

Technical Research Study

PROWESS

When the Architecture Drifts from the Vision: Evaluating VAST Data's Platform Claims Against Enterprise Data Realities

What VAST Data's platform change reveals about enterprise AI requirements, and why Dell AI Data Platform fits established data estates.

Executive Summary

This paper evaluates how VAST® AI OS has evolved beyond its original architectural story and why Dell AI Data Platform offers established enterprises a clearer path for production AI. VAST Data launched with a vision of a disaggregated, tierless, software-defined alternative to conventional enterprise storage.[1] In practice, supported deployments remain bounded by specified hardware, minimum configurations, qualification requirements, and support constraints. Furthermore, its current platform is broader and less dogmatic in adhering to that original vision. This changing platform includes consolidated EBox deployments, rack-aware resilience, defined media behavior, published limits on hardware-generation compatibility, and other data and compute services that stray from its original vision.

Existing customers might have to bear the cost of adapting deployments, operations, and support processes to these changes, while newer components require separate maturity and support validation.

This pattern of evolution matters because each response points to the same conclusion: production AI platforms cannot escape enterprise requirements regardless of the vision behind them. They must define failure domains and support envelopes, data-movement behavior, lifecycle commitments, governance controls, and recovery paths. A clean architecture diagram can postpone these questions, but it can't retire them.

TOP FINDING

VAST's platform evolution demonstrates that the deployment, media, lifecycle, and data-management realities of enterprise environments did not disappear; they ultimately had to be addressed. While VAST has added capabilities to accommodate those requirements, the changes also reflect a meaningful departure from its original architectural narrative and introduce uncertainty for organizations that adopted the platform based on that earlier vision. By contrast, Dell AI Data Platform was designed around these enterprise operational requirements from the outset, providing a more established and modular foundation for evolving data estates.

This paper treats VAST as a case study in architectural convergence. For a greenfield deployment or a new AI environment with little incumbent data, consolidation can reduce integration work and deliver a consistent operating surface. For an established enterprise, however, the same consolidation-first logic can turn data movement, migration sequencing, and platform dependence into recurring costs.

Dell AI Data Platform (part of the broader Dell AI Factory) provides a useful comparison because it makes a different set of choices explicit.[2] It separates storage engines, data engines, and orchestration so that platform layers can evolve at different rates. It also builds on storage systems with documented mixed-generation operation, media choice, and cyber-resilience controls. Meanwhile, VAST AI OS has expanded well beyond the architectural story that launched the platform. This study examines what that evolution means for established enterprises, and it compares VAST's current Disaggregated Shared Everything (DASE) and EBox deployment models with Dell AI Data Platform.

Why Established Enterprises Need Flexible AI Storage

AI infrastructure discussions start with accelerators, but many established organizations encounter a harder constraint one layer below those accelerators. These organizations' useful data sits across file systems, object stores, databases, cloud services, edge locations, and archives. The data has owners, retention rules, application dependencies, and recovery requirements. Building an effective production AI service means reaching that distributed data without stripping away the controls that made it usable in the first place.

Enterprise AI draws on multimodal data across file and object storage, clouds, data centers, and remote sites. The strongest architecture isn't simply the one that runs the first workload fastest. Rather, it lets teams start small, connect existing data, and scale infrastructure without creating new silos or forcing every layer to change at once.

Security also belongs in the architectural evaluation from the start. Proprietary and regulated data moves through retrieval, training, inference, and agentic workflows, often crossing systems and organizational boundaries along the way. The National Institute of Standards and Technology (NIST)'s **AI Risk Management Framework** emphasizes governance across the AI lifecycle, while the NIST Cybersecurity Framework treats govern, protect, detect, respond, and recover as concurrent responsibilities.[3]

Prowess Consulting's April 2026 **technical research brief** comparing Dell AI Data Platform and VAST AI OS surfaced a pattern that warranted deeper examination: the gap between VAST's foundational architectural narrative and its current shipping platform.[4] VAST's drift from its original vision is pronounced because its original proposition was starkly clear. The 2019 launch promised an end to hard-disk drives (HDDs) and conventional storage tiering through Universal Storage and the DASE architecture. VAST did deliver on the HDD portion of that vision and automated data placement within the platform rather than exposing traditional tiers to administrators. However, its architecture still relies on different media classes using storage-class memory (SCM) for writes and quad-level cell (QLC) NAND for capacity. Those distinctions remain relevant because the media differ in cost, endurance characteristics, and supply dynamics.

