

Fuel Analytics and AI Workloads with the Right Data

Discover, query, prepare data for AI with Dell Data Engines

3-5x

Faster queries¹

Up to

53%.

Reduction in the cost of data analytics¹

3x

faster deployment with ProDeploy or ProDeploy Plus³

What are data engines?¹

Purpose-built for data engineers and architects, these engines unify search, integrate seamlessly with AI and data tools, and drive transformative outcomes for enterprise-scale AI and analytics.

Why is object storage recommended?

A best-in-class object storage is essential to realize all the benefits of data engines. Dell's Object Storage offerings - experience 76% lower TCO than public cloud, and the world's most powerful object storage appliance for Kubernetes² - combined with the power of the Dell Data Engines.

It allows organizations to modernize their data strategy to prepare for an AI-driven future.

Learn more about:

- Dell ObjectScale
- Dell PowerScale

In a world driven by information, the ability to manage and leverage data is your greatest advantage. However, data is often scattered across various systems and formats, creating challenges for data engineers and IT architects to unify it for critical AI and analytics workloads.

Dell's [AI Data Platform \(AIDP\)](#) is purposefully designed to tackle these challenges head-on. It focuses on enabling three key outcomes when working with data:

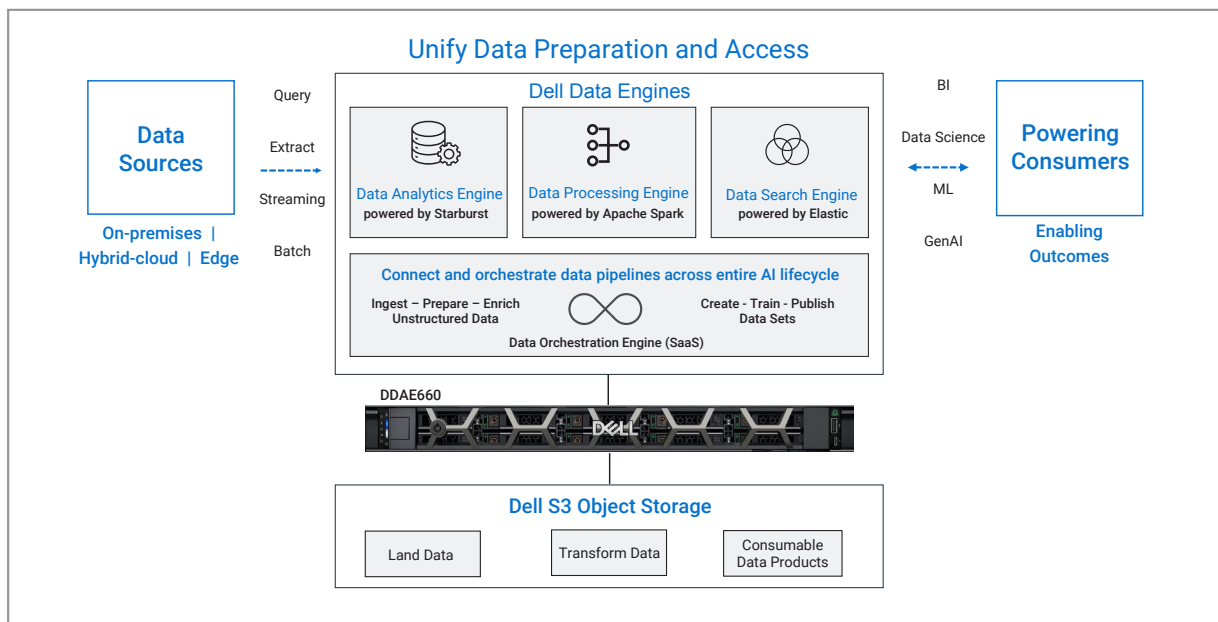
- **Placement of Data:** Store data on high-performance storage to deliver efficiency at scale—supporting data ingestion, AI model training, fine-tuning, and inferencing.
- **Processing Data:** Ensure seamless, unified access to all enterprise data—empowering analytics and AI by transforming raw data into ready-to-use assets.
- **Protecting Data:** Enable secure, compliant operations with built-in data protection and cyber resiliency you can trust.

Core for **processing data are Dell Data Engines**, which empower organizations to query, search, enrich, analyze, and gain insights from both enterprise-created and externally generated data.

Dell Data Engines: Powering Analytics and AI

Unlike proprietary engines that limit flexibility, Dell Data Engines are built with open, cutting-edge tools to handle the complexity of modern data environments. They offer a suite of integrated solutions designed for data engineers and architects, focused on three key functions:

- **Data Analytics:** Enable powerful, federated queries across all your structured data, whether it's in the cloud, on-premises, or in various databases. This allows you to pull together insights from disparate sources quickly and efficiently.
- **Data Search and Enrichment:** Maximize your ability to search across all unstructured data, including documents, logs, and transcripts. This is critical for powering large language models (LLMs) with vector search and AI-driven discovery.
- **Data Ingestion and Processing:** Support both real-time data streams from sources like IoT devices and traditional batch updates. This ensures your data is always ready for analysis, no matter how it's ingested.
- **Data Orchestration:** Effortlessly manage and orchestrate enterprise multimodal data for generative AI applications, seamlessly running AI workflows across cloud and on-prem environments, while simplifying deployment and scaling with NIMs and AI blueprints.



Data Analytics Engine

The Data Analytics Engine accelerates your journey from raw information to actionable insight. Built to scale, it delivers rapid analysis across complex datasets, empowering your teams to discover trends and make informed decisions.

