



Research Report:

ADDRESSING ENTERPRISE AI DATA CHALLENGES AT SCALE WITH THE DELL AI DATA PLATFORM

STEVE MCDOWELL, CHIEF ANALYST

OCTOBER 2025

ADDRESSING ENTERPRISE AI DATA CHALLENGES AT SCALE

While enterprises rush to operationalize artificial intelligence, many struggle with the most fundamental ingredient: their data. The core challenge is not a lack of models or compute power. The issue lies in the sprawl of data silos across legacy applications, departmental tools, and hybrid cloud environments.

These silos undermine AI strategies. They delay innovation, complicate compliance, and expose organizations to security risk. Executives understand that AI is now a board-level strategy, not just an IT ambition. But they cannot deliver AI-driven growth, efficiency, and customer impact without a data foundation designed for the demands of modern AI workloads. This is the problem the Dell AI Data Platform sets out to solve.

The Dell AI Data Platform addresses these challenges with a purpose-built, modular approach that integrates file and object storage into a unified system optimized for AI. Key components include data engines that enable aggregation, in-place querying, and searching across both structured and unstructured data sets, as well as storage engines that provide the necessary performance, scale, and security across on-premises, edge, and cloud environments.

Dell's solution also features deep ecosystem alignment with NVIDIA and leading AI software providers.

NEW DEMANDS ON ENTERPRISE INFRASTRUCTURE

Artificial intelligence is driving a fundamental shift in the way enterprise infrastructure is designed, deployed, and operated. Unlike traditional workloads, AI pipelines demand a radically different approach to data management, one that can sustain extreme throughput, support unpredictable access patterns, and scale seamlessly across distributed environments.

Each phase of the AI lifecycle places distinct demands on infrastructure, requiring different performance characteristics and protocols to achieve optimal results:

- **Data Ingestion and Preparation** require high-throughput capabilities to handle streaming data from multiple sources simultaneously. Organizations must process everything from real-time sensor data to massive batch uploads of historical information, often requiring storage systems to maintain consistent performance across varying access patterns.
- **Training** demands exceptional bandwidth and parallelism to keep GPU clusters fully utilized. Multi-GPU training scenarios can require sustained throughput measured in hundreds of gigabytes per second, with any I/O bottleneck capable of idling expensive compute resources and dramatically extending training times.
- **Inference Operations** present different challenges, prioritizing ultra-low latency over raw throughput. Real-time AI applications may require storage fetches for model weights, key-value caches, and prompt data to complete within microseconds, meeting user experience requirements.
- **RAG and Fine-tuning Pipelines** combine elements of both training and inference workloads, requiring rapid access to vector databases and source documents while managing metadata-rich queries at scale. These workloads often exhibit unpredictable access patterns that challenge traditional caching strategies.

THE DELL AI DATA PLATFORM

The Dell AI Data Platform addresses the limitations of traditional approaches to storage when building an infrastructure ready for enterprise AI.




FIGURE 1: DELL AI DATA PLATFORM (DELL TECHNOLOGIES)

Dell starts with a set of modular “engines” that can be combined to address different enterprise AI requirements. Unlike traditional storage-centric architectures, Dell’s approach enables customers to create a storage infrastructure tailored to their specific environment’s needs.

The storage layer offers multiple AI-ready engines: unified file/object for hybrid workflows, S3 object storage for massive unstructured datasets, and parallel file systems optimized for GPU-intensive training.

Above storage sits Dell Data Engines, which include federated SQL through Starburst, large-scale processing through Apache Spark, unstructured data search, and upcoming capabilities for AI pipeline orchestration.

These components are integrated under a Dell-built Kubernetes control plane, simplifying lifecycle management and eliminating the need for executives to fund expensive in-house integration projects.

Cyber-resilience is embedded throughout. Access controls, encryption, and isolation features provide intrinsic protection, supported by Dell’s broader PowerProtect suite. In practice, this means enterprises can scale AI workloads with confidence that governance, compliance, and security requirements are fully aligned.

