Human–AI Systems

Third Circle

STOP PROMPTING, START DIRECTING

Your Questions, Answered

A follow-up to the joint webinar from Human-AI Systems and Third Circle.

Thank you for joining us on Wednesday 24 June 2026. We promised to follow up with answers to the questions that came in during the session, and here they are. We have grouped them by theme, expanded the ones we only had time to touch on live, and added a few that come up often enough to be worth including. Mike McKeown took the questions on building and running the AI; John Clifton took the ones on strategy, adoption and leadership.

GETTING STARTED

Tools, projects and assistants

I already use AI most days. When should I set up a "project", and when do I need an "agent"?

Think of it as a ladder, and the words mean slightly different things as you climb it.

At the assistant stage, rung two, a project does two jobs at once. It is the container that organises everything to do with a piece of work, for example a project holding all my HSBC material, and it is also where the agent lives. At this stage an agent, sometimes called a skill, is a focused set of instructions and context that does a particular job, and it sits inside that project.

At the worker stage, rung three, those two things come apart. The project becomes purely the container, the place all the HSBC work happens, while the agents become separate workers in their own right that can be called into any project. In the worker demonstration, the HSBC work ran inside the HSBC project, but several different agents did the actual tasks: a knowledge agent did the research, a CRO agent prepared the meeting and later wrote it up, and a proposal agent produced the proposal document.

So a rule of thumb: set up a project the moment you find yourself re-explaining the same background, because it pays for itself straight away. Reach for separate agents once you have distinct jobs you want done well and want to reuse across different pieces of work.

Which tool should I actually start with: ChatGPT, Claude or Copilot?

Start with whatever you already have. The skill that matters, briefing well and judging the output, carries across all of them, so the effort is not wasted whichever you pick.

If you are choosing fresh, our honest view on the day was that Claude has the edge right now for this kind of work, which is the Cowork setup you saw in the worker demonstration. ChatGPT is excellent and the most familiar starting point for most people. If your organisation runs on Microsoft, Copilot is the sensible choice because it sits inside the tools and the security your business already uses. The method is deliberately tool-agnostic. We start from what you are trying to do, not from a favourite platform.

One thing worth knowing for the worker stage specifically. The full worker setup you saw runs in a desktop app, not the ordinary chat window. For Claude that is Claude Cowork. For ChatGPT it is the Codex desktop app, which works in a very similar way. And Microsoft has just launched Copilot Cowork, which under the bonnet is built on the same Claude Cowork technology. So whichever of the three you are on, there is now a route to the kind of setup we demonstrated.

How should I set up a project so it gives me its best, and can the AI help me build it?

Give it three things to begin with: who you are and your role, how you like to work and write, and the handful of documents it should always know about. Those documents are the context. In a project or assistant they live inside that workspace, so you load them once rather than attaching them every time. Keep them current and the quality of everything that follows improves.

And yes, use the AI to help build it. Ask it to interview you about your role, to draft its own instructions, or to tell you what context it is still missing. Some of the best setups are the ones where you let the AI help design how you work with it.

I am not especially technical. Can I really build something like the worker you showed, or do I need a developer?

You do not need a developer, and you do not need to write code. The real skill is leadership, not engineering: briefing clearly, judging what comes back, and giving good feedback so it improves. If you can manage a capable new joiner, you have what this needs.

The honest part is that it takes time. Building a full worker setup from scratch is realistically thirty to sixty hours of evenings and weekends. The barrier is time and a starting point, not talent. That gap is exactly what the programme closes: we build the core with you, so you begin near the top of the ladder instead of spending months climbing it.

RUNNING IT

What it costs

What are the cost implications of the more advanced setups? I was partway through some work, needed to buy more tokens, and the cost started to climb quickly.

A fair concern, and worth being straight about: cost can be a real issue if it is not managed. It is driven by three things: how powerful a model you use, how much work you put through it, and whether things run constantly or only when you ask. Hitting a wall mid-task and topping up usually means a heavy job ran on a premium model with no ceiling on it.

The biggest single lever is matching the model and the thinking level to the task. Use a light, fast model for routine work, and save the heavyweight, deep-reasoning models for the jobs that genuinely need them. Get that wrong and the cost climbs quickly. Get it right and it stays sensible. It is one of the things we cover properly on the programme, because it makes such a difference.

On the plans themselves, most individuals get a long way on a standard monthly subscription, in the region of twenty to thirty pounds, before usage-based costs come in. If you end up using it very heavily, you may want a higher tier. Mike runs his whole business on it and is on the eighty pounds a month Claude Max plan. That is a deliberate choice: the value he gets back is well worth it, which is exactly why he uses it so heavily. The principle is that spend should track the value you are getting, not creep up by accident.

Do the agents run full time, or only when you ask them? And what does that mean for daily cost?

By default they run when you give them a task. The real power, though, comes from scheduled tasks: work that runs on a timetable without you starting it. In Mike’s own setup a morning brief runs once a day, some tasks fire when an email arrives, and others he simply starts when he needs them.

