



AI governance and responsible AI

Creating practical controls, ownership, and
evidence for safe organizational use.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 13 / EMBER & MINT

EXECUTIVE WORKING EDITION - 08 OCT 2026

13

AT A GLANCE

Creating practical controls, ownership, and evidence for safe organizational use.

AI governance is the set of decisions, owners, controls and records that let an organization use AI with confidence rather than hope. This playbook shows how to build a proportionate system: an inventory of what is in use, a risk-tiered review path, named accountable owners, monitoring and an incident process. It is written for leaders who need governance that people actually follow, and that produces evidence a customer, auditor or regulator can read. It is general guidance, not legal advice.

WHO IT IS FOR

Executives, risk and compliance leads, legal counsel, IT and data leaders, and business owners who must approve, oversee or defend how AI is used in their organization.

TIME TO USE

75 minutes to read; two working sessions to stand up a first version

What you will have at the end

- A clear definition of what AI governance must achieve in your organization and what it should leave alone
- A risk-tiered intake and review path that approves low-risk uses quickly and examines high-risk ones properly
- A named owner, reviewer and escalation route for every AI use case in production
- A starter AI inventory, vendor due diligence checklist and incident process
- A 30/60/90 plan that produces usable evidence within a quarter

Contents

01	Governance is how you say yes safely	03
02	What governance covers, and where it should stay out of the way	04
03	The building blocks of a working governance system	06
04	Roles, ownership and decision rights	07
05	The risks that matter and the controls that answer them	09
06	Evidence, metrics and how to know it is working	11
07	Standing up governance in 90 days	12
08	Common failures and a worked scenario	14
09	A first quarter that produces evidence	15
10	Terms used in this publication	16

EXECUTIVE BRIEF

Governance is how you say yes safely

01



Good AI governance speeds up responsible adoption by making the rules, owners and evidence clear before someone has to ask.

Most organizations meet AI governance in one of two bad ways. Either there is no process and people quietly use whatever tools they like, or there is a heavy committee that treats every chatbot request like a nuclear plant license. Both fail. The first creates unmanaged risk. The second pushes work underground and teaches staff to avoid asking. Practical governance sits between: a small set of clear rules, proportionate review and a visible record of decisions.

Responsible AI is the outcome governance is meant to produce: systems that are lawful, fair enough for their purpose, secure, transparent about what they are, and overseen by accountable people. It is not a slogan or a values poster. It is a set of behaviors you can inspect: who approved this use, what data it touches, how we know it works, who watches it, and what happens when it goes wrong.

- Governance answers four questions for every AI use: who is accountable, what could go wrong, what controls apply, and what evidence exists.
- It should be proportionate: the review for a meeting-notes assistant and for a credit decision support tool must not be the same.
- It covers AI you build, AI you buy, and AI that is already inside tools your people use every day.
- It needs an executive sponsor, because trade-offs between speed and risk are business decisions, not technical ones.

SCOPE NOTE

This playbook gives operating guidance. It does not provide legal advice. Laws and standards on AI, data protection, employment and sector rules change and differ by jurisdiction. Have qualified counsel confirm current obligations for your situation.



WORKING QUESTION

If a customer asked tomorrow how your organization decides which AI uses are acceptable, what could you show them today?

WHERE IT FITS

What governance covers, and where it should stay out of the way

02

Governance effort should rise with the consequence of a wrong output, not with how impressive the technology sounds.

The useful test is consequence. Ask what happens if the AI output is wrong, biased, leaked or acted on without checking. A tool that drafts internal meeting summaries has a low consequence of error. A tool that screens job candidates, sets prices for individuals, triages medical or safety reports, or influences who gets credit has a high one. Governance should be light for the first and rigorous for the second.

USE CASE	TYPICAL TIER	WHY
Drafting internal emails and summaries with approved tools	Low	A person reads and owns the result; no sensitive data leaves approved boundaries
Customer-facing assistant answering product questions	Medium	Wrong answers reach customers; needs testing, disclosure and monitoring
Summarizing contracts for legal review	Medium	Errors can mislead; a lawyer must verify before reliance
Screening or ranking job applicants	High	Affects individuals; employment and discrimination rules and fairness testing apply
Automated decisions on credit, insurance, benefits or health	High	Significant effect on people; strict legal and oversight requirements likely

USE CASE	TYPICAL TIER	WHY
Using a public chatbot with confidential client data	Prohibited unless approved	Data exposure risk; usually blocked by policy and tooling



What governance should not do

- Require a full review for every low-risk productivity use, which only creates shadow AI.
- Replace the engineers, product owners and managers who actually run the system with a committee that only meets monthly.
- Produce policy documents nobody reads while real usage goes unrecorded.
- Promise that risk is eliminated. The aim is known, owned and reduced risk.

