




Weka

Shimon Ben David – CTO

One Architecture : Zero Compromises



LIMITLESS DATA PLATFORM

SIMPLICITY

Easy to Install
Any Protocol
On-Prem or Cloud

SPEED

Resiliency without Performance Impact
3X Faster than Local SSD
10X Faster than All Flash NAS

SCALE

From Terabytes to Exabytes
Performance Across Large and Small
Files
Extend to Cloud for Maximum Elasticity

WEKA Market Traction

Magic Quadrant for Distributed File Systems and Object Storage 2021



FORTUNE 50

12 customers in the Fortune 50

AI/ML

40% running production AI instances on Weka

CLOUD

25% customers running Weka in public cloud

Pushed 230PB of data to AWS S3 in 2020

Gartner Peer Insights 64 of 5/5 Star Reviews

Highest ranking of any infrastructure company¥product



**“Great Performance
In An Easy Package”**



**“Delivered on all
benchmarking requirements
with excellent service”**



**“Happy Customer, Leading
Edge Technology With Great
Support Service”**



**“The Ferrari of
The Storage World”**



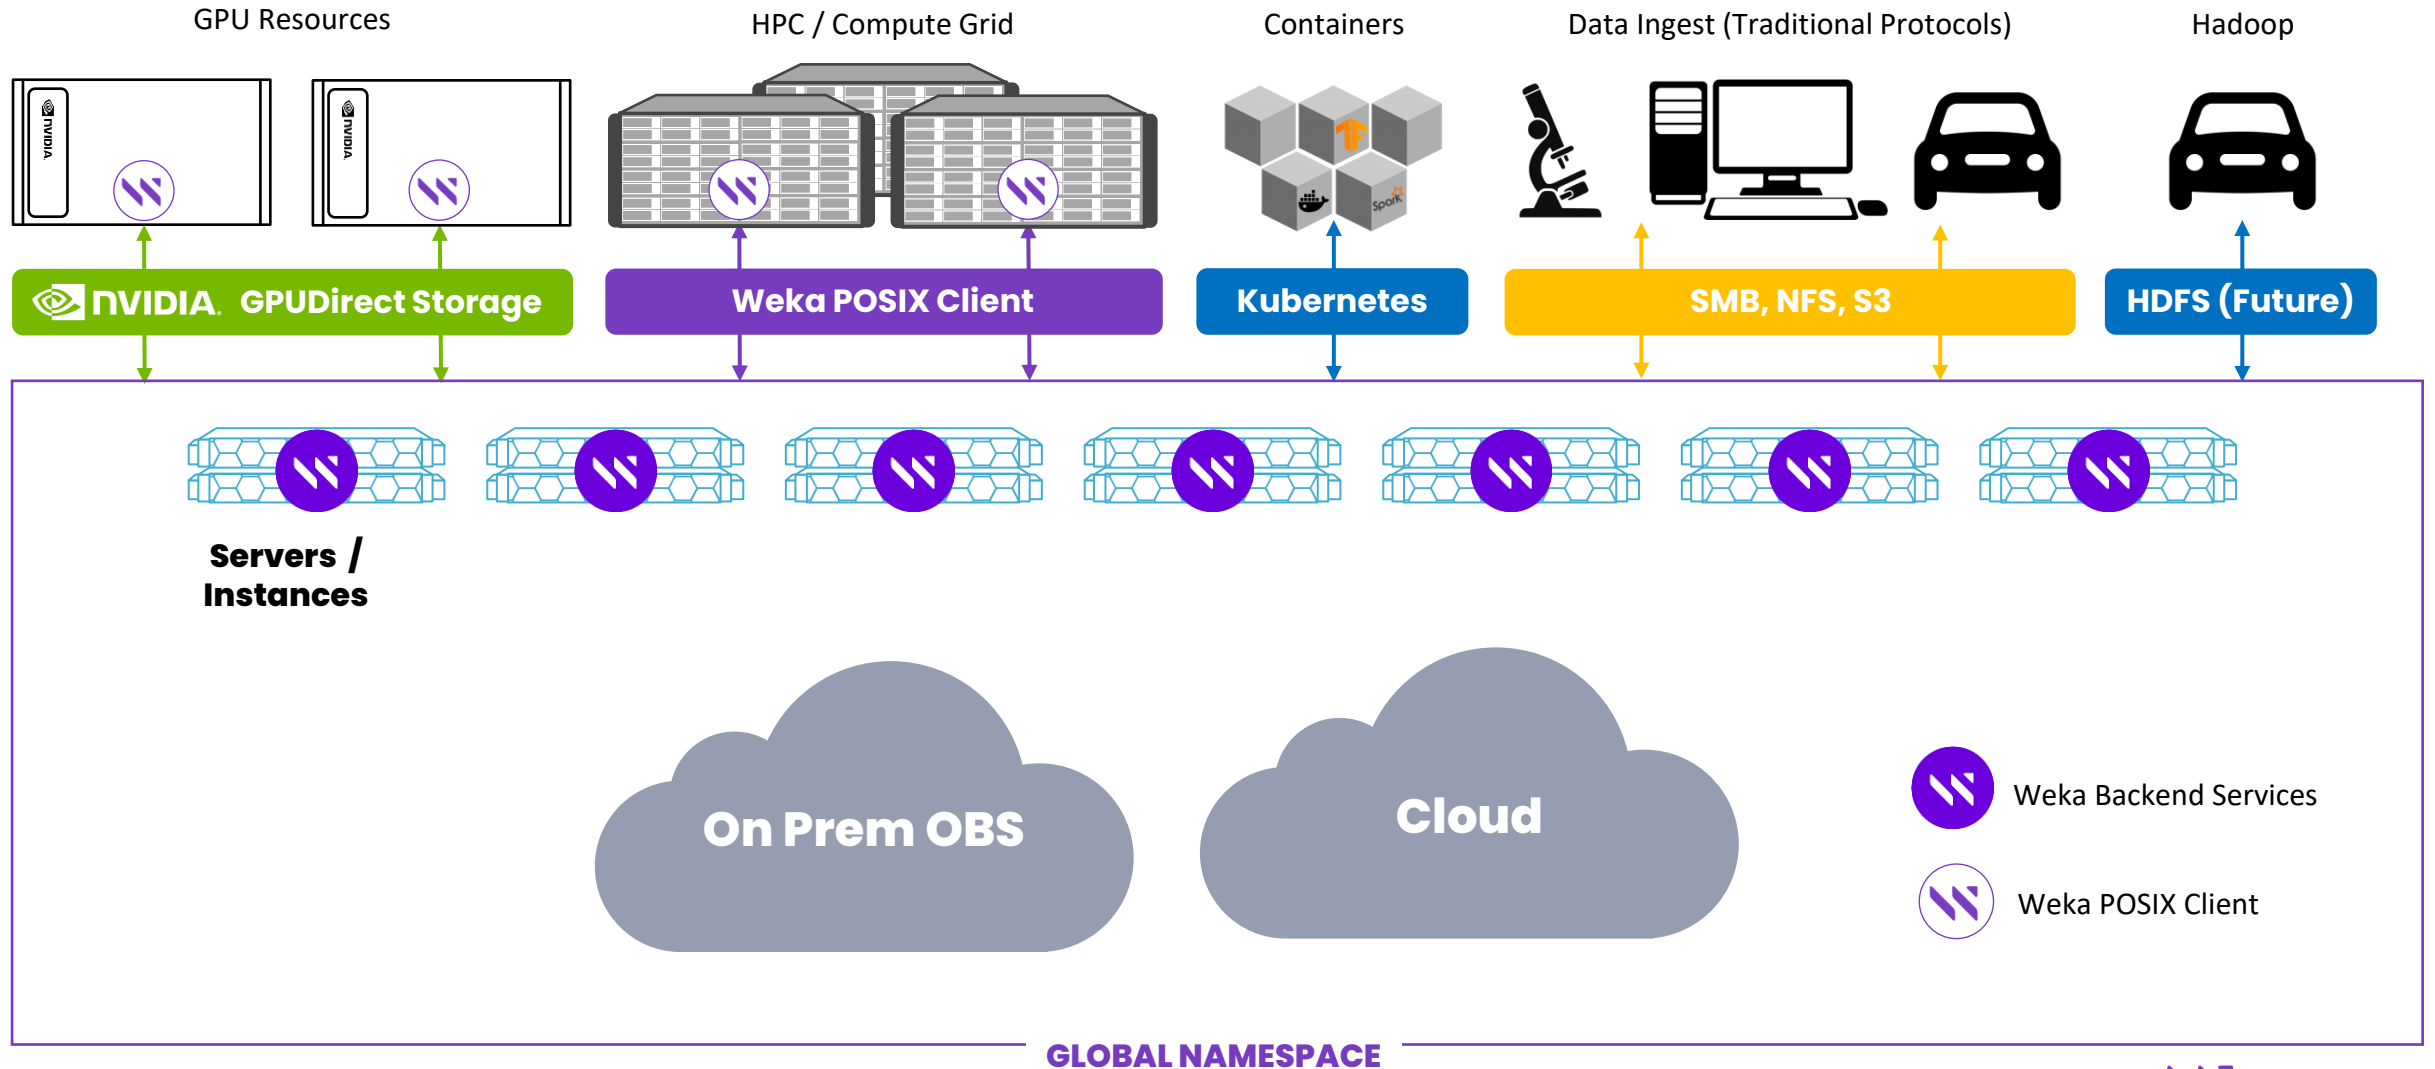
**“Care About Scaled ML
Performance?
WekaFS Delivers.”**



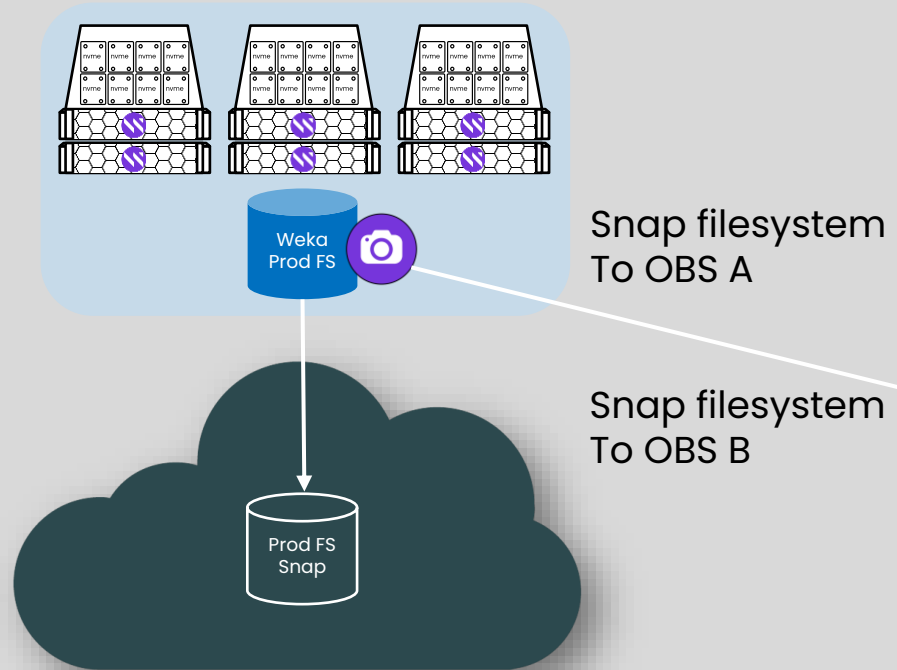
**“Excellent Product Technically,
Backed-Up By
First Class Support”**