The current VAST AI OS still rests on DASE, but it now describes a much wider operating system (OS) that spans storage, database, and data-processing functions whose components have reached production at different times.[5] It also documents the concrete limits and operating models that appear as a platform meets real customers, including deployment options, media behavior, resilience mechanisms, and lifecycle boundaries.

Dell AI Data Platform offers a useful comparison because it makes different choices explicit by separating storage engines, data engines, and orchestration so platform layers can evolve at different rates. Its architecture also provides an enterprise-mainstream reference point for contrasting and comparing VAST's evolution.

Prowess Consulting based the analysis in this study on public documentation, with no lab testing performed or implied. Performance numbers appear only where vendors have published methods and results, and those numbers remain secondary to architecture, lifecycle, and operating fit. This analysis gives IT leaders a transparent basis for comparing how each platform's topology affects fault isolation and scaling, how resilience and media choices shape capacity and cost, and how support boundaries, component maturity, data movement, and lifecycle limits affect deployment risk and long-term planning (see Table 1).

Table 1 | Four enterprise requirements where VAST's original story and what it ships today no longer align—and why the drift can be challenging for customers

Enterprise Requirement	VAST's Original Proposition	What VAST Sells Now
Deployment flexibility	VAST positioned disaggregated CNodes and highly available DBoxes as a way to scale compute and capacity independently.	VAST now also offers EBox, which colocates compute and storage and adds additional deployment options.
Why VAST's current model can be problematic for organizations	VAST's move from a predominantly disaggregated model to multiple deployment models can shift complexity to the customer. DASE and EBox may require different scaling increments, network designs, failure domains, qualification rules, and support paths.	
Media economics	VAST positioned Universal Storage as a tierless, all-flash alternative that abstracted placement decisions from administrators.	VAST's architecture still relies on distinct media classes, including SCM for writes and QLC for capacity, with movement managed within the platform.
Why VAST's current model can be problematic for organizations	VAST's tierless positioning can obscure the loss of a practical workload-placement lever. Tiering allows customers to place hot data on faster media while using more economical media for warm, cold, or archival workloads. VAST's all-flash model may force premium media economics across workloads that do not require premium performance, while customers still absorb endurance, replacement, supply, and lifecycle exposure. (Moreover, VAST's Amplify program for extracting more value from existing customer solid-state drive [SSD] investments implicitly acknowledges that storage economics and workload-specific tradeoffs still matter, which undercuts its longstanding claim that all-flash storage makes traditional tiering economically obsolete.)	
Component-level maturity	VAST presented one unified platform story, implying a consistent operating model across the platform.	VAST describes its DataStore as long-shipping but VAST DataBase, DataEngine, streaming, message-broker, and agentic functions have different release histories and current limits.
Why VAST's current model can be problematic for organizations	VAST AI OS requires customers to operate components with uneven maturity, upgrade dependencies, production references, and support boundaries.	
Distributed data operations	VAST positioned consolidation into one universal platform as a way to simplify the data environment.	VAST DataSpace federates VAST clusters, while VAST SyncEngine indexes and synchronizes external sources into the VAST environment.
Why VAST's current model can be problematic for organizations	Customers may have to manage copies, indexes, synchronization frequency, freshness, authoritative sources, bandwidth, governance, permissions, failure handling, and recovery. Avoiding one migration can become an ongoing synchronization program.	

