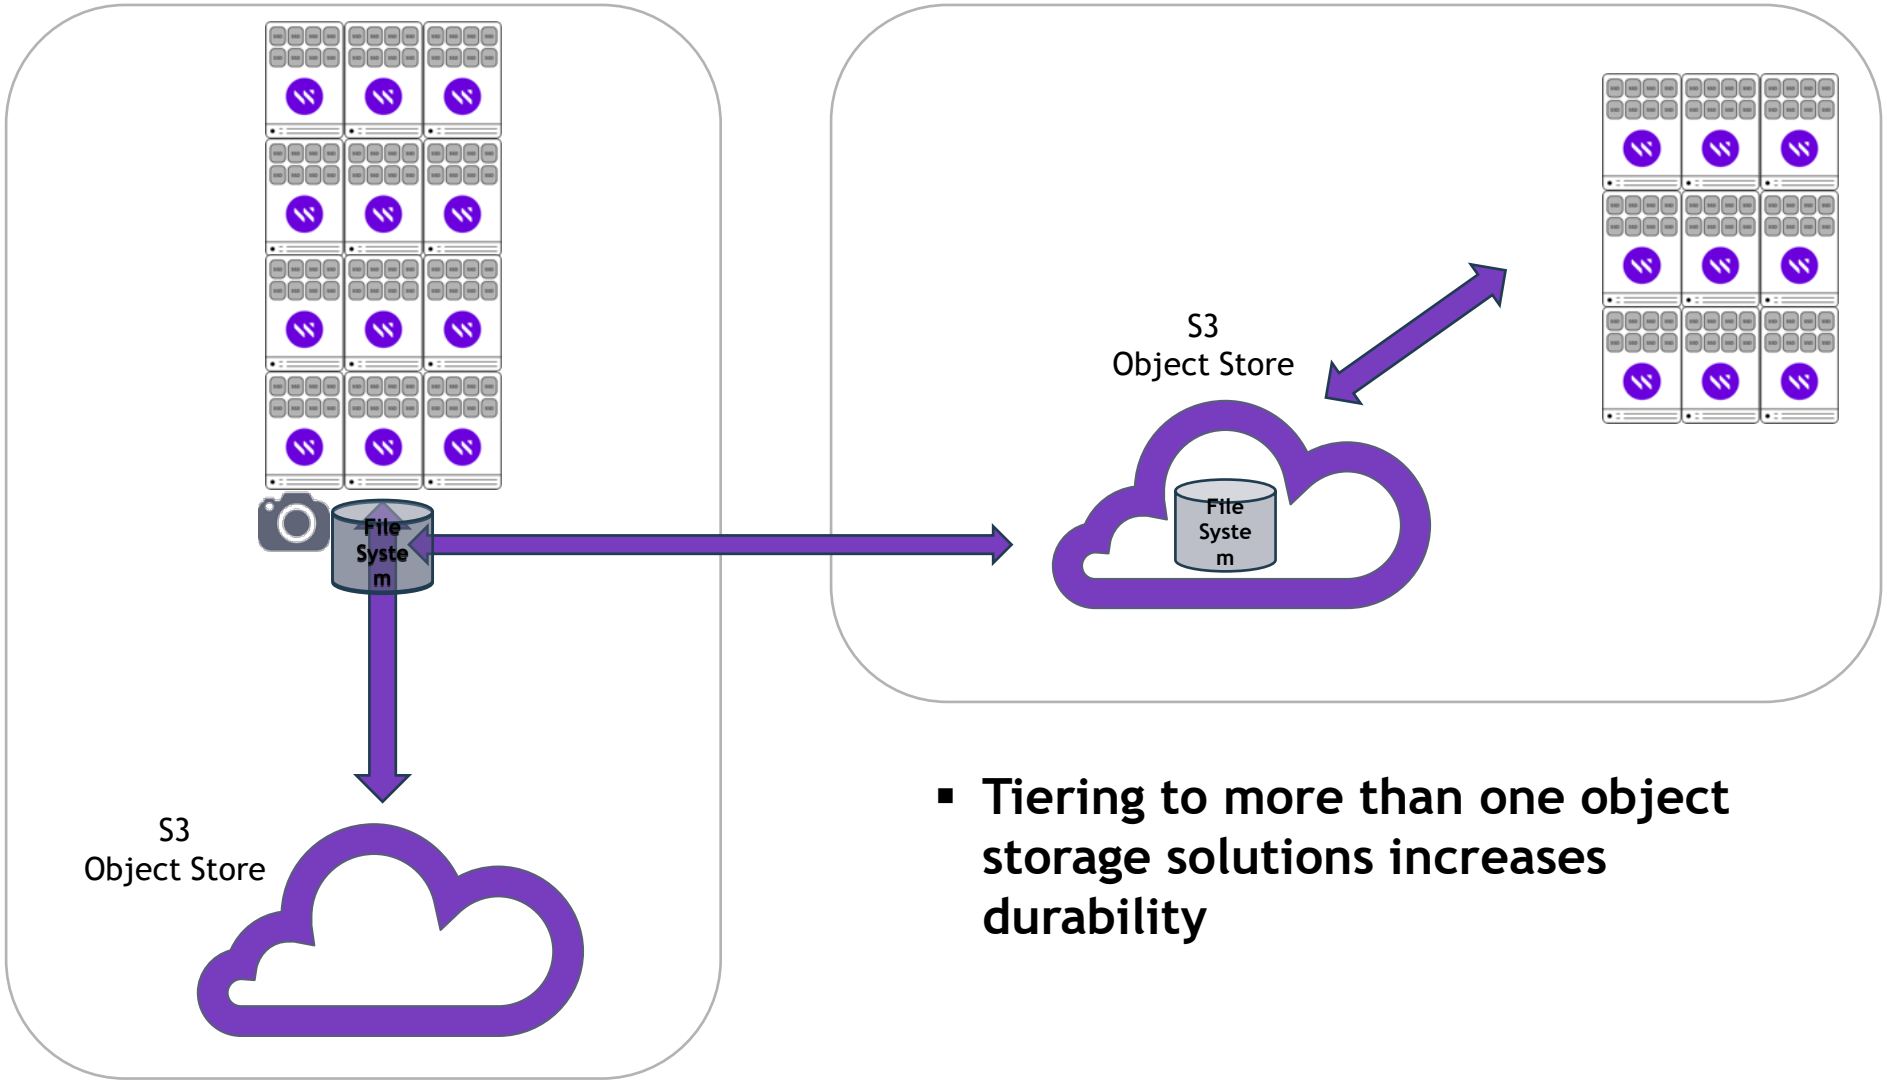

- Tiering to more than one object storage solutions increases durability

Atomwise Case Study

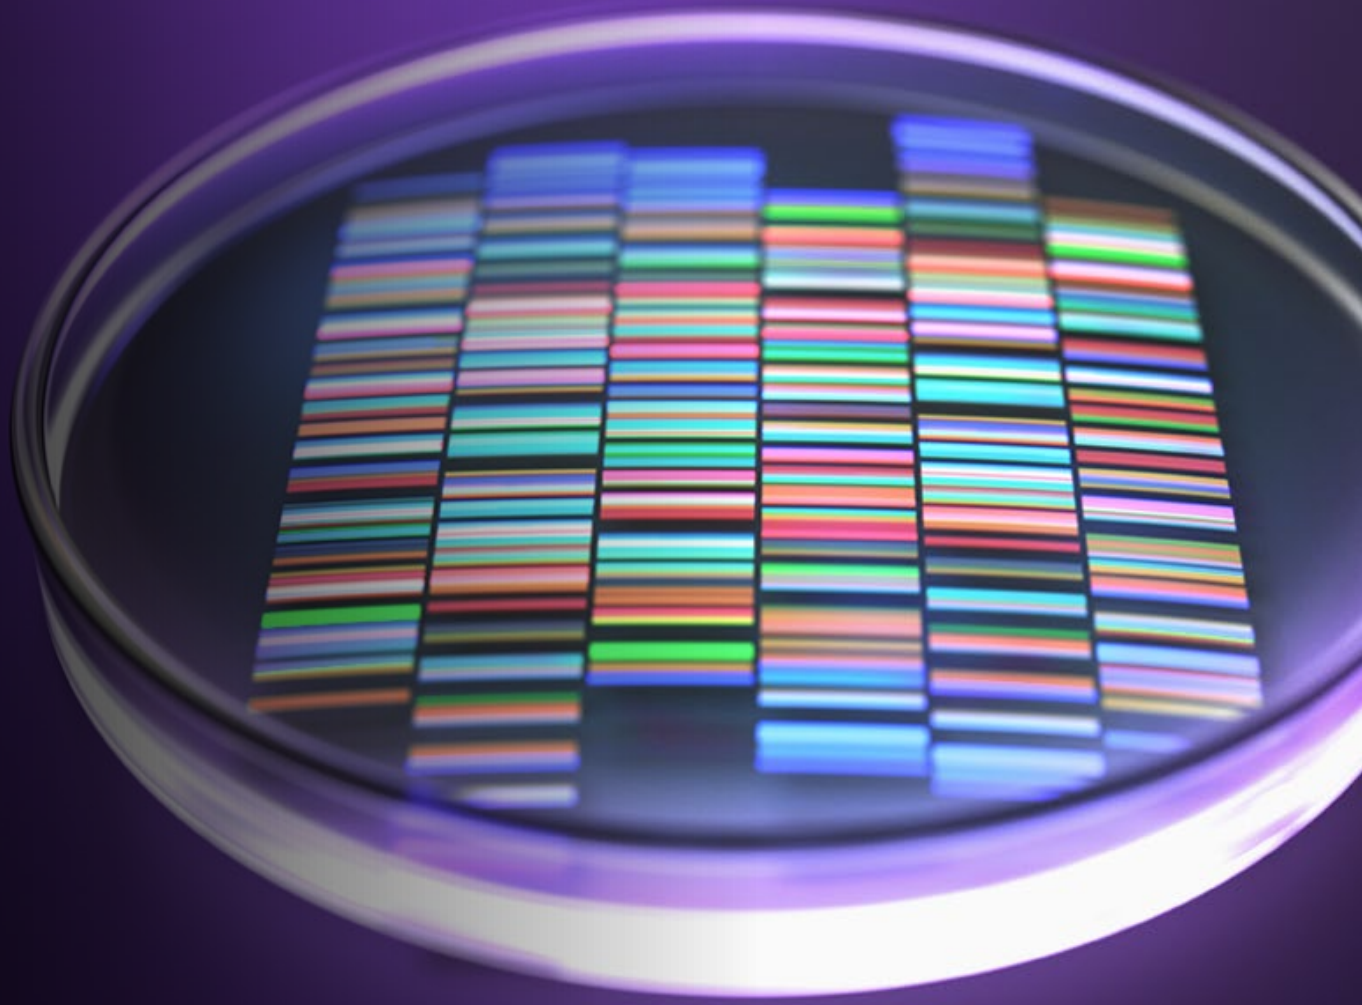

WEKA

Life Science and AI in AWS

Atomwise

- Problem
  - Needed to serve 200 million small 10KB-sized files to GPU nodes
  - GPU servers were starved of I/O causing massive budget overruns
  - Multiple copies of data to mange on EBS to improve performance
  - Had to be cloud native in AWS
- What They Tried
  - Used EBS to run the workloads direct to each GPU node
  - Ended up with multiple copies of the same data set on EBS which proved prohibitively expensive

