Solution Brief

# VAST DataEngine

## Real-Time Execution and Orchestration for Data & AI




Modern AI and analytics pipelines break under the weight of fragmented orchestration, slow ETL, and constant data movement between storage, compute, messaging, and workflow engines. VAST DataEngine changes the model by providing a unified execution and orchestration layer inside the **VAST AI Operating System.**

Using event-driven pipelines, serverless Python functions, managed execution environments, and integrated analytics engines, DataEngine executes workflows where data already resides, eliminating unnecessary infrastructure while enabling real-time, governed pipelines for analytics and AI.



# How It Works

**Unified Execution Where Data Lives**

The VAST DataEngine provides a managed execution fabric that runs directly on the VAST AI Operating System. Using VAST Native Compute, DataEngine executes integrated analytics engines like Trino and Spark, serverless Python functions, managed execution environments, and containerized services directly on VAST compute nodes (CNodes), bringing compute to the data without requiring external Kubernetes clusters or unnecessary data movement.
By unifying execution with the underlying data platform, DataEngine reduces infrastructure complexity, accelerates analytics, and provides a common foundation for streaming pipelines, SQL analytics, AI inference, and data processing.

**Automated, Event-Driven Pipelines**

With DataEngine, data becomes active the moment it arrives. The VAST Event Broker captures file, object, and streaming events and triggers downstream workflows automatically, without job schedulers, cron management, or fragile ETL. Pipelines can automatically transform unstructured content into structured tables, enrich metadata, initiate vectorization and embedding generation, populate VAST DataBase, or launch downstream AI workflows. Whether processing streaming data, ingesting enterprise documents, or preparing AI-ready datasets, DataEngine keeps intelligence flowing continuously and in real time.

**Managed Execution Environments for Python and Containers**

Developers can author lightweight serverless Python functions or package complete containerized applications with their own dependencies, libraries, and container images. These workloads execute on VAST Native Compute, providing a fully managed Kubernetes environment within the VAST cluster while eliminating the need to provision or operate separate compute infrastructure. These workloads can parse, transform, enrich, route, or analyze data while initiating AI inference or downstream processing without moving data into external compute clusters. Execution environments are versioned, reusable, and governed, enabling consistent deployment across development and production. Engineers can work in a code-first model using their preferred IDE, while DataEngine also provides an intuitive low-code interface for visually building, managing, and updating pipelines. This blend of programmable flexibility and visual simplicity accelerates production automation across streaming, batch, analytics, and AI workflows.

**Unified Pipeline Automation Across Streaming, Analytics & AI**

DataEngine orchestrates end-to-end pipelines spanning streaming ingestion, metadata enrichment, batch processing, SQL analytics, vectorization, inference, and event-driven AI workflows without external orchestration tools or workflow engines.
Using integrated Spark and Trino execution together with the VAST InsightEngine, organizations can automate everything from document ingestion and embedding generation to retrieval pipelines and real-time analytics, all on a single platform without copying data between systems.

**Operational Visibility and Pipeline Traceability**

The VAST DataEngine provides full observability and auditing across every pipeline, giving teams complete insight into how data is triggered, processed, and transformed. Each workflow step is tracked, making it simple to see who accessed data, which functions operated on it, and how it progressed through the pipeline. Built-in execution history, tracing, performance telemetry, and pipeline lineage simplify debugging, optimization, compliance, and operational monitoring. This unified operational view eliminates the need to stitch together logs from multiple services and ensures that data processing remains trustworthy, compliant, and easy to manage throughout its lifecycle.

