Obtainium.ai

ENGINEERED INTELLIGENCE. REAL IMPACT.

Who Should Own AI Governance At Your Business?

Obtainium.ai LLC

Engineered Intelligence. Real Impact.

The Question Every Business Needs To Get Right

When a small business decides to use AI tools — whether that's a chatbot on your website, an AI assistant for your team, or automated workflows that handle customer data — one question usually goes unanswered: **who is responsible when something goes wrong?**

Most businesses stumble into AI adoption without assigning clear ownership. The marketing team uses one AI tool. Operations uses another. IT approves a third. And if a customer's data gets mishandled, or the AI produces a bad recommendation, or a new regulation kicks in, nobody's sure who owns the problem.

This guide breaks down the five organizational models for AI governance — who typically takes charge, what they do well, and where they fall short. Then it gives you a practical framework for deciding what makes sense for your business right now.

> *What is AI governance? It's the set of practices and policies that let you get real value from AI systems — while managing the risks those systems introduce. It answers three questions: What AI are we using? Who's accountable for each tool? And what do we do when something breaks?*

Why AI Governance Matters for Small Businesses

You might think governance is a Fortune 500 concern. It's not.

Small businesses are increasingly in the crosshairs of AI-related risk:

- **Regulatory exposure**: Laws like CCPA (California) and emerging federal AI rules apply to businesses of all sizes that collect customer data.
- **Vendor liability**: Most AI tools shift legal responsibility to the business using them, not the vendor providing them.
- **Customer trust**: A misstep — an AI chatbot sharing the wrong data, sending an inappropriate message, or making a discriminatory recommendation — can damage your reputation faster than a bad Yelp review.
- **Internal confusion**: Without clear ownership, teams adopt AI tools inconsistently, creating security gaps and compliance blind spots.

The good news: small businesses don't need a full governance program. They need **clear**

ownership and a **basic policy**. That's achievable in weeks, not months.

The Five Models: Who Usually Owns AI Governance

In organizations that have thought about this, five teams typically end up owning AI governance — each with real strengths and real limitations.

Model 1: The Security Team

The pitch: Security teams already evaluate new technology. They have tools for risk assessment, vendor review, and incident response. Putting AI governance under security feels like a natural extension of what they already do.

What they do well:

- Protect the confidentiality, integrity, and availability of data — exactly what's at stake with most AI tools.
- They have established processes for onboarding new technology safely.
- They speak the language of risk and can document threats systematically.

Where it breaks down:

- Security teams are trained to see threats, not opportunities. Their instinct is to slow things down, and AI adoption often requires moving at business speed.
- Security risk is only one dimension of AI risk. Environmental impact, reputational damage, and business value considerations tend to fall outside their scope.
- Many security professionals are newer to AI specifically. Without hands-on experience with large language models or machine learning systems, they may misjudge where the real risks sit.

Best fit for: Businesses where data protection and compliance are the primary AI risk drivers — healthcare, finance, legal services.

Model 2: Legal and Compliance

The pitch: AI regulation is real and growing. Legal and compliance teams understand regulatory environments and know how to read the fine print in vendor contracts.

What they do well:

-

Navigate complex regulatory requirements better than any other function.

- Translate dense terms-of-service and privacy policies into plain risk summaries.
- Build documentation trails that protect the business if regulators come knocking.

Where it breaks down:

- Legal teams often lack the technical depth to understand how AI systems actually work.
- Like security, legal is wired toward caution. Excessive risk-aversion can block beneficial AI adoption.
- Without technical fluency, legal teams can't assess real-world system impacts.

Best fit for: Regulated industries or businesses with significant contractual exposure.

Model 3: Privacy

The pitch: Most AI tools handle customer data in some form. Privacy teams are expert at data handling, consent, and customer rights.

What they do well:

- Address the intersection of AI and consumer data regulation directly.
- Build customer trust by ensuring AI tools handle personal information responsibly.
- Spot privacy-by-design gaps that technical or legal teams might miss.

Where it breaks down:

- Privacy-first thinking can underweight business value.
- Privacy teams may underestimate organizational risks beyond consumer data.
- Privacy governance alone doesn't cover AI risks unrelated to personal data: bias, accuracy, automation failures.

Best fit for: Consumer-facing businesses where data trust is a brand differentiator.

Model 4: Data Science or AI Teams

The pitch: Who understands AI better than the people building and deploying it?

What they do well:

- Evaluate AI capabilities and limitations with precision.
-

Identify deployment risks early, before systems go live.

- Assess vendor claims honestly.

Where it breaks down:

- Technical fluency doesn't automatically translate into compliance knowledge or security depth.
- AI and data science teams are often evaluated on deployment and performance, not risk avoidance.
- These teams can be siloed — optimizing the AI without sufficient input from business units.

Best fit for: Technology-forward businesses where AI performance is the primary concern.

Model 5: A Dedicated AI Governance Function

The pitch: If AI is strategic enough to invest in seriously, it's strategic enough to govern seriously.

What they do well:

- Build specialized expertise across the full spectrum of AI risk.
- Act as a neutral party between departments.
- Develop frameworks and policies that keep up with AI's pace of change.

Where it breaks down:

- Cost. Dedicated governance functions add overhead.
- Siloing risk. Isolating AI governance can disconnect it from day-to-day decisions.
- Coordination overhead with every other department.

Best fit for: Mid-size businesses with significant AI investment and regulatory exposure.

Practical Framework: Choosing the Right Model

Step 1: Inventory Your AI Tools

List every AI tool your business currently uses or is evaluating.

Step 2: Identify Your Primary Risk Driver

For most small businesses, one risk category dominates. Match your primary risk to a starting model.

Step 3: Assign a Single Owner

Even if a committee advises, one person should own AI governance.

Step 4: Write Three Sentences of Policy

1. Which AI tools are approved for use, and who approves new ones.
2. What data employees are prohibited from entering into AI tools.
3. Who to contact if an AI tool produces a result that seems wrong, biased, or harmful.

Step 5: Review Quarterly

Schedule a 30-minute quarterly review of your AI inventory and policy.

Key Takeaways

- **AI governance is not just for big companies.** Any business using AI tools that touch customer data needs basic governance.
- **Every ownership model has a tradeoff.** Know the tradeoff you're accepting.
- **Start small.** An AI tool inventory, a single owner, and three sentences of policy is a working governance program.
- **Review regularly.** AI governance is a process, not a document.