

● Sample guide for Example Organisation Ltd

Ethical AI Framework

What an Ethical AI Strategy engagement delivers: principles, decision rules, bias testing, incident response and a rollout plan, shown for a fictional organisation.



“The question isn’t just ‘Can we build it?’ It’s ‘Should we build it and if so, how do we do it right?’”

Prepared for
Example Organisation Ltd
a fictional organisation

Prepared by
AI Assist
part of Cloud Assist Group

Date
October 2026
sample edition

Engagement
Ethical AI Strategy
8 to 12 weeks

About this sample. This guide shows an extract from each part of a custom Ethical AI Framework. Example Organisation Ltd is fictional: a UK employer of about 2,400 people, with a small team in Dublin. Its examples, roles and thresholds are invented to show the format. A real framework is written with your leaders, your staff and their representatives.

1 Principles

Six principles, written for the way Example Organisation works

This framework sets out how Example Organisation decides whether to use AI, how it builds or buys it, and how it checks AI once it's in use. It covers systems built in-house, AI bought from suppliers, and staff use of general AI assistants.

#	Principle	What it means here	The question we ask
1	Purpose-driven	Every AI use starts with a named problem, an owner and a measure of success. Ideas that are AI for AI's sake are turned down.	Can we say in one plain sentence what problem this solves, and for whom? Would a simpler fix work?
2	Explainable	Anyone affected by an AI-assisted decision can get a plain-English explanation of how it was reached.	Could the manager using it explain the output to the person it affects?
3	Fair and unbiased	We test for unfair outcomes before launch, and keep testing while the system is in use.	Have we compared outcomes across groups, and acted on any gaps?
4	Secure and private	Personal data is used only where it's needed, with a data protection impact assessment (DPIA) wherever the law or this framework requires one.	Do we know what data goes in, where it goes, and who can see it?
5	Accountable	A named person owns every AI system and every significant decision it supports.	If this goes wrong, who answers for it, and do they know?
6	Augmentation	AI supports people's judgement. A person makes the final call on decisions about people.	Does a person decide anything that affects someone's job, pay or access to a service?

The six principles build on published frameworks for ethical AI from leading AI and HR researchers, reimagined by AI Assist for each client.

What's in the full framework with working templates for every part

Deliverable	In this sample
Custom Ethical AI Framework document	Principles, this page
Decision-making protocols and approval workflows	Page 3
Implementation roadmap with governance structure	Governance, page 3; roadmap, page 5
Bias testing and mitigation playbook	Page 4
Incident response procedures	Page 4
Training programme outline	Page 5
Stakeholder engagement toolkit	Page 5

Find your way

Proposing an AI use?

Start with the six stations and risk tiers on page 3.

Building or buying one?

Use the bias testing playbook on page 4.

Something gone wrong?

Follow the incident procedure on page 4, the same day.

Leading a team?

See the training and engagement tools on page 5.

2 Decisions and governance

How an AI idea becomes an approved system


Every AI use passes the same six stations. The depth of review depends on the risk tier, so low-risk uses move quickly and high-risk uses get the scrutiny they need.



Risk tiers and approval routes

Tier	Typical uses here	Review	Approver	Evidence before launch
Tier 1: low	Drafting, summarising and searching with the licensed assistant	Self-assessment checklist	Line manager	Completed checklist. Reviewed yearly.
Tier 2: medium	Customer email drafting with staff review; knowledge search over service guides	AI working group	AI product owner	Purpose statement, data map, test results, user guidance. Reviewed six-monthly.
Tier 3: high	Anything that informs decisions about people's jobs, pay or access to services, or that profiles people	Full ethics review, including bias testing, a DPIA and consultation with employee representatives	AI Ethics Council	All of tier 2, plus a bias test report, DPIA, human oversight plan and EU AI Act classification. Reviewed quarterly.
Not permitted	Inferring staff emotions at work, social scoring of staff or customers, and covert monitoring are banned by this framework. Emotion recognition at work (outside narrow medical and safety reasons) and social scoring are also prohibited practices under Article 5 of the EU AI Act.			

