

CUSTOM AI SOLUTIONS

Built for the problem, not the hype

Four illustrative solution briefs that show how we decide what to build, what to fix, and when to say “don’t build anything new”.

Written for Example Organisation Ltd, a fictional UK employer. The briefs, names, scores and figures are illustrative; none describes a client project.



CONTENTS

01	How we decide	2
02	Briefs: when we build	3
03	Briefs: when we fix	4
04	Inside a brief	5
05	About AI Assist	6

***“Sometimes you need something built.
Sometimes you need something fixed.
Sometimes you just need someone to tell
you the truth.”***

AI Assist, Custom AI Solutions

Sample edition, October 2026. Prepared by AI Assist, part of Cloud Assist Group.

01 HOW WE DECIDE

Three paths, one honest recommendation

Off-the-shelf tools don't always fit, and AI isn't always the answer. We start with the business problem, not the perceived solution, test three paths, and build only when AI is genuinely the right tool.

Path 1
Process improvement
 Can better process design solve this?

 Right when the problem is how work flows between people, not what the tools can do.


Path 2
Configuration optimisation
 Can your existing tools do this with better setup?

 Right when you already own the capability and it is switched off, half set up or hard to find.

Path 3
Custom AI solution
 Is this genuinely a problem AI should solve?

 Right when the work needs judgement at a scale people cannot reach, the data supports it, and people can check the output.

Our process



EVERY BUILD IS

- **Explainable:** the people who use it, and the people it affects, can see why it said what it said.
- **Testable:** tested against real scenarios, including bias and edge cases, before go-live.
- **Auditable:** inputs, outputs and human decisions are logged.
- **Human in the loop:** a person makes every decision that matters.

WHAT WE DON'T BUILD

- ✗ "AI for AI's sake" vanity projects
- ✗ Systems that can't be explained
- ✗ Solutions where simpler alternatives exist
- ✗ Tools that replace human judgement on critical decisions

If a simpler alternative exists, we recommend it, even when that means a smaller piece of work.

Examples of what we build

HR and people

Skills taxonomy and matching engines; internal mobility recommendations; interview bias detection; employee sentiment analysis; onboarding assistants

Compliance and risk

Automated policy change tracking; contract risk assessment; regulatory requirement mapping; audit trail automation

Operations and efficiency

Document classification and routing; meeting insights and action tracking; customer query triage; workflow optimisation

The briefs in this pack

Brief	Request	Area	Path chosen	Indicative timeline
01	Internal mobility matching	HR and people	Custom AI solution	16 weeks
02	Policy change tracking	Compliance and risk	Custom AI solution	12 weeks
03	Customer query triage	Operations	Configuration	6 weeks
04	Onboarding assistant	HR and people	Process improvement	4 weeks

Timelines are illustrative. Engagements typically run 8 to 24 weeks, depending on the problem.

02 BRIEFS: WHEN WE BUILD

When AI is genuinely the answer

BRIEF 01 | HR AND PEOPLE

Custom AI solution

Indicative timeline: 16 weeks

Internal mobility matching

The problem, as defined in week 1. People leave to find roles they could have done here. Internal vacancies are posted, but employees don't see the ones that fit them, and hiring managers can't see who has the skills.

PROCESS IMPROVEMENT

A monthly vacancy digest and manager nominations. Helps visibility; does not match skills. Started anyway, as a quick win.

CONFIGURATION

The core HR platform's marketplace is licensed, but skills data is too sparse for its matching to work yet.

CUSTOM AI CHOSEN

A skills taxonomy and matching engine that suggests open roles and next steps to each employee.

WHAT WE BUILT

A shared skills taxonomy drawn from job descriptions and skills employees confirm themselves, and a matching engine that suggests roles with plain-English reasons: "matches 6 of 8 listed skills; gap: stakeholder management". Designed to hand over to the platform's own marketplace when skills data matures, so nothing locks you in.

HUMAN IN THE LOOP

Suggestions go only to the employee, who chooses whether to apply. Managers never see a ranked list of people. Hiring decisions are made exactly as before.

HOW WE TESTED

Suggestion rates compared across groups (sex, age band, ethnicity, working pattern) where that data is held lawfully; edge cases including career returners, part-time staff and sparse profiles; user acceptance with 40 employees at three sites.

