

Hammerspace and Hitachi Vantara

Supercharge Your Hybrid Cloud AI Data Workflows with a Full Stack AI Infrastructure Solution




SOLUTION BRIEF

Consolidate and Orchestrate File and Object Data, Deliver Performance and Scale for AI/HPC Workloads

The opportunity of AI is forcing Enterprise IT to address data and infrastructure constraints that have been in place for decades. This is driving a major build out in new AI infrastructure – both on-premises and in the cloud – that requires a new stack of GPU computing, high speed networking, and storage that can meet the performance and scale demands of Large Language Model (LLM) training and tuning, RAG, and inferencing. Data orchestration – in this case the ability to bring large unstructured data sets to the GPUs that need them no matter where those GPU resources are located – has also emerged as a critical requirement for hybrid-cloud AI infrastructure.

To address these new requirements, Hammerspace and Hitachi Vantara have partnered to offer an integrated solution for AI infrastructure that combines the Hammerspace Global Data Platform for hybrid-cloud file and object workloads with Hitachi Virtual Storage Platform (VSP) One as a flexible, resilient, and scalable data storage platform.




Benefits:

- **Eliminate Data Silos** by consolidating file and object data onto a single Global Data Platform that spans sites, clouds, and storage
- **Deliver Performance and Scale for AI Workloads** with linearly scalable performance to feed a few to tens of thousands of GPUs
- **Faster Time to Value** with a full stack infrastructure solution that combines Global Data Platform software with flexible, resilient, scalable data storage
- **Reduce Storage Footprint and Go Green** with Hitachi Advanced Data Reduction (ADR), intelligent power consumption, and a sustainability dashboard



Simplify Deployment and Speed Time to Value with Validated Reference Architecture that Scales to Meet Your Needs

Hammerspace and Hitachi have partnered to define a scalable architecture for AI infrastructure. In this architecture, Hammerspace provides high performance data access using industry-standard protocols including NFSv4.2, NFSv3, SMB, S3, and NVIDIA GPUDirect Storage (GDS), and delivers linearly scalable performance based on the Hyperscale NAS architecture.



Hitachi Vantara servers are used for the Hammerspace metadata and data servers, while storage is provided by the Hitachi Vantara VSP One Block 28 NVMe storage platform. The same basic design can be adjusted for use by other VSP One Block platforms to meet the required performance and cost profile.

Capacity for Hammerspace with Hitachi storage can be provided from the VSP One Family of products: VSP One Block, VSP One File, and VSP One Object.