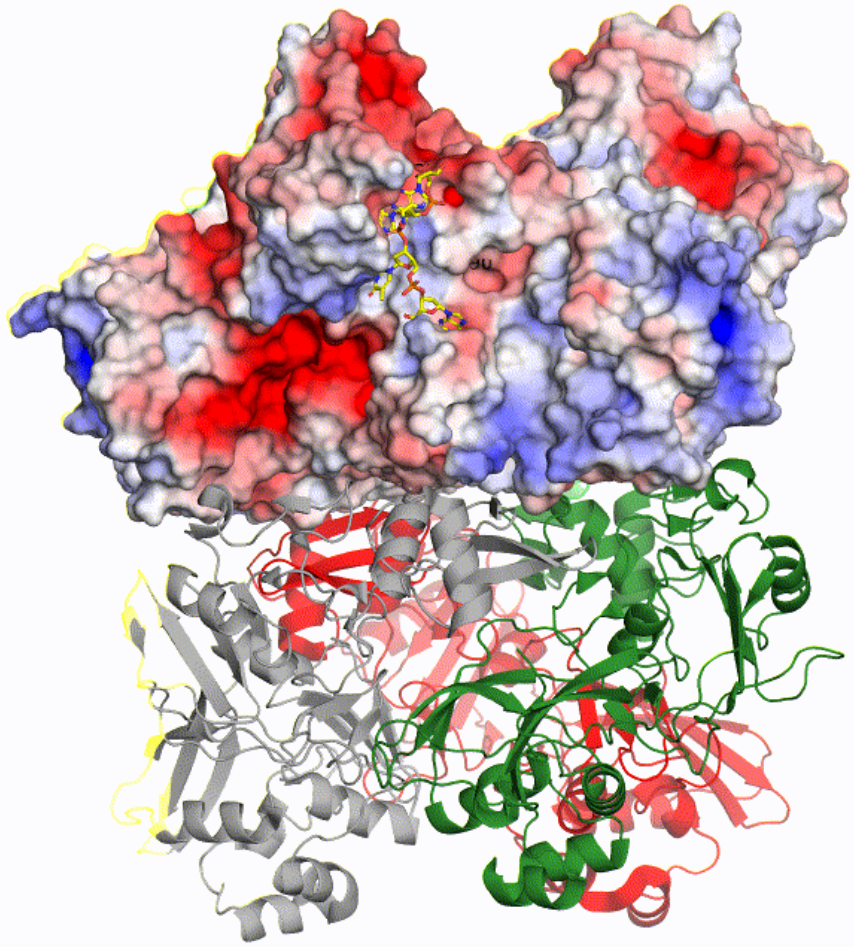

nsp15

(SARS-CoV-2019 endoribonuclease)

