Hammerspace Tier 0

SOLUTION BRIEF

Turn GPU Server-Local NVMe Into a New Tier of Ultra High-Performance Shared Storage

AI training, checkpointing, inferencing, and agentic AI demand high-throughput, low-latency access to large volumes of unstructured data. To meet performance demands, organizations deploy expensive external flash storage arrays and high-speed networking—delaying AI initiatives and consuming significant budget, power, and rack space.

Hammerspace Tier 0 solves this by activating the local storage you already own within a cluster of GPU servers and turning it into a new tier of high-performance shared storage - managed and protected by Hammerspace.

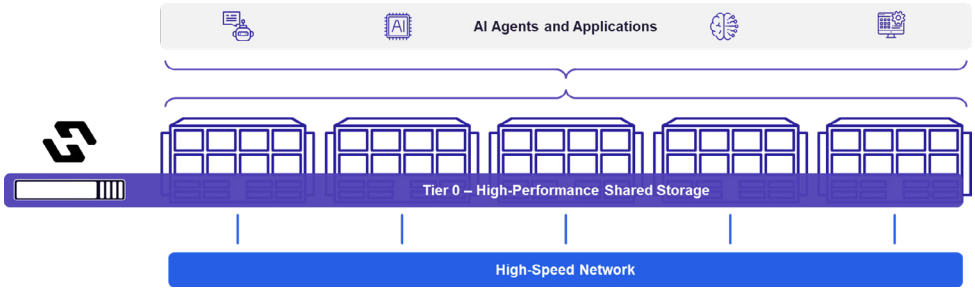

Benefits

Deploy in Hours, No Storage Required
Activate Tier 0 capacity in hours, using the storage and networking infrastructure you already have in place. No complex agent installs, no exotic networking, just standards-based simplicity.

Feed On-Prem and Cloud GPUs up to 10x faster
Avoid network bottlenecks and unlock local NVMe speed to increase GPU utilization, complete checkpoints faster, and improve response times for inferencing and agentic AI.

Cut Storage Costs, Power, and Rackspace
Leverage existing GPU server NVMe to reduce external storage footprint, power consumption, and rack space.

Unlock Hybrid-Cloud Agility with Consistent High-Speed Access
Achieve seamless, low-latency data access across both on-prem and cloud GPU environments.

HAMMERSPACE

Activate Tier 0 In Hours, No New Storage, No Complex Agents to Install

Tier 0 capacity can be brought online quickly, and you can be serving data to your users and applications from Hammerspace in half a day. Deploy Hammerspace software, add Tier 0 volumes to Hammerspace, then connect your applications and users. Put your GPUs to work quickly and minimize costly GPU depreciation expenses.

Deploy Hammerspace Tier 0 and Connect Your Users and Apps in Three Easy Steps

1 | Deploy Hammerspace

- Bare metal, VM, cloud
- Hammerspace-Ready Nodes
- No client software to install
- Takes less than 30 minutes

2 | Add Tier 0 Volumes

- Add Tier 0 vols to Hammerspace
- Assimilate metadata from storage
- Create and configure shares
- Set data policies (“Objectives”)

3 | Connect Apps and Users

- All data visible in minutes
- Unified Global Namespace
- Multi-protocol access
- Automate protection and mobility

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

Using Hammerspace Tier 0 means data is written at local NVMe speed, which can be up to 10x faster than sending data over a 100Gb/sec network to external storage. And the advantage in the cloud is even more pronounced. Tier 0 uses the local NVMe drives that are included in many cloud GPU instance types, and also has the added benefit of using the fast East-West inter-node networking between compute nodes.

Avoid network bottlenecks, increase GPU utilization, complete checkpoints faster, and improve response times for inferencing and agentic AI.

On-Premises [1]

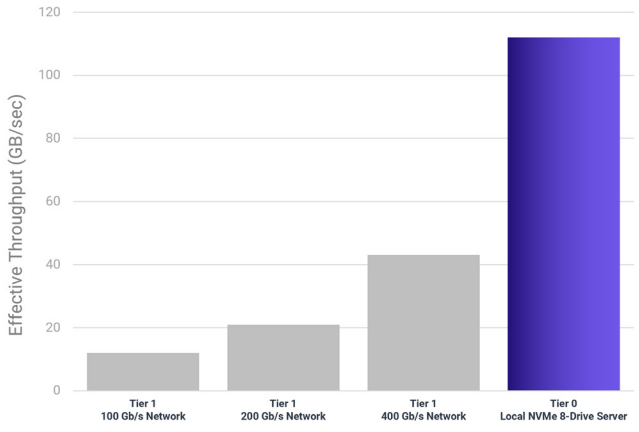

Cloud [1]

