

AI governance · Ownership

The AI Plan Is Funded. Who Governs It?

The roadmap is approved and funded, yet the enterprise cannot produce the rules and evidence that govern it. Six terrains, each named for the question leaders must answer, and one ownership test show where capability is outpacing control.

What you finish with: A one-page AI Governance Ownership Record for your organization: a named owner and a control state of Cleared, Evidence Pending, or Not Cleared for each of the six terrains, with the evidence, any accepted exception, and the re-examination date recorded.

[Open this research online](#)

GOVERNANCE / OWNERSHIP

Funding does not fill the ownership gap.

01

One accountable owner

A person who answers for the terrain.

02

A rule in force

Scope and effective date; applied in the work.

03

Current evidence

A record, log or test that can be produced.

Apply all three to each material terrain. States do not average across terrains.

01 Overview

What matters. A recurring AI governance failure appears after the plan is approved. The roadmap is funded, delivery is underway, and governance coverage does not match the capability being deployed. This note focuses on a less visible exposure: an organization with a credible strategy and no owned rules underneath it.

What it means. AI expands enterprise technology from executing predefined tasks into drafting, recommending, and taking bounded actions. Traditional technology controls still matter, but they do not by themselves resolve accountability for those outputs and actions. Governance here means named owners, rules in force, and evidence that can be produced on demand, not a compliance document.

What to do. Work the six terrains in this note. Convene one working session chaired by the executive accountable for the AI roadmap, with the finance, technology, data, security, and people leaders whose terrains are in scope. The method is small enough to hold in memory: six terrains, three questions, three control states. Each terrain carries diagnostic questions designed to expose whether ownership, rules, and evidence exist before an oversight review requires them, and each resolves to the same three-question ownership test. A terrain is either Cleared, Evidence Pending, or Not Cleared. No averaging: a strong answer to What Runs does not offset an unowned What It Costs.

02 The pattern in the field

These are practitioner archetypes from documented executive advisory work, not empirical maturity bands. This note addresses the middle one.

Tool-first. The organization buys capability, then looks for the problem. No strategy, no sequence, no rules. This note does not address that posture.

Intermediate. The organization has done real work: a multi-year roadmap, a governed data platform, named delivery teams, an executive narrative, often a quantified baseline justifying the investment. Its governance coverage has not kept pace with the capability being deployed, and asked directly whether the rules exist, the honest answer is no.

Governed. Rules are in force, owners are accountable, and current evidence is available. The test below is how to find out whether that describes you, rather than assuming it does.

Figure 1 places the three postures on one path, with the governance gap concentrated in the middle. The intermediate posture creates a less visible exposure precisely because it looks like progress. The plan survives every review that asks what the organization will do. It struggles in the first review that asks how the organization governs what it does, and that review arrives on the board's schedule, the auditor's, or a regulator's, not the roadmap's.

FIGURE 1 / PRACTITIONER ARCHETYPES

The funded middle can still be exposed.

Tool-first

Capability bought while the problem is still being sought. No strategy. No rules.

Intermediate

Funded roadmap, data foundation and delivery teams. Governance coverage lags capability.

Governed

Rules in force. Accountable owners. Current evidence available on demand.

Conceptual advisory archetypes, not measured maturity bands. Governance in force determines the posture, not spending.

03 Six terrains, one seam

M. Pact organizes AI governance into six terrains, each named for the question leaders must answer. They run in the order the work runs: first what is in scope, then who is accountable, then the information, the created artifacts, the resilience, and the economics.

The six terrains at a glance

Terrain	What it covers
1. What Runs	Every AI capability operating on company work: purchased, embedded by vendors, or adopted by employees without asking. Scope comes first, because you cannot govern an entity you do not know is operating.
2. Who Uses It	Use and disclosure, role definitions, performance assessment, capability development, and succession for the roles that will govern the rest. Accountability is human before it is technical.
3. What It Keeps	What data enters the capability, what providers or systems retain from its use, and what business records its outputs create. Depending on use and jurisdiction, those inputs and outputs may be subject to access, classification, retention, legal hold, discovery, or disclosure obligations; confirm the applicable requirements with counsel.
4. What It Builds	Machine-written software entering production, including through low-code environments used outside traditional development teams, with the quality, security, provenance, and human-review-capability questions that follow.
5. What Attacks	The threat surface generative capability opens on both sides: adversarial use against your systems, and fabricated content manufactured with the same tools.
6. What It Costs	Workflow-level economics per unit of work, the routing of work between deterministic automation and premium AI, and renewal exposure.

