


Aura Intelligence

Aura Intelligence is the set of real-time AI capabilities built into SingleStore Helios. It runs analytics, models, and machine learning on your live operational data, in the same system that stores it. Business users get answers in plain language, developers call AI straight from SQL and their own applications, and none of it needs a separate stack to feed.

SingleStore is trusted by the world's most demanding enterprises



Why AI stalls on enterprise data

What teams run into

The volume of data inquiries from business users far outpaces the capacity of BI and analytics teams.

Heavy AI analysis on a live legacy database, an Oracle, SQL Server, MySQL, or Postgres system, competes with the transactions it was built to serve.

Models lack shared context, forcing agents to reconstruct the business landscape on every call.

Moving data to an external platform to run AI alongside it compromises security and disrupts audit controls.

A pilot operating smoothly for a small test group is suddenly scaled across the entire organization.

What it costs the business

Data requests either pile up in backlogs or go entirely unasked, leading to **critical decisions** being made without the insights already available within existing data (live and historical).

Analyzing live data compromises the performance of customer-facing systems, whereas utilizing delayed copies defers threat detection until transactions have already cleared and the opportunity to intervene is lost.

Pipeline and compute costs continue to escalate, yet models still **fail to deliver trustworthy answers** to actual business questions.

With approval withheld by legal and security teams, the intended users are left waiting indefinitely for the tool.

When subjected to a genuine production load, the **deployment that performed flawlessly in testing starts to fail**: simultaneous queries surge, response latency increases, and operational costs spiral.

Every one of these traces back to the same fault line. Your live data, the context that explains it, and the AI meant to act on it were built to live in separate systems. So nothing can answer a governed question while the answer still matters. Live intelligence needs three things to be true at the same time. Most architectures give you two.

Context

The right data, with your definitions and permissions intact, served by the SingleStore Context Engine.

1

Concurrency

Thousands of users, apps, and agents querying at enterprise scale without contention.

2

Real time

Answers on data as it is now, not on a copy exported in the past.

3

Aura Intelligence brings all three into one engine.

One system for your data and the AI that uses it

Aura Intelligence is the set of AI capabilities built into SingleStore Helios. They run directly on your SingleStore databases, the instances holding your live data, so there is nothing to copy or keep in sync. The **SingleStore Context Engine** gives every capability a governed source of context, combining the definitions your team supplies with what the platform learns on its own. A **Domain** scopes the data and context each user group can reach. The result is one place to ask a question, call a model, run a function, or generate an embedding, all on current data under your control.

Five ways to put AI on your data

Analyst

An AI data analyst that turns a question into SQL, runs it, and explains the result. Business users work with it in the portal at analyst.singlestore.com, or embed it in your application through the API.

AI Functions

Call LLMs and embedding models directly in SQL for sentiment analysis, summarization, classification, and semantic search, choosing the model you want at run time.

ML Functions

Train and run traditional machine learning next to your data, including anomaly detection and classification, with no export to a separate ML stack.

Models

One place to host or connect LLM and embedding models through inference APIs, and to train, register, and serve your own machine learning models.

Embedding Service

Generate vector embeddings inline as data moves through pipelines, and run vector search from SQL, so retrieval never leaves the database.

Use it in the portal, in SQL, or in your product

How teams use it

Business users work in the Analyst portal. Developers call Analyst, AI Functions, and embeddings from SQL and APIs inside their code, while administrators manage them in one place.

Where it runs

Aura Intelligence runs inside SingleStore Helios, the fully managed cloud service, available on AWS, Azure, and Google Cloud. Helios handles provisioning, scaling, and management, so there is no complex infrastructure for your team to maintain.

What it changes for the enterprise

Revenue growth

Faster decisions on live data improve **customer experience and conversion** while the moment still matters.

Analyst embedded in your product opens **new intelligence-as-a-service revenue**.

Cost reduction

Business users self-serve on live data, **clearing analyst cycles and reporting backlogs**.

In-database AI and one context layer **replace fragmented pipelines, caches, and copies**.

Risk mitigation

Answers grounded in **governed enterprise data** reduce hallucination risk.

Access control means every user, app, and agent **sees only what it is permitted to see**.

See Aura Intelligence answer your questions, on your data.

Book a working session with your SingleStore team.