ALTIMETRIK®

Ai FIRST

Governing Custom GPTs

Building Responsible AI Foundations for the Enterprise

Author: Martin Gaida, AI Solutions Architect
Contributor: Frankie LaCarrubba, AI Solutions Architect

Introduction

Generative AI is a transformative force redefining how enterprises innovate, operate, and compete. With the advent of Custom GPTs, organizations can now rapidly deploy domain-specific assistants that embed knowledge, workflows, and integrations directly into business processes. This democratization of AI empowers employees at every level to design solutions once reserved for specialized teams, unlocking new frontiers of productivity and creativity.

Yet with scale comes complexity. Recent studies show that over 60% of enterprises cite governance and compliance as their top concern in deploying generative AI.

An EY study in June published that while 75% of enterprises are already using GenAI, only about one-third have responsible controls in place - a stark gap between adoption and governance.

Left unmanaged, risks such as behavioral drift, regulatory non-compliance, and uncontrolled data exposure threaten to erode trust and limit adoption.

The opportunity is clear: enterprises must embrace the speed and creativity of Custom GPTs while embedding the governance structures needed to safeguard accountability, resilience, and long-term scalability. This paper proposes a governance framework designed to responsibly unlock AI's potential, balancing innovation with oversight to ensure sustained business value.

The Governance Imperative

Custom GPTs allow employees to build powerful AI solutions with minimal friction: natural-language instructions, file uploads, and tool integrations. However, this ease of creation brings enterprise-wide risks:

Lack of a single source of truth (SSOT) for ownership, configuration, and usage.

Unmonitored exposure to higher-risk capabilities such as browsing or API integrations.

Behavioral drift as GPT configurations evolve without structured review.

Compliance and reputational risks tied to unvetted outputs and data handling.

In the absence of governance, these risks compound, creating shadow AI that mirrors the pitfalls of shadow IT. For example, a recent McKinsey study found that 55% of organizations experimenting with Generative AI report concerns about shadow AI, where tools are deployed without governance oversight.

A Three-Pillar Governance Framework

Altimetrik recommends a modular framework that organizations can adapt and scale, ensuring responsible AI adoption across the enterprise.

Pillar 1: Central Registry – Establishing Control

A central registry becomes the enterprise's authoritative inventory of all Custom GPTs. Through automated metadata collection (e.g., GPT ID, ownership, instruction sets, tools enabled) combined with owner-supplied context (e.g., business purpose, risk tier, compliance regimes), the registry provides:

- **Visibility** into GPT deployment and lifecycle.
- **Auditability** aligned to compliance obligations such as GDPR, HIPAA, and SOC 2.
- **Foundation for downstream controls**, from access restrictions to risk assessments.

Example: A global bank might classify GPTs supporting fraud detection as high-risk assets requiring quarterly compliance reviews, while customer service GPTs may fall into a lower-risk tier. This mirrors how enterprises already manage structured data catalogs, reinforcing a familiar governance discipline.

Pillar 2: Evaluation & Monitoring – Enabling Continuous Oversight

Governance must be dynamic, not static. Continuous monitoring ensures Custom GPTs remain safe, effective, and aligned with business intent. This includes:

- **Static Analysis** – scanning for unsafe prompts, unapproved domains, risky tool enablement, or tier mismatches.
- **Activity Monitoring** – analyzing usage telemetry such as adoption rates, safety scores, and drift signals.
- **Adaptive Review Cycles** – calibrating monitoring intensity to declared risk tiers.

Example: In healthcare, a Custom GPT that drafts patient communications could be monitored weekly for HIPAA compliance, while an e-commerce GPT generating product descriptions may only need quarterly checks. Gartner research shows that **enterprises with automated monitoring reduce mean-time-to-remediation by 35%,** underscoring the importance of real-time visibility.

Pillar 3: Guardrails & Enforcement – Operationalizing Trust

When risks emerge, organizations need clear response mechanisms. Enforcement should be **policy-based, scalable, and proportionate**:

- **Access Restrictions** to isolate risky GPTs.
- **Escalations** to owners for remediation and metadata updates.
- **Deactivation** of dormant or non-compliant GPTs.

Incident response flows (detection, containment, resolution, retrospective) ensure issues are logged, remediated, and inform future guardrail refinements.

Example: A pharmaceutical company detected that a Custom GPT was attempting to connect to unapproved domains through a custom action. Enforcement workflows isolated the GPT immediately, preventing a potential HIPAA violation. This mirrors cybersecurity incident response playbooks, where detection, containment, resolution, and retrospective analysis are standard.

Regulatory Alignment

Governance is not just good practice, it is regulatory necessity.

