

Govern AI with Confidence

Enterprise controls built into every AI transformation. Reduce AI execution risk with built-in security, governance, and policy enforcement.

Built for mission-critical AI

The Unframe platform embeds security and governance rules AI cannot override — reducing execution risk at every layer of your AI program. You define the infrastructure, data access, agent behavior, and user controls. The AI operates within your governance framework, not around it.



Data doesn't leave your perimeter

Process and store everything within your enterprise boundary.



Granular access

Manage with RBAC, SSO, and adaptive access policies.



Deployment flexibility

Deploy in cloud, private VPC, or fully on-premises.



Model agnostic

Choose the models you already use and trust.


Control every layer of AI

There are no defaults in your governance model — the platform is configurable at every layer. Choose your infrastructure, define agent rules, and configure user and data governance to match your requirements.

1	2	3
Infrastructure	Agent guardrails	User & data governance
On-premises, Unframe-hosted, or your own cloud — you choose the perimeter. In customer-hosted mode, Unframe has zero access to your environment. Your data never moves unless you decide it does.	Every agent operates within a typed data contract that defines what it can access, what actions it can take, and what requires human approval. Rules are enforced at the platform layer — not delegated to model judgment.	The most granular layer of control, configurable independently of infrastructure and agent settings. Data - what an agent can see Agent - what an agent can do User - who directs it, and the proof



One governed path for every request



Control what AI can see

- ✓

Scoped, credentialed access
Agents see only what you explicitly grant, with no ambient access to the wider environment.
- ✓

End-to-end encryption
All data is encrypted at rest and in transit, with controls that mitigate exfiltration risk across every connected source.
- ✓

PII boundaries & masking
Sensitive data is detected and redacted automatically before it can be exposed.
- ✓

Knowledge Fabric & managed memory
Fragmented data is modeled into governed, agent-readable entities; each agent's memory stays isolated.

Control what AI can do

- ✓

Deterministic steps & gates
Critical steps run as fixed, auditable logic; actions beyond an agent's approved gate route to review.
- ✓

Governed registry
Every agent and tool is registered, certified, and audited; new capability enters only by admin approval.
- ✓

Human-in-the-loop
Gated actions wait for a named person or role, with an escalation ladder; nothing gated proceeds on its own.
- ✓

AI security
Built-in defenses against prompt injection, jailbreaks, and adversarial inputs — the model cannot be manipulated into bypassing platform controls.

Control who can use AI - and how

- ✓

Role-based access & two-tier admin
RBAC binds every user; super-admins set tenant policy, unit admins govern their own scope.
- ✓

Policy enforcement in the moment
PBAC blocks out-of-scope actions as an explicit, logged decision, with no silent override.
- ✓

Audit-ready evidence
Each run exports a machine-readable bundle: policies evaluated, PII boundary, and the human approval.

About Unframe

Unframe is the Applied AI Lab dedicated to closing the gap between AI capabilities and your business outcomes. We meet organizations where they are: end-to-end managed delivery, dedicated expertise for teams that self-build or strategic advisory. All on a proven platform built for speed, governance and enterprise scale. Unframe solutions run in the customer’s preferred environment, integrating with existing AI programs and core systems.

Unframe is headquartered in Cupertino, California, with a global presence in Tel Aviv and Berlin.