



WEKA Data Platform

Jonathan Poppe

Sr. Systems Engineer

Welcome to WEKA

○ Founded in 2014

Engineering HQ in Tel Aviv
Corporate HQ in California

○ Finance

300% growth in 2022
Investments from Nvidia, Mellanox, HPE, Seagate,
Western Digital, Micron, Cisco

○ Partners

AMD, AWS, Hitachi Vantara, HPE, Lenovo, Nvidia,
Supermicro, Dell, Quantum, Scality, and IBM

○ Goals

Address scalability & performance
challenges of next-generation workloads

○ Customers

8 of the Fortune 50
(Finance, life sciences, manufacturing, retail,
oil & gas, telco)
Government, defense, academia/research






For Those Who Solve Big Problems

WEKA has built a software-defined data platform that leverages NVMe storage, fast networks and a parallel file system to unleash the value of your data

WEKA Delivers Multi-Dimensional Performance

**Data Platform designed
for NVMe, and modern
networks**

**Broad multi protocol
access to data on-
premises or in the cloud**



**Performance for
bandwidth, IOPS, low
latency and meta data**

**Faster than local-disk for
mixed workloads and
small files with no tuning**

Real World Example:

100 microsecond latency, 6.5m IOPs and 135 GiB/s

The Secret Sauce




Cloud Native, Datacenter Ready

Common software stack deployable
across Core, Cloud or Edge
Datacenters



Unlimited Linear Scale

Seamless, non-disruptive scaling
from Tb's to 10's Eb



Zero-tuning Mixed Workload Support

World's fastest filesystem, supports
data pipelines with high I/O, low
latency, small files, auto-tuning,
mixed workloads & data portability




Class Leading Economics

Smart tiering between high-
perf/high-cost Flash and S3 deliver
class leading economics