These capabilities are underpinned by specific Dell products and partner integrations that deliver enterprise-grade performance, scalability, and security:

- **Data Placement** capabilities enable the effective ingestion and storage of streaming data, large batch datasets, and complex multimodal information, handling the massive data volumes required by modern AI while maintaining the performance characteristics necessary to keep downstream processing systems fully utilized.
- **Data Processing** functionality supports advanced data preparation, metadata enrichment, and curation of datasets required for training, fine-tuning, and inference operations.
- **Data Protection** ensures that personally identifiable information and proprietary data remain secure, resilient, and compliant throughout the AI lifecycle. Dell’s approach provides comprehensive

data protection addressing both security and regulatory requirements without compromising the accessibility that AI workflows require.

Dell's approach provides several intrinsic benefits:

- **End-to-End Pipeline Integration:** AI data pipelines can be developed, secured, and optimized in a single ecosystem, avoiding the need to stitch together disparate storage, processing, and governance tools.
- **Built-In Cyber-Resilience:** Cyber protection and compliance enforcement are embedded at the storage and metadata level, minimizing risks associated with AI model leakage, hallucination, or data corruption.
- **Scalability without Complexity:** Independent scaling of compute, storage, and data pipelines allows customers to right-size their AI infrastructure without architectural rewrites.
- **AI-Ready by Design:** High-throughput I/O, RAG compatibility, metadata extraction, and AI-centric reference designs reduce time to value for model development.
- **AI Ecosystem Alignment:** Deep partnerships with NVIDIA, Intel, and leading AI software vendors ensure continuous alignment with the AI tooling landscape.

With its AI Data Platform, Dell delivers an orchestrated strategy to bring data gravity, GPU acceleration, and enterprise control into a unified solution stack. In contrast to cloud-native-only platforms like Snowflake or Databricks, Dell offers a hardware-optimized, AI-first approach that aligns with enterprise needs for security, performance, and control, whether on-premises or in a hybrid environment.

DATA ENGINES: UNIFIED ANALYTICS PLATFORM

Dell's Data Engines provides a comprehensive platform that combines the benefits of data lakes and data warehouses while addressing the specific requirements of AI and analytics workloads.

Dell's Data Engines are modular, best-of-breed software components that make enterprise data accessible, usable, and AI-ready. They sit above Dell's storage engines and provide specialized capabilities for different types of data challenges.

Instead of forcing organizations into a one-size-fits-all solution, Dell delivers composable engines that can be deployed individually or combined as needs evolve. Each engine is specialized for a different workload, but together they give enterprises a flexible, composable framework that can grow with their AI ambitions.

These capabilities remove the friction that has long slowed enterprise AI adoption:

- **Eliminate silos** by connecting data sources without costly migrations.
- **Accelerate insight** by turning raw data into curated products ready for analytics.
- **Enable GenAI** by unlocking unstructured and vector search capabilities.
- **Simplify operations** by automating complex, multi-step data pipelines.

DATA PROCESSING ENGINE (APACHE SPARK)

Raw data rarely comes in a form that's ready for AI. The Data Processing Engine, built on Apache Spark, is the workhorse that transforms that data. It supports both batch and real-time processing, enabling enterprises to handle a wide range of tasks, from traditional ETL jobs to streaming IoT feeds.

By enriching, tagging, and transforming raw inputs into structured, usable data products, this engine ensures that AI and analytics models are trained on reliable, high-quality data.

DATA ANALYTICS ENGINE (STARBURST)

This engine tackles one of the biggest enterprise challenges: data sprawl. Using Starburst technology enables teams to run SQL queries across multiple, diverse data sources, whether databases, data lakes, cloud storage, or legacy systems.

Instead of moving or duplicating data, the data analytics engine queries data in place, reducing cost and complexity. For enterprises migrating off Hadoop or consolidating reporting platforms, this capability accelerates time-to-value.

DATA SEARCH ENGINE (ELASTIC)

Today's most valuable enterprise data often comes in unstructured forms: documents, transcripts, log files, images, and more. Dell's Data Search Engine, powered by Elastic, delivers full-text and vector search across these sources.

This is essential for modern GenAI use cases, such as RAG, semantic search, and document intelligence. By making unstructured data searchable and accessible, this engine opens new possibilities for customer engagement, compliance, and knowledge management.