Convolutional neural networks for image recognition

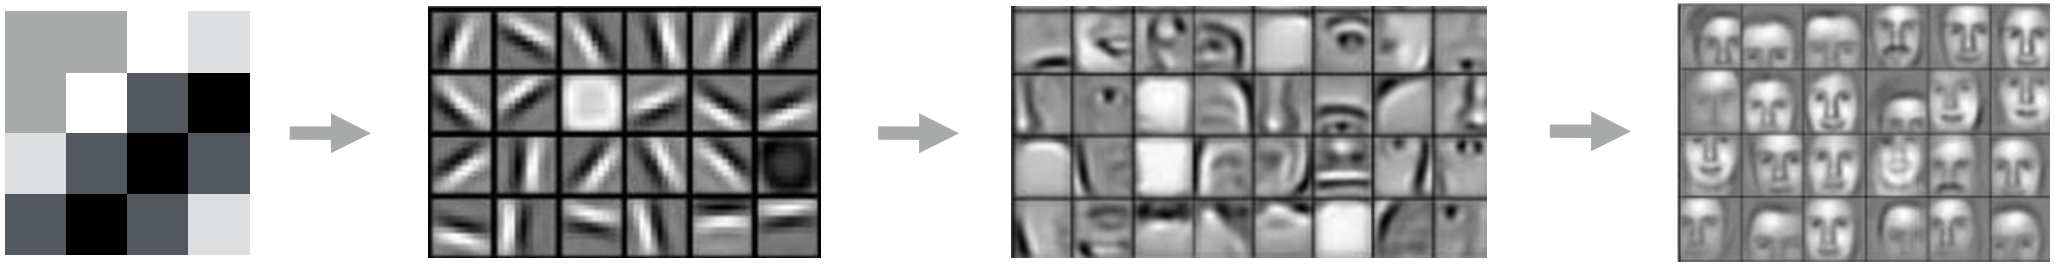

Pixels → Edges → Eyes, Noses, Mouths → Faces

Convolutional neural networks for molecular recognition

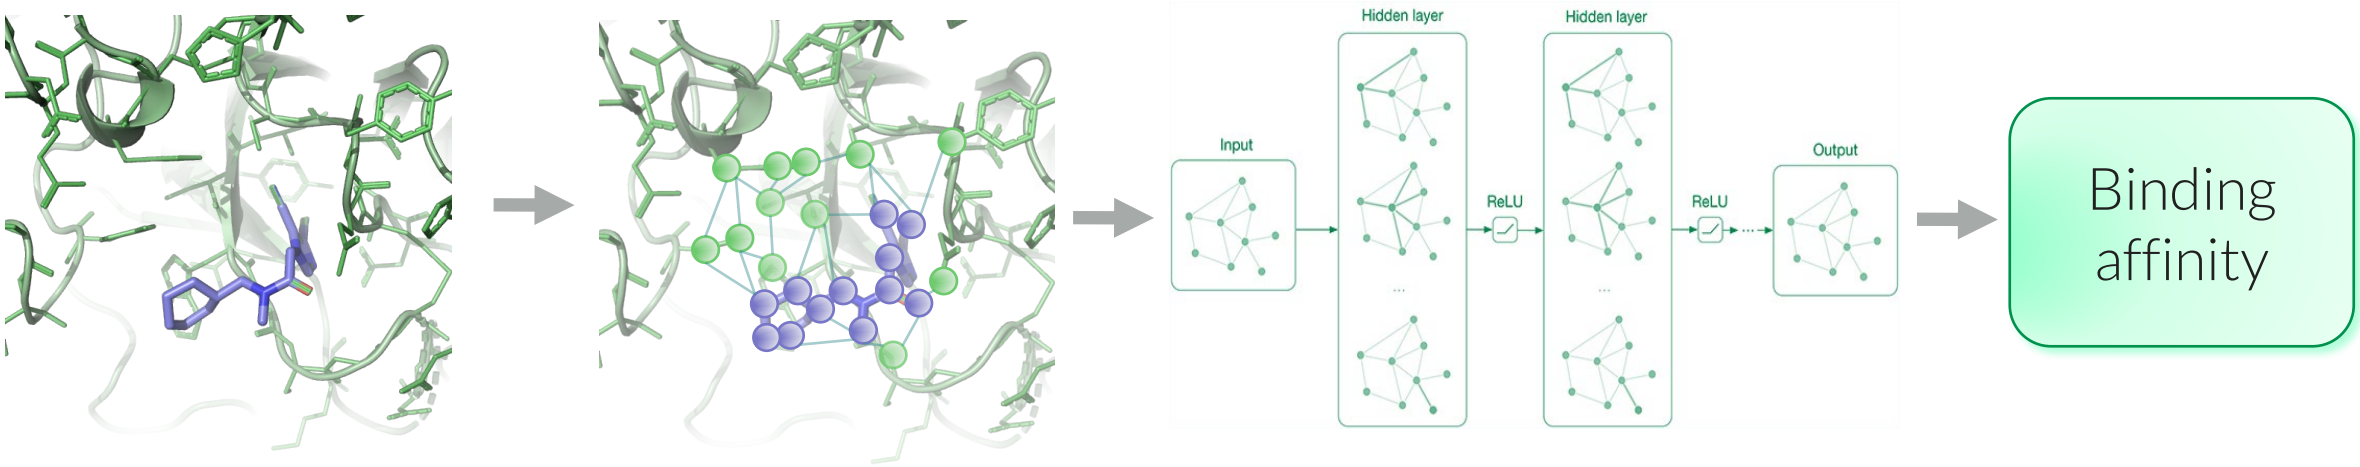

>4,000
PROTEIN TARGETS

>3M
COMPOUNDS

>15,000,000
EXPERIMENTAL MEASUREMENTS

Parting remarks

- It was quick to set up a server and sync our S3 data to it
- IO performance is **phenomenal** for shared file access
- We saw model training times drop by up to **2x**
- Incorporating WekaIO changed our solution space
  - Some experiments went from requiring **3 months to 1 week**
- It's an ideal match for distributed deep learning with large file inputs
- The team has been extremely supportive

Thank You!

- Weka + AWS + Atomwise Webinar On-demand

https://www.weka.io/resources/accelerating-ai-training-models/

- Weka + AWS + Atomwise Case Study

https://www.weka.io/wp-content/uploads/files/2021/10/2021-Atomwise-Weka-AWS-Partner-Case_Study_WekaIO_FINAL.pdf

- Weka Named a Visionary in Magic Quadrant™ for Distributed File Systems and Object Storage

https://www.weka.io/2021-gartner-mq-distributed-file-systems-block-storage/

- **Contact Us**

https://www.weka.io/contact-us/

NetApp

Esteban Rubens

Principal, Healthcare AI Practice

www.netapp.com

Enterprise Imaging storage

- Data creation
  - Traditional areas
    - Radiology
    - Cardiology
  - Other specialties
    - All the other “ologies”
    - Visible-light imaging
- Data consumption
  - Screening
  - Diagnostic
  - “AI”
  - Clinical trials
- Data access
  - File vs object

Digital pathology examples

- Most scanners support resolutions of 0.5 µm/pixel (effective viewing magnification: 20X) or 0.25 µm/pixel (effective viewing magnification: 40X)
  - Following image compression at the scanner, each image file produced is over 1 GB in size
- Academic medical center
  - Annual Volume 137k cases
  - 87k anatomic pathology, 30k cytology, 20k hematopathology
  - 7-10 slides/case @ 1.3 GB/slide
  - 1.2-1.7 PB/year storage required in a single institution
- Hospital with 880 beds
  - 500 TB/year for digital pathology

AI in Enterprise Imaging

- Computer-Aided Detection (CADe)
- Computer-Aided Diagnosis (CADx)
- Case prioritization
  - Worklist management
  - Triage
- Automated segmentation
  - Automated labeling
- Automated measurements
  - Pre-populate structured reports
- Automated reading of normal cases
- Automated patient positioning

The NetApp Data Fabric

NetApp Data Fabric

Have the right data in the right place, at the right time, at the right price, with the right access characteristics.

Seamlessly move data

- To and from any clouds
- Between different clouds
- Between ONTAP systems (edge, core, and cloud)
- ONTAP on AFF for high-performance model training paired with an S3 data lake

