

On-Demand AI/ML and Cloud Native Workloads with Liquid and Weka Composable Disaggregated Infrastructure



Composable Disaggregated Reference Architecture for optimal utilization of data center resources with Weka Data Platform and Liquid



COMPOSABLE, ELASTIC AND HYBRID

- Disaggregated and composable compute, networking and storage
- Avoid vendor or technology lock-in
- Scale as you grow with 1024 mount points
- Run on-premises or hybrid



ACCELERATE TIME TO VALUE

- Better performance than local NVMe Flash
- 162 GB/s of throughput to single mountpoint with NVIDIA GPUDirect Storage
- Meet rigid workload requirements
- Built-in data management – protection, DR, cloud-bursting with snap2obj



INCREASE EFFICIENCY

- Mixed workload handling
- Resolve poor utilization of expensive GPU resources
- Purchase only what you need and reduce datacenter costs
- Maximize investment utilization



SIMPLIFY

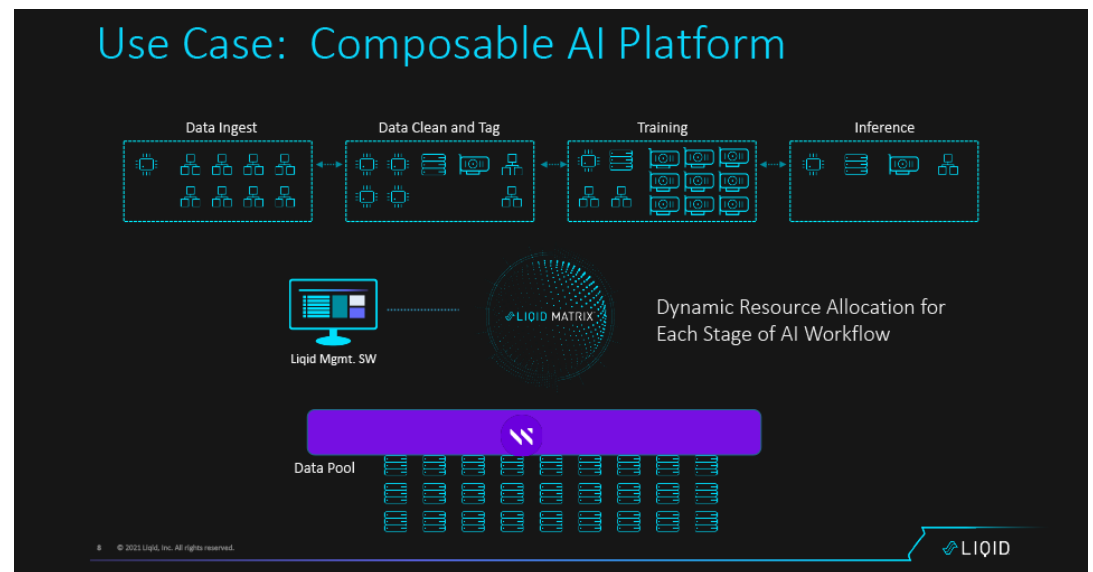
- Multi-protocol access
- Shared high performance composable storage
- End to End Security
- Unified management

Adjusting to real-time business requirements is paramount in the modern data center. Liquid empowers IT administrators to deploy, manage, and scale physical bare-metal server systems in seconds and then reallocate core data center devices on-demand as business needs evolve. Pools of compute, storage, storage class memory, and accelerator (GPU) devices are interconnected over PCIe, Ethernet and /or InfiniBand to deliver dynamically configurable bare-metal servers perfectly sized with the exact physical resources required by the application being deployed. Liquid's intelligent fabric switching technology allows core system resources to be instantly interconnected into physical servers via a PCI-Express (PCIe) fabric, and dynamically reconfigured as needed. Liquid Command Center™ is the powerful management software that enables IT to compose physical computer systems from pools of baremetal resources. Liquid enables IT administrators to deliver adaptive compute infrastructure and scale core resources on demand for the world's most data-intensive workloads.

WITH NEW DATA-INTENSIVE WORKLOADS AT MASSIVE SCALE COME NEW CHALLENGES

Weka's Limitless Data Platform helps accelerate outcomes and improve productivity in the hybrid cloud. Delivered as a broad solution for many use cases, the Platform is built on the Weka File System (WekaFS™), a software-defined architecture that delivers the industry's best performance and efficiencies by leveraging the latest technologies in storage such as NVMe, networking technologies like NVMe-oF, NVIDIA Mellanox InfiniBand, 100Gb Ethernet, and advances in computing technologies like GPU acceleration and GPU Direct storage.

WekaFS, the modern POSIX compliant parallel file system is used by eight of the Fortune 50 enterprise organizations to uniquely solve the newest, biggest problems holding back innovation and discovery. WekaFS, the world's fastest and most scalable parallel file system for data-intensive workloads, addresses the shareability, performance, and portability challenges by providing stateful and parallel storage, allowing seamless deployment on-premises and easy migration to the cloud, for cloud bursting. Using the WekaFS Kubernetes CSI plugin, organizations now have increased flexibility in how and where they deploy containers while delivering local storage performance and latency.



