

AI Governance Setup

Build the operating system your organization needs to make AI decisions consistently

AI governance is more than a policy. A useful governance system defines who can approve what, which information may be used, how new use cases enter review, what vendors must prove, how pilot evidence is evaluated, when a system can move into production, and what happens when performance or conditions change.

The CAVA Suite AI Governance Setup is a complete, fixed-scope engagement designed to move an organization from scattered AI activity to a working governance and decision system that internal owners can operate after handoff.

Best fit	Organizations where AI use or planned adoption is material enough that leadership needs a durable operating model rather than isolated rules.
Typical engagement	Approximately 12 calendar weeks for a standard single-organization scope.
Core result	A complete governance package covering current state, strategy, roles, policy, intake, vendor review, assurance, roadmap, and executive decisions.
Commercial model	Fixed-scope project. A representative example is \$50,000 for a 38-person professional-services organization; actual scope and fee are confirmed for each client.

The Problem It Solves

Organizations often reach a point where basic AI rules are no longer enough. Leadership needs a repeatable answer to questions such as:

- Who owns AI governance and who has authority to approve a tool, use case, exception, pilot, or production deployment?
- Which uses are low risk, which require formal review, and which should not proceed?
- What data can be used in which systems and under what conditions?
- How should legal, security, privacy, records, finance, HR, procurement, and business owners participate?
- What evidence must a vendor or internal team produce before management approves an AI system?
- Which opportunities deserve investment, and which should be deferred or rejected?
- How will the organization monitor approved systems and respond to incidents or material changes?

Move from Policy to an Operating System

The objective of AI Governance Setup is not to leave the organization with a binder of policies. It is to establish a working decision system that managers, IT, counsel, business owners, and other specialists can actually operate.

Complete Deliverable Package

Package	What it covers	Representative artifacts
1. Strategy and current state	Current tools, uses, readiness, gaps, and immediate management direction.	Executive strategy, assessment, inventory, readiness scorecard, gap register
2. Opportunities and business cases	Bounded use cases, benefit hypotheses, risk gates, prioritization, and initial economics.	Use-case portfolio, prioritization matrix, business cases
3. Governance framework	Authority, roles, tiers, RACI, escalation, records, and meeting cadence.	Charter, decision-rights matrix, RACI
4. Policy and operating standards	Rules for approved services, data, human oversight, generative AI, acquisition, and use.	Policy suite and operating standards
5. Review and approval toolkit	How new tools and use cases enter review and receive a decision.	Intake form, risk method, checklist, approval workflow
6. Vendor and procurement toolkit	Data, security, performance, control, contract, and change evidence expected from vendors.	Questionnaire, evaluation checklist, RFP requirements, demo scorecard
7. Assurance and validation framework	Risk-tier evidence, acceptance criteria, testing, monitoring, change, and revalidation rules.	Assurance framework, templates, pilot candidate
8. Roadmap and executive handoff	Owners, sequencing, budget priorities, training, and first-year execution.	12-month plan, three-year roadmap, executive decision brief

How the 12-Week Engagement Works

Phase	Weeks	Primary result
Discovery and current state	1-3	Evidence baseline, inventories, readiness assessment, gap register
Opportunity portfolio	4-5	Prioritized use cases and initial business cases
Governance and policy design	6-8	Operating model, policy standards, intake and approval workflow
Vendor and assurance design	9-10	Procurement toolkit and governance-level validation framework
Roadmap and transfer	11	12-month plan, three-year roadmap, owners, budget priorities
Executive decisions and handoff	12	Final package, owner walkthrough, acceptance and decision record

What Makes the Engagement Different

- Evidence first. Policies are designed after the team understands actual tools, workflows, data, benefits, and control gaps.
- Decision rights are explicit. The package assigns real owners and defines when specialist consultation is required.
- Risk is proportionate. Lower-risk use can move through a lighter path while higher-consequence uses require stronger evidence and review.
- Business value is part of governance. Use cases are evaluated for measurable benefit, not only risk.
- Vendor claims are tested against evidence needs. Marketing statements do not substitute for data, contract, control, and performance evidence.
- The system is designed to transfer. Editable tools, owner walkthroughs, tabletop exercises, and acceptance criteria reduce dependence on the consultant.

What Is Not Included

- Legal opinions, breach determinations, or regulatory conclusions.
- Cybersecurity testing, penetration testing, or independent privacy assessment.
- Software configuration, migration, integration, or production implementation.
- Full production validation of a selected AI system, and large-scale training or change-management programs.

Where It Fits in the CAVA Suite Service Path

1. AI Exposure Audit. Use when management first needs visibility and a diagnostic baseline.
2. AI Governance Setup. Use when the organization needs a complete operating model and transferable governance system.
3. AI Validation. Use when a selected system or workflow needs measurable evidence before production, scale, or continued use.
4. Fractional Chief AI Officer. Use when leadership wants ongoing ownership of roadmap execution, governance operation, vendors, pilots, and recurring decisions.