

From **Caversham House** this month: **3 strategic insights** on the trends reshaping how organisations think about AI, **3 tools** we've been testing at the frontline and **3 training courses** to help your teams build real capability.

1. The Month in View: 3 Strategic Trends

Curated insights to help you read the room and lead the shift.

Trend 1: AI can replace tasks before it can replace roles

AI is getting better at performing discrete tasks, but a role is more than the sum of its task list.

Jobs also include judgement, coordination, exception handling, context, accountability and the tacit knowledge that lets people know when the standard process no longer applies. Strong performance on a bounded task is therefore weak evidence for eliminating the role around it.

There is a second risk too: if AI removes too much of the work through which junior staff learn, organisations may weaken the pipeline that produces future expertise.

Call to action: Separate task automation from role redesign, and require evidence at the level of the whole workflow before changing headcount.

1. Map which parts of a role are genuinely automatable and which depend on judgement or context.
2. Test AI across real workflows, including exceptions and hand-offs, before making workforce decisions.
3. Redesign junior work so AI removes low-value effort without removing the experiences that build expertise.

Relevant courses:

[GenAI Foundation Course](#) | [Data, Privacy and AI: Practical Decisions Every Day](#)

Read more:

[AI Replacing People at Work: What the Evidence Shows](#) | [Why Humans Stay at the Centre of AI Skill Transfer](#)
| [Sleepwalking into a Soft Skills Crisis Caused by AI](#)

Trend 2: Middle managers and culture decide whether Gen AI sticks



Once people have access to AI, adoption becomes less a technology problem and more a management problem.

Frontline managers shape whether people experiment, share what works, raise concerns and change established routines. But managers differ in confidence, risk tolerance, incentives and the kind of evidence they trust, so a single enablement approach will produce uneven results.

Call to action: Treat frontline managers as the operating layer of AI adoption, not simply another audience for training.

1. Segment managers by confidence, risk appetite and evidence needs.
2. Give them practical ways to lead AI-enabled work in their teams.
3. Measure changes in behaviour and workflow, not course completion or licence activation.

Relevant courses:

[Leading Teams Through New AI Workflows](#) | [AI Steering Committee and Governance](#)

Read more:

[Middle Managers Will Make or Break AI Adoption](#) | [Preparing Leaders to Build an AI-Ready Culture Now](#)

Trend 3: When everyone uses AI, everyone can start thinking alike

AI increasingly influences how problems are framed, which options are surfaced and which evidence appears important.

The risk is convergence. If teams use similar models in similar ways, they can end up exploring the same assumptions, options and lines of reasoning. The result can look like faster thinking while actually reducing the diversity of ideas being considered.

A human in the loop does not solve this if people only review the final answer. By then, the range of possibilities may already have narrowed.

Call to action: Protect diversity of thought before AI starts collapsing the option space.

1. Define the problem and key assumptions before asking AI to solve it.
2. Use AI to generate competing frames and counterarguments, not only recommendations.
3. Build deliberate challenge points into meetings before consensus forms.

Relevant courses:

[AI Steering Committee and Governance](#) | [Leading Teams Through New AI Workflows](#)

Read more:

[How AI Is Rewriting the Decisions Leaders Need to Make](#) | [AI Can Enhance Every Stage of Teamwork Under Two Conditions](#)

2. Tools: three ways to use always-on agents

This month we've been looking at a new class of AI tool: agents that can take on ongoing work rather than waiting for the next prompt.

The products are multiplying quickly, but we think there are really three distinct ways to use them.

1. Give an AI a job

This is the closest thing to the idea of an AI employee.

Instead of defining a workflow, you give the agent an ongoing responsibility: manage our sales pipeline; watch for relevant tenders; research prospects; keep this system up to date.

The important difference is that you don't necessarily know every step in advance. You know the outcome you want someone to own, and the agent works out how to get there using a browser, files, software and other tools.

This is most useful for messy operational work that crosses systems and changes from case to case.

Use it when: you can describe the responsibility more easily than you can describe the workflow.

Example: Grok Bot.



2. Build an AI into a process

Here, you're not really hiring an AI employee. You're building an automated worker into a defined business process.

You control its instructions, tools and environment. It can run code, manipulate files, call APIs and execute scheduled or repeatable work.

This is a better fit when the process is reasonably well understood and you care about consistency, integration and control: generating reports, processing documents, checking systems for exceptions, performing compliance reviews or coordinating operational workflows.

There is usually more setup involved, but you also get much tighter control over how the work is done.

Use it when: you know roughly how the process should work and want AI to execute it repeatedly inside your business.

Example: Claude Managed Agents.

3. Put an agent inside the ecosystem where you already work

For many organisations, this may ultimately be the most important category.

Rather than introducing a separate agent platform, you add autonomous capability to the productivity environment your people already use. The agent already sits close to your email, documents, meetings, data, identity and permissions.

That makes the buying decision surprisingly simple: start with the ecosystem you're already in.

If your organisation lives in Microsoft 365, there is a strong argument for starting with Microsoft's agent tools before adding another platform. The same logic applies to organisations already centred on ChatGPT or Google Workspace.

The underlying model matters, but integration, permissions and access to the systems where work actually happens may matter more.

Use it when: most of the work you want to automate already happens inside one productivity ecosystem.

Examples: OpenAI Dots and Microsoft Copilot Studio. Google Gemini Spark is heading in the same direction, but we're not reviewing it yet because it isn't currently available for business accounts or in the UK/EEA.

The distinction we find useful is:

- Give it a responsibility.
- Build it into a process.
- Put it into your existing work ecosystem.

Those boundaries will blur, but for now this is a much more useful way to choose between agent platforms than comparing long feature lists.

3. Training Update: Learning Spotlight

We're rebuilding our AI Steering Committee course into short stand-alone modules; here are the first three.

AI steering committees

The AI Steering Committee

Whether you already have an AI steering committee or are setting one up, this course helps you structure it, define roles and close the gap between widespread AI use and the governance needed to manage it responsibly.

AI steering committees

The Business Risk

Understand exactly where your organisation is exposed and use a clear framework to assess your AI risk.

AI steering committees

What AI Is In Business

Identify the highest-impact AI opportunities in your own organisation.

Ready to take the next step?

Already a subscriber? These modules are live in your training portal now.

Thinking about subscribing? Take a look at our [2026 Training Subscriptions](#) to see what's included and how it works for your team.

Want to talk strategy? [Book a briefing](#) or [contact us](#) and we can work through your 2026 AI priorities together.

