

Built on LibreChat · Enterprise-Grade

Enterprise Gen AI chat, without the **shadow** **AI risk.**

Obichat gives every employee a governed, policy-enforced AI workspace — powered by LibreChat, the open-source alternative to ChatGPT — so your teams get productive while security and legal stay in control.

- *Zero shadow AI. Full audit trail.*
- *Add-on to Obiguard's Governance AI*

Bring your own

MODELS SUPPORTED

Redaction & block-lists

GUARDRAILS

Cloud / On-prem

DEPLOYMENT OPTIONS

Built in

WORKSPACE POLICY

Shadow AI is your next data breach.

01

Shadow AI is already inside your org

Employees use personal ChatGPT accounts, browser extensions, and third-party AI tools — sending customer PII, source code, and deal terms to unknown endpoints every day.

02

No visibility, no audit trail

When a regulator asks “what did your staff send to an AI last quarter?”, consumer tools can't answer. Obichat keeps every session in your own tenant, fully logged.

03

Policy is impossible to enforce at scale

Blanket AI bans kill productivity and drive shadow usage underground. You need allow-lists, department-level policies, and real-time guardrails — not a firewall rule.

Everything security needs. Everything teams want.

GOVERNED WORKSPACES

One AI front-door, organized by workspace

- Per-workspace AI provider connections
- Per-workspace model allow-lists
- Workspace visibility controls (private / team)
- Admin-only access to workspace settings

REAL-TIME GUARDRAILS

Policy enforced before the prompt leaves your network

- PII / PCI redaction
- Jailbreak & prompt-injection detection
- Keyword and regex block-lists
- Policy sets assigned per workspace

AUDIT & COMPLIANCE

A record of activity for every workspace

- Per-workspace audit log, exportable to CSV
- Conversation and event history
- Framework tagging (NIST AI RMF, ISO/IEC 42001)
- Redacted-message visibility for reviewers

IDENTITY & ACCESS

Fits the identity setup you already have

- SAML 2.0 / OIDC sign-in
- Admin-gated workspace settings
- Private or team-scoped workspaces
- Segregated by org and project

Obichat, from the inside.

A real session — governance enforced inside the chat itself, not bolted on after.

Jailbreak, caught instantly


A “Do Anything Now” prompt-injection attempt is blocked by the workspace's Policy Set — rule control-jailbreak, confidence score 0.87 — before it ever reaches GPT-4.1.

Sensitive requests, declined

A follow-up message referencing payment card details is refused under the same governed session — the model itself won't act on it.

The interface your team already knows

Built on LibreChat, so employees get a familiar ChatGPT-style workspace, with governance running invisibly underneath.



Connect any model your enterprise already approves.

Obiguard doesn't lock you into one provider. Point it at the models your organization has already vetted and approved — Obiguard governs access; it doesn't gatekeep the model.

Microsoft Azure OpenAI

Anthropic Claude

Google Gemini

AWS Bedrock

Mistral AI

Llama 3

Any OpenAI-compatible endpoint works, including self-hosted models via vLLM or Ollama.

Built with governance in mind.

Obichat's controls are mapped to the AI governance frameworks your team is already working toward.

NIST AI RMF

Govern, Map, Measure, Manage

ISO/IEC 42001

AI management system

OWASP LLM Top 10

Mitigate top vulnerabilities

COMPLIANCE

Compliance Frameworks

See how your existing policies, controls, and activity align with recognized AI governance frameworks.

 This is a self-assessment, not a certification.

Coverage is estimated from your platform activity — policies, controls, criteria, use cases, and violations. With the exception of ISO/IEC 42001, none of these frameworks have a certification body attached to them; this view is meant to help you prepare for an audit, not substitute for one.

NIST AI RMF

ISO/IEC 42001

OWASP LLM Top 10

Voluntary guidance across four functions — Govern, Map, Measure, Manage. Not a certifiable standard; this view reflects platform-evidenced alignment, not an audited attestation.

 Export CSV



GOVERN

Governance policies and processes are in place Reference

Policy Sets carry a draft → pending → active lifecycle with rule-level review.

EVIDENCE

3 active Policy Set(s) with rule-level review history

12 approved Control(s) backing those policies

Covered

Accountability structures and roles are defined Reference

Controls enforce segregation of duties — a submitter cannot approve their own control.

EVIDENCE

Self-approval is blocked on all Control approvals

0 Access Key(s) scoped and active

Covered

Third-party and value-chain risks are managed Reference

Not covered

Actual Compliance Frameworks view: coverage estimated automatically from your policies, controls, and violations

Common questions.

What is LibreChat, and why does Obichat use it?

LibreChat is the leading open-source alternative to ChatGPT, with a familiar interface and support for 20+ AI providers. Obichat is built on LibreChat so you get a battle-tested UX — plus Obiguard's enterprise security and governance layer on top.

Can we run Obichat in our own cloud account?

Yes. Obichat supports private deployment in AWS, Azure, and GCP via Helm chart or Terraform. Your prompts, responses, and files never leave your infrastructure.

How does Obichat prevent employees from bypassing it?

Obichat pairs with Obiguard's network-layer proxy. You can block direct egress to consumer AI endpoints at the firewall level, so Obichat becomes the only approved on-ramp.

Does Obichat support our existing models or fine-tunes?

Yes — any OpenAI-compatible endpoint works, including Azure OpenAI, AWS Bedrock, Vertex AI, or self-hosted models via vLLM or Ollama.

When is Obichat generally available?

Obichat is available now. Launch the app to sign in, or talk to sales to set up your organization.

NOW AVAILABLE

One policy layer. Every AI surface in your company.

Governance AI and Obichat share the same Gateway, Guardian, and Policy Sets — so coverage never has a gap between your APIs and your employees.

GOVERNANCE AI

Every API, agent, and model call

- Gateway + Guardian enforce policy in real time
- Policy Sets, allow-lists, and jailbreak detection
- Mapped to NIST AI RMF, ISO 42001, OWASP LLM Top 10

+

OBICHAT

Every employee's everyday AI chat

- A governed, ChatGPT-style workspace — built on LibreChat
- Same Policy Sets — no separate stack to stand up
- SSO, per-workspace audit log, exportable evidence

= one governed AI surface, one policy layer, one audit trail — across agents and employees alike.

Email: hello@obiguard.com →

<https://obiguard.ai> ↗