Regulation shapes the high tier. The EU AI Act takes a risk-based approach with specific duties for high-risk systems and transparency duties for some others, and many jurisdictions are developing or applying their own rules. Existing law on data protection, consumer protection, employment, copyright and sector regulation already applies to AI use today. Check current requirements for the places you operate, and treat any framework as a structure for your own judgment, not a checklist that guarantees compliance.



WORKING QUESTION

Which two AI uses in your organization would land in the high tier, and does anyone with authority know they exist?



HOW IT WORKS

The building blocks of a working governance system

03

A governance system is a small number of linked parts: policy, inventory, intake, review, monitoring and incident response.

Think of governance as a lifecycle rather than a document. A use case is proposed, classified by risk, reviewed to the depth its risk deserves, approved with conditions, monitored in operation and eventually changed or retired. Each step leaves a record. The record is the evidence that turns a claim of responsible AI into something a third party can verify.

01 Set the policy

Write a short acceptable-use policy and a principles statement in plain language. State what is allowed, what needs approval, what is prohibited, and who to ask.

02 Build the inventory

Record every AI system, model, vendor feature and agent in use: purpose, owner, data, users, vendor, risk tier and review date. You cannot govern what you cannot list.

03 Run intake and tiering

Give teams a short form that asks what the use does, what data it touches, who is affected and what happens if it is wrong. The answers set the tier.

04 Review to the tier

Low tier is approved on a checklist. Medium adds testing evidence and a named owner. High adds legal, privacy, security and fairness review plus sign-off by an accountable executive.

05 Monitor and report

Track quality, incidents, complaints and drift. Report a short dashboard to the sponsor on a fixed cadence.

06 Respond and retire

Have a tested incident process, and a way to pause or switch off a system fast. Retire systems that are no longer used or reviewed.

Frameworks you can borrow from

You do not need to invent structure. The NIST AI Risk Management Framework organizes the work into govern, map, measure and manage functions. ISO/IEC 42001 describes an AI management system that can be audited and certified, in the same spirit as other management system standards. OWASP publishes guidance on security risks specific to applications built on large language models. Use these as references to shape your own controls, and check which versions and obligations currently apply.

PRACTICAL TIP

Start by mapping your existing risk, privacy, security and procurement processes. Most AI governance works best as an added set of questions inside those processes, not as a parallel bureaucracy.

WORKING QUESTION

Which step of this lifecycle exists in your organization today, and which exists only in people's heads?

OPERATING MODEL

Roles, ownership and decision rights

04

Governance fails when everyone is consulted and no one is accountable, so name the roles before you write the rules.

Keep the structure small. A governance forum that has fifteen members and meets quarterly cannot approve anything in time. Aim for a lightweight core group with real authority, delegated reviewers for each risk area, and a named business owner for every AI system. The business owner is the most important role, because someone must answer for the system when something goes wrong.

ROLE	RESPONSIBILITY	TYPICAL HOLDER
Executive sponsor	Sets risk appetite, resolves trade-offs, owns the annual report	COO, CRO, CIO or general counsel
AI governance lead	Runs intake, inventory and cadence, keeps the evidence file	Risk, compliance or an AI program manager
System owner	Accountable for one AI system's outcomes, users and ongoing quality	Business line manager
Technical reviewer	Checks architecture, testing, data flows and security	Engineering or data lead
Privacy and legal reviewer	Checks lawful basis, contracts, notices, sector rules	Privacy officer, counsel
Frontline reviewers	Check outputs in daily work and report issues	Users and team leads

Decision rights by tier

- Low tier: the system owner approves against the checklist, and the governance lead records it in the inventory.
- Medium tier: system owner plus technical and privacy reviewers approve, with conditions such as testing thresholds and disclosure wording.
- High tier: the forum reviews, the executive sponsor signs, and the decision records the residual risk accepted and who accepted it.
- Anyone can pause a system they believe is causing harm; restarting needs the system owner and the governance lead.

Write the decision rights down and publish them. A staff member who wants to try a new tool should be able to find, in one page, what they may do alone, what needs a form, and who answers within a stated time. A service level for review, such as an answer in five working days for medium tier requests, does more to reduce shadow AI than any warning memo.

WORKING QUESTION

For your most consequential AI use, can you name the single person who would answer for it if it harmed a customer?