40+ granted patents, 95 pending patents

The Big Problems




Finance

Increasing profitability for financial trading

Reducing fraud

Reducing risk

Improving business performance through predictive / prescriptive analytics




Life Sciences

Accelerating drug & vaccine discovery

Accelerating precision medicine

More accurately detecting diseases



AI / ML

Winning the race to level 4 vehicle autonomy

More accurately detecting anomalies in medical images

Increasing sales & profitability with recommendation systems

Increasing manufacturing yields



HPC


Increasing national security

Discovering & maximising profitability from oil & gas reservoirs

Manufacturing safer, greener & more profitable products

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



5X Performance

0.1% of the Space

I/O per Client is off the charts

2020



NVIDIA DGX-A100
6 Rack Units, 5 petaFLOP

One Architecture Delivers on Three Promises

File semantics

Simple to install
Simple to manage
Cloud integration

Simplicity

Speed

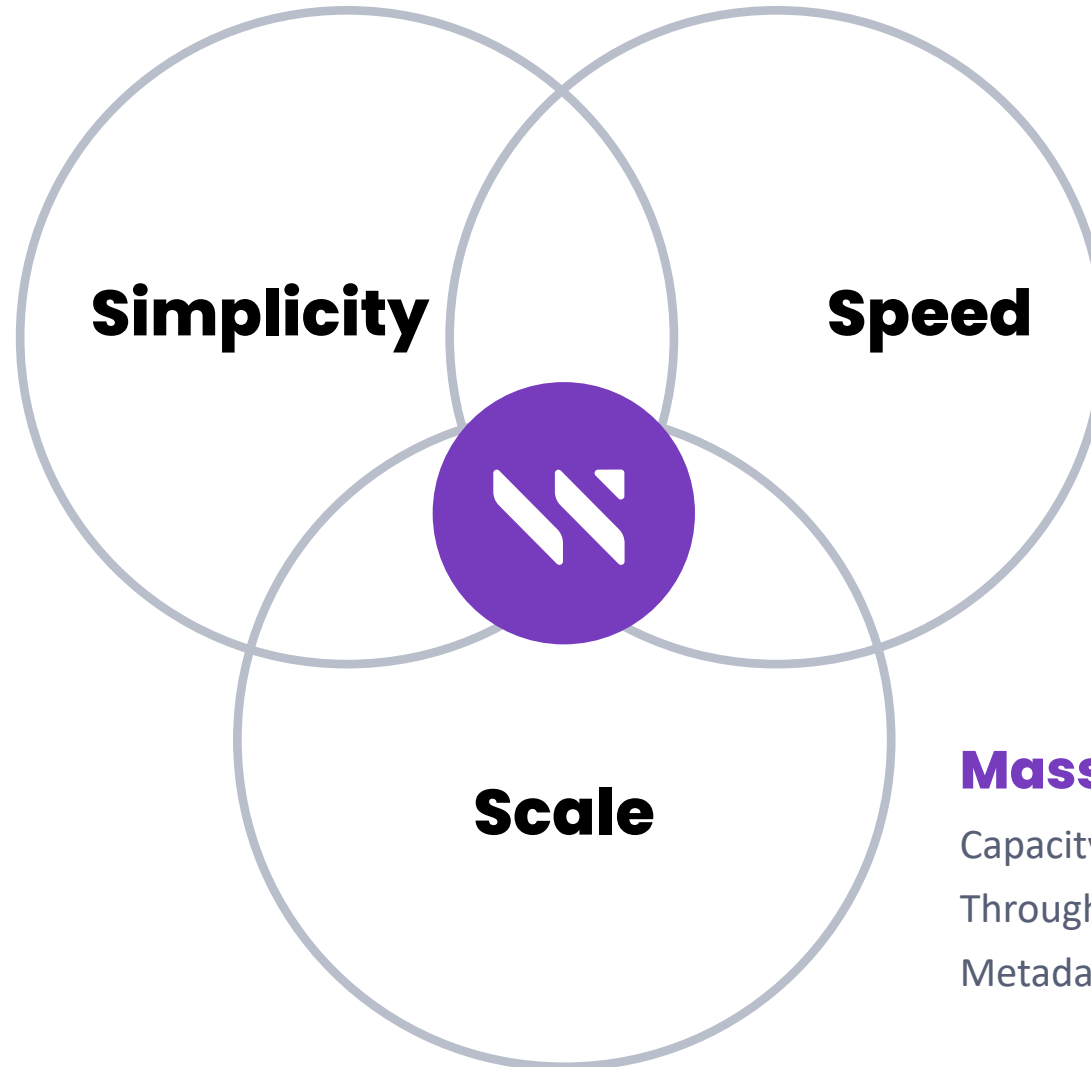
Lowest latency

Faster than a local filesystem
or All Flash array

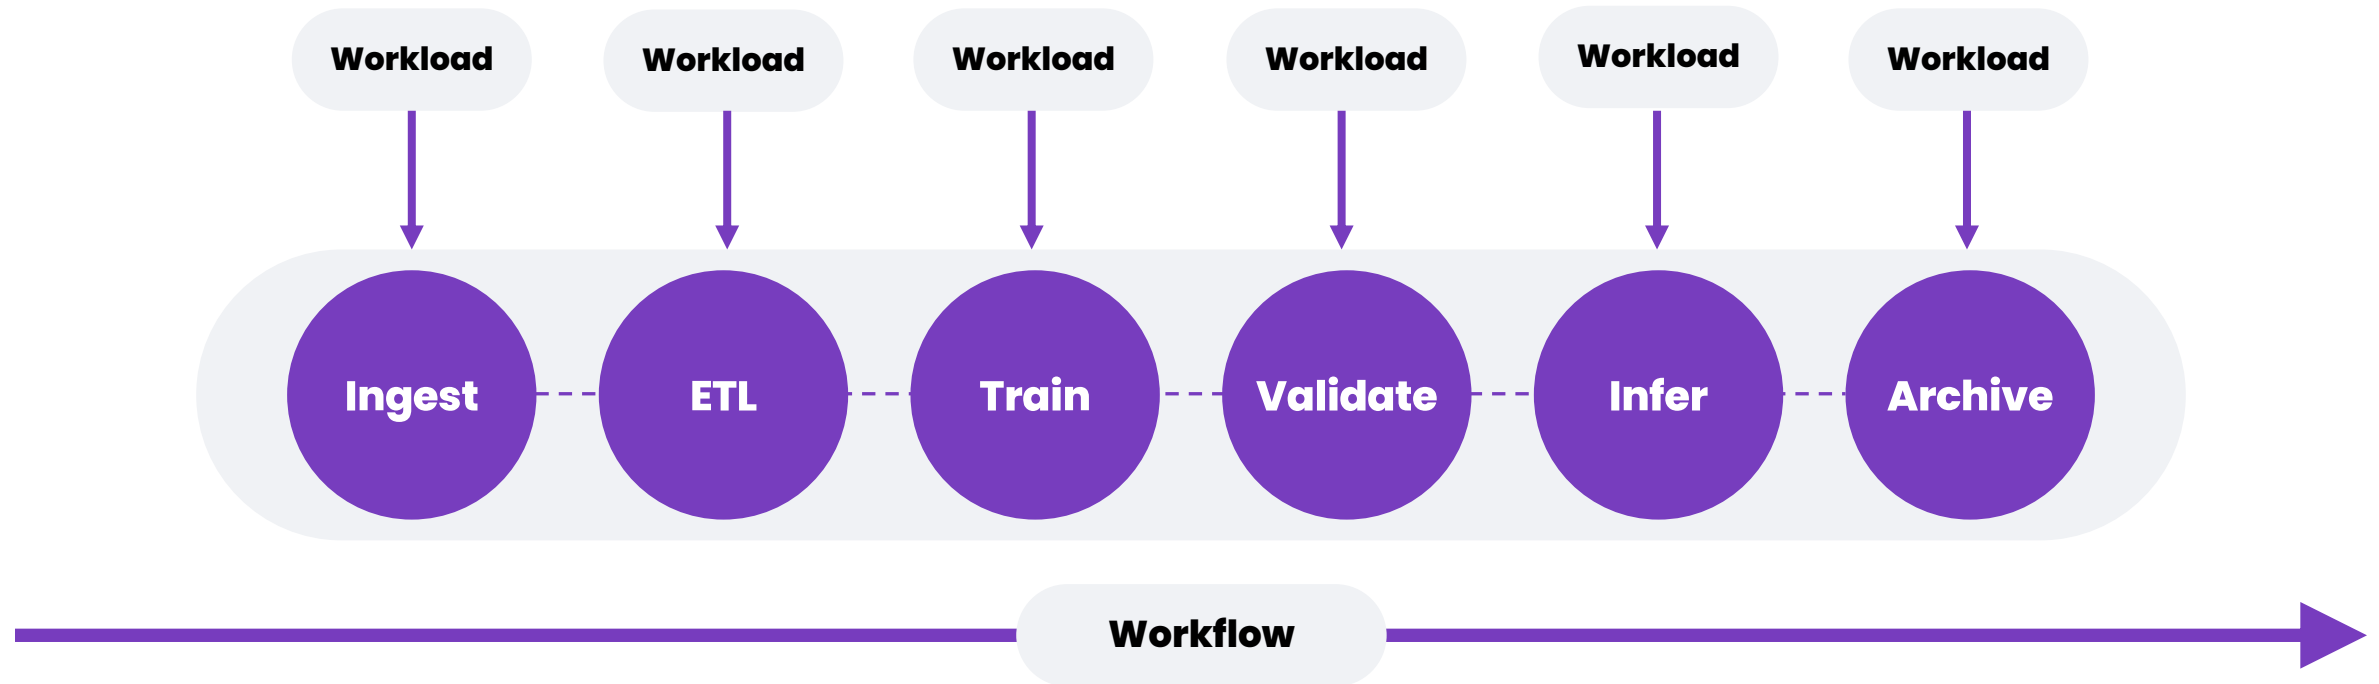
Scale

