

Organisational integration is key



This issue covers:

- aligning AI with regulatory priorities
- why and how to keep humans-in-the-loop

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

Most organisations are still in an early stage of understanding what AI can and cannot do in their specific context. The gap between having access to AI tools and using them to meaningfully improve regulatory outcomes is real, and closing it requires more than making tools available.

A new case management system changes how people do their work. AI can change what work people do, how decisions get made, and where accountability sits. Getting genuine organisational buy-in is what makes the difference. Without it, AI tools tend to be adopted at the margins, used by some and avoided by others, without delivering consistent value or being properly governed.

AI use needs to be integrated into how regulatory work is designed and delivered, not developed in isolation. This means aligning with legal obligations and working across organisational functions including policy, operations, data, privacy, procurement, and digital. Involving these perspectives early helps ensure risks are understood, responsibilities are clear, and systems remain accountable in practice.

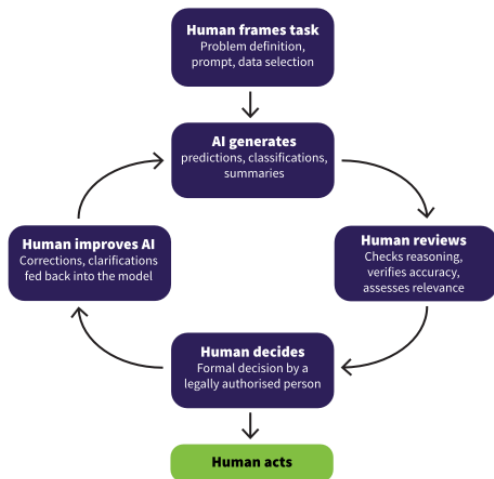
AI adoption needs to be aligned to regulatory priorities, supported by clear governance, and backed by real investment in people [\(see the Empowering your people issue\)](#).

Regulatory leaders will need to make considered decisions about where to direct limited resources; whether that is maintaining business-as-usual capability, improving existing systems, or investing in AI that could meaningfully change how work gets done. Those decisions are easier to make, and easier to defend, when the whole organisation understands the direction and why it matters.

Why and how to keep humans-in-the-loop (HITL)

Human-in-the-Loop (HITL) means ensuring people remain actively involved in decisions that are supported by AI. In a regulatory context this is critical. Decisions that affect rights, obligations, or public trust must be made by authorised decision-makers, not automated systems [\(see the AI in Regulation issue\)](#).

HITL is therefore not simply about having a human present in the process. It requires practitioners who understand what the AI system is doing well enough to genuinely question its outputs, identify where those outputs should not be trusted, and make decisions they can explain and defend independently of what the AI suggested [\(see the Opportunities and risks issue\)](#).



Humans should shape AI long before it produces an output. People choose the tools, define the problem, design governance settings, select data, and interpret results. Human judgement should influence AI at every stage, not just at the point of review.

But, HITL matters most at points of consequence. This includes licensing decisions, enforcement action, compliance assessments, and any decision that may be reviewed, appealed, or challenged. AI is most appropriate in lower-risk regulatory functions where statutory decision-making or judgement about a person's or entity's suitability is not required.

Examples include:

- triaging applications based on risk indicators
- checking applications for completeness and missing information
- validating data against known rules or thresholds
- flagging inconsistencies for human review.

Where AI has influenced a decision affecting a member of the public, that person has a right under section 23 of the Official Information Act 1982 to seek information about how that decision was reached. HITL processes should be designed with this obligation in mind. The role AI played in any decision must be clearly documented and capable of being explained. Accountability remains with the organisation and the legally authorised decision-maker **(see the AI in Regulation issue)**.

HITL is not a complete safeguard on its own. Automation bias can undermine human review even when it is formally in place. Regulators should design HITL processes that actively require reviewers to interrogate outputs rather than simply confirm them **(see the Opportunities and risks issue)**. Oversight mechanisms should periodically review how human reviewers are engaging with AI, not just how the AI system itself is performing.

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