FOR AI AND DATA GOVERNANCE

Guardrail AI

AI governance, risk and compliance: inventory, shadow AI, assessments, controls

Know every model, agent and AI tool in use, and show that each one is governed.

HOW IT WORKS

01
Discover
Gain complete AI visibility
Discover every model, agent and AI tool, including shadow AI.

02
Assess
Understand risk and impact
Continuously assess risk and business impact.

03
Govern
Establish policy and control
Policies, frameworks and model governance at scale.

04
Protect
Secure AI and sensitive data
Enforce controls, prevent misuse and data leakage.

05
Respond
Detect, investigate and resolve
Detect, investigate and resolve AI incidents fast.

CAPABILITIES

EXECUTIVE & STRATEGY

- **AI Guardrails** Risk index, control coverage and open incidents for the whole AI estate in one view.
- **Reports & Analytics** Board, regulator and risk reports drawn from every module, delivered on schedule.

DISCOVER

- **AI Inventory** Every model, agent, MCP server and RAG pipeline in one register, from discovery to retirement.
- **Shadow AI Discovery** Unsanctioned AI tools found across your estate and held until someone approves them.

ASSESS

- **Assessment Hub** DPIA, FRIA and conformity assessments run from templates, with each sign-off on record.
- **Risk Intelligence** One risk index across risk domains and asset types, re-scored as conditions change.
- **Impact Assessments** Assessment coverage for each high-risk system, with missing evidence flagged.
- **Vendor & Supply Chain** Third-party AI rated by risk tier, with AI-BOMs and attestations on file.

GOVERN

- **Compliance & Frameworks** Controls mapped to the EU AI Act, NIST AI RMF and GDPR, with evidence attached to each.
- **AI Security Frameworks** MITRE ATLAS and OWASP threats mapped to the controls that address them, with gaps shown.
- **Policy Library** AI policies taken from draft to enforcement, with exceptions reviewed and every change logged.
- **Data Governance & Lineage** Which datasets feed which models, and whether consent and transfer rules allow it.

PROTECT

- **Security & DLP** Sensitive data in prompts detected and blocked before it reaches a model.
- **GenAI Controls** Only approved prompt templates reach production, and answers are checked against their sources.
- **Agentic AI Controls** Each agent's tool calls checked against its capability boundary before they run.

RESPOND

- **Unified Alerts Center** Alerts from every module in one queue, correlated and triaged with playbooks.
- **Incidents & Cases** AI incidents tracked from open to closed against their SLAs, with evidence grouped by case.
- **Serious Incidents** Serious incidents reported to regulators and executives within the required deadline.

WHY GUARDRAIL AI

Built for AI, not adapted to it
Controls are written for models, prompts, agents and tool calls. It is not a general GRC tool with an AI tab added.

Models and agents in one register
Credit models, GenAI assistants and autonomous agents sit in one inventory, each with an owner, a risk tier and an autonomy level.

Enforces, not just records
Policy runs at the point of use: prompts are inspected, unsanctioned tools are held and agent calls are checked before they execute.

Evidence a regulator can follow
Each control links to its evidence, each assessment to its sign-off and each serious incident to the report that was filed.

FRAMEWORKS AND REGULATION

EU AI Act	Risk-tier classification, conformity and fundamental rights impact assessments, and serious incident reporting.
NIST AI RMF	Govern, Map, Measure and Manage functions mapped to controls and evidence.
GDPR	DPIAs, records of processing and cross-border transfer checks for AI that uses personal data.
Bangladesh DPA	Personal data controls for AI systems that process data in Bangladesh.
MITRE ATLAS	Adversarial AI techniques, such as prompt injection and data poisoning, mapped to controls.
OWASP Top 10 for LLMs	LLM risks such as sensitive information disclosure and excessive agency, mapped to controls.

DEPLOYMENT, DATA AND AI

Deploy your way
On-premises, SaaS or hybrid. Self-managed, co-managed, or operated by Raptas 24/7.

Your data stays yours
On-premises and hybrid deployments process data inside your infrastructure; SaaS runs in a dedicated environment in your region. Never shared, never used to train models for anyone else.

AI you can account for
Risk scores and classifications show the factors behind them. Unsanctioned tools stay blocked until a person approves them, high-severity fixes need sign-off, and every decision is written to an immutable audit log.

Connects with: AWS · Azure · GCP · MLflow · Kubeflow · LangChain · CrewAI · AutoGen · ServiceNow · Archer · OneTrust

See Guardrail AI on your own scenario.
Request a demo at raptas.ai/contact · connect@raptas.ai