

# Liquid ThinkTank LQD8360 Composable GPU System

Increase AI infrastructure agility and maximize resource utilization.

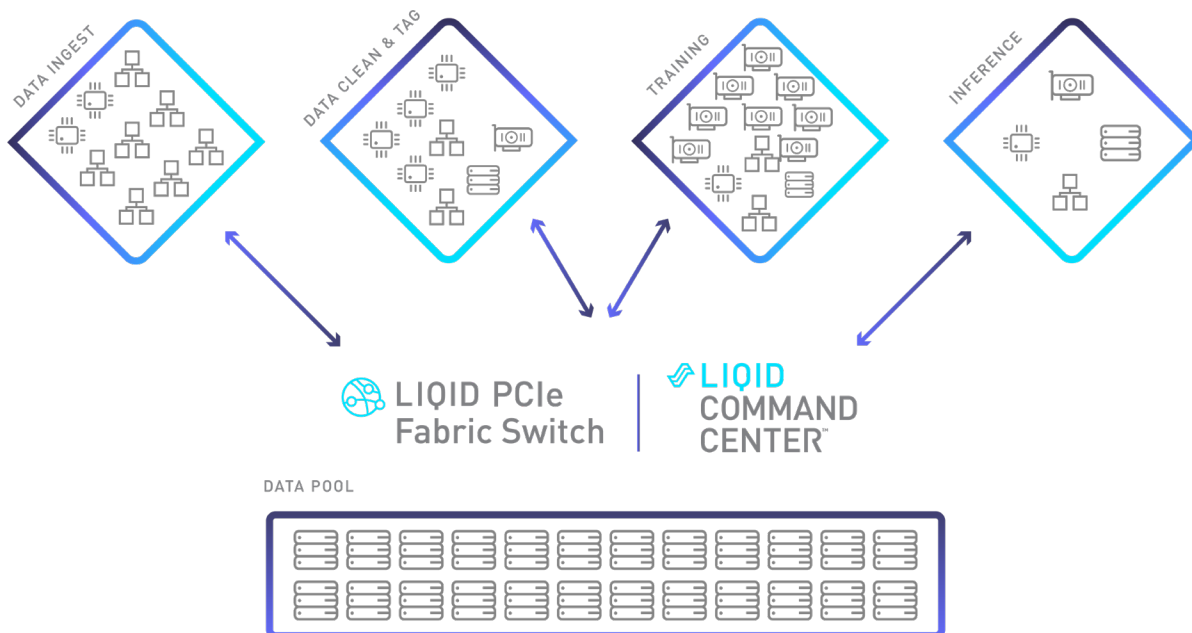
Composable AI Infrastructure from Liquid and Dell EMC expand the promise of software-defined composability, allowing IT users to create and maintain a more efficient data center infrastructure to meet the crushing and inconsistent data performance requirements of today's AI-driven compute environments,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

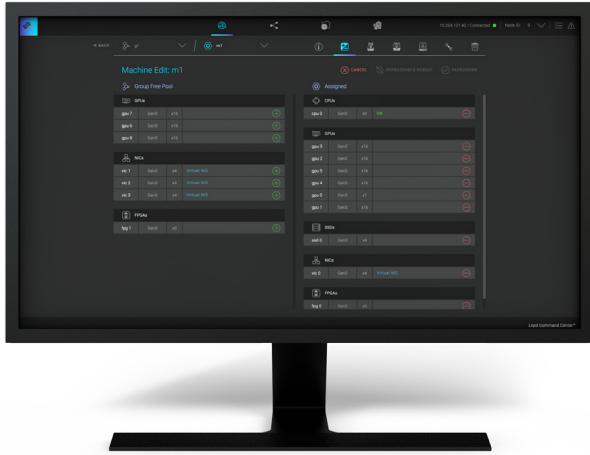
AI workloads represent a data-intensive, powerfully uneven series of compute processes. Infrastructure requirements significantly differ, for example, between the NIC-centric data ingest phase versus the GPU-oriented data training phase. Traditional static data center architectures have required that 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.

#### Solution Bundle:

CPU		GPU		SSD		NIC	
	1x, 2x or 4x Nodes		16x NVIDIA RTX8000		60TB NVMe Storage		4x 100GbE Networking


## Liquid ThinkTank LQD8360 Composable GPU System




#### Configuration Overview

<b>Solution</b>	Liquid ThinkTank LQD8360
<b>Compute</b>	1x, 2x or 4x Compute Nodes (1TB DRAM)
<b>GPU</b>	Up to 20x NVIDIA Quadro RTX 8000 (768GB RAM)
<b>Storage</b>	60TB NVMe SSD
<b>Networking</b>	4x 100Gb/s (IB or Eth)
<b>Video</b>	Yes – Video Out Available
<b>Architecture</b>	Disaggregated Composable Infrastructure
<b>Multi-Node</b>	Yes

### Multi-Node Configurations Optional:



Up to 20xGPU – 1xCPU  
L8360-001GSP-B30 (single node)



Up to 20GPU – 2xCPU  
L8360-002GSP-B30 (dual node)



Up to 20GPU – 4xCPU  
L8360-004GSP-B30 (quad node)

For quotes and/or inquiries, please contact:

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