

A six-step method for mapping shadow AI use across your organization

The Shadow AI Register

Most AI governance work begins with a policy draft. But a policy written without knowing what people are actually doing can’t possibly address the practical risks that will appear.

The solution? Look first. his guide walks through six steps for building a register of actual AI use across your organization: how to open the exercise so people tell you the truth, how to ask in a way that surfaces tools people don't think of as AI, how to find what self-reporting misses, how to structure what you collect, how to keep it current, and how to turn it into a decision your leadership can act on.

FEATURED ON ETHICSVERSE DAY 2026

“The Great AI Debate: Ethical Nightmare or Competitive Advantage?” aired 7/23/26



Reid Blackman
Founder & CEO
Virtue Consultants




Andrew McBride
Founder & CEO
Integrity Bridge LLC



Kay Firth-Butterfield
CEO Good Tech
Advisory



Wendell Wallach
Emeritus Chair Technology and Ethics Studies,
Formerly Yale University and the Carnegie
Council for Ethics in International Affairs



Before You Start..

How to Use This Resource



This is a working method rather than a reference document. Read it through once, then run it. In a mid-sized organization the full exercise takes four to six weeks from launch memo to read-out.

WHO RUNS IT

Ethics and compliance owns the method and does the drafting. A business sponsor (CEO or COO) sends the launch memo and opens the read-out. That split is deliberate and Step 1 explains why.

WHO YOU NEED LINED UP FIRST

Step 3 depends on data only IT, procurement, finance, and legal can pull. Recruit them before the survey launches, not when you reach that step — the discovery checks have to run inside the same window as the survey to be useful.

SEQUENCE

Steps 1, 2, and 3 form the collection phase, with Steps 2 and 3 running concurrently. Steps 4, 5, and 6 form the response phase and run in order. The one sequencing error worth avoiding: do not close the survey before starting discovery, or you lose the ability to compare what people reported against what was actually running.

WHAT YOU END UP WITH

A populated register of actual AI use, a standing intake path that keeps it current, and a leadership decision on authority and resources.

DECIDE THIS BEFORE YOU LAUNCH

Who reviews the findings, and by when. The exercise generates documented knowledge of unapproved use, and that knowledge carries an obligation to act on it.

Section 01

Write the launch memo

Who signs it matters more than what it says

The launch memo is a one-page message that opens the collection window. Compliance drafts it; the CEO, COO, or another business sponsor sends it. A census announced by compliance reads as an audit. The same census announced by the business reads as planning. Planning gets more honest responses.

Five elements to cover:

- 01

The ask, stated plainly. One survey, five minutes, everyone. Name the deadline.
- 02

Why now. Frame it as enabling adoption, not restricting it. You cannot approve tools, buy licenses, or negotiate enterprise terms for use you cannot see.
- 03

The amnesty terms. This is the load-bearing element. State that disclosure during the window carries no disciplinary consequence, and that after the window closes, undisclosed use of unapproved tools is a policy violation. Both halves are necessary. Without the first, people conceal; without the second, you have given up your ability to act later.
- 04

What happens to the responses. Who reads them, whether individual answers are reported upward, and what the output will be used for.
- 05

A named contact. One person, one email address, for questions people would rather not put in a survey field.

Time-box the amnesty explicitly — a date, not a vague "during this period." The clean line between protected disclosure and later violation is what makes the second half enforceable.

Decide before you launch: who reviews the findings, and on what timeline. Running the census creates documented awareness of unapproved use. An organization that collects the data and does nothing with it is in a worse position than one that never looked.

Section 02

Run The Survey

Ask about tasks, not tools

Never open with "do you use AI tools." A large share of the people who would answer no are using a summarizer in their meeting software, an assistant inside a spreadsheet, or a browser extension they have never classified as AI. Ask what they do, then ask how they do it. Eleven questions. Five minutes. Distribute enterprise-wide.

A

Task screener

1.

