



A REFERENCE CHECKLIST

AI Governance

Checklist

Twenty-four questions covering where AI is in use, who approves it, what data it receives, and how problems get noticed.

J. Lynne & Co.

The patience to see it. The drive to move it.

HOW TO USE THIS

The ways organizations run into difficulty with technology are well established: a decision made by someone who should not have been making it, a process nobody owns, information that turned out to be wrong, a supplier whose interests diverged. AI does not add to that list.

What it changes is how problems become visible. When a person makes a poor decision there is usually a trail — a message, a conversation, a name attached. When a system makes the same decision at scale, the output looks much the same on a good day and a bad one. So most of governing AI is about restoring visibility: knowing where it is, knowing who approved it, and being able to tell when it stops working well.

Mark each statement Yes, No, or Not sure, answering for the organization rather than for yourself. Section 01 is where most organizations find the largest gap, and the other three depend on it.

About the three columns

A No marks something known and not yet done. A Not sure marks something believed but not tested.

They lead to different next steps. A No can be assigned to someone. A Not sure usually needs someone to go and check first.

01 INVENTORY

Where it is

STATEMENT	YES	NO	NOT SURE
We have a current list of where AI is used here, including tools staff started using on their own.			
We know which of our existing vendors have added AI features to products we already bought.			
We know which of those uses affect decisions about people — hiring, lending, eligibility, discipline, or service.			
We know which of them receive confidential information.			
One named person is responsible for keeping that list current.			
We know which uses we would want to explain carefully if asked about them publicly.			

02 APPROVAL

Who decides

STATEMENT	YES	NO	NOT SURE
We have written down what AI may and may not be used for here.			
Someone specific approves a new AI use before it goes live.			
Every AI-assisted decision that affects a person has a named human who can overturn it.			
Staff know when they are required to disclose that AI was used.			
We know what we would do, and who we would tell, if a tool produced a harmful or clearly wrong result.			
Our board or governing body knows our current position on AI.			

03 DATA

What we have agreed to

STATEMENT	YES	NO	NOT SURE
We know what information each AI tool receives and where that information goes.			
We know whether our data is being used to train someone else's model.			
We have checked our existing contracts for AI terms added after we signed them.			
We know which rules and regulations apply to the decisions we are automating.			
Staff have clear, specific guidance on what must never be entered into an AI tool.			
We can retrieve or delete our data if we stop using a tool.			

04 DETECTION

How problems get noticed

STATEMENT	YES	NO	NOT SURE
We know how we would notice if a tool gradually started producing worse results.			
We keep enough of a record to review a specific AI-assisted decision months later.			
Someone reviews a sample of outputs on a schedule, not only when there is a complaint.			
We have defined what working correctly means for each use, in terms we can measure.			
Staff and customers have a way to report a bad outcome, and we act on it.			
We revisit our approved uses at least once a year and remove those we no longer need.			

READING YOUR ANSWERS

Count the No and Not sure answers together, out of 24.

OPEN ITEMS	WHAT THAT USUALLY MEANS
0–3	Few open items Unusual at present. Keeping the inventory current is the main ongoing task.
4–9	A common range Where most organizations sit. These are usually questions of policy and habit rather than technology, and they close quickly once someone owns them.
10–16	A wide spread

OPEN ITEMS	WHAT THAT USUALLY MEANS
	Adoption is running ahead of the arrangements around it. Section 01 is the usual starting point, since the rest depends on knowing what is in use.
17+	Most items open The organization cannot yet describe its own use of AI. This is common, and building the inventory is normally the first useful step.

One question underneath all twenty-four

If one of these tools began producing worse results next month, how long would it take to notice, and who would raise it?

Where the honest answer is that it would surface through a complaint, detection is the area with the most to gain.

ABOUT J. LYNNE & CO.



J. Lynne & Co. is an independent governance advisory firm working with organizations to lead major transformations with confidence. We take no vendor commissions and hold no reseller agreements.

We publish reference material of this kind for executives, boards, and program sponsors. If it is useful, you are welcome to adapt it for your own organization.

If you would like to talk

We are happy to answer questions about anything on this list, or to help build a first inventory.

Website: jlynne.co

Book a conversation: jlynne.co/contact

This document is general reference material. It is not legal, accounting, or regulatory advice, and it is not a substitute for professional counsel on your organization's circumstances.

© 2026 J. Lynne & Co. All rights reserved.