HiBob

HORIZONS 2026 · IN GOOD COMPANY CAFÉ

Trust by design

Where AI, People, & Privacy Meet

THE PLAYBOOK · THE WHO, THE WHAT & HOW

Kesiena Ogefere & Natalie Neilson

Horizons 2026 · Protein Studios, Shoreditch · 18 Nov 2026

✦ THE WHO · SLIDES 04–06 · START WITH OURSELVES

Change is happening around us. And within us.

Change is happening everywhere and more rapidly than ever before. It can become overwhelming and often we can be so focused trying to control or 'keep up' with what's happening around us that we forget ourselves in the process.

We adapt our organisations, systems, and processes to meet a changing world. But how often do we stop to consider how we're evolving ourselves to meet it? What if we were to redirect some of that energy that we put out into the world, back into ourselves? What would change?

While we can't necessarily control the pace of change, we do have greater control and influence over how we meet it. When we focus within, practise metacognition, and build our internal capacity, we build and expand our 'toolbox' so that we have greater internal 'resources' to meet the twists and turns that work and life has in store.

HERE ARE A FEW OF MY FAVOURITE SELF-REFLECTIVE QUESTIONS TO TAKE AWAY WITH YOU

01
What is this experience trying to teach me about myself?

02
What might I see differently if I took a step back?

03
What am I currently energising? Is it within my control?

For a deeper reflection, read Natalie's article:

→ Why does scaling ourselves feel like the most challenging task of all?

How to use this playbook

Trust in AI is built by the choices you make before a tool goes live: which data it touches, who it affects, and who stays accountable.

Those choices start with us. When we notice our own response to change and build our capacity to lead through it, we make steadier decisions for everyone else. The pages that follow take the questions we worked through together and turn each one into something you can run with your team next week.

Each part follows the same pattern:

THE PROBLEM
in plain English

HOW TO SPOT IT
in your own organisation

WHAT TO DO
step by step, with a template where it helps

You don't need a data science team to use any of it. You need **one use case, one owner and an honest conversation.**

1 Find where AI creates opportunity, and where it creates friction

Most trust problems with AI in HR come from one of three tensions. Name the one you're in and the fix becomes much clearer.

Privacy *vs* insight

Which data would people expect you to use, and which would surprise them?

HOW TO SPOT IT

A tool pulls in data collected for a different purpose: chat logs, badge swipes, calendar data, sickness notes. Or nobody can list exactly which fields it reads.

WHAT TO DO

Map every data field the tool uses. Mark each one "expected" or "would surprise people". Remove the surprises, or explain them clearly before launch.

Augmentation *vs* surveillance

Is the tool helping people do the work, or watching how they do it?

HOW TO SPOT IT

Output goes to managers before it goes to the employee. Dashboards rank individuals. The language shifts from "support" to "monitor" or "flag".

WHAT TO DO

Show the output to the person it's about first. Report at team level unless there's a clear, stated reason not to.

Speed *vs* governance

Where is the pace of rollout outrunning the checks you'd want in place?

HOW TO SPOT IT

A tool is switched on because it came free in a system update. Staff are using AI tools nobody approved. Nobody owns the tool once it's live.

WHAT TO DO

Keep a simple AI register (tool, owner, data used, review date). Nothing goes live without a named owner.

A QUICK TEST YOU CAN RUN THIS WEEK

Pick one AI feature already live in your HR system. Ask three employees what they think it does with their data. If their answers don't match reality, **you have a trust gap to close before you add anything new.**

WORTH KNOWING FOR UK TEAMS

If a tool processes employee data in a way that's likely to be high risk, UK GDPR requires a Data Protection Impact Assessment before it goes live. The ICO's guidance on monitoring workers is a good starting point. If you have staff in the EU, AI used in recruitment and people decisions is treated as high risk under the EU AI Act. Bring your DPO or legal adviser in early, not at the end.

2 Before launch: three questions for any AI use case

If you can't answer all three in writing, the use case isn't ready to launch.