AI DATA PIPELINE ORCHESTRATION

AI workflows rarely involve just one step. Data must move through a sequence: ingestion, cleaning, labeling, processing, and model execution. The upcoming AI Data Pipeline Orchestration Engine will provide an end-to-end automation layer that connects these steps into a seamless flow. This will simplify operations, reduce manual intervention, and accelerate the deployment of AI models into production.

DELL POWERSCALE: SCALE-OUT FILE STORAGE FOR AI

Dell PowerScale delivers enterprise-grade, scale-out network-attached storage engineered to meet the most demanding AI and analytics requirements. Purpose-built for performance, scale, and security, PowerScale provides the data backbone modern enterprises need to unlock AI value at any stage of their journey.

- **Unmatched Performance at Scale:** PowerScale delivers linear performance scaling from terabytes to multiple petabytes, with each additional node contributing both capacity and throughput, ensuring that storage performance keeps pace without bottlenecks.
- **Multi-Generation, Unified Namespace:** PowerScale uniquely supports multiple generations of nodes in a single namespace, protecting existing investments while enabling organizations to scale capacity and performance without disruption. Enterprises can expand incrementally, leveraging both legacy and latest-generation PowerScale nodes within a single unified system.
- **Flexible, Multi-Protocol Access:** With simultaneous support for NFS, SMB, HDFS, and S3, PowerScale consolidates diverse workloads into a single, secure storage environment. AI frameworks, big data pipelines, and analytics tools all operate against the same datasets, eliminating silos and simplifying infrastructure management.

- **Secure and Resilient by Design:** PowerScale incorporates enterprise-grade security at every level, including encryption, fine-grained access controls, secure snapshots, and built-in ransomware protection.
- **Automated Data Management:** Intelligent tiering, policy-driven automation, and integrated backup and recovery capabilities reduce operational overhead while ensuring availability and resilience at scale. This allows enterprises to focus on deriving insights from data rather than managing storage complexity.

DELL OBJECTSCALE: HIGH-PERFORMANCE OBJECT STORAGE

Dell ObjectScale provides containerized object storage engineered explicitly for the most demanding AI and analytics workloads, delivering enterprise-grade performance and scalability:

- **Massive scaling capabilities:** support multi-petabyte deployments with global namespace functionality, enabling organizations to store and manage vast amounts of unstructured data across distributed environments. The platform scales seamlessly from single-site deployments to global installations.
- **S3 Compatibility and Performance:** industry-leading object storage performance with full S3 API compatibility, ensuring seamless integration with AI frameworks and tools that rely on object storage protocols. Advanced features include S3 over RDMA for ultra-high-performance scenarios.
- **AI-Optimized Architecture:** specialized optimizations for AI workloads, including enhanced metadata management, parallel data access patterns, and integration with popular AI development frameworks. The platform supports both training and inference workloads with consistent performance.
- **Enterprise Security and Governance:** comprehensive data protection, encryption, and compliance capabilities designed for enterprise environments. The platform provides granular access controls and audit capabilities, essential for regulated industries and sensitive AI applications.
- **Software-Defined Storage (SDS) Flexibility:** Dell ObjectScale can be deployed consistently across on-premises, edge, and hybrid cloud environments. This enables customers to adapt to changing workload demands without being locked into a single hardware configuration.
- **PowerEdge Integration:** Dell's industry-leading server platform delivers the compute performance, GPU acceleration, and reliability required for optimal ObjectScale deployments. Together, ObjectScale and PowerEdge deliver a tightly integrated, high-performance foundation for AI and analytics at scale.

DELL PROJECT LIGHTNING: ULTRA-HIGH PERFORMANCE PARALLEL FILE SYSTEM

Dell Project Lightning is Dell's breakthrough approach to addressing the most demanding AI infrastructure requirements, delivering the world's fastest parallel file system specifically designed for large-scale AI deployments:

- **Unprecedented Performance,** with Project Lightning delivering up to 2x greater throughput than competing parallel file systems, enabling organizations to support the most demanding AI training workloads involving tens of thousands of GPUs.
- **Massive Scale Architecture** is purpose-built for the largest AI deployments, supporting configurations that can serve thousands of concurrent GPU nodes while maintaining consistent

performance characteristics. The platform scales linearly as compute resources are added to AI clusters.