RISKS AND CONTROLS

The risks that matter and the controls that answer them

05



Map each real risk to a control, an owner and a form of evidence, and skip controls that answer no risk.

Risk lists tend to be long and vague. Make them specific to the use case. For each risk, decide what control reduces it, who operates the control, and what evidence proves the control ran. A control without evidence is an intention. Evidence can be simple: a signed review form, a test report, an access log, a screenshot of a configuration, or a monitoring chart.

RISK	EXAMPLE CONTROL	EVIDENCE
Inaccurate or fabricated output	Grounding in approved sources, human review of consequential output, test set before launch	Test report; review sample logs
Unfair or discriminatory outcomes	Representative testing across groups, human decision-maker, ability to contest	Fairness test results; appeal records
Confidential data exposure	Approved tools only, data classification rules, vendor terms that restrict training on your data	Contract clauses; configuration records
Over-reliance and loss of skill	Training on limits, required checks, periodic spot audits	Training records; audit results
Unclear provenance and intellectual property	Vendor indemnity review, usage rules for generated content, disclosure rules	Vendor review; policy acknowledgment
Agents taking unintended actions	Least-privilege access, approval gates for irreversible actions, logging	Permission map; action logs



Human oversight that is real

Saying a human is in the loop is easy. Making it effective is not. Reviewers who see hundreds of outputs per hour approve almost everything, a pattern called automation bias. Design oversight so the reviewer has the time, information and authority to disagree. Give them the

source evidence, show confidence or uncertainty where possible, sample for quality rather than reviewing everything shallowly, and track how often reviewers override the system. An override rate of nearly zero can mean the AI is excellent, or that nobody is really checking.



Human oversight design check

- The reviewer can see the evidence behind an output, not only the output itself
- The reviewer has authority and time to reject or change it without penalty
- Override and rejection rates are tracked and discussed
- The affected person can contest a decision and reach a human
- Reviewers are trained on the system's known weak spots



TRANSPARENCY

Tell people when they are dealing with AI, when AI-generated content is used in a way that matters, and how to reach a person. Check current disclosure requirements for your market, since some jurisdictions mandate it for certain uses.



WORKING QUESTION

For your highest-risk AI use, which control would you be least comfortable explaining to a regulator, and what is the evidence it actually ran last month?



MEASURING GOVERNANCE

Evidence, metrics and how to know it is working

06



Measure whether governance is covering reality and reducing risk, not how many policies you have published.

Establish a baseline first. Count the AI systems you know about, survey teams about the tools they use, and review network or expense data for AI services nobody registered. The gap between the official inventory and what you find is your first and most honest metric. Then track a short set of indicators every month and share them with the sponsor. Keep the numbers few and tied to decisions.



METRIC	WHAT IT TELLS YOU	HOW TO MEASURE
Inventory coverage	Share of AI use that is known and registered	Registered systems divided by systems found via survey and spend review
Time to decision	Whether governance blocks or enables	Days from intake submission to decision, by tier
Controls with evidence	Whether controls actually run	Sampled controls with dated evidence divided by controls sampled
Incidents and near misses	Whether staff report problems	Count, severity, time to contain, time to fix
Review overrides	Whether human oversight is genuine	Share of AI outputs changed or rejected by reviewers
Overdue reviews	Whether systems stay in date	Systems past their scheduled review date



Do not treat a rising incident count as failure. In a young program, more reported incidents often mean people trust the process and use it. Watch the trend in severity and in time to contain. Run a tabletop exercise at least once a year: invent a plausible failure, such as an assistant giving a customer wrong terms, and walk through who decides, who is told and how the system is paused.



ILLUSTRATIVE NUMBERS ONLY

If a survey finds 40 AI tools in use and the inventory lists 12, coverage is 30 percent. Those figures are an example, not a benchmark. Your own gap is the number that matters.

WORKING QUESTION

What is the gap between the AI tools you have approved and the AI tools your people actually use, and how would you find out?

IMPLEMENTATION GUIDE

Standing up governance in 90 days

07

Build the smallest system that gives you visibility and control, then tighten it as use grows.

01 Appoint the sponsor and lead

Name an executive sponsor and a governance lead with time allocated. Agree the risk appetite in one page: what you will do, what you will not do, and what needs special approval.

02 Discover current use

Run a short survey, review software spend and talk to department heads. Ask openly and promise amnesty for past use, so people tell the truth.

03 Publish a short acceptable-use policy

Cover approved tools, data that must never go into unapproved tools, human review expectations and how to request something new. Keep it to two pages.

