

AI IN LEARNING TECHNOLOGY

The Good, the Bad and the AI.

Why read access is safe, and write access needs a human.

Some of what's arriving in learning platforms is genuinely excellent. Some of it will cost you later. The line between them is more useful than any feature list — and it's easier to draw than you'd think.

THE DISTINCTION

Read or write

Almost every argument about AI in learning technology gets easier once you stop asking "is this good?" and start asking "is it reading, or is it writing?"

AI that **reads** takes information that already exists in your platform and makes it easier to get at. Nothing new is created. Nothing changes. If you switch it off tomorrow, or leave the platform entirely, you lose a convenience and nothing else.

AI that **writes** creates something — content, a learning path, an enrolment, a skills profile, a decision about who needs what. That thing then lives in your system, forms part of your record, and may need to be explained to somebody one day.

READS — LOW REGRET

Answering questions. Surfacing content.
Summarising. Searching. Reporting. Nothing is created, nothing is decided.

WRITES — NEEDS A HUMAN

Authoring content. Building paths. Assigning training.
Inferring skills. Updating existing material. All of it becomes your record.

This isn't an argument against write-access AI. Some of it is genuinely useful and I'd happily use it. It's an argument that the two categories deserve completely different levels of scrutiny, and most organisations are applying the same level to both — usually the lower one.

| Read access is safe. Write access needs a human in the loop, every time.

Everything below follows from that.

THE GOOD

What I'd switch on tomorrow

All of these read. None of them create anything you'd have to explain later.

Asking your platform questions from outside it

The most useful development in this space is the least glamorous. An open standard called the **Model Context Protocol** — MCP — lets an AI assistant read your learning platform directly. A manager can ask, in their own words, who in their team is overdue on a mandatory course, and get an answer without logging in, learning the reporting interface, or asking an administrator to build something.

Several major platforms now offer this or are shipping it. It's worth understanding why it's the good kind:

- It reads data you already own. It doesn't create an asset you might not be able to take with you.
- It's an **open protocol**, not a proprietary format, so it doesn't deepen your dependence on any one vendor.
- It removes a genuine burden — the endless queue of "can you pull me a report on..." requests.
- If you leave, you lose a convenience, not an asset.

For anyone running a platform, this is the single most useful thing on the list, because it takes work off your plate without adding anything to your risk register.

Learner-facing assistants that help people find things

Instead of scrolling a catalogue, a learner asks for what they need and gets pointed at it. The better implementations don't just list results — they suggest: *did you mean this? would this be a better fit?* and link straight through.

This is read-access AI doing exactly what it's good at. Your content is already there; the problem was always that nobody could find it.

Drafting, with a person reviewing

Course outlines, question banks, translations, captions, summaries. All useful, all quick, and all safe **provided a human signs off before it goes live**. The responsible platforms position this explicitly as a first draft for a subject-matter expert to validate. That's the right framing and it's worth holding vendors to it.

Building an audience or a path from a description

Describing what you want in a couple of lines and having the system draft it is a real time-saver, and it lands squarely in the 10% of platform work that stays human. I have no problem with this — as long as an administrator reviews what was built before it goes anywhere near a learner. Which brings us to the difficult part.

THE BAD

What costs you later

Content you can't take with you

This is the one I'd want every buyer to understand, because it's invisible until the day it matters.

I've used an in-platform AI authoring tool in production, and the good parts are genuinely good. It builds a usable course very quickly. You can edit what it produces. You can drop questions into the middle of a video to check people are actually watching rather than leaving it running in another tab. For getting material live quickly, it works.

What you cannot do is take it with you. Content authored inside a platform's own AI tooling frequently cannot be exported in any usable form. So if a significant share of your library has been built that way, then on the day you decide to move — and most organisations move eventually — you discover that your content isn't yours to relocate. You rebuild it, or you stay.

The picture is improving unevenly. Some vendors have added standards-based export for content built in their authoring tools; others haven't; and within a single platform, some AI-generated assets export while others don't. Nobody publishes this clearly.

THE ONLY TEST THAT SETTLES IT

Ask your vendor for a **production-quality SCORM export of a real course built with their AI tooling** — before you commit. Not a roadmap item, not an assurance in an email. A file you can open.

The ownership question nobody asks

Separately from whether you can move it: your ownership claim over AI-generated content is weaker than over content a person wrote. UK law covers computer-generated works, but whether purely AI-generated output attracts copyright at all is genuinely unsettled, and a generic prompt is unlikely to make you the author. If content underpins a certification, gets licensed to customers, or carries your brand into someone else's organisation, that uncertainty belongs on your risk register rather than in a footnote.

AI locked behind the top licence tier

A newer pattern, and an expensive one. Platforms increasingly reserve their best AI capability — assistants, agents, role-play, sometimes MCP itself — for their highest tiers or a separate premium licence type.

The practical consequence: the feature that sold you the platform may not be the feature you can afford to give your whole population. You end up buying enough premium licences for a subset of users, or upgrading the entire estate to reach one capability.

Ask which tier each AI feature requires, in writing, before you sign — and ask what it costs to give it to everyone who'd actually use it.

THE DIFFICULT PART

Write access and the audit trail

This is where I'd want anyone to slow down, and it isn't because the technology is bad. It's because of a question you'll be asked eventually.

