

AI GOVERNANCE & RISK MANAGEMENT.

Every AI model an organization builds, buys, or embeds — **from vetted platforms to shadow tools employees adopt on their own** — carries risk traditional software review was never built to catch. Governance turns unmanaged AI use into an accountable program.

- MODEL INVENTORY
- SHADOW AI DISCOVERY
- RISK TIERING
- FRAMEWORK MAPPING

01 THE MECHANISM



Model Inventory
Cataloging **every AI system in use** — internal builds, vendor platforms, and embedded features — before risk can be assessed.



Risk Classification
Tiering each use case by **impact, autonomy, and data sensitivity** to focus oversight where it matters most.



Human Oversight
Keeping a **human accountable for high-impact decisions**, especially where a model's output isn't fully explainable.

02 THE FLOW

The Intake Cycle PER MODEL

1

A new AI use case — **build, buy, or embedded feature** — is registered before deployment.

2

The model is scored against **risk criteria** — data sensitivity, decision autonomy, and potential harm.

3

Higher-risk use cases route to deeper review before approval.

The Oversight Cycle ONGOING

1

Approved models are **monitored for drift, bias, and performance degradation** over time.

2

Usage and outputs are **periodically audited** against the original approval scope.

3

Material changes trigger re-assessment, not silent continuation.

03 CORE PILLARS & COVERAGE

PILLAR	WHAT IT COVERS
Model Inventory & Discovery	Surfaces sanctioned tools, shadow AI, and vendor-embedded models across the environment.
Risk Tiering & Classification	Scores use cases by data sensitivity, autonomy, and potential impact.
Human Oversight & Accountability	Assigns an accountable owner for every high-impact automated decision.
Monitoring & Drift Detection	Tracks performance, bias, and behavior changes after deployment.

04 OVERSIGHT APPROACHES COMPARED

SIGNAL	SUFFICIENT ALONE?	WHAT IT MISSES
Vendor Compliance Certification	No	Confirms the vendor's controls, not how your organization actually uses the model.
One-Time Model Review	No	Model behavior and training data can change well after initial approval.
IT Security Review Alone	No	Focuses on infrastructure risk, not decision-quality or bias risk.
Employee Self-Reporting	No	Shadow AI adoption routinely goes unreported until an incident surfaces it.
Combined AI Governance Program (Inventory + Risk Tiering + Oversight)	Yes	Continuously accounts for every model in use and who is accountable for its output.

AI GOVERNANCE.

INVENTORY · TIER · OVERSEE

05 PROGRAM OPERATIONS

You can't govern what you haven't inventoried. Shadow AI adoption is now a leading blind spot in enterprise risk registers.

Not every model needs the same scrutiny. Risk tiering focuses deep review on high-impact, high-autonomy use cases.

Explainability matters as much as accuracy. A high-performing model with unexplainable outputs still carries accountability risk.

Governance is a program, not a launch gate. Approved models must be re-evaluated as data, usage, and regulations evolve.

06 CONTROL MAPPING

NIST AI RMF 1.0

Foundational **risk-management framework** for trustworthy AI.

ISO/IEC 42001

International **standard for AI management systems**.

EU AI ACT

Risk-tiered legal obligations for AI systems by impact category.

NIST CSF 2.0

Govern function coverage for organizational risk oversight.

07 WHAT THIS DOES NOT SOLVE

Model bias baked into training data
Governance flags and monitors bias; it doesn't **remove it from historical training data**.

Adversarial manipulation of inputs
Prompt injection and data poisoning fall **outside a pure governance review**.

Vendor model changes without notice
Third-party providers can **update models with limited advance disclosure**.

Cultural resistance to reporting
Policy alone doesn't guarantee employees **disclose the tools they actually use**.

08 MYTHS VS. REALITY

MYTH "We don't use AI, so this doesn't apply to us."	REALITY Employees adopt AI tools informally long before an official program exists.
MYTH "A vendor's compliance badge covers our risk."	REALITY Certifications cover the vendor's controls, not your specific use case .
MYTH "Only customer-facing models need oversight."	REALITY Internal decision-support models carry the same bias and accountability risk .
MYTH "Accuracy is the only metric that matters."	REALITY Explainability, fairness, and accountability matter just as much as performance.
MYTH "Governance slows innovation."	REALITY A clear intake process moves lower-risk use cases through faster , not slower.

09 SPEAK THE PROGRAM — KEY TERMS

AI GOVERNANCE
The policies, processes, and accountability structures that oversee **how AI is built, bought, and used**.

SHADOW AI
AI tools adopted by employees or teams **without formal review or approval**.

RISK TIERING
Classifying AI use cases by **potential impact** to prioritize oversight effort.

MODEL DRIFT
Gradual degradation in a model's accuracy or behavior as **real-world data shifts** from training data.

EXPLAINABILITY
The degree to which a model's decision-making process **can be understood by a human reviewer**.

HUMAN-IN-THE-LOOP
A control requiring **human review or approval** before a high-impact AI decision takes effect.

SOURCES: NIST AI RMF 1.0 · ISO/IEC 42001 · EU AI ACT · NIST CSF 2.0

AI GOVERNANCE & RISK MGMT · FIELD REFERENCE