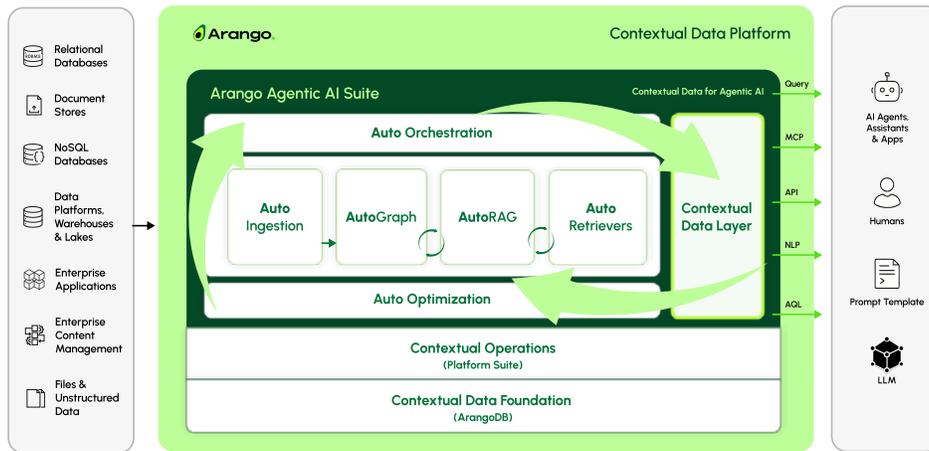


ArangoDB provides a native multi-model database engine that unifies graph, document, key-value, vector, and full-text search capabilities within a single operational system and query language.

As organizations increasingly rely on connected enterprise data to power analytics, AI applications,

and agentic workflows, production-ready systems require flexible data modeling and consistent access across structured and relationship-rich datasets. The value drivers below summarize how ArangoDB enables teams to store, connect, and retrieve enterprise data within a single operational environment.



Outcomes and Benefits

Once deployed and operationalized, the Arango Data Platform helps teams standardize data operations and reduce platform overhead while enabling reliable scale for production workloads.



Outcome	Benefits
Unified Multi-Model Data	Manage vector, graph, document, and key-value data in one operational engine—removing the complexity of integrating and maintaining separate data stores.
Relationship-Aware Retrieval	Model real-world dependencies between entities using native graph capabilities for improved investigation and decision support.
Consistent Querying Across Data Types	Use a single declarative query language (AQL) to retrieve and analyze structured and connected data within one query.
ACID (Atomic, Consistent, Isolated, Durable) Transactions and Data Integrity	Ensures reliable and consistent data operations with ACID-compliant transactions, supporting multi-document transactions in single-server deployments and strong consistency for operational workloads in clustered environments.
Flexible Data Structures	Store JSON documents with adaptable schemas to support evolving enterprise datasets and application requirements.
Integrated Search and Aggregation	Perform graph traversals, aggregations, geo-spatial queries, and full-text search within the same execution engine.
Scalable Distributed Operations	Scale horizontally through clustering while maintaining transactional guarantees for enterprise workloads.

Key Value Propositions

ArangoDB is a unified multi-model database that stores and queries documents, graphs, key-value data, and search indexes in a single engine. By modeling entities and relationships natively, it gives applications, analytics, and AI systems direct access to rich business context without multiple databases. This architecture provides a scalable foundation for connected data, combining flexible data management, strong transactions, and efficient queries across diverse datasets.



How It Fits in the Platform

ArangoDB serves as the core multi-model data engine within the Arango portfolio—storing and querying enterprise data across graph, document, key-value, and search models. It operates within the Arango Data Platform for deployment and lifecycle management and provides the contextual data foundation used by AutoGraph for retrieval workflows and the Agentic AI Suite for AI agent execution.

Together, this architecture enables AI applications, agents, and co-pilots to **retrieve and reason over connected enterprise context in production environments**.



Why Choose ArangoDB?

ArangoDB is designed for organizations that want to store and query connected enterprise data within a single operational system — without maintaining separate databases for different workloads.

Multi-Model by Design: Unifies documents, key-values, and graphs in one database core for flexible data modeling.

Single Query Language: Access structured and relationship-rich data using AQL without switching tools or pipelines.

Relationship-Aware Storage: Model real-world dependencies between entities to support analytics and investigation workflows.

Scalable Foundation: Supports distributed clustering for enterprise-scale applications and data services.

Ready to activate
AI with context,
trust, and scale?

Schedule a demo to see how ArangoDB jumpstarts intelligent applications and agents.

arango.ai