VAST's current platform should be evaluated not only as an expansion beyond core file and object storage, but also as a potential change in what organizations must purchase, operate, govern, refresh, and eventually migrate. These ongoing additions are more pragmatic responses than graceful maturation; they are indications of a platform still working out its architectural strategy in production and retrofitting the deployment flexibility, federated data operations, explicit support boundaries, resilience mechanisms, and managed data-placement models it once dismissed. Each addition concedes a principle that enterprise-grade platforms adopted years ago, and architecture still in this kind of flux is a risk for buyers to underwrite rather than a feature to celebrate. Dell AI Data Platform provides a useful contrast because it illustrates how enterprise-grade platforms adopted these principles from the start. Dell AI Data Platform did not have to reverse course as it was engineered around these customer operational requirements from the start and has long rewarded enterprises with proven reliability rather than architectural purity.

Deployment Reality and the Limits of Pure Disaggregation

VAST's original DASE design separates stateless compute nodes from persistent storage enclosures.[5] That separation lets customers scale compute and capacity independently and gives every CNode access to a shared system state. However, this independence still depends on validated components, network design, and the scaling boundaries of the selected DASE configuration. It is a design that can work in environments whose architects are willing to build the fabric and operate the validated components around it.

The introduction of the EBox broadens that model beyond VAST's original vision.[6] According to VAST's product documentation, the design serves hyperscalers, cloud providers, and OEM channels that could not or did not want to deploy highly available DBoxes.[7] EBox places compute and storage functions on the same server, while metadata triplication and mixed CNode/EBox clusters preserve VAST's software model (though it is unclear to what extent VAST tests interoperability, performance, or support implications across mixed clusters of this changing architecture). The change shows that deployment flexibility and a validated server footprint can matter as much as the ability to scale every resource independently and brings VAST's architecture closer to the mainstream that its original vision eschewed.

Enterprise customers rarely replace every platform layer at once. Servers, storage media, data services, and orchestration follow different budget and refresh cycles. EBox is therefore more than a form-factor change, reflecting a wider preference for modular adoption paths over architectural purity. That modular, multi-layer assembly also raises a question the architecture itself doesn't answer: who stays accountable for the full stack as each layer is sourced, validated, and refreshed on its own schedule and keeps running long after Day 1?

Unified Architecture and Accountability Are What Survive Day 2

Day 1 proves a platform can be installed; Day 2 is everything after: operating, scaling, and refreshing it. Alongside performance and scalability, organizations should assess failure domains, support ownership, and the degree to which the solution depends on technologies and roadmaps outside a single provider's control. Dell AI Data Platform is built as an end-to-end portfolio with unified accountability, whereas architectures spanning multiple hardware vendors may increase coordination requirements and introduce uncertainty as underlying technologies evolve. Over time, those architectural dependencies can have a meaningful impact on operational flexibility, support experience, and long-term planning.

Availability and Media Economics Under Production Conditions

Real data centers fail in groups. A breaker trips, a rack loses cooling, a chassis goes dark, or a maintenance action removes more hardware than a device-level model assumed. VAST's recent DBox-HA and rack-level-resilience work responds directly to those conditions. The company documents how broader protection consumes capacity and how changes rebuild behavior rather than pretending that resilience is free.[8] That work quietly reintroduces the failure-domain planning that VAST's original resilience story implied it had engineered away: resilience was never free, only deferred. The same discipline applies to media.

VAST launched Universal Storage as a tierless alternative, yet its architecture has always used two distinct flash classes: SCM as a write buffer and QLC NAND for capacity. Current documentation explains how CNodes move data from the SCM buffer to QLC media.[9] The system hides that movement from applications, but the media classes still differ in endurance, cost, availability, and replacement risk. An organization's task is to model the mechanism honestly. Resilience overhead, write amplification, rebuild time, and flash supply exposure exist whether the product calls them tiers, buffers, pools, or data-placement policies.

Lifecycle Flexibility Becomes a Written Support Boundary

Long-lived infrastructure decisions rest on contracts and compatibility matrices. VAST's Gemini terms market an "Infinite Storage Lifecycle" in which data is "never bound to hardware generations" and customers "never again need to migrate."[10] The written commitment is far narrower: VAST states it will support any VAST-branded appliance (CBox or DBox) for up to 10 years. That gap between an "infinite," "never migrate again" slogan and an up-to-10-year written per-appliance support term is what buyers making 10-plus-year decisions should weigh, not the marketing language.

