Hammerspace and NVIDIA

Make your storage GPUDirect Storage to accelerate data pipelines and speed up parallel processing

SOLUTIONS BRIEF

Overview

NVIDIA GPUDirect is a family of technologies that provides for more direct read/write access between NVIDIA data center GPUs and connected storage systems. NVIDIA GPUDirect Storage uses Remote Direct Memory Access (RDMA) to streamline the path between GPUs and storage.

Specific to Hammerspace, NVIDIA GPUDirect enables a direct path for data exchange between NVIDIA GPUs and Hammerspace to enable faster data access.

Because of Hammerspace's unique architecture, Hammerspace can make existing storage systems GPUDirect Storage, simply by making those storage systems part of a Hammerspace Global Data Environment.

There are three options to make storage GPUDirect Storage, including:

- Any RDMA-capable NAS storage
- Any block storage (typically NVMe flash) attached to a Hammerspace Data Services (DSX) node
- Any non-RDMA NAS storage fronted by a DSX node

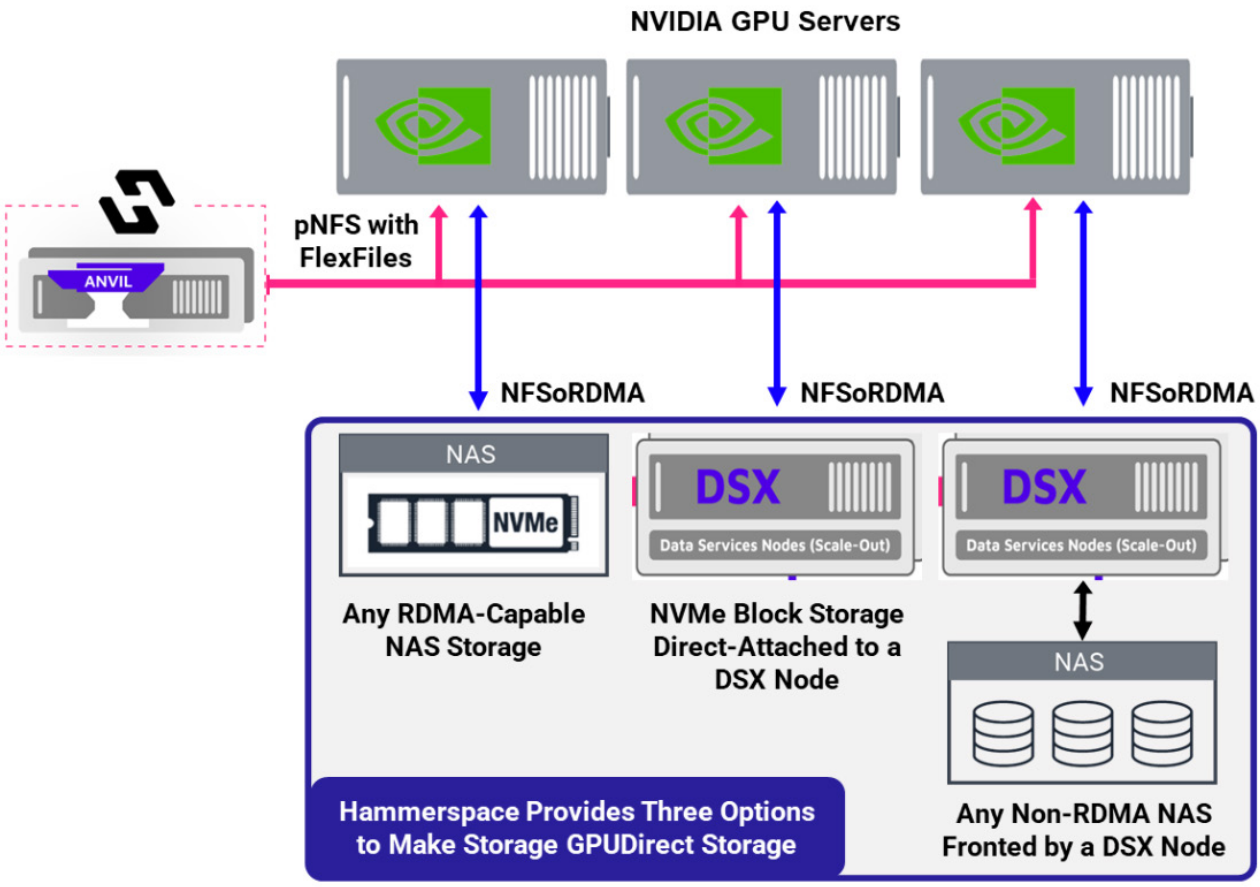
Solution Benefits

- **Make Any Storage NVIDIA GPUDirect Storage:** Deploy Hammerspace on top of any existing storage to effectively make that storage GPUDirect Storage.
- **Improve Throughput and Reduce Latencies:** Use RDMA technology to streamline the data path between GPUs and storage to improve throughput and reduce latencies
- **Maximize GPU Resource Utilization at Any Scale:** Serve thousands of users and tens of thousands of GPUs concurrently. Scale linearly from a few nodes to tens of thousands of storage nodes.
- **Easily Move Data to GPU Resources:** Seamlessly and transparently move directories and files to different data center locations, and even different clouds, to move data to available GPU resources.

Because of Hammerspace's unique architecture, Hammerspace can make existing storage systems GPUDirect Storage, simply by making those storage systems part of a Hammerspace Global Data Environment.

© 2023 - Hammerspace

Use Case Diagram



About Hammerspace

Hammerspace software makes data available to different foundational models, remote applications, decentralized compute clusters, and remote workers to automate and streamline data-driven development programs, data insights, and business decision-making. Organizations can source data from existing file and object storage systems to unify large unstructured data sets for AI training.

And now, Hammerspace's Hyperscale NAS architecture can scale out without compromise to saturate the performance of even the most demanding network and storage infrastructures. It is the first to combine the performance and scale of parallel HPC file systems with the simplicity of enterprise NAS, bringing new levels of speed and efficiency to data storage in order to address emerging AI and GPU computing applications.

HAMMERSPACE DATA IN MOTION | NVIDIA

© 2023 - Hammerspace