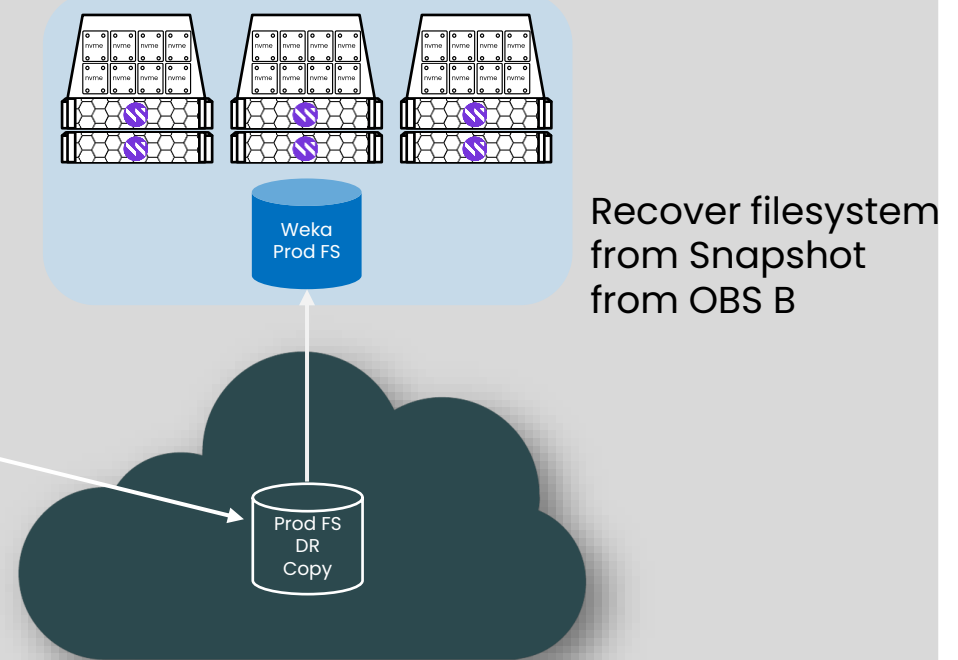
WEKA Data Platform Architecture



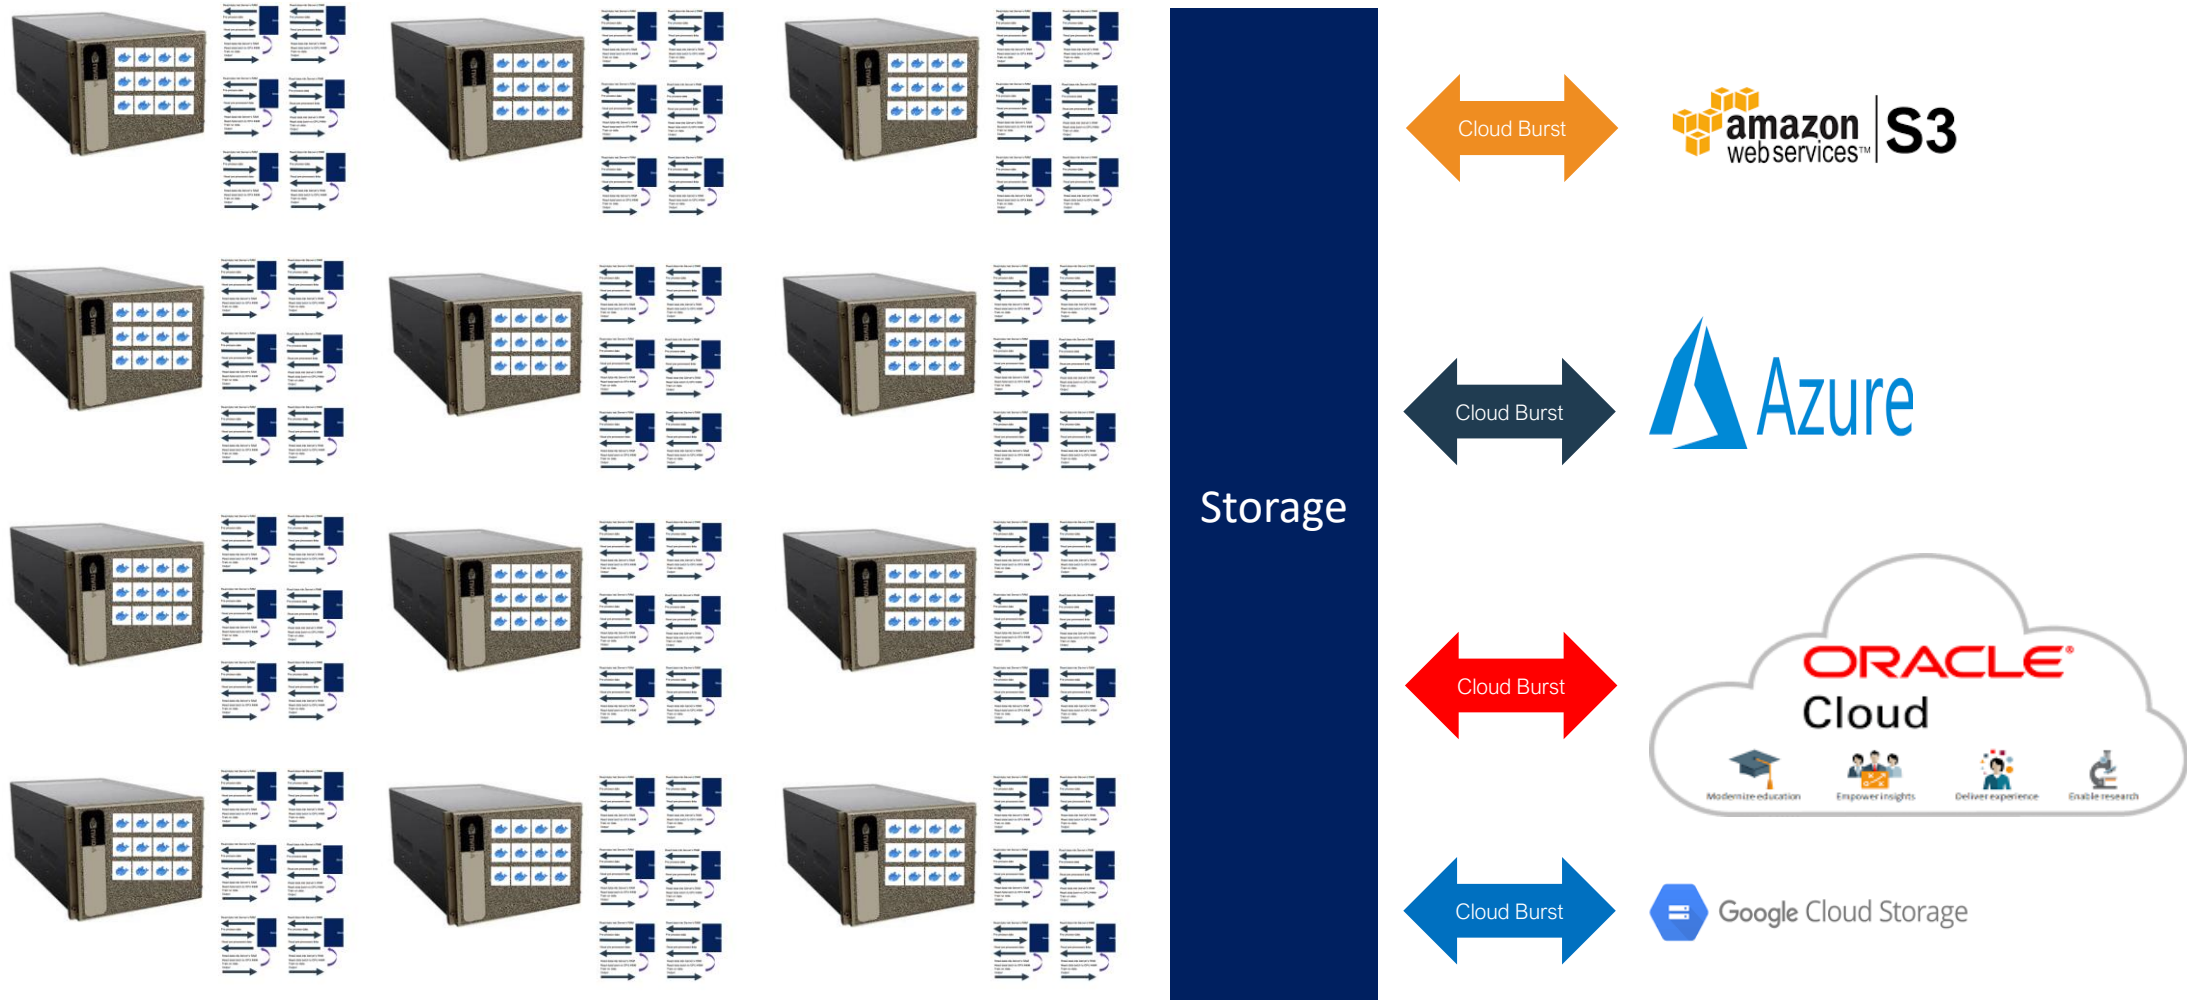
Datacenter A



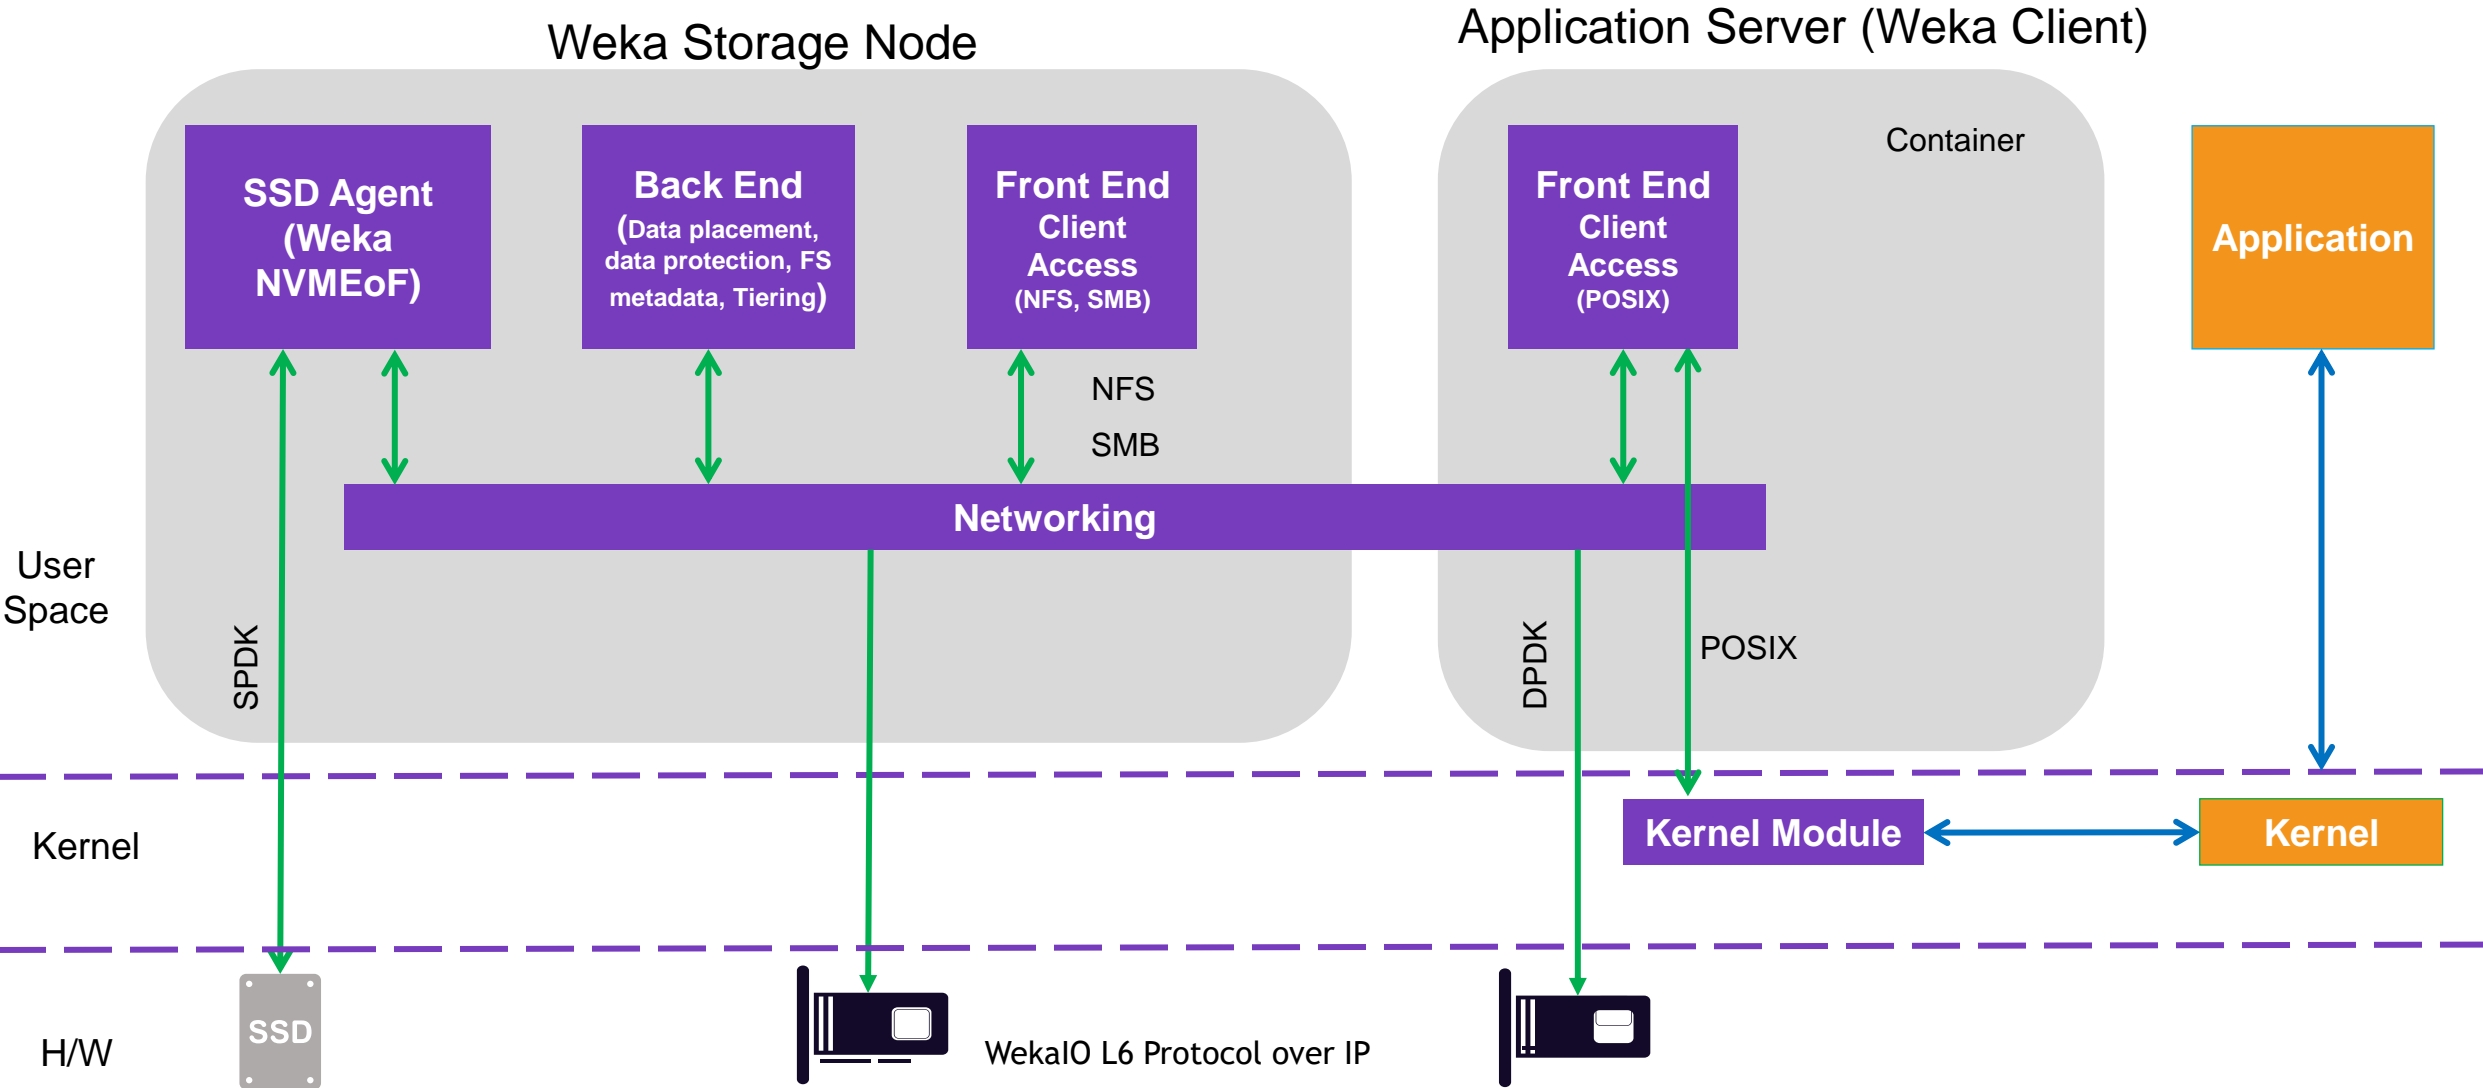
Datacenter B



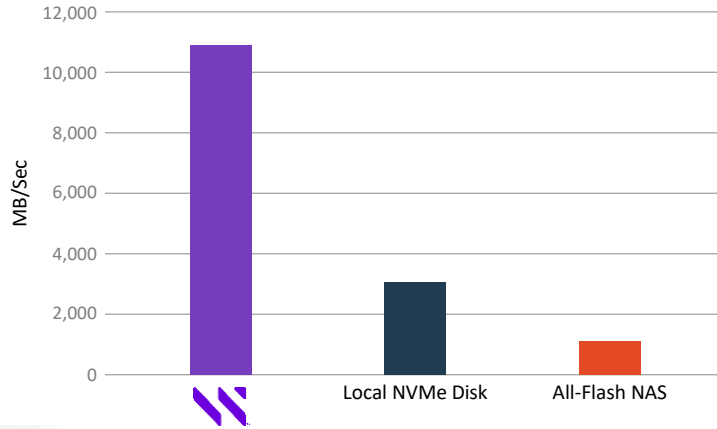
Multi DC/Cloud AI Training challenge



Weka Software Architecture



Single Client Performance Across 1 x 100Gb




HIGH-SPEED NETWORK SATURATION

162GB/Sec & 2 million IOPs Performance to a Single GPU client

MASSIVE SINGLE CLIENT PERFORMANCE

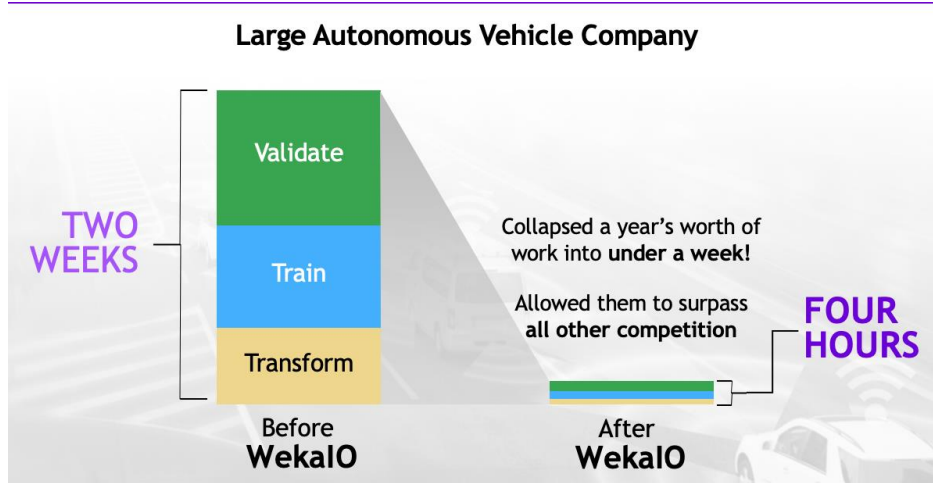
Scaling Performance From 1 to 9 DGX-1 Systems



