

The Jonesi Consultancy | Houston, TX

EXECUTIVE AI ADOPTION · GOVERNANCE GUIDE

AI GOVERNANCE STARTER GUIDE



8 Questions to Answer Before You Write an AI Use Policy

A policy should document decisions your organization has already made.

WHY GOVERNANCE COMES BEFORE POLICY

Most organizations approach AI governance backwards. Somebody is assigned the policy, finds a template, fills in the blanks, and circulates it for acknowledgment. The document is competent. Then it turns out that nobody had actually decided which tools were approved, where human judgment was required, or who answers the question when an employee is unsure. A policy is a document. Governance is the set of decisions, responsibilities, boundaries, and review practices that make the document usable. Answer the eight questions in this guide with the right people in the room and the policy becomes a transcription exercise. Skip them and you will write a document that governs an imaginary organization.

THE CENTRAL PRINCIPLE

An AI use policy should document decisions your organization has already made. It should not be the place where you discover what those decisions are.

1 The Governance Stack

Eight layers, each depending on the one above it. Most policies address the middle and skip the ends.

#	LAYER	THE QUESTION IT ANSWERS
1	 PURPOSE	Why are we using AI at all?
2	 BOUNDARIES	What is allowed, restricted, and prohibited?
3	 DATA	What information may be used, and with what?
4	 HUMAN OVERSIGHT	Where must judgment remain human?
5	 VERIFICATION	What has to be checked before output is used?
6	 ACCOUNTABILITY	Who owns the outcome?
7	 ESCALATION	What happens when something is unclear or goes wrong?
8	 REVIEW	How will this evolve as the tools do?

WHERE ORGANIZATIONS USUALLY START

At layer two. Boundaries are the visible part of governance and the easiest to write, so they get written first. A boundary without a stated purpose behind it reads as an arbitrary restriction, and arbitrary restrictions are the ones people work around.

1 What Are We Actually Governing?

The word “AI” now covers a chatbot, a summarization button inside software you already own, an automated workflow, a custom assistant somebody in marketing built, and the analytics feature your finance platform added last quarter without telling anyone. A policy that says “AI tools” without saying which of those it means will be read narrowly by anyone who wants to.

WHAT YOU ARE ACTUALLY DECIDING

- › Whether this covers general-purpose assistants only, or also the AI features embedded in software you already own.
- › Whether custom assistants and agents built by employees are in scope.
- › Whether AI-assisted analytics and automation count.
- › Whether the policy follows the person or the tool, meaning whether personal accounts used for work are covered.

WATCH FOR: Embedded features are the most commonly missed category, and the most likely to already be in use across your organization without a decision having been made.

2 What Kinds of Use Are Acceptable, Restricted, and Prohibited?

“Employees may use AI responsibly” is not a boundary. It transfers the entire judgment to the individual and gives you nothing to point at afterward. People need categories they can recognize on a Tuesday afternoon without calling anyone.

WHAT YOU ARE ACTUALLY DECIDING

- › Which categories of work are fine without a second thought.
- › Which require an approved tool, an organizational account, or a review step.
- › Which are prohibited outright, and whether the list is short enough to be remembered.
- › Who has authority to approve an exception, and how fast they can answer.

WATCH FOR: A prohibition list longer than about six items stops functioning as a rule and starts functioning as a document nobody finishes reading.

3 What Information Can Employees Put Into AI Systems?

This is the section that prevents the incident. Categories like “sensitive information” get interpreted generously by people under time pressure, so the useful version names actual categories in your actual organization.

WHAT YOU ARE ACTUALLY DECIDING

- › Where the line falls for client, customer, patient, and case information.
- › How personally identifiable, financial, and health-related information are treated.
- › What happens with proprietary material, unpublished results, and anything under a confidentiality agreement.
- › Whether the rule changes for an enterprise account with different data-handling terms.

WATCH FOR: Give people a category that is explicitly permitted. A policy that only lists what is forbidden leaves cautious employees assuming everything is.

4 Where Must a Human Remain Accountable?

Some decisions cannot be delegated to a system regardless of how good the output looks. The work here is deciding which, in your organization, before somebody decides it for themselves at four o'clock on a Friday.

WHAT YOU ARE ACTUALLY DECIDING

- › Which decisions about employment, compensation, and discipline require documented human judgment.
- › Where financial, legal, clinical, and regulatory work requires a qualified person to do the substantive part.
- › What client-facing work cannot be released without a named reviewer.
- › Whether any category of work should not be AI-assisted at all.