Massive scale

Capacity
Throughput
Metadata



Workload vs. Workflow





Weka FS Architecture

Shared File System Myths that Weka Breaks

Myth 1

Local File Systems are faster than shared FS

We can deliver 3x the performance of NVMe local disk

Myth 2

File Systems Do Not Scale Well

100's of Petabytes of NVME and Exabytes of OBJ storage in a namespace

Myth 3


It takes a PhD to manage a parallel file system

Simple, intuitive software with minimal setup and tuning

I/O intensive Applications



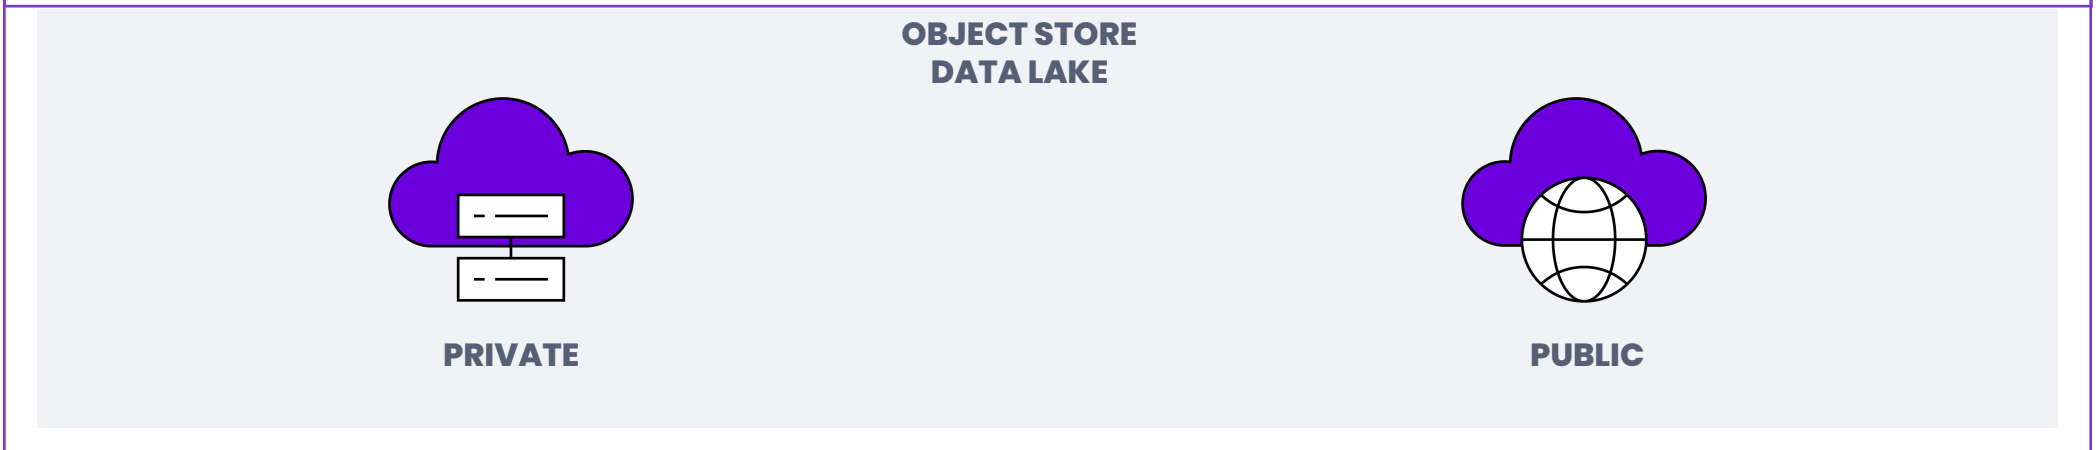
Clients With data access Via POSIX, NFS, SMB, S3, and GPU Direct



WekaFS Hosts



Automatic Tiering and Data Protection via S3 to On- Prem or Public Cloud Object Store




NFS = Not For Speed

A protocol invented in 1984 trying to solve a 2020 problem

pNFS tried to fix NFS but failed when metadata workloads exploded

Legacy parallel file systems like Lustre and GPFS cannot handle billions of small files

- And they require a PhD to operate



POSIX Filesystem

Enabling “local FS” Use Cases

The Challenge with NFS

- Protocol was defined when networking was the bottleneck (10 MbE)!
- We are now 3-4 ORDERS of magnitude faster with 10GbE– 100 GbE networking
- Not suitable for low latency operations or metadata heavy workloads
- Does not include file locking as part of the protocol


Before WekaIO, the only way to achieve good FS performance was to local SSD on a single server (not shared) or overlaid on an All-Flash SAN.

WekaIO’s networking stack and native VFS POSIX binding into the OS enable shareable filesystem with performance comparable to a local FS

- with coherency across the cluster

Networking Stack

- Low latency is critical in order to waive “locality” concerns away enabling us to use distributed coding for protection
- Architected for standard Ethernet or IB, runs on AWS
- Utilizes SR-IOV and DPDK to run stack from user-space
- Optimized for virtualized environments
- Reduces load and dependency on Linux kernel
- RDMA-like semantics from user space providing high performance
- Fully featured stack, traffic shaping supported



Fully Distributed Metadata Scaling

A must for

- Linear performance scalability
- Huge directories, with many concurrent metadata operations
- Unlimited 4k random IO that scales out at consistent latency

Metadata plane is sharded into small “buckets”

