

YASH BHATT CONSULTING

Responsible AI Governance for Mid-Market Companies

A pragmatic governance framework that doesn't require a dedicated ethics team.

AI Strategy - Enterprise Solutions - Analytics
yash121r@gmail.com · [linkedin.com/in/yash-bhatt-anand](https://www.linkedin.com/in/yash-bhatt-anand)

The Mid-Market Governance Problem

Enterprise AI governance frameworks assume committees, ethics boards, and dedicated staff. Mid-market companies need the same risk protection with roughly two hours per week of overhead. This whitepaper describes a three-layer framework proven in growth-stage deployments.

Layer 1 - The AI Register (30 minutes to start)

A single shared document listing every AI system in use: what it does, what data it touches, who owns it, and its risk tier. Most companies discover 2-3x more AI usage than leadership expected - shadow ChatGPT usage included.

Risk Tiers

- Tier 1 (Low): internal drafting, summarization, no sensitive data. Governance: usage guidelines only.
- Tier 2 (Medium): customer-facing content, internal analytics, PII-adjacent. Governance: human review before external use, quarterly output audits.
- Tier 3 (High): decisions affecting people or money - pricing, hiring screens, credit, medical. Governance: documented review workflow, bias testing, named accountable owner, kill switch.

Layer 2 - Guardrails, Not Gates

Policies that block work get bypassed. Effective mid-market guardrails are defaults, not approvals:

- Approved-tools list with data boundaries per tool (what may and may not be pasted/uploaded).
- A standard disclosure line for AI-assisted external content.
- Prompt and output logging on Tier 2+ systems, sampled monthly.
- A 15-minute intake form for new AI use cases - triage into tiers within 48 hours.

Layer 3 - The Quarterly Review

One 60-minute meeting per quarter: review the register, sample Tier 2/3 outputs, check incidents, and retire what isn't delivering value. Attendees: the AI owner, one operations leader, one legal/compliance contact. Output: an updated register and a one-page summary for leadership.

Metrics Worth Tracking

- Coverage: % of known AI systems on the register.
- Incident rate: outputs escalated per 1,000 uses.
- Review latency: days from intake to tier assignment.
- Value check: each Tier 2/3 system re-justified annually against its ROI model.

Getting Started This Month

Week 1: create the register. Week 2: tier everything and publish the approved-tools list. Week 3: stand up the intake form. Week 4: schedule the first quarterly review. For a facilitated governance setup workshop: calendly.com/yash121r/30min