

No one wants to be the organisation that got AI wrong

The L&D Guide to AI: Conversations from the Field



Investment in AI across regulated sectors is accelerating well beyond most organisations' capacity to govern it. Few have articulated a clear position on staff usage, acceptable risk, or accountability when a tool fails.

Totara's new eBook, *The L&D Guide to AI: Conversations from the Field*, draws together perspectives on where AI is delivering measurable value in L&D, and where the sector is moving faster than the evidence warrants.

Deployment before strategy

AI tooling now features in every conference agenda and every vendor pitch across government, healthcare, and other regulated industries. L&D leaders consistently describe the same structural problem: procurement outpaces implementation planning, leaving staff to absorb new tools without a defined operating model.

This tension is particularly visible in regulated environments. Some organisations decline to adopt AI at all, judging the exposure unacceptable. Others confine experimentation to controlled environments with defined parameters. In most cases, deployment decisions require sign-off from IT security, compliance, legal, and procurement before any tool reaches end users.

That caution reflects a legitimate constraint. AI operates on statistical inference; it has no grasp of regulatory intent or the reasoning behind a given policy. Instructional design, subject-matter judgement, and final review remain human responsibilities. As AI absorbs more production work, the quality of that human oversight becomes more consequential, not less.

Where the value shows up

The eBook does not position AI as a substitute for L&D expertise. It identifies specific functions where AI reduces load on constrained teams: converting static policy documents into structured, trackable learning content, flagging compliance gaps and lapsed certifications, and surfacing an organisation's own approved material to learners rather than the open internet.



Contributors return to a consistent caution. Procuring an authoring tool is not equivalent to having an AI strategy. Durable progress depends on embedding AI into existing workflows and developing capability across the team, not on the tool itself.

The skills question

A central section addresses the capability gap in L&D teams, and it is not a call to develop AI specialists. The requirement is to understand how an AI agent operates, to map a workflow with precision before automating it, and to identify the point at which human review becomes non-negotiable. Organisations under pressure to upskill rapidly, while operating under strict compliance obligations, frequently bypass this step before a tool is selected.

A framework for cautious adoption

The eBook sets out an adoption framework suited to regulated environments. It begins with a foundational audit: whether existing content is accurate and well-structured before AI is layered on top. It specifies where human review must remain mandatory, and challenges the assumption that greater AI usage correlates with better outcomes.

For leaders building organisational AI capability while the technology's reliability is still being established, this offers a disciplined, sector-specific starting point rather than a generic adoption playbook.

Who's in the eBook

The eBook is built on interviews with industry analysts, members of Totara's board and product leadership, and partners operating in regulated sectors. It addresses what is proving effective, what is not, and the considerations government, healthcare, and private-sector L&D teams should weigh before their next AI investment.

[Download the eBook here](#)

