LIQID

Liqid is the leader in software-defined composable infrastructure, delivering flexible, high-performance, and efficient datacenter and edge solutions for AI Inferencing, VDI, and HPC, solutions.

Liqid enables customers to manage, configure, reconfigure, and scale essential compute, accelerators (GPU, TPU, FPGA), memory, storage, and networking into physical bare metal server systems in seconds. Liqid customers can optimize their IT infrastructure and achieve up to 100% GPU utilization for maximum tokens per watt and dollar.

Top Reasons Why Liqid for Higher Education and Research

Academic institutions lead the way in discovering next generation emerging technologies, accelerated computing, data science, and AI. With limited budgets, institutions must balance the demand for HPC infrastructure with scarce resources. Liqid's software-defined composable infrastructure delivers flexibility, performance, and cost efficiency that deliver discoveries without limits.

1. Transform Campus Infrastructure with Software-Defined Composability

Liqid's software-defined composable infrastructure platform gives higher education and research institutions the ability to break free from static hardware constraints. By disaggregating, pooling, and sharing compute, GPU, storage, and memory resources, Liqid enables dynamic, software-defined reallocation based on real-time workload demands. This means your infrastructure can adapt as research priorities shift, departments evolve, or new projects launch—without having to purchase entirely new systems.

2. Support a Broad Range of Scientific and Academic Workloads

Whether it's genomics research, physics simulations, AI inferencing, or large-scale data analysis, academic institutions run some of the most diverse and compute-intensive workloads. Liqid provides a solution that can support all of these use cases simultaneously—dynamically shifting resources to wherever they're needed. This ensures researchers, faculty, and students always have the compute power they need to innovate.

3. Unlock and Share GPU Resources Across the Campus

GPUs are often siloed within specific departments or trapped in underutilized nodes, creating resource shortages in other areas. Liqid enables full GPU composability, allowing you to share GPUs across multiple systems, labs, and users to achieve 100% utilization. You get maximum ROI on every GPU purchased, and departments can collaborate more effectively by accessing a shared pool of high-performance accelerators as needed.

4. Maximize Investments in NVIDIA and High-Performance Hardware

Over 90% of higher ed and research departments are incorporating AI into their research, and 92% leverage GPUs and other specialized compute resources. Liqid amplifies the value of these assets by allowing them to be disaggregated and reallocated on demand. You can dynamically provision GPUs to different users or labs, scale GPU clusters quickly for timesensitive projects, and eliminate the need for every department to overprovision hardware.

5. Deliver Cloud-Like Agility On-Premises

Public cloud costs can spiral quickly in research settings, and many workloads require low-latency, high-bandwidth, or data-sovereign environments. Liqid helps you build a private cloud with the flexibility and automation of public cloud—right in your own datacenter. Resources can be spun up in seconds, and reallocated with ease, enabling rapid experimentation and innovation without sacrificing control or budget.

6. Accelerate Time-to-Insight for Research and Innovation

Getting new infrastructure online in academic settings often takes months—delaying critical discoveries and slowing progress. Liqid dramatically reduces time-to-deployment for complex environments like AI labs, VDI classrooms, or high performance research clusters. Faster provisioning means researchers can get to work quicker, and institutions can better compete for grant funding and academic prestige.

7. Enable Central IT to Do More with Less

With Liqid, central IT teams can become powerful service providers to the rest of campus. By pooling resources into a composable fabric, IT can support a wider range of projects, departments, and research centers without ballooning operational overhead. This efficiency means your IT team can handle more users and more compute-intensive workloads—without increasing staff or complexity.

8. Adapt Quickly to Changing Academic Priorities

In higher education, priorities shift frequently—from expanding AI programs to launching new scientific initiatives or responding to sudden spikes in usage. Liqid gives you the agility to react in real-time. Need to reallocate GPUs to a new data science course? Spin up more memory for a genomics lab? Expand storage for a growing archive? Liqid lets you do it all with a few clicks; no forklift upgrades required.

9. Protect Your Infrastructure Investment for the Long Term

Liqid's composable approach extends the lifecycle of your infrastructure. You can upgrade one resource type—like GPUs or storage—without touching the rest of the system. This protects your capital investment while keeping your environment current with the latest research and compute demands. It's a smarter, more sustainable approach to campus IT modernization.

10. Partner with an Innovator Trusted by Leading Research Institutions

Liqid is already helping more than two-dozen global, top-tier universities and research labs optimize their infrastructure for AI, HPC, and data-intensive science. When you work with Liqid, you're gaining a partner who understands the unique challenges of academic computing—and who's committed to supporting your mission of discovery, innovation, and learning. From deployment to ongoing support, we're with you every step of the way.