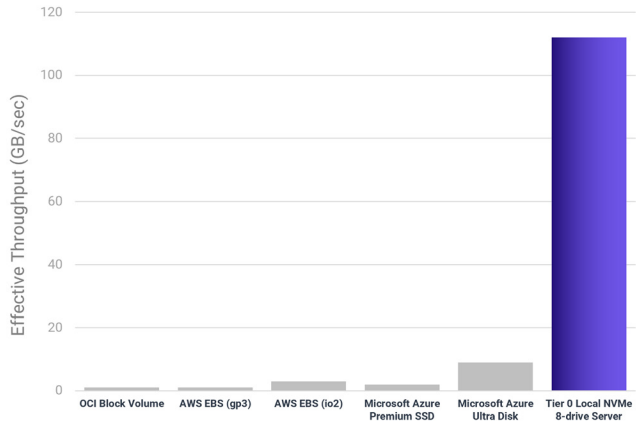

Tier 0 Reduces Storage Costs, Power and Rackspace

Activating Tier 0 capacity offsets and reduces the need for external storage.

This in turn reduces the storage and networking costs, power consumption, and rack space that those storage servers and switches would otherwise consume.

This chart shows the order of magnitude savings in a 20-node GPU cluster, and in large clusters with hundreds or thousands of GPU servers, you can save millions.

Cost, Power, and Rack Savings Based on Typical Tier 0 Capacity in a 20-node GPU Cluster

$149.9K
Storage and networking cost savings[2]

$16.5K
Energy costs saved over 3 years[2]

109,766 KwHs
Power consumption saved over 3 years[2]

6U
Rackspace saved[2]

Enable Hybrid-Cloud and Multi-Cloud Agility with Hammerspace

Join edge, core, and multiple cloud regions together in a single Global Namespace using the Hammerspace Data Platform. Bring your data to the GPUs, no matter where they are located, with Hammerspace Data Orchestration policies, and serve that data at local NVMe speed using Hammerspace Tier 0 - on-prem or in-cloud.

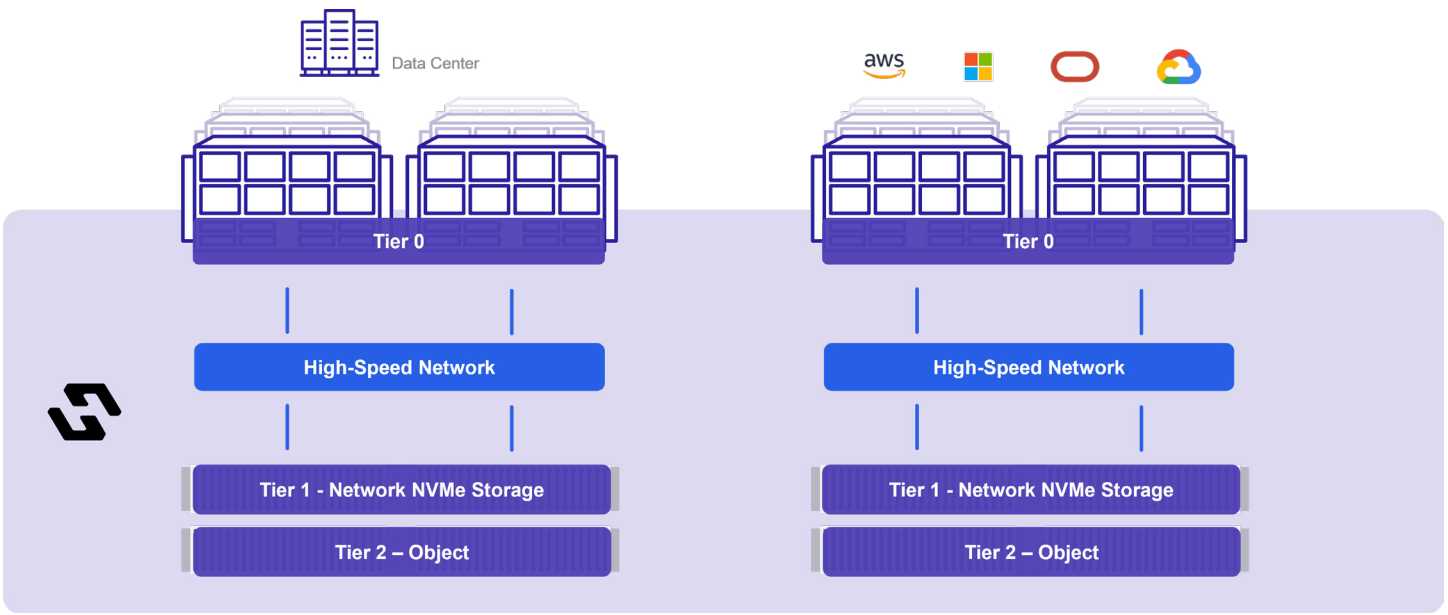

1 Source: Activating Untapped Tier 0 Storage Within Your GPU Servers

2 Source: Tier 0 Savings Calculator based on a 20-node GPU cluster with 487 TBs of usable capacity.