



EQUITY RESEARCH

UPDATED

08/04/2025

VAST Data

TEAM

Jan-Erik Asplund
Co-Founder
jan@sacra.com

Marcelo Ballve
Head of Research
marcelo@sacra.com

DISCLAIMERS

This report is for information purposes only and is not to be used or considered as an offer or the solicitation of an offer to sell or to buy or subscribe for securities or other financial instruments. Nothing in this report constitutes investment, legal, accounting or tax advice or a representation that any investment or strategy is suitable or appropriate to your individual circumstances or otherwise constitutes a personal trade recommendation to you.

This research report has been prepared solely by Sacra and should not be considered a product of any person or entity that makes such report available, if any.

Information and opinions presented in the sections of the report were obtained or derived from sources Sacra believes are reliable, but Sacra makes no representation as to their accuracy or completeness. Past performance should not be taken as an indication or guarantee of future performance, and no representation or warranty, express or implied, is made regarding future performance. Information, opinions and estimates contained in this report reflect a determination at its original date of publication by Sacra and are subject to change without notice.

Sacra accepts no liability for loss arising from the use of the material presented in this report, except that this exclusion of liability does not apply to the extent that liability arises under specific statutes or regulations applicable to Sacra. Sacra may have issued, and may in the future issue, other reports that are inconsistent with, and reach different conclusions from, the information presented in this report. Those reports reflect different assumptions, views and analytical methods of the analysts who prepared them and Sacra is under no obligation to ensure that such other reports are brought to the attention of any recipient of this report.

All rights reserved. All material presented in this report, unless specifically indicated otherwise is under copyright to Sacra. Sacra reserves any and all intellectual property rights in the report. All trademarks, service marks and logos used in this report are trademarks or service marks or registered trademarks or service marks of Sacra. Any modification, copying, displaying, distributing, transmitting, publishing, licensing, creating derivative works from, or selling any report is strictly prohibited. None of the material, nor its content, nor any copy of it, may be altered in any way, transmitted to, copied or distributed to any other party, without the prior express written permission of Sacra. Any unauthorized duplication, redistribution or disclosure of this report will result in prosecution.



VAST Data

Data platform unifying storage, compute, and database for large-scale AI workloads

#ai

[Visit Website](#)

Details

HEADQUARTERS

New York, NY

CEO

Renen Hallak



REVENUE

\$200,000,000

2025

FUNDING

\$531,000,000

2023

Revenue




VAST Data

ARR & ARR Growth Rate

\$200.0M

YoY



Sacra estimates that VAST Data hit \$200 million in annual recurring revenue (ARR) in January 2025. The company currently projects reaching \$600 million ARR by 2026.

VAST Data reports that it has achieved free cash flow positive status. The company's top 100 new customers average \$1.2 million in commitments, with three customers representing over \$100 million in total commitments each.

Valuation

VAST Data is in talks to raise funding at a valuation of up to \$30 billion from Alphabet's CapitalG and existing investor Nvidia. This represents a significant jump from the company's \$9.1 billion valuation in December 2023.

The company has raised \$380 million in total funding to date. The most recent round was a \$118 million Series E in December 2023 led by Fidelity Ventures. Key investors include Nvidia, Dell Technologies Capital, Goldman Sachs, Tiger Global, NEA, BOND Capital, Drive Capital, Next47, 83North, Norwest Venture Partners, and Mellanox Capital.

The potential new funding round could close within weeks and would make VAST Data one of the most valuable private technology companies globally, reflecting investor appetite for AI infrastructure plays as data center buildouts accelerate.

Product

VAST Data is a unified data platform that collapses traditional storage, database, and compute infrastructure into a single all-flash system optimized for AI workloads. Instead of organizations having to manage separate storage tiers, feature stores, data warehouses, and compute clusters, VAST provides one integrated platform where data can be stored, indexed, queried, and processed without moving it between systems.

The core architecture separates stateless compute controllers from dense flash storage enclosures, all connected by high-speed RDMA networking. This allows any compute node to access any piece of data with nanosecond-scale latency across the entire cluster. A machine learning team can store raw training data as files, automatically index and catalog that data for discovery, run SQL queries to subset datasets, and execute preprocessing jobs directly on the same hardware without copying data elsewhere.

