

Your AI Should Remember, Follow the Rules, and Show Its Work.

Cognous is a control and continuity layer for enterprise AI: it preserves the context that matters, applies your policies, and records the evidence behind every result.

ORDINARY AI

Starts over with every request

It resends old context, forgets earlier decisions, buries policy inside prompts, and leaves little proof of why a result was allowed.



WITH COGNOUS

Governed work that carries forward

It keeps approved context, checks each result against your rules before release, and keeps a record your team can review later.

WHY THIS MATTERS NOW

AI work is getting longer. Teams now rely on it for research, planning, customer operations and agents that run for weeks, and every new step costs more when the model has to reread the entire history.

It is also harder to prove what the AI relied on and why an answer was permitted. Cognous addresses both problems. It does not make models infallible; it makes them easier to govern, verify, afford and move between providers.

HOW IT WORKS

BETWEEN YOUR MODELS AND YOUR BUSINESS

Your applications and agents sit above Cognous, while models, tools and data sit below it. Cognous gives the model only the context it needs, evaluates each draft against your policies, and then releases it, sends it back for revision, or refuses it.

A REQUEST
ARRIVES



COGNOUS
CHECKS AND
RECORDS IT



YOU RECEIVE A
RESULT YOU CAN
USE

WHAT YOU GET

- **Lower operating costs**, because Cognous carries forward only the context that matters instead of replaying the whole history.
- **Governance that travels with the work**, since identity, policy and source status stay attached to each task.
- **A record for every result**, showing what was used, what was decided, and why.
- **Freedom to change models**, because your rules and memory remain in place when providers change.
- **Controlled autonomy**, where AI can propose actions but your policies decide what actually runs.

EARLY TEST RESULTS

Reduction in input sent to the model in preliminary synthetic testing:

53.9%
at 100 turns

75.5%
at 200 turns

89.9%
at 500 turns

These figures measure input context only; total cost savings are still being validated.

LET'S TALK

Start with one real workflow that runs over time. Together we will measure tokens, quality, rework and governance side by side, and you expand only when the results hold up.

FOR YOUR TECHNOLOGY TEAM

ONE API

Integrate through a versioned REST API or SDKs for TypeScript and Python.

ANY MODEL

Switch approved providers without rewriting governance logic.

EVIDENCE ON FILE

Every result links to a record you can retrieve later.

FAILS SAFE

If required evidence cannot be saved, the result is not released.

ABOUT COGNOUS

Cognous builds governance-native infrastructure for enterprise AI. Our platform, ISS, keeps AI work governed, traceable and portable, so organizations can rely on what AI delivers.

COGNOUS