1
What problem does it solve?

Name the task or decision it improves, in one sentence. "Speed up HR" fails. "Cut the time to answer payslip queries from two days to two hours" passes.

2
Who is affected?

List everyone who sees the output and everyone who is judged by it: employees, managers, candidates. These are often different people, and the second group is the one whose trust you need most.

3
What stays human?

Name the decisions where a person stays accountable, and name that person. A role title is fine. "The system" or "HR" is not.

The one-page use case card

Fill this in before any AI tool goes live. Keep it with the tool's owner and review it at every check-in.

QUESTION	YOUR ANSWER
Use case name	
The problem, in one sentence	
Data it uses (list the fields)	
Who sees the output	
Who is judged by the output	
Decisions that stay with a person	
Named owner	
What employees will be told, and when	
Review date	

WHERE TEAMS USUALLY GET STUCK

Question 1. Most stalled AI projects started with the tool and went looking for a problem. If the problem is vague, fix the process first. AI on top of a messy process gives you faster mess.

3 Guardrails: five you can put in place this quarter

Governance works best as a set of small, repeatable decisions your team actually makes. A 40-page policy nobody reads doesn't build trust. These five guardrails do, and none of them needs new software.

GUARDRAIL	WHY IT BUILDS TRUST	HOW TO SET IT UP
1 Human sign-off on people decisions	Nobody is hired, rejected, rated or exited by a tool alone.	List the decisions the tool feeds into. For each, write the role that signs off. Add a sign-off step in your HR system workflow so it can't be skipped.
2 Plain-English notice to employees	People trust what they understand. Surprises break trust fastest.	One short paragraph per tool: what it does, what data it uses, what it doesn't do, who to ask. Publish it where people already look, such as the intranet or HR portal.
3 Consent before meetings are recorded	AI note-takers are the most common way AI enters a team without anyone deciding.	Set a default: recording and AI notes are off unless everyone agrees at the start. Add a line to meeting invites. Check the settings in Teams, Zoom or Google Meet.
4 A named owner for every AI tool	An unowned tool is the one that drifts, breaks or gets misused.	Add an owner column to your AI register. The owner checks the outputs, handles questions and calls the review.
5 A review date before scaling up	It forces a decision on evidence instead of momentum.	Set the date at launch, usually 6 to 12 weeks out. At review: keep, change or stop. Write down which, and why.

START WITH ONE

If you do nothing else, give every AI tool a named owner.

Most of the other guardrails follow once someone is accountable.

DISCUSSION PROMPT FOR YOUR TEAM

Which of these five would have prevented the last AI-related concern someone raised with you?

SLIDE 10 · START SMALL

4 Start small: design a test you can learn from

The smallest useful test is one team, one process, one tool, a short window and a clear end date. **A small test that ends on time builds more trust than a big rollout that drifts.**

Step by step

1. **Pick the use case.** Choose a process that is high volume, low risk and already well defined. Good first candidates: answering policy questions, drafting job adverts, summarising survey comments, cleaning duplicate records.
2. **Fix the process first.** Map the current steps on one page. Remove anything that exists only because the old system needed it. If the process is unclear, the AI output will be too.
3. **Check the data.** Confirm the data the tool needs is accurate, current and allowed to be used for this purpose. Bad data is the fastest way to lose trust in a good tool.
4. **Set the boundaries.** One team of 5 to 20 people. A 4 to 8 week window. A fixed end date in the diary.
5. **Agree two measures before you start.** One for value, one for trust (see below). Take a baseline now so you can compare later.
6. **Tell people what's happening.** Use the plain-English notice from section 3. Give them a way to raise concerns that doesn't go through their manager.
7. **Review and decide.** At the end date: keep, change or stop. Share the result with the people who took part, even if the answer is stop.