PERFORMANCE SCALES LINEARLY TO MULTIPLE CLIENTS

AI/ML, DEEP LEARNING DATA PIPELINE FOR AUTONOMOUS DRIVING

- REDUCED DATA EPOCH RUNTIME FROM 2 WEEKS TO 4 HOURS VIA MODERNISATION OF NETWORKING & STORAGE WITH WekaFS



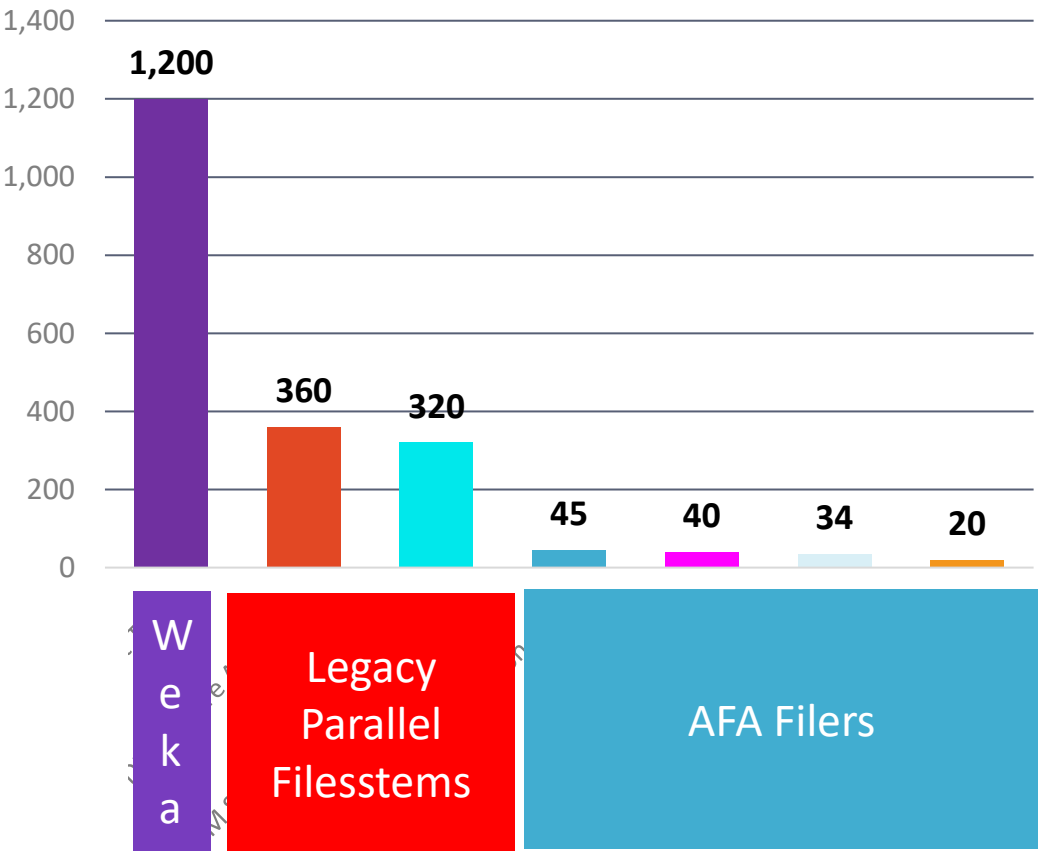
Benchmark	DAS (Optane SSD Server)	SAN (NVMe-oF)	NAS (All-Flash)	WekaFS (HPE NVMe Servers)
100T.YR1VWAB-12D-HO	15633	1886	4183	1028
100T.YR2VWAB-12D-HO	18114	1418	3294	892
100T.YR3VWAB-12D-HO	20730	1910	4773	1141
100T.YR4VWAB-12D-HO	24741	3317	7037	1550
100T.YR5VWAB-12D-HO	36888	22389	11376	4808
10T.YR2-MKTSNAP	176	355	6898	655
10T.YR3-MKTSNAP	176	358	7855	675
10T.YR4-MKTSNAP	149	375	8531	711
10T.YR5-MKTSNAP	155	393	8684	726
1T.2YRHIBID	645	374	1419	309
1T.3YRHIBID	1129	630	2737	480
1T.4YRHIBID	1957	1082	4881	804
1T.5YRHIBID	3234	1804	8589	1234
1T.OLDYRHIBID	61	46	129	48
1T.YR1VWAB-12D-HO	334	226	545	294
1T.YR2VWAB-12D-HO	394	268	632	355
1T.YR3VWAB-12D-HO	462	347	750	430
1T.YR4VWAB-12D-HO	553	517	928	547
1T.YR5VWAB-12D-HO	841	769	1298	732
50T.YR1VWAB-12D-HO	1089	1748	4302	2300
50T.YR2VWAB-12D-HO	1988	1774	4798	1971
50T.YR3VWAB-12D-HO	2865	2278	6253	2409
50T.YR4VWAB-12D-HO	4195	3118	8840	3077
50T.YR5VWAB-12D-HO	6731	4625	13597	4111
Average Result (lower is better)	5968.33	2166.958	5097.041	1303.625

FINANCE, TICK DATA ANALYTICS APPLICATION PERFORMANCE


- 4X FASTER THAN ALL-FLASH NAS
- 4.5X FASTER THAN DAS WITH OPTANE
- 1.6X FASTER THAN ALL-FLASH ARRAY

Comparative Storage Performance

Throughput Performance GB/s
(1PB Usable Capacity)



IOPS Performance
(1PB Usable Capacity)



Customer Stories

Genomics england

26x

reduction in storage
cost-per-genome

90PB Flash +
Object

Completely seamless to system users

50x

improvement in compute
workload speeds

Large Autonomous Car
company



1 week

on Weka was equivalent to
a year on the competition!

