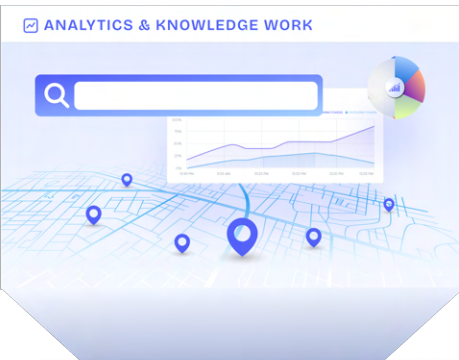


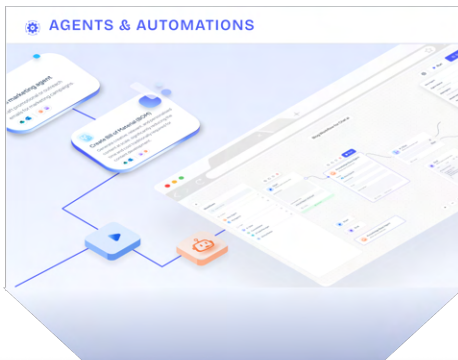
Trusted AI, Grounded In Your Enterprise Data

Synoptix AI is the platform for building, deploying and monitoring agentic AI on your own data: ontology-based, sovereign and defensible by design. No data movement. No hyperscale lock-in. Every answer is grounded, cited and auditable.

