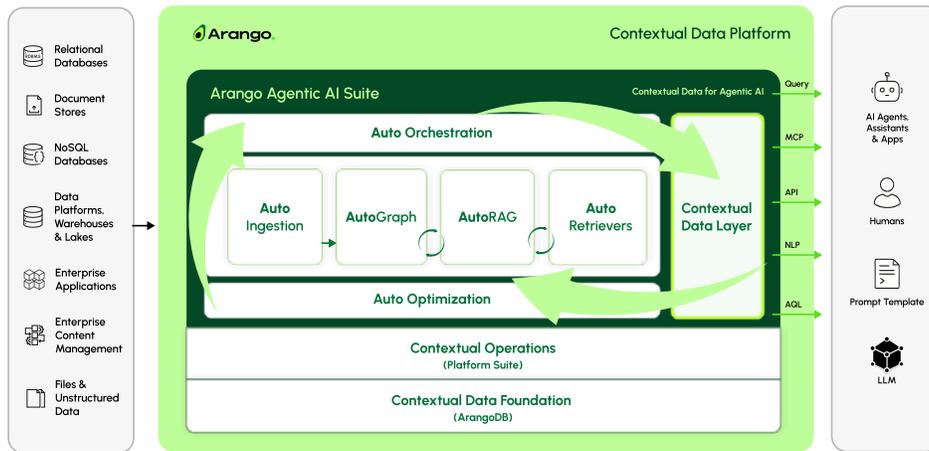


# Arango AutoGraph

## Arango AutoGraph is a contextual retrieval architecture for enterprise AI agents & assistants.

It ingests structured and unstructured data that automatically organizes enterprise data into a connected Context Graph — linking documents, entities, and operational signals into domain-aware knowledge shards for AI agents and assistants.

By unifying structured, semi-structured, and unstructured data into a multi-model graph with vector embeddings, Arango AutoGraph builds the graph to improve response accuracy, consistency, and explainability across enterprise AI applications.



## Outcomes and Benefits

Arango AutoGraph organizes enterprise data into connected context and applies domain-aware retrieval before inference, enabling AI agents and assistants to deliver more relevant, consistent, and explainable outcomes.

It structures documents, entities, and relationships into knowledge models that power ContextRAG, AutoRAG, and the Arango Co-Pilot Builder — improving response relevance, reducing variability, and supporting governed AI workflows.



Outcome	Benefits
<b>Auto-Generated Context Graphs</b>	AI agents and assistants can reason over connected enterprise data, improving decision quality and reducing time spent reconciling fragmented information.
<b>Domain-Aware Retrieval</b>	Responses are tailored to domain-specific context, resulting in more relevant outputs and fewer incorrect or incomplete recommendations.
<b>Multi-Modal Data Integration</b>	AI systems can incorporate operational, transactional, and document-based inputs — enabling more comprehensive insights across business processes.
<b>Reduced Hallucinations</b>	LLM outputs are grounded in enterprise relationships and source data, improving response accuracy and reducing unsupported or speculative answers.
<b>Explainable AI Outcomes</b>	Responses can be traced back to source entities and relationships, supporting auditability, governance, and regulatory compliance.
<b>Consistent AI Responses</b>	AI agents deliver more predictable outputs across workflows by grounding decisions in a shared, domain-specific context.

# Core Arango AutoGraph Components and Capabilities

Arango AutoGraph includes the following native capabilities that structure enterprise data into contextual knowledge shards and assign domain-aware retrieval strategies prior to inference. These components enable AI agents and assistants to retrieve relevant context across structured and unstructured sources for production-grade decision support.

## ContextRAG Corpus Graph

Automatically organizes enterprise data into connected documents, chunks, and partitions — creating domain-aware knowledge shards for contextual retrieval across structured and unstructured sources.

## AutoRAG Retrieval Strategizer

Dynamically assigns retrieval strategies per domain to optimize chunk selection, detect inconsistencies, and improve response accuracy before inference.

## Arango Co-Pilot Builder

Provides a development environment for building AI co-pilots and agents that interact with contextual enterprise data using NLP-driven document processing and prompt orchestration.

## Knowledge Shard Partitioning

Segments enterprise data into domain-specific knowledge shards to improve retrieval speed, relevance, and response consistency across AI workflows.

## Multi-Model Context Storage

Combines vector embeddings, graph relationships, documents, and key-value data in a single operational cluster — eliminating the need for separate context stores.



## How It Fits in the Platform

Arango AutoGraph operates in the contextual retrieval layer of the Arango Contextual Data Platform, transforming fragmented enterprise data into a connected Corpus Graph used by AI agents, enterprise copilots, analytics, and natural language interfaces.

It works in conjunction with:

- The **Agentic AI Suite** for agent orchestration
- GraphRAG / ContextRAG workflows for contextual grounding
- AQLizer for natural language query execution
- Public or private LLM deployments

This architecture enables production-grade AI systems to retrieve governed, relationship-aware knowledge in real time — supporting human-guided decision-making at enterprise scale.

## Why Choose Arango AutoGraph?

AutoGraph helps organizations turn fragmented enterprise data into AI-ready context — without complex, multi-tool architectures.

**AI-Ready Instantly:** Converts enterprise data into graph + vector structures for RAG, agents, and AI assistants.

**Faster Time-to-Value:** Removes manual data modeling and pipelines to speed deployment.

**Production-Oriented:** Supports governed, explainable AI from pilot to production with relationship-aware data.

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

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

[arango.ai](https://arango.ai)