- **Advanced Integration Capabilities** include deep integration with leading AI frameworks and orchestration platforms, ensuring optimal data delivery patterns for diverse AI workloads. Project Lightning supports both training and inference workloads, with specialized optimizations tailored to each use case.
- **Future-Ready Design** goes beyond current demands by incorporating support for emerging AI technologies and deployment patterns, including agentic AI systems that require dynamic data access patterns and real-time processing capabilities. Project Lightning is designed with future-state memory semantics in mind, enabling seamless adoption of persistent memory and CXL-driven data models.
- **SDS Flexibility and Hardware Awareness** ensure that Project Lightning is not tied to proprietary appliances but delivered as software-defined storage. This enables enterprises to continually leverage advancements in hardware, particularly in network-bound environments where next-generation interconnects, SmartNICs, DPUs, and NVMe-oF will play a decisive role in sustaining AI performance at scale.

DELL & NVIDIA: BETTER TOGETHER

Integration with broader AI technology ecosystems is essential for enterprise AI success. Dell collaborates across the industry to ensure seamless integration with comprehensive enterprise software and hardware ecosystems.

One of Dell's most significant relationships is its deep, long-standing AI partnership with NVIDIA. The two companies have worked closely together to ensure seamless compatibility between the Dell AI Data Platform and NVIDIA's comprehensive enterprise software and hardware ecosystem. This strategic alignment creates a unified platform that maximizes the performance and efficiency of GPU-accelerated AI workloads.

The collaboration between Dell and NVIDIA extends beyond simple compatibility testing to include joint engineering efforts, shared roadmap development, and comprehensive validation programs that ensure optimal performance across the entire AI infrastructure stack.

The partnership enables enterprises to deploy AI solutions with confidence, knowing that their infrastructure has been specifically optimized for NVIDIA's leading AI technologies:

NVIDIA Platform Integration	<ul style="list-style-type: none"> • Comprehensive Platform Validation • DGX SuperPOD, NCP, and Enterprise Storage Certification • Performance Optimizations • Unified Support Experience
NVIDIA AI Data Platform Alignment	<ul style="list-style-type: none"> • NVIDIA NeMo Integration • RAPIDS Acceleration Support • Triton Inference Server Optimization
NVIDIA Hardware Portfolio	<ul style="list-style-type: none"> • NVIDIA GPUs • NVIDIA Mellanox DPUs • Spectrum X Ethernet • Connect X
NVIDIA Software Stack	<ul style="list-style-type: none"> • NVIDIA AI Enterprise Compatibility

	<ul style="list-style-type: none"> • Native support for NVIDIA Container Runtime & Kubernetes operations • MLOps Platform Support • Unified Management spanning Dell storage and NVIDIA compute resources
Workload Orchestration	<ul style="list-style-type: none"> • Intelligent Scheduling Coordination • Real-Time Performance Monitoring • Automated Optimization

This deep alignment positions Dell storage as a foundational pillar in NVIDIA-powered AI Factory implementations. The comprehensive integration ensures that enterprises can confidently build their AI strategies on Dell infrastructure, knowing that their platform will continue to deliver optimal performance as NVIDIA's AI technologies evolve.

DELL'S COMPETITIVE DIFFERENTIATION

The AI Data Platform offers a purpose-built, AI-optimized, and unified data foundation that delivers high performance, openness, and resilience, addressing the unique needs of modern AI pipelines.

In contrast, traditional data systems are often general-purpose, fragmented, and require extensive retooling to support high-throughput, AI-ready environments.

