



Platform Engineering for AI Governance: Turning Regulations into Runtime-Verified Policies

Mohamed ElBendary

Enterprise Architecture Manager, FasTrak SoftWorks, Inc.

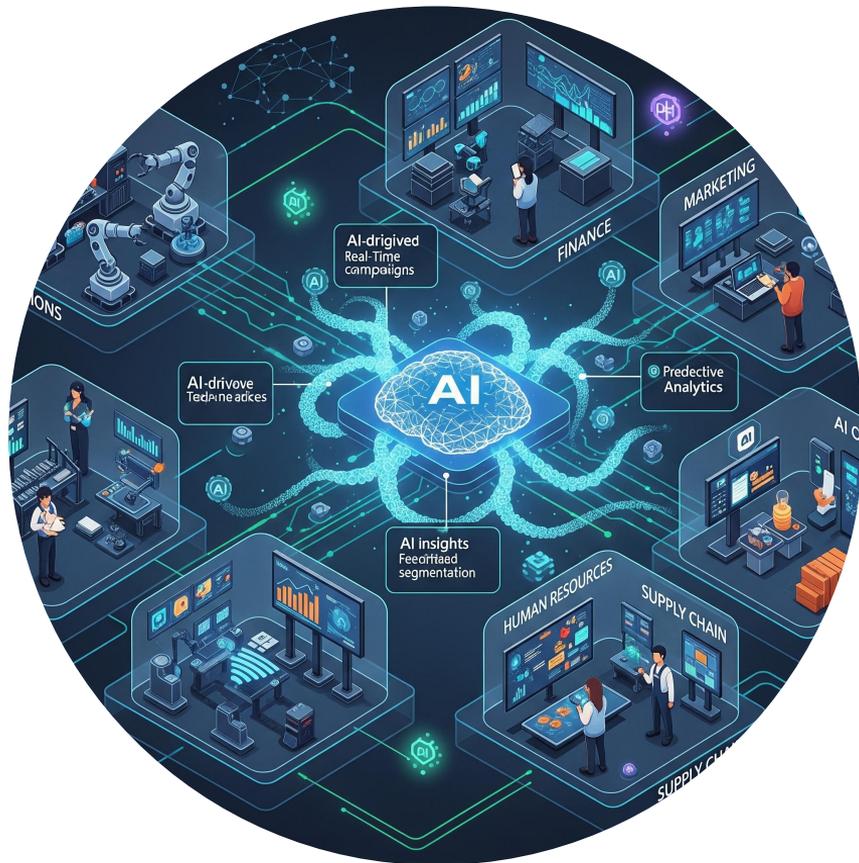


Takeaway

Platform Engineering is the most viable provider of the "Deterministic Layer" required for Agentic AI Survival

- 1. Enforce Model Independence:** Decouple business logic from foundation models to survive correlated upstream failures.
- 2. Guarantee Traceability:** Immutable, deterministic logs are the only way to satisfy the evidence requirements of post-2025 insurance policies and arriving regulations.
- 3. Implement Governance-as-Code:** Actively reject non-compliant agentic/model behavior





AI is a Capability, Not a Tool

- **Resources:** Data, Compute (GPUs), Models, Talent, Capital, etc.
- **Processes:** MLOps, RAG Pipelines, Red Teaming, Governance, training, fine-tuning, monitoring, etc.
- **Values (Priorities):** Determining acceptable trade-offs (e.g., Accuracy vs. Latency, Innovation Speed vs. Safety Compliance, etc.)



AI Creates a Novel Risk Profile

1

Non-Deterministic Execution: Probabilistic outputs break traditional QA and reduce insurability forcing a commercial correction that increases adoption costs

2

Ubiquity & Shadow AI: Low barrier to entry coupled with low organizational AI literacy expand security risks

3

Regulatory Complexity: Global and state-level patchwork of AI regulations increase the cost of compliance when present and increase innovation/adoption tax when uncertain

4

Correlated Failures drive Aggregated Risk: Foundation Model dependencies create failure synchronization at scale with unpredictable scope



Why Platform Engineering for AI Capability Implementation

1

Insulate business value streams from self-serve sanctioned AI infrastructure implementation

2

Define and enforce enterprise-wide AI interaction semantics (entitlements, observability, model routing, compliance enforcement, escalation, etc.)

3

Facilitate the creation of a shared language across functions for communicating risks and opportunities of AI integration

4

Build domain-specific runtime containers for deterministic execution with plug-in architecture for business units



The How: Four-Layer Contract Loops Architecture

- 1 Regulatory Layer:** encodes the non-negotiable legal constraints externally imposed on the system. Derived directly from laws like the EU AI Act or GDPR.
- 2 Governance Layer:** encodes the organization-wide internally imposed mandates governing the enterprise's involvement in AI development and deployment.
- 3 Domain Layer:** Defines roles, responsibilities, and operational semantics of end-to-end value streams.
- 4 Operation Layer:** Defines the semantics of a single operation on one or more value stream(s).



The How: Domain and Operation Contracts

1

Responsibility/Roles/Inputs/Outputs

3

Success/Failure End Conditions

2

Preconditions/Postconditions

4

Escalation Semantics



The How: Contract Loops Architecture Constraints

- 1 Assume-Guarantee scheme from upper layer to lower layer to support scaling oversight
- 2 Cryptographic Agent Identity
- 3 Agent Outputs are Proposals not Decisions that must be checked using multi-model antagonism
- 4 Tamper-proof logging of orchestration, delegation, execution, and escalation



The How: Contract Discovery Workflow

- 1** Map use case to value stream
- 2** Determine which operations on the value stream are going to be handled by AI agent(s)
- 3** Determine outer loop requirements for each operation if any
- 4** Define a contract for each operation in machine-readable format (e.g. JSON)

Benefits for the Enterprise

1. **Model vendor independence with hot swapping against contracts**
2. **Regulatory adaptability without taxing innovation agility**
3. **Non-repudiation lowers compliance cost and maintains insurability**
4. **Leverages existing capabilities reducing execution risks of transformation**
5. **Secure multi-agent orchestration across business units and externally**
6. **Enables Governance-as-code and Adaptive GovOps**



Thank You!

Email: me@prosterk.com

Linkedin:

<https://www.linkedin.com/in/mohamedelbendary/>

