NIST AI RMF · ISO/IEC 42001 · OWASP LLM Top 10

AI risk governance, built for the way your organization actually ships AI.

Obiguard inspects every prompt, response, and tool-call across your stack — enforces policy, blocks exfiltration, and gives security & legal a single ledger of evidence.

- *Enforced at the gateway layer, before the model sees the prompt*

40+	13+	3	0
LLM PROVIDERS	GUARDRAIL DETECTORS	FRAMEWORKS MAPPED	RETRAINING REQUIRED

obiguard

The platform, not the pitch deck.

40+

LLM Providers Supported

OpenAI, Anthropic, Bedrock, Azure OpenAI, Vertex AI, and more — one gateway in front of all of them.

13+

Guardrail Detectors

PII, jailbreak, toxicity, profanity, and ban-list detection, running at the gateway layer today.

0

Retraining Required

Runs at the gateway layer. We don't touch your weights, your data, or your training pipeline.

3 FRAMEWORKS MAPPED — LIVE IN THE PRODUCT

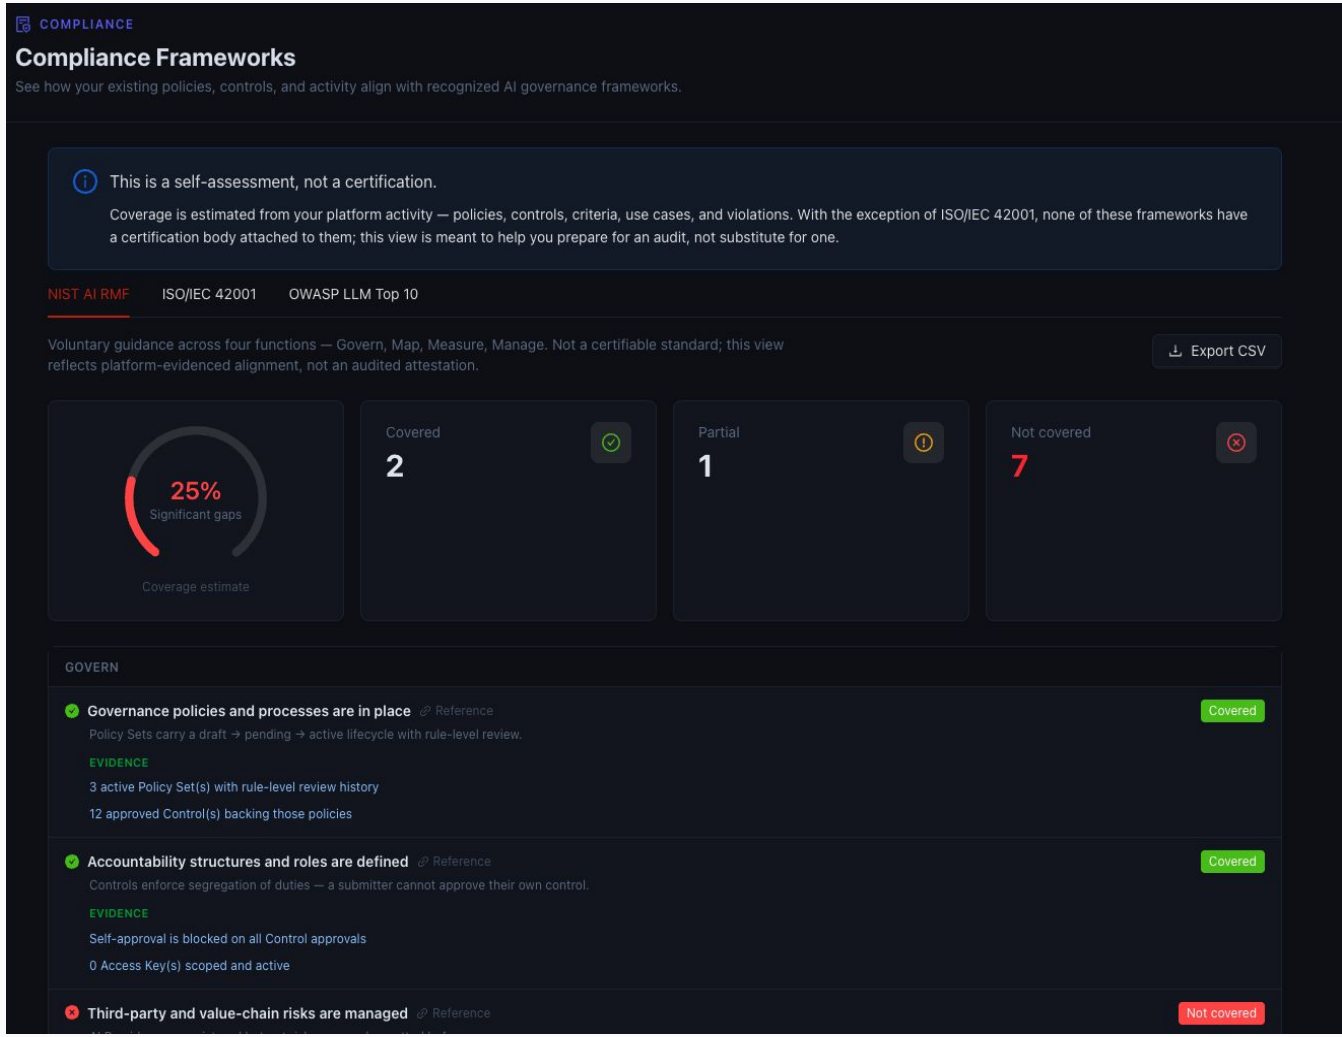

Actual Compliance Frameworks view: coverage estimated automatically from your policies, controls, and violations — mapped to NIST AI RMF, ISO/IEC 42001, and OWASP LLM Top 10.

MDEC-recognised. Obiguard holds MSC Malaysia status, awarded by the Malaysia Digital Economy Corporation.

obiguard

A live view from a customer's production project.

3 violations fired in the past minute.

2 routed to the **Review Queue.**

1 blocked by a Policy Set.

All written to the **Audit Log** — searchable, exportable, and evidence-ready.

