

FICTIONAL CASE STUDY

From Unmanaged AI Use to an Operating Governance System

How a 38-person law firm could move from scattered AI experimentation to controlled, measurable use in 12 weeks

Engagement	AI Governance Setup Engagement
Client profile	10 lawyers 28 support staff 38 total personnel
Illustrative fee	\$50,000 fixed fee
Duration	12 calendar weeks
Prepared by	CAVA Suite LLC

Case study purpose

This downloadable case study illustrates the type of work and management decisions that a CAVA Suite AI Governance Setup Engagement can produce. Harbor and Pine Law LLP is fictional. All findings, costs, scores, business cases, decisions, and outcomes are invented examples based on the governance delivery playbook and example deliverables package.

The engagement does not provide legal advice, cybersecurity testing, technology implementation, or full production validation. Client counsel and other qualified specialists retain responsibility for conclusions within their professional domains.

Executive Summary

The firm had already adopted AI before management had adopted AI governance. Lawyers and staff were using consumer tools and AI-enabled vendor features for research, drafting, summaries, administrative work, and experimentation. The activity created real potential value, but it also created uncertainty around data handling, verification, supervision, vendor terms, billing, incident response, and approval authority.

CAVA Suite's 12-week governance engagement was designed to turn that activity into a manageable operating system. The work combined current-state discovery, opportunity analysis, policy and decision design, vendor diligence, assurance planning, and an implementation roadmap.

Starting Point

Observed condition	Illustrative evidence	Management implication
AI use was widespread and unmanaged	18 of 31 survey respondents reported monthly use; 7 reported weekly use	Create an approved environment and register active uses
Matter content may have entered consumer tools	Four interviewees described pasting brief excerpts, discovery summaries, or case facts	Use confidential follow-up and counsel-directed incident triage
AI output supported substantive legal work	Associates used AI for issue outlines, counterarguments, and draft sections	Require source verification and responsible-lawyer review
Vendor diligence was incomplete	Individual and free accounts existed without an approved-service register or AI-specific terms review	Review vendor terms, configuration, and control evidence
Benefits were plausible but unmeasured	Users reported time savings, but the firm did not track AI-assisted work	Measure cycle time, correction effort, and pilot results

Readiness Snapshot

Dimension	Score	Illustrative basis
Leadership and accountability	1.5	No named owner or recurring forum
Policy and ethics	1.0	No AI-specific policy; obligations handled informally
Technology and security	1.5	No approved enterprise AI environment or configuration baseline
Data and records	1.5	Existing rules did not address prompts, outputs, or retention
People and skills	2.5	High interest, uneven ability to validate output
Measurement and assurance	1.0	No acceptance tests, monitoring, or change triggers

Overall illustrative readiness: 1.5 of 5, an ad hoc stage in which useful activity existed without management controls.

The 12-Week Engagement

The work was structured so each phase produced a management decision, an owner, and a usable artifact rather than a stack of disconnected recommendations.

Weeks	Focus	Representative outputs
1-3	Discovery and current state	Work plan, evidence register, use inventory, readiness score, gap register
4-5	Opportunity and business cases	Use-case portfolio, prioritization matrix, three initial business cases
6-8	Governance, policy, and approvals	Charter, RACI, risk tiers, acceptable-use rules, intake and approval toolkit
9	Vendor and procurement	Vendor questionnaire, evaluation checklist, RFP requirements, contract considerations
10	Assurance and validation design	Tier requirements, acceptance templates, monitoring and revalidation triggers
11	Roadmap and transfer	12-month implementation plan, three-year roadmap, budget priorities
12	Executive decisions and handoff	Executive brief, editable files, owner walkthrough, acceptance record

Governance Design

The proposed model was intentionally small. Routine low-risk decisions did not require a large committee, while higher-risk uses received stronger review.

