



# The Data Platform for the **Cloud & AI Era**

---

WEKA Korea

김승훈 이사장 ( [chris.kim@weka.io](mailto:chris.kim@weka.io) )

Jun 2023



# The Data Platform for AI.








**Data, Lots of Data**



**DNN Models**



**Compute Power**



**Cambrian  
Explosion  
of AI**



**“By 2025, AI will be the top category driving infrastructure decisions, due to the maturation of the AI market, resulting in a **tenfold growth** in infrastructure requirements”**

– Gartner, December 2021

## What Is Artificial Intelligence?

Artificial intelligence (AI) applies advanced analysis and logic-based techniques, including machine learning, to interpret events, support and automate decisions and take action.

### How does artificial intelligence work?

**Here at Gartner, we define artificial intelligence (AI)** as applying advanced analysis and logic-based techniques, including machine learning (ML), to interpret events, support and automate decisions and to take actions. This definition is consistent with the current and emerging state of AI technologies and capabilities, and it acknowledges that AI now generally involves probabilistic analysis (combining probability and logic to assign a value to uncertainty).

Other organizations and individuals may use different definitions. There


# The Problem



**Exponential growth** in modern AI/ML, Life Sciences, Financial Services and cloud applications is fueling unheard-of data growth



Organizations desire competitive advantage through digitalization by **10x faster processing** of 10x bigger datasets