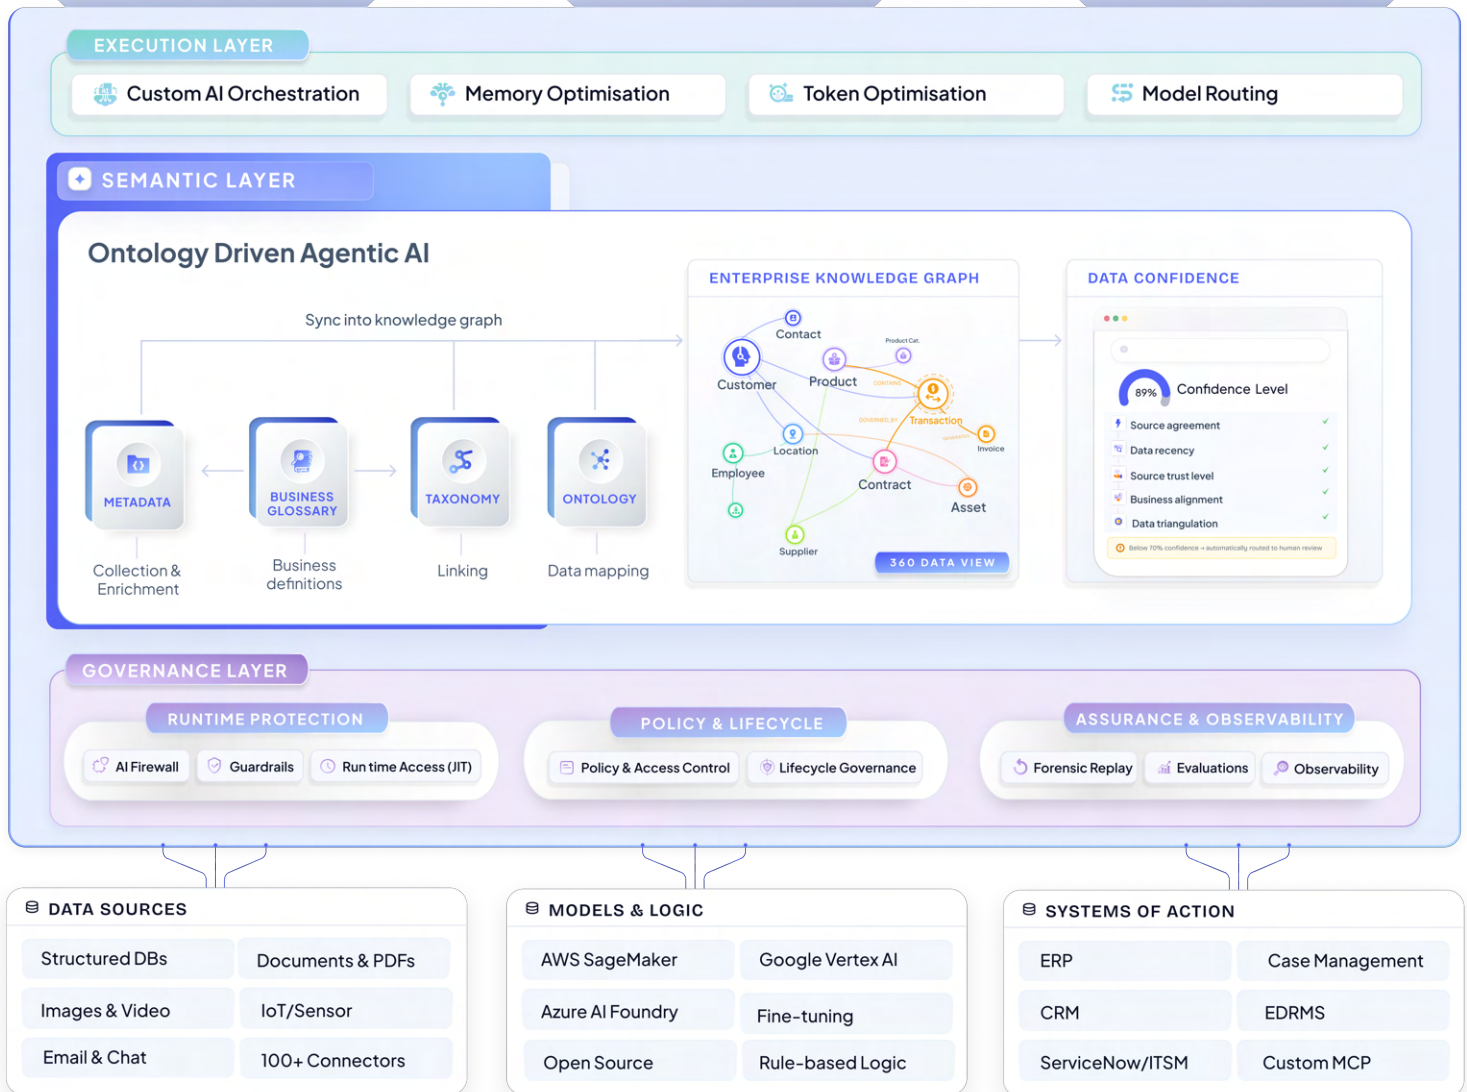
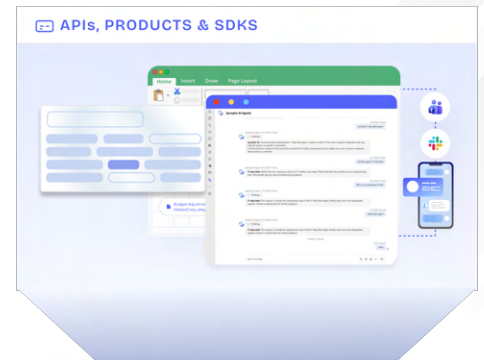
Cross-system Search Conversational AI Custom Tools & Skills



AI Workflow Builder No-Code Agent Builder Agent Library



Agent APIs Field Assistant Custom Application



ZERO Data Movement

Synoptix AI runs where your data lives. Zero ETL. Zero copies. Zero exfiltration risk

EVERY Answer Grounded

Citations and lineage on every response. Synoptix AI retrieves from your governed data

