



SmartStack MX Datasheet

Why Liquid

Liquid enables GPU-centric workload access for Dell PowerEdge MX compute sleds, ensuring maximum investment protection.

Key Advantages

- » Seamless GPU Expansion: Dynamically add up to 20 GPUs to a single compute sled, to enable new workload potential like AI.
- » Reallocate GPU resources on-demand between compute sleds to maximize resource utilization.
- » Advanced Peer-2-Peer Capabilities: RDMA communication between GPUs, significantly enhances throughput and latency reduction.

Key Features:

- » Dynamic GPU Provisioning
- » Bare Metal Connectivity
- » Granular Flexible Scalability
- » Multi-Vendor GPU Support
- » Simple UI, API, CLI Management
- » Automate Deployments
- » RDMA Performance Boost
- » GPU Hot-Plug / Hot-Remove
- » All-in-One Simplicity

Contact Information

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

SmartStack MX

Direct GPU Integration for Dell MX Compute Sleds

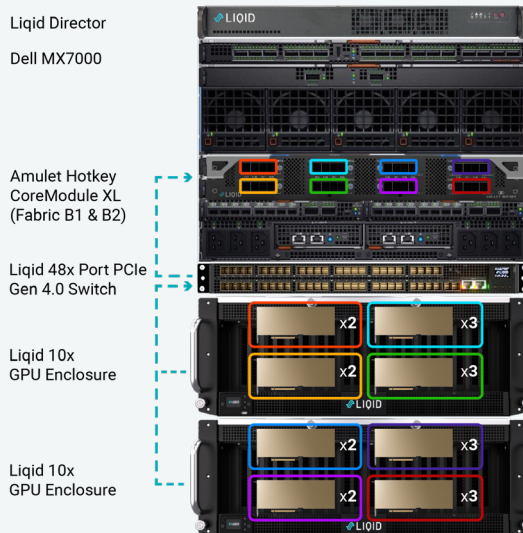
Overview

Blade server customers value the compact and efficient design of their systems, which maximizes space and reduces energy consumption in data centers. Despite these advantages, they lack GPU density, which significantly limits their capability to manage advanced AI, and graphics-intensive applications, and VDI.

Liquid UltraStack MX, designed specifically for the Dell PowerEdge MX7000 modular chassis, significantly enhances the value of MX compute sled investments. Utilizing a module designed by Amulet Hotkey, up to 8 hosts can be connected to the B fabric. This system enables customers to dynamically connect and scale enterprise-grade GPUs from NVIDIA, AMD, and Intel directly at the bare metal, facilitating workloads that were previously unattainable to modular systems.

By connecting pools of 10, 20, or 30 GPUs to the MX7000 chassis via our PCIe fabric, Liquid Matrix software enables customers to dynamically attach up to 20 full-height, full-length (FHFL) dual-slot GPUs to a single MX760c compute sled. Liquid's Peer-2-Peer RDMA across GPUs enhances throughput and reduces latency. Additionally, customers can maximize GPU resource utilization by reallocating GPUs between compute sleds as workload demands evolve.

Enable Dynamic GPU Allocation



Integrate GPUs into MX7000

- » Add and move GPUs between compute sleds
- » Scale up to 20 GPUs per MX compute sled
- » Leading density with up to 30GPU per MX enclosure
- » Supports heterogeneous FHFL 2W GPU types
- » Compatible with MX740c, MX750c, and MX760c
- » Connect GPUs to MX, R and C series servers simultaneously

Liquid not only enhances current server functionality but also future-proofs GPU investments (removed comma) by allowing GPUs to be provisioned not just to Dell PowerEdge MX compute sleds but also extending this capability to R-series and C-series rack servers. This adaptability enables seamless migration between platforms, enhancing utilization and agility while protecting investments.

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

Liquid SmartStack MX Series Technical Specifications



	SmartStack MX10	SmartStack MX20	SmartStack MX30	SmartStack MX30+
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™			
Key Features	Bare Metal GPU Provisioning, Granular Scale, Resource Reallocation, UI/ API/ CLI Mgmt., RMDA GPU-to-GPU and GPU-to-storage communication			
Management Appliance	1x Liquid Director 1U			
Maximum GPUs Per MX7000 Chassis	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
Supported Device Types	GPU, SSD, FPGA, DPU, NIC	GPU, SSD, FPGA, DPU, NIC	GPU, SSD, FPGA, DPU, NIC	GPU, SSD, FPGA, DPU, NIC
Maximum Host	4x Compute Sleds	8x Compute Sleds	8x Compute Sleds	16x Compute Sleds
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	Integrated Switch	1x 48 Port PCIe Switch	1x 48 Port PCIe Switch	2x 48 Port PCIe Switch
Rack Units	5U	10U	14U	15U
PCIe Host Bus Adapter	1x Fabric B CoreModuleXL w/ 4 PCIe Gen4 x16 HBAs	1x Fabric B CoreModuleXL w/ 8 PCIe Gen4 x16 HBAs	1x Fabric B CoreModuleXL w/ 8 PCIe Gen4 x16 HBAs	1x Fabric B CoreModuleXL w/ 16 PCIe Gen4 x16 HBAs
Warranty & Support	36 or 60 - Month Software and Hardware			