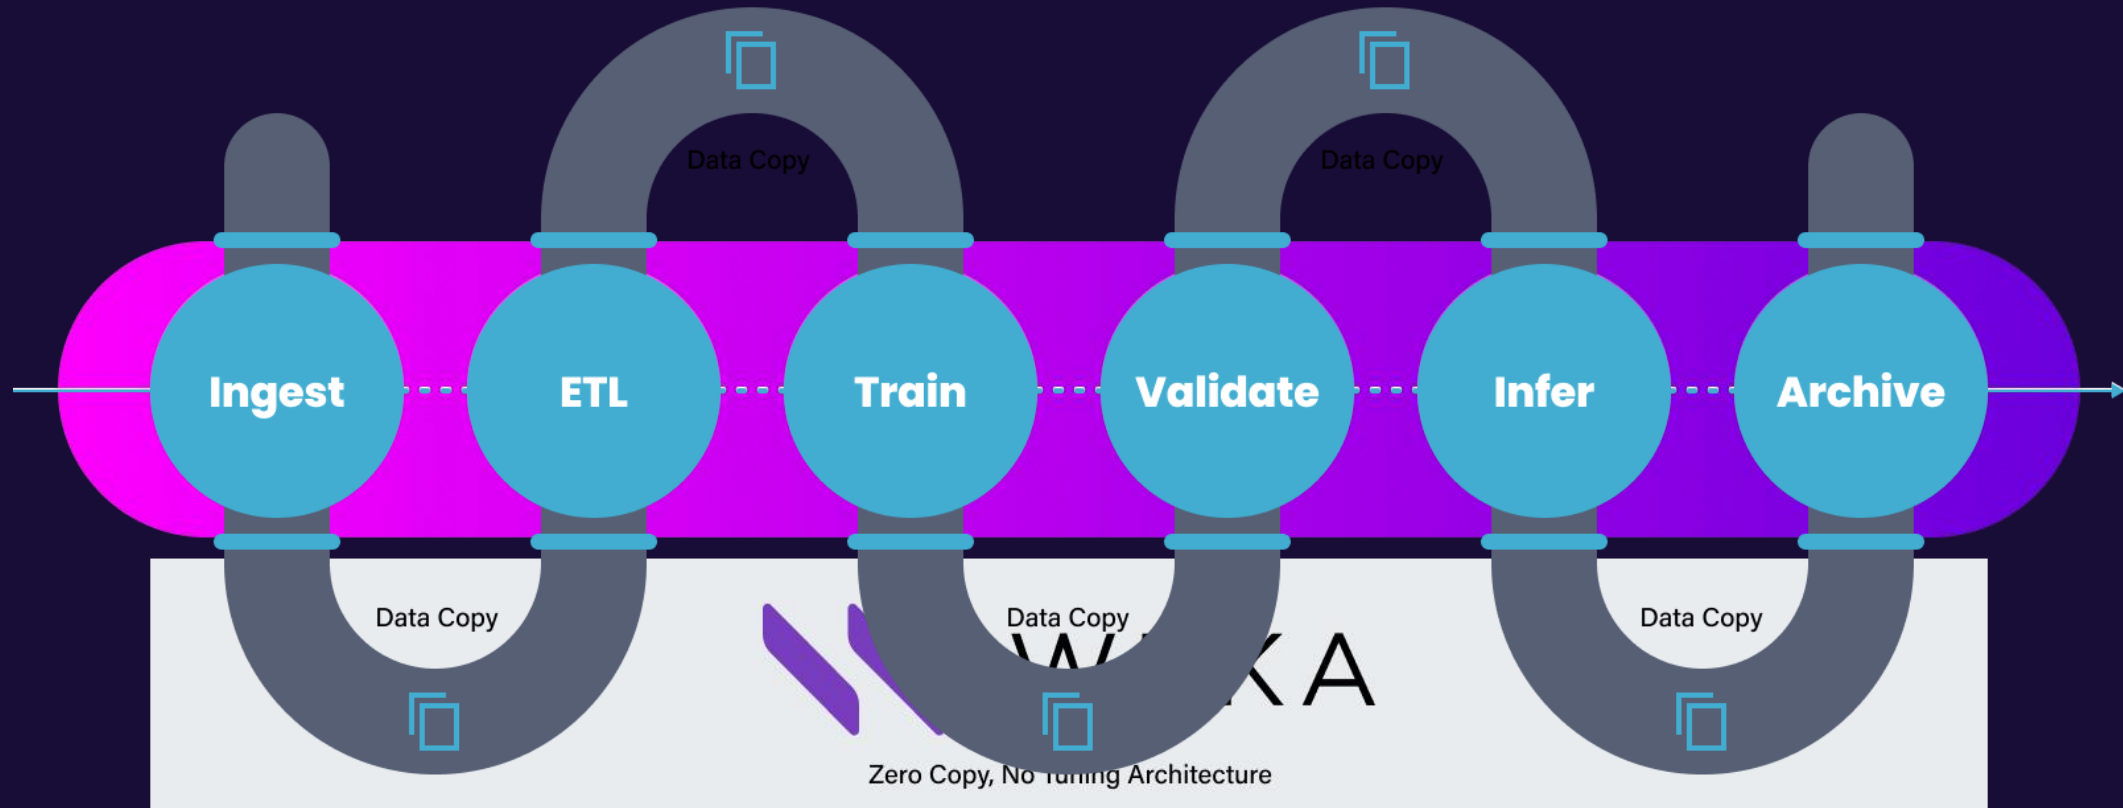


**Traditional hardware** and cloud storage cannot meet rapidly escalating requirements for performance AND scale across core, multi-cloud and edge data centers