BUILT-IN Compliance

Row-level security, audit trail, IRAP and Essential 8 enforced at every query

WEEKS To Production

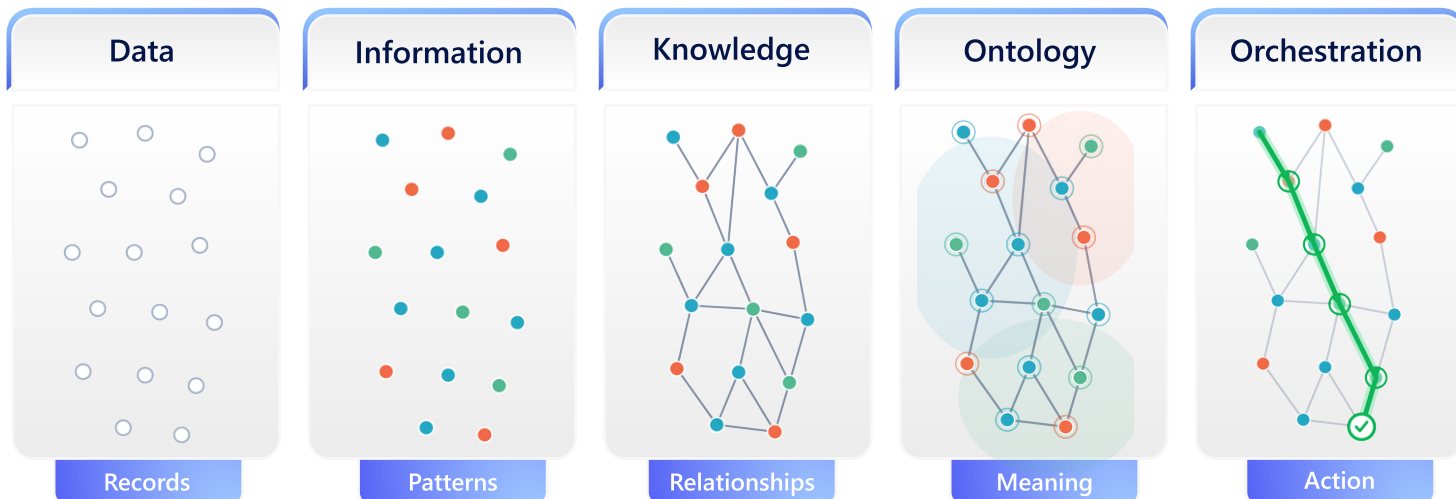
Joint reference architectures and accelerators that reach production in weeks, not months

ALWAYS Evolving

New models, new data, new threats. Swap LLMs and return agents without rebuilding

Most AI Stops At Data. Synoptix AI Starts At Meaning

The ontology turns connected data into governed knowledge your agents can reason over and act on.



Taxonomy Library

Your Industry, Already Modelled. Live In Weeks.

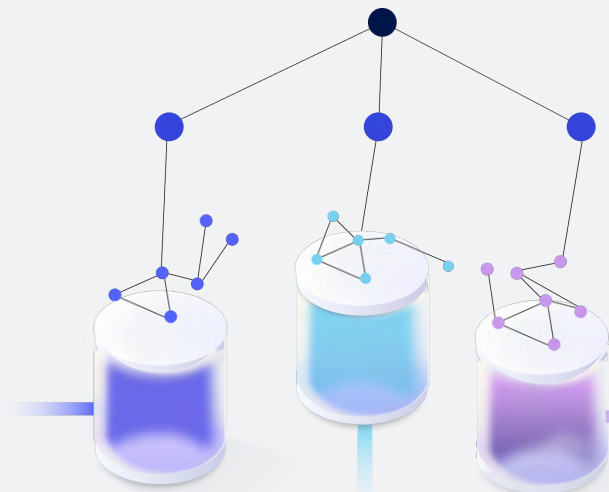
Start from a library of pre-configured taxonomies covering dozens of regulated industries, each backed by real business definitions.

What's a Taxonomy?