Governance structure



3 Testing and incidents

Bias testing playbook and incident response

Bias can enter at any stage. The playbook says where to look, how to test, and what to do when a test fails. Extract shown.


Stage	Where bias can enter	How we test	What we do about it
Data	Past decisions carry past bias; some groups are under-represented; fields such as postcode act as proxies	Compare the data with the population it describes. Screen fields for proxies of protected characteristics.	Rebalance or remove fields. Record known limits in the system file.
Design	What the system is told to optimise, and which features it uses	A diverse review panel challenges the target and the features.	Change the target. Drop features that act as proxies.
Before launch	Unequal outcomes or error rates between groups	Compare outcomes and error rates across groups, using protected characteristics under the Equality Act 2010 where the data is lawfully held.	Gaps above the Council's threshold are fixed or explained before launch.
In use	Drift as data changes; people relying on the output too much	Monthly outcome checks, the rate at which people override the system, and a feedback channel.	Retrain, add a review step, or retire the system.

Worked example: shift allocation tool Illustrative

Test	Measure	Review trigger set by the Council	Result
Part-time and full-time staff	Share of preferred shifts granted	Gap of more than 5 percentage points	Pass
Age bands	Share of preferred shifts granted	Gap of more than 5 percentage points	Review
Manager overrides	Share of allocations changed by managers	More than 1 in 5	Pass

What happened next: the age-band gap was traced to a rule that favoured staff with longer service. The rule was changed and the tool retested before launch.

Incident response procedure



Personal data breaches are assessed against UK GDPR: where a breach must be reported, the ICO is told within 72 hours of becoming aware. For EU AI Act high-risk systems, serious incidents also go to the provider and the relevant authority once those duties apply on 2 December 2027.

4 Rollout

From framework to everyday practice

The framework is built in 12 weeks and embedded over the following nine months. Governance starts early, so the Council is ready to approve real cases by week 11.

Implementation roadmap



Training programme outline

Audience	What they learn	Format
All staff	What AI is and isn't; acceptable use; how to spot and report a problem	45-minute e-learning and a team briefing
Managers	Using AI output in decisions about people; explaining decisions; when to escalate	Half-day workshop
System and product owners	Risk triage, bias testing, documentation and monitoring	Two-day practitioner course
AI Ethics Council and board	Oversight duties, the regulatory picture, reading test results	Two-hour briefing, refreshed yearly

Stakeholder engagement toolkit

Stakeholder map

Who each AI use affects, and how to involve them before it goes live.

Plain-English staff explainer

What the framework means for people's day-to-day work.

Consultation guide

How to involve employee representatives and works councils, and when.

Customer FAQ

How Example Organisation uses AI, and how to ask for a person instead.

Feedback channel

A simple route for anyone to raise a concern about an AI system.

"You said, we did" update

A quarterly template that shows people what changed because they spoke up.

5 About AI Assist

AI that works for people

AI Assist helps organisations use AI where it genuinely helps, and tells them honestly when it won't. Simple, ethical AI that puts people first. No hype. No jargon. Just clarity.

"Sometimes the answer is better process. Sometimes it's smarter use of existing tools. And sometimes, yes, it's AI. Our job is to tell you which, honestly."

No proprietary lock-in

You own everything we create. Our recommendations are vendor-agnostic.

Ethics as standard

Every engagement includes ethical assessment, bias testing, governance and transparency.

Plain English, always

If your finance director can't follow the recommendation, we haven't done our job.

How an Ethical AI Strategy engagement runs



Who it's for

- Organisations deploying high-risk AI systems
- Regulated industries, such as financial services, healthcare and the public sector
- Companies concerned about reputational risk
- Leadership teams committed to responsible innovation

Part of Cloud Assist Group



Book a discovery call

No pitch. No pressure. Just 30 minutes to understand your challenge and explore whether we can help.

AI Assist is part of Cloud Assist Group. ai-assist.cloud