



HR DELIVERY

Making Sure HR Delivers

HR DELIVERY WORKSHOP SERIES

AI in HR

An introduction to responsible AI use in HR

Explore where AI assistance may be useful in HR and where its limits matter. Practise defining a bounded use case, checking synthetic demonstration outputs and planning meaningful human review before considering any workplace pilot.



01 DAY

Introductory taster

IN-HOUSE

Your people. Your context.

PRACTICAL

Explore, discuss and apply

A PRACTICAL STARTING POINT

Understand the topic. Try it in practice.

An impressive AI answer can hide an incorrect source, an unsupported inference or a task that should remain with a person. This workshop starts with a modest HR communication task and examines the work needed to make assistance useful, including information boundaries, checking, escalation and the authority to reject outputs.

Participants explore the Human-Governed AI Lifecycle by framing a limited use case and testing illustrative outputs against approved source material. The purpose is to build informed judgement about value, controls and review, leading to questions for appropriate organisational specialists before any pilot.

WHO IS IT FOR?

HR leaders, practitioners and operations colleagues considering AI assistance in everyday work, together with managers who may sponsor a pilot. No coding or previous AI implementation experience is required.

BY THE END OF THE DAY, PARTICIPANTS CAN:

- 01 Frame a useful task**
Describe one bounded HR task, the intended benefit, affected users and a simpler alternative worth considering before using AI.
- 02 Set information boundaries**
Identify permitted synthetic inputs, excluded uses and questions about access, retention and supplier arrangements that need specialist review.
- 03 Check an AI output**
Compare a demonstration answer with its source, identifying unsupported details, omissions and misleading confidence before deciding whether it is usable.
- 04 Define meaningful review**
Specify what a competent reviewer must check, the evidence they need and their authority to reject, correct or escalate outputs.
- 05 Outline a limited test**
Draft task-specific acceptance criteria, difficult examples and a stopping rule for a proposed pilot using approved information and oversight.
- 06 Identify approval needs**
Explain which owners and specialists must consider the use case before authorisation and what evidence they would need to review.

A useful first application


This day introduces responsible use through synthetic demonstrations and a bounded exercise. It does not authorise deployment, certify compliance or support autonomous consequential people decisions; proposed workplace use needs organisational and jurisdiction-specific review.



THE FRAMEWORK / WORKSHOP FOCUS

HR AI Framework: Human-Governed AI Lifecycle

The Human-Governed AI Lifecycle connects six stages: Frame, Bound, Test, Authorise, Operate and Review. It keeps a defined purpose, accountable human judgement and evidence of performance central to HR assistance.



SIX QUESTIONS TO EXPLORE IN THE WORKSHOP

01 Frame

What specific task and benefit justify considering assistance from AI?

03 Test

What evidence would reveal incorrect answers, omissions or unsafe behaviour?

05 Operate

Can reviewers challenge outputs and provide a reliable human alternative?

02 Bound

Which inputs, outputs, users and actions are permitted or excluded?

04 Authorise

Who must review the evidence and approve any limited pilot?

06 Review

What changes or failures should trigger retesting, narrowing or stopping?

These six questions are workshop focus questions within the existing lifecycle. Exercises use synthetic information, require meaningful human review and exclude autonomous decisions about hiring, dismissal, promotion, pay, discipline or comparable employment consequences.

[Explore the full model >](#)



LEARN BY DOING

Explore. Challenge. Make the ideas useful.

Short explanations lead into lively exercises, paired discussion and small-group challenges. Participants test ideas, share experience and build a draft together. The tone is engaging and enjoyable, with time to reflect and connect learning to real work.

01 Sort the proposed uses

Teams examine HR task cards, compare AI assistance with simpler options and separate a modest drafting task from uses needing specialist scrutiny or exclusion from practice.

02 Inspect a flawed demonstration

Participants compare synthetic policy extracts with a prepared AI answer containing deliberate mistakes, mark unsupported statements and decide which corrections, rejection or escalation are justified.

03 Design the reviewer role

Pairs specify a reviewer's evidence, time, competence and authority, then rehearse rejecting an attractive answer when its source does not support the proposed HR communication.

04 Build a small test pack

Groups draft ordinary and difficult test questions, acceptance criteria and a stopping rule for the case, then challenge whether the proposed controls could work in practice.

DIAGNOSE BEFORE YOU ACT

Evidence > shared feedback > agreed priorities

Use selected diagnostic prompts to distinguish known facts from assumptions. Discuss the findings and agree their meaning before proposing a solution. Feedback is part of diagnosis, not an afterthought.

ONE CASE, DEVELOPED THROUGH THE DAY

An HR operations team wants AI to draft replies to routine questions about a fictional training booking process. A demonstration blends two policy versions and invents an approval step. The team has not named a source owner or reviewer, and its estimated time savings exclude the work of checking answers.

Bound the task to supplied synthetic material, identify demonstration failures and specify reviewer checks. Recommend what evidence and approvals would be needed before a limited organisational pilot could be considered.

THE FACILITATOR WILL

Invite different perspectives.

Make assumptions visible.

