



Liquid Composable Artificial Intelligence

Complete AI Infrastructure Platform that increases flexibility and maximizes resource utilization.

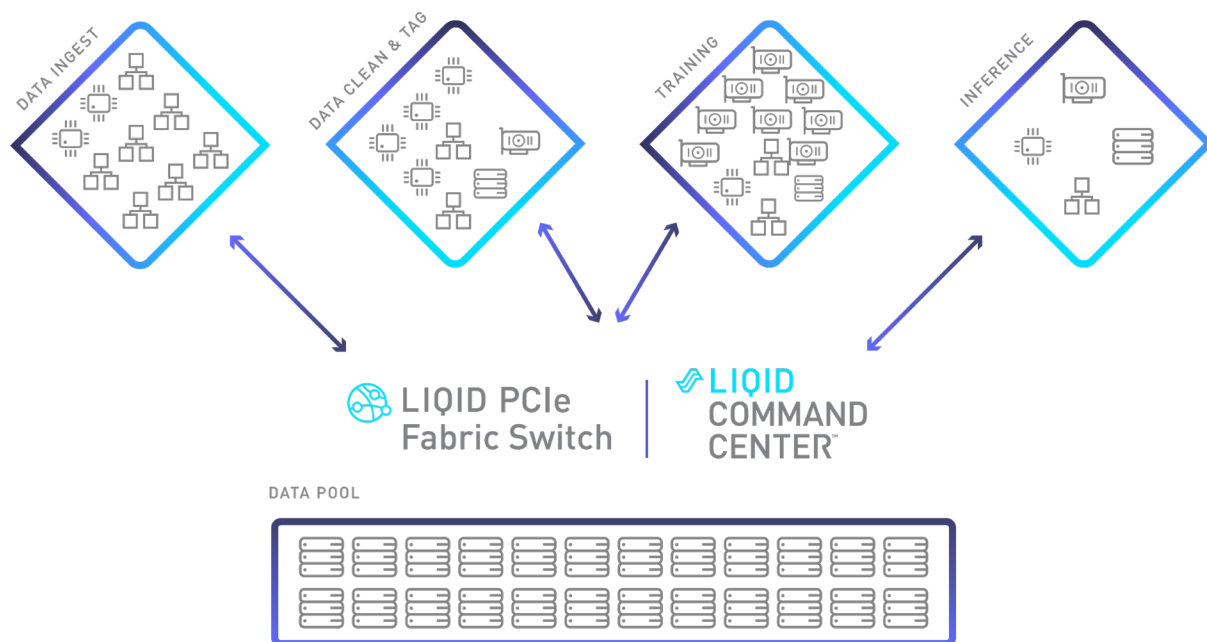
Composable AI from Liquid and Dell EMC expand the promise of software-defined composability, allowing IT users to create and maintain a more efficient data center infrastructure to meet the crushing and inconsistent data performance requirements of today's AI-driven compute environments, as well as prepare for emerging, high value applications.

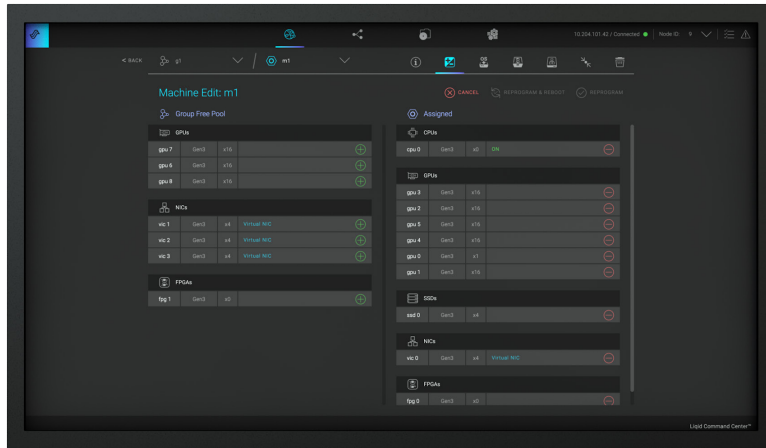
AI workloads represent a data-intensive, powerfully uneven series of compute processes. Infrastructure requirements significantly differ, for example, between the NIC-centric data ingest phase versus the GPU-oriented data training phase. Traditional static data center architectures have required that individual systems be built to handle each phase. The data is then moved from one system to the next for data ingest, cleaning/tagging, training, and inference.

This approach leaves resources sitting idle while simultaneously increasing the need for physical real estate, heating, cooling and maintenance. Rather than move heavy data from system to system creating inefficiency and hardware sprawl, Liquid has devised an approach to infrastructure capable of adapting the same resources to each phase in the AI workflow as is required.

Liquid delivers software-defined, composable resource allocation across bare metal via Liquid's ultra-low-latency PCIe fabric to ensure significant increases in data agility, capacity and bandwidth. Leveraging pools of disaggregated GPUs, CPUs and NVMe storage, IT users can compose balanced systems for each AI phase of data ingest, cleaning/tagging, training, and inference, while minimizing the data center footprint.


Dynamic Resource Allocation for Each Stage of A.I. Workflow





Liquid Command Center (Software)

Extensible composable infrastructure management software that automates, orchestrates, and dynamically composes bare-metal machines from pools of disaggregated bare-metal elements.



Liquid Fabric

The ultra-low latency, intelligent, and managed PCIe Gen3 switching fabric that electrically interconnects pools of disaggregated system elements.

Compute Elements (Nodes)

Dual processor Intel Xeon Scalable compute nodes with up to 24x DDR4 DIMMs per node. Each compute node will leverage a PCIe x16 Adapter that enables connectivity to the Liquid Fabric.

Expansion Chassis

PCIe-connected expansion chassis that hold up to 8x GPU devices and 8x NVMe SSDs. Flexible topology design supporting multiple GPU types and Liquid Flash.

Contents of the Solutions Bundle include:

Device	Quantity	Description
Software	1pcs	Liquid Command Center Management SW
Fabric	1pcs	Liquid Grid - PCIe Gen3 Fabric Switch
Expansion	2pcs	Liquid PCIe Expansion Chassis
Compute	4pcs	Dual Intel® Xeon® Scalable Processors
GPU	8pcs	Nvidia V100 GPU 32GB DDR
SSD	8pcs	NVMe SSD - 4TB per Drive (32TB Total)

For quotes and/or inquiries, please contact:

Liquid, Inc.
329 Interlocken Pkwy., Ste 200
Broomfield, CO 80021
office: +1 303.500.1551
email: sales@liquid.com