The same discipline should apply at the component level. VAST DataStore was launched in 2023.[11] DataBase, DataEngine, event-streaming, and agent-oriented functions entered the platform later and continue to acquire release-specific capabilities and limits.[12] The components share the VAST AI OS platform and firmware foundation, but organizations should still validate the production history, release requirements, and limits of the exact services they plan to deploy.

VAST's Gemini program lets customers purchase a term based on duration, capacity, and compute cores. Organizations should still clarify what happens when the purchased term ends or requirements exceed its licensed capacity and compute entitlements. Ask about renewal terms, support for older hardware, data portability, and responsibility for any eventual migration.

Why the Pattern of Evolution Matters More Than Any Single Change

The issue is the direction of the changes that VAST felt it had to make. That direction reveals the requirements that recur when AI infrastructure moves from a clean-sheet design into production: deployment flexibility, correlated-failure planning, explicit media behavior, lifecycle limits, validated configurations, component-level maturity, and distributed data operations. For buyers, an architecture still assembling these fundamentals in production is a risk to underwrite, not a milestone to celebrate.

Dell AI Data Platform is a useful comparison because it addresses many of those requirements through a different architectural boundary. Dell Technologies separates data services, storage services, and orchestration rather than asking one storage-centered platform to mature every function together. The comparison is not innovation versus convention; it is two ways of absorbing the same enterprise constraints.

Why Enterprise Data Estates Change the Evaluation Criteria

A data estate is the full distributed body of an organization's data across warehouses, lakes, object stores, file systems, edge sites, cloud services, and legacy platforms. Established enterprises start with this estate, plus the policies and applications attached to it. These enterprises need a platform that can discover and govern data where it resides, only moving or synchronizing data when a workload requires. Organizations should determine whether each service catalogs source data in place, creates another copy, or depends on continued synchronization.

Greenfield AI clouds can treat consolidation as a feature but an established enterprise has less freedom. Moving petabytes of data can take months, consume network and staff capacity, duplicate governance work, and create synchronization pipelines that persist after the initial migration. VAST's current DataSpace and SyncEngine seek to address this reality.[13,14] DataSpace federates VAST clusters across locations, while SyncEngine indexes and synchronizes external sources into the VAST environment. That is a pragmatic expansion of the original consolidation story, but it is nevertheless a drift from VAST's original vision and it requires organizations to examine what moves, what remains authoritative, and how many copies must be governed.

Federated access changes the cost model. A query engine that reaches existing databases, lakes, and object stores can reduce wholesale relocation. However, it trades data-movement cost for catalog, policy, and performance complexity. Dell's **Data Analytics Engine**, based on Starburst® and Trino, is positioned to query distributed data sources.[15] That approach can preserve existing systems longer, although it does not erase latency differences or inconsistent metadata.

Federation Is Not Free

Federation can help organizations avoid large-scale data migrations, but it is not free. The resulting costs may include increased platform dependence, reduced flexibility in technology selection, and closer alignment with a single vendor's roadmap and release cycle. Dell AI Data Platform was designed to balance end-to-end platform integration with customer choice, enabling organizations to leverage proven best-of-breed technologies and adopt future innovations without being tied to a single architectural path.

Supply Cycles Turn Media Choice into an Operating Requirement

The 2026 NAND supply cycle offers a useful example of why enterprise architecture extends beyond benchmark performance. **TrendForce®** projected NAND contract price increases of 70 to 75% quarter-over-quarter in the second quarter of 2026 as AI and data-center demand tightened supply.[16] Its **July update** estimated a 4 to 5% NAND supply deficit for 2026, with the market returning to balance only in the second half of 2027.[17]

An architecture almost exclusively built around QLC exposes organizations more directly to the availability and price of that media class. A platform that supports all-flash, hybrid, and archive capacity offers more procurement options when the market moves. **Media flexibility** does not guarantee lower cost; workloads still need the right performance tier.[18] However, flexibility with storage media gives operators room to maneuver when flash prices spike, instead of forcing a choice between delaying the project or accepting the quoted flash price.