In the last 90 days, have you used any software tool to help with any of the following? (Multi-select)

☐

Draft or edit writing

☐

Summarize a document, meeting, or thread

☐

Translate

☐

Write or review code

☐

Create images or slides

☐

Answer questions about company information

- ☐ Research a topic
- ☐ None of these

2. How often? (Single-select)

- ☐ Daily
- ☐ A few times a week
- ☐ A few times a month
- ☐ Rarely

B Tool identification

1. Which tools did you use? (Multi-select — populate with your known and approved tools)

☐ [Approved tool 1]

☐ [Approved tool 2]

☐ [Approved tool 3]

☐ Something built into software we already use

☐ Other — please name

2. For each tool you selected, where did it come from? (Single-select, repeated per tool)

☐ Provided by the company

☐ I signed up for it myself

☐ Built into software we already use

☐ Not sure

3. Are you paying for any of these personally? (Single-select)

☐ Yes

☐ No

D Data touched

1. What kinds of information have you entered into these tools? (Multi-select)

- ☐ Publicly available information
- ☐ Internal but not sensitive
- ☐ Customer or personal data

- ☐ Employee or HR data
- ☐ Financial data not yet public
- ☐ Legal or privileged material
- ☐ Source code or technical specifications
- ☐ None of the above

2. Have you uploaded a file, document, or dataset? (Single-select)

- ☐ Yes
- ☐ No
- ☐ Not sure

D Output use

1. Does anything the tool produces reach someone outside the company? (Single-select)

☐ Yes, directly	☐ Yes, after edits
☐ No	☐ Not sure

2. Does it inform a decision about a person — hiring, performance, pay, discipline, credit, or eligibility? (Single-select)

☐ Yes	☐ No
☐ Not sure	

3. Before the output is used, does a person check it? (Single-select)

☐ Always	☐ Usually
☐ Sometimes	☐ No

E Open field

1. What would you like to use these tools for that you can't today? (Free text)

Section 03

Run Discovery Parallel To The Survey

Self-reporting has a floor that no amount of good framing will fix

Some use-cases will never appear in a survey because the person forgot, because they don't consider it AI, or because they are using a personal tool on a personal device for work. That last category is the one that produces the incidents. Run these six checks concurrently with the survey window, not after it.

Source	Owner	What to Pull	What it Reveals
SSO / identity provider logs	IT	Applications authenticated against corporate credentials, last 90 days	Tools signed up for with a work email
Expense reports	Finance	Reimbursement lines for software subscriptions under a set threshold	Personally purchased tools claimed back
Procurement and AP	Procurement	Vendor payments and renewals not routed through standard software review	Departmentally purchased tools
Endpoint and browser inventory	IT	Installed applications and browser extensions across managed devices	Assistants and plug-ins embedded in daily workflow
Network egress	IT/Security	Traffic to known AI service domains, including from unmanaged sessions	Use that bypassed both procurement and SSO
Vendor contracts	Legal/Procurement	Existing agreements where the vendor processes your data	AI in your supply chain that you did not deploy

Section 04

Build The Register Schema

Twelve fields. Each one earns its place.

The register is a spreadsheet. One row for every distinct thing someone is doing with AI, and one column for every fact you want to know about it. "Schema" is simply the word for the set of column headers.

This step is not the spreadsheet, it is the decision about what the columns should be. That decision is cheap to make now and expensive to change later, because revising the columns after two hundred rows are filled in means going back and re-interviewing people.

Column Headers	Type	Definition and Example
Use case ID	Auto	Unique identifier
Use case	Short text	What the tool is being used to do (drafting first-pass vendor questionnaires)
Business unit	Select	Function or department
Owner	Name	The person accountable for this use, not the person who reported it
Tool and model	Short text	Product and version where known
Provisioning	Select	Company-provided / self-signed-up / embedded in existing software / vendor-side
Data classes	Multi-select	From the survey question 6 list
Output leaves the company	Yes / no / unsure	Directly, after edits, or not at all
Affects a person	Yes / no / unsure	Hiring, pay, discipline, credit, eligibility, access
Human review	Select	Always / usually / sometimes / never
Consequence of error	Select	Negligible / contained / material / severe
Discovery source	Select	Survey / SSO / expense / procurement / endpoint / network / contract
Sanctioned status	Select	Approved / provisional / under review / prohibited
Next review	Date	Set at logging, not left blank

Worked Example

One row, filled in. Someone in procurement is using a personal AI account to summarize vendor questionnaire responses.