100sPB NVMe

Constantly adding live capacity

12K GPUs

World's largest
GPU Supercomputer



80x

Accelerate training epoch time,
enable increasing training
quality

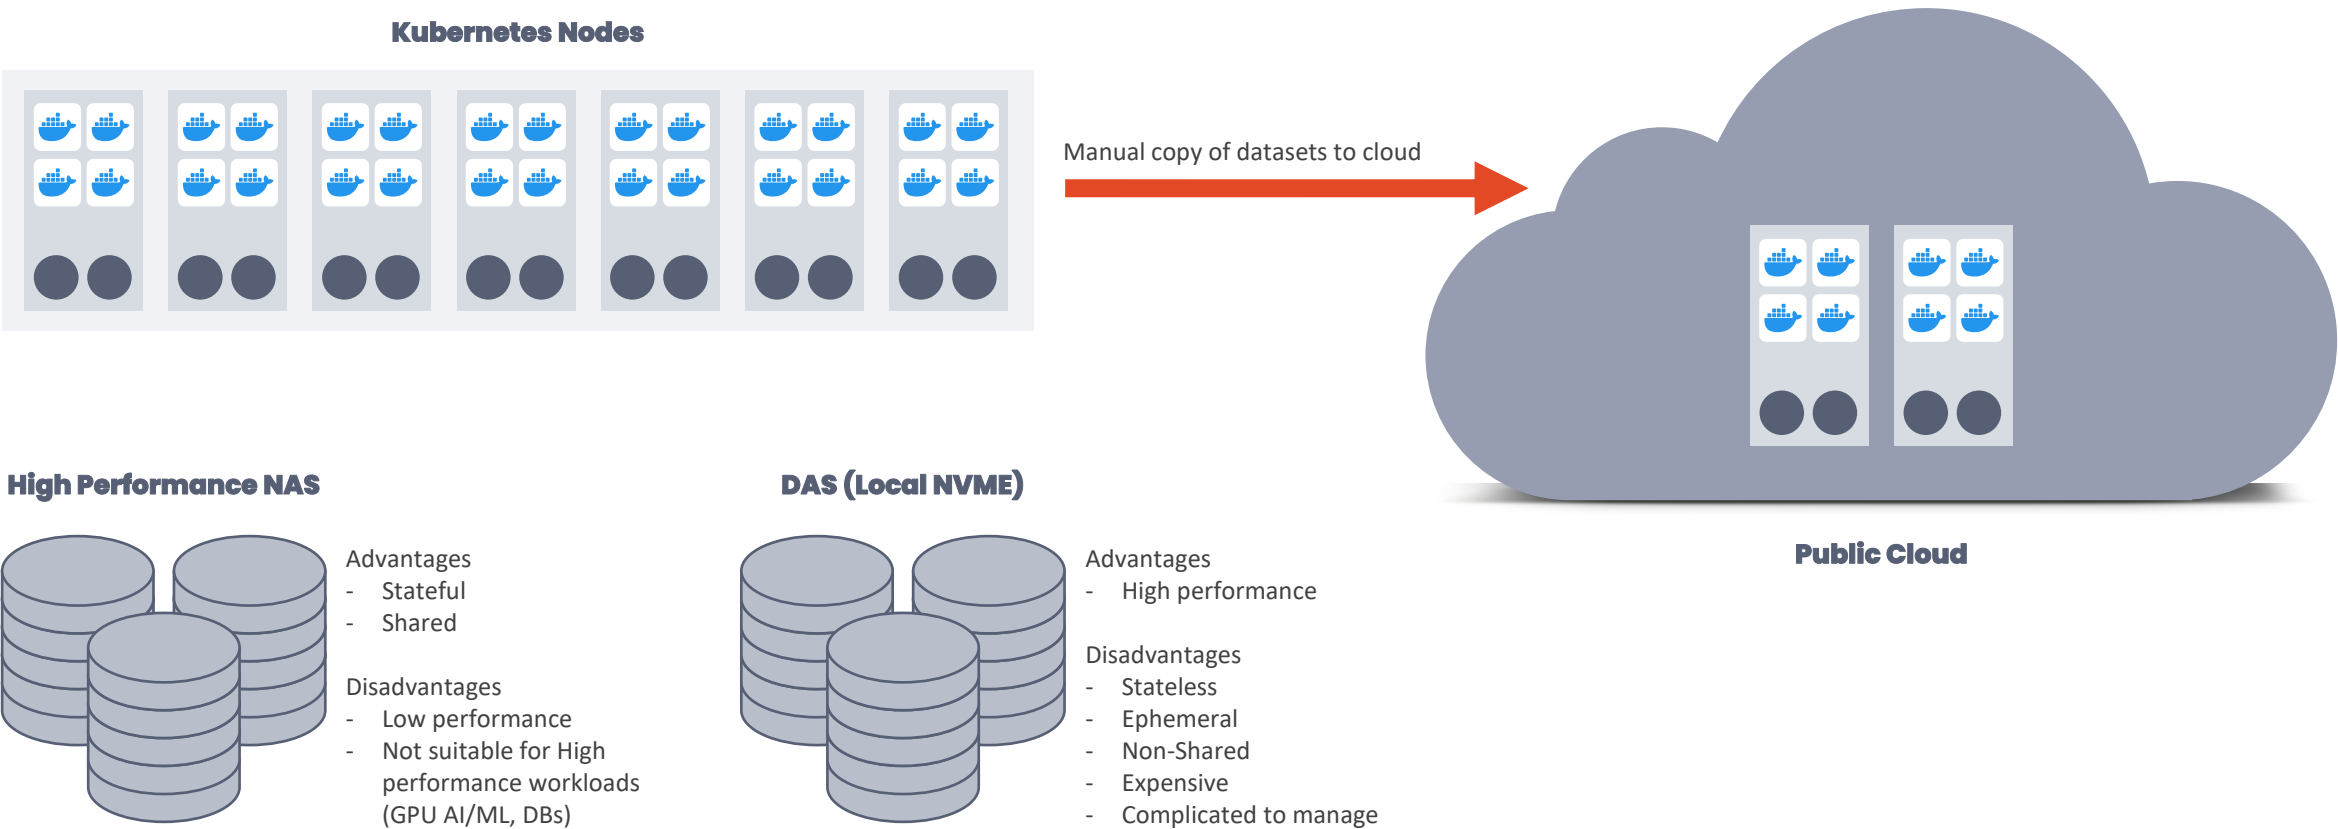
3PB Flash +
Object

reduction in the number of GPU's
required

5 Datacenters

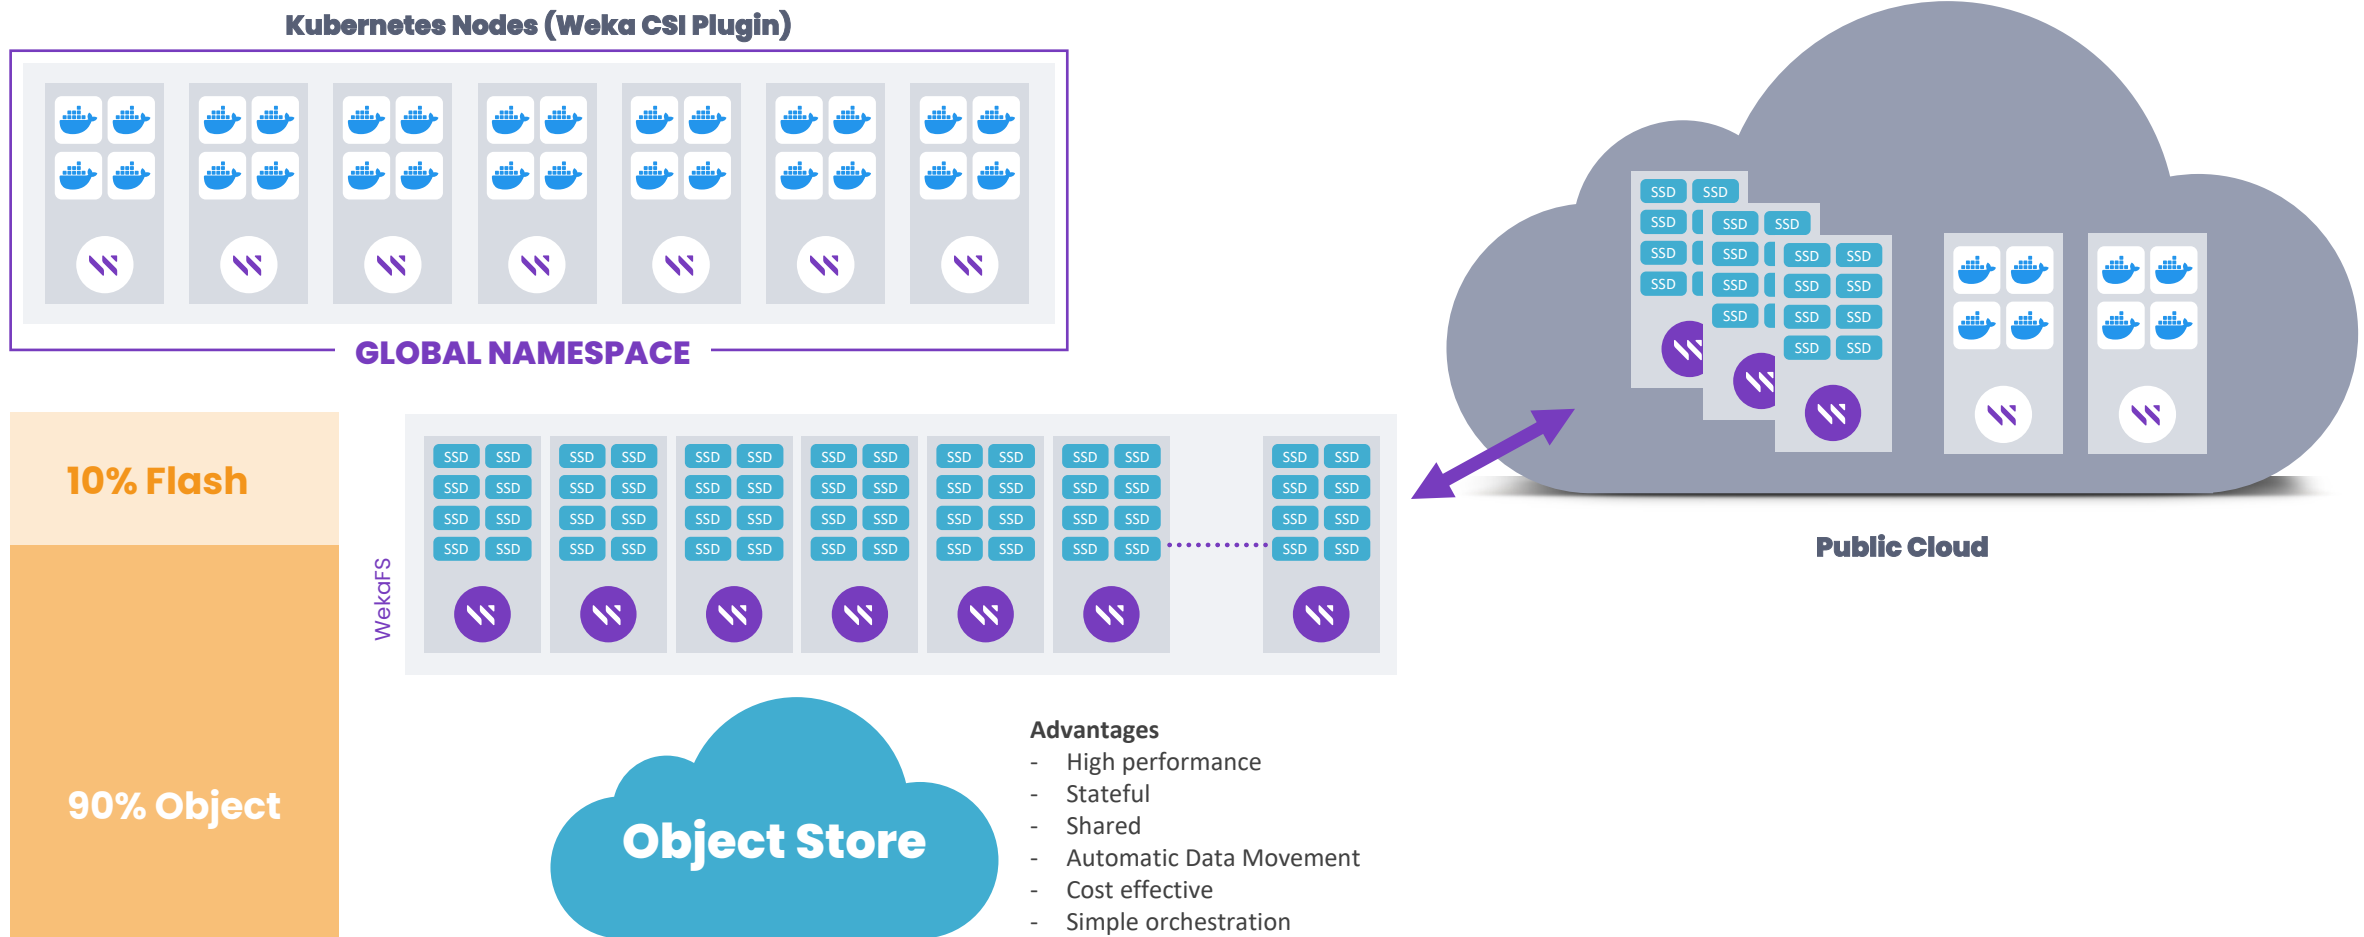
Traditional Kubernetes Deployments

Compromise between Performance and Agility



WEKA for Kubernetes


Performant and Agile – avoid data gravity



Computational Drug Discovery



Atomwise

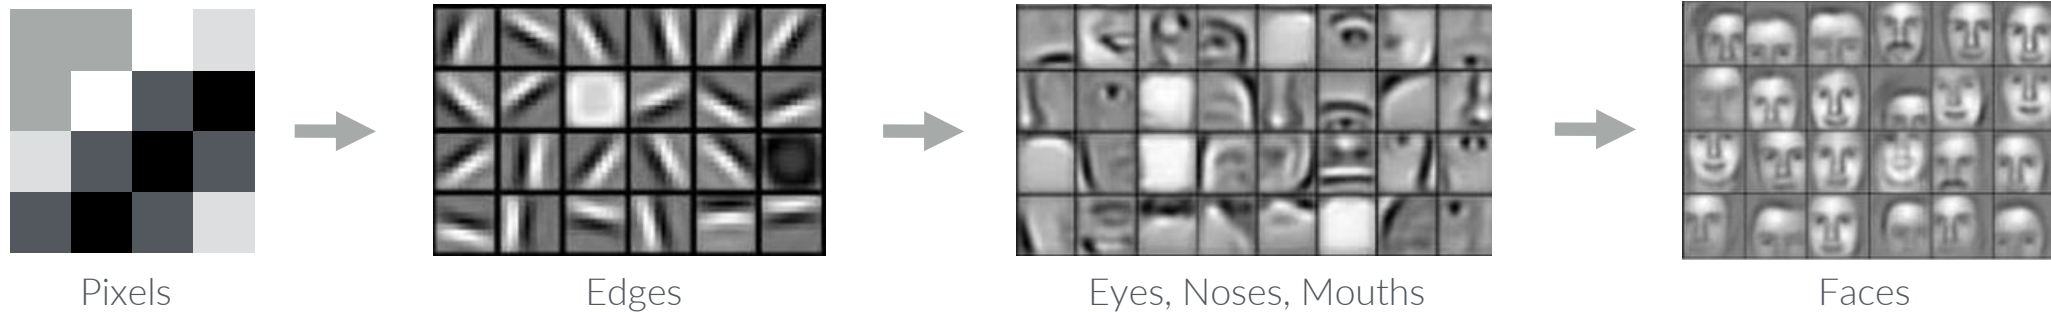


nsp15


(SARS-CoV-2019
endoribonuclease)

Structure-based drug design with deep learning

Convolutional neural networks for image recognition



Convolutional neural networks for molecular recognition



Weka Reduced Atomwise Epoch Time by 92% on AWS



Needed to feed billions of tiny files (<10KB) into expensive GPU instances



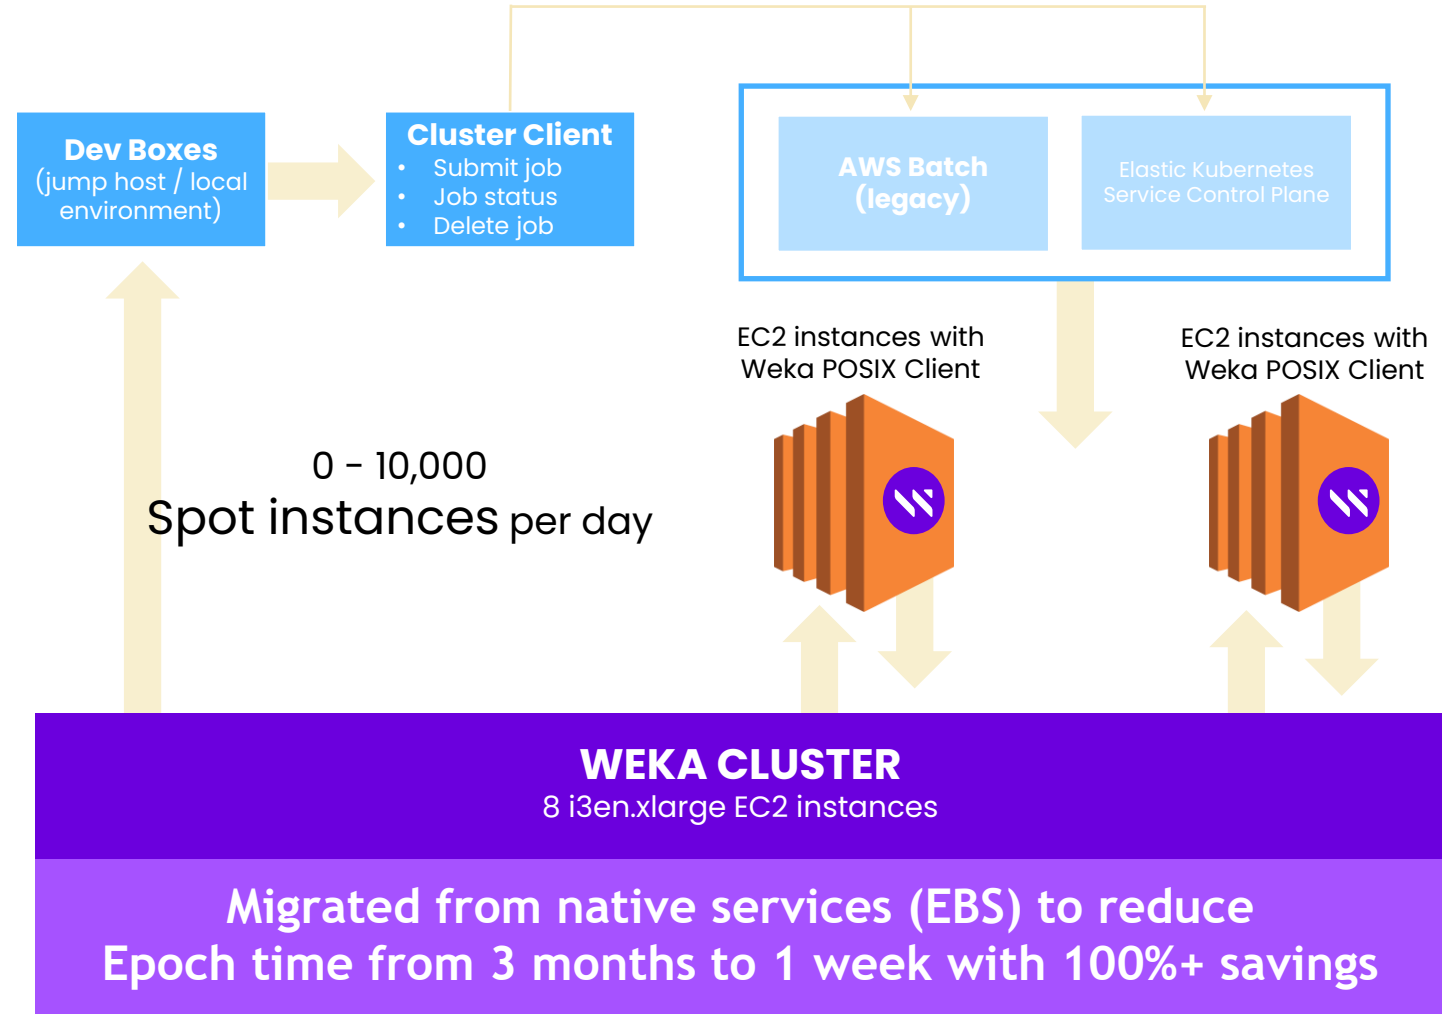
Needed to improve AI Epoch times (was taking 3 months)



Needed a less costly solution that eliminated the need to copy the same data to multiple EBS instances