WATCH FOR: “A human reviews it” is not the same as “a human exercises judgment.” A person clicking approve satisfies the first and not the second.

5 What Must Employees Verify Before Using AI Output?

Verification is where most policies become decorative. “Check AI output for accuracy” is an instruction nobody can follow, because it does not say what to check or how much checking is enough for the work in front of them.

WHAT YOU ARE ACTUALLY DECIDING

- › What must be verified in every case: factual claims, figures, citations, names, and dates.
- › Which categories of work require verification proportionate to the stakes rather than uniform across the board.
- › What a reasonable verification effort looks like, described concretely enough to be done.
- › Whether verification gets recorded anywhere, and for which work.

WATCH FOR: Uniform verification requirements produce over-checking on low-stakes work and under-checking on high-stakes work, because people calibrate to the average.

6 When Should AI Use Be Disclosed?

Ambiguity here is what makes people conceal their usage, which is the outcome you least want. The question is not whether to disclose but where a reasonable person would expect to know.

WHAT YOU ARE ACTUALLY DECIDING

- › Whether disclosure is required for external publications, client deliverables, and funder or grant submissions.
- › How images, voices, and likenesses are handled, which is usually a different answer than text.
- › Whether internal drafting, summarizing, and formatting need disclosure, and the reasoning behind that answer.
- › What the standard wording is, so five people do not invent five versions.

WATCH FOR: Check your client contracts and funder agreements before writing this section. Several of them probably already say something, and it will bind you regardless of what your policy says.

7 What Happens When Something Goes Wrong, or Someone Is Unsure?

Assume a mistake will happen. What determines the cost is whether people report it quickly or hide it, and that is decided entirely by what they expect to happen when they tell you.

WHAT YOU ARE ACTUALLY DECIDING

- › Who receives a report, by name or role, and how quickly they can respond.
- › What timeframe is expected for reporting, and whether people know it.
- › Whether reporting in good faith is explicitly protected, in writing.
- › Where an employee goes with a question that is not yet a problem, which is the more common case and the one usually left out.

WATCH FOR: A reporting route nobody has used in a year is not evidence of a flawless organization. It is evidence that people do not believe reporting is safe.

8 Who Owns Governance, and How Will It Evolve?

A policy written today describes a tool landscape that will have changed within two quarters. Without a named owner and a review date, it becomes accurate on the day it was signed and quietly wrong thereafter.

WHAT YOU ARE ACTUALLY DECIDING

- › Who owns the policy by name, not by department.
- › Who approves new tools, and what a request looks like.
- › How often it is reviewed, and what triggers an off-cycle review.
- › How employee experience feeds back into the policy, given that they will find the unworkable parts before anyone else does.

WATCH FOR: Governance assigned to legal or IT alone reliably produces rules that are technically sound and operationally unworkable, because the people who do the governed work were not in the room.

LEADERSHIP CHECKPOINT

Before drafting begins, leadership should be able to explain the scope, the boundaries, the data rules, and the decisions that still require human authority in plain language.

A USEFUL TEST

If these eight questions generate disagreement, that is not a policy-writing problem. It is a governance decision that should be resolved before the policy is finalized.

What the document should make clear once governance decisions have been made

This is the anatomy rather than the text. Once the eight questions are answered, a policy covering these components will be a fair record of what you decided.

COMPONENT	WHAT IT NEEDS TO ESTABLISH
PURPOSE AND SCOPE	Why this exists and who it covers, including contractors and volunteers.
APPROVED AND PROHIBITED USES	Recognizable categories rather than a standard of reasonableness.
TOOLS AND APPROVAL AUTHORITY	Which tools, and who decides about new ones.
DATA AND CONFIDENTIALITY	What may and may not be entered, in your own categories.
HUMAN REVIEW REQUIREMENTS	Where a person must decide, and what deciding means.
ACCURACY AND VERIFICATION	What gets checked, proportionate to the stakes.
BIAS AND APPROPRIATE USE	Where output could affect people unevenly, and what to do about it.
DISCLOSURE	When AI involvement is stated, to whom, and in what words.
ACCOUNTABILITY	Who owns the output. Usually the person who sent it.
INCIDENT REPORTING AND ESCALATION	Where to report, how fast, and what protection applies.
TRAINING RESPONSIBILITIES	Who is trained, on what, and by when.
POLICY OWNERSHIP	A person, not a function.
REVIEW AND REVISION	A cadence and a version history that is visible.

