

Executive Brief

Probabilistic Models Require Deterministic Governance

Why AI architecture — not just AI policy — must be deterministic.

THE PROBLEM

Generative and agentic AI systems are probabilistic by nature. They hallucinate, drift, and produce outputs that cannot be fully predicted or guaranteed. Yet enterprises deploy them in consequential contexts — customer interactions, financial decisions, autonomous operations — where errors carry legal, reputational, and operational risk.

Current governance approaches treat AI as a model to be documented, not as a system to be controlled at runtime. Risk assessments are static. Compliance evidence is manual. Controls are advisory, not enforceable. The gap between policy and execution grows wider with every model update and every new agent deployment.

**We cannot make probabilistic systems deterministic.
We can constrain their authority and bound their consequences.**

THE ARCHITECTURE: DETERMINISTIC CAGE™

The solution is not better models alone — it is a deterministic control perimeter surrounding the probabilistic core. The Deterministic Cage constrains the model's authority, data access, action scope, transaction size, execution time, and reversibility. The model proposes; the system evaluates; humans retain authority.

Authority bounded	The agent may propose; it may not authorise.
Data bounded	Scope of data access is policy-enforced, not model-decided.
Action bounded	Permitted actions are explicitly enumerated; all others are denied.
Time bounded	Execution windows have hard deadlines and automatic timeout.
Financial impact bounded	Transaction-size caps prevent catastrophic exposure.
Reversibility enforced	Every action is logged with sufficient context to reverse.

ILLUSTRATIVE SUCCESS MEASURES

-40% reduction in audit preparation time	-60% faster incident detection and response	>99.5% illustrative target: policy enforcement coverage at runtime
--	---	---

Illustrative objectives, not historical performance claims. Calibrable by sector, maturity, and risk appetite. Evidence demonstrated in public repos: KS drift detection with configured threshold and JSON evidence packs with SHA-256 hashing.

THE OPERATING MODEL: DGOM™

Dual Governance Operating Model

DGOM operationalizes the Deterministic Cage thesis through a four-phase lifecycle that connects corporate governance, regulation, and engineering through executable controls and continuous evidence.

I	PLAN	II	BUILD	III	DEPLOY	IV	MONITOR
Risk classification, regulatory mapping, control objectives, and accountable assignment before any model is built.		Evidence-based model selection, adversarial testing (GARAK, PyRIT), bias audits, and tamper-evident validation.		Runtime authorization gates, human approval workflows, veto logs, and tamper-evident audit trails at every trust boundary.		Drift detection, continuous assurance, boundary recalibration, and independent audits across the model lifecycle.	

REGULATORY INTEROPERABILITY

Framework	How DGOM maps
EU AI Act	Art. 9 risk management, Art. 13 transparency, Art. 14 human oversight → PLAN + DEPLOY
NIST AI RMF 1.0	Govern, Map, Measure, Manage → all four phases
ISO 42001	AI management system clauses → PLAN + MONITOR

PUBLIC EVIDENCE

The architecture is not theoretical. Two public reference implementations demonstrate executable controls, drift detection, and evidence generation:

- Insurance Compliance Hub — regulated-insurance governance gate with KS drift detection, machine-readable evidence (JSON evidence packs with SHA-256), and CI/CD blocking on threshold breach.
github.com/tshapedconsultant/porto-seguro-compliance-hub
- Predictive Maintenance System — Evidence-based model selection (EBM vs Random Forest), GARAK/PyRIT adversarial testing, dual-approval workflow with signed veto log.
github.com/tshapedconsultant/Predictive-Responsible-Maintenance

Next step: AI Governance Diagnostic (2–3 weeks) — assess your DGOM 1–4 maturity across PLAN, BUILD, DEPLOY, and MONITOR. Receive a prioritized 90-day roadmap with executable controls.
andreslage@tshapedconsultant.com · [LinkedIn](#) · [GitHub](#)

Full thesis: Probabilistic Models Require Deterministic Governance (22 pp.) · Operating model: DGOM™ — Dual Governance Operating Model (12 pp.)

© 2026 Andrés Lage Freire. All rights reserved. Provided for information purposes only. No part may be reproduced, adapted or commercialised without written permission.