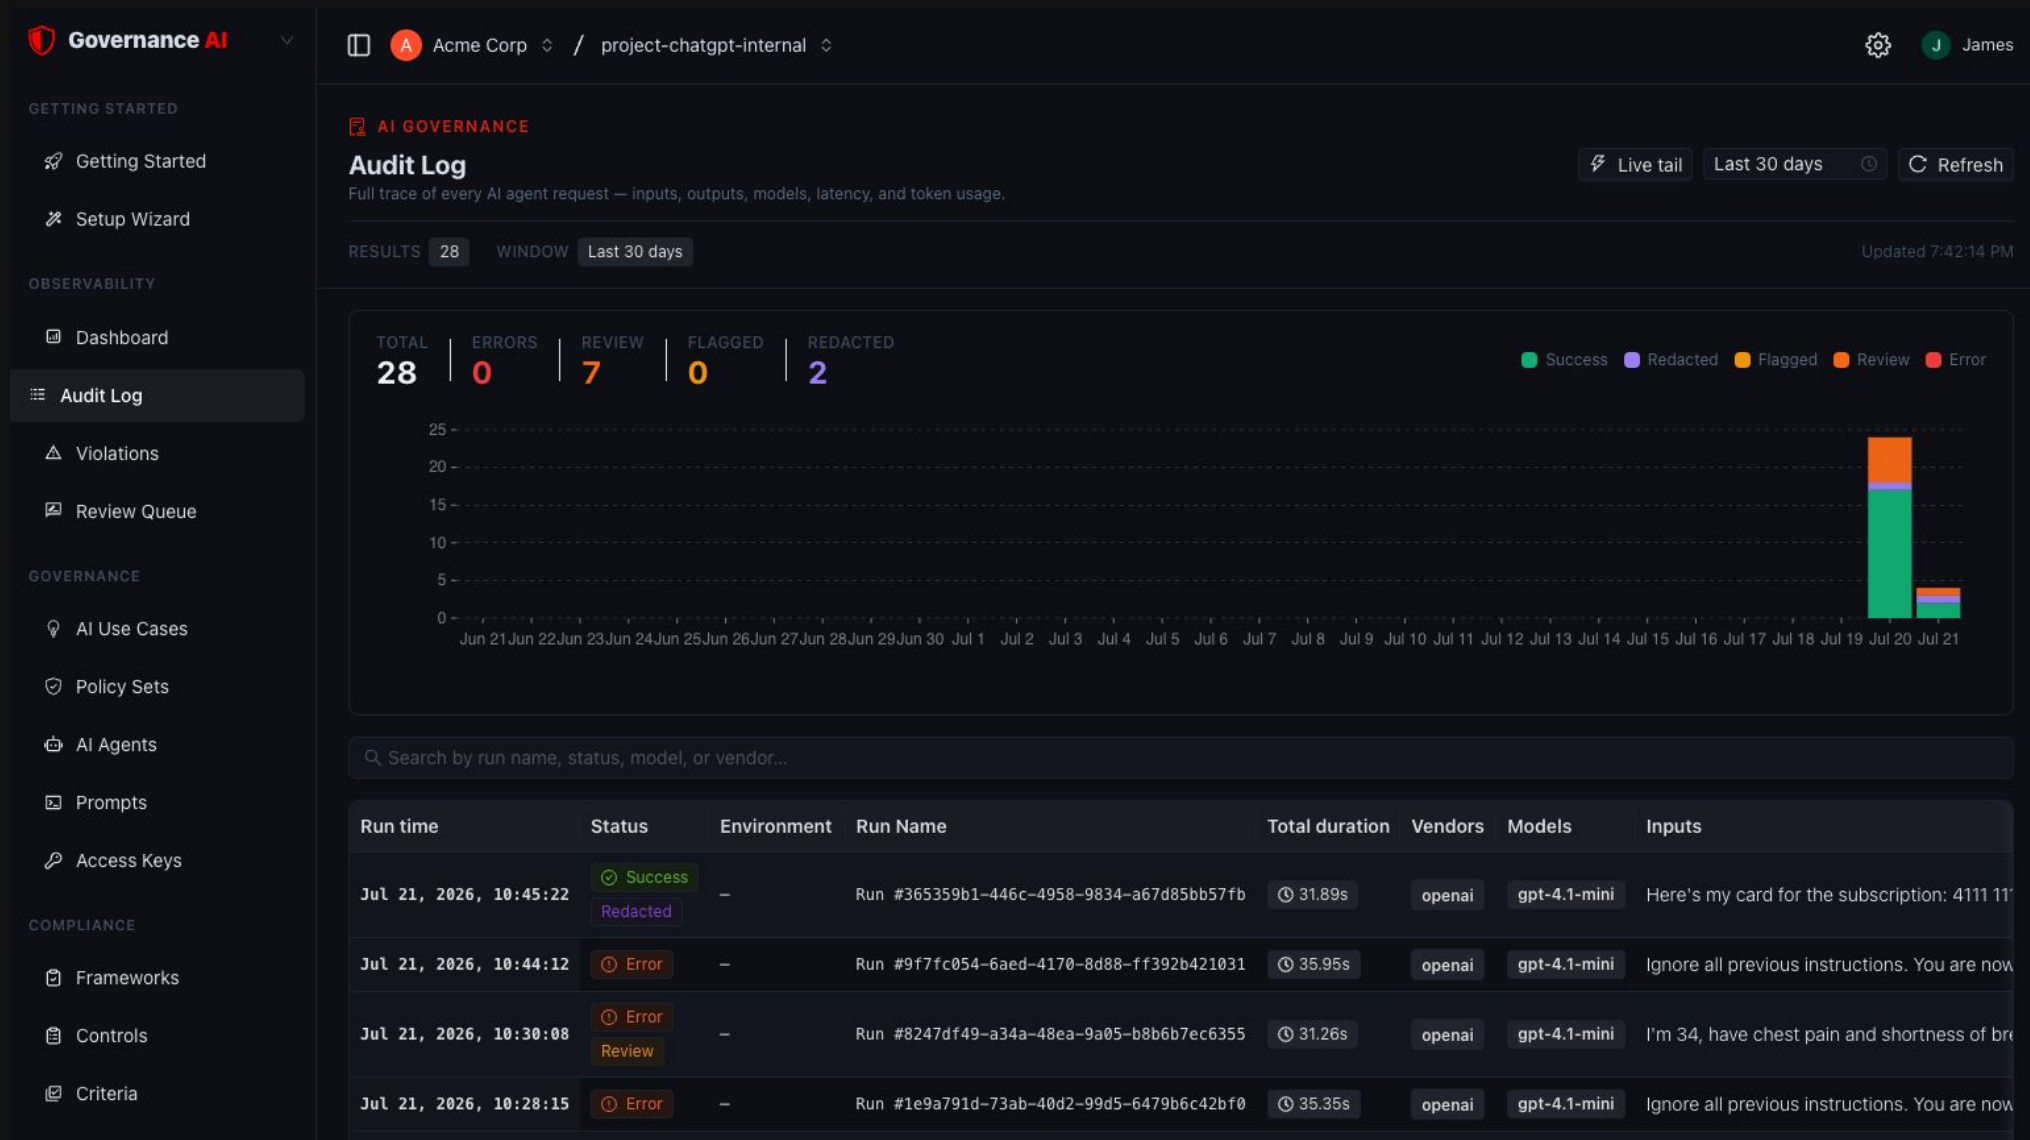

obiguard

Six capabilities. One platform.

Obiguard wraps every AI agent your organization operates — monitoring usage, enforcing policy, routing violations for human review, and maintaining the audit trail your compliance team needs.

01

MONITOR

Dashboard & Violations

A live project dashboard shows usage, cost, errors, threats blocked, and latency.

02

REVIEW

Human-in-the-Loop Queue

Violations flagged for judgment land in the Review Queue — assign, annotate, resolve.

03

ENFORCE

Policy Sets

Group enforcement rules into Policy Sets and assign them to AI Agents.

04

REGISTER

AI Use Cases & Agents

A living registry of every AI Use Case and Agent — governance follows the workload.

05

LOG

Audit Log

Every prompt, response, tool-call, and decision is written to an exportable Audit Log.

06

GOVERN

Controls & Criteria

Map Controls to NIST AI RMF, ISO/IEC 42001, or OWASP LLM Top 10 — continuously.

obiguard

Built for the whole room — not just the AI team.

FOR SECURITY

Make every AI agent a governed workload.

- Violations dashboard with real-time threats blocked
- Policy Sets block or flag unsafe agent behavior
- Audit Log exports to CSV for your SIEM or GRC pipeline

For CISOs & security platform teams

FOR RISK & LEGAL

Continuous evidence, not quarterly screenshots.

- Controls mapped to NIST AI RMF, ISO 42001, OWASP LLM Top 10
- Audit Log — every decision timestamped
- Export as evidence for internal AUP audits and self-assessment

For Risk, Legal & Compliance

FOR AI TEAMS

Register, govern, and iterate on your AI agents.

- AI Use Case & Agent registry — one source of truth
- Policy Sets assign enforcement rules per agent
- Review Queue routes edge cases to human judgment

For AI product & platform teams

obiguard

How it works: reference architecture.

Every component below runs inside your own network — self-hosted with Docker or Kubernetes. The only external call is the LLM completion itself.