The platform provides multiple data access methods simultaneously. Data scientists can access the same dataset through standard file protocols like NFS, object storage APIs like S3, or SQL queries through the integrated database engine. The VAST Catalog automatically extracts metadata from files and makes everything searchable, so users can run queries like finding all video files from 2023 with specific processing tags. The DataEngine component allows Spark jobs and Python workloads to run directly where the data lives, eliminating the traditional extract-transform-load bottlenecks that slow down AI development cycles.

For AI training specifically, VAST optimizes for the high-throughput, parallel data access patterns that GPU clusters require. Multiple GPU servers can simultaneously stream training data from the same storage pool without performance degradation, while the system handles the complex data management tasks like versioning, lineage tracking, and multi-site replication automatically.

Business Model

VAST Data operates as a B2B infrastructure software company that sells unified data platform licenses to large enterprises and cloud service providers. The company monetizes through annual software subscriptions that include the full stack of storage, database, and compute orchestration capabilities, typically priced based on raw storage capacity deployed.

The business model centers on displacing multiple point solutions with a single platform purchase. Instead of customers buying separate storage arrays, data warehouse licenses, ETL tools, and compute clusters, VAST provides all these capabilities in one integrated system. This consolidation approach allows VAST to capture significantly more wallet share per customer while reducing the customer's total cost of ownership and operational complexity.

VAST's go-to-market strategy focuses on large enterprise accounts and cloud infrastructure providers that need to process massive datasets for AI workloads. The company targets customers with petabyte-scale data requirements where the performance and operational benefits of consolidation justify premium pricing. Average deal sizes exceed \$1 million annually, with the largest customers committing over \$100 million in total platform spend.

The company's cost structure benefits from software-centric margins while leveraging commodity flash storage hardware. VAST develops the software stack in-house but partners with hardware manufacturers and system integrators for deployment, creating an asset-light model that scales efficiently. The platform's ability to run multiple workload types on the same infrastructure creates operational leverage, as customers expand usage across different teams and use cases without requiring separate infrastructure investments.

Competition

AI-first storage specialists

WEKA represents VAST's most direct competitor in the high-performance AI storage market. WEKA focuses on extreme throughput optimization for GPU workloads, offering software-only deployments that can deliver faster token streaming for large language model training. WEKA has demonstrated significant performance advantages in specific benchmarks, particularly for KV-cache prefill operations. However, WEKA lacks the integrated database layer that VAST provides, requiring customers to still manage separate systems for data cataloging and SQL analytics.

DDN targets the high-performance computing market with turnkey AI appliances, leveraging strong relationships with supercomputing centers and research institutions. DDN's strength lies in pre-integrated hardware-software bundles that simplify procurement for large-scale deployments, but this approach limits flexibility compared to VAST's software-defined architecture.

Enterprise storage incumbents

Pure Storage has aggressively moved into AI infrastructure with FlashBlade systems certified for NVIDIA's AI Data Platform. Pure leverages existing enterprise relationships and channel partnerships to compete for AI workloads, but lacks the integrated compute and database capabilities that VAST provides natively. Pure's premium pricing model also creates opportunities for VAST to win on total cost of ownership arguments.

Dell Technologies combines PowerScale storage with broader infrastructure offerings, using their massive installed base and services organization to compete for large AI deployments. Dell's advantage lies in their ability to provide complete data center solutions, but their legacy storage architecture creates performance bottlenecks compared to VAST's purpose-built design for parallel GPU access patterns.

NetApp and IBM compete primarily in hybrid cloud scenarios where enterprises want to extend existing storage investments into AI workloads. These incumbents face architectural limitations from legacy file systems that weren't designed for the parallel, high-throughput access patterns that modern AI training requires.

Cloud-native analytics platforms

Databricks and Snowflake represent a different competitive threat as they move down-stack to manage unstructured data directly. These platforms could potentially abstract away storage vendors entirely by providing data management capabilities as part of their analytics offerings. However, their cloud-centric architectures create challenges for customers who need on-premises or hybrid deployments for data sovereignty or latency requirements.

TAM Expansion

Full-stack AI operating system

VAST has evolved from a high-performance storage array into a complete data-compute platform that handles storage, database, and serverless execution in a single system. This transformation expands VAST's addressable market from the \$20 billion enterprise storage segment into the combined \$90 billion market encompassing data warehousing, analytics platforms, and serverless compute infrastructure.