The NetApp Data Fabric

Scale

Protect

Analyze

Secure

Optimize

Store

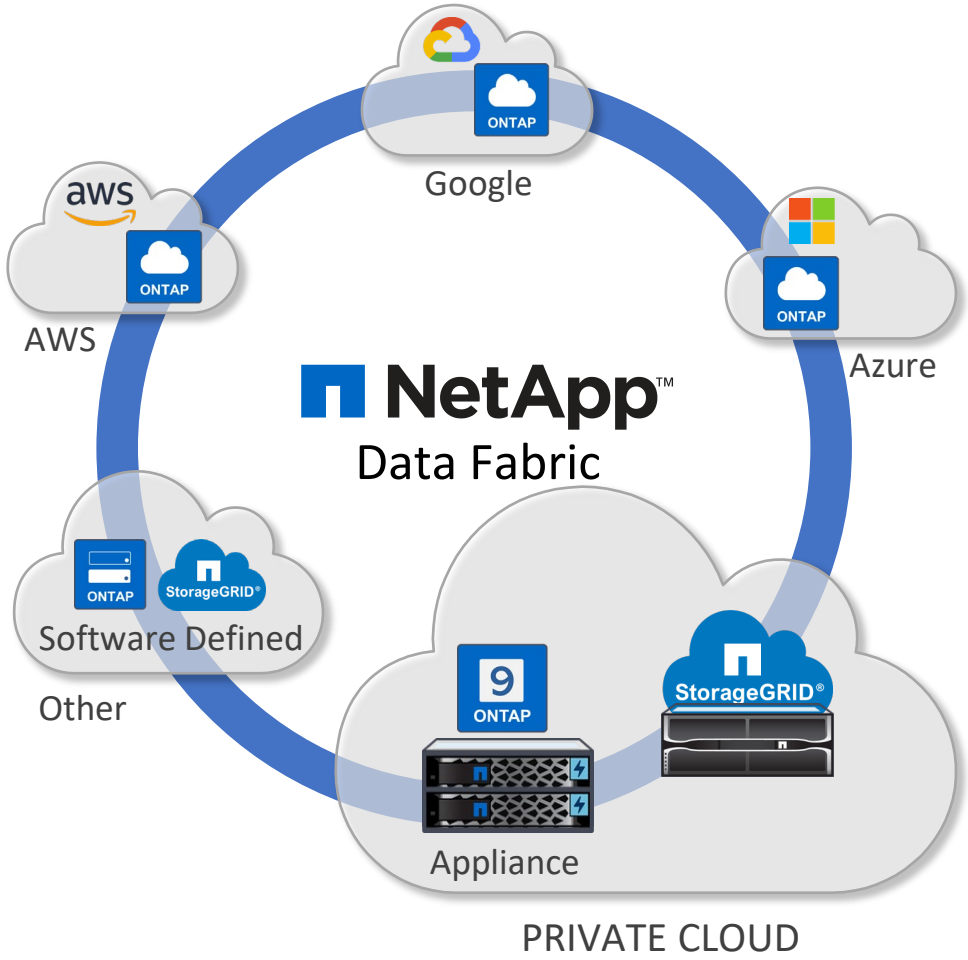

NetApp Cloud Tiering

NetApp Cloud Tiering

- Tiering infrequently-used data automatically and seamlessly
- From on-prem ONTAP SSD or CVO to low-cost public cloud object storage automatically
- Create custom tiering policies and we take care of the rest
- No impact to the application layer
- Available via Cloud Manager