FIGURE 2 / SIX TERRAINS

Give the economic seam an owner.

TERRAIN	QUESTION IN THE WORK	CANDIDATE OWNER
What Runs	AI scope and approval control	CIO
Who Uses It	Disclosure and evaluation of work	CHRO
What It Keeps	Retention of outputs and business records	CDO
What It Builds	Code provenance and review standards	CTO
What Attacks	Adversarial use and fabricated content	CISO

06 What It Costs

Workflow economics: one accountable finance or technology executive.

Candidate roles are starting points to test, not universal assignments. Name the actual accountable person for every terrain.

Figure 2 places a candidate owner against five of the six terrains, offered as a starting point for your own test rather than as observed prevalence. What It Costs is shown without one, because in our advisory observation it sits in a seam. A common failure mode recurs. Finance treats workflow economics as a technology matter, technology treats it as a budget line finance approved, and no single accountable owner holds the answer.

The pressure behind that seat is documented. Gartner forecasts that AI inference costs per agentic workflow will increase more than fivefold through 2028, and separately puts the provider inference cost of routing a task to an agentic reasoning model at least five times that of a basic chatbot interaction. The forecast states no dollar baseline and does not address enterprise subscription or per-seat pricing.¹ M. Pact's inference is narrower: upstream cost pressure strengthens the case for an owner who can see workflow-level consumption, routing, contract terms, and renewal exposure. Without that owner, the organization may first confront workflow economics when a price change, renewal, or consumption report forces a budget decision.

The discipline that fills the seat is straightforward to state. Not every task deserves premium AI. Define which steps of a repeatable process run on deterministic automation and which need AI for analysis, recommendation, or drafting, and specify the human review at the point of judgment. Organizations

sometimes arrive at this pattern by accident. The point of governance is that accident is not a control. Document the routing criteria and assign an owner, so teams can apply them consistently and explain them in a budget review.

04 The ownership test

Three questions, applied to every terrain. Figure 3 shows the reading they produce.

Who owns it? One accountable name. A committee may decide; a person answers for it.

What is written? A rule in force today, with an effective date and a scope, not a policy under development.

What would you show an auditor tomorrow? Current evidence: a register, a log, a signed policy, a test result. If the answer is a meeting that happened, the terrain is not cleared.

FIGURE 3 / EXAMPLE

Control states never average into a score.

Terrain	Owner	Rule	Evidence	State
What Runs	Held	Held	Held	Cleared
Who Uses It	Held	Held	Unmeasured	Pending
What It Keeps	Held	Held	Cannot produce	Pending
What It Builds	Held	Held	Held	Cleared
What Attacks	Held	Held	Untested	Pending
What It Costs	Absent	Absent	Absent	Not Cleared

Pending means Evidence Pending. Scale only when every material terrain is Cleared, or a named risk owner accepts a bounded exception.

An exception retains the risk, containment and deadline; it does not turn the underlying control state into Cleared.

Diagnostic questions gather the evidence. The ownership test determines the terrain state. Any confirmed absence produces Not Cleared. Any unresolved required element produces Evidence

Pending. Only complete, current support across all material questions produces Cleared. An accepted exception does not change the underlying state; it authorizes bounded action against a named residual risk.

A terrain is **Cleared** only when all three hold. It is **Evidence Pending** when an element may exist but cannot currently be produced. It is **Not Cleared** when an element is absent. The states do not average across terrains: scale a workflow only where every terrain the workflow actually touches is cleared, judged by whether that terrain's failure could stop, reverse, or expose the workflow, or where an explicitly accepted exception names the risk owner, the containment, and the deadline.

One diagnostic question per terrain, with its evidence prompt. These six are enough to run the first pass: ask each one, and record the evidence produced or the reason it could not be.

What Runs. Who approves or rejects AI use cases? Show the approval conditions, deferrals, restrictions, or rejections produced by review. A register proves inventory, not decision control.

Who Uses It. What does the acceptable-use policy require people to disclose? Show the policy and the last measured compliance rate, not the signature page.

What It Keeps. Who decides how long AI-generated output lives and when it is destroyed? Show the rule that clearly covers AI-generated output, whether explicitly or through a defined record class, and evidence that retention or disposal has been applied when required.

What It Builds. How does anyone disclose that a machine wrote the code they shipped? Show the provenance standard in the workflow and a review log with it applied.

What Attacks. Who responds the day convincing fabricated content about your organization moves through the market? Show the runbook and its last test date.

