



EXECUTIVE GUIDE

Responsible AI Operating Controls

A concise control model for accountability, risk, human oversight and continuous monitoring.

BEST FOR

Executives, governance teams, risk and compliance leaders, product owners, security teams and architects.

01

Define accountability

Assign clear business, technical and risk ownership for every material AI solution.

02

Match controls to risk

Use stronger assurance where impact, autonomy, sensitivity or regulatory exposure is higher.

03

Keep humans in control

Define review, intervention and escalation points for consequential or uncertain outcomes.

04

Monitor after launch

Treat quality, safety, incidents, drift and material change as ongoing operating responsibilities.

FRAMEWORK

Core control areas across the AI lifecycle

Responsible AI becomes operational when principles are translated into specific owners, controls, evidence and review points.

Purpose & acceptable use

1

Document intended purpose, users, prohibited uses, decision boundaries and when the solution must defer to a person.

Ownership & risk classification

2

Assign accountable owners and classify impact using autonomy, sensitivity, scale, affected users and consequence.

Data, knowledge & access

3

Use approved sources, least-privilege access, data minimisation, retention rules and controlled knowledge updates.

Quality, robustness & transparency

4

Evaluate expected and edge cases, document limitations and provide appropriate user guidance and traceability.

Human oversight & approvals

5

Define approval gates, review thresholds, escalation paths, intervention rights and evidence required for release.

Monitoring, incidents & change

6

Track quality, policy adherence and incidents; reassess controls after material changes to models, data, prompts or integrations.

ACTION CHECKLIST

Responsible-AI control checklist



Purpose, users, boundaries and prohibited uses are documented.



Business, technical and risk owners are explicitly assigned.



Risk classification determines the level of assurance required.



Data, knowledge and access controls follow approved permissions.



Human review is retained where consequences or uncertainty require it.



Testing covers quality, robustness, misuse and failure scenarios.



Logging and evidence support traceability and incident review.



Material changes trigger reassessment, testing and approval.

NEXT STEPS

Make controls part of delivery, not an afterthought

- Apply proportionate controls early in discovery and design.
- Require evidence at pilot, production and material-change gates.
- Review operational feedback and incidents to improve controls continuously.

Important

This summary supports internal governance discussions. Applicable laws, sector requirements and customer policies should be assessed separately for each deployment.