A set of concepts organised into a hierarchical structure covering a topical domain. It organises and labels information, data and objects so they are easier to find, analyse and use.

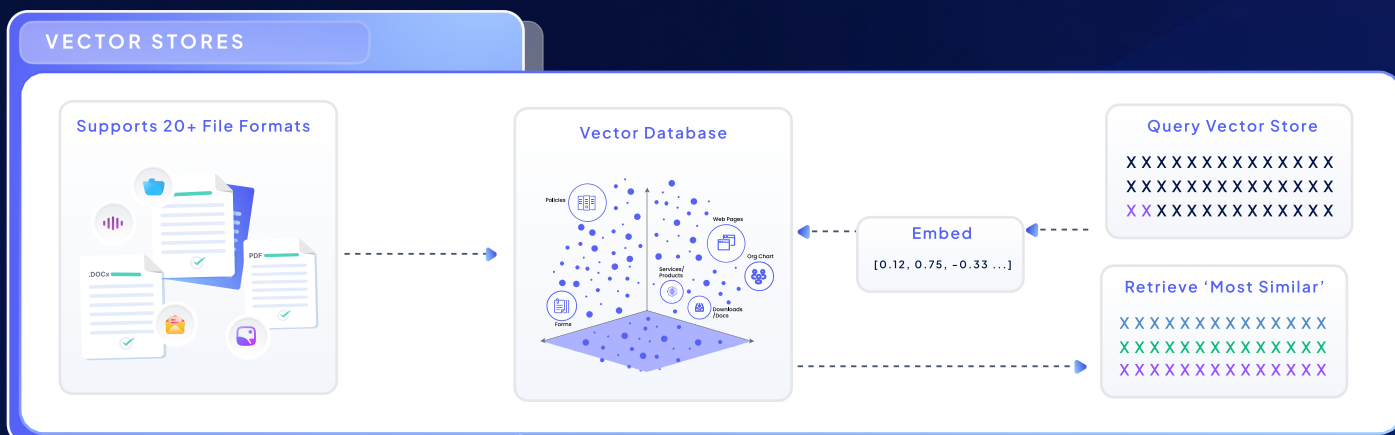
Pre-built for regulated industries

- Government
- Banking
- Healthcare
- Telecom
- Insurance
- & More



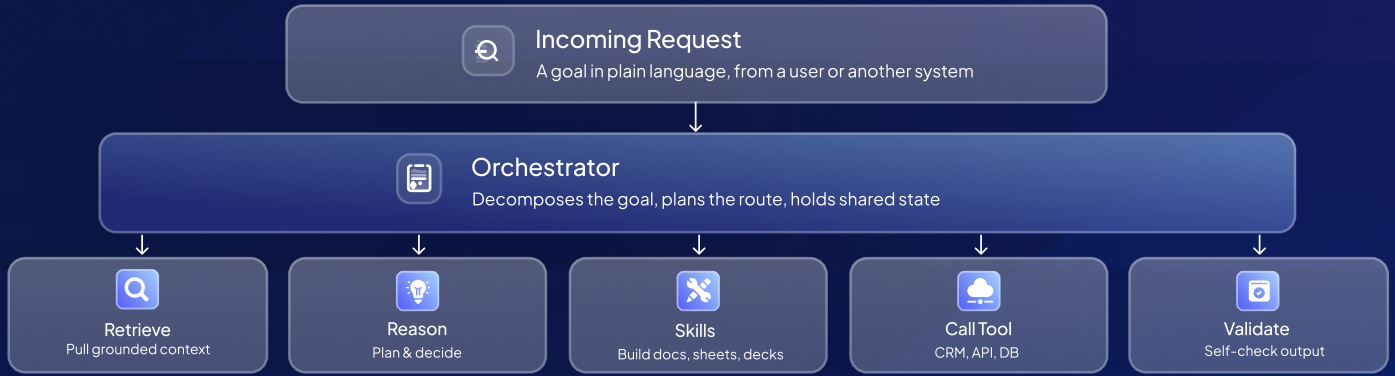
Your Documents, Searchable By Meaning

Ask in plain language; get back the right passage from thousands of documents ranked by meaning, not keywords



No Single Flow Fits Every Use Case

Every use case carries different risk, latency and cost. Route them all through one flow on one flagship-priced model and you over-pay, over-expose and under-perform.