Needed to integrate with their Kubernetes pipelines

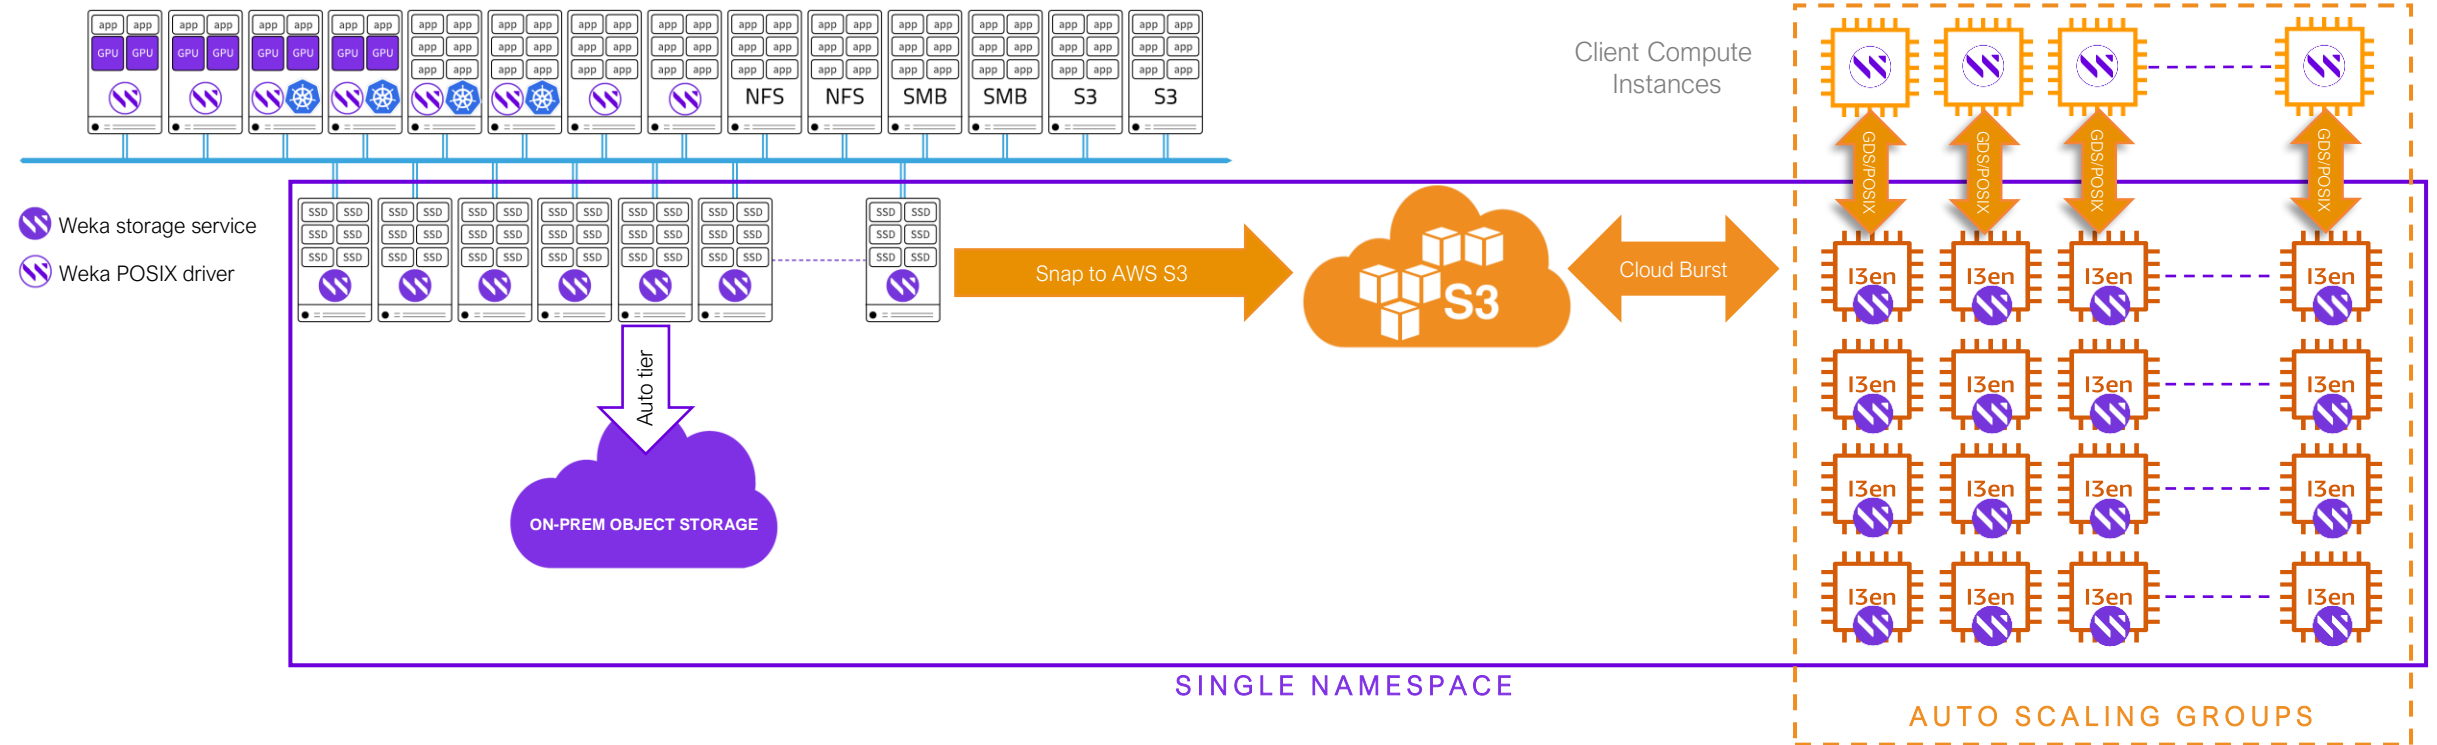


<https://youtu.be/kAYzLoIWTJY>

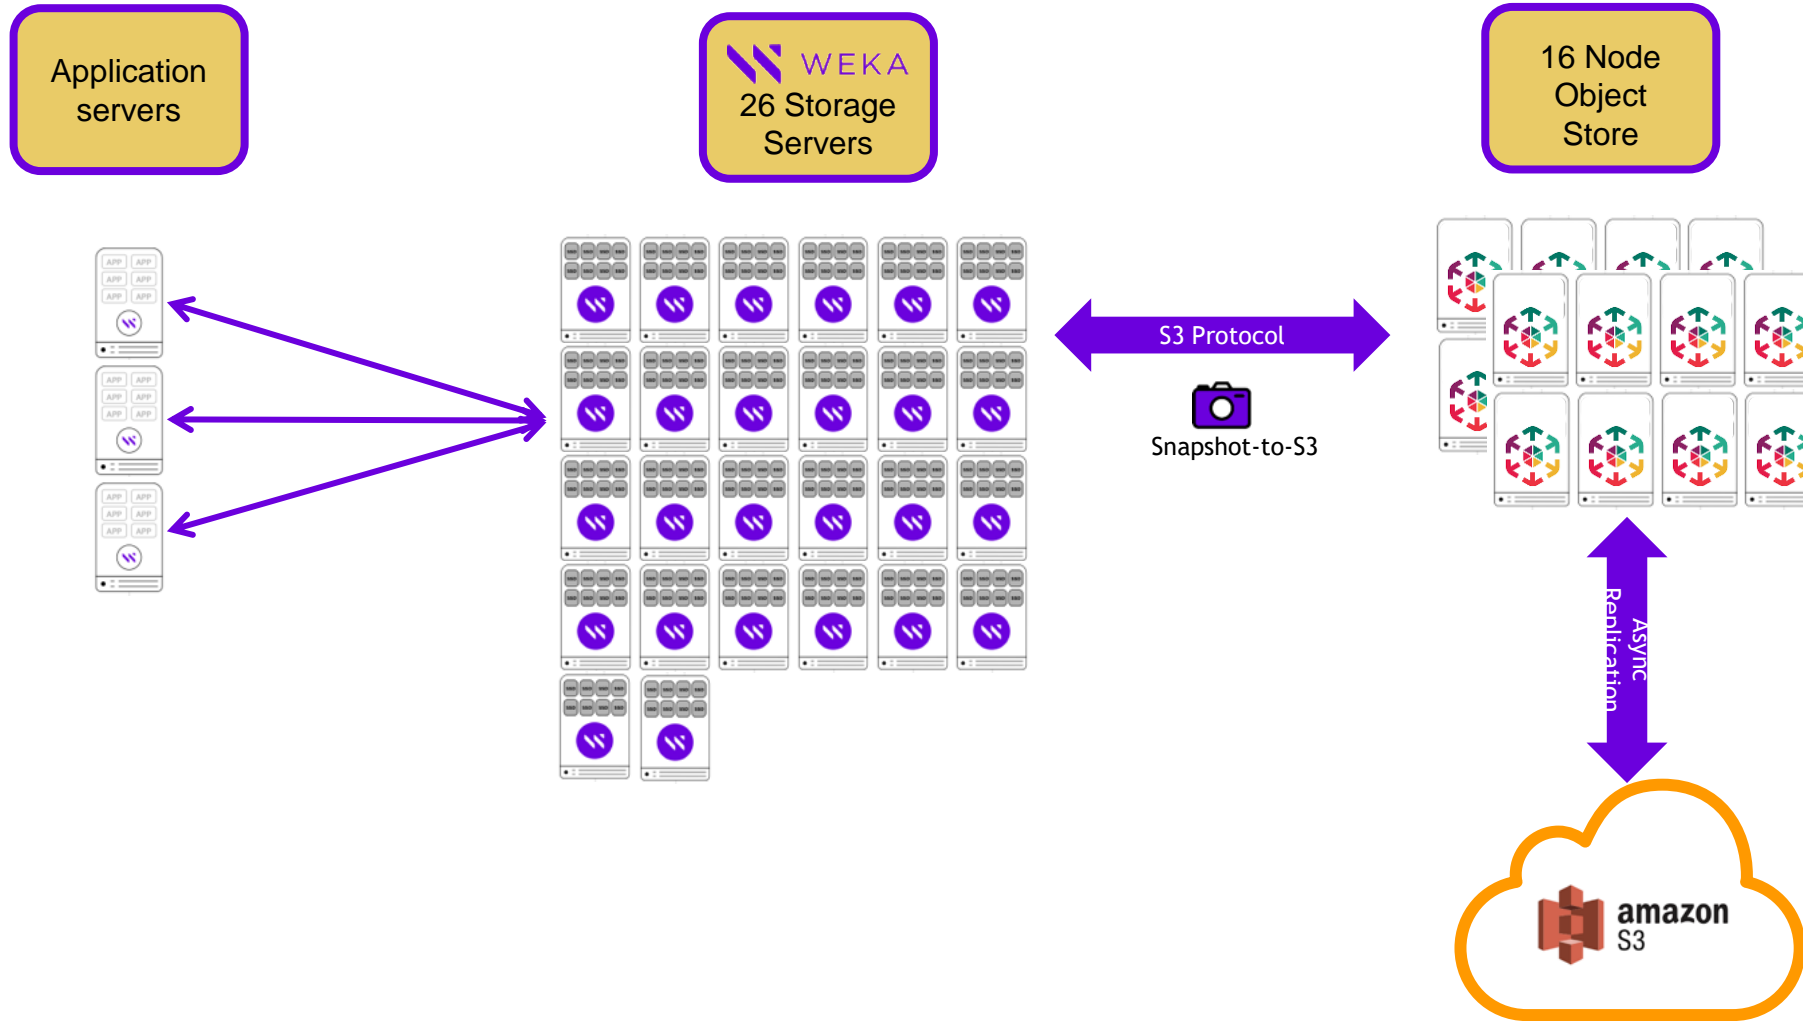
Flexible Deployment Options

DEDICATED INFRASTRUCTURE

PUBLIC CLOUD INFRASTRUCTURE



Boston Pharmaceutical HPC Solution



One solution for many applications

Scientists never wait for technology

Integrated backup

AWS is now DR and BC

Drastic improvement in availability and performance

Technology is now an enabler



Genomics England : WekaFS Scale-Out to Data Lake



Single shared namespace to users and applications




All data accessed through NVMe



Authenticated mounts, LDAP, Active Directory, Extended ACLs, Encryption in Flight



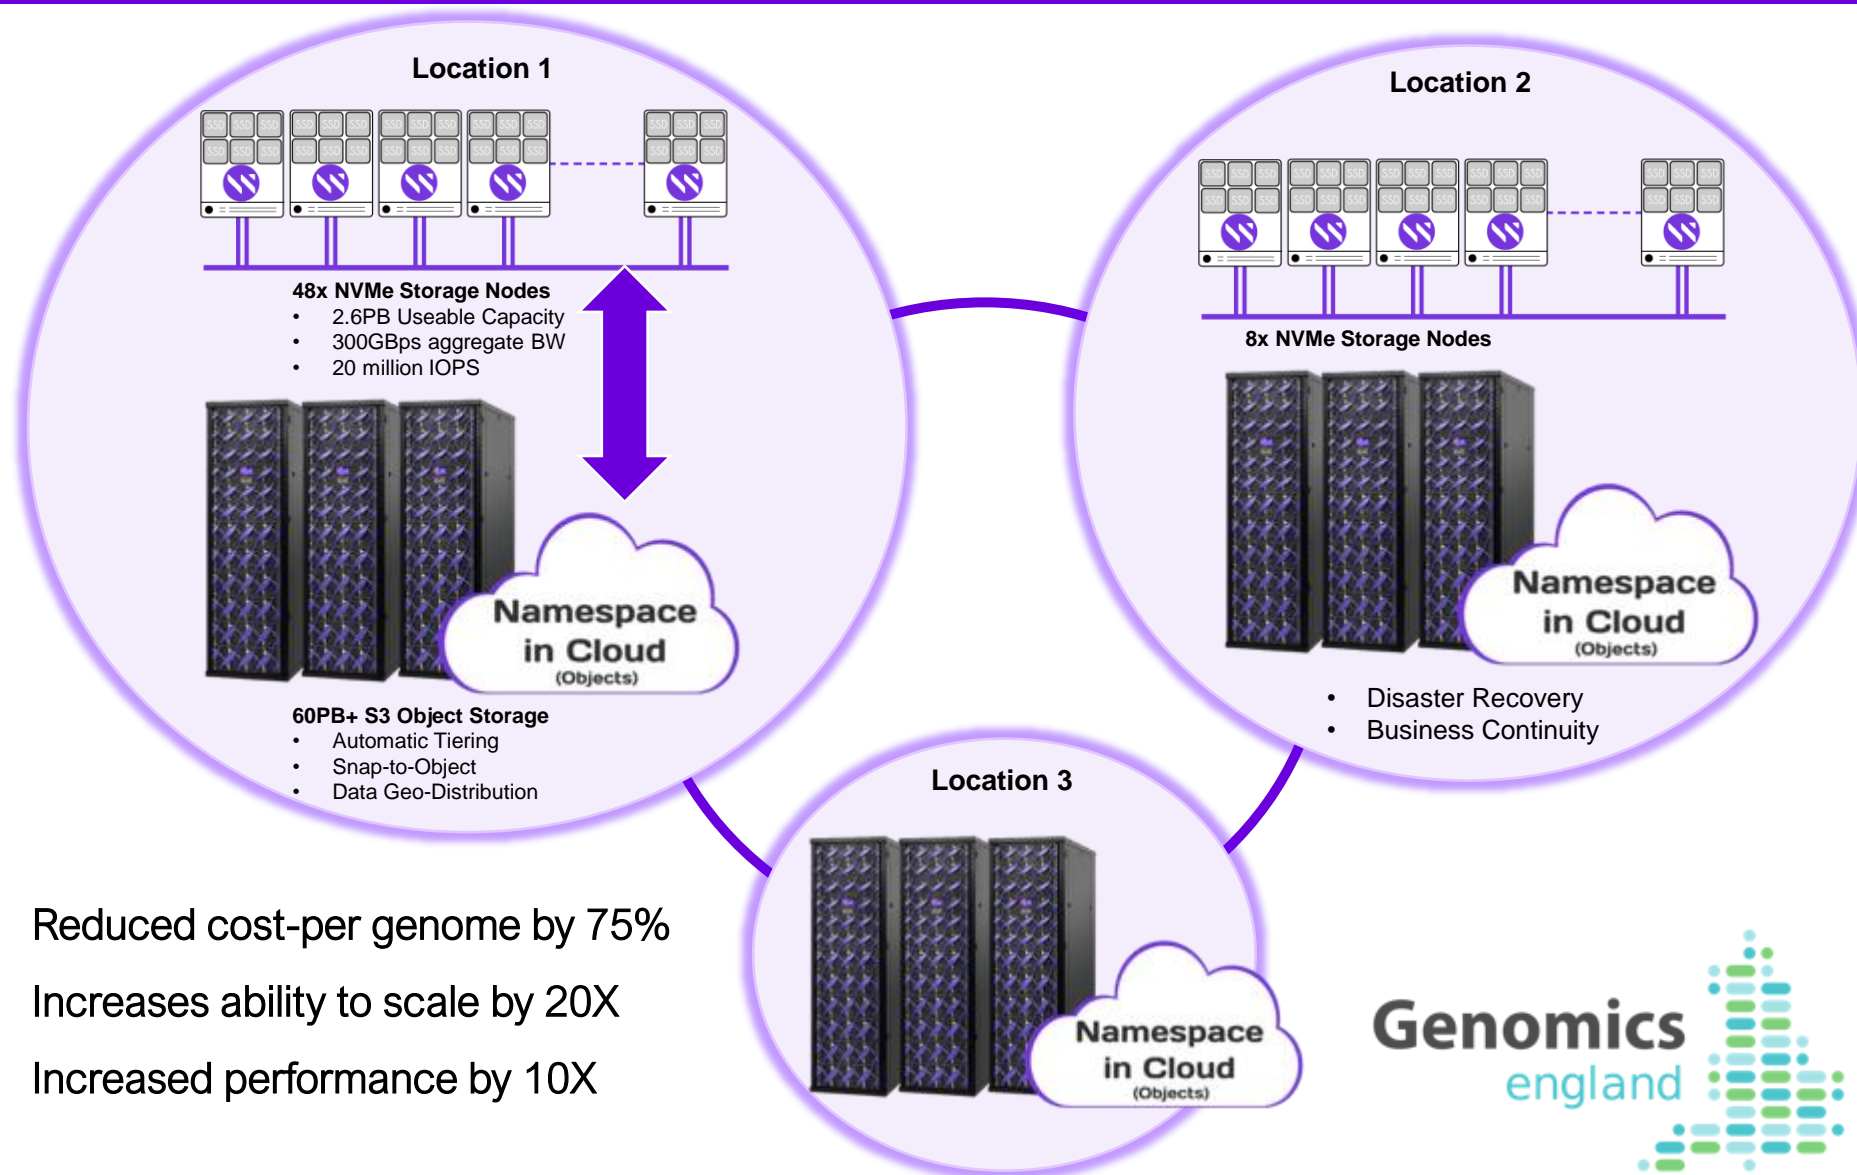
Weka cluster at site 2 for DR and business continuity



