

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

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



AI work is getting longer — and harder to govern

Teams now rely on AI for research, planning, customer operations and agents that run for weeks. Every new step costs more when the model has to reread the entire history, and it is harder to prove what the AI relied on and why an answer was allowed.

ORDINARY AI

Starts over with every request

- ✗ Resends old context on every turn
- ✗ Forgets earlier decisions
- ✗ Buries policy inside prompts
- ✗ Leaves little proof of why a result was allowed




WITH COGNOUS

Governed work that carries forward

- ✓ Keeps approved context, not the whole history
- ✓ Carries decisions and source status forward
- ✓ Checks each result against your rules before release
- ✓ Keeps a record your team can review later

BOTTOM LINE Cognous does not make models infallible; it makes them easier to govern, verify, afford and move between providers.

One control layer between your models and your business



WHAT HAPPENS IN THE MIDDLE

- 1 Gives the model only the context it needs
- 2 Evaluates each draft against your policies
- 3 Releases it, sends it back for revision, or refuses it
- 4 Records the evidence and lineage for review

Seven parts of one governed platform



Governed AI Runtime

Applies identity, policy, state, admissibility and evidence requirements around model execution.



Idea Space Substrate (ISS)

Carries task context, claims, uncertainty, source status, lineage and constraints across work.



Memory & Continuity

Persistent, identity-scoped state so long-running work resumes without replaying the whole history.



Governed Dispatch

Approves, requires revision, refuses or fails — rather than treating every output as deliverable.



Evidence & Reconstruction

Append-only records, lineage and negative outcomes that show what happened and why.



Clearframe Inspection

Inspects evidence, state, lineage and reliance conditions — separate from authority to certify or execute.



APIs & SDKs

A stable integration surface, so governance logic is not tied to a single model vendor.

What leadership gets from governed AI



Lower long-run AI cost

Carries forward bounded state instead of replaying entire histories. In testing, the input sent to the model fell by up to 89.9% on long conversations.



Governance that travels

Identity, policy, source status and restrictions live in state, not scattered prompts.



Faster adoption

One runtime and evidence layer, so each team need not build its own.



Model portability

Change providers behind a controlled interface; governance and state stay put.



Less ambiguity

Separates output, state, inspection, certification, acceptance and refusal.



Stronger audit & response

A durable record of decisions, transitions, denials and source relationships.




Controlled autonomy


AI can propose and reason; execution stays bounded by authority and policy.

Measured: far less input sent to the model

Cognous keeps the last 20 messages word for word and a short summary of older history, so the cost of each turn stays roughly flat as a conversation grows.



Reduction in input tokens versus sending the full history, published benchmark, July 2026.



SCOPE Input tokens only, on synthetic conversations. Total savings, including output and rework, are being measured on live providers now.

How a governed request flows

Downstream applications receive a governance verdict and a correlation ID that retrieves the full evidence bundle.



GOVERNED MODE Outputs are buffered and evaluated before anyone sees them, so known blocking violations are never released as qualified output.

Core components of the platform



Identity & authority

- API keys and JWT, tenant and workspace identity
- Capability or credentials are never treated as authority



Governed state & policy

- Structured state for claims, sources, uncertainty and lineage
- Enterprise rules compiled into runtime constraints
- Load management contracts reasoning under pressure



Runtime control

- Readiness and admissibility decided at run time
- Governed emission buffers output before release
- One exclusive, reversible write path for canonical state



Evidence & proof

- Append-only evidence, audit references and negative records
- Signed evidence and certification where required
- Metering for cost per accepted deliverable



Integration

- Versioned REST API
- TypeScript and Python SDKs
- Model portability behind the platform boundary

DESIGN PRINCIPLE Cognous sits across applications and models rather than being rebuilt inside each one.

Principles built into the platform



Fail closed

No governed-success result is returned if required evidence cannot be saved.



Append-only evidence

Evidence records are create-only and immutable on the governed path.



Chain integrity

Lineage chains detect material change and truncated or replaced history.



No hidden reasoning

Governance relies on inspectable state, policy and decisions — not private chain-of-thought.



Controlled configuration

Material assurance settings are versioned, authorized and attested.



Separate roles

Inspection, certification and execution authority remain distinct.



Negative outcomes kept

Refusals, blocks, revisions and revocations stay in the record.



Claims match evidence

Product claims track implemented controls and executed tests, not diagrams.

Designed to sit across applications and models



Integrate once

Use the versioned REST API or TypeScript/Python SDKs instead of rebuilding governance in each application.



Separate by workspace

Workspace-scoped policy and state separate tenants, business domains and applications.



Swap models freely

Provider configuration stays behind the platform boundary, so changing models does not mean rewriting governance.



Connect enterprise systems

Identity, data sources, logging, key management and policy or security systems, as assurance requires.



Buffer when it matters

Use governed buffered mode for reliance-bearing work; treat streaming as a weaker posture.




Trace every output

Put evidence correlation IDs in application logs, so any AI output can be joined back to its evidence.

STAGED ASSURANCE Ordinary governed use, higher-assurance use and certified use carry different proof obligations, so no request pays maximum overhead by default.

A measured path from one workflow to many



COMMERCIAL OFFERS

Platform API Core governed runtime and integration surface Subscription / usage	ISS Runtime & State State, continuity, context planning, lineage Included; advanced by tier	Governance & Evidence Policy, admissibility, evidence bundles, proof Higher-assurance tiers	Clearframe Inspection of governance records and evidence Add-on or bundled	Implementation Architecture, policy setup, benchmarks, testing Professional services	Assurance Extensions Signing, attestation, identity and security Optional package
---	---	---	--	--	---



THE BOTTOM LINE

Govern the work, not just the model.

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

LET'S TALK

HONEST BOUNDARIES

Cognous provides bounded process assurance — not a guarantee of truth, compliance or outcomes.

- Compression is not certification.
- Logging is not authorization.
- Inspection is not certification.
- A completed workflow is not automatically reliance-bearing.
- A model's capability is not an execution grant.
- Fewer tokens only count if quality holds.