

Astrapia.AI — GenAI Governance & Enablement Platform

Astrapia.AI is a centralized platform to govern, scale, and operationalize GenAI across the enterprise. It unifies use case intake, prompt assets, policy guardrails, evaluation & monitoring, cost controls, and integrations—so teams can move from experimentation to production with confidence.

Why Now

GenAI is being deployed rapidly across functions, but scaling responsibly is hard: fragmented tools, limited visibility, inconsistent standards, integration gaps, and security/compliance risks. Astrapia.AI addresses these with a single, governed operating layer for GenAI.

Key Capabilities

- Use Case Catalog & Intake — standardize discovery, prioritization, and approvals.
- Prompt & Asset Library — shared, versioned prompts and components to reduce duplication.
- Model & Connector Hub — plug and play with OpenAI, Azure OpenAI, Google Gemini, and custom LLMs.
- Guardrails & Policies — data masking/PII redaction, safety filters, RBAC, and audit trails.
- Evaluation & Monitoring — quality, drift, hallucination and latency metrics with test sets and golden prompts.
- Observability & Cost Control — usage dashboards, budget caps, per team cost allocation.
- Workflow Orchestration — promote from sandbox to staging to production with checks and approvals.
- Integrations — Jira/ADO, Slack/Teams, Git, ServiceNow for tickets, chat, and CI/CD hooks.

Business Benefits

- Faster Time to Value — move pilots to production quickly with standard rails.
- Reduced Risk — enforce policy, security, and compliance centrally.
- Higher Quality — evaluate and improve prompts/models continuously.
- Lower Cost — control usage, prevent sprawl, and reuse shared assets.
- Enterprise Visibility — leadership level oversight of initiatives and outcomes.

Deployment Options

Available on cloud or on prem with SSO, SCIM, and granular RBAC. API first architecture with webhooks for easy integration into existing tools and data platforms.