Role	Primary responsibility
AI Steering Group	Approves policy, Tier 3 uses, exceptions, significant incidents, and annual priorities
AI Program Lead	Runs intake, inventory, approvals, training, monitoring, records, and routine Tier 1 decisions
Practice Leaders and System Owners	Own the business need, Tier 2 approval, acceptance evidence, human review, and performance
Lawyers and Support Personnel	Use approved workflows, verify outputs, protect data, and report issues
Specialist support	Counsel, IT, security, privacy, records, HR, and finance decide issues within their domains

What Changed in the Operating Model

The engagement translated broad concerns such as 'use AI safely' into specific operating rules and decision rights.

Illustrative Policy Direction

- Use only approved AI services and features through firm-managed accounts for firm work.
- Do not place client-confidential, privileged, protected, health, financial, or other restricted information into an AI service unless the service and workflow are specifically approved for that data.
- Treat generated legal content as a draft or research lead. A responsible lawyer must retrieve authorities, confirm quotations, verify factual statements and legal propositions, and approve final work product.
- Require lawyer supervision of nonlawyer AI use as part of ordinary matter supervision.
- Report accidental data entry, materially false output, security concerns, policy violations, or harmful reliance through a defined AI incident route.
- Require written approval and recorded conditions for exceptions.

Decision Rights

Trigger	Decision owner	Expected response
New Tier 1 use on approved platform	AI Program Lead	Routine review and record
New Tier 2 legal-work use	Practice Leader	Program Lead and IT consultation; counsel as needed
Tier 3 use or policy exception	AI Steering Group	Specialist consultation and formal approval
Suspected confidential-data disclosure	Managing partner delegate	Immediate escalation to counsel and relevant owners
Repeated material output failure	AI Program Lead may suspend	Immediate suspension and review

Opportunity Portfolio

The firm did not treat every AI idea as equally valuable or equally safe. Risk operated as a gate, not as a positive score.

Use case	Value	Risk	Illustrative decision
Internal procedure assistant	High	Low	Proceed
Billing narrative cleanup	Medium	Low	Proceed
Research starting points	High	Medium	Controlled pilot
Brief development	High	Medium	Controlled pilot
Case-file chronology	High	Medium	Pilot after vendor approval
Client email drafting	Medium	Medium	Limited pilot
Intake chatbot	Medium	High	Defer
Outcome prediction	Unknown	High	Do not pursue

From Opportunity to Business Case

Three candidate uses were developed far enough for leadership to compare expected value, cost, controls, and evidence requirements. The point was not to promise ROI. It was to define what would have to be true before the firm invested.

Candidate	Benefit hypothesis	Illustrative cost	Key controls	Decision
Approved legal research assistant	Reduce nonbillable search-framing effort	\$18,000 annual license and training allowance	Approved accounts, source retrieval, lawyer verification, monthly error review	60-day pilot with four lawyers
Brief development pilot	Reduce first-outline time by 25% without increasing corrections	\$9,000 licenses, configuration, training, and measurement	Matter eligibility, minimum necessary data, retention rules, supervising partner approval	Proceed after vendor diligence and policy approval
Internal procedure assistant	Save 12 staff hours per week and improve consistency	\$12,000 setup plus existing platform licensing	Approved source library, source citations, no matter data, owner review	Preferred first pilot

Vendor and Procurement Controls

The governance package also changed how the firm would evaluate AI vendors. Marketing claims were not treated as evidence. The review asked for concrete answers about model providers, retention, training, tenant isolation, logs, deletion, incident notification, change control, customer testing, output rights, liability, and known limitations.

- Identify all models, providers, hosting locations, and subprocessors.
- Document whether customer inputs or outputs are retained, reviewed by humans, or used for model improvement.
- Confirm administrative controls, logging, deletion, feature disablement, export, and rollback capabilities.
- Require task-specific performance evidence rather than generic accuracy claims.
- Route contractual, legal, ethics, privacy, and security conclusions to qualified client owners.

