AI agents for the work around the query and the figure

Agents that read your tables, your queries, your reports and your definitions, prepare the work, and stop for a named person. On infrastructure you control.

THE PROBLEM

An independent academic test in 2026 found the best AI systems right about one time in ten on questions from real private data warehouses. Of 135 public interface descriptions, 126 import, 8 are over the limit and 1 is refused; the 126 give 21,994 operations, 12,252 sent as changes.

HOW IT WORKS, IN THREE STEPS

STEP 01

Connect what the agent may read.

Metric definitions, the business glossary, data standards, table documentation and past analyses become knowledge sources, with each reader's permissions kept. The warehouse is queried through a connection that can only read, with the asker's permissions. The scheduler, the catalogue and the reporting tool are reached through the description of an interface you already run.

STEP 02

Decide where it stops.

Every tool call passes your policy. A tool that would write to a table, rerun a pipeline, publish a report or grant access waits for a named person.

STEP 03

Run it where the tables live.

In your data centre or your private cloud, on a model you serve.

THE RULE EVERY AGENT FOLLOWS

It prepares. A named person decides.

Thirty-five agents in nine areas

Questions and analysis	Data Analyst. Analysis Check. Analysis Request Intake. Analysis Write-up Draft
Reports and dashboards	Report Commentary Draft. Report Reconciliation. Report Request Triage. Report Usage Review
Metrics and definitions	Metric Definition Lookup. Metric Difference Explanation. Metric Definition Draft. Glossary Gap
Data quality	Data Quality Check Summary. Data Quality Rule Draft. Data Issue Triage. Reference Data Mismatch
Pipelines	Pipeline Failure Summary. Change Impact. Transformation Review Preparation. Schema Change Notice Draft
Catalogue and documentation	Table Documentation Draft. Data Finder. Lineage Summary. Documentation Gap
Access and privacy	Data Access Request Preparation. Personal Data Field Check. Access Review Preparation. Data Sharing Check
Models	Model Record Draft. Model Performance Summary. Experiment Summary. Training Data Record
The data office	Data Office Report Draft. Platform Cost Summary. Knowledge

Also in the pack: the full list of what each agent reads and where it stops, and the white paper “AI agents inside the data and analytics department”, with sources.

Seven questions to ask before an agent touches a table, a pipeline or a figure in a report

01 Where does it run, and where is the model?

02 Does it query as the person asking, with that person’s permissions?

03 Which tools write to a table, rerun a pipeline, publish a report or grant access, and does each wait for a person?

04 Are the read operations and the change operations separate tools?

05 Does every figure come with the query behind it, and who checks it before it is used?

06 Data holds text that outsiders wrote. Which tools could act on that text with no person?

07 What is recorded, and can we stop one agent or all of them at once?

WHERE THE PLATFORM RUNS TODAY

In production at Taiwan Cooperative Bank, air-gapped, inside the bank’s own perimeter.

MODON built its AI Centre of Excellence on Katonic inside the Kingdom of Saudi Arabia.

ePLDT, the ICT arm of the Philippines’ largest telco, selected Katonic to power Pilipinas AI, a sovereign AI cloud hosted in the Philippines.

Where to start

One team, one report, invented data first. An analysis check agent can start first, because it reads and changes nothing.

KATONIC.AI/SOLUTIONS/DATA-ANALYTICS