AI Governance & Risk Checklist

Set proportionate rules for every AI system: who owns it, how risky it is, and the checks it must pass before and after go-live.

InflexionIQ

Good governance speeds AI up. When the rules are clear and proportionate, teams know what they must do and can get on with it.

> *Match the checks to the risk. A drafting assistant is not a credit decision.*

HOW TO USE IT

01 **Name an owner** and set the risk tier for each AI system.

02 **Work through the checklist** (page 2) at each stage, recording the evidence.

03 **Log the specific risks** and controls in the register (page 3), alongside your RAID log (template 21).

04 **Review it** at least yearly, and whenever the model, data or supplier changes.

THREE RISK TIERS

Low risk

Internal productivity: drafting, search, summaries. No decisions about people.

Needs: Usage policy, staff training, spot checks.

Medium risk

Supports decisions or touches customers, with a person making the final call.

Needs: Owner, testing, monitoring, sign-off by risk.

High risk

Makes or strongly shapes decisions about customers or colleagues, such as credit, pricing, claims or hiring.

All of the above, plus fairness testing, explainability, human review and executive sign-off.

ROLES

- Executive sponsor: accountable for the outcomes.
- System owner: runs it day to day and owns the risks.
- Risk and compliance: sets the rules and challenges.
- Data protection officer: personal data and individuals' rights.

UK RULES TO CHECK WITH YOUR COMPLIANCE TEAM

- UK GDPR, including the rules on solely automated decisions.
- FCA Consumer Duty, for firms serving retail customers.
- Equality Act 2010, for unfair treatment of protected groups.
- Sector rules, such as PRA model risk expectations for banks, and the EU AI Act if you serve EU customers.

WORKED EXAMPLE: COMPLAINT TRIAGE AT A UK LENDER

Risk tier: medium	Owner	Biggest risk	Control
It shapes how complaints are handled, but a person makes every decision.	Head of customer care, with the chief risk officer as sponsor.	Missing signs that a customer is vulnerable.	A person reviews every case the model is unsure about, and a weekly sample of the rest.

This checklist is a practical starting point, not legal advice.

CHECK	DONE	EVIDENCE	OWNER
BEFORE BUILD			
The purpose and the users are written down and agreed.	☐		
A named owner is accountable for the system.	☐		
The risk tier is set (see page 1).	☐		
A data protection impact assessment is done if personal data is used.	☐		
We know how we'll test accuracy and fairness.	☐		
BEFORE GO-LIVE			
Testing shows accuracy is good enough, for every customer group.	☐		
People affected can get an explanation and challenge the outcome.	☐		
Human oversight is designed into the process, with clear authority.	☐		
Customers are told when AI is used, where that matters.	☐		
Risk, compliance and data protection have signed off.	☐		
IN LIVE USE			
Performance and drift are monitored, with alert thresholds.	☐		
Complaints and incidents involving the system are tracked.	☐		
Supplier and model changes trigger a review.	☐		
The system is reviewed at least once a year.	☐		
There is a plan to switch it off safely if needed.	☐		

AI risk register

AI SYSTEM

DATE

InflexionIQ

Risks and controls

Categories: bias · privacy · accuracy · security · explainability · conduct · supplier

#	RISK	CATEGORY	LIKELIHOOD	IMPACT	CONTROLS	OWNER
01			1–5	1–5		
02			1–5	1–5		
03			1–5	1–5		
04			1–5	1–5		
05			1–5	1–5		
06			1–5	1–5		
07			1–5	1–5		
08			1–5	1–5		

Sign-off

Name, role and date for each approver.

Next review

Date, or the change that triggers it.

Switch-off plan

How we'd stop safely, and who decides.