Lifecycle horizons compound the issue of relying exclusively on all-flash storage. Regulated, government, and research environments often keep data longer than the first hardware generation that stores it. Organizations should compare a vendor's written generation limits with their own retention, expansion, and decommissioning plans. The relevant question is not whether a platform can live forever, but whether the next transition is documented, supported, and budgeted.

The emergence of programs such as VAST Amplify underscores that the data estate is ultimately a risk-management concern as much as a storage-economics concern. Reusing existing flash investments may reduce acquisition costs, but it can also increase exposure to hardware aging, support limitations, drive failures, and data-protection challenges. Organizations should evaluate whether those tradeoffs align with the operational and resilience requirements of the data they are entrusting to the platform.

Architecture determines more than whether the first workload runs. It determines how the organization connects the next data source, expands capacity, survives a supply constraint, and refreshes hardware without rebuilding the operating model.

What Happens After the First Workload Goes Live?

The first workload proves that a platform can run. The next workload tests whether the architecture can grow without forcing the organization to redesign its operating model.

From Day 2 to Day 365, teams add data sources, expand capacity, introduce new services, and prepare for hardware refreshes. Each step exposes details that a product vision can obscure. Organizations must understand not only what VAST AI OS promises as a unified platform, but also what the deployed configuration requires over time.

Expansion Changes the Cost Model

VAST has two deployment models. DASE separates CNodes from DBoxes so organizations can scale compute and storage independently. EBox packages compute and storage in a more integrated configuration. Both are VAST offerings, but they create different scaling increments, network requirements, and failure domains. Offering both models is itself a concession: VAST is still converging on the balance between disaggregated and converged designs that Dell has engineered for years, and an architecture that is still deciding its own model can leave buyers to absorb the uncertainty.

Those differences matter after the initial purchase. An organization that adds capacity or performance must determine which components to expand, whether the existing network can support the change, and how the new configuration affects resilience. The first workload rarely reveals these costs.

A Unified Vision Can Still Require Data Movement

VAST AI OS brings storage, database, data processing, and orchestration capabilities under one platform vision. Production environments, however, rarely begin with all data inside one system. Critical data might remain on existing file and object storage, in clouds, or at remote sites.

DataSpace can federate VAST clusters, while SyncEngine can index and synchronize external data into the VAST environment. These capabilities help VAST work across a distributed estate, but they don't eliminate the work. Teams still need to decide which data moves, which copy remains authoritative, how often synchronization occurs, and where governance policies apply. Those decisions can add network consumption, duplicate capacity, and operational overhead as the environment grows.

Lifecycle Terms Become Operational Decisions

Organizations should ask what happens when licensed capacity or compute requirements change, when hardware reaches its supported horizon, or when supply conditions affect the preferred media. A platform centered on QLC flash can deliver density, but it gives infrastructure teams fewer media alternatives than a portfolio that also supports other flash types, hard-disk storage, cloud resources, and software-defined deployment.

Architecture Must Absorb Change

The first workload tests performance. Day 2 and every day after it tests whether the architecture can absorb change without compounding cost or constraining choice. VAST's current offerings expose a gap between a unified platform vision and the operational decisions customers must continue to manage. Enterprise reality calls for an architecture designed around those decisions from the start.

A More Credible Architecture for Enterprise Reality

Dell AI Data Platform is an engine-based architecture.[2] Storage engines hold and protect data; data engines process, search, and analyze it; and the Dell Data Orchestration Engine coordinates workflows across those engines. Dell describes this separation as deliberate, and it matters because enterprise AI environments rarely refresh every layer on the same schedule.

Platform modularity is broader than infrastructure modularity. VAST's DASE design lets compute and capacity scale independently within the VAST system. Dell's model extends the boundary upward. An organization can change a data-processing or search engine without necessarily replacing the storage layer, and a distributed SQL engine can reach data sources beyond a single storage system. The **orchestration layer** in Dell AI Data Platform is meant to coordinate those engines across on-premises, hybrid, and software-as-a-service (SaaS) deployment models.[19]

That modularity supports incremental adoption. An organization can begin with a mature storage engine, add a search or analytics capability for a defined workload, and expand orchestration as operating experience grows. This approach also creates integration seams. Dell's engines share governance, metadata, support escalation, and lifecycle management and function as a platform rather than a catalog.