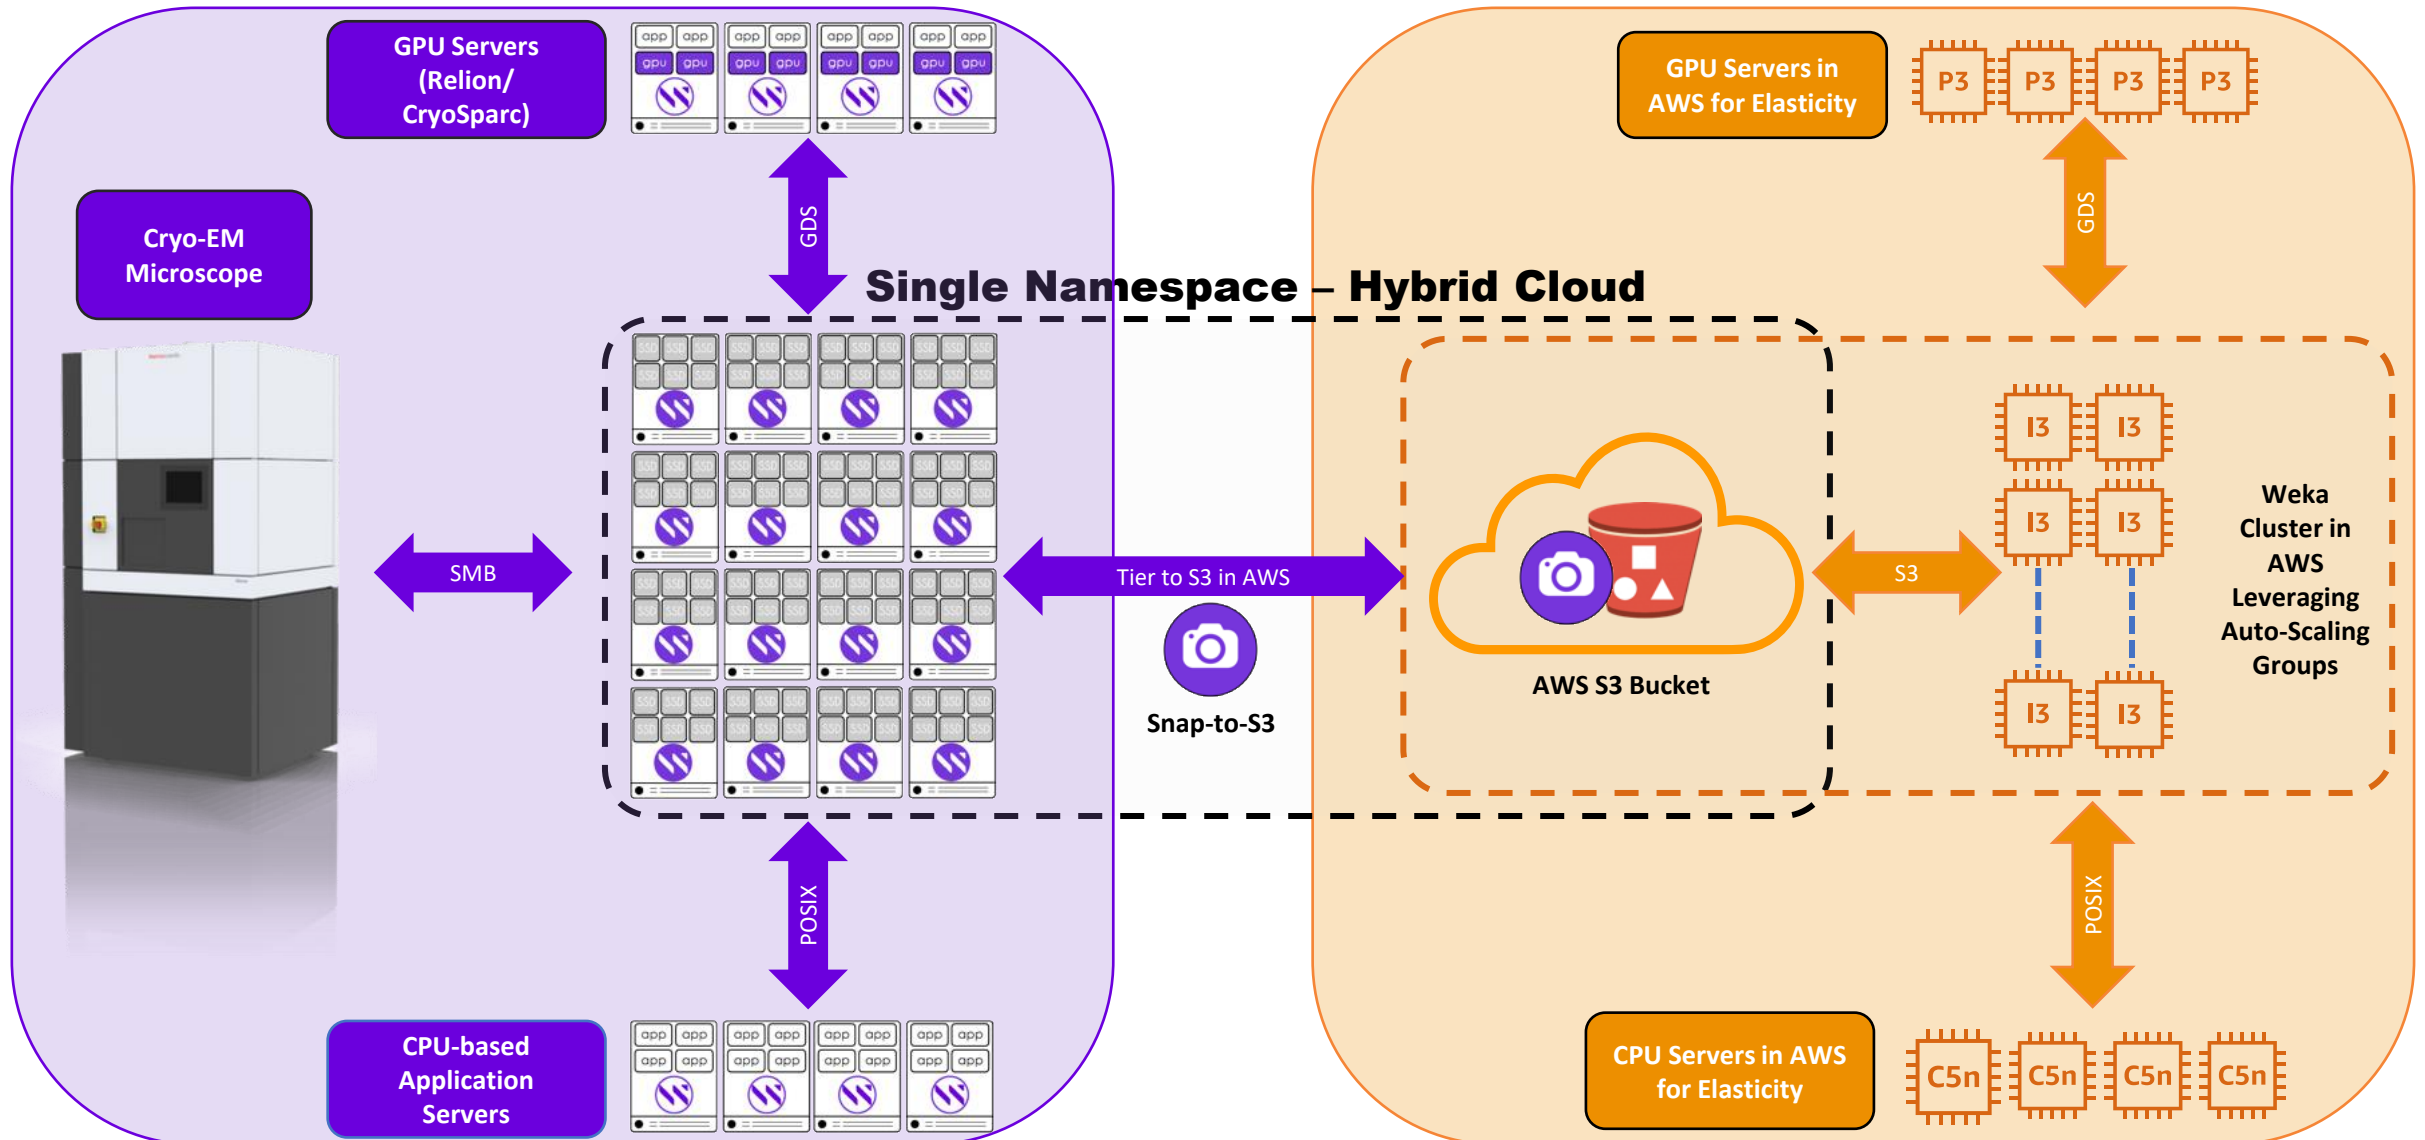
NVMe tier will scale to 5PB & Object tier to 160PB within project lifetime



Over 2,500 Researchers access data for drug discovery & COVID research



HLS Hybrid Cloud Workflow






WEKA




Simplifying AI With Weka

Shimon Ben David
Chief Technology Officer



Top Use Cases That Companies are Investing In

n=1181




Companies are investing in multiple initiatives in AI

GPUs Have “Densified” Compute into a Single Server Creating a Huge Data Bottleneck

Click to edit Master text styles




5x Performance
0.1% of the Space
I/O per Client is Off the Charts




10,400 Rack Units
1 PetaFLOP
CPU-Only Servers
100's of servers with CPUs
100's of low bandwidth network connections
No one server was particularly demanding on storage