04 Launch intake and tiering

Use one form and a decision rule. Review the first ten requests live with the forum to calibrate tier boundaries.

05 Add vendor due diligence

Add AI-specific questions to procurement and renew existing vendors against them on their next contract cycle.

06 Start monitoring and incident handling

Define what each owner reports, how incidents are raised and who can pause a system. Run one tabletop exercise.

Vendor due diligence questions

- Is our data used to train the vendor's models, and can that be switched off in the contract?
- Where is data processed and stored, and which sub-processors are involved?
- What security certifications or audit reports can the vendor share, and are they current?
- What testing has the vendor done for accuracy and bias on tasks like ours, and can we see it?
- How are model changes announced, and can we pin or test versions before they affect us?
- What are the incident notification terms and time limits?
- What are the terms on intellectual property, indemnity and output ownership?
- How can we export our data and exit?

START WITH WHAT EXISTS

Reuse your current risk register, change process, privacy impact assessments and vendor onboarding. Add AI questions to them. A parallel system doubles the work and halves the compliance.

WORKING QUESTION

Which existing process in your organization could carry the AI intake questions starting next month?

PITFALLS

Common failures and a worked scenario

08

Most governance failures are design failures, and each has a predictable fix.

- Policy without inventory: rules exist but nobody knows what they apply to. Fix by building the inventory first.
- One-size review: every request goes through the same slow path. Fix with tiers and a fast lane.
- Committee ownership: a group is accountable, so nobody is. Fix with a named system owner.
- Paper compliance: forms are completed but controls never run. Fix by sampling evidence.
- Vendor trust by default: assuming the supplier handled risk. Fix with due diligence and contract terms.
- Ignoring embedded AI: features switched on inside existing software update the risk without any review. Fix by tracking vendor release notes.

ILLUSTRATIVE SCENARIO (HYPOTHETICAL)

A regional services firm bans all generative AI by memo. Within months, staff paste client documents into personal accounts to keep up. A survey after an incident finds many unregistered tools. The firm replaces the ban with an approved tool, a two-page policy and a five-day review service level. Visible use rises, risky use falls.

The scenario shows a pattern worth remembering. Bans without alternatives move risk out of sight. Governance earns trust by offering a safe, fast path that is easier than the workaround. When designing yours, ask what the person with a deadline will do if your process is slow, and make sure the answer is not the thing you wanted to prevent.

Warning signs your governance is performative

- The inventory has not changed in three months.

- No request has ever been declined or sent back for changes.
- Nobody can name the owner of a live AI system.
- Incident reports are zero, in an organization with hundreds of AI users.

WORKING QUESTION

If your approval process is slower than a workaround, which workaround are your people using right now?

30 / 60 / 90

A first quarter that produces evidence

09

The first ninety days should produce visibility, owners and one reviewed high-risk use, rather than a perfect framework.

First 30 days: ground it

- Name the executive sponsor and governance lead
- Survey current AI use and compile a first inventory
- Publish a two-page acceptable-use policy
- Agree the three risk tiers and who decides each

Days 31 to 60: test it

- Launch the intake form and review the first ten requests
- Add AI questions to vendor due diligence
- Assign a named system owner to every live system
- Run one tabletop incident exercise

Days 61 to 90: decide

- Complete a full review of one high-tier use with sign-off
- Report the first governance metrics to the sponsor
- Sample controls for evidence and fix the gaps
- Set the annual review calendar and training plan

GLOSSARY

Terms used in this publication

10

TERM	MEANING
AI inventory	A maintained register of AI systems in use, with purpose, owner, data, vendor and risk tier.
Risk tier	A category such as low, medium or high that sets how deep review and monitoring must be.
System owner	The business person accountable for one AI system's outcomes and ongoing quality.
Human oversight	Design that lets an informed person review, challenge and override AI outputs.
Automation bias	The tendency to accept machine output without enough checking.
Shadow AI	AI tools used for work without the organization's knowledge or approval.
Model drift	A change over time in model behavior or input data that degrades performance.
NIST AI RMF	A voluntary framework for managing AI risk, organized around govern, map, measure and manage.
ISO/IEC 42001	An international standard for an AI management system that organizations can implement and be audited against.
Residual risk	The risk that remains after controls are applied, and which someone must knowingly accept.
Tabletop exercise	A discussion-based rehearsal of an incident to test roles and decisions.



TERM	MEANING
Vendor due diligence	Structured questions and evidence reviewed before and during use of a supplier's AI.



A note on use

This playbook is general guidance and not legal advice. Confirm current legal and regulatory requirements with qualified advisers.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.