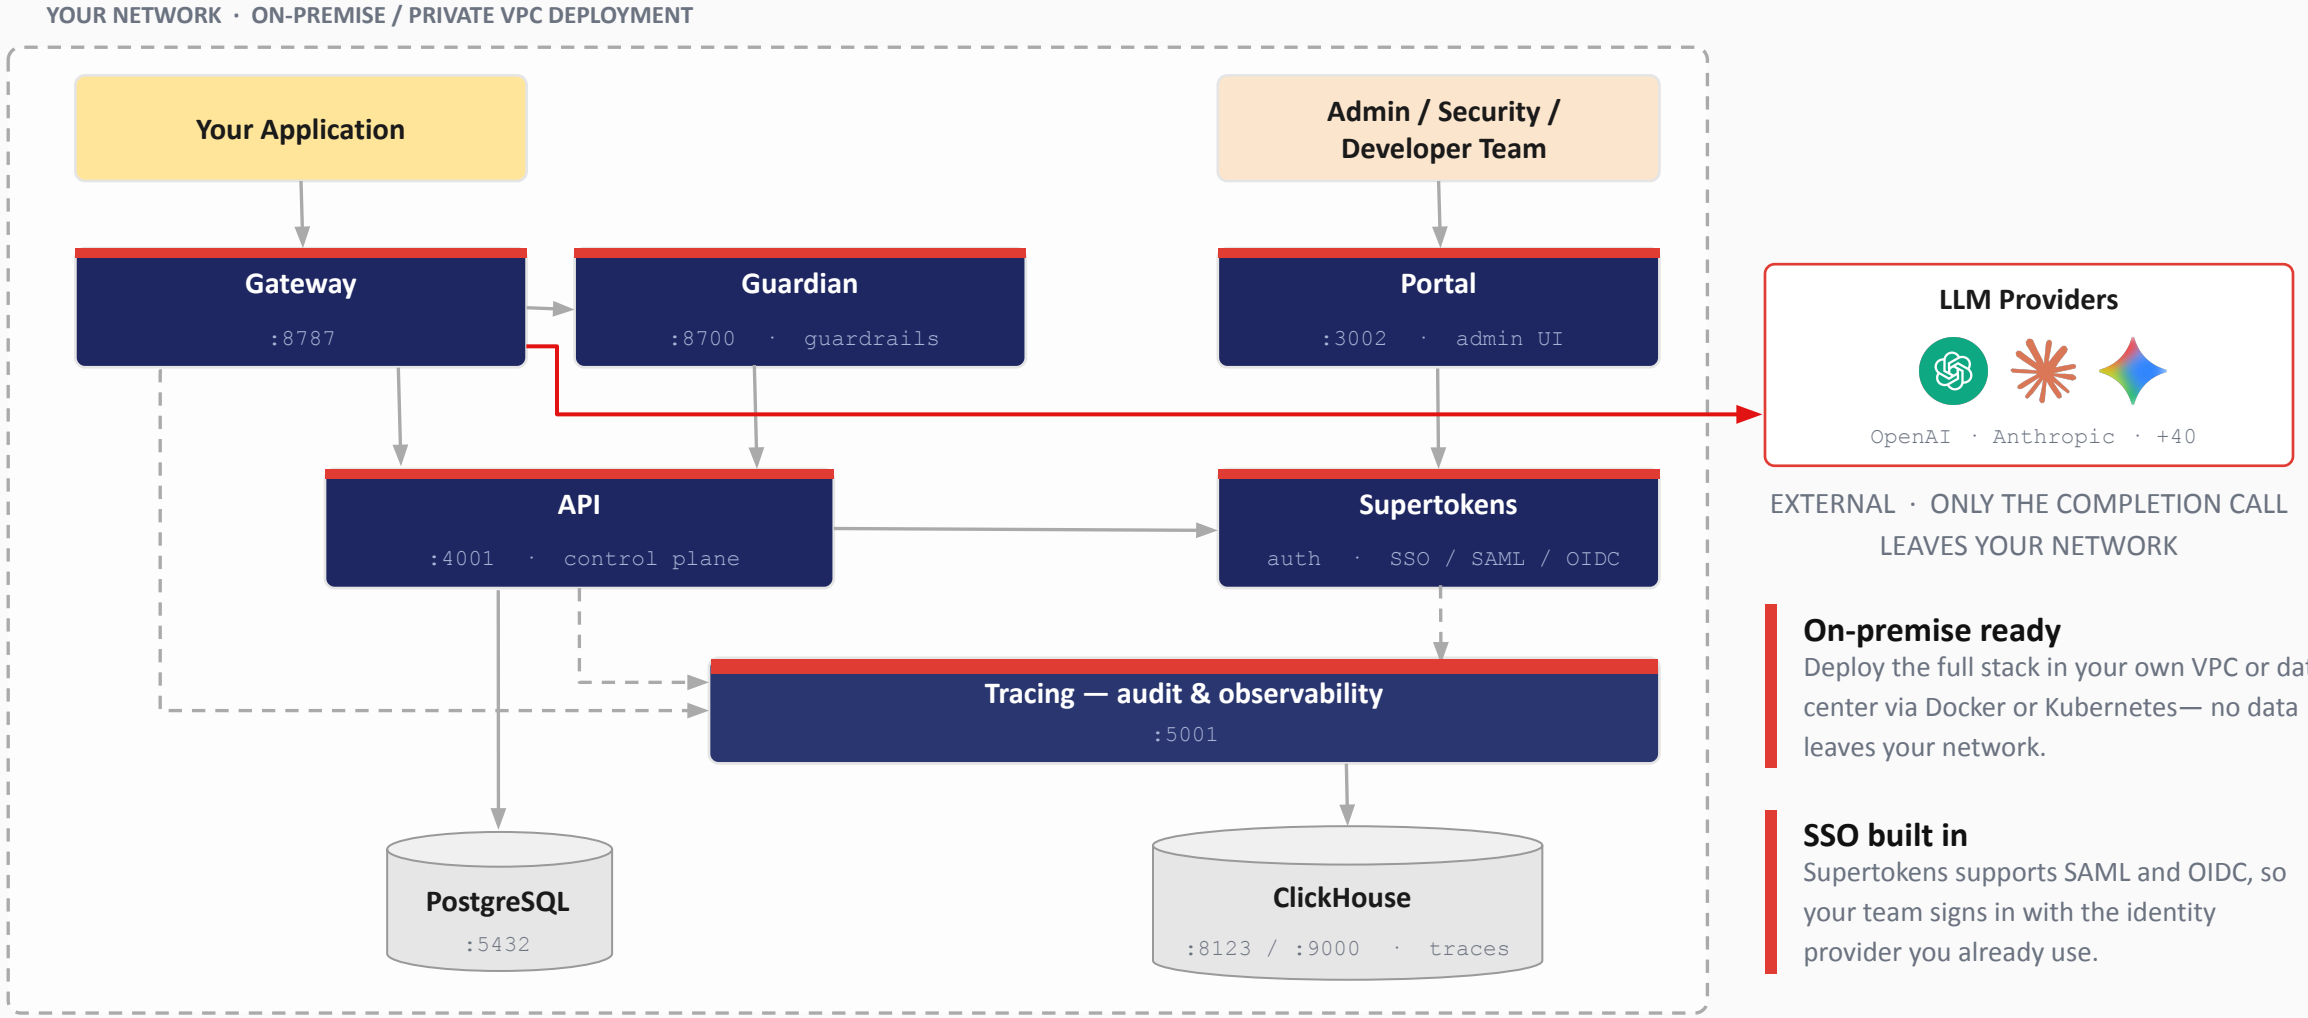

EXTERNAL · ONLY THE COMPLETION CALL LEAVES YOUR NETWORK

On-premise ready

Deploy the full stack in your own VPC or data center via Docker or Kubernetes— no data leaves your network.

SSO built in

Supertokens supports SAML and OIDC, so your team signs in with the identity provider you already use.

obiguard

Common questions.

Don't see what you're looking for? Our solutions engineers respond within one business day.

01

Does Obiguard see our customers' data?

In SaaS mode, content is processed in-memory to make a policy decision — only the decision and metadata are persisted to the Audit Log. Private and on-prem options are available.

02

Does this add latency to model calls?

Inspection runs on the request path, so it adds some latency — how much depends on which detectors are enabled for your policy.

03

How is Obiguard different from a model gateway?

Gateways route traffic. Obiguard governs it — adding policy enforcement, violation tracking, review workflow, and audit trail.

04

Which compliance frameworks do you support?

NIST AI RMF, ISO/IEC 42001, and the OWASP LLM Top 10 today, as a self-assessment aid — plus your internal AUP.

obiguard

Bring the same rigor you apply to data and identity — to the AI in your stack.

A 20-minute call with a solutions engineer is enough to scope your pilot. Most customers are enforcing policy in production within two weeks.

Email: hello@obiguard.com →

https://obiguard.ai ↗

AI risk governance for organizations shipping LLMs in production. Built by veterans of security, identity, and ML platform.