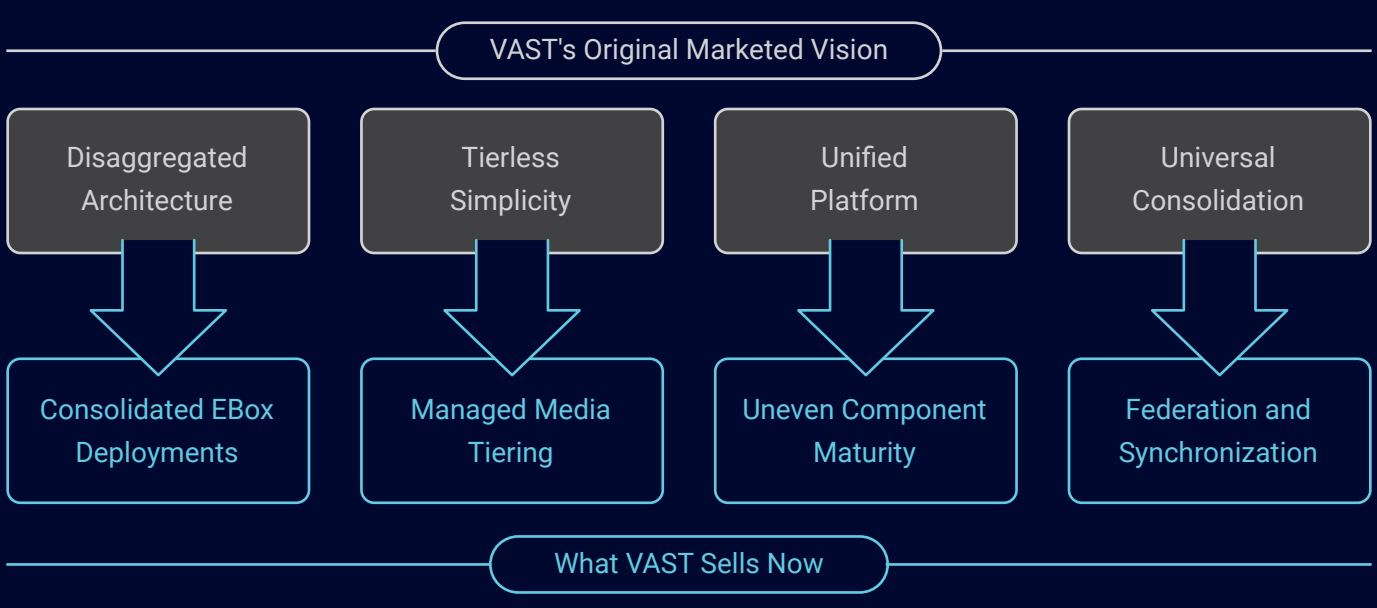
PROWESS

Production AI Must Adapt to Enterprise Reality

VAST Data entered the market with a vision: a disaggregated, tierless, software-defined platform designed to simplify enterprise storage and AI infrastructure. Once production deployments began, they exposed requirements that every enterprise AI data platform must eventually address: requirements around which Dell AI Data Platform was engineered from the start.

Enterprise Reality Forced VAST to Pivot

Enterprise realities eventually force product pivots, and VAST is no exception. Customer requirements have forced VAST to change its deployment form factors, media behavior, platform components, and distributed data services as the architecture met real production environments.



VAST's pivots are not a graceful evolution but a series of course corrections toward enterprise production requirements that it once dismissed.

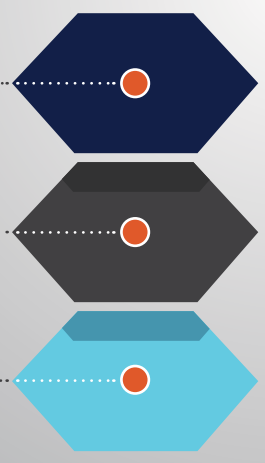
MODULARITY | FEDERATED DATA ACCESS | MEDIA AND LIFECYCLE FLEXIBILITY | CYBER RESILIENCE | OPERATIONAL CLARITY

Dell AI Data Platform Was Designed Around These Enterprise Realities From the Start

Storage Engines
File • Object • Block
Parallel File System

Data Engines
Processing • Analytics
Search

Data Orchestration
Data Management
Human Feedback
Pipeline Automation



Architecture, Backed by Results

Dell storage engines deliver better time-to-first-token (TTFT) performance, up to:

19x better on NVIDIA® H100 GPUs[1]

13.6x better on NVIDIA GB300 GPUs[2]

Production AI platforms must absorb change. The course corrections behind VAST® AI OS show how real-world deployments forced VAST's hand. For established enterprises with distributed data, governance obligations, and long infrastructure lifecycles, Dell AI Data Platform offers a clearer path forward.

Read the full report to compare Dell AI Data Platform to VAST AI OS.

Endnotes

[1] TTFT with key-value (KV) cache retrieved from Dell PowerScale™ and Dell ObjectScale™, compared to standard vLLM recomputing the KV cache on the GPU in Dell testing. Source: David Noy. Dell Technologies. "**Dell Storage Engines: Accelerating AI inferencing with PowerScale and ObjectScale**." October 2025.

[2] Mean TTFT with KV cache retrieved from PowerScale, compared to recomputing the KV cache on the GPU in Dell testing. Source: Ugur Kaynar. Dell Technologies. "**Faster GPUs and Smaller Caches Raise the Bar for Storage**." September 2026.

The analysis in this document was done by Prowess Consulting and commissioned by Dell Technologies.
Claims have been verified against publicly available information and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.
Prowess Consulting and the Prowess logo are trademarks of Prowess Consulting, LLC.
Copyright © 2026 Prowess Consulting, LLC. All rights reserved. Other trademarks are the property of their respective owners.

1126/260087