WHITEPAPER · RESPONSIBLE AI SERIES

CONNECTIVE BUSINESS SOLUTION

STRATEGY RE-DEFINED

AI GOVERNANCE IN THE ENTERPRISE

A practical framework for establishing AI governance policies, risk management protocols, and ethical guidelines, built for organizations that want the benefits of AI without inheriting its unmanaged risks.

Written For

Executive Leaders · CIOs and Technology Leaders · Risk and Compliance Officers

2026 · connectivebusinesssolution.com

AI adoption in most organizations is running ahead of AI governance. Teams are already using generative tools, vendors are already embedding AI in products, and procurement contracts are already writing AI restrictions into terms and conditions. The question facing leadership is no longer whether to govern AI, but whether governance will be designed deliberately or assembled after an incident.

This paper presents a practical five-pillar framework for enterprise AI governance: Accountability, Use-Case Governance, Data and Security Controls, Human Oversight, and Workforce Enablement. It includes a risk-tier model for classifying AI use cases, a 90-day implementation roadmap, the metrics that indicate whether governance is working, and the contract considerations that increasingly determine what AI use is permitted at all. The framework is designed to be right-sized: enough structure to manage real risk, not so much process that teams route around it. That balance, more than any single policy, is what separates governance that works from governance that exists on paper.

01 / THE GOVERNANCE GAP

Three forces make AI governance urgent now rather than eventually.

Adoption is bottom-up

Unlike most enterprise technology, AI tools arrive through individual initiative. Employees adopt generative assistants because the tools are useful, often before any policy exists. The result is "shadow AI": organizational AI use that leadership cannot see, inventory, or control. The risk is not that employees are acting in bad faith; it is that well-intentioned use without guardrails puts confidential data, work product quality, and compliance obligations in the hands of individual judgment.

Vendors are embedding AI faster than contracts are catching up

AI capability now arrives inside products an organization already licenses. Each embedded feature raises the same questions: what data does it access, where is that data processed, and who is accountable for its outputs? Organizations that cannot answer these questions for their own vendor stack cannot honestly certify their compliance posture to customers, regulators, or auditors.

Counterparties are writing AI terms first

Governments and enterprise customers are moving from guidance to contract language. As a working example: the State of Florida's current enterprise standard contract terms define AI technology broadly and prohibit its use in deliverables, or on government data, unless specifically described in the contract or approved in writing by the customer. An organization delivering services under such terms without an internal mechanism to know where AI touches its work is not merely ungoverned; it may be in breach without knowing it.

> **THE PRACTICAL POINT**
>
> Governance is often framed as an ethics exercise. In practice it is increasingly a contractual and operational necessity: the organizations asked to certify their AI practices will be the ones that can, and they will win work that others cannot.

02 / THE FIVE-PILLAR FRAMEWORK

Effective AI governance answers five questions: who decides, what is allowed, what protects the data, who checks the output, and how people learn the rules. Each pillar below exists to answer one of them.

PILLAR 1

Accountability and Ownership

Governance fails first when no one owns it. Establish a cross-functional AI governance group with real decision rights: typically technology, legal or compliance, security, and a business owner, meeting on a defined cadence.

- Name a single accountable executive for AI governance, with authority to approve, restrict, and retire AI use.
- Maintain a living AI inventory: every tool, embedded feature, and internal model in use, with its owner and risk tier.
- Define decision rights in writing: what teams may adopt on their own, what requires review, and what is prohibited.

PILLAR 2

Use-Case Governance

Govern uses, not tools. The same assistant can be low risk for drafting an internal summary and high risk for producing client deliverables. An intake and classification process keeps review effort proportional to actual risk.

- Create a lightweight intake: what the use case does, what data it touches, who consumes the output, and what happens if the output is wrong.
- Classify each use case into a risk tier (Section 03) that determines the required controls and approval level.
- Re-review when the use changes, not just when the tool changes. Scope creep is the most common path from low risk to high.

PILLAR 3

Data and Security Controls

Most AI risk is data risk wearing a new interface. The controls are familiar; the discipline is applying them to a new channel.

- Define what data classes may and may not be provided to AI systems, aligned to your existing data classification.
- Require security and legal review of AI vendor terms: training rights, data retention, processing location, and confidentiality.
- Prefer enterprise deployments with administrative controls over unmanaged consumer accounts for any business use.
- Log and review AI system access the way you review access to any other system that handles sensitive data.

PILLAR 4

Human Oversight and Quality

AI output is a draft until a qualified person accepts responsibility for it. Oversight requirements should scale with the risk tier of the use.

- Set review requirements by tier: high-risk outputs require documented review by a qualified owner before use.
- Prohibit fully automated decisions that materially affect people (hiring, credit, discipline, benefits) without human review.
- Validate before scaling: pilot new use cases with defined success criteria and error checks before broad rollout.
- Establish an incident path: how employees report AI errors, data exposure, or misuse, and who responds.

PILLAR 5

Workforce Enablement

A policy no one has read is a liability document, not a control. The final pillar turns governance into behavior.