One underrated benefit of scheduling is getting work done overnight. You can set a task running before you finish for the day and have the results waiting when you log on in the morning. It also keeps daily cost modest and predictable, because you are paying for work done rather than for something idling in the background. Deciding what runs on a schedule, what runs on a trigger, and what only runs on demand is, in effect, deciding your cost.

ACCURACY AND TRUST

Keeping it right and up to date

We have a Copilot agent answering questions over documents in a SharePoint folder. When we add new documents, does it re-index automatically? And when it gets something wrong, how do we correct it?

We checked this, because the exact behaviour matters. Copilot relies on SharePoint's semantic indexing. An update to a document that is already indexed is generally picked up very quickly, often more or less immediately. A brand new file behaves differently: new documents added to a SharePoint site are typically indexed on a longer cycle rather than instantly, and a long or complex file can take longer still before it is fully searchable. So if something you have just added does not show up straight away, that lag is usually the reason rather than a fault.

When it gets an answer wrong, the fix is almost always one of three things: update the agent's instructions so it behaves differently, update or add to the underlying files so the knowledge base is right, or adjust the prompt you are using to ask the question. Most real corrections are a combination of those three. It is less about retraining the model and more about managing what it reads and how you ask.

Either way, treat the knowledge base as something you curate over time. This iterative side, keeping the source clean and coaching the system as it goes, is the part self-builders most often neglect, and it is where a lot of the long-term value sits.

How do I stop it making things up, and how do I know I can trust what it gives me?

Start from the right expectation: it drafts, you direct and approve. The human stays responsible for what goes out of the door. Used that way, the occasional confident error gets caught the same way you would catch a mistake from anyone who works for you, by reviewing it.

A few habits help. Give it good source material to work from rather than relying on its general memory, ask it to show where an answer came from, and keep a person in the loop on anything that matters. It is at its most convincing exactly when it is wrong, so the discipline is to keep your judgement switched on. That judgement is the skill, and it is one you already have.

FIT

Connecting it to your business

What other systems do you need around it: a CRM and so on?

Less than you might expect. The point of this approach is that the AI connects to the tools you already use rather than adding another system to run. It can work with your email, calendar, documents and the common business tools, so it fits into your existing setup instead of replacing it.

As an example, Mike runs his own CRM inside the AI system rather than buying a separate package. That will not suit everyone, and where you already have a tool you like, the AI connects to it. The principle is the same either way: start from the work you need done, then connect only what that work actually requires.

What about tools like Warp that route each task to the "right" model automatically, rather than keeping you in one model?

Good question, and it helps to separate a couple of things out, because two different distinctions tend to get mixed together.

First, the models versus the platform. There are the AI models themselves, such as ChatGPT 5.5 or Claude Opus 4.8, and then there is the platform you access and use them through, such as Claude Cowork, Microsoft Copilot Cowork, Warp or Replit. The same model can sit inside very different platforms.

Second, two broadly different kinds of platform. There are general AI platforms for your own day-to-day work, such as Claude Cowork, ChatGPT Codex and Microsoft Copilot 365 Cowork. And there are agent-builder platforms for constructing specific workflows, often ones that are external, touching customers or partners, or genuinely complex. That is where tools like Warp, Make and Replit sit. There is overlap, but they largely solve different problems, and for most leaders the day-to-day platform is the place to start.

On the routing idea itself, there is real value, especially in more complex workflows, in choosing different models for different tasks or sub-tasks based on what each one is best at. You can also build in checking, where one agent reviews and improves another's work. That is sometimes called looping. The proposal step in the demonstration did exactly this: one agent wrote the proposal and another checked it and, where needed, updated it. So picking the right engine for each job is genuinely useful, but the bigger gains still come from good context and a well-designed workflow, not from routing on its own.

We are a Microsoft organisation. Do we have to move off Copilot to do this?

No. If your business runs on Microsoft, you can build this inside Copilot and your existing governance, and that is often the right call. Copilot sits within the security and compliance your organisation already trusts, which removes a real barrier to adoption.

There is also a recent development worth knowing about. Microsoft has just added Cowork to Copilot, and it is built on the same Claude Cowork technology we demonstrated. In practice that means the kind of worker setup you saw is now available to Microsoft organisations inside Copilot itself, which opens up far more than was possible even a few months ago. Staying on Microsoft no longer means missing out on this.

One word of caution we raised on the day: running something important on a platform nobody else in the organisation uses tends to have a shelf life. If the wider business lives in Microsoft, building there usually beats a separate tool that only you understand. The method works across Claude, ChatGPT and Copilot, so we fit the platform to your organisation rather than the other way round.

CONTROL

Risk, governance and responsibility

Isn't this risky with company or client data?

Yes, and it should be taken seriously rather than waved away. The risks are real, but they are manageable, and they are mostly about set-up and discipline rather than the technology itself.

In practice it comes down to a few things: use business-grade tools with the right data terms rather than consumer accounts, be deliberate about what information goes where, and set clear boundaries on what the system can see and do. Done properly, a well-governed AI setup can be safer than the unmanaged copying and pasting into free tools that is almost certainly already happening across your organisation.

