

Fresh Gravity's Agentic intelligent Data Engineering (AiDE) Solution

Artificial intelligence is evolving from assistive tools to autonomous, goal-driven systems capable of acting with minimal human intervention. Agentic AI represents this next phase, enabling intelligent agents to collaborate, adapt, and execute complex workflows at scale.

Fresh Gravity's Agentic intelligent Data Engineering (AiDE) solution applies this paradigm to modern data engineering, helping enterprises improve speed, scalability, governance, and trust across cloud data platforms.

Background

Traditional data engineering relies heavily on manual coding, static rules, and fragmented toolchains. As data volumes, domains, and expectations grow, these approaches result in long development cycles, brittle pipelines, inconsistent data quality, and high operational overhead.

Despite investments in modern platforms like Databricks and Snowflake, teams often struggle to operationalize data efficiently due to limited intelligent automation, and end-to-end governance often requiring a significant human effort investment.

These challenges are especially visible in initiatives that require high data consistency and traceability, such as analytics modernization, regulatory reporting, and master data programs.

Agentic AI addresses this gap by enabling autonomous, goal-oriented agents that collaborate across the data lifecycle, reducing manual effort while improving scalability, consistency, and trust.

Fresh Gravity's AiDE Solution

Our Agentic intelligent Data Engineering (AiDE) solution is an AI-powered data engineering platform built on a network of specialized agents that collaborate across the entire data lifecycle. AiDE accelerates data discovery, ingestion, pipeline generation, orchestration, data quality validation, reconciliation, and observability.

Instead of engineers manually building and maintaining code-heavy pipelines, AiDE assigns intelligent agents with clearly defined responsibilities, guardrails, and auditability.

Engineers interact with AiDE using simple prompts, providing context for data tasks while retaining control through built-in human-in-the-loop mechanisms. All artifacts generated by AiDE follow a configuration-driven framework, enabling extensibility, lineage, and long-term flexibility across evolving data ecosystems.

Enterprise Platform Readiness & Delivery Model

Production Ready

AiDE is designed to be production-ready on modern cloud data platforms, with native execution support for Databricks (including Mosaic AI capabilities) and Snowflake (including Cortex). It leverages compute, storage, security, and orchestration features to generate and execute pipelines without introducing proprietary runtime dependencies.

Service-Led Accelerator Model

AiDE is **not positioned as an off-the-shelf product**. Instead, it is a service-led accelerator framework delivered by Fresh Gravity, combining a reusable library of production-ready AI agents, engineering standards, and orchestration patterns. These agents can be deployed as-is or extended based on client-specific use cases, data platforms, and governance requirements.

Seamless Enterprise Integration

By integrating seamlessly into existing enterprise architectures, AiDE enables agent-driven data engineering while preserving current platform investments, security models, and governance controls, delivering product-level consistency with the flexibility of services-led execution.

Key Features

AI Agents & Automation

- Specialized AI agents for data discovery, ingestion, pipeline creation, orchestration, execution, and optimization
- Prompt-based execution with automated generation of production-ready artifacts
- Central controller layer orchestrating agents end-to-end with built-in guardrails and auditability
- Native execution on Databricks (Spark, PySpark, Delta) and Snowflake (SQL, tasks, streams, Cortex-enabled workflows)

Gold Layer Modeling & Load Automation

- Prompt-driven creation of Gold-layer modeled data objects, including ER-style data models
- Automatic generation and execution of DDLs to create target tables
- Agent-driven development and execution of stored procedures to load and refresh Gold-layer tables
- Support for multi-table modeling and loading in a single workflow, sourcing schema and data from multiple upstream tables and systems

Governance, Lineage & Observability

- End-to-end lineage and traceability maintained across data processes, including modeled, and aggregated layers
- Centralized metadata repository managing source-to-target mappings
- Observability dashboards available out-of-the-box for job monitoring, compute utilization, alerting, and anomaly detection

Data Quality & Reconciliation

- Dedicated agents performing automated data quality checks and reconciliation
- Detection and reporting of variations, threshold breaches, and inconsistencies
- Support for high-trust pipelines feeding analytical, operational, and master data systems

Engineering Standards by Design

- Configuration-driven pipeline generation ensuring consistent structure, naming conventions, and reusability across all artifacts
- Built-in support for Bronze–Silver–Gold architectures, including dependency management, and multi-table orchestration
- Automatic generation and execution of DDLs, DMLs, and stored procedures aligned with platform best practices

- Embedded engineering patterns such as full load, incremental load, CDC-style processing, loading facts, dimensions, and performing aggregations
- First-class support for schema evolution, including detection, impact analysis, and controlled propagation
- Native handling of PII and sensitive data through encryption, hashing, masking, and role-based access controls
- Integrated audit logging, error handling, and retry mechanisms for operational resilience
- Compatibility with enterprise security, identity, and governance models in Databricks and Snowflake, including role-based access and audit trails