The FEs can talk directly to the correct bucket for each operation,

- All I/O is parallelized for highest performance

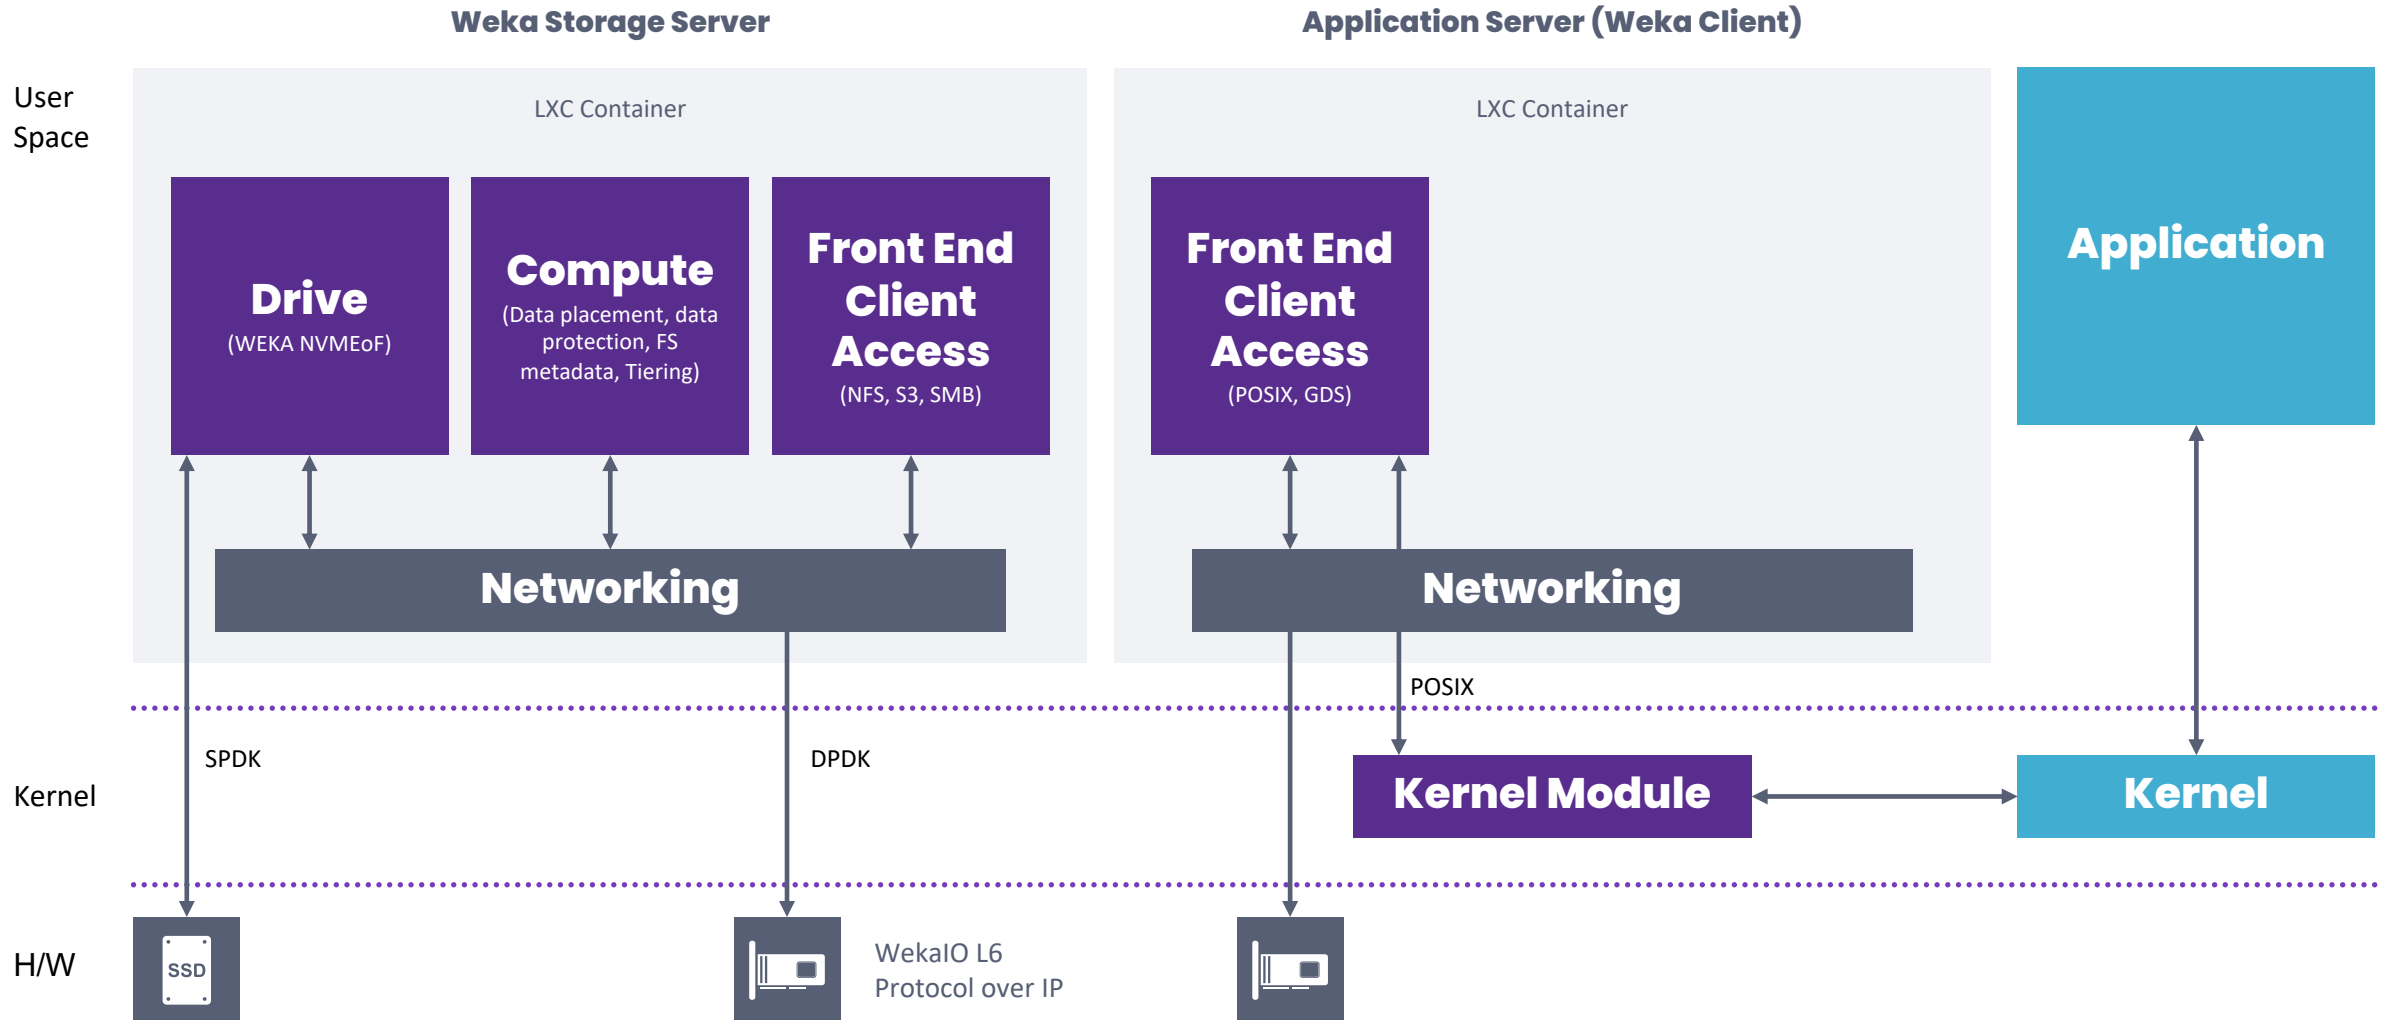
Load balancing across nodes prevents hot spots