Consider how an enrolment normally happens. You build an audience — everyone in this department, in this role, at this location. You attach training to it. Six months later somebody asks why a particular person was assigned a particular course, and the answer is in the rule. You can point at it. You can show it hasn't changed. You can demonstrate the same thing would happen to anyone else in the same position.

If you build the audience yourself, you know exactly why a learner was enrolled.

Now consider the same enrolment made by an agent that assessed a skills gap and acted on it. The outcome may be entirely sensible — better, even. But the reasoning is much harder to reproduce, and in some implementations it isn't recorded at all in a form you could hand to an auditor.

For discretionary development, that's a reasonable trade. For **mandatory and regulated training**, it isn't, because the evidence requirement isn't just "did they complete it" — it's "can you show your process assigns it correctly and consistently."

THREE QUESTIONS WORTH ASKING YOUR VENDOR

1. When an agent creates, assigns or changes something, is that recorded in the audit log with the reasoning — or only the outcome?
2. Can an administrator review and approve agent actions before they take effect, rather than after?
3. Can write-access features be switched off selectively, so agents can suggest but not act?

These are fair questions and good vendors will answer them. Several are actively building better oversight and approval controls, which is exactly the right direction. The point isn't that anyone is behaving badly — it's that this capability arrived faster than the governance around it, and the gap is yours to manage rather than theirs.

One regulatory note, briefly

If AI is inferring employees' skills and that inference influences assessment, development opportunities or progression, it may fall within the high-risk category of the EU AI Act — relevant to UK organisations with EU employees or operations. That would bring obligations around human oversight, transparency to staff, and giving people access to their own inferred profiles.

Be careful with the timing. High-risk obligations were originally set for August 2026, but proposals under the Digital Omnibus package may make them conditional on harmonised standards and push deadlines into 2027 or 2028. Consultation duties under existing law apply regardless. If you're deploying skills inference across an EU workforce, this is a conversation to have with someone qualified rather than a paragraph to rely on.

IN PRACTICE

How this fits the 90/10 rule

A well-run platform runs itself about 90% of the time. AI changes where that automation comes from — and the difference matters more than it first appears.

The 90% has always been **rules**. Everyone in this department gets this course. This certification expires in twelve months, so trigger the refresher at ten. Rules are unglamorous, and they have three properties that matter enormously: they're explainable, they behave identically every time, and you can test them.

An inference has none of those properties. It may well be smarter. It is not more auditable.

	A RULE	AN INFERENCE
Can you explain it?	Yes, read it	Not really
Same result every time?	Yes	Not guaranteed
Can you test it?	Yes	Only by observation
Fit for compliance?	Yes	Be careful

So the sensible position isn't "no AI". It's:

WHERE EACH BELONGS

- Use rules for the 90%.** Enrolment, assignment, chasing, expiry, evidence. Deterministic, auditable, boring, effective.
- Use AI to help people do the 10%.** Drafting, translating, answering questions, finding content, auditing the library. A human is in the loop by design, because the human is doing the work.

And one thing worth saying plainly

The organisations most excited about AI agents are quite often the ones that never configured basic enrolment rules in the first place.

An agent won't fix that. It adds a second thing nobody owns — and this one makes decisions. If your platform isn't doing the straightforward automation it was capable of five years ago, that's the work. AI is a poor substitute for configuration you skipped.

BEFORE YOU SWITCH ANYTHING ON

A short checklist

- ☐ **Is this feature reading or writing?** If it writes, everything below applies.
- ☐ **Can I get the output back out?** Ask for a real export file, not a promise.
- ☐ **Is agent activity in the audit log** — with the reasoning, not just the result?
- ☐ **Can a human approve before it takes effect**, rather than reviewing afterwards?
- ☐ **Can I turn write access off** and keep read access on?
- ☐ **Which licence tier does this need**, and what does it cost for everyone who'd use it?
- ☐ **Is this touching mandatory or regulated training?** If so, keep it on rules.
- ☐ **If it infers things about employees**, have we checked our obligations?

THE SHORT VERSION

Let AI read anything. Let it draft. Be slower about letting it decide, and slowest of all about letting it decide anything you'd have to defend to an auditor.

FINALLY

Who wrote this

I'm Marc Denman. I've spent more than fifteen years administering, configuring, rebuilding and connecting learning platforms — for charities, government bodies and private companies, on estates from 300 learners to 45,000, in the UK and overseas.

I've deliberately not named platforms here. Every observation comes from using these tools or reading the documentation, and naming names would turn a useful piece into a scorecard — unfair to the ones doing this well, and out of date within months either way. If you want a specific view on the platform you're actually running, ask me directly.

Red6 sells no software. No partner badges, no licences, no vendor commission.

Not sure what your platform is doing with AI?

The Red6 LMS Health Check is free: a scored review of your platform across experience, operations, compliance and reporting — including which AI features are switched on, what they're writing, and whether you could get any of it back. Written summary, no obligation.

marc.denman@red6digitallearning.co.uk

red6digitallearning.co.uk

© 2026 Red6 Digital Learning. This paper may be shared freely within your organisation. Regulatory comment is general information, not legal advice.