



SmartStack Datasheet

Why Liqud

Liquid unlocks new levels of GPU speed and efficiency, enabling accelerated time-to-value for modern workloads like AI and HPC.

Key Advantages

- » Accelerate time to results with up to 20 GPUs per host, and GPU-to-GPU and GPU-to-Storage direct communication.
- » Meet real-time workload demands with zero-touch, dynamic GPU-to-Server configurability.
- » Drive new levels of efficiency with superior resource utilization and an on-demand approach to infrastructure.

Key Features:

- » Dynamic GPU Provisioning
- » Bare Metal Connectivity
- » Granular Flexible Scalability
- » Simple UI, API, CLI Management
- » Cloud Orchestration Integration
- » Automated Deployments
- » Massive Scale-out/Scale-up
- » Peer-2-Peer Performance
- » GPU Hot-Plug / Hot-Unplug
- » Multi-Vendor GPU Support
- » All-in-One Simplicity

Liquid SmartStack

Deliver Unmatched GPU Performance and Flexibility Anytime, Anywhere

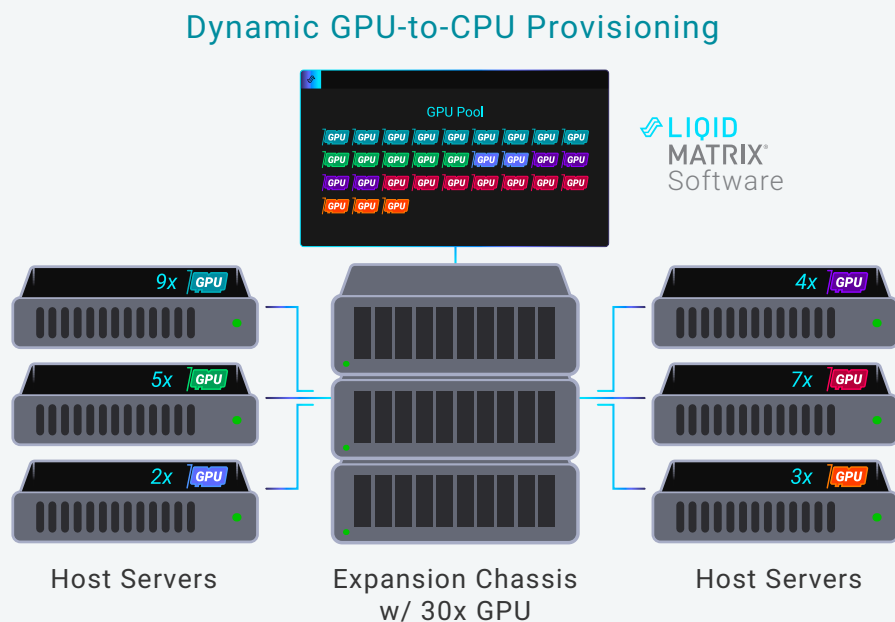
Overview

Traditional servers struggle to support the surging GPU demands of today's most challenging AI and HPC workloads. Limited GPU density reduces performance potential, resulting in poor utilization and wastes money.

Organizations need a flexible and cost-effective way to deploy GPUs. Liquid GPU On-Demand allows you to provision, scale and migrate physical GPUs via software for new levels of GPU density and flexibility. Instead of being limited to a fixed number of GPUs per host, you can now dynamically deploy as many GPUs as needed – anytime, anywhere. When GPUs are no longer being utilized, move them to a new host, zero-touch.

With Liquid GPU On-Demand, your GPUs are installed into a PCIe expansion chassis and connected to hosts over PCIe Gen4 fabric. Physical GPUs can then be quickly provisioned from GPU pools to bare metal host via Liquid Matrix software, delivering faster time-to-results, plus greater agility and efficiency.

Liquid SmartStacks are fully validated GPU On-Demand solutions designed to meet your most challenging GPU requirements. Available in four sizes, with a maximum capacity of thirty GPUs and sixteen servers per system, each SmartStack includes everything you need to deploy Liquid GPU On-Demand.