- Publish a short, plain-language responsible-use policy: what is encouraged, what requires approval, what is prohibited.
- Train by role, not by memo: the finance team, the delivery team, and the marketing team face different risks.
- Make the approved path easier than the workaround: sanctioned tools, clear guidance, fast intake decisions.
- Revisit training as tools change. Annual is a floor, not a target, in a field that moves quarterly.

03 / CLASSIFYING RISK: A FOUR-TIER MODEL

Risk tiers keep governance proportional. The classification depends on three factors: the sensitivity of the data involved, the consequence of a wrong output, and who relies on the result.

TIER	TYPICAL USES	REQUIRED CONTROLS
PROHIBITED	Uses barred by contract, law, or policy; automated consequential decisions about individuals; inputting restricted data into unapproved tools.	Not permitted. Communicated explicitly in policy and training, with an intake path for exception requests where a contract allows written approval.
HIGH	Client or public-facing deliverables; regulated processes; uses touching confidential or personal data.	Governance group approval; documented human review of outputs; approved enterprise tools only; contract terms verified before use.
MODERATE	Internal analysis, drafting, and decision support consumed by informed internal audiences.	Manager awareness; sanctioned tools; outputs labeled as AI-assisted where relied upon; spot-check quality reviews.
LOW	Individual productivity on non-sensitive content: brainstorming, formatting, summarizing public material.	Standing policy applies; no per-use approval. Monitored through the AI inventory and periodic review.

> **DESIGN PRINCIPLE**
>
> If most of your organization's AI use does not land in the Low tier, your tiers are miscalibrated or your adoption is genuinely risky. Either finding is worth knowing.

04 / A 90-DAY IMPLEMENTATION ROADMAP

Days 1 to 30: See clearly

- Name the accountable executive and stand up the governance group with a defined charter.
- Build the first AI inventory through surveys, expense data, and vendor contract review. Expect it to be incomplete; publish it anyway.
- Issue interim guidance covering the two urgent items: what data may never be entered into AI tools, and which uses require approval today.

Days 31 to 60: Decide deliberately

- Adopt the risk-tier model and classify the inventory. Resolve the high-tier and prohibited items first.
- Publish the responsible-use policy in plain language, with the intake path attached.
- Review the contracts that govern your work for AI clauses, and reconcile current practice against them.

Days 61 to 90: Make it operational

- Deliver role-based training, starting with teams whose work lands in the High tier.
- Stand up sanctioned enterprise tools so the approved path is the convenient path.

- Define governance metrics (Section 05) and set the review cadence: quarterly for the framework, continuously for the inventory.

05 / MEASURING WHETHER GOVERNANCE WORKS

Governance that cannot be measured degrades into ceremony. Six indicators tell leadership whether the framework is functioning.

- **Inventory coverage:** the share of known AI use captured in the inventory, trending toward complete.
- **Classification currency:** the share of inventoried uses with a current risk tier and named owner.
- **Intake velocity:** median days from use-case submission to decision. Slow intake manufactures shadow AI.
- **Training reach:** the share of the workforce, and of high-tier teams specifically, with current training.
- **Incident visibility:** reported AI incidents and near misses. Early on, an increase usually signals trust, not deterioration.
- **Contract conformance:** the share of active engagements with AI terms identified and practice verified against them.

06 / COMMON PITFALLS

- **The blanket ban.** Prohibition without an approved alternative does not stop use; it hides it, and converts a governance problem into a visibility problem.
- **The 40-page policy.** Length is not rigor. A policy people can hold in their heads outperforms one they cannot finish reading.
- **Governing tools instead of uses.** Tool lists age in weeks. Use-case tiers remain stable as the tool market churns.
- **Ethics without operations.** Principles unattached to intake, review, and training decorate a slide and change nothing.
- **One-time governance.** A framework reviewed annually in this field is a framework that spends most of the year out of date.

07 / THE CONTRACT DIMENSION

For organizations that sell to governments or enterprises, AI governance now has an external audience. Procurement terms increasingly define AI technology broadly, restrict its use in deliverables and on customer data, and require written approval for exceptions. Meeting these terms takes exactly the machinery this framework builds: an inventory that knows where AI touches each engagement, a tier model that flags customer-facing use, and an approval trail that can be shown on request. Organizations that build this capability convert a compliance obligation into a differentiator, because they can say yes, in writing, to questions their competitors cannot answer.

ABOUT CONNECTIVE BUSINESS SOLUTION

Connective Business Solution is a Business Performance and Advisory Firm based in Orlando, Florida, helping organizations adopt AI responsibly through strategy, readiness assessments, governance design, and workforce training. Our leadership has managed live AI/ML programs inside a $48M Fortune 500 technology portfolio and structures public sector engagements around state AI contract provisions. Connective is a certified WOSB and MBE and holds two active State of Florida term contracts.

connectivebusinesssolution.com · contact@connectivebusinesssolution.com · 407-541-0808