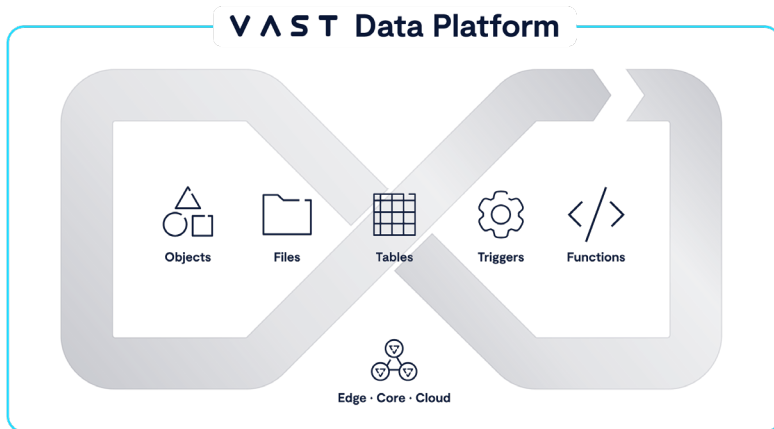


# Accelerate AI pipelines with the VAST Data Platform & Supermicro



## The Data Platform For The AI Era

Designed for Cloud Service Providers and Enterprises, the VAST Data Platform is a breakthrough approach to data-intensive computing that serves as the comprehensive software infrastructure required to capture, catalog, refine, enrich, and preserve data through real-time deep data analysis and deep learning.



## VAST Data Platform + Supermicro Hyperscale Servers

VAST Data, the leading AI Data Platform provider, collaborated with Supermicro, a leader in AI and GPU servers, to deliver a full-stack AI solution for LLM and Generative AI workloads. This strategic alliance equips Cloud Service Providers and large data-centric organizations with a scalable, high-performance AI infrastructure through the VAST Data Platform, augmented by Supermicro's robust hardware. With a starting point of 1.4PB, the solution allows seamless scaling to exabyte levels, anchored by an impressive 99.999% uptime and the cutting-edge VAST DASE™ (Disaggregated Shared Everything) architecture, enabling customers to affordably house data in a single tier to achieve the highest levels of GPU utilization.



Automatically bring structured context to all unstructured data



Accelerate all transactional and analytical workloads



Eliminate data silos with a unified global namespace from edge-to-cloud



Enable AI-based workloads on all your data



Secured by Zero Trust and Multi-Tenant architecture

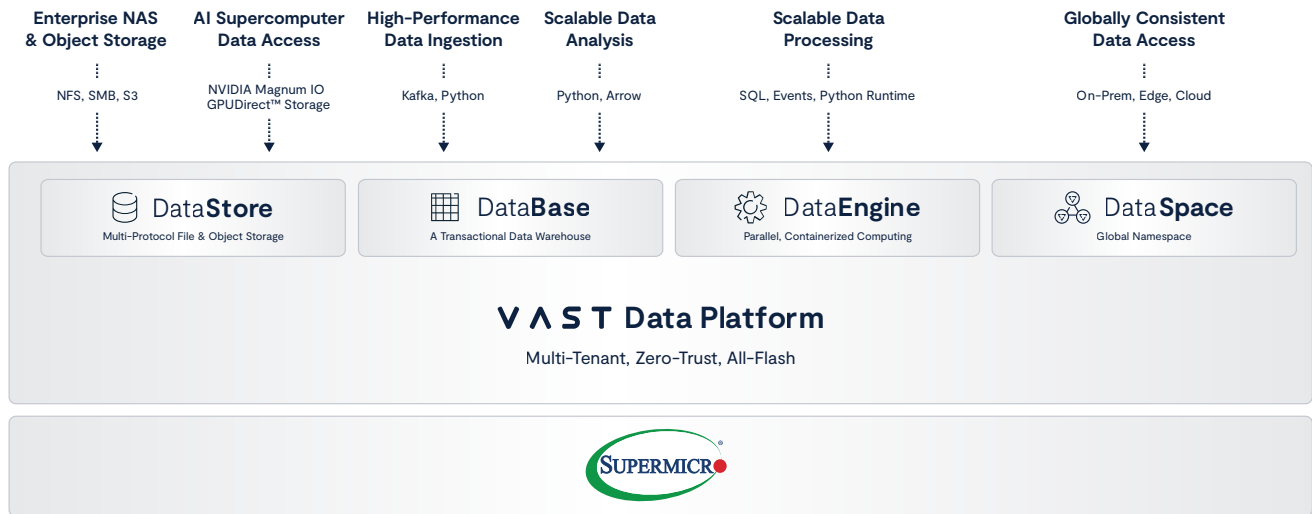
**Limitless Scale and High GPU Utilization.** Efficiently scale to exabyte levels with VAST + Supermicro, while ensuring 99.999% uptime - with pre-configured systems that enable rapid deployment.

**Linearly-Scalable Parallel Architecture.** Unlock performance with VAST DASE™, the world's first truly parallel system architecture, which makes it possible to feed the world's largest GPU clusters.

**Unified Support for Data Prep and Data Training.** Accelerate feature engineering with VAST + Spark while simultaneously feeding large training and inference workloads.

**Edge to Cloud Data Access.** Achieve instant data availability across locations using VAST's global namespace with strict write consistency and the performance of local systems.

# The VAST Data Platform Accelerates AI Pipelines



With its ability to natively communicate with applications, the VAST Data Platform is designed to handle almost any type of workload. This enables data to be consumed or augmented by different applications without additional data copies. This has a significant improvement on data pipelines and drastically reduces the total time traditionally required for AI pipelines.

Aligned with our vision of seamless scalability, VAST Data and Supermicro offer a rack-scale design for quick, effortless expansion. This innovative approach features pre-configured systems, enabling rapid deployment from installation to operation. This strategic advancement empowers Cloud Service Providers and data-centric enterprises to easily enhance their AI infrastructure, ensuring scalability is as intuitive as it is efficient.

Data Access	Enterprise Features	Data Management	Multi-Tenancy
<ul style="list-style-type: none"> <li>• <b>Unstructured</b>—Multi-Protocol: NFS, NFS/RDMA, SMB, NFS, S3, NVIDIA Magnum IO GPUDirect™ Storage, VAST Multipath Client, K8S CSI</li> <li>• <b>Structured</b>—Support for Spark, Trino, VAST SQL, and more</li> <li>• 1 million transactions per second with terabytes of query throughput at exabyte scale</li> </ul>	<ul style="list-style-type: none"> <li>• Snapshot and Object immutability</li> <li>• Async replication with automated failover</li> <li>• Access Audit</li> <li>• Advanced Data Reduction to reduce TCO</li> <li>• Similarity-based Data Reduction</li> </ul>	<ul style="list-style-type: none"> <li>• Global Namespace across private and public clouds</li> <li>• Local write performance with decentralized lock management</li> <li>• Synchronization policies for optimized data placement</li> <li>• Consistency modes for global read/write or point-in-time copies</li> </ul>	<ul style="list-style-type: none"> <li>• Quality of Service (QoS)</li> <li>• Per-Tenant encryption at rest and flight</li> <li>• Tenant-level reporting</li> </ul>

## Infrastructure

Scalable to Exabytes · Performance for AI workloads · Archive Economics · Full ACID transactions

## Trusted By Industry Leaders