SUCCESS MEASURES

Internal applications per vacancy; share of roles filled internally; employees who say they can see a next step (targets set with the client, illustrative here).

HOW IT WORKS RED STATION: A PERSON DECIDES



BRIEF 02 | COMPLIANCE AND RISK

Custom AI solution

Indicative timeline: 12 weeks

Policy change tracking

The problem, as defined in week 1. The compliance team learns about regulatory and policy changes from newsletters. Changes reach policy owners late, and there is no record of who reviewed what, or when.

PROCESS IMPROVEMENT

A weekly review rota and a shared log. Tried before; it fell behind within months.

CONFIGURATION

The existing risk platform's alerts cover only the sources it subscribes to, a minority of those needed.

CUSTOM AI CHOSEN

A monitor that summarises changes from named sources and maps each to the internal policies it affects.

WHAT WE BUILT

A change monitor for an agreed list of official sources. Each change gets a short summary that quotes the passage it relies on and links to the original, plus a suggested list of affected internal policies from a requirement-to-policy map.

HUMAN IN THE LOOP

The policy owner confirms relevance and decides any change. Nothing is published automatically. Every step is logged for audit.

HOW WE TESTED

Replayed against six months of past changes already reviewed by the compliance team; every summary checked against its source for omissions; edge cases such as consultations versus final rules, and guidance versus law.

SUCCESS MEASURES

Days from a source change to owner review; share of changes with a complete audit record; owner-rated summary accuracy.

HOW IT WORKS RED STATION: A PERSON DECIDES



03 BRIEFS: WHEN WE FIX

When the honest answer is simpler

BRIEF 03 | OPERATIONS

Configuration optimisation

Indicative timeline: 6 weeks

Customer query triage

The problem, as defined in week 1. Customer emails land in one shared inbox. The service team spends the first hour of each day sorting, and urgent complaints wait behind routine requests.

PROCESS IMPROVEMENT

Agreed categories and a triage rota. Part of the answer.

CONFIGURATION CHOSEN

The service platform you already own has classification and routing rules that were never switched on.

CUSTOM AI

A bespoke classifier. Not needed: it would duplicate a capability you already own.

WHAT WE ADVISED

Configure the existing platform: twelve categories agreed with the service team, routing rules by category, and priority rules for complaints and for customers who may be vulnerable. No new build.

HUMAN IN THE LOOP

Advisers can re-categorise in one click. Anything flagged as a possible vulnerability goes straight to a person.

HOW WE TESTED

Replayed four weeks of past emails through the new rules and reviewed every misroute with the service team before go-live.

SUCCESS MEASURES

Time to first response; share of emails routed correctly first time; adviser time spent sorting.

HOW IT WORKS RED STATION: A PERSON DECIDES

BRIEF 04 | HR AND PEOPLE

Process improvement

Indicative timeline: 4 weeks

Onboarding assistant

The problem, as defined in week 1. The request was an onboarding chatbot. The real issue: new starters ask the same questions because the answers differ by site and by manager.

PROCESS IMPROVEMENT CHOSEN

One onboarding checklist and one set of answers, owned by HR.

CONFIGURATION

The core HR platform's onboarding module handles tasks and reminders once the content is agreed.

CUSTOM AI

A chatbot built on inconsistent content would give confident, wrong answers.

WHAT WE ADVISED

Don't build a chatbot now. Agree one checklist and one set of answers, then configure the onboarding module. Revisit in six months: if questions still flood in, an assistant grounded only in the agreed answers becomes a sound option.

HUMAN IN THE LOOP

Managers own each new starter's first-week plan. HR owns the answers and reviews them every quarter.

HOW WE TESTED

Piloted the new checklist with two intakes of new starters and logged their questions to find gaps.

SUCCESS MEASURES

Questions to HR per new starter; new starters ready on day one; 30-day new-starter feedback.

HOW IT WORKS RED STATION: A PERSON DECIDES

Why include briefs where we built nothing? Because they are often among the most useful work we do. A recommendation not to build is a finished deliverable: it comes with the reasoning, the simpler fix, and the point at which it would be worth looking again.

04 INSIDE A BRIEF

What you receive: excerpts from Brief 01

Each brief comes with the working behind the recommendation. These excerpts show three of the documents we hand over, and the requests we have advised against.