The piece worth singling out is governance. A properly researched Responsible AI policy, covering which tools are approved, what data can go where, and who is accountable, turns all of this from ad hoc caution into something the organisation can actually rely on. Putting that in place is one of the things we do. The question is not whether to take data seriously, it is whether you do it by design or by accident.

Could it handle something like completing a tax return, and who is responsible if it gets something wrong?

It can do a great deal of the work: gathering the information, organising it, drafting the figures and preparing the return. What it should not do is file it unchecked. For anything with a legal or financial consequence, a competent human stays accountable and signs it off, ideally the relevant professional.

That is the right model for this kind of work generally. The AI takes the heavy lifting off your plate and gets you most of the way in a fraction of the time. You, or your accountant, own the final check and the responsibility. It does not transfer the accountability, it removes most of the effort that comes before it.

PEOPLE AND LEADERSHIP

Adoption, careers and the human side

Where do businesses waste the most money on AI?

Buying tools before understanding the problem. The common pattern is to roll out licences across the organisation, tick the box, and assume adoption will follow. It rarely does. The licences go in, the change does not, and the spend shows up with very little to show for it.

The businesses that get value start the other way round. They work out where time and value are actually leaking first, then apply AI to that specific bottleneck. A tool pointed at a vague ambition burns money. A tool pointed at a real, understood problem pays for itself.

I run a small business and I am stretched thin. Where do I even start?

Not with AI. Start with the bottleneck that is costing you the most: the task that eats your week, the work you keep putting off, the thing that does not need your judgement but still lands on your desk. Name that first.

Then point AI at that one thing and get it genuinely working before you do anything else. The mistake we see most often is treating AI as a toy to experiment with, rather than aiming it at a real job that needs doing. One solved problem you can feel beats ten half-built experiments, and it buys you the time and the confidence to tackle the next one.

How do I get my team or my leadership to actually adopt this, rather than just nod at it?

This is the part almost everyone underestimates, and it is where most AI efforts quietly stall. You can hand someone the best assistant in the world, and if nothing about how they work changes, nothing changes. Adoption is a change-management challenge, not a technology one.

What works is starting small and visible: pick a real workflow, show a concrete result on work people recognise, and let one team feel the time it gives back. Bring people in by solving a problem they actually have, rather than announcing a tool and hoping. And leaders carry this by directing it themselves, not delegating it entirely. People copy what their managers do, not what they are told to do.

What are the most common pitfalls of a poor AI implementation?

Two stand out. The first is bad data in, bad data out. If the information the system works from is messy, out of date or wrong, the output will be confidently wrong, and at speed. Getting the source material right is unglamorous and it matters more than almost anything else.

The second is treating mass adoption as a technology problem when it is really a change-management problem. Organisations buy the licences and expect transformation to follow. It does not, because the hard part is people changing how they work. The implementations that fail tend to get the tool right and the people wrong. The ones that succeed put as much effort into adoption and governance as into the technology.

Is AI going to replace jobs on my team?

The more useful way to look at it is that roles change rather than simply disappear. The routine, repetitive work gets absorbed, and the human time moves up to judgement, relationships and the things that genuinely need a person. Most roles end up reshaped, not removed.

What is clearly changing is what makes someone valuable. After sixteen years placing senior leaders, John’s read is that the people pulling ahead are the ones who can brief AI well, judge its output and direct it. Being an AI-capable leader is becoming part of the job description rather than a nice-to-have. The risk is less about AI replacing people, and more about people who use it well outpacing those who do not.

What are you seeing in the market that most people are missing?

The gap between buying AI and actually using it. Almost every organisation now has the licences. Very few have changed how work happens. When we polled the room, two-thirds said their organisation had no working AI strategy, and that matches what we see everywhere. The tools are in. The transformation is not.

That gap is the whole opportunity. The organisations that close it are not the ones with the most tools, they are the ones that treated adoption as a leadership and design problem and saw it through. Most people are still waiting for the technology to be the answer. It is only ever the enabler.

NEXT STEPS

Putting it into practice

What was the one thing you said we could do this afternoon?

Stop starting from a blank box. Open a project in ChatGPT or Claude and give it three things: who you are and your role, how you like to work and write, and the handful of documents it should always know about. That single step moves you off the Tool rung and onto the Assistant rung. It is the cheapest, highest-return move in all of this, and the setup guide we are sharing walks you through it.

If you are on Microsoft, Copilot can do the same thing, though it is a slightly more involved combination of agents and Teams rather than a single project. The guide covers that route too.

Can I get the recording, the demonstrations and the setup guide?

Yes. Everyone who registered receives the recording, the three demonstrations from the session, and the one-page guide to setting up your first assistant. They are useful on their own, with no obligation attached.

If you would like to go further, you can register your interest in the programme we run together, where we build a personalised AI system with you and then teach you to run and extend it yourself. The aim is for you to finish owning it, not depending on us.

Still have a question?

If something here raised more questions than it answered, that is usually a good sign. Get in touch with either of us directly, or register your interest in the programme using the link we shared, and we will pick it up with you.

Mike McKeown, Human-AI Systems · mike.mckeown@human-ai-systems.com

John Clifton, Third Circle · john@thirdcirclepartners.com