

The Jonesi Consultancy | Houston, TX

EXECUTIVE AI ADOPTION · POLICY GUIDE

THE ANATOMY OF A PRACTICAL AI USE POLICY



12 Components, What Each One Has to Accomplish, and
Where They Usually Fail

*A policy should document governance decisions, not substitute for
them.*

AN ANATOMY, NOT A TEMPLATE

Most AI use policies contain the right sections and still fail to change behavior. The sections are present, correctly titled, and written in language that satisfies a legal review. What is missing is the thing each section was supposed to accomplish. Use this to audit a policy you already have, to check whether a draft covers what it needs to, or to know what to ask for when somebody else is writing it. Each component includes the common weakness, because that is where the useful information sits. There is no sample language here, because the language is the easy part.

1 Three Tests for Practicality

A practical policy gives employees enough clarity to make responsible everyday decisions. It should be readable by the people expected to follow it, not only by the people who drafted it. Three tests decide whether it is.

FINDABLE

Can an employee locate the answer to a common AI question quickly?

UNDERSTANDABLE

Can a non-technical employee explain the rule in plain language?

ACTIONABLE

Does the policy tell the employee what to do when the answer is unclear?

A WRITING TIP WORTH FOLLOWING

Describe the decision employees must make, not only the principle behind it. “Protect confidential information” is a principle. “Do not enter client-identifiable information into an unapproved AI tool” is an actionable boundary.

BEFORE THE LANGUAGE IS DRAFTED

The organization should already have made core decisions about acceptable use, data handling, human oversight, verification, accountability, and who owns governance over time. A policy records those decisions. It cannot make them on your behalf.

Define the boundaries before you define the consequences. These four decide what the policy is about and what it actually permits.

1 PURPOSE AND SCOPE

Why the policy exists and exactly who it binds.

WHEN IT WORKS

It opens with a purpose rather than a warning, and it names whether contractors, volunteers, interns, and board members are covered. It also settles whether personal devices and personal accounts used for work are in scope.

COMMON WEAKNESS

Scope is left implicit. Six months later a contractor does something the policy would have prohibited, and nobody can say whether it applied to them.

2 APPROVED AND PROHIBITED USES

Recognizable categories a person can apply without calling anyone.

WHEN IT WORKS

Three tiers work better than two: fine without thinking, allowed under conditions, and never. The never list stays short enough to be remembered, and the categories describe work situations employees can identify.

COMMON WEAKNESS

"Employees may use AI responsibly." This transfers the whole judgment to the individual and gives the organization nothing to point at afterward.

3 APPROVED TOOLS AND APPROVAL AUTHORITY

Which tools are sanctioned, and who decides about new ones.

WHEN IT WORKS

A named list, a stated home for that list, an organizational account requirement where it matters, and a route for requesting something new that gets answered in days rather than quarters. It also says what to do when an embedded AI feature appears in software you already own.

COMMON WEAKNESS

No list at all. The most common single gap, and the one that most reliably produces unsanctioned use, because people who are not told what is approved assume nothing is.

4 DATA AND CONFIDENTIALITY EXPECTATIONS

What information may and may not be entered, in your own categories.

WHEN IT WORKS

Actual categories from your actual organization: client records, case notes, unpublished results, employee files, credentials, and anything client-restricted. Plus an explicitly permitted category, so cautious people know where they can stand.

COMMON WEAKNESS

Only prohibitions. A policy that lists what is forbidden and nothing else leaves careful employees assuming everything is, which is why your most conscientious staff are often your least engaged.

3

Components 5–8: Accuracy, Fairness, Disclosure

These four clarify what happens once AI is inside the workflow. They are also where policies most often become decorative.

5 HUMAN REVIEW REQUIREMENTS

Where a person must decide, and what deciding actually involves.

WHEN IT WORKS Review effort scaled to stakes, with the highest tier naming what a qualified person must do rather than merely that someone must look. It also names which decisions must remain human-led outright.

COMMON WEAKNESS “Human review required” with no definition. A person clicking approve satisfies it. Within a month that is exactly what it becomes.

6 ACCURACY AND VERIFICATION EXPECTATIONS

What has to be checked before output counts as work.

WHEN IT WORKS Named things: factual claims, figures, citations, quotations, names, dates. Plus an honest statement that models produce confident wrong answers as a property rather than a defect, and heavier expectations for externally facing content.

COMMON WEAKNESS “Check for accuracy.” Nobody knows what that means in practice, so in practice it means reading it once.

7 BIAS, FAIRNESS, AND APPROPRIATE USE

Where output could affect people unevenly, and what to do about it.

WHEN IT WORKS It names the places this actually bites in your organization: hiring, eligibility, service allocation, lending, discipline. It also says which uses require additional scrutiny and where a concern gets raised.

COMMON WEAKNESS A paragraph of principle with no application. It signals awareness and changes no decision.

8 DISCLOSURE EXPECTATIONS

When AI involvement is stated, to whom, and in what words.

WHEN IT WORKS A clear line across internal work, client deliverables, public communications, and grant or regulated content, with standard wording so five people do not invent five versions, and a check against existing client and funder agreements that may already say something.

COMMON WEAKNESS “Use your best judgment.” Sounds empowering, delivers nothing, and guarantees inconsistency across the organization.

9 ACCOUNTABILITY*Who owns the output.*

WHEN IT WORKS One sentence, stated plainly: the person who sends it owns it, exactly as with work produced any other way. Manager and subject-matter approvals stay where they already are, and AI holds neither authorship nor liability.

COMMON WEAKNESS Left unstated, which lets people develop private theories. The most common private theory is that responsibility is shared, which in practice means nobody feels like the last line.

