

The Jonesi Consultancy | Houston, TX

EXECUTIVE AI ADOPTION · GOVERNANCE GUIDE

7 GOVERNANCE MISTAKES THAT CREATE SHADOW AI



How Well-Intentioned Rules Push Usage Somewhere You
Cannot See It

Shadow AI is not a discipline problem. It is what people do when the sanctioned path is slower, unclear, or unsafe to ask about.

SHADOW AI IS A SIGNAL

Almost every organization has shadow AI: employees using unapproved tools for work, usually on personal accounts, usually without telling anyone. The instinctive response is a stricter policy and a reminder of it. That response reliably makes the problem worse, because it addresses the behavior rather than the reason for it. Shadow AI tells you that somewhere between what people need to do their job and what the organization has sanctioned, there is a gap. The seven mistakes in this guide are the most common ways that gap gets created, and every one of them is committed by people acting responsibly.

1 What Shadow AI Is, and What It Costs

THE DEFINITION

The use of AI tools, features, or workflows outside an organization's approved, visible, or governed practices. It includes unapproved tools, personal accounts, undisclosed AI-assisted work, and workarounds created because employees do not know how to use the approved options effectively.

THE COST, STATED PLAINLY

An employee using an unapproved tool on a personal account puts your information somewhere you have no agreement, no retention terms, no audit trail, and no way to respond if something goes wrong. You carry the entire exposure and none of the visibility. Prohibition does not remove that exposure. It removes your knowledge of it.

2 The Shadow AI Cycle

The cycle is self-reinforcing. Each stage produces the conditions for the next, and the final stage restarts the first.

1 AMBIGUITY

Employees are unsure what is actually allowed.

2 FRICTION

Approved paths feel slow, unclear, or impractical.

3 WORKAROUND

Employees find another way to get the task done.

4 HIDDEN USE

AI activity moves outside normal visibility.

5 LESS LEARNING

Leaders lose the feedback needed to improve governance.

6 TIGHTER RULES

The response is more restriction, restarting the cycle.

3

Mistakes 1–3: Restriction and Friction

The first governance failures usually happen before an employee ever makes a risky AI decision. They happen when the rules are difficult to interpret, or the approved path is harder than the workaround.

1 Writing rules so restrictive that compliance is impractical

A policy that prohibits more than people can work without.

WHY IT CREATES SHADOW AI

Employees still have deadlines. When the sanctioned path cannot meet them, the work does not stop, it relocates. The most restrictive policies produce the highest rates of unsanctioned use, and the organizations that have them are usually the last to know.

WHAT TO DO INSTEAD

Ask what work the prohibition is preventing. If the honest answer is that people need to do it anyway, the rule is generating shadow AI rather than reducing risk. Restrict the uses that create real risk and preserve a clearly permitted space for low-risk, role-relevant work.

2 Never naming a single approved tool

A policy that describes acceptable use without saying what to use.

WHY IT CREATES SHADOW AI

This is the single most common gap. In the absence of an approved list, careful people freeze and confident people choose their own. Both are rational responses to the same missing decision, and only one of them is visible to you.

WHAT TO DO INSTEAD

Publish a list, however short. Three named tools with an organizational account and a route for requesting a fourth removes most of the problem.

3 Making tool approval slow enough that nobody uses it

A request process measured in quarters when the need is measured in days.

WHY IT CREATES SHADOW AI

A formal review that takes eight weeks is functionally a prohibition, and everyone learns that within one cycle. The second person does not bother submitting, which means you also lose the signal about what people actually need.

WHAT TO DO INSTEAD

Set a service level for the answer, even a conservative one, and make clear who reviews what. A definite no in five days generates less shadow AI than a maybe in eight weeks.

LEADERSHIP CHECKPOINT

If responsible AI use requires employees to navigate more uncertainty than the work itself, the governance system is creating friction rather than reducing it.

Governance becomes harder to follow when the people doing the work cannot see themselves in it. Rules written entirely from the top miss the task pressures and judgment calls employees face every day.

4 Leaving no safe path for experimentation

Governance that allows approved uses and provides nowhere to try something new.

WHY IT CREATES SHADOW AI

People discover what these tools are good for by trying things that were not on anyone's list. Without a sanctioned space for that, curiosity does not disappear. It moves onto a personal account with a personal login and your organization's information.

WHAT TO DO INSTEAD

Define a low-stakes lane: approved tool, no confidential information, nothing leaves the team without review. Most experimentation fits inside that and becomes visible.

5 Writing the rules without asking the people who do the work

Governance designed by legal, IT, or an executive committee alone.

WHY IT CREATES SHADOW AI

Rules written at a distance turn out to be unworkable in ways that are obvious to frontline staff and invisible upward. Workarounds appear quietly, and the governance body is the last to hear about them because nobody wants to report that they are not following the policy.

WHAT TO DO INSTEAD

Include at least one manager who does the governed work, and one frontline employee. They will find the unworkable clause before it is published rather than six months after.

WHY WORKFORCE INPUT MATTERS

- › Employees know where AI is already entering real work, including places leadership may not see.
- › Frontline feedback exposes rules that sound clear on paper but become ambiguous in practice.
- › Managers know which decisions carry meaningful consequences and where review already occurs.
- › Visible participation makes governance feel like an operating system rather than a compliance document imposed from elsewhere.

A USEFUL DISTINCTION

Employee input does not mean employees decide the risk standard. It means governance decisions are informed by the work they are intended to govern.

5 Mistakes 6–7: Keeping Shadow AI Invisible

The final two mistakes prevent the organization from learning. When employees cannot surface uncertainty safely, or governance never reaches the workflow level, leaders receive an incomplete picture of how AI is actually being used.

6 Inconsistent leadership behavior

Rules that apply to staff and visibly not to leadership.