6 Rack Units
5 PetaFLOP
GPU Servers
8 NVIDIA A100 Tensor Core GPUs


The AI Challenges



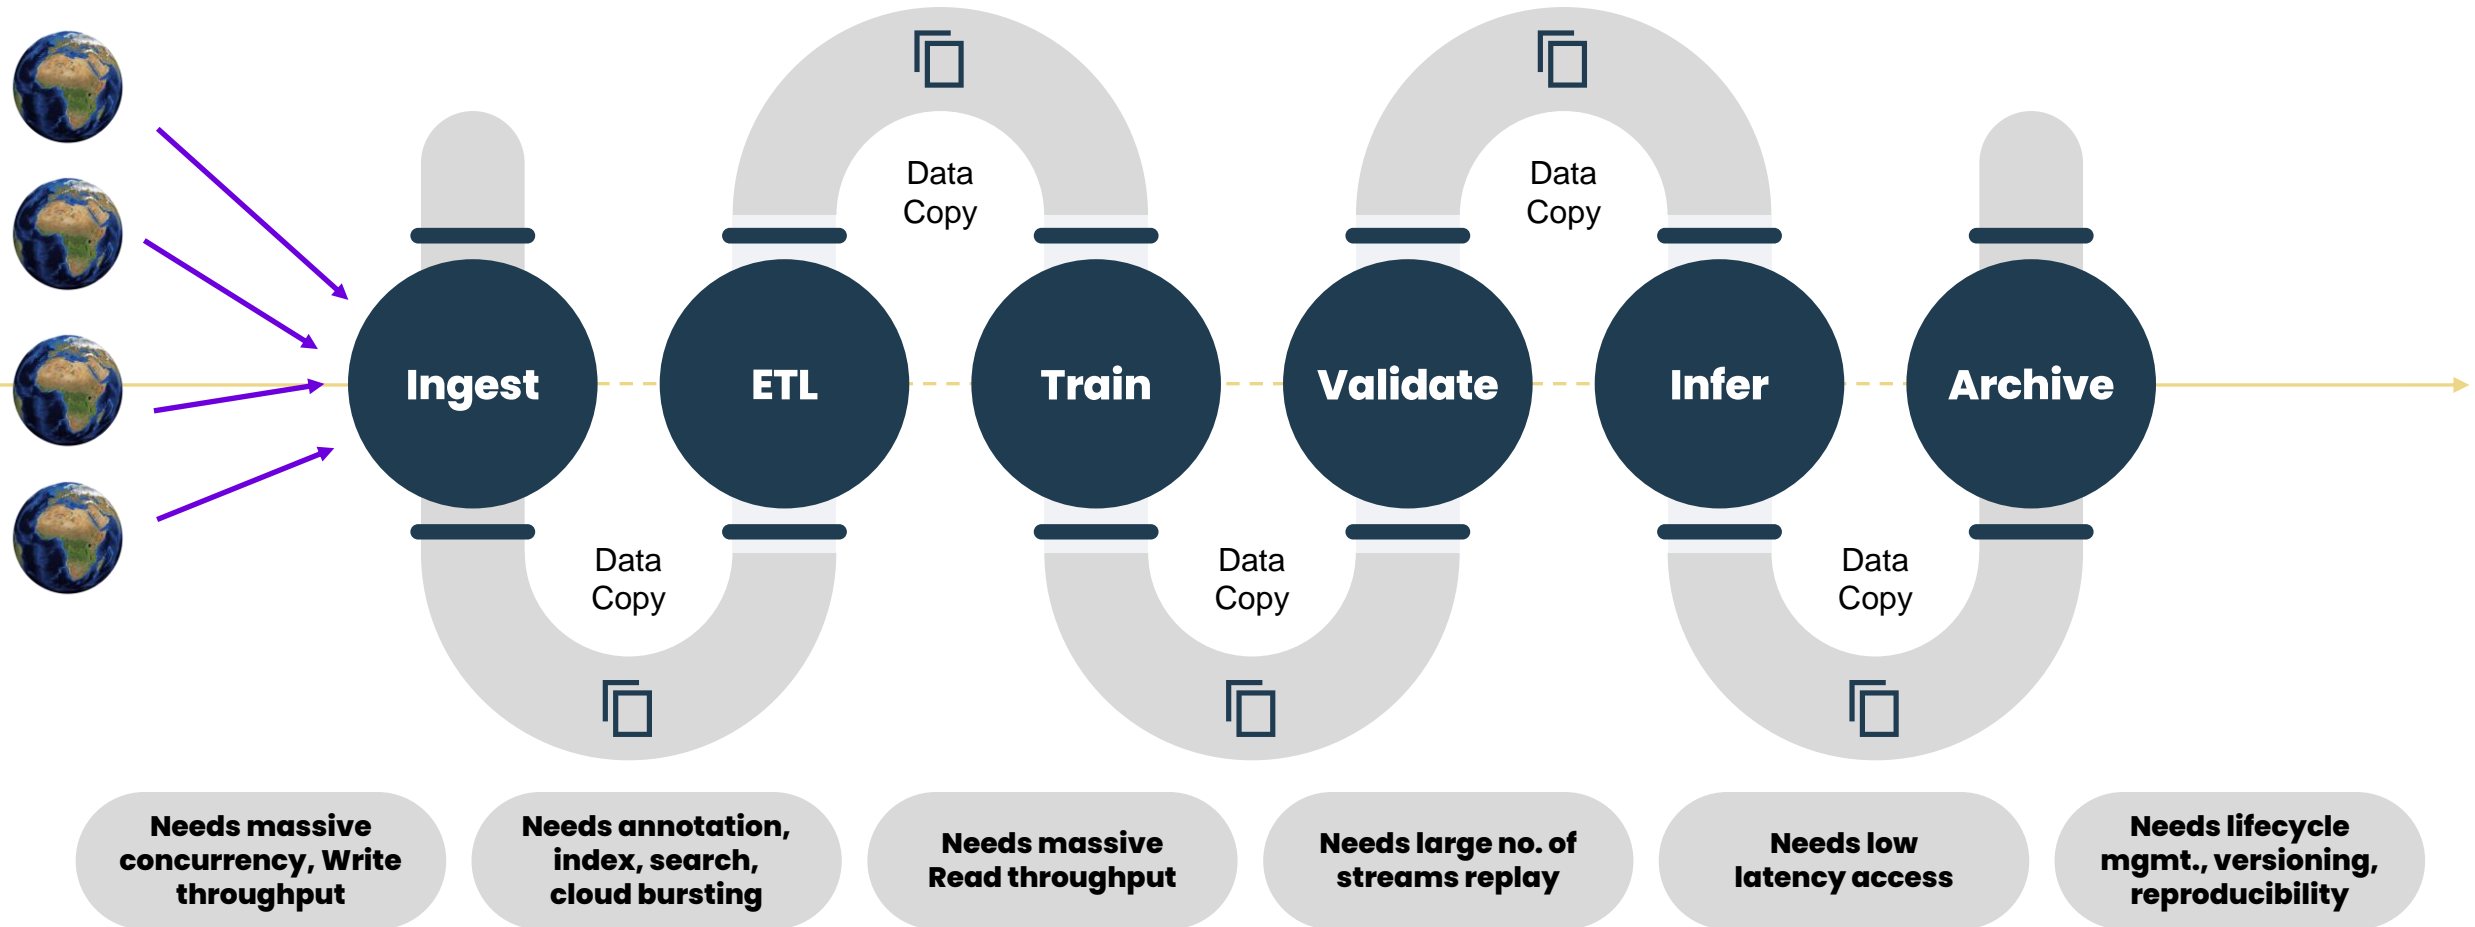
- Compute environments improved considerably in order to accelerate AI
 - Nvidia
 - Cerebras
 - Graphcore
 - SambaNova Systems
 - Tenstorrent
- MLOPs environment and frameworks improved considerably
 - Ubiops
 - Run:ai



Neural Network




Ingest Challenge



Pre-Processing – the best known secret of AI

Can take 50% of a training epoch's time

- Pre-Processing is manipulation of the accumulated data to a state that is expected by the AI Model (e.g. image resize, contrast changes, etc...)
- Within an organization there are huge variations between the pre-processing steps each organization performs – even within an organization every researcher might need to pre-process the data differently
- Pre-Processing can consist to 50% or more of the training epoch time
- **The IO Implications are massive reads and writes with varying changes to raw accumulated data – WRITE CACHE MISS**

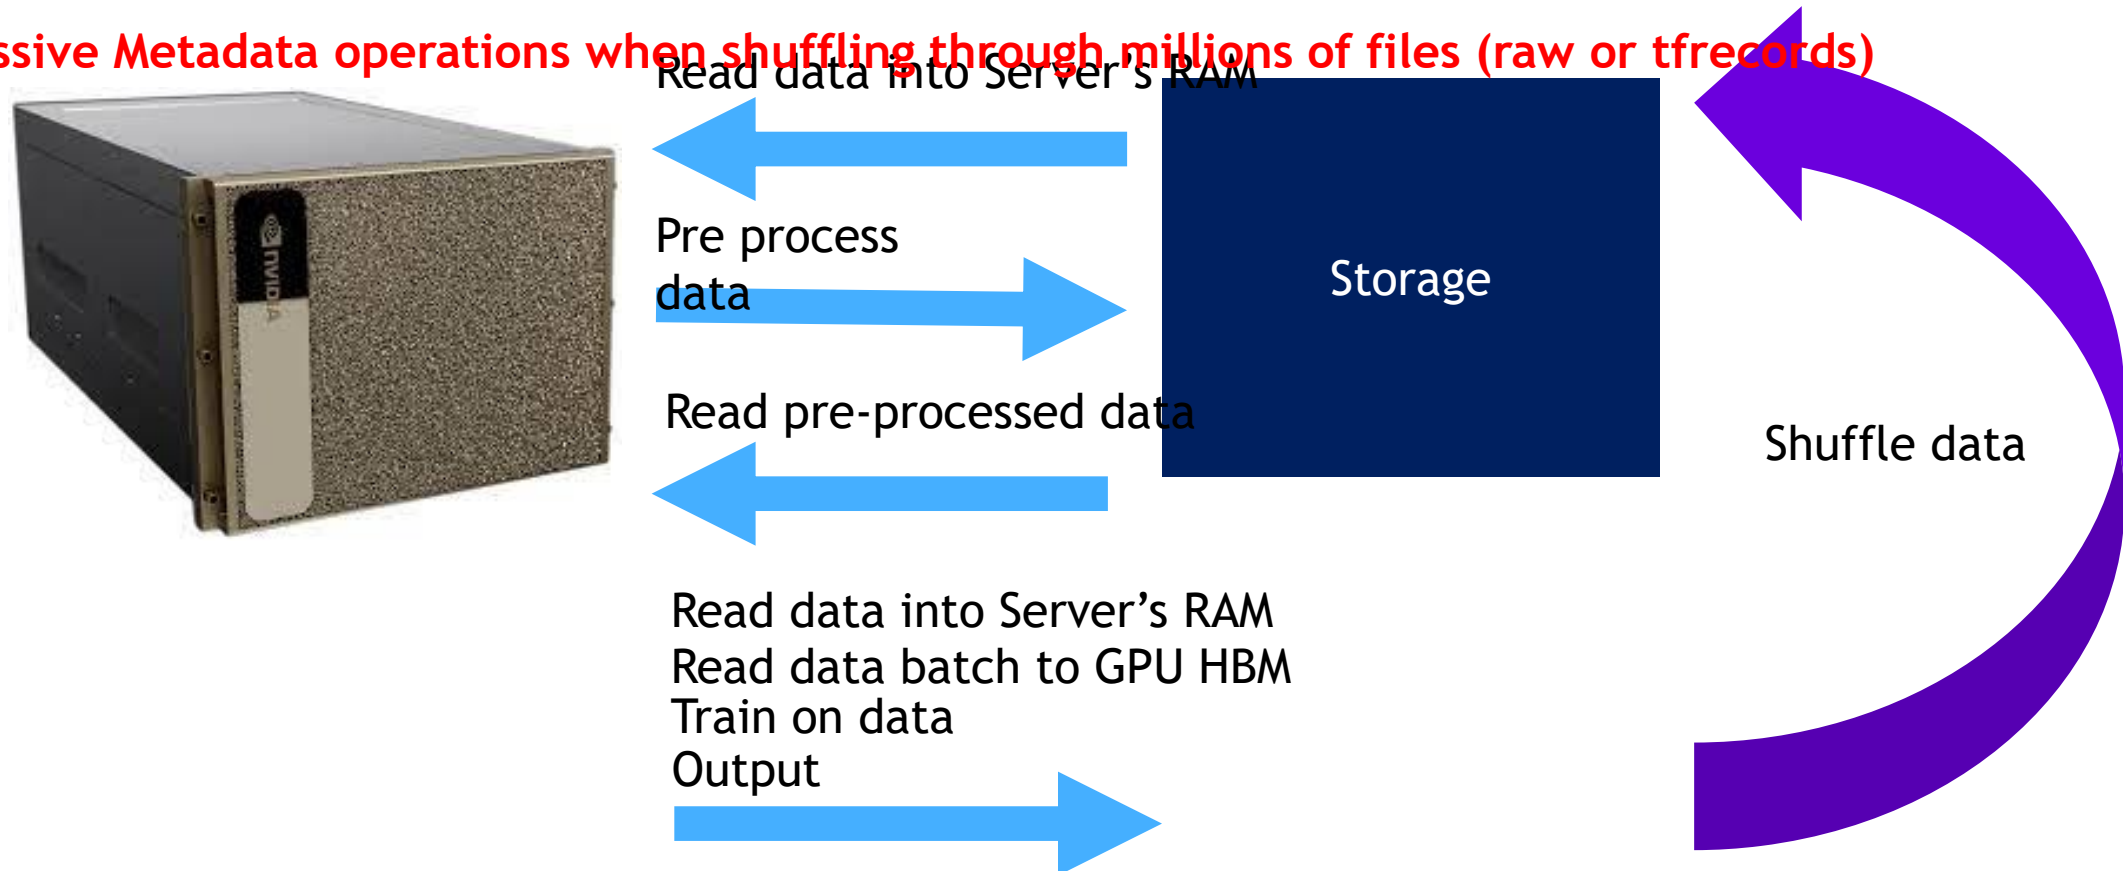


Training your model – Small data sets (Day 1)

Small data sets can fit into GPU(s) memory and eliminate storage latency

Small data sets pre processing will still require reading and writing from and to the storage

Massive Metadata operations when shuffling through millions of files (raw or tfrecords)

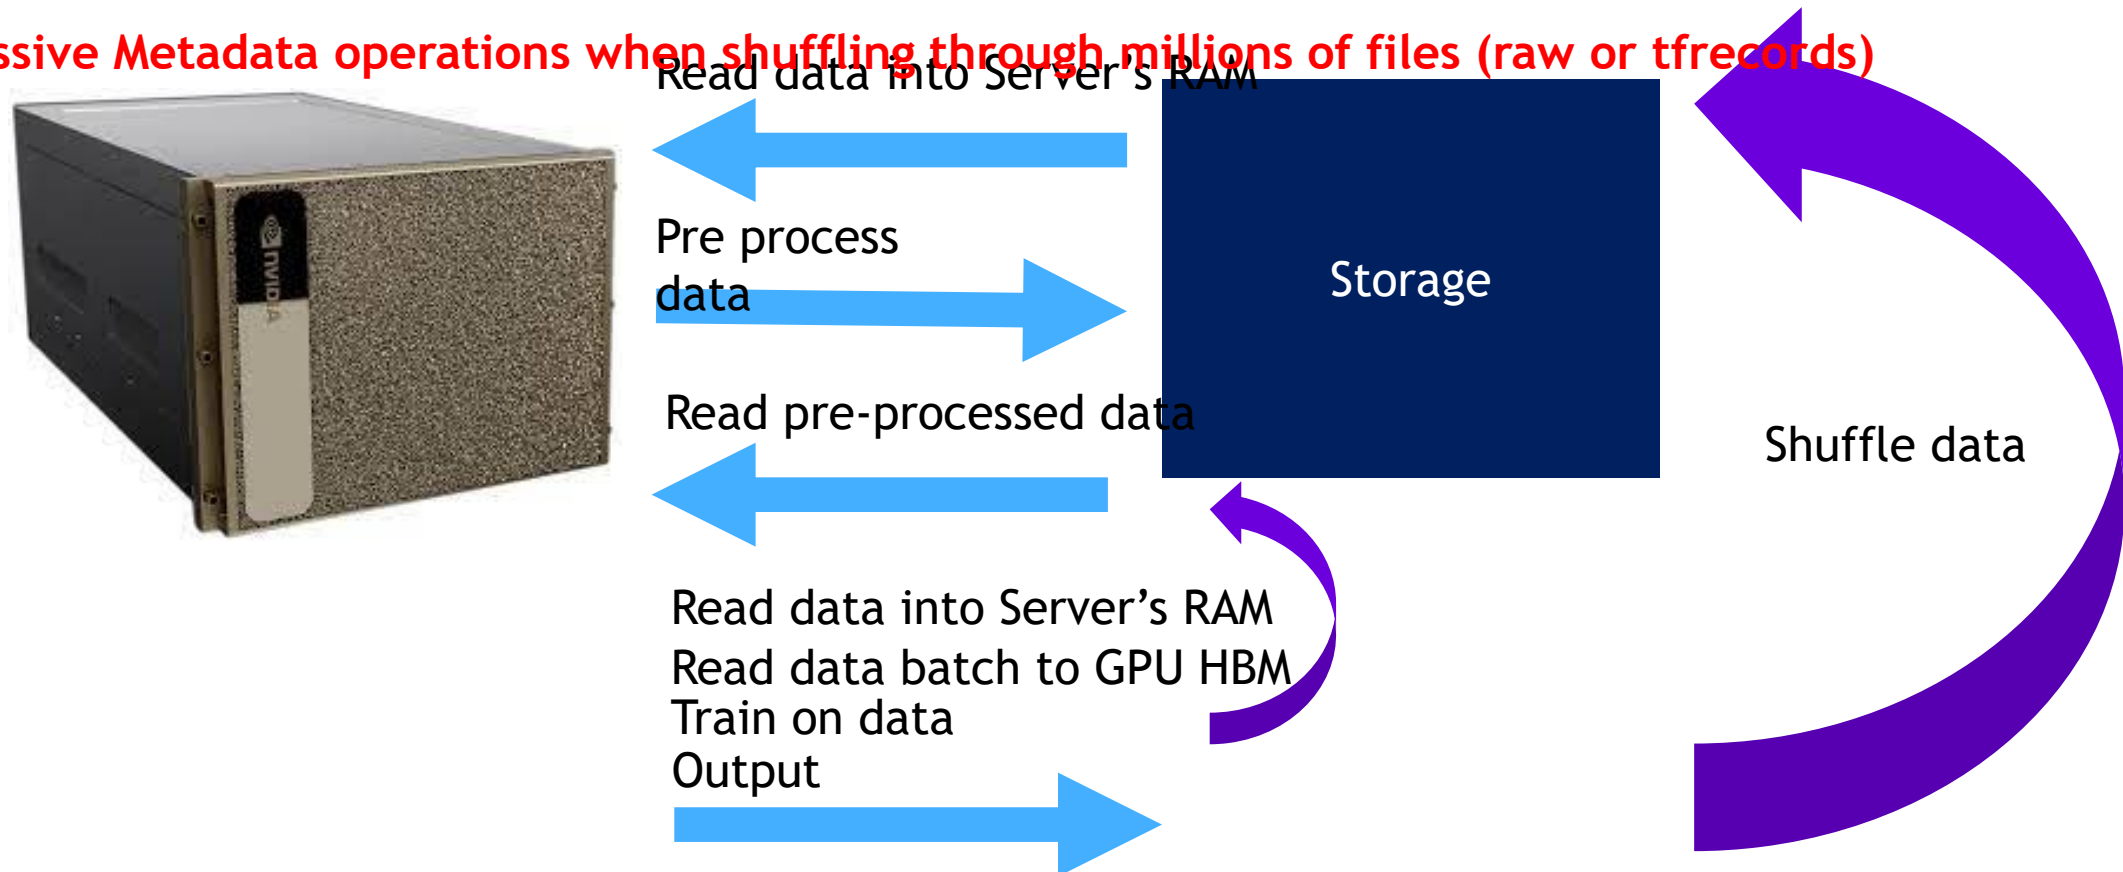


Training your model – Large data set (Day

Larger data sets can not fit into GPU(s) memory - double buffering will try to keep up with reads

