



## EX-5410C SmartStack Datasheet

### Why LIQID

Liquid unlocks new levels of memory performance, agility, and utilization, enabling accelerated time-to-value for modern workloads like AI inferencing, RAG pipelines, VectorDBs, GraphDBs, and HPC.

### Key Advantages

- » Accelerate time-to-results with scalable, high-bandwidth CXL memory, up to 20TB per chassis and 100TB+ per host fabric.
- » Meet real-time workload demands with zero-touch, dynamic DRAM-to-server provisioning.
- » Drive new levels of efficiency with near-100% memory utilization and an on-demand approach to capacity delivery.

### Contact Information

LIQID Inc.  
 11400 Westmoor Circle, Suite 225  
 Westminster, CO 80021  
 office: +1 303.500.1551 email: sales@liquid.com

# LIQID EX-5410C SmartStack

## Composable CXL Memory, Purpose-Built for AI Inferencing, RAG, Graph Databases, and Vector Databases

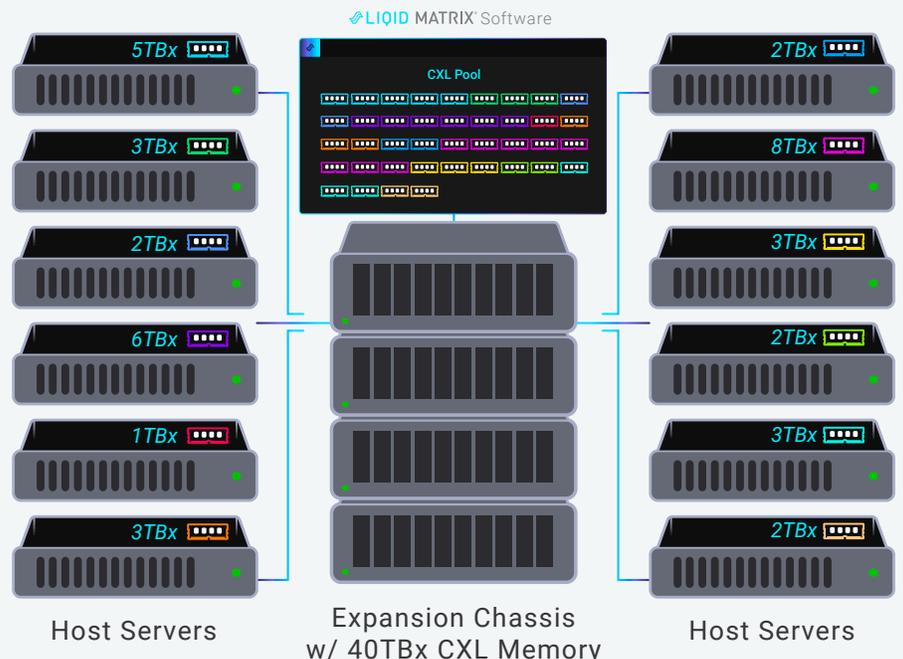
### Overview

Organizations need a flexible, high-performance way to deliver large-scale memory capacity exactly where and when it's required in order to keep up with the massive and unpredictable memory demands of modern AI inferencing, retrieval-augmented generation (RAG), GraphDBs, and VectorDBs. Fixed DRAM configurations limit embedding sizes, slow index builds, restrict model concurrency, and force unnecessary sharding, resulting in poor performance, higher latency, and wasted infrastructure spending.

With Liquid EX5410C SmartStack, composable CXL memory enables you to allocate and dynamically provision terabytes of memory. Expand, pool, and share DRAM across servers to meet the needs of your memory-bound AI workloads on demand, instantly, via software, wherever workloads run.

For AI inferencing and RAG pipelines, this means dramatically larger context windows and faster embedding lookups. For GraphDB and VectorDB workloads, CXL memory eliminates sharding, reduces replication overhead, and provides ultra-low-latency access for traversal-heavy and similarity-search operations.

### Dynamic CXL Memory Allocation



Liquid EX5410C SmartStack is a fully validated, production-ready solution designed for the most demanding AI and database environments while significantly reducing power and software licensing expense. Offering up to 20TB per chassis and scalable to 100TB+ in a single fabric, it provides everything needed to unlock breakthrough performance for inferencing clusters, real-time RAG systems, and memory-driven GraphDB and VectorDB architectures.

## LIQID SmartStack 5410C Series Technical Specifications



	SmartStack 10TB (Small)	SmartStack 20TB (Medium)	SmartStack 40TB (Large)
<b>Description</b>	10x Modules / 5 Host Capacity	20x Modules / 8 Host Capacity	40x Modules / 12 Host Capacity
<b>Composable Software</b>	Liquid Matrix™ Software		
<b>Key Features</b>	Bare Metal DRAM Provisioning, Granular Scale, Intuitive UI, API, CLI Management, Cloud Automation Integration		
<b>Management Appliance</b>	1x Liquid Director 1U		
<b>Memory Modules</b>	10 modules - 8x 128GB DDR5 DIMM	20 modules - 8x 128GB DDR5 DIMM	40 modules - 8x 128GB DDR5 DIMM
<b>Maximum Host</b>	5x Host Servers	8x Host Servers	12x Host Servers
<b>CXL Expansion Chassis</b>	1x Liquid EX-5410C CXL Chassis	2x Liquid EX-5410C CXL Chassis	4x Liquid EX-5410C CXL Chassis
<b>CXL Fabric Switch</b>	None	1x CXL ToR Switch 15x Uplinks @ x16 Gen5	2x CXL ToR Switch 15x Uplinks @ x16 Gen5
<b>Rack Units</b>	5U	10U	19U
<b>CXL Host Bus Adapter</b>	Liquid CXL 2.0 HBAs (Purchase one per host server)		
<b>Warranty &amp; Support</b>	36 or 60 - Month Software and Hardware		
<b>Add-ons</b>			
<b>IO Accelerator</b>	Liquid AX-5500 IO Accelerator - Up to 64TB - NVMe PCIe Gen 5.0 x16 FHFL Add-in Card		