STATIC INFRASTRUCTURE CHALLENGES IN THE ENTERPRISE


Organizations are continually seeking new ways to streamline and transform IT infrastructure in order to gain competitive advantage, improve time to value, and overall operational effectiveness. With the first true bare-metal architecture breakthrough in 25+ years, Liquid's CDI allows organizations transform infrastructure delivery at the edge, cloud, and on-prem.

ESG recently conducted research on data center modernization investment that supports this premise. Of 664 respondents surveyed, they uncovered that 34% of organizations anticipate increasing their use of on-premises hyperscaler cloud solutions in the next 12-18 months, while 32% anticipate a significant investment in implementing a software-defined data center strategy¹.

Traditional on-premises infrastructure has remained relatively unchanged for decades. Only resource speeds and capacities have increased. Software innovations like virtualization and hyperconverged infrastructure have improved operational efficiencies and utilization. However, it still takes weeks or months to deploy a physical server. Workloads are limited to the resources that can be jammed into the box and are difficult to scale. If valuable resources like GPU, NVMe SSD, FPGA, or Storage Class Memory are sitting idle in a server, they're trapped.

Organizations are increasing deploying public cloud options, as they accelerate time to deployment and simplify scaling. However, not all applications are cloud-ready or should run in the cloud (i.e. security or latency). Also, it can be difficult to control costs in cloud environments, leading some organizations to repatriate workloads on-premises.

COMPOSABLE DISAGGREGATED SOLUTION FROM LIQID AND WEKAIO



Composable Rackscale Reference design

Liquid and Weka and unlock the flexibility and agility of public cloud for datacenters with new levels utilization and operational efficiency that can drive down costs. Composable Disaggregated Infrastructure (CDI) enables organizations to deploy bare-metal servers and all of their resources including GPU and NVMe flash storage on demand via software. With Liquid, customer disaggregated their critical resources into pools of hardware resources, including compute, GPUs, NVMe SSDs, storage-class memory (SCM), FPGAs, and NICs. Using Liquid's intuitive Command Center UI, API or CLI, organizations can quickly add the exact amount of GPU, accelerator, storage, and/or networking resources needed to their existing servers via software, eliminating the need to physically install or remove components from a server chassis.

Resources are disaggregated into PCIe expansion chassis and interconnected via a PCIe (Gen3 or Gen4) and/or Ethernet fabric switch. Host servers are seen as compute resources and are connected via PCIe HBA or smart NIC. LiquidOS resides on the fabric and intelligently connects and manages all resources without the need for drivers or agents. In some cases, resources can be added to running systems without a reboot. Should the workload require additional GPU processing or storage capacity, can be added via software in seconds. If the workload is no longer needed or its resources are idle, any or all resources can be returned to their respective free pools future use. Since these are bare-metal servers that support common operating systems, hypervisors, or container engines.

With off the shelf high-speed Ethernet switches organizations can leverage Liquid CDI to compose servers with the required amount of NVMe SSD storage capacity. By leveraging Liquid's support for multiple switch fabric types, organizations can compose a single host server with high performance storage from WekaFS interconnected with both Ethernet and IB fabrics.

“ The ability to compose a system on demand from a shared pool of compute, networking and storage resources should mean that expensive hardware such as GPUs or FPGAs can be shared across high performance composable storage and allocated as and when needed, instead of being permanently installed in every system that may need to run a workload that needs an accelerator. Weka’s software defined offering works extremely well in a disaggregate composable architecture like Liquid’s CDI.

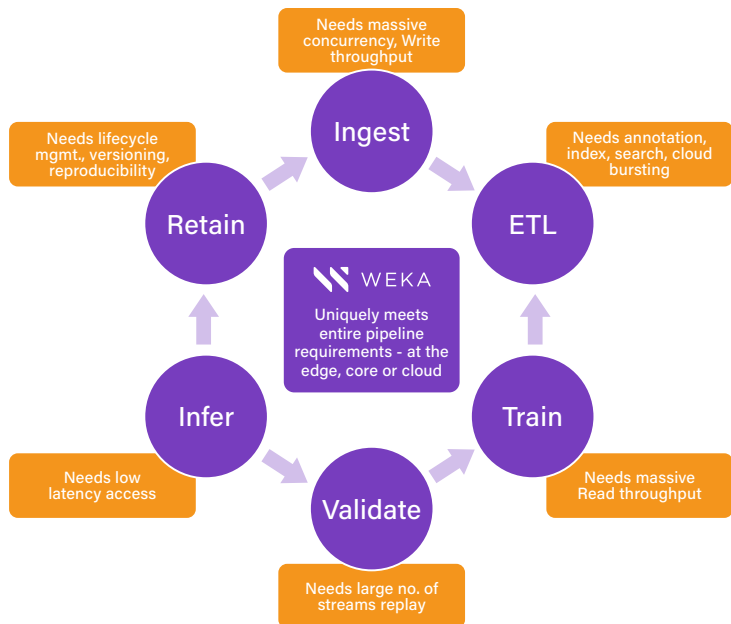
Earl Dodd, HPC & HPDA Architect, Alchemist & Business Practice Leader, WWT