**...A NEW APPROACH IS NEEDED.**

# Infrastructure is Slowing Down AI Data Pipelines

**70% of an epoch** happens before the data ever reached the GPU\*





# WELCOME TO WEKA



## FOUNDED IN 2014

ENGINEERING HQ IN TEL AVIV  
CORPORATE HQ IN CALIFORNIA

## GOALS

ADDRESS SCALABILITY &  
PERFORMANCE CHALLENGES  
OF **UNSTRUCTURED DATA**

## FINANCE

MORE REVENUE IN Q3 FY22  
THAN ALL OF FY2021. (**635%  
GROWTH**)  
SERIES D RAISED \$135M IN  
2022

## CUSTOMERS

**8 OF THE FORTUNE 50**  
FINANCE, HEALTHCARE,  
MANUFACTURING, OIL & GAS,  
GOVERNMENT, DEFENSE,  
ACADEMIA/RESEARCH

## PARTNERS

AMD, AWS, HITACHI  
VANTARA, HPE, LENOVO,  
NVIDIA, SUPERMICRO, DELL  
TECHNOLOGIES,  
QUANTUM, SCALITY, IBM

# 외부 조사기관 평가(Gartner Magic Quadrant)

Figure 1: Magic Quadrant for Distributed File Systems and Object Storage

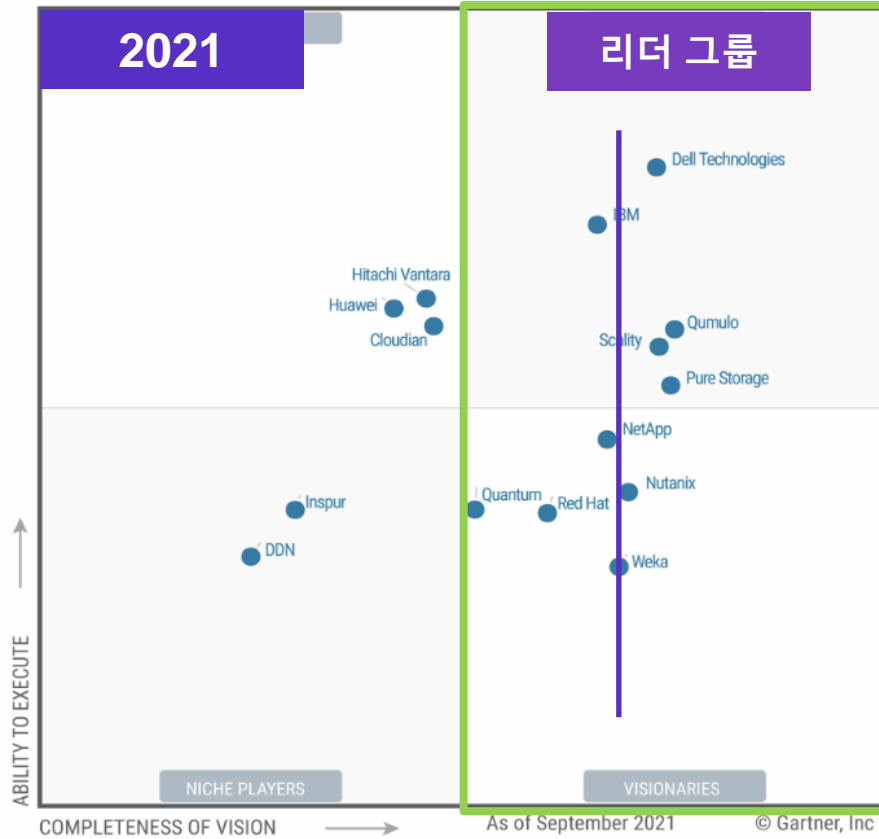
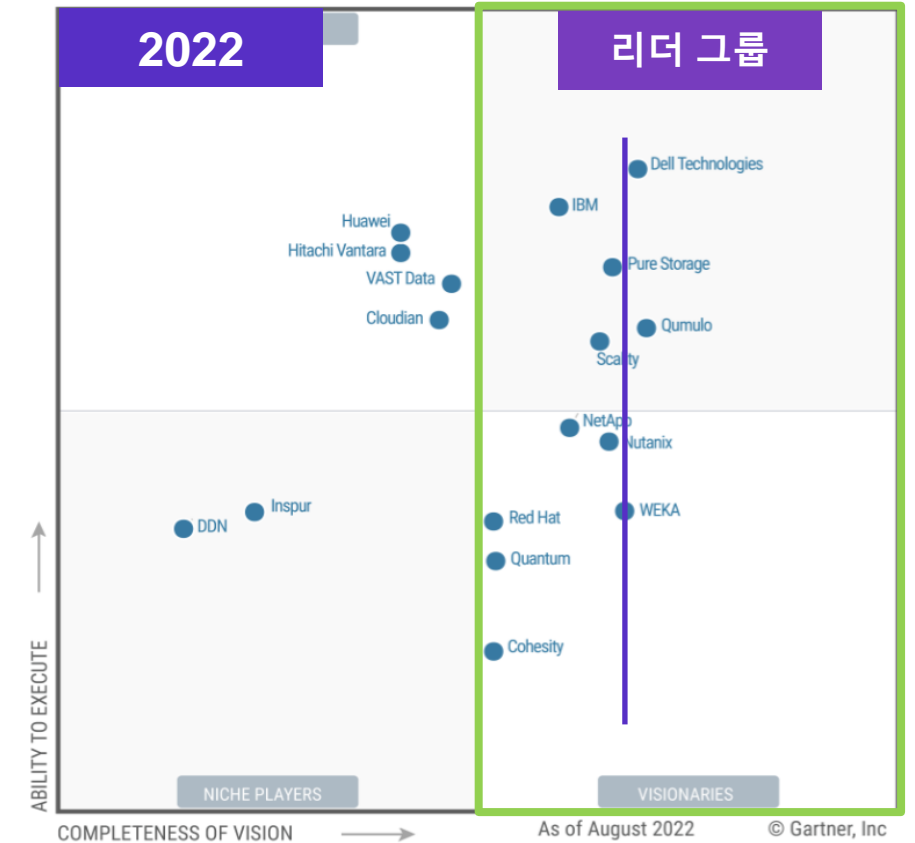


