

Empower your people



This issue covers:

- onboard before starting
- AI literacy

Read this alongside other issues in the bite-sized AI guidance series

Getting people involved early matters. Consulting staff before significant adoption decisions are made helps identify practical concerns that leaders may not have anticipated. It can surface sticking points in existing workflows, highlight where proposed AI use sits uncomfortably with regulatory obligations, and occasionally reveal that the problem being solved does not actually require AI at all. Finding these things early is considerably better than finding them after implementation.

Engaging staff early also builds the confidence and shared understanding needed for AI to be used well. People who understand why a tool is being introduced, what it will and will not do, and how their role fits within the new process are better placed to exercise the judgement that keeps AI-supported decisions sound ([see the *Organisational integration is key issue*](#)). The goal is to make a genuinely informed decision about whether, and how, AI can improve regulatory practice, with the people who will be closest to it.

Why organisational buy in matters

1. Cultural shift

Research consistently shows that organisations underestimate the groundwork required to bring staff along, and that return on AI investment is very hard to achieve without it. Building AI literacy (more on that below) is an ongoing part of maintaining the critical thinking that keeps AI-supported decisions sound ([see the *AI in Regulation issue*](#)).

2. Role redefinition

How significant that shift will be, and how quickly it happens, will depend on the tools adopted and the regulatory context they are used in. Clear communication about what is changing, and what is not, helps reduce uncertainty and keeps teams engaged. Role-specific support is more useful than general reassurance ([see the *Leadership mindset issue*](#)).

3. Capability building

AI literacy is not simply a new digital skill to add to an existing training programme. Understanding how AI systems work, why they produce the outputs they do, and

how to critically evaluate those outputs requires a different kind of learning, one that is as much about judgement and critical thinking as it is about using a tool. Effective capability building is practical, targeted to specific roles and use cases, and built into day-to-day work.

4. Ethical and governance alignment

Principles such as human-in-the-loop oversight, transparency, and the right to appeal to a human decision-maker only work when people understand what they mean in practice. Embedding these expectations consistently across teams and decision points is an important part of getting people ready for AI ([see the **Organisational integration is key issue**](#)).

Regular review of AI pilots should look not just at what the AI is producing, but at how human reviewers are engaging with it. Designing oversight processes that actively guard against automation bias is as important as the technical performance of the AI system itself.

AI literacy is a key capability

People who do regulatory work need the confidence, curiosity, and practical know-how to test, oversee, and apply AI systems safely within their authorising environment. AI introduces new ways of working, including supervising AI-supported processes and actively questioning outputs rather than accepting them at face value. The foundation for this is AI literacy ([see the **Opportunities and risks issue**](#)).

Staff do not need to be technical experts, but they do need a working understanding

of how AI systems produce their outputs, why those outputs can sometimes be wrong, and what the limitations of the tools they are using are. Without this, human review can become superficial. People tend to defer to outputs that look authoritative, particularly under time pressure, and that deference is precisely what good oversight is meant to prevent.

In a regulatory context, AI literacy means understanding:

- how AI models are trained and why that affects the reliability and fairness of their outputs
- what AI can and cannot do, and where human judgement must lead
- how to question outputs for bias, accuracy, and relevance
- how to recognise hallucination and other common failure modes
- how human oversight and the right to appeal to a human decision-maker are preserved
- how AI fits within risk-based and responsive regulation.

Regulators who develop genuine AI literacy are better placed to identify appropriate use cases, ask the right questions of vendors ([see the **Procure with purpose issue**](#)), recognise when outputs should not be trusted, and maintain accountability for AI-supported decisions.

Starting small gives workers space to learn and adapt

AI implementation is as much about people as technology; building the habits of questioning AI outputs, surfacing risks early, and learning from what does not work. Working iteratively helps organisations separate genuine usefulness from hype. It gives staff direct experience of AI's strengths, such as speed, pattern recognition, and triage, alongside its weaknesses, including hallucination, bias,

and lack of contextual understanding. Early pilots build the staff confidence and organisational knowledge that are prerequisites for using AI well in higher-stakes settings ([see the *Grow as you learn* issue](#)).

To get more practical steps for how regulatory leaders can lead AI innovation with confidence, check out the full guidance: [Responsible AI in Action](#).