- A must have when considering “noisy neighbor” metadata issues

Supports very large clusters – up to 64K nodes

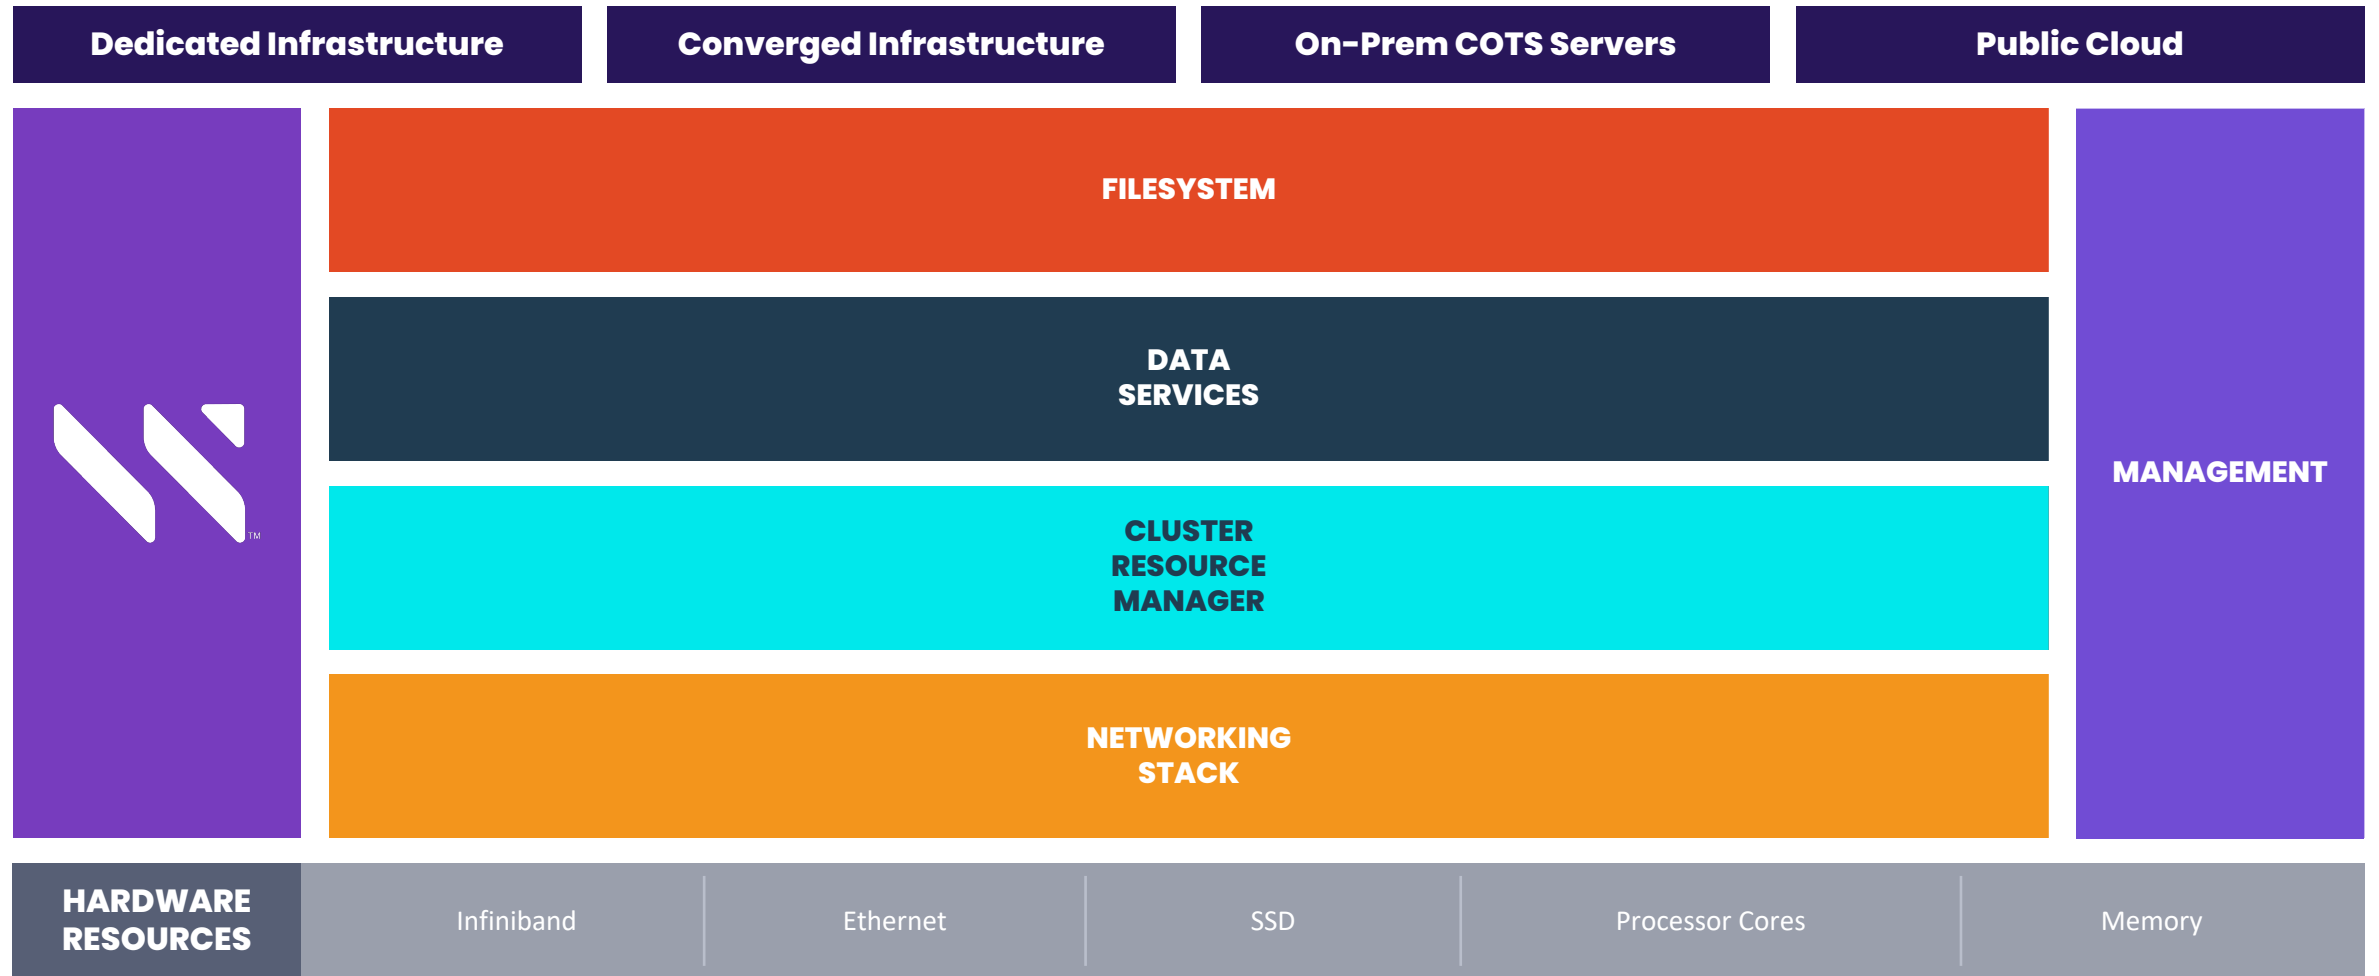


Software Architecture - Dedicated



WekaFS™

Full Featured And Flexible



WekaIO Scaling (v4.2)

Single name space

Up to 6.4 trillion(T) files/directories

Up to 14EB with up to 512PB on SSD

Up to 6.4B files in a single directory

Single files up to 4PB in size

Up to 1024 Filesystems

Up to 4096 snapshots

File System Scales Linearly with Cluster Size

WekaFS performance in AWS

6.6GB/sec

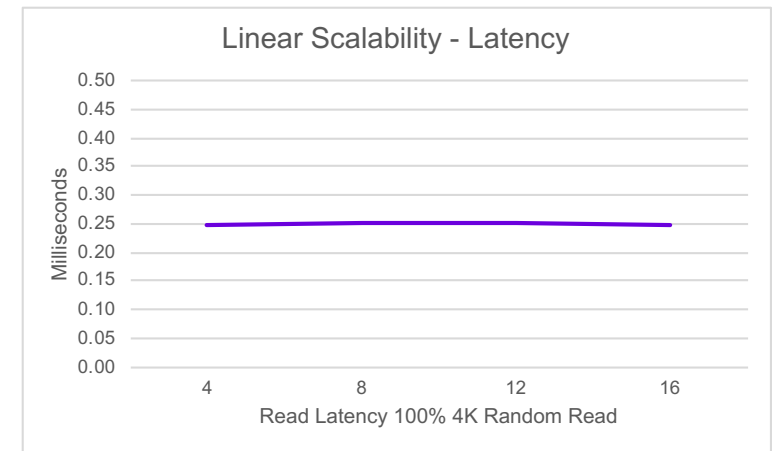


AWS i Instance

308K

OPS/AWS i Instance

<250


microsecond i latency



WEKA on OCI Delivers 2 TB per Second Performance

Maximum performance at cloud scale

- Run at petabyte scale in a high-performance file system
- NVMe SSDs for hot data and object storage for warm or cold data
- High-performance computing (HPC) bare metal Compute shape (BM.Optimized3.36)
- 100-Gbps RDMA over converged ethernet (RoCEv2) and 3.8 TB of local NVMe SSD



<https://blogs.oracle.com/cloud-infrastructure/post/weka-on-oracle-cloud-infrastructure-delivers-2-terabytes-per-second-performance>

Thank You!

 @wekaio

 /wekaio

 @wekaio