10 INCIDENT REPORTING AND ESCALATION*Where to report, how fast, and what protection applies.*

WHEN IT WORKS A named recipient, a stated timeframe, explicit protection for good-faith reporting, a note on when work should pause pending review, and a separate route for questions that are not yet problems.

COMMON WEAKNESS A reporting route with no protection clause. People weigh the personal cost of reporting against the chance of being found out, and conceal accordingly.

11 TRAINING RESPONSIBILITIES*Who is trained, on what, and by when.*

WHEN IT WORKS A commitment with dates, covering new hires as well as current staff, a statement of what managers are responsible for reinforcing, and a refresher cadence tied to policy changes.

COMMON WEAKNESS Omitted entirely. The policy assumes a level of competence the organization has not funded, which is how a reasonable rule becomes an unreasonable expectation.

12 OWNERSHIP, REVIEW, AND REVISION*Who owns this document and how it stays current.*

WHEN IT WORKS A person by name, a review cadence, a visible version history, and a route for employees to report that a rule is unworkable. New tools, incidents, and lessons learned all trigger an update.

COMMON WEAKNESS No owner and no date. The policy is accurate the day it is signed and quietly wrong within two quarters, usually discovered by an incident it should have covered.

ONE MORE WORTH INCLUDING: DEFINITIONS

A short shared vocabulary, not a technical glossary. Define only the terms employees need in order to understand the expectations: generative AI, AI-assisted work, approved tool, confidential information, human review, and consequential decision.

5

Reading a Policy Quickly

If you have been handed a policy and have ten minutes, these five checks will tell you most of what you need to know.

1 Search for a Tool Name.

If the document never names a specific tool, there is no approved list, and unsanctioned use is near certain.

2 Look for a Category That Is Explicitly Allowed.

Prohibitions without permissions produce cautious paralysis in exactly the people you most want engaged.

3 Find the Word “Review” and Read the Sentence Around It.

If it does not say what reviewing involves, the requirement will decay into a signature.

4 Look for a Person's Name or a Role.

Ownership, approval authority, and incident reporting all need a recipient. A policy with no names has no operating mechanism.

5 Check the Date and the Version History.

A policy older than about two quarters was written for a different tool landscape.

WHY THERE IS NO SAMPLE LANGUAGE HERE

Policy language that fits your organization comes from decisions your organization has made, and a clause borrowed from a template describes somebody else's decisions. The wording is also the part your counsel should shape, since obligations vary by sector, jurisdiction, and contract. What this document gives you is the structure, which is the part that transfers.

6

The Order These Belong In

The twelve are listed in the order they usually appear in a policy, which is not the order in which they should be decided. Purpose, data, and accountability get settled first, because every other component depends on them. Tools and disclosure come next.

Training, ownership, and review are settled last and skipped most often, which is why so many policies are structurally complete and operationally inert. If you are auditing rather than writing, start at the bottom of the list instead.

Use this before publishing a new policy or during the next refresh. The objective is not to score the organization. It is to find where the language is still too vague, difficult to apply, or disconnected from real work.

AREA	PRACTICALITY QUESTION	YES	REVIEW
TOOL CHOICE	Can an employee tell whether a tool is approved, and who to ask if it is not listed?	☐	☐
DATA BOUNDARY	Can an employee recognize what information should never be entered into an AI system?	☐	☐
USE BOUNDARY	Can an employee distinguish an acceptable use from one requiring approval or one that is prohibited?	☐	☐
HUMAN REVIEW	Can an employee tell when AI output needs stronger human review because the stakes are higher?	☐	☐
VERIFICATION	Does the employee know what must be checked before AI-assisted work is relied on or shared?	☐	☐
DISCLOSURE	Can the employee tell when AI involvement should be disclosed?	☐	☐
ACCOUNTABILITY	Is it clear who owns the final work and decision?	☐	☐
ESCALATION	Does the employee know where to take uncertainty, mistakes, or suspected incidents?	☐	☐
MANAGER CONSISTENCY	Can managers explain the policy in roughly the same way across teams?	☐	☐
POLICY UPKEEP	Is there a named owner and a defined review trigger or cadence?	☐	☐

PROBLEMS TO REMOVE BEFORE LAUNCH

- › Rules that rely on undefined terms such as “sensitive,” “appropriate,” or “high risk.”
- › Sections that assign conflicting responsibilities or approval authority.
- › “Human review” language that does not explain what reviewers are expected to evaluate.
- › Requirements that cannot be followed with the tools, access, or time employees actually have.
- › Prohibitions with no safe alternative, which encourage hidden or shadow AI use.
- › No process for questions, edge cases, or policy updates.

Why There Is No Sample Language Here

Policy language that fits your organization comes from decisions your organization has made, and a clause borrowed from a template describes somebody else's decisions. The wording is also the part your counsel should shape, since obligations vary by sector, jurisdiction, and contract.

What this document gives you is the structure, which is the part that transfers. A policy with no owner, no review date, and no training commitment will not stay accurate regardless of how good the first nine components are.

DECIDED FIRST

Purpose, data, and accountability. Everything else depends on them.

DECIDED NEXT

Tools, human review, verification, and disclosure.

DECIDED LAST

Training, ownership, and review. Skipped most often.

AUDITED FIRST

The bottom of the list, where the upkeep lives.

THE FINAL PRINCIPLE

A strong AI use policy is not the longest policy. It is the one employees can understand, managers can reinforce, and the organization can keep current as tools and workflows change. A policy cannot stay useful by standing still, which is why it names an owner and creates a routine for learning, revision, and communication.

The structure is what transfers. The decisions are yours.

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

© 2026 The Jonesi Consultancy. All rights reserved.

Auditing a policy, or about to commission one?

We help organizations make the decisions first, then write the document that records them.