What It Costs. Who owns the per-workflow economics of the AI estate? Show the latest workflow cost report with named routing decisions. Make gaps explicit in the first report, and assign an owner and a due date to resolve each one.

05 What It Costs, worked

What clearing this terrain looks like in practice. The owner is named. One finance or technology executive accepts accountability for workflow economics, whichever side of the seam the organization chooses, so long as the seat is filled. The rule is in force. Work routes to deterministic automation by default, premium AI is applied at the judgment steps on stated criteria, and any workflow whose unit cost moves past a set threshold triggers review. The evidence exists. A monthly per-workflow report showing consumption, the routing boundary for the three most heavily automated processes, the renewal calendar with concentration flagged, and the answer to one forward question: if a vendor repriced tomorrow, which workflows could move within a quarter, and which are constrained by integrations, configuration, workflow history, contract terms, or data portability. An organization that can produce that page has cleared What It Costs. The difficult work is deciding who owns it. Until that seat is filled, the rule and evidence have nowhere to live.

06 What not to do yet

Do not commission a governance framework document as the first move. A new document will not close an ownership gap by itself. Do not assume the transformation program absorbs this; in our observation, the assumption that governance lives inside some adjacent program is precisely how the gap survives. And do not treat the audit as the deadline. An audit can make the gap expensive; it did not create it.

07 Signals you are getting this right

Each terrain has a named owner who has accepted accountability. At least one rule has been tested against a live incident or a real budget decision. The What It Costs owner can state, for one process, where automation ends and AI begins, and why. The accountable owner can produce current evidence without reconstructing it after the request. And if no rule has ever changed, delayed, or stopped an action, test whether governance is operating or merely documented.

08 Run it now

A first pass of the test can be run from this note in one cross-functional working session: six terrains, three questions each, three states. Record the pass on one page, six rows, one per terrain, with these columns: the accountable name; the rule and its effective date; the evidence produced, or the reason it could not be; the resulting state; any accepted exception, with its risk owner, containment, and deadline; the remediation owner where the state is not Cleared; and the date the terrain is re-examined. That page is your AI Governance Ownership Record. Keep it, version it, and bring it to the next review rather than rebuilding it. When the record surfaces a gap, take first the terrain whose failure would be hardest to contain if it were tested next week by a board question, an incident, or a renewal: assign the owner, put one rule into force, test one evidence request, and record what re-opens the question.

09 The basis

This note presents a practitioner governance method, not an empirical maturity benchmark. The ownership test, the terrain architecture, and the accompanying record are M. Pact Advisory synthesis developed from cross-functional practitioner observation; public standards inform the dimensions of responsible AI governance, and the NIST AI Risk Management Framework and ISO/IEC 42001 reflect that AI governance spans inventory, data, security, development, and workforce dimensions. Neither prescribes this structure. The six terrains and the ownership test are M. Pact's synthesis for assigning enterprise accountability. Practitioner patterns are stated as observations, and external evidence supports only the claims specifically cited.

¹ Gartner, ["Gartner Predicts AI Inference Costs Per Agentic Workflow Will Increase More Than Fivefold Through 2028."](#) Gartner Newsroom, 17 August 2026. The release is addressed to product leaders building and operating agentic AI. It states no dollar base figure and makes no forecast of enterprise subscription or per-seat pricing.

This note is general executive guidance, not legal, accounting, or regulatory advice.

About this research

© 2026 M. Pact Advisory Services, LLC. All rights reserved.

This material is proprietary to M. Pact Advisory Services, LLC and is provided for executive advisory, strategic decision support, and internal planning purposes. It reflects practitioner judgment informed by experience across enterprise strategy, finance, technology, and operating models.

The frameworks and perspectives in this research are M. Pact Advisory's independent analysis and professional judgment at the time of publication. They are intended to inform executive thinking and decision-making, not to prescribe specific actions for any individual organization.

This research does not constitute legal, accounting, investment, or regulatory advice and should not be relied upon as such. Consult appropriate professional advisors before making decisions based on the information contained here.

The analysis is developed independently by M. Pact Advisory and is not influenced by third parties. M. Pact Advisory accepts no vendor compensation for placement, coverage, or rating. References to market practices, categories, or external concepts are illustrative and do not imply endorsement or affiliation.

This material may not be reproduced, distributed, or transmitted, in whole or in part, without the prior written consent of M. Pact Advisory Services, LLC.

This research may not be used as input for the training, fine-tuning, or development of generative artificial intelligence systems, machine learning models, algorithms, or related technologies without prior written permission.