Each SmartStack supports heterogeneous GPUs from major vendors. For a list of supported GPUs and hosts, please refer to [Liquid's hardware compatibility list](#). For enhanced data performance, a single Liquid LQD4500 IO Accelerator can be added to each PCIe expansion chassis.

Liquid SmartStack 4410 Series Technical Specifications



	SmartStack 10	SmartStack 20	SmartStack 30	SmartStack 30+
Description	10 GPU / 4 Host Capacity	20 GPU / 8 Host Capacity	30 GPU / 6 Host Capacity	30 GPU / 16 Host Capacity
Composable Software	Liquid Matrix™ Software			
Key Features	Bare Metal GPU Provisioning, Granular Scale, Intuitive UI, API, CLI Management, Cloud Automation Integration, Automated OS Deployment, ioDirect GPU-to-GPU and GPU-to-storage communication			
Management Appliance	1x Liquid Director 1U			
Maximum GPUs	10x Full-height, full-length (FHFL) 10.5", dual-slot	20x Full-height, full-length (FHFL) 10.5", dual-slot	30x Full-height, full-length (FHFL) 10.5", dual-slot	30x Full-height, full-length (FHFL) 10.5", dual-slot
Maximum Host	4x Host Servers	8x Host Servers	6x Host Servers	16x Host Servers
PCIe Expansion Chassis	1x Liquid EX-4410 PCIe Chassis	2x Liquid EX-4410 PCIe Chassis	3x Liquid EX-4410 PCIe Chassis	3x Liquid EX-4410 PCIe Chassis
PCIe Fabric Switch	None	1x 48 Port PCIe Gen 4.0 Switch	1x 48 Port PCIe Gen 4.0 Switch	2x 48 Port PCIe Gen 4.0 Switch
Rack Units	5U	10U	14U	15U
PCIe Host Bus Adapter	Liquid Gen 4.0 HBAs (Purchase one per host server)			
Warranty & Support	36 or 60 - Month Software and Hardware			
Add-ons				
IO Accelerator	Liquid LQD4500 IO Accelerator - 7.68TB, 15.36TB or 30.72TB - NVMe PCIe Gen 4.0 x16 FHFL Add-in Card			

Liquid SmartStack 4408 Series Technical Specifications



	SmartStack 8	SmartStack 16	SmartStack 24	SmartStack 24+
Description	8 GPU / 4 Host Capacity	16 GPU / 8 Host Capacity	24 GPU / 6 Host Capacity	24 GPU / 16 Host Capacity
Composable Software	Liquid Matrix™ Software			
Key Features	Bare Metal GPU Provisioning, Granular Scale, Intuitive UI, API, CLI Management, Cloud Automation Integration, Automated OS Deployment, ioDirect GPU-to-GPU and GPU-to-storage communication			
Management Appliance	1x Liquid Director 1U			
Maximum GPUs	8x Full-height, full-length (FHFL) 10.5", dual-slot	16x Full-height, full-length (FHFL) 10.5", dual-slot	24x Full-height, full-length (FHFL) 10.5", dual-slot	24x Full-height, full-length (FHFL) 10.5", dual-slot
Maximum Host	4x Host Servers	8x Host Servers	6x Host Servers	16x Host Servers
PCIe Expansion Chassis	1x Liquid EX-4408 PCIe Chassis	2x Liquid EX-4408 PCIe Chassis	3x Liquid EX-4408 PCIe Chassis	3x Liquid EX-4408 PCIe Chassis
PCIe Fabric Switch	1x 24 Port PCIe Gen 4.0 Switch	1x 48 Port PCIe Gen 4.0 Switch	1x 48 Port PCIe Gen 4.0 Switch	2x 48 Port PCIe Gen 4.0 Switch
Rack Units	6U	10U	14U	15U
PCIe Host Bus Adapter	Liquid Gen 4.0 HBAs (Purchase one per host server)			
Warranty & Support	36 or 60 - Month Software and Hardware			
Add-ons				
IO Accelerator	Liquid LQD4500 IO Accelerator - 7.68TB, 15.36TB or 30.72TB - NVMe PCIe Gen 4.0 x16 FHFL Add-in Card			

[Learn More at Liquid.com](https://www.liquid.com)