Figure 1: Magic Quadrant for Distributed File Systems and Object Storage



고성능 병렬 파일시스템 벤더 중 유일한 리더 그룹, 리더 그룹에서도 유일한 포지션이 약진한 고속 성장 벤더

# WEKA's Play ground



**AI/ML & DEEP LEARNING**



**LIFE SCIENCES**



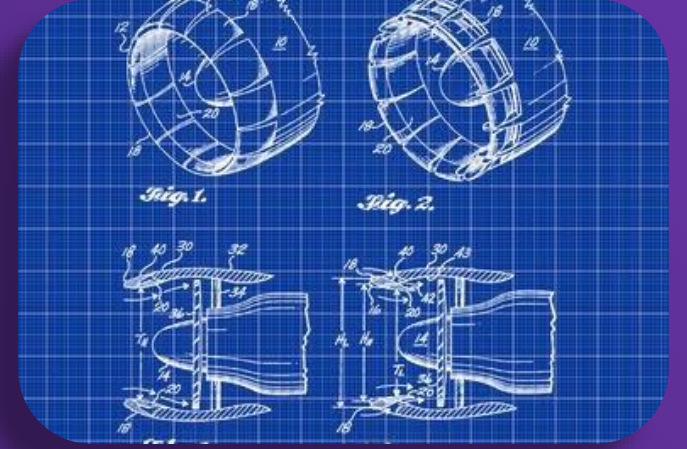
**TRADITIONAL HPC**



**MEDIA & ENTERTAINMENT**



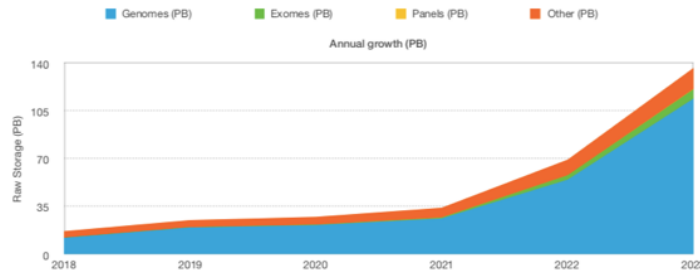
**FINANCIAL SERVICES**



**MANUFACTURING**

# ...And Their Common Pain Points


## IT Budgets Are Being Squeezed



DO MORE WITH LESS!!!