Assurance Before Scale

For higher-consequence uses, the firm received a governance-level assurance framework. It defined the task, users, operating conditions, failure consequences, required evidence, acceptance criteria, testing approach, monitoring, and revalidation triggers before production approval.

Requirement	Tier 1	Tier 2	Tier 3
Functional test	Basic	Documented	Independent or specialist-reviewed
Data and vendor evidence	Register check	Full review	Full review plus counsel/security approval
Acceptance criteria	Owner confirms	Written before pilot	Formally approved before testing

The Roadmap Leadership Could Execute

The final deliverable did not end with policy. It transferred ownership to the firm and sequenced the next year of work.

Period	Priority actions	Exit evidence
Days 0-30	Issue interim rule; appoint Program Lead; complete confidential follow-up; inventory tools and uses; adopt policy; freeze new AI purchases	Policy approved; inventory baseline; incident decisions recorded
Days 31-60	Complete vendor reviews; configure approved workspace; train personnel; open intake workflow; approve pilots	Approved register; training completion; signed pilot charters
Days 61-90	Run internal procedure and research pilots; monitor corrections, use, time, and incidents	Pilot scorecards and decision records
Months 4-6	Scale successful low-risk use; test brief-development controls; review billing and communication guidance	Quarterly governance report
Months 7-12	Consider chronology validation; address vendor gaps; refresh training; audit sample matters; complete annual review	Validation decisions; annual governance report; next-year plan

Executive Decisions at Handoff

The fictional engagement closes with a small number of explicit leadership decisions rather than an open-ended recommendation list.

- Approve the governance charter, acceptable-use policy, and AI Program Lead appointment.
- Require firm-managed approved accounts and prohibit personal or free accounts for firm matters.
- Authorize the internal procedure assistant and approved legal research assistant pilots, subject to vendor and acceptance gates.
- Defer client-facing automation, outcome prediction, and other Tier 3 uses.
- Fund first-year licensing, training, and pilot measurement only where the business case remains sound.

What the Engagement Produced

Work product	Management use
Current-state assessment and AI inventory	Know what is in use, by whom, and under what conditions
Opportunity portfolio and business cases	Decide where AI may create enough value to justify controlled evaluation
Governance charter, roles, RACI, and risk tiers	Assign authority and escalation
Policy and operating standards	Set enforceable rules for accounts, data, review, supervision, and reporting
Intake and approval toolkit	Route new uses through a repeatable decision process
Vendor and procurement toolkit	Demand evidence before purchase or expansion
Assurance and validation framework	Set acceptance, monitoring, and revalidation requirements
12-month plan and three-year roadmap	Sequence implementation, ownership, and investment

Business Value

The value of governance is not the number of policies written. It is the ability to make AI decisions consistently, preserve human accountability, stop uncontrolled use, and invest only where the evidence supports expansion.

Before	After the governance setup
No named AI owner	Program Lead and Steering Group with defined authority
Free and personal accounts used for firm work	Approved-account rule and managed-use register
Different users applied different verification practices	Defined human-review and source-verification requirements
Vendor AI features entered through ordinary procurement	AI-specific vendor questionnaire, evidence gates, and change review
AI ideas competed informally	Prioritized portfolio, risk tiers, business cases, and pilot gates

A Practical Boundary

This engagement establishes governance and transfers an operating model to the client. It does not configure software, conduct cybersecurity testing, issue legal or privacy opinions, or perform full validation of a production AI system. Those activities remain separate work packages when the evidence shows they are needed.

Is This the Right Next Step?

The AI Governance Setup Engagement is designed for organizations that have moved beyond the question of whether employees are using AI and need a durable way to decide what is permitted, what should be piloted, what should wait, and what evidence is required before scale.

CAVA Suite LLC

AI Governance Setup Engagement

Move from unmanaged experimentation to controlled, evidence-based AI adoption.

Next step: schedule a governance discovery call to determine whether your organization needs a full governance setup, a narrower roadmap, vendor evaluation, or validation engagement.