The integration with NVIDIA BlueField DPUs positions VAST as the data layer for GPU-centric AI factories, where every server becomes a self-contained processing unit. This architecture enables VAST to compete directly with cloud platforms like Databricks and Snowflake while offering on-premises deployment options that cloud providers cannot match.

Cloud and hyperscaler partnerships

VAST's standardization as the storage platform for CoreWeave's global GPU cloud creates a distribution channel to hundreds of AI startups that consume CoreWeave capacity. Similar partnerships with G42 Cloud and HPE expand VAST's reach beyond traditional enterprise sales into cloud service provider and sovereign cloud markets.

The company's OEM strategy with hardware vendors could replicate the hyperconverged infrastructure playbook that made Nutanix successful, allowing VAST to ride other vendors' sales channels and field organizations to reach mid-market customers who wouldn't typically buy storage platforms directly.

Geographic and hybrid cloud expansion

VAST DataSpace extends the platform into AWS, Azure, and Google Cloud environments, creating unified global namespaces that span on-premises and public cloud infrastructure. This hybrid capability opens up international markets where enterprises prefer cloud-first architectures but need on-premises components for data sovereignty or latency requirements.

The ability to provide consistent data services across multiple regions and cloud providers addresses the needs of multinational enterprises that want to standardize on a single data platform globally rather than managing different solutions in each geography.

Risks

Commoditization pressure: As NVIDIA's AI Data Platform creates reference architectures for AI infrastructure, storage vendors face increasing commoditization as hardware differentiation diminishes. If the market shifts toward standardized configurations where storage becomes an interchangeable component, VAST's premium pricing model could face pressure from lower-cost alternatives that provide similar performance within NVIDIA's certified ecosystem.

Flash memory supply: VAST's all-flash architecture creates dependency on QLC NAND flash memory supply chains that are subject to cyclical pricing and availability constraints. Significant increases in flash memory costs or supply shortages could compress VAST's margins or force the company to pass costs through to customers, potentially making their solutions less competitive against hybrid storage approaches that use cheaper disk-based tiers.

Cloud platform competition: Major cloud providers like AWS, Google Cloud, and Microsoft Azure are rapidly expanding their managed AI infrastructure services, potentially reducing demand for on-premises solutions like VAST's platform. If enterprises increasingly prefer consuming AI infrastructure as managed cloud services rather than deploying their own hardware, VAST's total addressable market could shrink significantly, particularly among mid-market customers who drive volume growth.

DISCLAIMERS

This report is for information purposes only and is not to be used or considered as an offer or the solicitation of an offer to sell or to buy or subscribe for securities or other financial instruments. Nothing in this report constitutes investment, legal, accounting or tax advice or a representation that any investment or strategy is suitable or appropriate to your individual circumstances or otherwise constitutes a personal trade recommendation to you.

This research report has been prepared solely by Sacra and should not be considered a product of any person or entity that makes such report available, if any.

Information and opinions presented in the sections of the report were obtained or derived from sources Sacra believes are reliable, but Sacra makes no representation as to their accuracy or completeness. Past performance should not be taken as an indication or guarantee of future performance, and no representation or warranty, express or implied, is made regarding future performance. Information, opinions and estimates contained in this report reflect a determination at its original date of publication by Sacra and are subject to change without notice.

Sacra accepts no liability for loss arising from the use of the material presented in this report, except that this exclusion of liability does not apply to the extent that liability arises under specific statutes or regulations applicable to Sacra. Sacra may have issued, and may in the future issue, other reports that are inconsistent with, and reach different conclusions from, the information presented in this report. Those reports reflect different assumptions, views and analytical methods of the analysts who prepared them and Sacra is under no obligation to ensure that such other reports are brought to the attention of any recipient of this report.

All rights reserved. All material presented in this report, unless specifically indicated otherwise is under copyright to Sacra. Sacra reserves any and all intellectual property rights in the report. All trademarks, service marks and logos used in this report are trademarks or service marks or registered trademarks or service marks of Sacra. Any modification, copying, displaying, distributing, transmitting, publishing, licensing, creating derivative works from, or selling any report is strictly prohibited. None of the material, nor its content, nor any copy of it, may be altered in any way, transmitted to, copied or distributed to any other party, without the prior express written permission of Sacra. Any unauthorized duplication, redistribution or disclosure of this report will result in prosecution.

Published on Aug 04th, 2025