Key Capabilities:

- **Seamless Multi-Source Querying:** Query across diverse data sources using standard SQL without needing to move or consolidate data.
- **Broad Connectivity:** Connect to relational databases, data lakes, object stores, and NoSQL systems.
- **High-Performance Analytics:** Execute distributed queries in parallel for fast, unified analytics.

Data Search Engine

The Data Search Engine enables swift, accurate discovery across large volumes of enterprise data. With advanced algorithms and native Elasticsearch integration, this engine helps teams find the information they need, when they need it.

Key Capabilities:

- **Scalable Search Capabilities:** Supports full-text and vector search across large volumes of unstructured data.
- **Fast and Relevant Results:** Utilizes inverted indices for quick, relevance-ranked search results.
- **AI-Driven Discovery:** Leverages vector embeddings for semantic and AI-powered search.

Data Processing Engine

The Data Processing Engine is purpose-built to accelerate data preparation and transformation at scale for both batch and streaming data. It streamlines data ingestion, cleansing, and enrichment, ensuring your initiatives are powered by high-quality information.

Key Capabilities:

- **Distributed In-Memory Processing:** Provides high-performance processing for both real-time and batch data.
- **Comprehensive Data Workflows:** Supports ETL, analytics, and machine learning through a unified API.
- **Low-Latency Streaming:** Enables structured streaming for real-time data from sources like Kafka.

Data Orchestration Engine Definition

The Data Orchestration Engine is specifically designed to unify and streamline the AI development lifecycle, enabling seamless management of multimodal unstructured data, automated pipelines, and collaborative model iteration. Engineered for enterprise-grade performance and scalability, it supports cloud, on-prem, and hybrid environments with robust governance, intelligent data curation, and developer-centric extensibility to accelerate the creation of advanced AI applications.

Key Capabilities:

- **Unstructured Data Management:** Visualize and curate enterprise multimodal data for generative AI applications.
- **Hybrid-Native Orchestration:** Run AI workflows seamlessly across cloud and on-prem environments.
- **Accelerate Deployment with NVIDIA Hub:** Access NIMs and AI blueprints directly within the orchestration engine to simplify, deploy, and scale AI workflows.

Dell DDAE660:

The Dell DDAE660 is a high-performance appliance that simplifies IT operations while delivering scale and efficiency required to run the Data Engines. As a key foundation for this ecosystem, the DDAE660 offers:

- **Enterprise-Ready Features:** Integrates seamlessly into your environment with built-in redundancy for high availability, active/passive disaster recovery, and robust security measures like SSO and encryption.
- **Built for Scalability:** Allows your organization to easily expand by adding nodes as your data grows.
- **Simplified Management:** Benefits from managed upgrades and embedded Kubernetes, OS, middleware, and database support.

With its out-of-the-box features, the DDAE660 empowers IT teams to achieve faster time to production, reduced downtime, and significant operational cost savings.

Fast-Track Your Insights with Our Help

Dell Services provides the guidance and expertise you need to modernize your data environment and accelerate your AI outcomes.

- **Implementation Services for Dell AI Data Platform:** Deploy and optimize your Dell Data Engines to accelerate AI-driven insights and data analytics outcomes.
- **Optimization Subscription for Dell AI Data Platform:** Access a flexible, dedicated expert to continuously enhance your Dell Data Engines and unlock more value from your AI and analytics initiatives.
- **ProDeploy for Infrastructure Suite:** Deploy hardware with expert assistance to reduce complexity and accelerate implementation.
- **ProSupport for Infrastructure Suite:** Gain access to critical support to ensure you are up and running as quickly as possible when issues arise.

Technical Components

Dell Data Engines are deployed in a Dell's DDAE660 which includes Dell Data Lakehouse System Software that provides lifecycle management and tailor-made compute hardware all integrated into one. For storing and processing large datasets in open table formats, Dell's leading S3-compatible storage platforms such as ObjectScale and PowerScale offer exceptional performance, reliability and security.

Software	Dell Data Analytics Engine, powered by Starburst Dell Data Processing Engine, powered by Apache Spark Dell Data Search Engine, powered by Elastic Dell Data Orchestration Engine (SaaS) Dell Data Lakehouse System Software
Scale-out Compute	Dell DDAE660: compute for data engines (based on Dell PowerEdge)
Scale-out S3 Storage	Recommended for use with Dell PowerScale or Dell ObjectScale

¹ Cloud Data Warehouse vs. Cloud Data Lakehouse: A Snowflake vs. Starburst TCO and Performance Comparison, published by GigaOm

² ESG Economic Validation sponsored by Dell Technologies, "Analyzing the Economic Benefits of Dell ECS: Economic Benefit Analysis of On-premises Object Storage versus Public Cloud," by Tony Palmer, July 2022. Cost savings based on ESG comparison of ECS to a leading public cloud in active storage scenarios.

³ Based on a May 2023 Principled Technologies study "Using Dell ProDeploy Plus Infrastructure can improve deployment times for Dell technology"

⁴ Based on a July 2022 internal analysis of support data from February 2022 through July 2022.

⁵ Based on an internal analysis of service requests from the prior two years for Dell Technologies data protection and high-end storage products covered by ProSupport Plus for Infrastructure vs. products with Basic coverage, February 2024. Actual results may vary.



[Learn more](#) about
Dell solutions



[Contact](#) a Dell
Technologies Expert



[View more](#) resources



[Join the conversation](#) with