Welcome constructive debate.

Capture practical learning.



SAMPLE WORKSHOP AGENDA

A working day, with room to think.

09:00-16:30 | Six hours of facilitated learning, plus lunch and breaks. The agenda, emphasis and timings are always agreed with the client, including local scheduling needs.

09:00-09:20	Starting with the HR task Compare possible uses, affected people and simpler alternatives, establishing the workshop boundary against autonomous consequential people decisions from the outset.
09:20-10:15	Following the human governed lifecycle Explore Frame, Bound, Test, Authorise, Operate and Review, then define a modest task and an observable benefit for the case.
10:15-10:30	Break
10:30-11:30	Setting information and use boundaries Specify permitted synthetic inputs, excluded actions and supplier questions, and identify the specialists needed to resolve workplace information and access concerns.
11:30-12:30	Checking the demonstration answer Inspect a prepared AI response against fictional source extracts, identify unsupported content and practise correction, rejection or escalation with clear reasons.
12:30-13:30	Lunch
13:30-14:30	Making human review meaningful Design the reviewer role, including source access, relevant competence, time and authority to challenge outputs or pause the workflow when needed.
14:30-14:45	Break
14:45-15:30	Planning tests and stopping rules Draft task-specific criteria, ordinary and difficult examples, and review questions that assess usable quality alongside the time required for human checking.
15:30-16:10	Reviewing a proposed pilot Present the bounded use case and test plan to peers acting as an approval panel, recording unanswered questions and required organisational owners.
16:10-16:30	Choosing a responsible next step Identify a suitable organisational contact and decide whether to investigate, revise or defer an idea using the evidence developed during the day.

Follow one synthetic drafting task from a proposed benefit through boundaries, output checks and an accountable pilot decision.

FROM THE SESSION TO YOUR DESK

Useful models. Practical working records.

The day combines reference materials with drafts developed in the exercises. Each has a clear use beyond the workshop; the work reflects the evidence and perspectives available in the room.

01 Use case framing sheet

A worked outline of the task, users, intended benefit and simpler alternatives, including mistakes that would matter and evidence needed to judge practical usefulness.

Clarify what an AI proposal is actually intended to improve.

02 Information boundary draft

A case record of permitted inputs, outputs and actions, excluded uses and questions about permissions, supplier arrangements, access and retention requiring specialist consideration.

Prepare focused questions for privacy, security, procurement and HR colleagues.

03 Output checking example

An annotated synthetic answer showing unsupported statements, omitted conditions and source conflicts, with explanations of why a reviewer would correct, reject or escalate it.

Demonstrate why a plausible answer still requires evidence and judgement.

04 Human review role outline

A practical description of the reviewer's checks, competence, source access, available time and authority to reject outputs, pause use or seek appropriate specialist assistance.

Discuss whether meaningful review can be resourced in a proposed pilot.

05 Pilot test pack starter

An exercise pack containing ordinary questions, difficult cases, acceptance criteria and stopping triggers, with space to record errors and the effort required for review.

Develop a limited test proposal with the appropriate organisational specialists.

06 Approval and review questions

A concise set of prompts about ownership, residual risk, human alternatives, monitoring and material changes that may require renewed authorisation or additional testing before continuation.

Support a clear conversation about whether an idea should proceed.

CONTINUE EXPLORING

Continue with the models, diagnostic manuals and development resources on the HR Delivery website. Workshop exercises provide a first application; a full organisational diagnostic or implementation assignment can be agreed separately.

[Open the framework >](#)

[Explore diagnostic resources >](#)



MAKE THE LEARNING COUNT

A better conversation. A practical next step.

Participants should take one bounded idea to the relevant organisational owners before using workplace information. The workshop outputs support a conversation about purpose, risk and testing; any pilot needs authorised tools, permitted data, competent reviewers and explicit approval.

YOUR FIRST 30 DAYS

01

Days 1-7

Select and discuss

Share one modest use case with an HR owner and compare it with simpler alternatives. Confirm that autonomous consequential people decisions remain excluded and identify relevant privacy, security and procurement contacts.

02

Days 8-21

Clarify the controls

Discuss information permissions, supplier arrangements and reviewer capacity with appropriate specialists. Refine the synthetic test pack and record unanswered questions; do not introduce employee records while approval remains unresolved.

03

Days 22-30

Decide the next step

Present the revised purpose, boundaries, test criteria and ownership proposal to the appropriate decision owner. Agree whether to investigate further, seek pilot authorisation, narrow the task or stop the proposal.

Follow-through is participant-led. Additional facilitation, a full diagnostic or specialist support can be agreed separately.

AVAILABLE IN-HOUSE

Shaped around your organisation

Agree the in-house date and confirm organisational rules for AI demonstrations. The facilitator provides synthetic documents and prepared outputs; participants should bring task ideas without employee information. Any live demonstration uses an approved environment and non-personal material.

PLAN YOUR IN-HOUSE WORKSHOP

Talk to Dr Phillip Welsh

drphillip@hr.delivery

+974 7226 1481

Discuss your workshop



www.hr.delivery