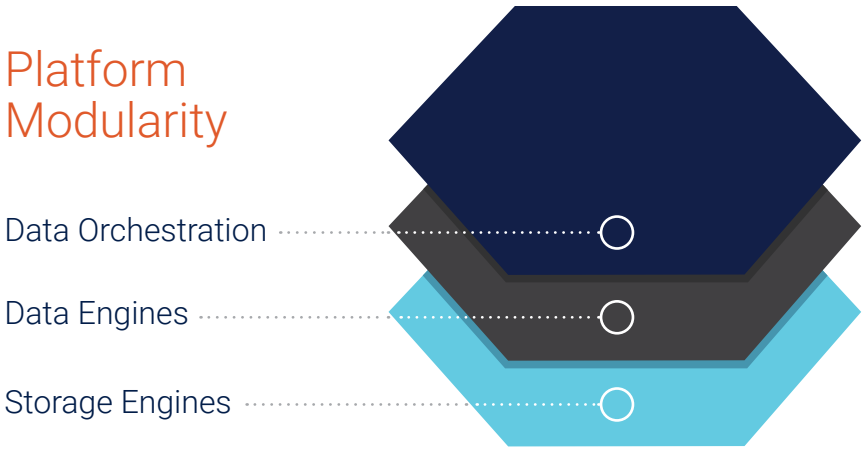


Figure 1 | Platform modularity means that different layers have the flexibility to evolve and refresh on their own cadence

Security and Recoverability as First-Order Design Criteria

Dell's enterprise advantage is clearest where AI data operations meet existing security and recovery practices. The **documentation** for the storage engine in Dell AI Data Platform describes ransomware detection, auditing, snapshots, air-gapped vault patterns, and recovery workflows.[20] The broader Dell AI Data Platform architecture positions cyber resiliency and governance alongside storage and data transformation rather than outside the platform boundary.

The contrast with VAST is evaluating how complete the operating model is today: which controls apply across file, object, table, stream, and function execution; which events enter the enterprise security stack; and which recovery point can restore an AI dataset with its metadata and lineage intact?

Lifecycle and Media Flexibility Are Verifiable Behaviors

Dell AI storage has its own architectural tradeoffs, including narrower erasure-code stripes and Dell hardware requirements. Its value as a reference point comes from long-shipping behaviors that customers can verify. For example, OneFS™ documents **node-class compatibility** that allows newer-generation nodes in pools with earlier generations.[21]

Dell **SmartPools™** provides a similarly explicit media model.[22] **OneFS** can present NVM Express® (NVMe®), SSD, software-attached storage (SAS), and capacity-oriented disk pools through one namespace, and it can move data according to policy.[22] Administrators can see the tiering decision and adjust it. That visibility carries management overhead, but it also lets organizations model the media mix against workload and supply conditions.

Performance Evidence Should Support, Not Replace, the Architecture Review

In testing described in this **Dell blog post**, Dell storage engines delivered **up to 19x better time to first token (TTFT) performance** than a standard vLLM serving engine without KV cache offloading on NVIDIA H100 GPUs.[23] Separate testing covered in another **Dell blog post** achieved **up to 13.6x better TTFT performance** on NVIDIA GB300 GPUs.[24] The results show that Dell's storage engines can participate directly in AI inference optimization, not merely hold model data.

These tests are still vendor-run and workload-specific. Organizations should inspect model size, context length, cache-hit behavior, network, client configuration, software versions, and concurrency before generalizing. A single inference result cannot establish platform maturity, database performance, recovery behavior, or total cost.

Dell AI Data Platform Carries Its Own Maturity Gradient

Dell's comparison advantage should not become a claim that every Dell AI Data Platform component has the production depth of Dell AI storage. The Data Orchestration Engine, GPU-accelerated Data Analytics Engine, search integrations, and other AI-adjacent services are evolving quickly. Dell's **2026 announcements** describe new integrations and future expansion across the storage portfolio.[25,26] Organizations should apply the same standard used for VAST: evaluate the exact engine, version, deployment model, and written support boundary.