Fewer Personnel



Budgets Squeezed

## Data Is Growing Exponentially



Real-time Analytics



Financial Processing & Analysis



Fraud Detection



Geospatial Research



HPC



Microscopy




Life Sciences



Imaging Device

## Compute Densification



2008

First Top500 Supercomputer  
10,400 Rack Units, 1 petaflop

5X Performance

0.1% of the Space

I/O per Client is off the charts


2020



NVIDIA DGX-A100  
6 Rack Units, 5 Petaflop

# Uniquely Designed for Next-Gen Workloads

Software Data Platforms *Replace* Traditional *Hardware* Storage Systems



**Zero Tuning** for  
**World-Record**  
**Performance**



**Exabyte++**  
**Scalability**



**Cloud, Edge &**  
**On-Prem**



**Modern Data**  
**Management**



**Class Leading**  
**Economics**

# One Architecture Delivers on Three Promises

## File semantics

Simple to install  
Simple to manage  
Cloud integration

**Simplicity**

**Speed**

## Lowest latency

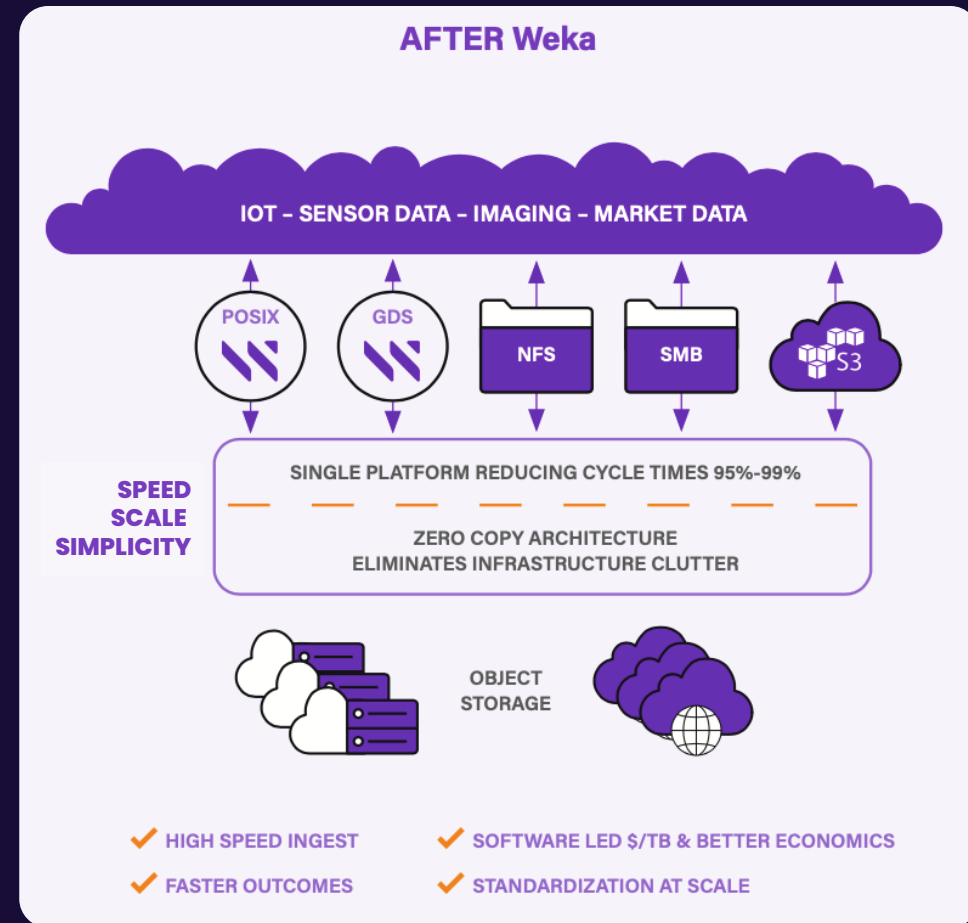
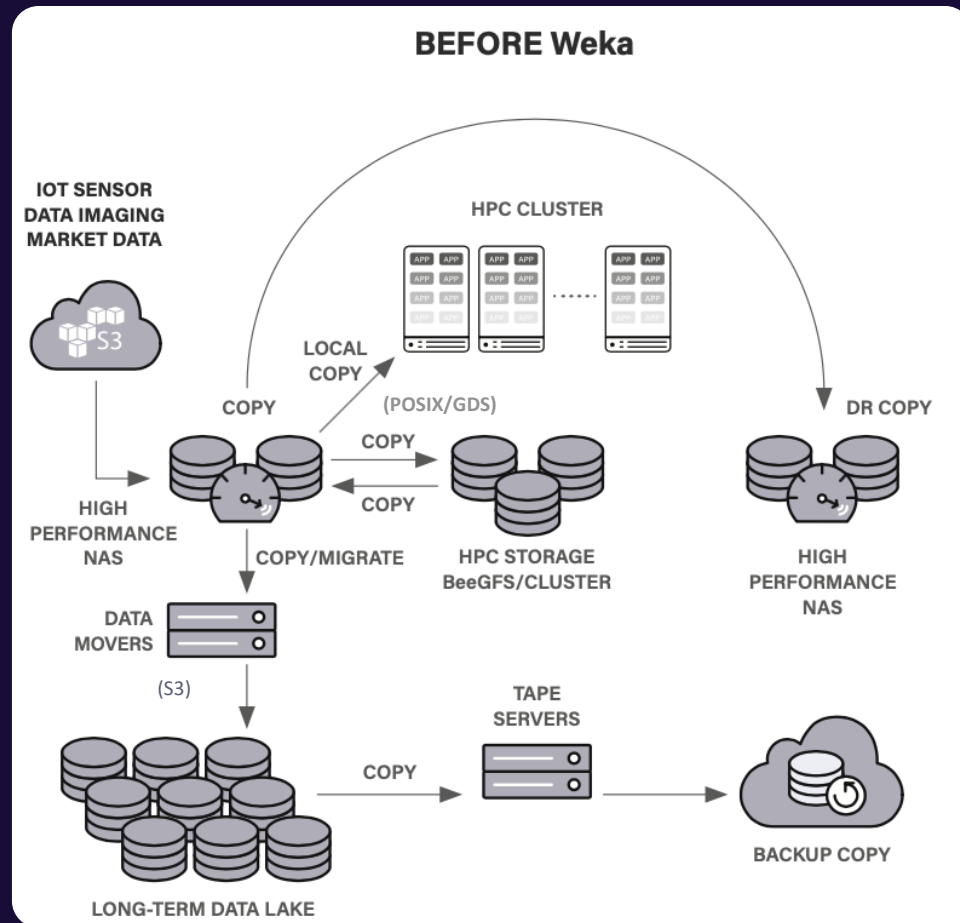
Faster than a local filesystem  
or All Flash array