WHY IT CREATES SHADOW AI

If an executive uses an unapproved tool in a meeting, or requires AI use while never demonstrating it, staff draw a conclusion about which rules are real. That conclusion generalizes to the rest of the policy immediately.

WHAT TO DO INSTEAD

Leadership follows the same list, publicly, including the inconvenient parts. Where an exception is justified, explain the principle behind it so employees understand the rule rather than assuming there are two systems.

7 No safe way to say “I already did this”

A reporting route with no protection, or no route for questions that are not yet problems.

WHY IT CREATES SHADOW AI

The employee who pasted something they should not have has two options: report it and find out what happens, or hope. Without an explicit protection clause and evidence that it holds, hoping is the rational choice. The mistake then stays hidden for exactly as long as it takes to become an incident.

WHAT TO DO INSTEAD

State in writing that good-faith reporting will not result in disciplinary action, provide a route for questions separate from the route for incidents, and then honor both the first time somebody tests them.

6 What Stronger Governance Looks Like

CLEAR

Employees can explain the boundaries in plain language.

USABLE

Approved paths are practical enough to use during real work.

SAFE TO QUESTION

Employees know where uncertainty can be surfaced without punishment.

WORKFLOW-CONNECTED

Expectations are tied to tasks, decisions, data, and review points.

CONSISTENT

Leaders and employees operate from the same visible principles.

ADAPTIVE

Feedback and new use cases change the governance over time.

You cannot survey your way to this answer. Asking employees directly whether they use unapproved tools is asking them to report a policy violation, and the results will tell you what people are willing to admit rather than what is happening. These indicators are more reliable.

✓	INDICATOR
☐	Output quality varies in ways that do not track with individual skill or seniority.
☐	Work is being produced faster than the sanctioned tooling would explain.
☐	People describe a process using a tool nobody has heard of, then correct themselves.
☐	Nobody has requested a new tool in six months, in a market where new tools appear weekly.
☐	Your approved list has not changed since it was written.
☐	An unapproved tool appears in a screenshot, a shared document, or a calendar invite.
☐	Nobody has reported a mistake in a year.

THE AMNESTY CONVERSATION

The fastest way to find out is to ask without consequence, once, and mean it. Say plainly that you want to know what people are using, that nothing said in this conversation will be held against anyone, and that the purpose is to make the sanctioned path good enough that nobody needs a workaround.

This works exactly once and only if it holds. If a single person is disciplined on the basis of what they disclosed, you will never get an accurate answer from that organization again, and the shadow use will simply become better hidden.

WHAT THE ANSWERS ARE WORTH

Every unapproved tool somebody names is a requirement your governance did not meet. Treated that way, the amnesty conversation is the most efficient tool-selection research available to you, and it is free.

The Shadow AI Governance Test

Use this as a leadership discussion tool, not a compliance score. A “no” or “not sure” identifies an area where ambiguity or friction may be pushing AI activity outside normal visibility.

AREA	QUESTION	YES	REVIEW
APPROVED TOOLS	Can employees name which AI tools are approved for routine work?	☐	☐
DATA BOUNDARY	Can employees explain what information must never be entered into an AI tool?	☐	☐
EXPERIMENTATION	Is there a practical path for low-risk experimentation?	☐	☐
ANSWER SPEED	Can employees get a timely answer when a use case falls outside the written policy?	☐	☐
MANAGER RESPONSE	Do managers know how to respond when someone discloses a mistake or uncertain AI use?	☐	☐
LEADERSHIP EXAMPLE	Do leaders visibly follow the same governance expectations they ask employees to follow?	☐	☐
HUMAN REVIEW	Are human-review expectations tied to actual tasks and consequential decisions?	☐	☐
REPORTING ROUTE	Do employees have a safe way to report suspected data exposure, inaccurate output, or other concerns?	☐	☐
FEEDBACK LOOP	Does the organization collect feedback about where governance is confusing or impractical?	☐	☐
UPKEEP	Is governance reviewed as tools, workflows, and organizational experience change?	☐	☐

THREE LEADERSHIP QUESTIONS TO DISCUSS

- › Where are employees most likely to work around our current rules, and what does that tell us about the rules?
- › Which governance expectations are clear in principle but difficult to apply during real work?
- › What would make responsible AI use easier to do openly than to do privately?

9

The Reframe

Governance that produces shadow AI has failed at the thing governance exists to do, which is to make the safe path the easy path. A rule that people follow only when it is convenient is not protecting the organization, it is documenting an intention.

The organizations with the least shadow AI are rarely the ones with the strictest policies. They are the ones where the approved tools are good enough, the approval route is fast enough, the experimentation lane is real, and reporting a mistake has visibly cost nobody anything.

THE TEST WORTH APPLYING TO ANY RULE YOU ARE ABOUT TO WRITE

If an employee under deadline pressure follows this rule, can they still do their job? If the answer is no, you have not written a control. You have written the reason your next incident will be discovered late.

10

Where to Start If This Describes You

FIRST

Publish an approved list this week, even a short one. It closes the largest gap in this guide.

SECOND

Set a response time for new tool requests and hold to it.

THIRD

Define a low-stakes lane where trying something new is explicitly permitted.

FOURTH

Write the good-faith reporting protection into the policy, then honor it the first time it is tested. That first test is the whole thing.

CLOSING PRINCIPLE

The goal of governance is not maximum restriction. It is enough structure to reduce avoidable risk, enough clarity to support good judgment, and enough psychological safety that employees do not need to hide how they are working.

AI governance works best when responsible use can happen in the open.

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

© 2026 The Jonesi Consultancy. All rights reserved.

Finding out what is actually running in your organization.

We help leadership teams get an honest answer and then close the gap that created it.