How AiDE Agents Transform Access

Without AiDE Agents

- Manual, code-heavy pipelines
- Long wait times for data engineering support
- Technical silos limiting cross-functional insights
- Static validation and limited transparency

With AiDE Agents

- Natural language interaction for data engineering and modeling tasks
- Instant, self-service pipeline and Gold-layer table generation
- Cross-domain data integration with lineage and governance by design
- Interactive data quality checks and guided validation

How AI Agents Enable Data Democratization

AiDE reduces technical barriers across the organization, enabling engineers, analysts, and data stewards to collaborate effectively. By abstracting complexity through agents, AiDE improves data literacy and confidence while supporting consistent, governed data usage across domains, including master and reference data.

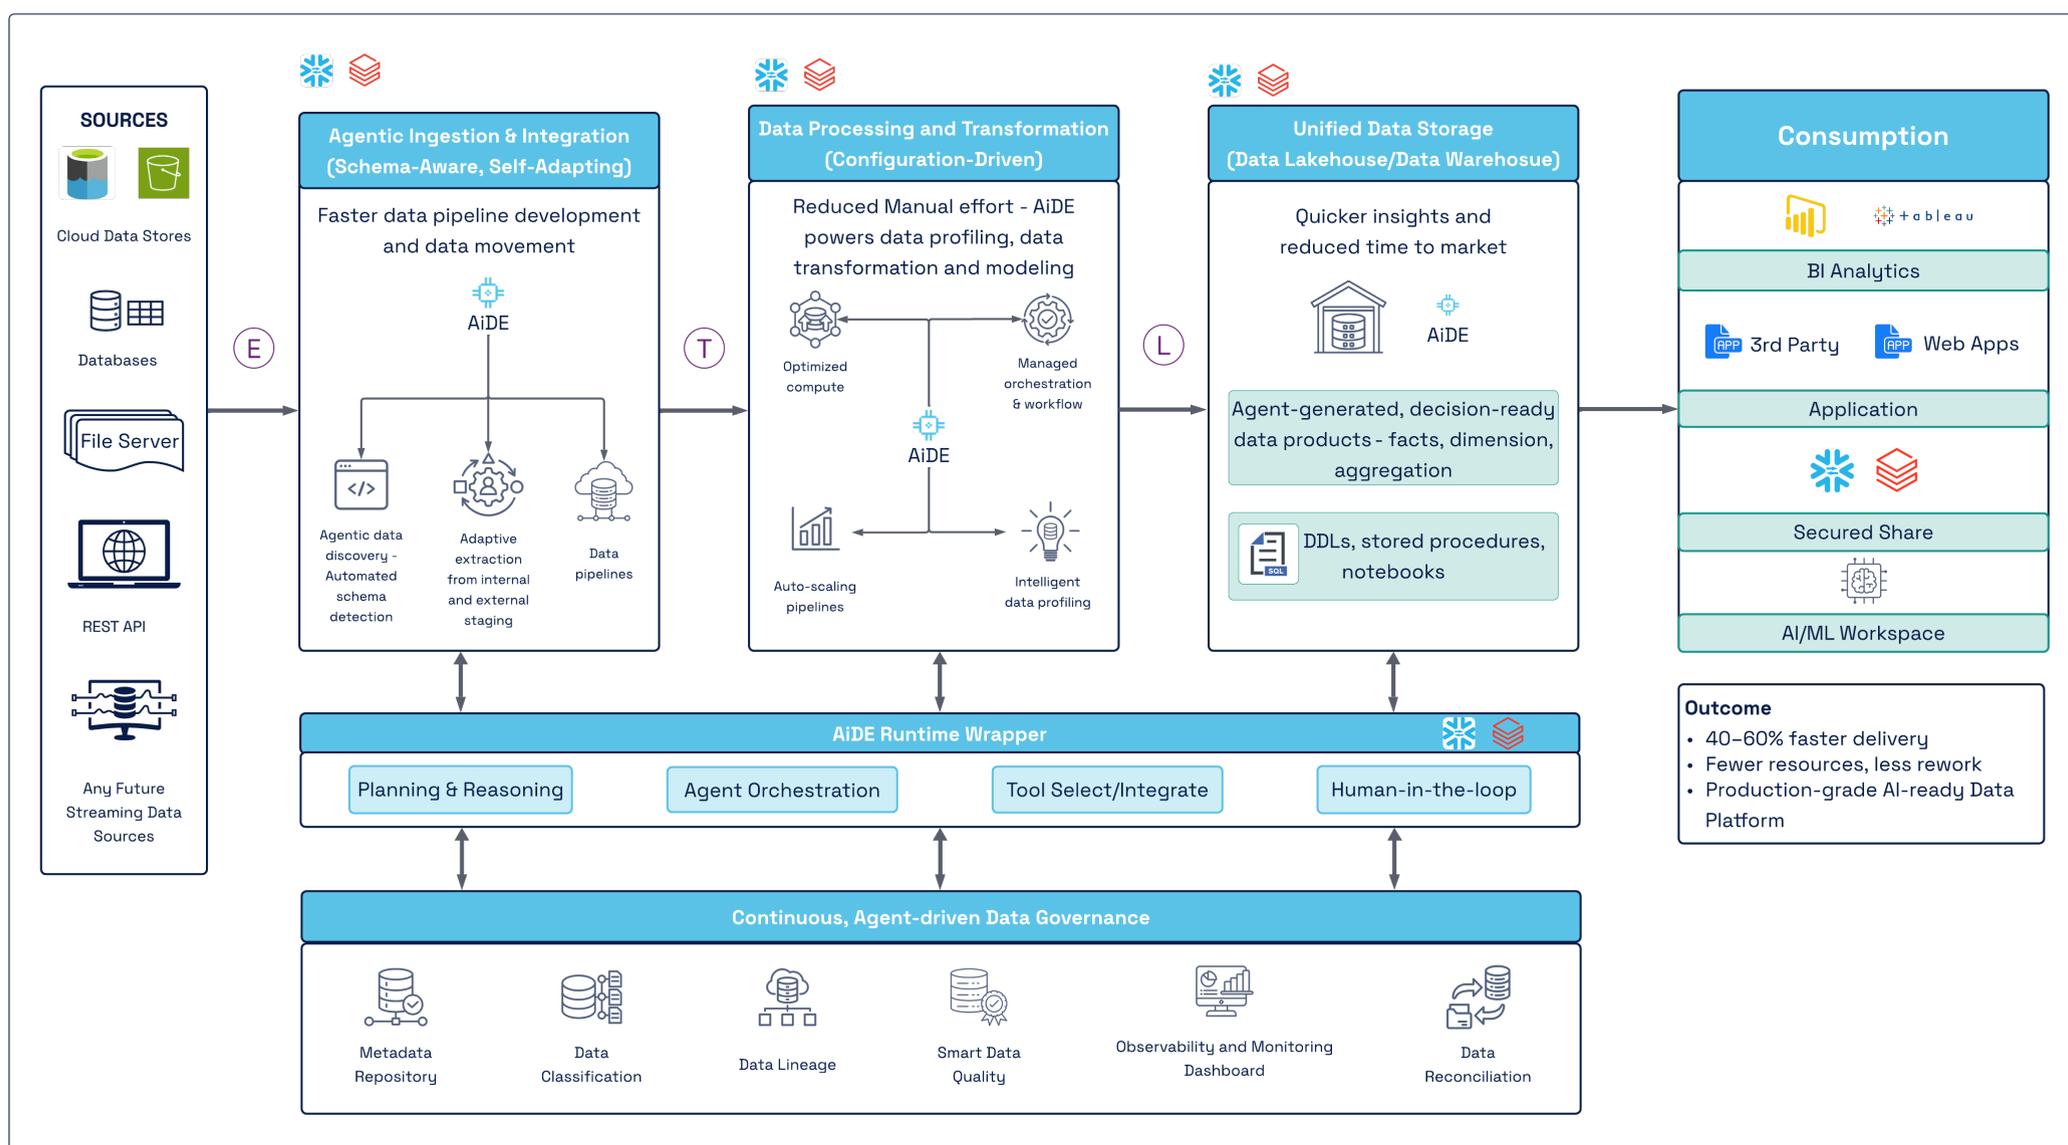
Business Impact*

By leveraging agent-driven execution and platform-native capabilities in Databricks and Snowflake Cortex, AiDE delivers measurable business impact:

 <p>Faster Decisions *78% reduction in time to insights for business users</p>	 <p>Innovation *3x increase in data-driven ideas from non-technical teams</p>	 <p>Data Literacy *92% of users report higher confidence in using data</p>	 <p>Collaboration *64% increase in cross-functional data sharing and analysis</p>
--	---	--	---

*Metrics based on standard data engineering delivery for a CDP over a 20-week implementation window.

AiDE Platform Agnostic Diagram



About Fresh Gravity

Founded in 2015, Fresh Gravity helps businesses make data-driven decisions. We are driven by data and its potential as an asset to drive business growth and efficiency. Our consultants are passionate innovators who solve clients' business problems by applying best-in-class data and analytics solutions. To know more about us and our offerings, contact us at info@freshgravity.com.

US Corporate HQ – Dallas, TX

2201 Spinks Road, Suite 205
Flower Mound, TX, 75022

Other Offices



USA
Santa Clara



Canada
Toronto



United Kingdom
London



India
Pune