A modular platform also creates more interfaces between components. Changes in an engine, metadata service, or orchestration layer can shift responsibility between Dell, an ecosystem partner, and the customer. An established enterprise might accept that tradeoff because it preserves existing data systems and allows staged change.

Plan for Change, Not Just Day 1

VAST's current platform extends beyond its original disaggregated, tierless proposition. Its platform evolution reflects a series of course corrections toward enterprise production requirements that it once minimized. It also shows why organizations should evaluate the current platform rather than rely on its original architectural story. The company has adjusted deployment form factors, resilience models, media behavior, lifecycle commitments, and distributed-data services as customers have pushed the architecture into real environments. The risk lies in evaluating that 2026 platform through a 2019 vocabulary that no longer captures the full product or its boundaries.

For an established enterprise with distributed data, governance obligations, and long asset lifecycles, Dell AI Data Platform presents a more credible architecture. It makes modularity, media choice, federated access, lifecycle behavior, and cyber resilience visible as separate operating concerns. It also states where the platform is still evolving.

The durable buying standard is operational clarity. Look for the vendor that describes the product it can support, names the tradeoffs it has chosen, and gives the enterprise a clear path beyond the first successful AI workload.

For a deeper comparison of Dell AI Data Platform and VAST AI OS across data processing, retrieval-augmented generation (RAG), performance, and operational overhead, read the **Prowess Consulting report**.

To learn more about why modularity and federation matter for enterprise AI, visit the **Dell AI Data Platform page**.

FAQ

Q: Who is this paper for?

A: This paper is for enterprise IT, infrastructure, storage, and AI platform decision-makers, including chief data officers (CDOs), data scientists, and engineers at organizations with distributed data estates, governance obligations, and long lifecycle requirements. It also supports technical evaluators, enterprise architects, procurement teams, and business leaders comparing vendor claims, architecture tradeoffs, and operational fit for production AI deployments.

Q: What is the central argument?

A: VAST AI OS has evolved beyond the simple architectural story that first distinguished it. Those changes are practical responses to production requirements, but they also expose deployment, lifecycle, and data-movement tradeoffs that buyers should evaluate directly. Dell AI Data Platform offers established enterprises a stronger alternative because its modular architecture makes those tradeoffs more explicit.

Q: Why does Dell AI Data Platform compare favorably for established enterprises?

A: Dell AI Data Platform compares favorably because it separates storage engines, data engines, and orchestration, allowing different layers to evolve and refresh on different schedules. The paper says this modularity supports staged adoption, distributed data access, media flexibility, and clearer lifecycle planning. It also acknowledges integration seams and newer components that organizations should validate before production deployment.

Q: Should AI infrastructure use file storage or object storage?

A: AI environments often need both. Scale-out file storage can serve performance-sensitive workflows, while object storage can hold large, diverse data collections efficiently. The key question isn't whether a platform supports file and object data, but how consistently it manages governance, access controls, lifecycle policies, and operations across both.

Q: How can an organization start small and scale its AI data infrastructure?

A: Organizations can choose an architecture that can expand capacity and data services in stages. Modular platforms reduce the risk that a successful first workload will trigger a larger, premature infrastructure replacement.

Q: What should enterprises do when AI data is scattered across many systems?

A: Enterprises can connect and govern data before deciding to move it. A federated architecture can preserve existing systems while making selected data available to AI workflows.

Q: What should enterprises evaluate for agentic AI and RAG?

A: Enterprises should evaluate data discovery, permission enforcement, index freshness, retrieval latency, and production support. Model performance alone doesn't show whether the underlying data architecture can sustain a governed enterprise deployment.

Q: What makes an AI data platform suitable for multimodal data?

A: The platform should provide governed access to documents, images, video, application data, and machine-generated data across file and object interfaces. It should not require every format or workload to use the same storage model.

Endnotes

1 VAST. "VAST Data Unveils its Universal Storage Architecture to Bring an End to the Enterprise HDD and Storage Tiering Era." February 2019. www.vastdata.com/press-releases/news-universal-storage-architecture-to-bring-and-end-to-the-enterprise-hdd-and-storage-tiering-era.

