MyConsultancy

Human-in-the-Loop Is an Operating Model Choice

The AI Operating Model Playbook

Manoj Tavarajoo

January 2026

MyConsultancy

Human-in-the-Loop Is an Operating Model Choice

The AI Operating Model Playbook

Manoj Tavarajoo

Opening context

Human-in-the-loop is often discussed as a technical safeguard. It is presented as a control mechanism that reduces risk by ensuring humans remain involved in AI decisions.

While this framing is comforting, it obscures a more important reality. Whether and how humans remain involved in AI systems is not a technical setting. It is an operating model choice.

That choice determines how authority is exercised, how accountability is assigned, and how learning occurs in practice.

Why this fails in most organisations

Many organisations treat human-in-the-loop as a default requirement rather than a deliberate design decision. Humans are inserted into decision flows without clarity about their role, authority, or accountability.

This creates friction. Humans are expected to oversee decisions without sufficient context or time. Responsibility blurs when outcomes are poor. AI systems are constrained by interventions that add little value but slow response.

In other cases, human involvement is minimised in the name of efficiency without redesigning governance or accountability. This shifts risk rather than managing it.

The underlying issue is not the presence or absence of humans, but the absence of intentional design.

The operating model insight

Human-in-the-loop is not about control. It is about where judgement belongs.

In an AI operating model, decisions must be deliberately allocated between humans and systems based on context, risk, and learning requirements. Human involvement should be purposeful, not symbolic.

This requires defining when humans intervene, what authority they have, and how accountability is shared. It also requires recognising that excessive intervention can undermine learning and performance, just as insufficient oversight can increase risk.

Treating human-in-the-loop as an operating model choice forces organisations to confront these trade-offs explicitly.

What this looks like in practice

Poorly designed human-in-the-loop models result in predictable problems. Humans are overwhelmed with low-value reviews. AI systems are slowed by unnecessary checks. Accountability becomes unclear when decisions are overridden without clear authority.

By contrast, organisations that design human involvement deliberately behave differently. Human judgement is applied where context, ethics, or uncertainty demand it. Automation is used where consistency and speed matter more. Authority and accountability align with the chosen design.

Importantly, these choices are revisited as systems and organisational maturity evolve.

Common mistakes to avoid

Treating human-in-the-loop as a universal safeguard regardless of context.

Removing humans entirely without redesigning governance and accountability.

Relying on technical controls alone while ignoring organisational design.

Using human oversight symbolically rather than purposefully.

What leaders must do differently

Leaders must recognise that human involvement in AI is a strategic and organisational decision. They must decide where judgement adds value, where automation is appropriate, and how accountability is maintained.

These decisions cannot be delegated entirely to technical teams. They shape how the organisation operates and how risk is managed over time.

Conclusion

Human-in-the-loop is not a technical switch to be turned on or off. It is an operating model decision with far-reaching implications.

Without intentional design, human involvement becomes either a bottleneck or an illusion of control. Designing it deliberately is essential for responsible, scalable AI.

MyConsultancy

Accelerating Digital Transformation and AI Maturity

www.myconsultancy.com.au