Purpose-built beats one-size-fits-all

COST EFFICIENCY

Route each use case to the right-sized model

Flagship reasoning only where it earns its price.

One model for everything \$5.00 / 1M Token

Right-sized routing \$1.40 / 1M Token

Upto

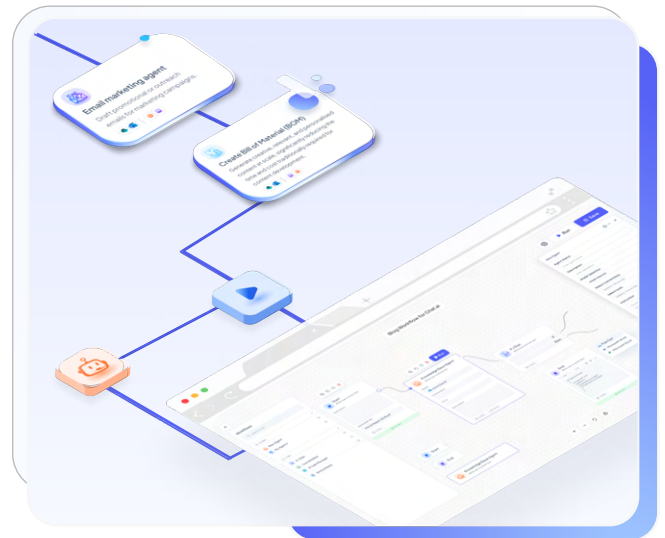
25x

Cheaper on simple task
3.5x lower across a realistic mix

Cost figures illustrative. Per-task figure reflects published input-token list prices for a flagship versus a small model as at June 2026. Blended-cost figure based on Route LLM (Ong et al., UC Berkeley, ICLR 2025), reporting up to 3.66x cost savings at 95% of flagship quality.

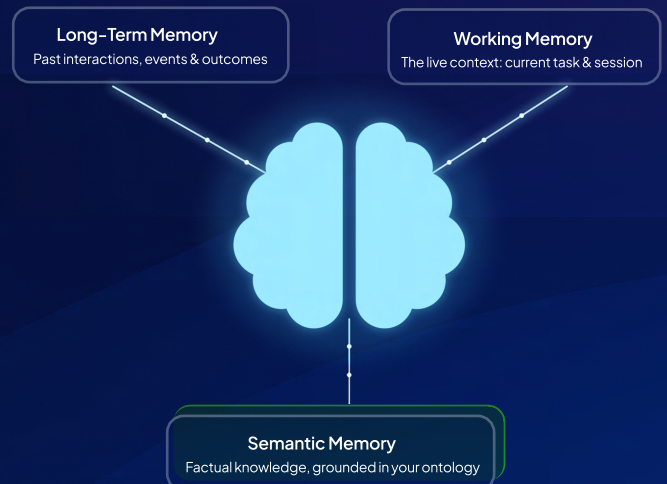
Ship Agents Without Engineering Tickets

- Build, test and ship multi-agent systems, with no engineering tickets required.
- Schedule tasks to run on their own, on the cadence the work demands.
- Save agents and workflows to a shared library, then reuse and improve them across teams.



Bigger Prompts Aren't The Answer. Governed Memory Is.

Large context windows cost more and still forget when the session ends. Synoptix gives agents layered memory across working, long-term and semantic, so recall is cheap, relevant and grounded in your ontology.