Weka SOLVES STORAGE PERSISTENCE CHALLENGES FOR COMPOSABLE INFRASTRUCTURE

WekaFS provides multi-protocol access to bare-metal, containerized and virtualized applications with POSIX, GPUDirect storage, NFS, SMB and S3 protocols. WekaFS also provides a unique way of data protection, mobility and disaster recovery in a bare-metal and Kubernetes environment, with its snap2object feature, where a persistence volume claim can match to a virtual filesystem within WekaFS global namespace. WekaFS supports 1024 of these virtual filesystems in a given global namespace, which can extend over NVMe flash tier and private or public S3 bucket.



WekaFS works with optional object storage to provide capacity tier under a single global namespace. Object storage enhances the durability, geo-replication and the economics needed of a capacity tier. WekaFS also provides integration with leading KMS (Key Management Systems) for end to end and at-rest security. AI/ML and new cloud native applications are inherently different, resulting in different stages within AI data pipelines that have distinct data (storage and I/O) requirements for massive ingest bandwidth, mixed read/write handling, and ultra-low latency, often resulting in a storage silos and challenges for each stage. This means that business and IT leaders must reconsider how they architect their storage stacks and make purchasing decisions for enabling these data pipelines.



WekaFS, addresses these challenges with its mixed workload handling and enables shareability, acceleration, and portability by providing stateful, parallel storage, while meeting the requirements of high performance and low latency. WekaFS also allows seamless deployment on premises and easy migration to the cloud as needed for cloud bursting.

Using the WekaFS Kubernetes CSI plugin organizations now have increased flexibility in how and where they deploy containers, all while delivering local storage performance and low latency. In fact, throughout the system WekaFS delivers the speed-to-market required of an AI-first solution. The WekaFS CSI plugin is deployed using a Helm Chart, along with the POSIX agent on Kubernetes worker nodes. WekaFS supports volume provisioning in both the dynamic (persistent volume claim) and static (persistent volume) forms with its own storage class.

[Find out more about Weka's offering, and explore the possibilities of accelerating your performance hungry financial services workloads.](#)

WekaFS™ Data Sheet: <https://bit.ly/36BRJdc>

Weka AI Data Sheet https://www.weka.io/wp-content/uploads/files/2017/11/Weka_AI_datasheet_W07R4DS201807_Web.pdf

* SPEC® SFS2014 Results: <https://bit.ly/35UHLdk>

† ESG Master Survey Results, 2021 Technology Spending Intentions Survey, December 2020.

TARGET MARKETS AND USE CASES

| ENTERPRISE VERTICAL SOLUTIONS |
|---|
| Banking |
| Government |
| Healthcare and Life Sciences |
| Manufacturing and Energy |
| Telco |
| Retail |
| ENTERPRISE USE CASES |
| HPC and AI/ML MLOps Data Pipelines |
| <ul style="list-style-type: none"> Autonomous vehicles, NLP/NLU, Lifesciences, Simulations |
| 5G Edge to Core to Cloud Pipelines |
| <ul style="list-style-type: none"> Cloud native applications |
| Modernizing IT Infrastructure |
| <ul style="list-style-type: none"> GPU accelerated compute farms |
| Big Data, Databases and Messaging Applications |
| <ul style="list-style-type: none"> Hadoop, Spark, NoSQL and single instance databases |
| Enterprise Business Applications |
| <ul style="list-style-type: none"> Oracle, SAP, Postgress |



Next-generation data center architecture, must become far more dynamic to accommodate the emerging AI/ML and cloud native workloads and often overwhelming new data reality. Composable infrastructure technologies solve this problem by unlocking resources from their final mechanical limitations and bringing agility and better utilization. Weka and Liquid's Composable reference architecture brings together high performance, low latency with cloud agility and economics needed by these next generation workload.

Steve Duplessie, Founder and Senior Analyst, Enterprise Strategy Group - now a TechTarget division



910 E Hamilton Avenue, Suite 430, Campbell, CA 95008 T: 408.335.0085 E: info@weka.io www.weka.io

©2019-2021 All rights reserved. WekaIO, WekaFS, WIN, Weka Innovation Network, the Weka brand mark, the Weka logo, and Radically Simple Storage are trademarks of WekaIO, Inc. and its affiliates in the United States and/or other countries. Other trademarks are the property of their respective companies. References in this publication to WekaIO's products, programs, or services do not imply that WekaIO intends to make these available in all countries in which it operates. Product specifications provided are sample specifications and do not constitute a warranty. Information is true as of the date of publication and is subject to change. Actual specifications for unique part numbers may vary.