



EX-5410P SmartStack Datasheet

Why LIQID

Liquid unlocks new levels of GPU performance, agility, and cost and power efficiency, enabling accelerated time-to-value for modern workloads like AI, HPC, and VDI.

Key Advantages

- » Deliver unmatched agility with up to 30 GPUs per host, and maximize performance without sacrificing bandwidth leveraging GPU-to-GPU 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

Contact Information

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

LIQID EX-5410P SmartStack

Deliver Unmatched GPU Performance and Utilization

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 GPU on-demand. Composing GPUs with Liquid Matrix™ 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 Matrix your GPUs are installed into a PCIe expansion chassis and connected to hosts over PCIe Gen5 fabric. Physical GPUs can then be quickly provisioned from GPU pools to bare metal host via software, delivering faster time-to-results, plus greater agility and efficiency.

Liquid SmartStacks are fully validated composable 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 Matrix.

Dynamic GPU Allocation



Each SmartStack supports GPUs from major vendors. For a list of supported GPUs and hosts, please refer to [Liqid's hardware compatibility list](#).

LIQID SmartStack 5410P 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	Liqid 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 Liqid 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 Liqid EX-5410P PCIe Chassis	2x Liqid EX-5410P PCIe Chassis	3x Liqid EX-5410P PCIe Chassis	3x Liqid EX-5410P PCIe Chassis
PCIe Fabric Switch	None	1x Gen5 ToR Switch 18-port Gen5x16	1x Gen5 ToR Switch 18-port Gen5x16	2x Gen5 ToR Switch 18-port Gen5x16
Rack Units	5U	10U	14U	15U
PCIe Host Bus Adapter	Liqid Gen 5.0 HBAs (Purchase one per host server)			
Warranty & Support	36 or 60 - Month Software and Hardware			
Add-ons				
IO Accelerator	Liqid AX-5500 IO Accelerator - Up to 64TB - NVMe PCIe Gen 5.0 x16 FHFL Add-in Card			