NetApp Cloud Tiering

Simple

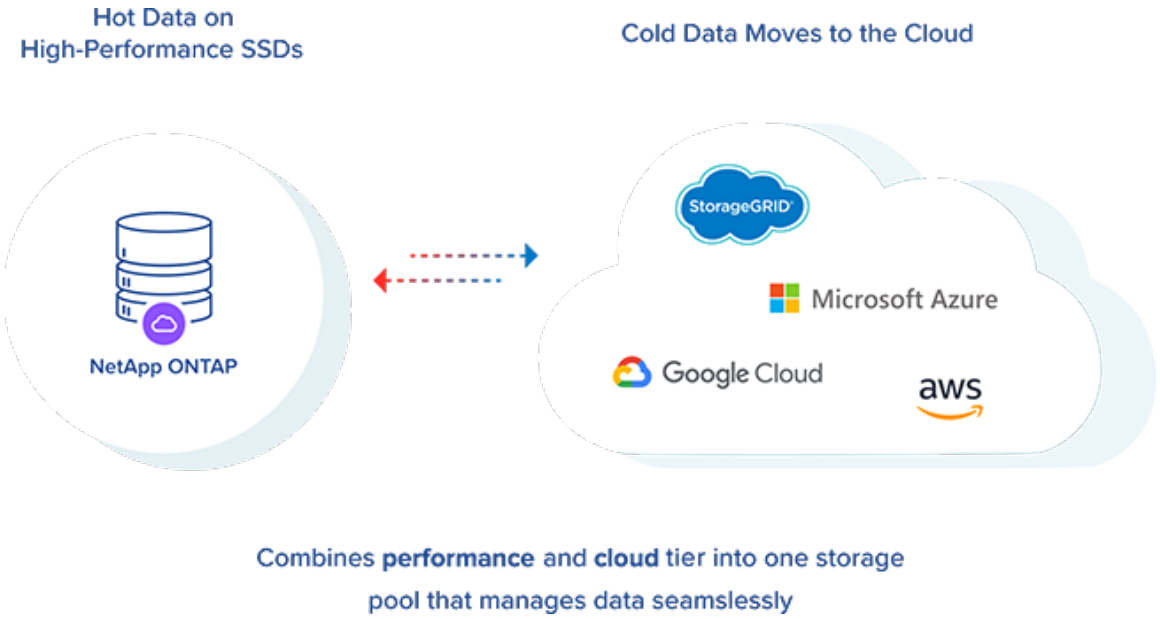

Detailed

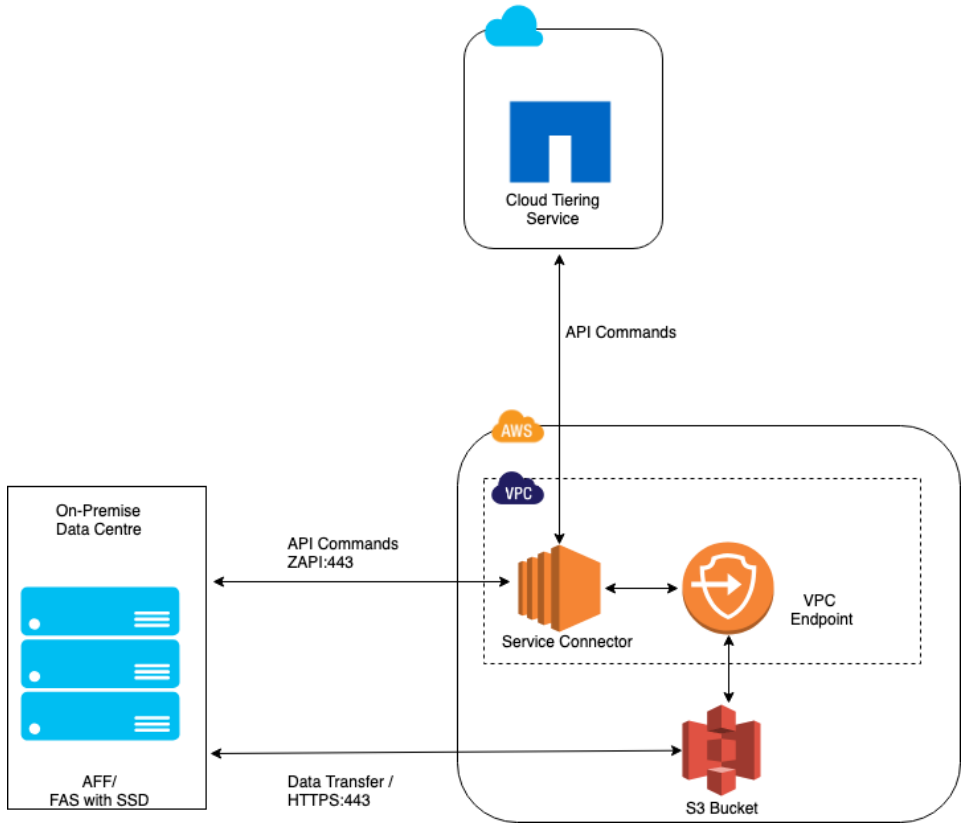

AWS FSx for NetApp ONTAP

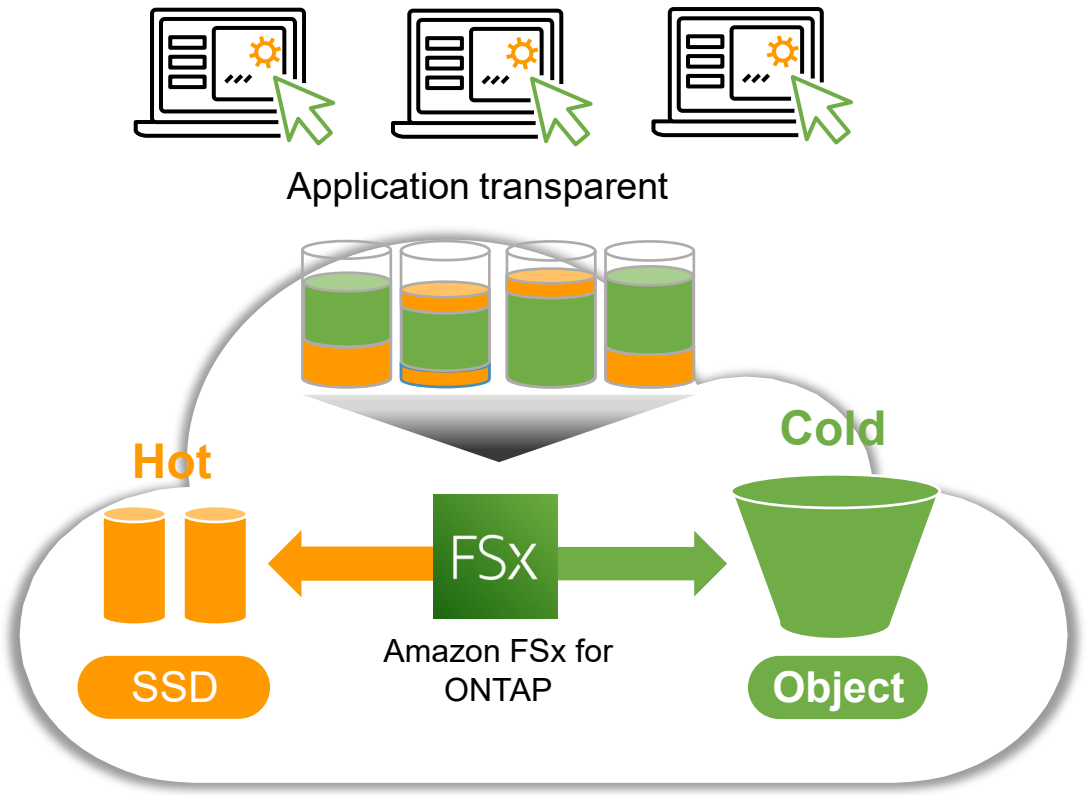

Bi-directional block-level data movement

Life sciences example

- Visible-light imaging of cell assays
- 10^8 to 10^9 images/year
- Image normalization & feature extraction in the cloud
- AI model training on-premises
  - NETAPP AI (NVIDIA DGX + NetApp AFF)
  - StorageGRID