Larger data sets pre processing will still require reading and writing from and to the storage

Massive Metadata operations when shuffling through millions of files (raw or tfrecords)



Training your model – Large data set – Real life

Larger data sets can not fit into GPU(s) memory - double buffering will try to keep up with reads

Multiple jobs are running concurrently at different stages - heavier RAM load - results in more READs from storage

Small data sets pre processing will still require reading and writing from and to the storage

Massive Metadata operations when dealing with millions of files (raw or tfrecords)



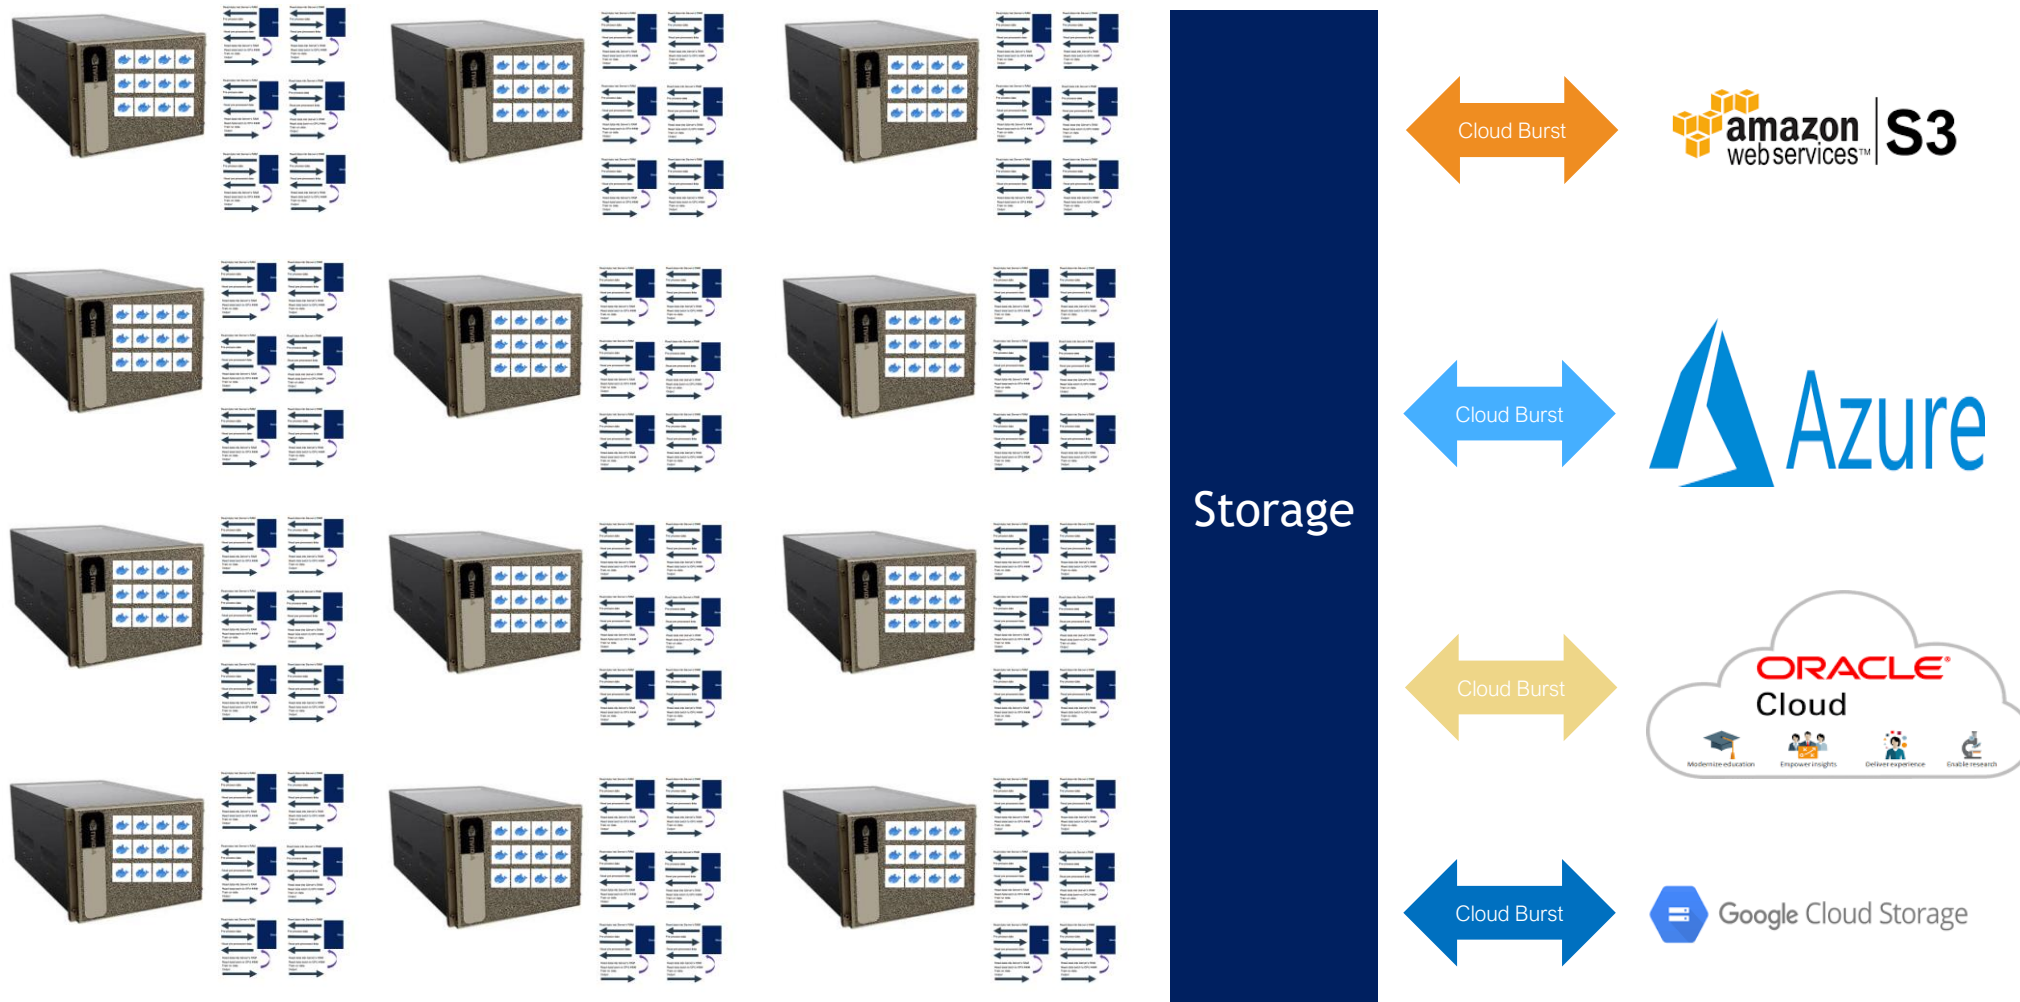
Storage

Training your model – superpod data set – SuperPod Like configurations

Compounding on the challenges of a single GPU server - The storage now needs to accommodate for thousands of Containers and GPUs




Multicloud Metadata challenge



Deep Learning IO Doesn't Look Like This:

Vendors, Emphasize Throughput, Which is only part of the problem



Deep Learning IO patterns



Lots of Random Reads of Lots of Small Files

■ The Deep Learning IO Process

- Mini batch - iterate over random subsets of training data
- Train on each mini batch
- Epoch - process the entire dataset in random order
- Hyper parameters control training (e.g. precision, #epochs, etc...)

Deep Learning IO Characteristics

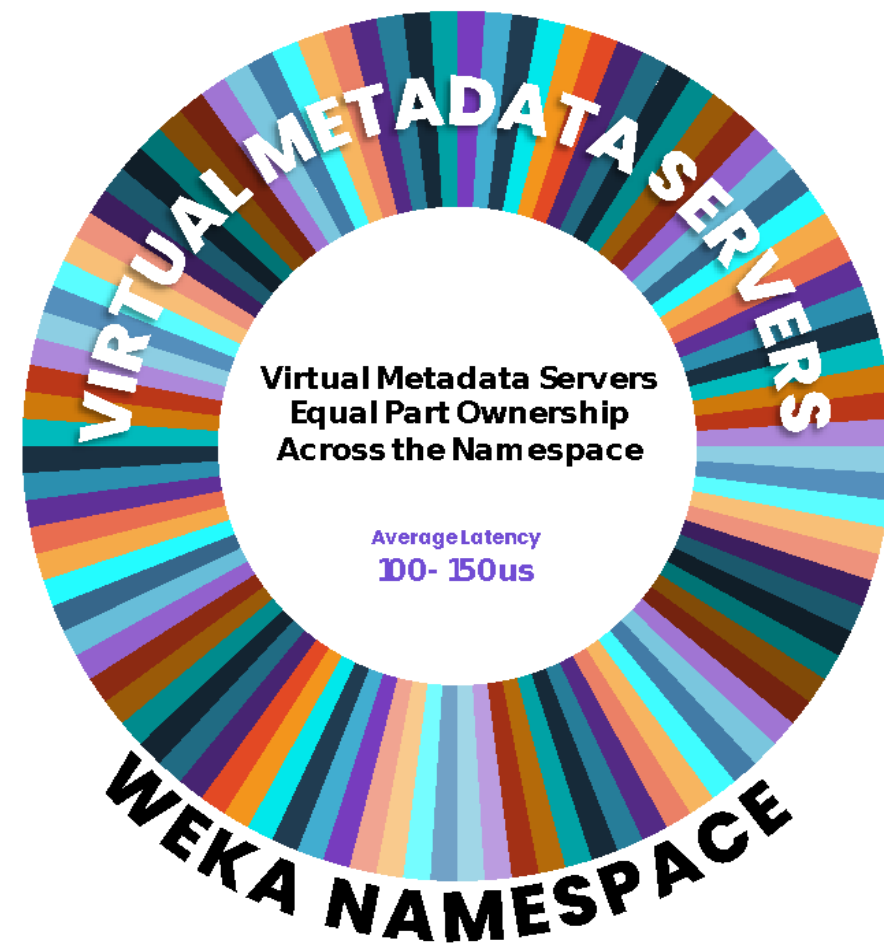
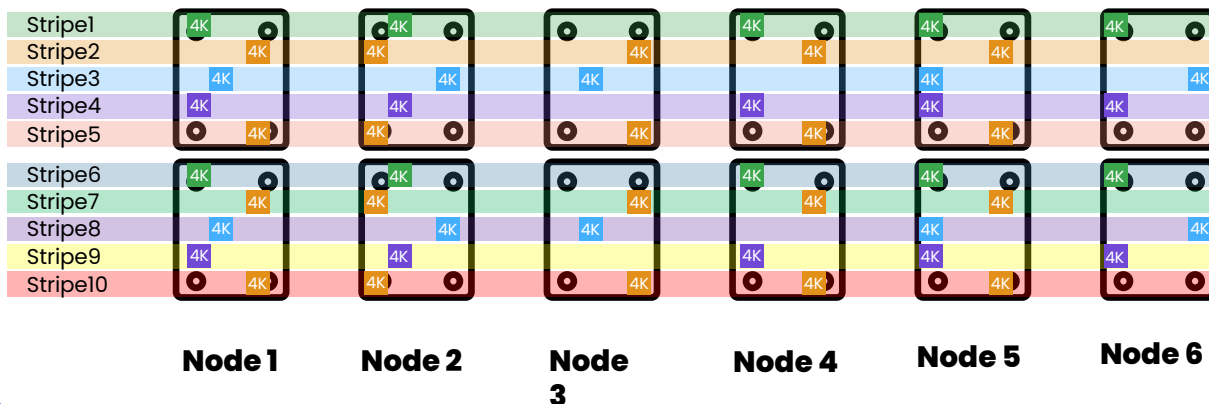
- Dominated by **many, many small IO requests**
- **Huge metadata overhead**



Data Distribution & Metadata Handling

- Every compute node runs some number of buckets - or virtual metadata servers (MDS)
- Each bucket is responsible for an equal shard (1/Nth) of Weka filesystem namespace
- Data fully distributed and written in 4K blocks across pool of NVMe SSDs - no 'Hot Spots'
- Environment is aware of SSD responsiveness, prioritizing writes to SSDs w/lowest queue depth
- Every 8MB-range of every single file handled by a different virtual metadata server
- Virtual MDS use Weka's block layout to persist the journaling of the metadata
- Buckets themselves can failover between compute nodes

1MB file = 256 x 4K blocks



Summary – Why Weka with AI/ML Workloads?

Enterprise AI workloads have new DATA challenges that are unlike previous enterprise workloads

- Performance - Train 40-80 times faster
 - Heavy large reads
 - Heavy writes
 - Small IOPs
 - Massive metadata
- Data retention
 - Explainable AI - Save models data points (AT SCALE)
- Data Mobility - integrate with MLOps environment
 - Avoid data gravity (AT SCALE)
 - Burst between data centers / clouds
 - More GPUs
 - Lower price

Source Text Placeholder

■ Weka is the most efficient platform to feed AI / GPU Pipelines



WEKA® proprietary and confidential | 2020