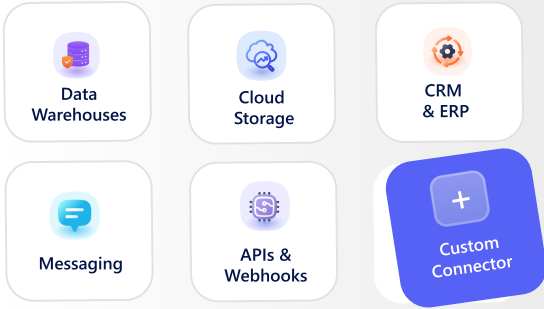
- ✓ Persistent recall
- ✓ Lower token cost
- ✓ Grounded retrieval
- ✓ Governed retention

Plugs Into What You Already Run

Connections

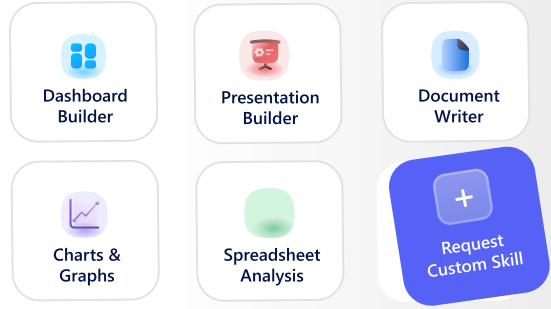
100+

pre-built connectors integrate with the systems you already use. For anything bespoke, we build custom connectors so your data stays where it is.



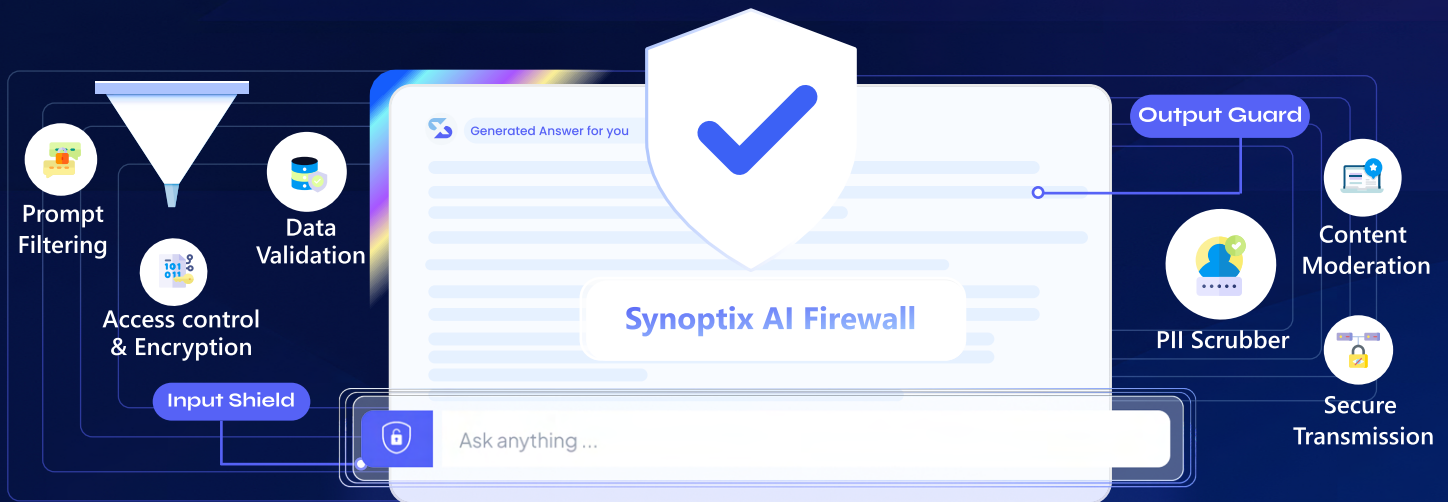
Skills

Information alone is not enough. Skills empower your team to take action with the right tools to complete any task.



