AGENTICOBJECTS · SECURITY AND AUDIT NOTE

You do not have to trust the language model. The architecture is built to check it.

This note is written for the IT, information security and internal audit teams evaluating AgenticObjects. It assumes that large language models (LLMs) can make mistakes, invent figures and misread a question. It explains which mechanism in the product answers each of those risks, where the evidence can be seen, and where each mechanism stops.

Version 1.0 · October 2026 · D-CAT Technologies

SUMMARY

Six core guarantees

01

The model does not produce figures

Every figure comes from a query that runs on your data warehouse. Before publishing, the engine checks every figure in the answer against the query result.

02

The model never touches data

The model writes no queries and has no connection to the warehouse. It sends a structured request; the product's engine builds the query and runs it through a read-only gateway.

03

What reaches the model is limited

The question, measure and dimension names, and a limited number of summary rows. Personal-data fields are masked. A copy of the warehouse is never sent.

04

Every step is recorded

Model calls, queries, permission and setting changes are recorded. Records can only be appended; they cannot be altered afterwards.

05

The model never changes on its own

There is no automatic fallback model. Choosing a model is a human decision, and it is recorded.

06

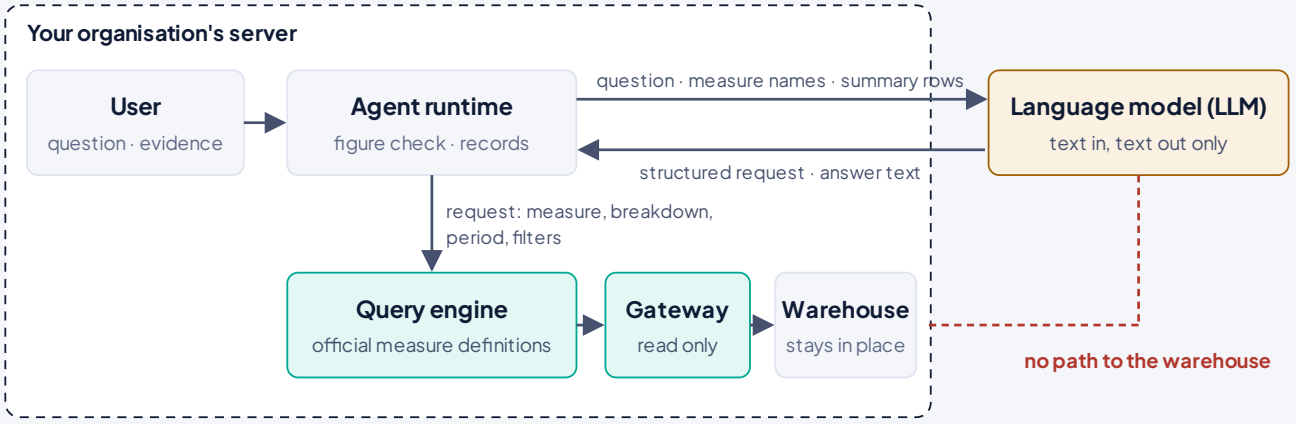
Everything runs on your server

The application, its records and its backups run on your organisation's server. Model calls can stay on in-country infrastructure if required.

ARCHITECTURE

The path of the data: where does the model stand?

The product's central design decision is to keep the language model **outside** the data path. The model understands the question and writes the answer; deterministic components do everything in between.



The model never connects to the data warehouse and never writes a query. The product's engine builds the query from the data model's official measure definitions, and the query runs through a gateway that lets only read queries through.

RISK AND CONTROL MAP

The risks an audit will ask about, and the product's answer

Risk	Control in the product	Where the evidence is seen
The model invents or miscalculates a figure	Every figure in an answer must trace to the result of the query that ran for that answer. A figure that does not trace is corrected, removed, or shown with an "unverified" mark.	The evidence card behind every figure: measure, period, filters, row scope, data date.
The direction of a change or a ratio is misread	Direction words such as "rose" or "fell" are produced by rule. Ratios state their denominator. A currency mismatch is refused before publishing.	The evidence card and the answer's verification state.
The model reaches data it is not permitted to see	The model has no access to the warehouse. The query runs with the permissions of the person who asked; row-level permissions are added to the query by the engine.	The permission scope on the run record; the audit log.
Data in the warehouse is changed or removed	The gateway passes only read queries; multiple statements, writes to files and any statement it does not recognise are refused. The connection also uses a read-only warehouse account.	The query log; for authorised people, the text of the query that ran.
Sensitive data reaches an external model	A limited number of summary rows per query reaches the model (at most 40 in chat, 80 in scheduled analyses). Values of fields marked as personal data are reduced to initials. Model calls can stay in-country.	The full text sent to the model, in the model-call record.
The model changes silently	The model is fixed at the start of each run and does not change during it. There is no automatic switch to a fallback model or on quality grounds.	The model used, on every run record; model choice in the audit log.

Risk	Control in the product	Where the evidence is seen
Cost runs out of control	Every analysis has a step, time and cost ceiling fixed before it starts; at the ceiling the analysis stops and is marked "cut off". Monthly budgets alert at 80% and 100%, and an unusual rise in daily spend raises an alert.	The spend screen: by user, model and agent.
The source connection or the model provider goes down	An unfinished analysis is never presented as complete; it closes as an error and the people concerned are notified. Completed steps are never re-run; the analysis resumes where it stopped.	The run page: steps, error reason, notifications.
The control layer is switched off	Figure checking is not a separate service; it runs inside the engine on every answer and cannot be turned off by any product setting. The only choice an administrator has is whether an unverifiable figure is hidden or marked.	Setting changes, in the audit log.
Someone bypasses the user interface	Permissions are enforced on the server: session, cross-site request forgery (CSRF) protection, action permission and resource scope are checked on every request. The default answer is refusal.	Refused requests; the audit log.
Source data is restated after the fact	Every answer is frozen at the moment it is produced: the query, filters, the date the data reached and the model. What an old answer rested on can always be seen. Recurring findings are marked "repeat" or "ongoing".	The permanent record of the answer and its evidence card.