WHAT POLICY ALONE CANNOT DO

- › Build employee skill or confidence.
- › Ensure managers reinforce the rules consistently.
- › Create psychologically safe experimentation by itself.
- › Make unclear organizational decisions disappear.
- › Redesign workflows around AI capability.
- › Measure whether AI use is producing value or reducing risk.

POLICY SETS BOUNDARIES. GOVERNANCE KEEPS THEM USABLE.

The strongest policies are connected to training, leadership behavior, approved tools, workflow practices, feedback channels, and ongoing review.

4 5 Governance Mistakes That Create More Risk

Common ways organizations unintentionally make AI use harder to govern

1 Writing the Policy Before Understanding How People Actually Use AI

The resulting policy governs an imaginary workplace. It prohibits things nobody was doing and says nothing about the thing half the organization does weekly. Find out first, without blame, because an accurate answer is worth more than a compliant one.

2 Making the Policy So Restrictive That People Hide Their Use

Prohibition does not stop usage. It moves usage onto personal accounts and personal devices, where you have all the exposure and none of the visibility. A policy that produces shadow AI has made the organization less safe than having no policy at all.

3 Saying “Human Review Required” Without Defining Meaningful Review

A person clicking approve is technically human-in-the-loop. It is not human judgment. If the policy does not say what reviewing involves, the review will become a signature within about a month.

4 Treating Governance as a Legal or IT Responsibility Alone

AI changes everyday work, which makes operations, HR, managers, and frontline staff stakeholders rather than audiences. Rules written without them turn out to be unworkable in ways that are obvious downward and invisible upward.

5 Publishing the Policy and Considering Governance Finished

Tools change, risks change, and your organization learns things. A policy with no owner and no review date degrades quietly, and the first sign is usually an incident it should have covered.

A MORE DURABLE GOVERNANCE POSTURE

Good governance does not mean eliminating uncertainty. It means creating enough shared structure that employees know the boundaries, understand their responsibilities, can ask questions safely, and know when a decision needs stronger review.

THE GOAL IS NOT MAXIMUM RESTRICTION

The goal is disciplined use: clear enough to reduce avoidable risk, practical enough to support real work, and adaptable enough to change as the organization learns.

5 The Policy Test

Could an employee answer these questions after reading your current policy?

Print this page and hand it to someone who did not help write the policy. Have them write the answer in their own words. Blank lines are the finding.

NAME _____ TEAM _____ DATE _____

1 Which AI tools may I use for work?

2 What information may I share with them?

3 What kinds of use are prohibited or require additional approval?

4 What tasks should I never delegate entirely to AI?

5 What am I responsible for verifying before I use or share AI output?

6 When must a human make the final decision, and when should I disclose that AI was involved?

7 Who do I ask when I am unsure?

8 How do I report a mistake, suspected data exposure, or other concern?

IF EMPLOYEES CANNOT ANSWER THESE QUESTIONS...

The organization may have a document, but it does not yet have usable governance. The next task is not necessarily to add more policy language. It is to clarify the decisions, responsibilities, training, and support that make the policy operable.

6 Before You Finalize the Policy

- › Validate the policy against real workflows and realistic employee questions.
- › Test whether managers can explain the policy consistently.
- › Plan training and reinforcement before launch.
- › Confirm that approved tools and data rules are operationally workable.
- › Create a safe route for questions, incidents, and edge cases.
- › Assign an owner and a review date.

7 What Comes After the Questions

Answering these eight will take one meeting with the right people or three meetings with the wrong ones. The right people are an executive sponsor, an operations lead, someone from IT or security, and at least one manager who does the work being governed. Include the last one and the policy survives contact with reality.

Once the answers exist, writing the policy is the easy part, and it is worth having counsel review the result before publication, particularly the data, disclosure, and confidentiality sections. What counsel cannot do is make the decisions for you. That is the work this guide is for.

EXECUTIVE SPONSOR

Owns the purpose and settles disagreement in the room.

OPERATIONS LEAD

Knows which workflows the rules will actually touch.

IT OR SECURITY

Knows what the tools do with your data and what is already in use.

A LINE MANAGER

Finds the unworkable parts before your employees do.

THE CLOSING PRINCIPLE

Governance works when your people can translate organizational expectations into everyday decisions. Policy is one part of that system. The larger work is creating the clarity, capability, accountability, and feedback that let people use AI responsibly in real workflows.

Policy sets boundaries. Governance keeps them usable.

General educational resource. Not legal, regulatory, cybersecurity, privacy, or compliance advice.

© 2026 The Jonesi Consultancy. All rights reserved.

Working through these eight with your leadership team.

We facilitate the decisions, then help you write and train to the result.