Two measures: value and trust

VALUE MEASURE (PICK ONE)	TRUST MEASURE (PICK ONE)
EXAMPLE Hours saved per week on the task	EXAMPLE % of participants who say they understand what the tool does with their data
EXAMPLE Time to answer a query	EXAMPLE % who would be happy for it to be used in their own team
EXAMPLE Error or rework rate	EXAMPLE Number of concerns raised, and how fast they were resolved
EXAMPLE Volume handled without HR stepping in	EXAMPLE % of outputs a person checked and changed

A pilot that saves time but loses trust has not worked. Measure both, from day one.

SLIDE 11 · TAKE IT BACK

5 Take it back: your first 30 days

On the day, we asked you to leave with **one question, one action and one conversation.** Here's a simple plan to turn those into progress.

WEEK 1

Find out what's already live

- [ ] List every AI tool or AI feature in use across HR, including ones switched on in system updates
- [ ] Start an AI register: tool, owner, data used, review date
- [ ] Ask three employees what they think one of those tools does with their data

WEEK 2

Close the obvious gaps

- [ ] Give every tool on the register a named owner
- [ ] Check meeting recording and AI note-taker settings
- [ ] Write the plain-English notice for your highest-risk tool

WEEK 3

Choose your test

- [ ] Pick one use case and fill in the use case card
- [ ] Map the process on one page and fix what's broken
- [ ] Agree one value measure and one trust measure, and take a baseline

WEEK 4

Launch small

- [ ] Brief the pilot team and share how to raise concerns
- [ ] Put the review date in the diary
- [ ] Have your one conversation: with IT, your DPO, a sceptical manager or your exec team

THE ONE QUESTION TO ASK BEFORE THE NEXT AI TOOL LAUNCHES

"Who is judged by this output, and do they know?"

THANK YOU & RESOURCES

Let's continue the conversation

Connect with us directly, and take the next step with your team.

We're in Good Company Chapter Co-Leads

Join us in the HiBob community to connect with other HR, L&D, and People professionals. Stay connected and up-to-date with future community events

→ Join the community



THE WHAT & HOW: HR TECH, DATA & AI

Kesiena Ogefere

IGC Reading Chapter Co-Lead

Founder, Copieux Group

Join the IGC Reading chapter →

Kesiena Ogefere is an In Good Company (IGC) Reading Chapter Co-Lead and the founder of Copieux Group, a UK consultancy that redesigns HR work so tech, data and AI deliver real results. She has spent more than 13 years in HR systems, working with teams at organisations including Vodafone, Mitsubishi Electric, Ocado and Martindale Pharma.

Copieux Group helps lean HR teams fix the process and data first, then set up the tech and AI to take the admin. The goal is simple: give HR teams time back to focus on work that moves the business.

Want to see how ready your HR work is for AI?
Take the free Work Design Readiness Assessment at copieuxgroup.com/assessment.

Want help applying this playbook to your own use case?
Book a free 30-minute consultation

Website: copieuxgroup.com
LinkedIn: linkedin.com/in/kesienaogefere



THE WHO · GROWTH & HUMAN CAPACITY

Natalie Neilson

IGC London Chapter Co-Lead

ICF Coach · Speaker · Founder, TTM Coaching

Join the IGC London chapter →

Natalie Neilson is an In Good Company (IGC) London Chapter Co-Lead, and an ICF-accredited coach, facilitator, speaker, founder of TTM Coaching, and co-host of The Power of You podcast. Her work sits at the intersection of individual growth and organisational performance, exploring how people and organisations build the capacity to navigate growth, change and complexity.

Natalie creates reflective spaces and experiences that help leaders and teams understand where they are, recognise what needs to shift, and build the capacity to move forward with greater clarity, confidence and sustainability.

Stay connected
Follow Natalie on Instagram / TikTok / LinkedIn / Substack / YouTube

Podcast: Tune into The Power of You Podcast, available on Apple and Spotify Podcasts.

Newsletter & sessions: Read growth insights via The Spark Newsletter or enquire about speaker sessions and facilitated growth experiences, helping leaders and teams navigate growth and transition: natalieneilson.com

At scale: Learn more about our growth readiness approaches for your teams and organisation through TTM Coaching: ttmcoaching.com