GDPR	Central registry supports Article 30 record-keeping.
SOC 2	Monitoring aligns with principles of change management and system oversight.
HIPAA	Guardrails map to administrative safeguards for PHI-related usage.
NIST AI Risk Management Framework & ISO 42001	Provide global benchmarks for responsible AI adoption, emphasizing risk-based controls and continuous monitoring.

Example: A European insurance provider may rely on GDPR Article 30 compliance obligations, while a U.S. healthcare organization prioritizes HIPAA safeguards. A unified governance framework ensures consistency across these diverse regulatory landscapes.

Implementation Roadmap

A phased approach allows enterprises to mature governance without slowing innovation:

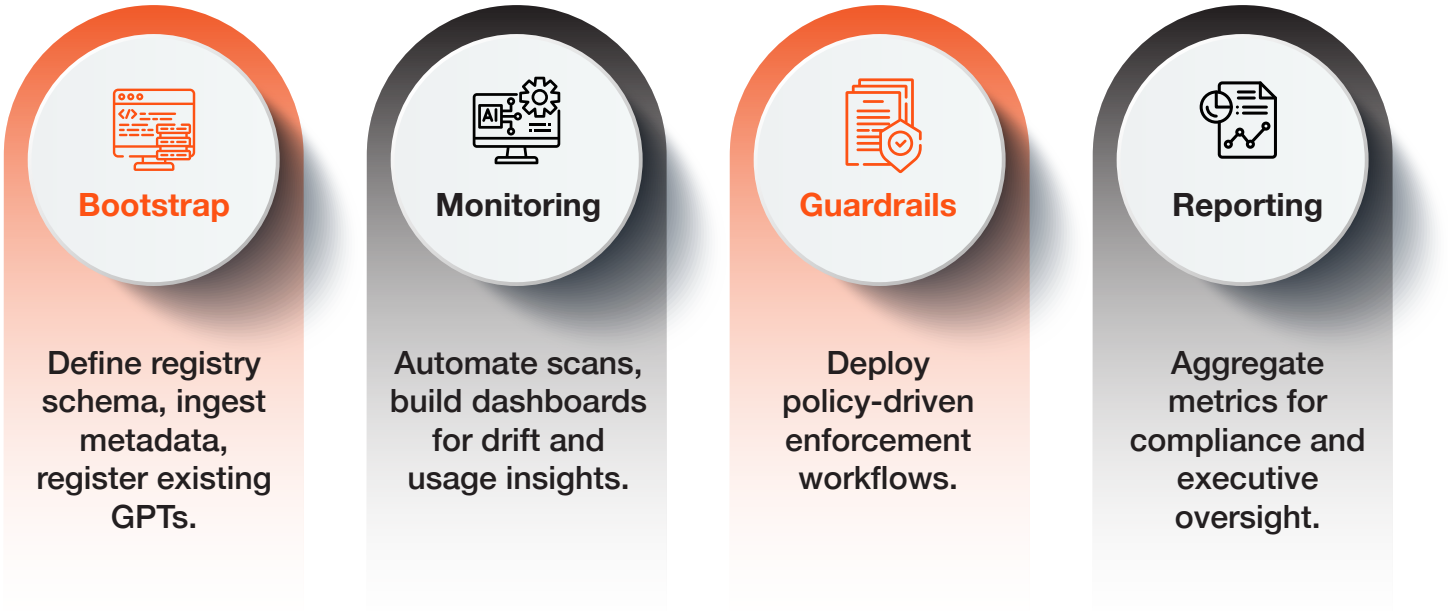

Benchmark: According to Gartner, enterprises that implement structured AI governance see **40% higher adoption rates of approved AI tools** compared to those with ad-hoc oversight.

Conclusion: Responsible AI as a Business Enabler

Custom GPTs represent a transformative opportunity for enterprises but only if paired with strong governance. By adopting this three-pillar model, organizations gain:

- **Innovation at scale** without exposing themselves to unmanaged risks.
- **Regulatory readiness** built into day-to-day AI usage.
- **Trust and accountability** across teams, customers, and regulators.

Governance is not a constraint but a catalyst for scaling AI with confidence.

Enterprises that treat responsible AI as a core business capability, rather than a compliance checkbox, position themselves to unlock innovation at speed and with trust. By embedding governance into the very DNA of the enterprise, organizations build resilience against risk, credibility with stakeholders, and adaptability in the face of evolving regulations.

In this way, governance becomes a **strategic enabler** driving not only safer adoption, but smarter adoption. It ensures AI systems remain aligned to purpose, accountable in their decisions, and scalable across diverse business functions. Enterprises that embrace this mindset will be the ones to lead in the new digital era: **innovating faster, safeguarding trust, and setting the standard for responsible AI at scale.**