Options analysis Illustrative scores, 1 to 5; 5 is best

Criterion	Process improvement	Configuration	Custom AI
GATE: DOES IT SOLVE THE PROBLEM AS DEFINED?			
Matches people to roles on skills	Partly: visibility only	Not yet: data too sparse	Yes
THEN: HOW WELL, AND AT WHAT RISK?			
Time to first value	■ ■ ■ ■ ■ 5	■ ■ ■ ■ ■ 3	■ ■ ■ ■ ■ 3
Explainable to employees	■ ■ ■ ■ ■ 5	■ ■ ■ ■ ■ 3	■ ■ ■ ■ ■ 4
Low risk to people	■ ■ ■ ■ ■ 5	■ ■ ■ ■ ■ 4	■ ■ ■ ■ ■ 4
You own it, no lock-in	■ ■ ■ ■ ■ 5	■ ■ ■ ■ ■ 2	■ ■ ■ ■ ■ 5
Light ongoing upkeep	■ ■ ■ ■ ■ 4	■ ■ ■ ■ ■ 4	■ ■ ■ ■ ■ 2
Recommendation	Do now, alongside the build	Revisit at month 12	Build with the safeguards below

Solution fact sheet

PURPOSE	Help employees find internal roles that fit their skills.
USERS	Employees who opt in; the talent team maintains it.
DATA USED	Job descriptions, open vacancies and skills employees confirm.
NOT USED	Performance ratings, absence, pay. Protected characteristics are used only, separately, to test for bias.
NEVER DOES	Rank people for managers, or make or recommend selection decisions.
EXPLAINS	Skills matched and gaps, with every suggestion.
MONITORING	Suggestion rates by group, monthly; accuracy and drift, quarterly.
OWNER	Head of Talent. Next review in six months.
CONCERNS	Raised with the talent team; a person replies.

REQUESTS WE ADVISED AGAINST

- 1 Flight-risk scores from email and chat activity**
Monitors people without their knowledge and can't be explained to them. Instead: stay conversations and a regular pulse question.
- 2 Automated performance ratings**
Replaces human judgement on a critical decision. Instead: help managers organise evidence from their own notes; the manager decides the rating.
- 3 A chatbot to replace intranet search**
Better page titles and a properly set-up search did the job in two weeks.

Test plan summary

Bias testing

Suggestion rates compared by sex, age band, ethnicity, disability and working pattern, where data is held lawfully. Review thresholds agreed with HR and the data protection lead before testing starts.

Edge cases

Career returners and gaps; part-time and job-share staff; profiles with few skills; unusual job titles; people who opt out (no suggestions, no penalty).

User acceptance

40 employees at three sites, across roles and grades. Passes when people can say why each suggestion appeared.

Go-live gate

DPJA complete; bias results within thresholds; acceptance signed off by the talent team and employee representatives; owner named; rollback plan ready.

05 ABOUT AI ASSIST

AI that works for people

AI Assist is an ethical-AI consultancy and part of Cloud Assist Group. We offer simple, ethical AI solutions that put humans first. No hype. No jargon. Just clarity.

We never do AI for AI's sake. 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.

Six principles we hold to on every engagement

1 Purpose-driven

AI must solve real problems, not create new ones. If we can't state the benefit in plain English, we don't proceed.

3 Fair and unbiased

We test for bias, design fairness in from the start and monitor to catch issues before they cause harm.

5 Accountable

AI assists decisions; humans make them. There is always a person in the loop for decisions that matter.

2 Explainable

If we can't explain it, we don't recommend it, to the people who use it or the people it affects.

4 Secure and private

Data security and privacy are foundations, not afterthoughts.

6 Augmentation over automation

The best AI makes people more effective. It enhances human judgement rather than replacing it.

Our services

☐ AI Readiness Assessment

☐ Compliance and Risk Management

☐ Ethical AI Strategy

☒ **Custom AI Solutions**

☐ HR Transformation Advisory

☐ Schools AI Readiness

Our promise

No proprietary lock-in. You own everything we create. Our recommendations are vendor-agnostic.

Ethics as standard. Every engagement includes an ethical assessment, bias testing and governance.

Plain English, always. Summaries that executives actually read; visual frameworks over dense reports.

Part of Cloud Assist Group



What happens next



Book a discovery call

Tell us about your challenge in a 30-minute conversation. We will tell you honestly whether AI is the answer, or whether there is a simpler way.

ai-assist.cloud