Let's look at in more detail:

Challenge	Dell AI Data Platform
Open, Flexible Approach that avoids vendor lock-in	Modular architecture that allows customers to build the solution right for their environment, with a rich partner ecosystem and adherence to open formats and open source-based engines.
Data Ingest	Optimized for high-throughput parallel I/O (up to 40 GB/s/client with RDMA and GPU-aware streaming)
AI Pipeline Integration	Built-in support for RAG, LLM prep, feature extraction, and NeMo connectors
Unified Data Architecture	Supports NFS, SMB, HDFS, S3, and streaming via one namespace.
Streaming Data Support	Native support via Streaming Data Platform.
Query & Analytics Layer	Includes Data Engines with federated query and metadata intelligence.
RAG & Vector Search Readiness	Integrated PowerScale connector for NVIDIA NeMo Retriever, AI-ready metadata indexing
AI/ML Framework Compatibility	API-first, open standards.
Cyber Resilience & Governance	Embedded snapshot immutability, RBAC, XDR, data masking, and air-gapped backup.
Scalability	Modular, software-defined, and independently scalable across data types and tiers
Data Discovery & Metadata Enrichment	Automated classification and tagging via MetadataIQ, integrated search indexing
AI Infrastructure Compatibility	Validated with NVIDIA DGX, Intel Gaudi 3, AMD Instinct; supports RDMA and GPUDirect.

Deployment Time

Faster time-to-value with validated designs, Omnia automation, and turnkey reference architectures

WRAPPING UP

Organizations investing in AI infrastructure today must balance current requirements with future flexibility. The most successful AI implementations are built on platforms that provide:

- **Open Architecture Foundations** that avoid vendor lock-in while enabling integration with emerging AI technologies and best-of-breed tools. This openness ensures that infrastructure investments remain valuable as AI technology continues to evolve.
- **Proven Enterprise Reliability**, built on decades of experience with mission-critical applications, ensures that AI infrastructure can meet the availability and performance requirements of production deployments. This foundation of reliability is crucial for AI initiatives that have a direct impact on business operations.
- **Comprehensive Ecosystem Support** through strategic partnerships and broad compatibility ensures that organizations can adapt to emerging AI technologies while protecting their infrastructure investments. This ecosystem approach provides access to innovation while maintaining operational stability.

A modern AI data platform represents more than data management infrastructure, it's an intelligent foundation that enables enterprises to transform their operations through AI while maintaining control over their data, costs, and competitive advantages. As AI continues to reshape the business landscape, organizations with robust, scalable, and intelligent data platforms will be best positioned to capitalize on emerging opportunities.

AI is rapidly evolving, and infrastructure investments must provide a set of essential capabilities to remain relevant as AI technology advances:

- **Modular, Container-Integrated Architectures** must support rapid deployment and scaling of AI applications while maintaining the stability and performance that production applications require. These architectures need to accommodate the continuous integration and deployment patterns that characterize modern AI development.
- **Edge and distributed intelligence support** become increasingly important as AI applications expand into industrial, healthcare, retail, and other specialized environments. These deployment scenarios require infrastructure that can support distributed processing while maintaining centralized governance and control.
- **Advanced Data Mobility Capabilities** enable seamless movement of data and models between different deployment environments. Organizations need infrastructure that can adapt to changing deployment patterns while maintaining performance, security, and compliance characteristics across diverse environments.
- **Sustainability and Efficiency Focus** addresses growing concerns about the environmental impact of AI workloads. Efficient storage and compute designs that reduce power consumption and cooling requirements are becoming critical factors in total cost of ownership calculations and environmental impact assessments.



Dell's AI Data Platform provides a purpose-built foundation that unifies performance, governance, and ecosystem integration into a single, future-ready stack. By aligning data infrastructure with the realities of modern AI pipelines, enterprises can accelerate time to insight, ensure resilience and compliance, and position themselves to take advantage of the next wave of AI breakthroughs with confidence.

For more information visit Dell's [AI Data Platform](#) page.



© Copyright 2025 NAND Research. NAND Research is a registered trademark of NAND Research LLC, All Rights Reserved.

This document may not be reproduced, distributed, or modified, in physical or electronic form, without the express written consent of NAND Research. Questions about licensing or use of this document should be directed to info@nand-research.com.

The information contained within this document was believed by NAND Research to be reliable and is provided for informational purposes only. The content may contain technical inaccuracies, omissions, or typographical errors. This document reflects the opinions of NAND Research, which is subject to change. NAND Research does not warranty or otherwise guarantee the accuracy of the information contained within.

NAND Research is a technology-focused industry analyst firm providing research, customer content, market and competitive intelligence, and custom deliverables to technology vendors, investors, and end-customer IT organizations.

Contact NAND Research via email at info@nand-research.com or visit our website at nand-research.com.