

Hammerspace Tier 0 ROI Analysis

EXECUTIVE BRIEF

Introduction

Hammerspace Tier 0 turns GPU server-local NVMe into high-performance shared storage, managed and protected by Hammerspace. Tier 0 storage is up to 10x faster than networked storage - so you can reduce checkpointing time for AI training and HPC, and improve response times for inferencing and agentic AI. It can be activated in hours - on-prem or in the cloud - with no clients or agents to install, so you can put GPUs to work immediately. And it enables optimal use of assets you already own so you can reduce the need for external flash storage - to reduce costs, and gain back the power and rack space those systems would otherwise consume.

Hammerspace Tier 0 delivers ROI and cost savings in two primary ways that are detailed in this analysis.

1. **GPU Depreciation Expense Avoidance:** Because Tier 0 capacity can be deployed in hours, not months, GPUs can be put to work quickly. This can result in avoiding months of GPU depreciation, which can equate to hundreds of thousands of dollars.
2. **Cost, Power, and Rackspace Savings:** By offsetting and reducing the need for external storage, Hammerspace Tier 0 reduces storage and networking costs as well as the associated power and rack space that those storage systems and switches would otherwise consume.

Note about Tier 0 Increasing GPU Utilization: Tier 0 also increases GPU utilization, on-prem and in-cloud, relative to external networked storage, because it delivers data to the GPUs at local NVMe speeds, and uses the faster east/west networking between compute nodes. Although this can represent a material benefit in terms of ROI on GPUs, this element is not quantified in this analysis.

Tier 0 GPU Depreciation Expense Avoidance

GPUs are typically depreciated over a 3-year schedule, reflecting their useful life in accelerated computing environments, and here we assume a 3-year straight-line depreciation schedule for GPUs, aligned with standard accounting practices for IT hardware.

The following table shows depreciation schedules for various GPU capital investment levels:



Initial GPU Capital Expense	Depreciation Schedule	Annual Depreciation	Quarterly Depreciation	Monthly Depreciation
\$1,000,000	3 years	\$333,333	\$83,333	\$27,778
\$2,000,000	3 years	\$1,000,000	\$250,000	\$83,333
\$3,000,000	3 years	\$1,666,667	\$416,667	\$138,889
\$4,000,000	3 years	\$3,333,333	\$833,333	\$277,778

For example, under a straight-line depreciation schedule, a \$3M GPU investment incurs a depreciation expense of \$250,000 per quarter over a 3-year useful life. This depreciation does not impact cash flow directly, but it reduces taxable income and helps spread the cost of the asset over its operational lifespan.

Organizations typically purchase and deploy their GPU computing infrastructure, then evaluate, purchase, and deploy external storage systems to feed data to those GPUs, and this can result in a delay of months between when GPU compute nodes are online and the external storage is online and available.

Then there is the time that is typically required to migrate data into a new storage system, and this can also be months. As shown in the table above, a delay of even one or two quarters can quickly exceed \$1M in depreciation expenses.

Conclusion: Hammerspace Tier 0 simplifies all of this complexity and delay. Terabytes, even Petabytes of Tier 0 capacity can be brought online in hours, not months, so you can put your GPUs to work quickly and avoid GPU depreciation expense.

Tier 0 Cost, Power, and Rackspace Savings

By offsetting and reducing the need for external storage, Hammerspace Tier 0 reduces storage and networking costs as well as the associated power and rack space that those storage systems and switches would otherwise consume.

Because of the power demands of large GPU computing clusters and the constraints that has put on data center resources, in some cases the power and rackspace savings are as important, or more important, than the storage CapEx savings.

Analysis Methodology and Assumptions

- Analysis uses the Hammerspace Tier 0 Savings Calculator, and looks at the cost, power, and rack savings at different sizes of GPU computing cluster: 10-node, 25-node, 50-node, 100-node, and 1000-node clusters
- Analysis assumes local NVMe storage is a sunk cost, in other words it is included with the purchase of the GPU servers and/or is already in-place.
- Assumes that Tier 0 storage useable capacity offsets external NVMe storage usable capacity 1:1, using three-way client-side mirroring as the protection scheme for Tier 0 storage, and erasure coding as the protection scheme for the external flash storage.



- Assumes 7.68TB SSD capacities used in the GPU servers and the external storage servers. (The SSD capacity used in both the GPU servers and in the external storage servers can have an impact on the comparison. The Tier 0 Savings Calculator allows for looking at comparisons across different capacity points.

Create your own custom ROI analysis at hammerspace.com/tier-0.

Analysis Results

The results of the analysis are shown in the table below.

Size of GPU Cluster	10-node	20-node	50-node	100-node	1000-node
Tier 0 Usable Capacity Available (TBs)	243	487	1,217	2,434	24,340
Storage and Networking Cost Savings	\$74,844	\$149,997	\$374,838	\$749,676	\$7,496,756
Power Cost Savings over 3 years	\$8,215	\$16,465	\$41,145	\$82,291	\$822,907
Reduced Power Consumption over 3 years (Kwhs)	54,770	109,766	274,302	548,605	5,486,047
Rack Space Saved (Rack Units)	4U	6U	14U	28U	266U (6-7 racks)

Conclusion: By offsetting and reducing the need for external storage, Hammerspace Tier 0 reduces storage and networking costs as well as the associated power and rack space that those storage systems and switches would otherwise consume. In large GPU clusters, the savings is substantial, and ROI is immediate.



Get Started

Want to see what savings you could achieve
in your environment?

[Schedule a 30-minute technical consultation](#)