Runtime Protection

Every prompt and response inspected, redacted and validated



Two Choices, One Governed Platform.

1 **DEPLOYMENT:** Where the Synoptix platform runs

Private Cloud

YOUR TENANT

Run Synoptix inside your existing cloud, under your own security and network policies.

DATA STAYS IN YOUR TENANT

Dedicated Cloud

FULLY MANAGED

Dedicated, isolated infrastructure managed by Synoptix, for enterprises that need separation without running it themselves.

ISOLATED TO YOU

On-Premise

SOVEREIGN AI

Deploy entirely within your own data centre for full control over security, compliance and data residency.

AIR-GAPPED CAPABLE

2 **MODEL & COMPUTE HOSTING:** Where your models and inference run, in your own cloud account

Azure AI Foundry

MICROSOFT AZURE

AWS SageMaker

AMAZON AWS

Google Vertex AI

GOOGLE CLOUD

Quality Measured In Production, Not In A One-Off Test

Once your agents are live, Synoptix keeps watching. Continuous evaluation against your own data, full-stack observability of every run, and forensic replay of any decision, so you always know what your AI did, and why.

Evaluations that never stop

Every response is scored against your ontology and your sources, not a benchmark someone else picked. Quality is measured continuously in production, not in a one-off test before launch.



See the whole run, not just the result

Full-stack observability across every agent, prompt, tool call, and token surfaced in one place, with alerts that reach you before your users feel a thing.



End-to-end traces:

Every prompt, tool call, retrieval, and token captured in one timeline you can step through.



Live cost & latency:

Spend, latency, and throughput tracked per agent and per workflow, live.



Quality alerts:

Drift, bias, and accuracy regressions monitored continuously and surfaced for review

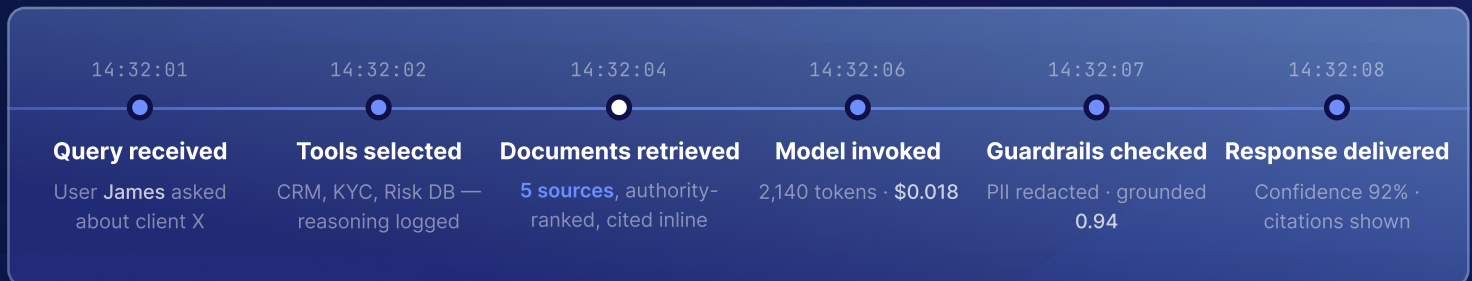


Safety on every step:

PII redaction, toxicity, and grounding checked at every hop, not just at the output.

Replay Any Agent. Any Day. Any Decision.

When the regulator asks "What did your AI do for that customer on 14 March, and why?" you answer in minutes, not in a forensic discovery project. Every run is reconstructable, step by step, exactly as it happened.



Every Agent has Rules and a Lifecycle



Policy & Access Control

Every person and agent sees only what their role permits, enforced on each query, not trusted to whoever built the agent.



Lifecycle Governance

Every agent and model is approved before production, with a full change history, so you always know what is live and who signed off.



Just-in-time Access

Access granted for one task, then revoked, so no standing credentials sit idle as a breach target.



Contact Us : hello@synoptix.ai