



Liquid Matrix Datasheet

Why Liquid

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

Key Advantages

- » Accelerate time to results with right-sized physical servers that can be deployed in minutes via software
- » Adapt in real-time to evolving business needs with a zero-touch, change-ready agility.
- » Drive new levels of efficiency with superior resource utilization and an as-a-service approach to infrastructure.

Key Features

- » Bare-metal Server Composability
- » Granular Scalability via Software
- » Simple UI, API, CLI Management
- » Cloud Automation Integration
- » Automated OS Deployment
- » Multi-Fabric Support (PCIe, Eth, IB)
- » Massive Scalability

Contact Information

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

Revolutionize IT Delivery

Unlock Infrastructure Flexibility And Efficiency With Liquid Matrix CDI

Overview

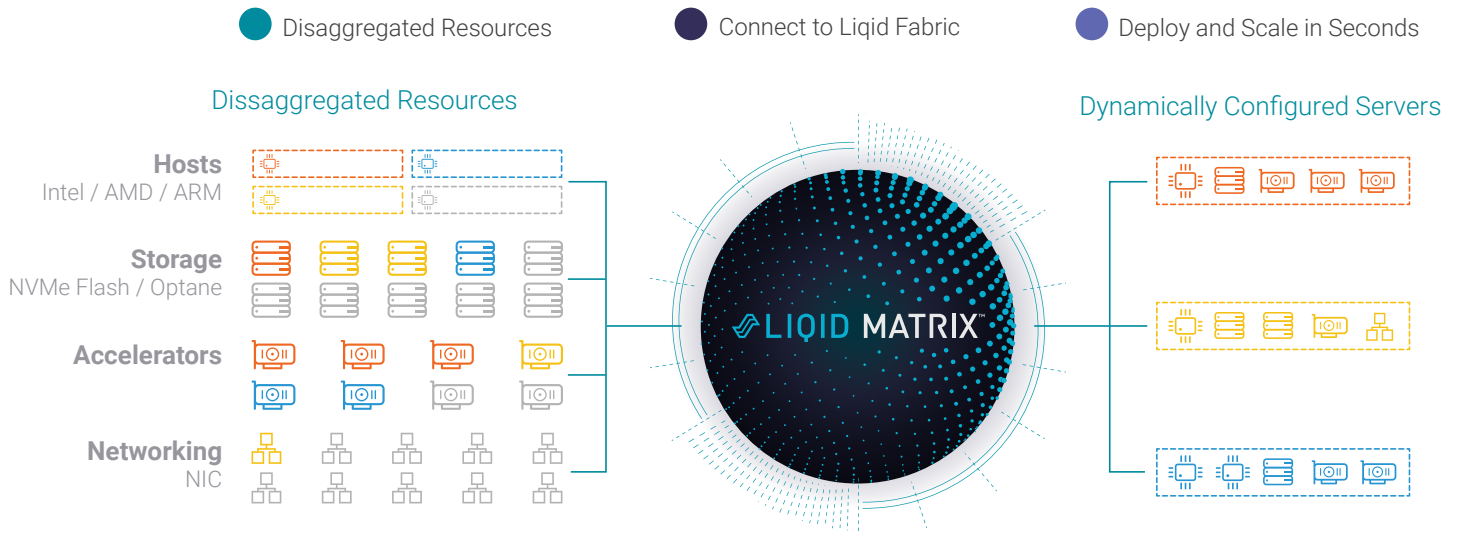
Stakeholder demand for fast results from modern workloads like AI and HPC continues to increase. Unfortunately, the conventional server infrastructure limitations delay workload time-to-value, limit configuration possibilities, and increase costs with underutilized resources and tedious manual tasks.

For organizations that seek to accelerate results while maximizing their IT investments, Liquid Matrix® is a new software defined approach to server configuration and management called composable disaggregated infrastructure (CDI) that accelerates application time-to-value, is change ready, and improves operations and capital efficiency at core and edge deployments.

With Liquid CDI valuable PCIe-connected resources like GPU, FPGA, NVMe SSD and Storage Class Memory are disaggregated from the server itself into PCIe enclosures. They are then interconnected to host servers over high-speed fabrics like PCIe, Ethernet and InfiniBand. Liquid Matrix software lives on the fabric, allowing IT to configure and deploy servers that meet explicit workload requirements in seconds via software without worrying if a server can physically support its GPU and or storage resources. If demand increases, add more resources on-demand. As business needs evolve, reclaim unused resources to be used by other servers.

After deploying and scaling servers the same way for 20+ years, Liquid Matrix enables IT organizations to reap cloud-like speed and flexibility in their own core and edge infrastructure.

Unlock New Levels of IT Speed and Efficiency



Accelerate Time to Value

Spend less time configuring and deploying servers and more time supporting strategic business goals with Liquid Matrix. Deploy bare metal servers that meet a workload's exact CPU to GPU requirements in seconds via Liquid UI, API or CLI, in addition to cloud automation tools and HPC schedulers like Slurm. Liquid is change ready, meaning as applications demand more resources, quickly add them in seconds. Fine-tune performance to meet application stakeholder needs via software without touching a single box or device.

Increase Agility

With compute, storage and GPU resources interconnected across the Liquid fabric, server configuration options are unconstrained. For example, while a 1U server with 10x GPUs is simply impossible with conventional server design, it can be composed and deployed with Liquid in a few short minutes. Dynamically add and remove resources online as demand evolves, without rebooting the OS. With resource-intensive workloads like AI and HPC on the rise resource agility is a very timely feature.

Improve Efficiency

Eliminate wasteful overprovisioning by composing only what's needed to meet a workload's current requirement, then scale granularly as needed. Maximize the resource utilization of valuable resource like GPU by assigning them to the free pool or even another server when underutilized. Liquid CDI allows organizations to decouple infrastructure purchasing decisions. Instead of replacing an entire server to upgrade a CPU, simply purchase a 1U server with new CPU and RAM and then compose the accelerator and storage resources from the old to new server. Liquid reduces OPEX and risk of component damage by eliminating manual adds/moves and changes with a true software defined approach.

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