**Scale**

## Massive scale

Capacity  
Throughput  
Metadata

# Eliminate A Mish-Mash Of Storage Systems



# 50%

Lower Cloud Costs



# 30min

Installation

# 4.2

# 1000s

of PVCs

# 8x

Data Reduction



# Business Benefits

**Increase  
productivity**

up to

**90%**

Improvement in  
deep learning  
epoch time

**Improve  
resource ROI**

**75%**

Reduction in  
workload wall clock  
time

**Lower  
storage costs**

**65%**

Savings over  
traditional high-  
performance  
computing

# One Data Platform That Does It All

✓ **1 Tier of Storage**

✓ **1 File System**



✓ **1 Point of Management**

✓ **1 Copy of Data**

**2 times performance guarantee program !!!**

# WEKA on Cloud



# Atomwise Reduced their Epoch Time by 92%



Feed billions of tiny files into expensive GPU instances



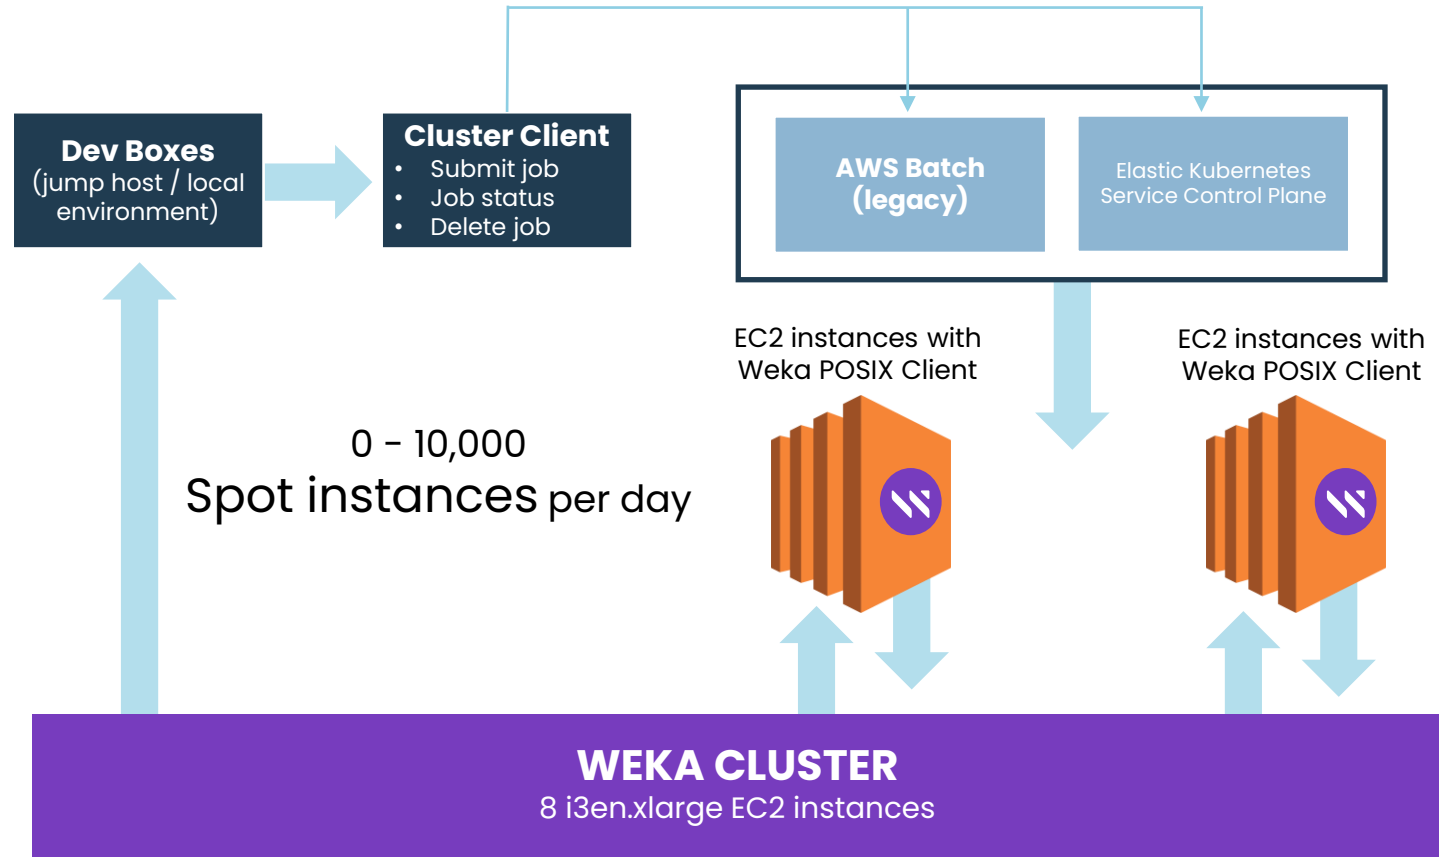
Improve AI Epoch times



Eliminate copies across multiple EBS instances



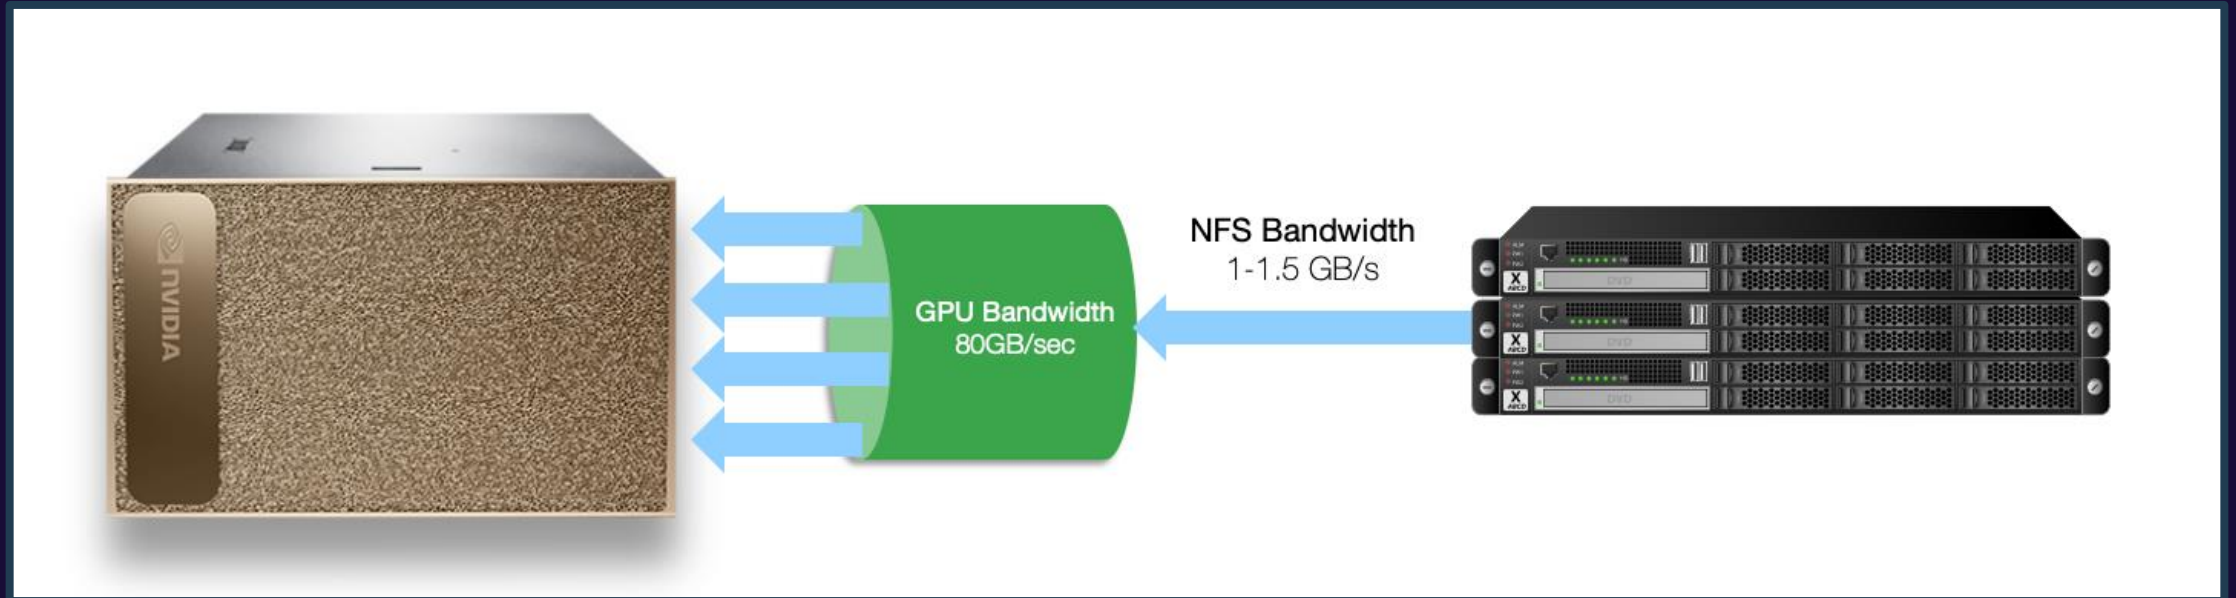
Integrate with **Kubernetes** pipelines



Migrated from native services (EBS) to reduce Epoch time from 3 months to 1 week with 100%+ savings