LANGUAGE MODEL ERRORS

"What if the model talks nonsense?"

AgenticObjects does not assume the model will not make mistakes; it assumes it will, and checks every figure before it reaches a person. We measure the effect of this approach. The breakdown of the 845 answers delivered in our development and test environments up to 5 September 2026:

845	663	129	53
answers delivered	clean: no correction needed	a wrong figure corrected by the engine before it reached the user	a figure shown with an "unverified" mark

These numbers do not show a weakness; they show the mechanism working. The model's mistakes are visible in the record, and none of them reaches a user looking verified.

Cases caught during development that became permanent safeguards

- **Reversed direction:** the words describing the direction of a change were found reversed while the figures were right. Direction words are now produced by rule.
- **Wrong currency:** an amount was found counted as verified under the wrong currency sign. This is now refused before publishing, and so are totals the agent adds up itself from rows on screen.
- **An attempt to slip past the read check:** a special space character was found to get past the read-only check. Such forms are now refused, and similar attack patterns were added to a test set that runs continuously.
- **Silent failure:** a scheduled briefing was found to fail without telling anyone. The owner of the analysis is now notified, or the administrators if there is no owner.

> BOUNDARY The product explains **where** the figures in a sentence came from; it does not explain **why** the model phrased the sentence as it did. In scheduled discovery analyses the model chooses which findings to bring forward, so two runs can bring forward different findings. The figures in every finding go through the same check on every run.

DATA SECURITY

Where is your data, who sees it, what reaches the model?

Installation and data location

The application, the agent runtime, the query engine and the record database run as Docker containers on your organisation's Linux server. Your data warehouse stays on its own server; the product keeps no copy of the data and receives only aggregated results. Outbound access is needed only for model calls and e-mail notifications; no inbound connection is needed.

What reaches the model

The question asked, the structure of the data model (measure and dimension names) and a limited number of summary rows per query. Values of dimensions marked as personal data are masked as "A*** K***"; personal fields never appear in sample value lists.

Keeping data in-country

The model is chosen from a catalogue your administrator approves. In-country options in Türkiye: the Bulutistan language model service and EVREN. With either, model calls stay on infrastructure in Türkiye.

Read-only access, two layers

The product's gateway passes only read queries and refuses any statement it does not recognise. Cube sources likewise accept only read queries. On top of that, we require the warehouse connection to use a read-only account.

Row-level permissions

In a model with row permissions switched on, every user gets answers only from the rows they are permitted to see. A user with no permission defined sees no rows. Every run records its permission scope.

Connection secrets

Warehouse connection passwords are stored in the database encrypted with AES-256-GCM. The key that decrypts them lives outside the database, in the server's settings file; a database backup alone cannot open them.

AUDIT AND RECORDS

What is recorded, and who sees it?

Everything the product does is traced through linked records. From a figure in an answer you can go down to the query that produced it and to the model calls in that run.

Record	What it holds	Who sees it
Model calls	For every model call that produces an answer: the full text sent to the model and the full reply; the model, the amount of text processed, cost and time. Each call shows which user's question or which scheduled analysis it served.	With a separate permission. Every viewing is written to the audit log.
Query log	For every query that runs on the warehouse: its purpose, the connection it went to, waiting and execution times, and the number of rows returned. Moved to an archive after 90 days, never removed.	Analysis owners and administrators.
Text of the query that ran	The SQL query actually run on the database (DAX/MDX for cubes).	Only people separately granted this permission.
Run record	Every step of each analysis, the model used, the frozen settings and ceilings, the permission scope, and the result or error reason. A run can be exported as JSON; every export is recorded too.	Management console, by permission.
Audit log	Who changed what, and when: granting and revoking permissions, the "see all rows" permission, model choice, setting changes, opening the evidence layer.	Administrators; detailed sign-in information with a separate permission.
Spend	Model usage cost by user, model and agent; budget alerts.	Administrators.

PRINCIPLE Records can only be appended; no operation in the product removes or alters a record. All records are kept on your organisation's server. Retention periods are set together with you, in line with applicable data-protection law and your organisation's policies.

OPERATIONS AND RESILIENCE

How the system stays up, and how it is monitored

Health check

Every 5 minutes: services and their capacity, the work queue and any stuck analyses, the scheduler, the database, the gateway, release consistency, the time left on the secure-connection certificate, backup freshness and the date the data reaches. An e-mail goes out when an item turns faulty or recovers; alerts carry no data.

Backups and upgrades

The product database is backed up automatically every night, and each backup is verified by a trial restore. An upgrade does not start without a backup; the previous release can be restored if needed. Backing up the server itself stays within your organisation's existing arrangements.

Behaviour during outages

On a temporary error a step is retried 3 times; on a permanent error it is not retried. If the model provider cannot be reached, scheduled analyses wait up to 4 hours, then close and notify. A stopped worker process is detected and its work is handed to another.

Management console

The overview screen shows failed analyses, spend, stopped services, scheduling problems, connection test status and data freshness. The services screen shows status and lets an administrator restart a service.

What your team should keep an eye on

- Alert e-mails and monitoring screens
- The management console's overview screen, once a week
- The read-only account that connects to the warehouse
- Keeping the encryption key file separately, in your organisation's password vault
- Model spend and budget alerts

This note describes the current version of AgenticObjects. Questions: hello@agenticobjects.ai