**Unified Workflow Orchestration**

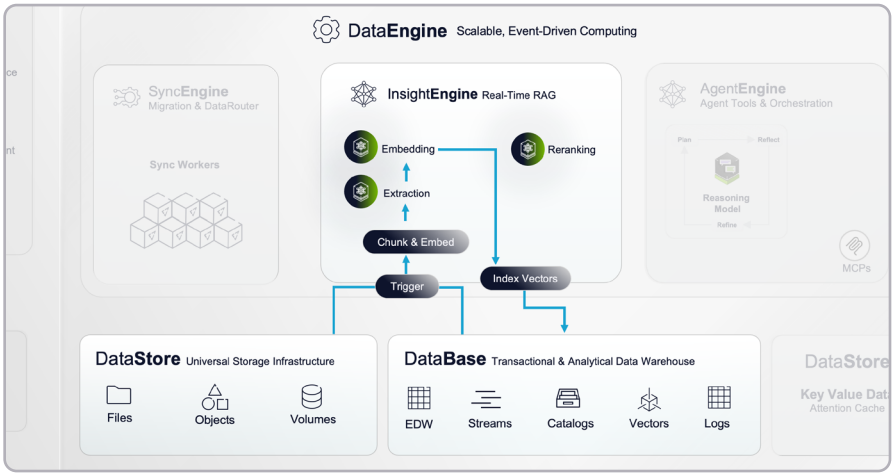
DataEngine orchestrates complex, event-driven workflows across Python functions, containers, SQL engines, AI services, and external applications. Pipelines can execute multiple stages of processing, chain dependent tasks together, and automatically respond to new data as it arrives.
Workloads can execute natively on VAST Native Compute or invoke services running in customer-managed Kubernetes environments, making it easy to integrate existing embedding models, chunking services, reranking pipelines, inference endpoints, and other AI components into a unified workflow. Rather than relying on external workflow engines to coordinate disparate systems, DataEngine orchestrates these pipelines through a common security, governance, and data model while enabling compute to run where it makes the most sense.

# The Foundation for Specialized AI Services

DataEngine serves as the execution layer for the VAST AI Operating System, powering specialized services that operationalize enterprise AI.

## InsightEngine

Provides the framework to build, automate, and govern AI data pipelines for document processing, metadata enrichment, vectorization, retrieval, and enterprise RAG workflows using customer-managed AI models.



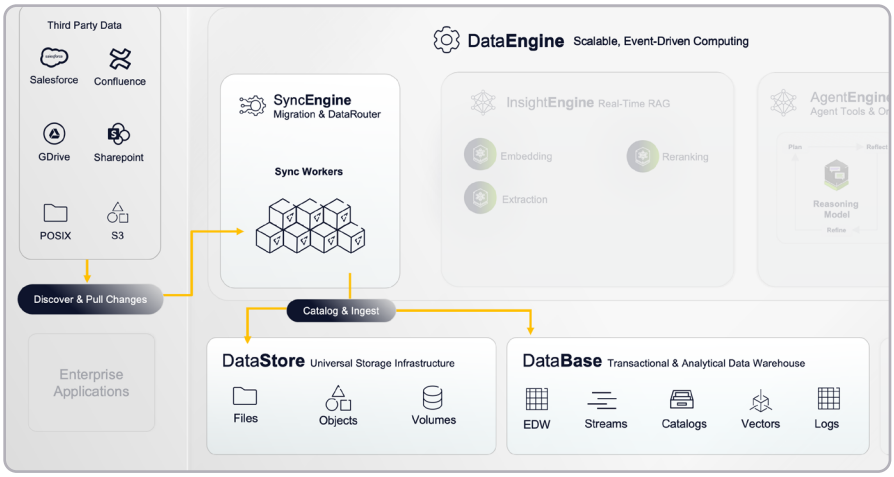

## AgentEngine

Provides the production infrastructure required to deploy, orchestrate, and govern long-running enterprise AI agents.




## SyncEngine

Discovers, ingests, and prepares enterprise data from virtually any source for analytics and AI.




Together, these services extend DataEngine into a complete execution platform for enterprise data, analytics, and AI.

# Why VAST




### Unified Execution

Execute analytics, AI, and data pipelines directly where data resides.



### Real-Time Automation

Trigger pipelines instantly as data changes.



### VAST Native Compute

Execute Python functions, containers, Spark, and Trino on a managed Kubernetes environment built directly into the VAST AI Operating System.



### Exabyte Scale

DASE architecture delivers predictable scale as data grows.



### Governed and Observable

Identity-aware execution with complete pipeline visibility and traceability.



### Execution Layer for AI

Powers InsightEngine, AgentEngine, and SyncEngine on a unified platform.

## About VAST Data

VAST Data is the AI Operating System company – powering the next generation of intelligent systems with a unified software infrastructure stack that was purpose-built to unlock the full potential of AI. The VAST AI OS consolidates foundational data and compute services and agentic execution into one scalable platform, enabling organizations to deploy and facilitate communication between AI agents, reason over real-time data, and automate complex workflows at global scale. Built on VAST's breakthrough DASE architecture – the world's first true parallel distributed system architecture that eliminates tradeoffs between performance, scale, simplicity, and resilience – VAST has transformed its modern infrastructure into a global fabric for reasoning AI. Learn more at **vastdata.com** and follow VAST Data on **LinkedIn**, **YouTube** and **X**.



©2026 VAST Data, Inc. All rights reserved. All trademarks belong to their respective owners.

06/26 I 00000