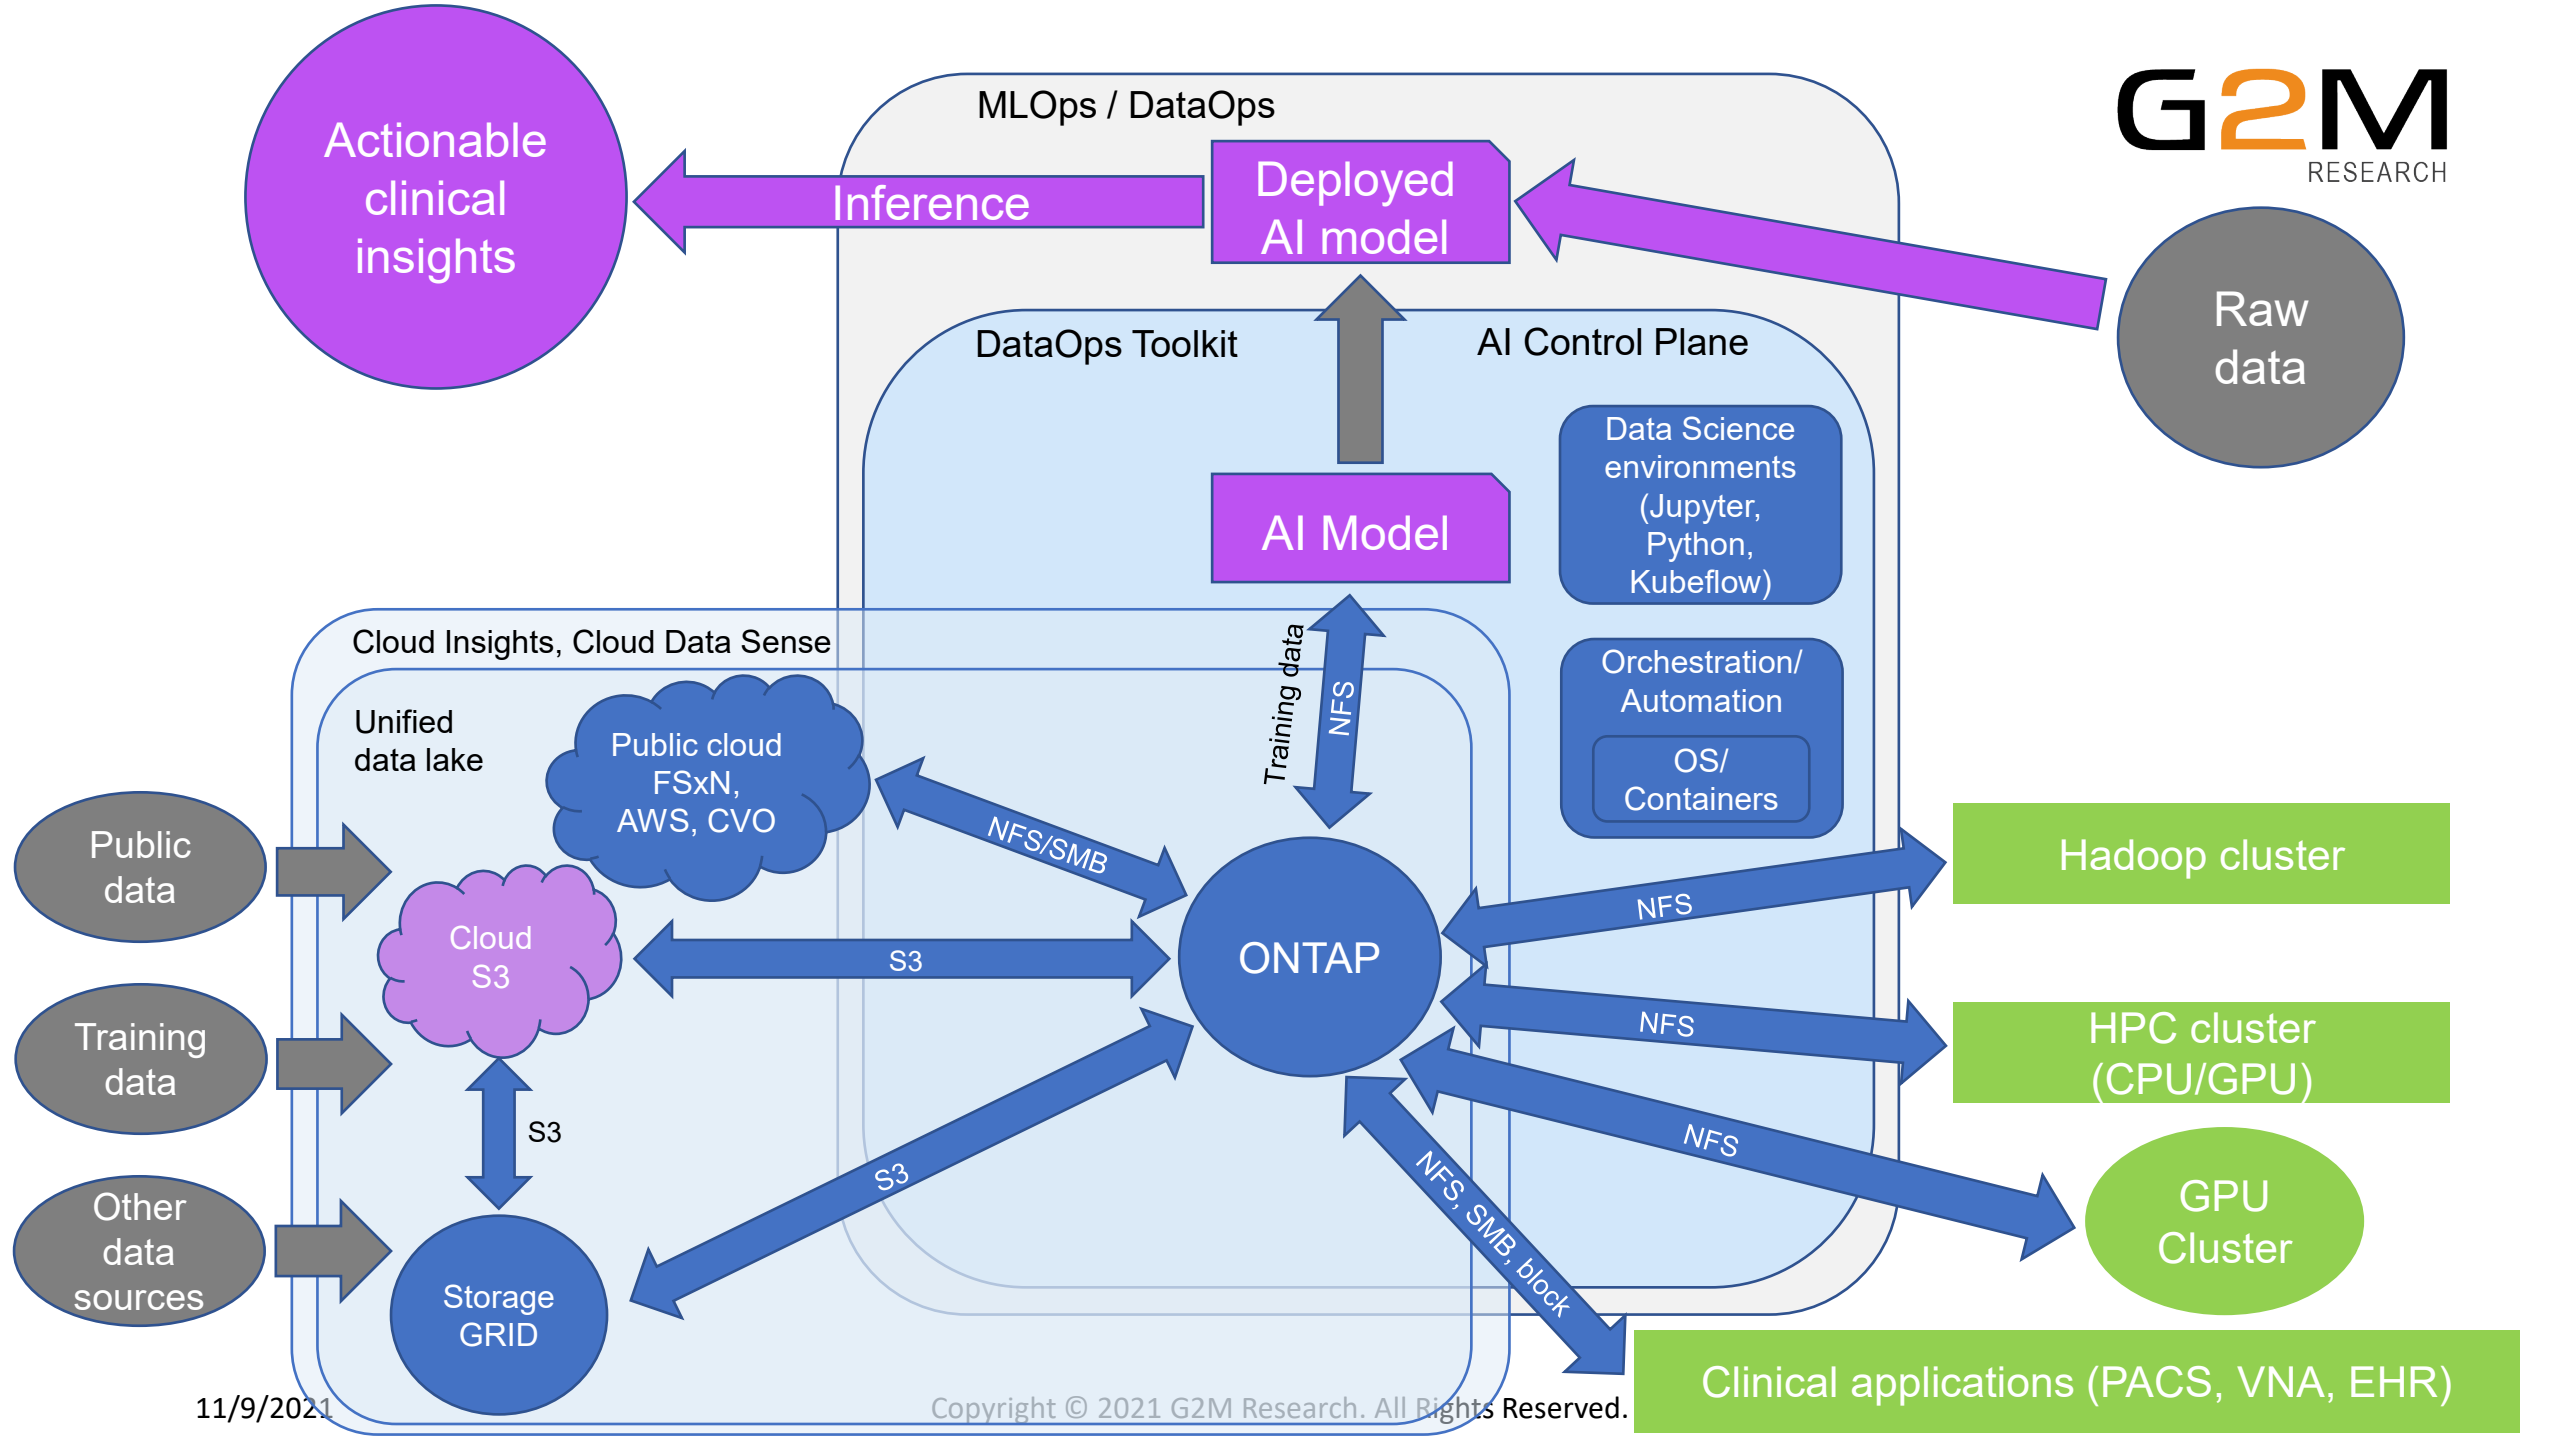

Panel Questions and Audience Surveys

Panel Question #1

The number of imaging technologies continues to explode, providing a great number of new sources for medical/bio imaging. Do these different technologies present the same or different issues for storage?

- Shimon Ben-David - Weka
- Matt Hallberg – Kioxia
- Esteban Rubens - NetApp

Audience Survey Question #1

How large are the imagery training data sets that you and/or your customers typically field (pick one answer):

- Greater than 5 PB: 10%
- Between 1PB and 5PB: 25%
- Between 250TB and 1PB: 35%
- Between 50TB and 250TB: 15%
- Less than 50TB: 5%
- Don’t know 10%

Panel Question #2

Are there specific areas where NVMe® and/or NVMe-oF™ provide unique benefits to large imagery storage solutions?

- Matt Hallberg – Kioxia
- Esteban Rubens – NetApp
- Shimon Ben-David - Weka

Audience Survey Question #2

When looking at storage solutions for medical/bio imagery, what are the critical factors that drive your technology and architecture choices (check all that apply):

- Performance (bandwidth and latency): 30%
- Performance consistency (especially latency consistency): 30%
- Storage networking performance: 20%
- Storage costs: 10%
- Ease of deployment/”turnkey” solutions 15%
- Don’t know/no opinion: 30%

Panel Question #3

Obviously, artificial intelligence and machine learning are being utilized to analyze all the images from these various sources. How can storage architectures be optimized to speed up the training/learning process for imagery?

- Esteban Rubens – NetApp
- Shimon Ben-David – Weka
- Matt Hallberg – Kioxia

Audience Q&A

G2M COMMUNICATIONS

Effective Marketing & Communications with Quantifiable Results