

THE 1-PAGE AI GOVERNANCE CHECKLIST

Global frameworks set a floor years away. These five controls protect your programmes this quarter. No treaty, no committee required.

01

AUDIT RIGHTS BEFORE SIGNATURE

Audit and evaluation rights go into every vendor contract before signing, so a model can be inspected, not just trusted.

02

KILL CRITERIA, SET IN ADVANCE

Define the specific failure that pulls a model from production. Decide it calmly, well before any incident forces the call.

03

ONE NAMED HUMAN PER DECISION

A single accountable owner for each automated decision. Accountability cannot sit with a system.

04

LOG EVERYTHING, AND THE WHY

Everything the model did, and why, is recorded, so any outcome can be explained later, to a regulator or a customer.

05

REVIEW THE BORING FAILURES FIRST

The smallest, quietest failures get reviewed before the loud ones. That is where silent drift hides.

None of this is glamorous, and none of it waits for regulation. It is the difference between saying you use AI responsibly and being able to show exactly how, on a Tuesday, when something breaks.