AI Governance Framework

A Practical Guide for Small and Mid-Size Businesses

Jerome Jayapal | Better Security | bettersecurity.ai

June 2026

AI adoption is accelerating faster than most organizations can manage the risks. This framework gives small and mid-size businesses a practical, plain-language roadmap to govern AI responsibly — without needing a team of lawyers or a Fortune 500 budget.

Why AI Governance Matters Right Now

In the past two years, AI tools have moved from experimental to essential. Employees at companies of every size are using ChatGPT, Microsoft Copilot, Google Gemini, and dozens of other AI platforms — often without formal policies, training, or oversight.

For small and mid-size businesses, this creates real and immediate risk:

Data Privacy: Employees may be pasting confidential customer data, contracts, or financial information into AI tools that retain and use that data for model training.

Compliance Exposure: Regulated industries — healthcare, finance, legal — face specific obligations around how data is handled. AI tools introduce new compliance gaps that auditors are beginning to flag.

Operational Risk: Over-reliance on AI outputs without human review creates decision-making risk, especially in customer-facing or high-stakes contexts.

Vendor Risk: Third-party AI tools are vendors. Without proper assessment, organizations have no visibility into how those vendors store, process, or protect their data.

Reputational Risk: A data exposure or AI-generated error that reaches a customer or the public can damage trust quickly and permanently.

The good news: governance does not have to be complicated. Most small businesses need five things — a policy, an inventory, a risk assessment process, a training program, and a monitoring routine. This framework gives you all five.

The Five Pillars of AI Governance

This framework is organized around five core pillars. Each includes practical actions your organization can implement without a dedicated compliance team.

Pillar	Name
1	**AI Policy & Acceptable Use**
2	**AI Inventory & Vendor Assessment**
3	**Risk Assessment & Classification**
4	**Training & Awareness**
5	**Monitoring & Continuous Improvement**

Pillar 1 — AI Policy & Acceptable Use

Every organization using AI tools needs a written policy defining what is allowed, what is prohibited, and who is responsible. Without a policy, employees make their own decisions — and those decisions may not align with your legal obligations or risk tolerance.

What your AI policy should cover:

– Which AI tools are approved for use and which are not

– What types of data employees may and may not input into AI tools (e.g., no customer PII, no confidential contracts, no financial records)

– Who owns the AI governance function — even if it is a part-time responsibility

– How employees should report concerns or suspected misuse

– How the policy will be reviewed and updated as AI tools evolve

Quick Start Action:

Draft a one-page AI Acceptable Use Policy and distribute it to all employees. Require acknowledgment. Review it every six months.

> *Tip: A clear, readable one-page document that employees actually understand is worth more than a 20-page policy no one reads.*

Pillar 2 — AI Inventory & Vendor Assessment

You cannot govern what you cannot see. Most organizations are surprised to discover how many AI tools their employees are using — many of which were never formally approved.

How to build your AI inventory:

– Survey employees to identify every AI tool currently in use, including personal accounts used for work purposes

– Document each tool: vendor name, purpose, data types accessed, cost, and whether it was formally approved

– Identify which tools have access to company data, customer data, or sensitive information

– Flag any tools operating under personal accounts with no enterprise data controls

Vendor Assessment Questions:

– Does the vendor use my data to train their models? If yes, can I opt out?

– Where is my data stored and who has access to it?

– Does the vendor have a SOC 2 report, ISO 27001 certification, or equivalent security validation?

– What is the vendor's data retention and deletion policy?

– Does the vendor's terms of service align with my industry's compliance requirements?

Quick Start Action:

Build a simple AI inventory spreadsheet: Tool Name, Business Purpose, Data Types Accessed, Approval Status, and Risk Level. Update it quarterly.

Pillar 3 — Risk Assessment & Classification

Not all AI use is equally risky. A simple risk classification system helps your organization focus governance effort where it matters most.

AI Risk Classification Model:

Risk Level	Description	Examples	Controls Required
Low	No sensitive data, internal use, human review of outputs	Grammar tools, internal summarization, scheduling assistants	Policy acknowledgment, basic training
Medium	Limited sensitive data, customer-adjacent, some automation	Customer email drafting, marketing copy, vendor research	Vendor assessment, data handling rules, human review required
High	Sensitive data, regulated context, or automated decision-making	Financial analysis, legal document review, HR decisions, healthcare	Full vendor assessment, legal review, audit trail, executive approval

Quick Start Action:

Apply this classification to every tool in your AI inventory. Any high-risk tool that has not gone through a formal vendor assessment and legal review should be flagged immediately.

Pillar 4 — Training & Awareness

Policies and inventories only work if employees understand them. AI training needs to be practical, relevant, and repeated regularly as tools evolve.

Core training topics for all employees:

– What AI tools are approved and how to access them

– What data is never acceptable to input into any AI tool (PII, confidential contracts, financial data, health information)

– How to critically evaluate AI outputs rather than accepting them as fact

– How to report a concern or suspected AI-related data incident

– Basic awareness of prompt injection and social engineering risks specific to AI tools

Additional training for managers:

– How to evaluate AI vendor risk before approving a new tool

– How to identify and escalate high-risk AI use cases

– Regulatory and compliance implications of AI in your industry

– How to build a culture of responsible AI use within their teams

Quick Start Action:

Deliver a 30-minute AI awareness session to all employees within 30 days of publishing your AI policy. Repeat annually.

Pillar 5 — Monitoring & Continuous Improvement

AI governance is not a one-time project. The tools change. Regulations evolve. New risks emerge. Your governance program needs a regular review cycle to stay effective.

Recommended monitoring activities:

Monthly: Review AI tool inventory for new tools; check vendor news for security incidents or policy changes

Quarterly: Update AI risk classifications; review and update the AI Acceptable Use Policy; check for new regulatory guidance

Annually: Conduct a full AI governance review; reassess all high-risk tools; update training content; report AI risk posture to leadership

As needed: Investigate suspected AI-related data incidents; evaluate any net-new AI tools before approving for use

Key metric to track: What percentage of AI tools in active use have been formally inventoried, risk-classified, and vendor-assessed? Your goal is 100%. Most organizations start below 20%.

Quick Start Action:

Assign a named owner for each pillar. Schedule a 30-minute quarterly review to check progress across all five pillars.

Getting Started: A 90-Day Roadmap

For most small and mid-size businesses, a basic AI governance program can be standing up within 90 days.

Phase	Timeframe	Key Actions
Foundation	Days 1–30	Draft and publish AI Acceptable Use Policy. Build initial AI inventory. Identify policy owner.
Assessment	Days 31–60	Complete vendor assessments for all medium and high-risk tools. Apply risk classification to full inventory. Identify any unapproved tools requiring immediate action.
Enablement	Days 61–90	Deliver AI awareness training to all employees. Establish quarterly review cadence. Document governance program for audit readiness.

About the Author

Jerome Jayapal is the founder of Better Security and a cybersecurity and GRC leader with 12+ years of enterprise experience at Verizon across third-party risk management, IT audits, security governance, security operations, and offensive security testing. He has conducted 80+ vendor security audits across multi-million dollar contracts and brings hands-on experience applying AI to accelerate security research, business operations, and strategic decision-making. He also founded and operates a fast-growing San Diego sports and community business, giving him firsthand perspective on the AI adoption challenges small and mid-size businesses face every day. CISA-certified and a former ISACA San Diego Board Member.

linkedin.com/in/jeromejayapal | jerome.jayapal@bettersecurity.ai | (619) 909-6844