

Five Ways to Accelerate AI Infrastructure Readiness

Unlock Maximum GPU Performance — Using the Infrastructure You Already Own

1. Idle GPUs Are Expensive—Don't Let Storage Be the Bottleneck

Your GPU servers likely have terabytes of local NVMe that sit unused. Activating this capacity with Hammerspace Tier 0 cuts time-to-value from months to hours—turning depreciating GPU assets into immediate productivity.

2. Performance at Scale Demands Ultra-Fast, Localized Data Access

AI and HPC workloads can't wait for external storage. Tier 0 feeds GPUs directly from local NVMe with shared access - avoiding network latency and accelerating performance by up to 10x versus external shared storage.

3. Why Buy New Hardware When You Already Own the Speed?

Tier 0 eliminates the need for costly new storage systems or high-speed networking. It transforms what's already in your racks into a low-latency, high-throughput storage tier—reducing both CapEx and power.

4. AI Ready Infrastructure IS Hybrid-Cloud Ready Infrastructure

Modern AI isn't just on-prem or cloud—it's both. Hammerspace Tier 0 lets you achieve seamless, ultra-low-latency data access across both on-prem and cloud GPU environments—without compromising performance or workflow portability.

5. Infrastructure Complexity Slows You Down - Tier 0 Deploys in Hours

Forget 12-week timelines. Tier 0 is agentless, non-disruptive, and deploys quickly across heterogeneous environments - minimizing operational drag and speeding up AI deployment.

The Hammerspace Tier 0 Advantage

**Activate In Hours,
No New Storage Required**

**Feed On-Prem and Cloud
GPUs up to 10x faster**

**Reduce Storage Costs,
Power, and Rackspace**

**Unlock Hybrid-Cloud
Agility with Consistent
High-Speed Access**



HAMMERSPACE