# Let's Talk About Performance That Matters

## Vendors Emphasize Throughput, Which is Wrong



Workloads Need Bandwidth, IOPS, Latency, and Metadata Performance

# Benchmarks on NVIDIA DGX Systems

## WEKA Data Platform for GPUDirect Storage



**DGX-A100 Benchmarks**



**DGX-2 Benchmarks**

# Solves the AI Storage Challenges

Improves productivity and **faster time to market and value**

Accelerate large scale data pipelines

- Reduced epoch times
- Fastest inferencing
- Highest images / secs benchmarks

Run entire pipeline on the same storage backend

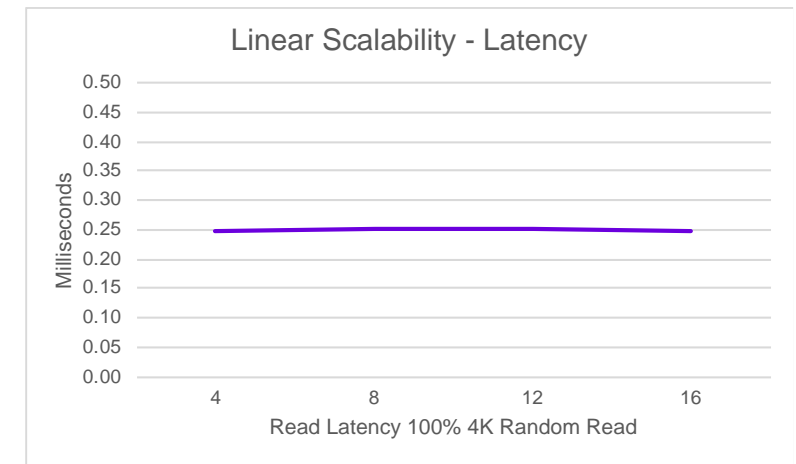
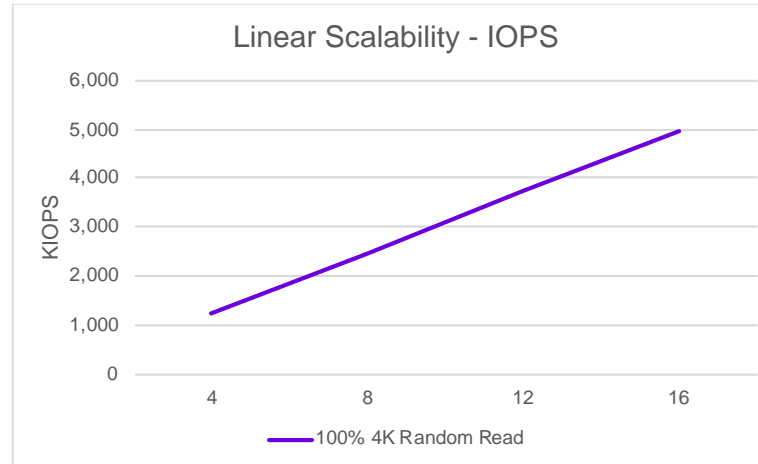
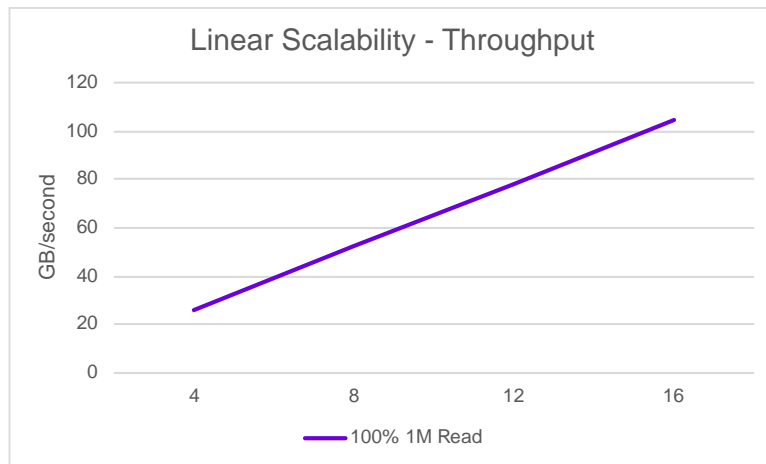
Multi Protocol Support

**Faster than local storage**



# File System Scales Linearly with Cluster Size

## WekaFS performance in AWS



- 308K OPS/AWS *i* Instance
- 6.6GB/sec/AWS *i* Instance
- <250 microsecond *i* latency



**Today, the world's data centers  
consume roughly 3%  
of the global energy supply**

**Without intervention,  
it will grow to 8% by 2025**

**That's about 40 megatons  
of CO<sub>2</sub>e worldwide...**

An aerial photograph of a lush green forest with a winding river. The text is overlaid on the image.

**Reforestation** is critical to  
**Reduce Carbon**

# Our Vision for Sustainable AI





WEKA