2 Dell. "**Dell AI Data Platform**." Accessed September 2026.

3 National Institute of Standards and Technology. "**AI Risk Management Framework**." Accessed September 2026.

4 Prowess Consulting. "**Why Open, Modular AI Data Platforms Win Over Closed, Storage-Embedded AI Data Stacks**." Accessed September 2026.

5 VAST. "VAST AI Operating System." Accessed September 2026. www.vastdata.com/whitepaper.

6 VAST. "Introducing the EBox: Compact Power and Resilience for Today's Data Centers." December 2024. www.vastdata.com/blog/introducing-the-ebox.

7 VAST. "VAST 5.2 Brings A Plethora of Enhancements." October 2024. www.vastdata.com/blog/vast-data-introduces-version-5-2.

8 VAST. "Bringing Rack-Level Resilience to VAST - Part 1: Accommodating DBox Losses." June 2025. www.vastdata.com/blog/bringing-rack-level-resilience-to-vast-part-one.

9 VAST. "The VAST AI Operating System: How It Works." Accessed September 2026. www.vastdata.com/whitepaper#how-it-works.

10 VAST. "Infinite Storage Lifecycle." Accessed September 2026. www.vastdata.com/infinite-storage-lifecycle. And VAST. "Gemini: Terms of Service Overview." Accessed September 2026. https://assets.ctfassets.net/2f3meiv6rg5s/3mhzg81F50u6AS9VTpdJM6/0d26993638c9c36bd81a1a93d20d6470/SB-GeminiTermsofService_VAST-Data.pdf.

11 VAST. "Hardware Agnostic, Enterprise Ready." Accessed September 2026. www.vastdata.com/platform/supported-platforms.

12 VAST. "VAST DataEngine: The Compute Fabric for Real-Time, Governed Data Operations." Accessed September 2026. www.vastdata.com/blog/vast-dataengine-the-compute-fabric-for-real-time-governed-data-operations.

13 VAST. "The VAST AI Operating System: The VAST DataSpace." Accessed September 2026. www.vastdata.com/whitepaper#the-vast-dataspace.

14 VAST. "The VAST AI Operating System: What is the VAST AI Operating System?" Accessed September 2026. www.vastdata.com/whitepaper#what-is-the-vast-ai-operating-system.

15 Jain, Vrashank. Dell. "**Navigating the modern data landscape: the need for an all-in-one solution**." March 2024.

16 TrendForce. "**AI Server Demand to Drive Memory Contract Price Increases in 2Q26 as CSPs Secure Supply via Long-Term Agreements**." March 2026.

17 TrendForce. "**NAND Flash Supply Growth to Outpace Demand in 2027, Easing Supply Constraints in 2H27, Says TrendForce**." July 2026.

18 Dell. "**Leading the AI Gold Rush Amid Supply Pressures**." Accessed September 2026.

19 Upalekar, Sunil and Vijayakumar, Prasanna. Dell. "**From Hybrid Data to AI-Ready Workflows**." May 2026.

20 Dell. "**PowerScale Cybersecurity Suite**." Accessed September 2026.

21 Dell. "**PowerScale OneFS Supportability and Compatibility Guide**." Accessed September 2026.

22 Dell. "**OneFS architectural overview**." Accessed September 2026.

23 Noy, David. Dell. "**Dell Storage Engines: Accelerating AI inferencing with PowerScale and ObjectScale**." October 2025.

24 Kaynar, Ugur. Dell. "**Faster GPUs and Smaller Caches Raise the Bar for Storage**." September 2026.

25 Vigil, Travis. Dell. "**AI at Scale Starts with Your Data: Introducing the supercharged Dell AI Data Platform with NVIDIA**." March 2026.

26 Noy, David. Dell. "**Dell AI Data Platform with NVIDIA: Full Throttle AI**." May 2026.



Legal Notices and Disclaimers

The analysis in this document was done by Prowess Consulting and commissioned by Dell Technologies.
Prowess 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