Column	Value
Use case	Summarizing vendor questionnaire responses
Business unit	Procurement

Column	Value
Owner	Procurement manager
Tool and model	[Tool name], personal account
Provisioning	Self-signed-up
Data classes	Internal; third-party business information
Output leaves the company	No
Affects a person	No
Human review	Usually
Consequence of error	Contained
Discovery source	Survey
Sanctioned status	Under review
Next review	[Date]

Once rows look like this, the table starts answering questions. Filter for every row where provisioning is self-signed-up. Pull every row where customer data meets a personal account. Hand each business unit only its own rows.

Section 05

Convert The Census Into A Live Register

A one-time snapshot goes stale in a quarter

The census tells you where you were on a given date. Without an intake path, you will be governing from a picture taken two quarters ago while the technology and the use cases keep moving. Publish a standing form, six fields, and require it before anyone adopts a new tool or extends an existing one to a new purpose.

- 1. What do you want to do? (One sentence — summarize incoming vendor questionnaires)
- 2. What AI tool do you want to use? (Product name; note if it is already approved for another purpose)
- 3. What information goes into it? (Data classes; note whether files are uploaded)
- 4. Who or what does the output affect? (Internal only, external parties, or decisions about individuals)
- 5. Who checks the output before it is used? (Name a role, or state that no review step exists)
- 6. Requester and business unit (Name, function, and the accountable owner if different)

Submissions route to owner or governance body. Response within [X] business days. Approved use cases are added to the register with a review date; declined requests receive a reason and, where possible, a sanctioned alternative.

Section 06

Present The Findings

The register is evidence. Use it to get a decision, not just to inform.

The read-out is short. Its purpose is not to describe the exercise but to convert what you found into authority and resources you did not have before. A general warning about AI risk is easy to acknowledge and set aside; a named list of what is already running inside the company is not.

Cover these points, in this order (each bullet should have its own slide):

- **Coverage.** Response rate, functions represented, and what the discovery checks added beyond self-reporting. Establishes that the finding is real rather than anecdotal.
- **Scale against expectation.** The total count of distinct use cases, stated next to whatever number leadership would have guessed. The gap is the argument.
- **Sanctioned versus unsanctioned.** The split, and how much of the unsanctioned share sits in functions with sensitive data.
- **Where risk concentrates.** Not a full risk ranking — the two or three functions or data classes where the exposure clusters.
- **Three named exposures.** Specific use cases, described concretely enough that a non-specialist understands what could go wrong. Specificity is what moves this from a briefing to a decision.
- **What is already resolved.** Anything shut down, brought under an approved tool, or given a review step during the exercise. Demonstrates the function acts on what it finds.
- **The ask.** Two parts: the authority to set and enforce the standard, and the resources to train the people expected to follow it. Tie each to a finding.
- **The standing commitment.** The intake path, who owns it, and when the register will next be reviewed.

OUTRO

Where This Leaves You

At the end of this exercise you have something most organizations do not: a factual account of how AI is actually being used inside the company, rather than an assumption about it. That changes what ethics and compliance is able to say. "AI carries risk" is a position anyone can hold and anyone can set aside. "Here are the forty-one things running in this company, here are the nine nobody approved, and here are the three that worry me" is a position that requires a response.

Be clear about what the register does and does not accomplish. It does not reduce risk. It makes risk visible and assignable, which is the precondition for every other piece of governance work you might do. The policy you write afterward will regulate what people are actually doing. The training you build will address the situations employees actually face. The triage you run will sort a known list rather than an imagined one.

Plan to re-run the full exercise annually, and sooner if something material changes — a major acquisition, a new enterprise tool deployment, or a shift in what your vendors are doing with your data. In between, the intake form carries the load.

The most common regret from teams who have done this is not that they found too much. It is that they ran the census before deciding who would act on it, and lost momentum in the weeks between the read-out and the first real decision. Name that owner in Step 1.