

VAST AI Operating System



The AI revolution is not merely an evolution; it's a fundamental paradigm shift demanding a new foundational approach to data and compute. Traditional computing architectures, born from a pre-AI era, are ill-equipped and struggling to meet the moment with their disparate silos of storage, databases, and compute. VAST Data introduces the VAST Operating System for AI (VAST AI OS), a platform reinvented from the ground up to unify and simplify the entire AI data pipeline. This is the next inevitable step in computing's evolution, built for the age of AI agents and beyond, empowering both AI researchers who demand extreme performance and C-suite executives driving business innovation. The moment to embrace this change is now, as the accelerating AI era waits for no one.

VAST The Operating System for the Thinking Machine

VAST Data AI OS is the first true Operating System for AI, engineered from the ground up to consolidate data storage, database, and AI compute into a single, coherent, unified layer. It eliminates the crippling complexity and need to piece together a patchwork of disparate tools never designed for the scale and nature of AI, providing a singular OS that manages everything from raw data to intelligent action in one place.

Built for today's AI and the intelligent, autonomous systems of the future.



End-to-end AI infrastructure in a single OS:

The first AI Operating System that unites storage, database, and compute. Eliminate stitching disparate tools; manage everything from raw data to intelligent action in one place. The AI era's foundational OS.



Infinite Scale. Lightning Speed. No compromises:

VAST's DASE (Disaggregated Shared Everything) architecture processes massive data in real time for demanding AI agentic workloads and large GPU clusters. No dataset too vast, no model too hungry for this platform. VAST delivers 70% lower latency and resource overhead for AI data operations.



AI Simplified. Powerful AI at Your Fingertips:

VAST's AI OS is radically simple. AI and IT teams focus on insights, not complex infrastructure. Eliminate laborious data management and silos. Lower burden, deploy AI faster.



Born for AI. Future-Proof. Ready for Next:

VAST AI OS is purpose-built by AI visionaries for AI and agentic computing. Supporting autonomous and generative AI, it's flexible across cloud, on-prem, and edge for lasting adaptability.

Why VAST Data AI OS?

The Differentiators

- **End-to-End AI Operating System:** A complete, tightly integrated software stack for data storage, management, and processing, replacing a patchwork of fragmented solutions.
- **Extreme Performance at Scale:** Co-designed with hardware innovators for massive throughput and parallelism, embedding services directly on GPUs/DPU to feed data at the rate models demand, scaling linearly.
- **Radical Simplicity & Unified Experience:** Obliterates the complexity of separate storage tiers, databases, and ETL pipelines with a unified namespace and API. One system to install, monitor, and scale, deployable on your choice of hardware or cloud.
- **Enterprise Security & Manageability:** Mature, built-in features including data protection, encryption, snapshots, multi-tenant isolation, and policy-based automation – AI-ready and enterprise-ready without compromise.
- **Agentic Computing Enablement:** Explicitly designed for emerging AI agents and autonomous processes with an event-driven DataEngine, supporting the next wave of AI applications out-of-the-box.



Inside the AI Operating System

InsightEngine | AgentEngine: Orchestrate the AI Ecosystem

Leverage VAST infrastructure services to build understanding from experiences, connecting AI agents for collaborative learning and action across your entire data landscape.

VAST DataStore: AI's Infinite Memory Vault

Holds every raw data point—images, video, text, events—forming the AI's foundational experiences. Exabytes of all-flash storage, instantly fueling AI learning and insight.

VAST DataBase: AI's Knowledge Base

This global database indexes all data and AI metadata. Search trillions of vectors and other tabular data in milliseconds. It enables your AI to reason over everything it knows.

VAST DataEngine: Real-Time Actions

A containerized compute engine that acts as the AI's central nervous system, instantly processing signals (new data/events) and dispatching instructions for immediate action.

VAST DataSpace: Unified & Global

Unifies components into a global AI brain, spanning edge to cloud. All data and compute share one namespace, security, and APIs. It's one secure cohesive AI OS, tying everything together.

Empowering AI Innovation

For AI Researchers & Developers:

- **Unleash Creativity at Scale:**
Access petabytes of data and thousands of GPUs. VAST's AI OS feeds models with endless data at high throughput, accelerating training and inference.
- **Faster Discovery:**
Run complex reasoning and search massive unstructured datasets in seconds. Instantly find insights with on-the-fly indexing and querying of AI-generated metadata.
- **Agentic Computing Ready:**
Build autonomous AI agents with integrated "memory" (DataStore), "knowledge" (DataBase insights), and "compute reflexes" (DataEngine), all in one platform.
- **No Ops Headaches:**
Automate data management, placement, and scheduling of AI workloads across edge-to-cloud environments, keeping your focus on innovation, not infrastructure.

For C-Suite & IT Executives:

- **AI at Business Speed:**
Drastically shorten time-to-deploy AI initiatives from months to weeks, turning rapid AI experimentation and productization into a competitive advantage.
- **Simplify & Save:**
Consolidate data lakes, databases, and special-purpose clusters into one platform, lowering costs, eliminating data copies, and reducing data center power/footprint.
- **Scalable to Ambition:**
Linearly scale from 100 to 100,000+ GPUs to meet any demand, ensuring your infrastructure won't bottleneck growth.
- **Enterprise-Grade Trust:**
Benefit from robust security (zero-trust isolation), high reliability, and governance features (smart tagging, audit trails) – an AI platform that meets CIO checkboxes while enabling CIO dreams.
- **Future-Proof Innovation:**
Position your organization at the forefront of the next tech revolution by democratizing access to AI. Embrace VAST's AI OS to lead in the age of agentic AI.



Proven Impact:



Trusted Scale

Hundreds of petabytes.
Hundreds of thousands of GPUs.



Unmatched Performance

Gigabytes/sec per GPU throughput in real customer environments.



Significant Efficiency

Reduced infrastructure sprawl by 70% for leading AI cloud providers.



Lower TCO

All-flash performance at 50% lower TCO than HDD storage systems.



Pioneering Adoption

AI leaders like CoreWeave & xAI power their most demanding workloads.

The AI Operating System: The Next Paradigm Shift

Your AI deserves an Operating System built for it. VAST Data AI OS is Unifying Data, Compute & Intelligence – All in One OS – built for the age of AI agents.

We don't run on an OS; **We are the OS**

About VAST Data

VAST Data is the leading data platform software company bringing businesses into the AI era. By accelerating time-to-insight for workload-intensive applications, the VAST AI Operating System delivers scalable performance, simplifies data management and enhances productivity. Launched in 2019, VAST has become the fastest-growing data infrastructure startup in history.

Looking to simplify your data center and unlock insights from all of your data?

Contact us at hello@vastdata.com

Trusted By Industry Leaders

verizon

NASA

CoreWeave

NIH

servicenow

Lawrence Livermore
National Laboratory

Crusoe

RESEARCH

Lambda

XI

Chan
Zuckerberg
Initiative

jump trading

agoda