Data Sheet

CORESTATION
AMULET HOTKEY

CoreStation CoreStack for GPU

Unlocking Infrastructure Flexibility with Disaggregated CPU Resources

Features and Benefits

Key Features:

- Software-based dynamic on-demand allocation and provisioning of GPU resources
- Scalability: from one GPU and one external enclosure to thirty GPUs and three external enclosures
- Bare metal connectivity via a PCIe Gen4 network fabric
- Automated deployment
- Peer-to-peer performance
- Supports multi-vendor GPUs
- Also supports FPGA, NVMe SSD storage, and networking

Designed for:

- AI-centric and data science
- High performance computing
- End user computing
- Server virtualization
- Edge computing

Why CoreStation SmartStack

- Adapt to demands in GPU performance requirements
- Improve efficiency, resource utilization, and optimization of costly infrastructure
- Deliver real-time right-sized GPU resources

Flexible and Scalable GPU Acceleration On-Demand

Unblock Server Limitations with External GPU Enclosures

For AI-centric workloads, or any other workloads such as VDI, that require multiple GPUs to deliver the resources required by those workloads, traditionally, these deployments are very static, inflexible, and therefore costly to deploy and manage.

Workstations and GPUs in this traditional deployment model are configured on a 1:1 basis, having a dedicated GPU installed, and depending on chassis size, a limited number of supported GPUs. Additionally, power requirements and heat generation must be considered when running at the deskside.

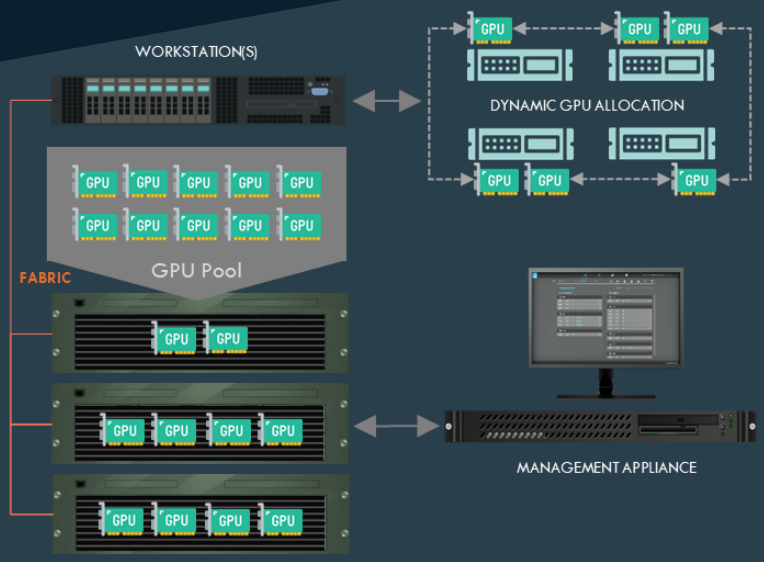

Scalability and Flexibility with Dynamically Composable GPUs

CoreStation CoreStack allows static barriers to be removed by abstracting the GPU resources from individual workstations by creating a datacenter-based centralized pool of GPU resource that can be dynamically allocated, using management software, to individual workstations as workloads and demands dictate.

GPU resources can be allocated to workstations dynamically on demand with the ability to deploy up to 30 GPUs per workstation!

Data Sheet

CORESTATION
AMULET HOTKEY

CoreStation CoreStack

Specifications

CoreStack10

Component	Configuration Options
Maximum number of GPUs per workstation	Up to 10 Full-height, full length 10.5", dual-slot
Maximum number of workstations supported	4
Composability software platform	Liqid Matrix Software
Management	1 x Liqid Director Management Appliance (1U)
PCIe External Expansion Chassis/Enclosure	1 x EX4410 PCIe Chassis (4U)
PCIe Host Bus Adapters	LQD1416 PCIe Gen4 (one per workstation)

CoreStack20

Component	Configuration Options
Maximum number of GPUs per workstation	Up to 20 Full-height, full length 10.5", dual-slot
Maximum number of workstations supported	8
Composability software platform	Liqid Matrix Software
Management	1 x Liqid Director Management Appliance (1U)
PCIe External Expansion Chassis/Enclosure	2 x EX4410 PCIe Chassis (8U)
PCIe Fabric Switch	1 x 48-port PCIe Switch (1U)
PCIe Host Bus Adapters	LQD1416 PCIe Gen4 (one per workstation)

CoreStack30

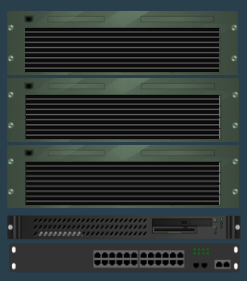

Component	Configuration Options
Maximum number of GPUs per workstation	Up to 30 Full-height, full length 10.5", dual-slot
Maximum number of workstations supported	16
Composability software platform	Liqid Matrix Software
Management	1 x Liqid Director Management Appliance (1U)
PCIe External Expansion Chassis/Enclosure	3 x EX4410 PCIe Chassis (12U)
PCIe Fabric Switch	1 x 48-port PCIe Switch (1U)
PCIe Host Bus Adapters	LQD1416 PCIe Gen4 (one per workstation)

Contact Us

EMEA Sales
+44 (0) 20 7960 2400
emeasales@amulethotkey.com

N America Sales
+1 212 269 9300
ussales@amulethotkey.com

APJ Sales
+61 409 930 884
apsales@amulethotkey.com