

Are We Asking the Right Questions About A.I. in L&D?

How leaders can navigate AI expectations without rushing into the wrong decisions.



Introduction

There is a concept in orbital science known as the Kessler effect. It describes what happens when too many objects accumulate in the same space. Each addition increases the chance of collision, and every collision creates more debris. Over time, the environment becomes so congested that movement becomes difficult, unpredictable, and eventually unsafe.

The important part is not the science itself, but the pattern. When enough material gathers without coordination or control, the system stops being usable, even if every individual addition once made sense.

Something similar is happening in many AI conversations today.

In learning and development and training management, discussions about AI often begin with clear intent. There is curiosity about potential value, pressure from leadership to explore what is possible, and a genuine desire to make informed progress.

As those conversations evolve, they start to attract more inputs. New tools are introduced. Expectations grow. Assumptions are layered in. External narratives shape internal thinking. All of this happens quickly, and often at the same time.

As the conversation expands, it becomes harder to hold onto a clear center. What began as a practical discussion about outcomes and improvement starts to feel crowded and reactive. Learning leaders are asked to respond more quickly, often without the time or clarity they would like.

The result is a continuing rise in expectations to both implement AI and demonstrate measurable results. At the same time, it becomes less obvious which assumptions are sound, and which choices may be difficult to undo later.

When that happens, the AI conversation can collapse under its own weight, leaving L&D leaders with growing pressure to act, but without a stable frame for doing so responsibly.

For many L&D leaders, they're expected to explain what is being done with AI, why it matters, and what value it is creating, often sooner than feels reasonable.

When that's the case, the conversation often jumps straight to solutions. Slowing it down means paying attention to the questions driving that urgency, and the AI decisions taking form around them.

About Administrate

Administrate is a learning technology company based in Edinburgh, Scotland. Administrate is also the name of our training management system, which is the only TMS dedicated to live training.

Our training management system helps enterprise training teams manage live training. We've helped customers like Maersk Training, Roche Diagnostics, and Siemens Healthineers train thousands of learners every year.

How A.I. is Being Adopted in Learning & Development

AI adoption in learning and development has reached a clear inflection point. According to the AI in L&D 2025: The Race for Impact report, more than half of surveyed L&D practitioners now say they are actively using AI in their work, rather than experimenting or piloting it. This represents a 14% increase year on year and marks what the authors describe as an “implementation inflection” for the profession.

While adoption is rising, the way AI is being used remains uneven. The same report shows that the most common applications are still concentrated in content and design activities, such as creating learning materials, drafting outlines, and supporting instructional design tasks. More complex uses, including analytics, performance support, and strategic insight, are growing but are far less widespread.

This pattern is echoed in Synthesia’s AI in Learning and Development Report 2026. The report found that most L&D teams using AI focus on efficiency gains, particularly faster content creation and localization. Fewer teams report AI being embedded into learning operations or used to support decision-making at a system level.

Importantly, increased usage has not been matched by equivalent clarity around ownership or direction. The Race for Impact report notes that many L&D teams are moving straight from limited experimentation to active use, often influenced by social proof and external success stories, rather than through a deliberate, staged approach. This

shift increases visibility and expectations, even when governance, data foundations, and success criteria are still forming.

Across the reports, a consistent theme emerges. AI is no longer optional or theoretical in L&D. Leaders are under growing pressure to show progress and explain value, often before the shape of that value is fully agreed. Adoption is happening quickly, but alignment around purpose, sequencing, and risk is still catching up.

The questions that help learning leaders stay in control of the A.I. conversation.

AI is already on the agenda for many learning and development leaders. It often surfaces as an executive question about efficiency or scale, or as an expectation that L&D can show how it is keeping pace with the rest of the organization. That pressure frequently sits alongside a broader push to position learning as a strategic partner to the business, rather than a delivery function.

In that environment, slowing the conversation down can feel risky. Hesitation may be interpreted as a lack of confidence, and asking too many questions can appear misaligned with the pace of expectation. At the same time, moving too quickly creates a different risk: being asked to stand behind AI decisions that were never fully examined once results are expected and scrutiny increases.

The questions that follow are intended to help L&D leaders stay in control of AI discussions as expectations rise. They create space to surface what is being assumed about AI before commitments are made, and to explain and defend AI-related decisions clearly once those decisions begin to carry weight.

The right questions to ask about AI in training.

- What problem are we actually trying to improve with AI right now?
- If this worked exactly as hoped, what would be meaningfully better for the business?
- Are we pursuing this because of a clear internal need, or because AI has become an expectation?
- Would this change still matter if AI were not the proposed solution?
- Are we introducing AI at a point where it can support existing operations, or where it will expose gaps we are not ready to address?
- What assumptions are being made about data, process, or ownership that have not been discussed explicitly?
- Who is responsible if this creates confusion, errors, or additional work?
- How will we know when this is working well enough, and when it isn't?

What problem are we actually trying to improve with A.I. right now?

AI conversations in L&D often start with a general sense that something should change, rather than a shared view of what is not working today. Early discussions can move quickly between different motivations, such as reducing manual effort, scaling delivery, or showing progress to executive leadership, without those shifts being made explicit.

When that happens, different stakeholders can end up reacting to different problems at the same time. Decisions may appear aligned on the surface, while in practice they are pulling in different directions. Over time, this makes it harder to judge whether AI is addressing a real operational issue or simply absorbing wider pressure to modernize.

Being clear about the specific problem AI is expected to improve keeps the conversation grounded. It creates a reference point that makes later decisions easier to explain, measure, and defend once expectations around outcomes start to increase.

“What is the specific outcome we want to improve, in plain terms? If we could only improve one part of the training operation with AI this quarter, what would it be and why? If we removed the word “AI” from the discussion, would we still prioritize this problem in the same way?”

If this worked exactly as hoped, what would be meaningfully better for the business?

AI initiatives are often discussed in terms of general benefits, such as efficiency, speed, or scale, without a shared understanding of which outcomes actually matter to the organization. Those improvements can mean very different things depending on whether the business is prioritizing cost control, risk reduction, consistency, growth, or credibility with customers and regulators.

When that distinction is not clear, it becomes difficult to assess whether AI activity is creating real value or simply visible activity. L&D may be asked to demonstrate impact against loosely defined expectations, or to justify trade-offs that were never explicitly discussed. Over time, this can leave learning teams accountable for outcomes that were implied rather than agreed.

Clarity on what would be meaningfully better for the business creates a standard against which progress can be judged. It also helps ensure that AI efforts support broader business priorities, rather than optimizing for metrics that look positive but do not move the organization forward in a material way.

“Which business outcomes would leadership recognize as genuine progress, rather than technical improvement? If this succeeded, what would change for the organization that matters beyond the L&D function? How would we know the difference between real value being created and activity that simply looks productive?”

Are we pursuing this because of a clear internal need, or because A.I. has become an expectation?

AI initiatives often begin before the operational problem is clearly articulated. There is an assumption that AI will reduce effort, increase scale, or improve insight, without first agreeing where that effort is currently concentrated, which processes are under strain, or where information regularly breaks down. When those details remain vague, AI becomes an abstract goal rather than a practical intervention.

When those details are not clear, AI can become an abstract goal rather than a targeted intervention. The conversation focuses on what AI could do in theory, rather than where it would make a practical difference in day-to-day operations. This makes it harder to distinguish between responding to genuine internal need and responding to external signals that suggest AI adoption is expected.

That distinction matters because AI tends to amplify the systems it touches. If it is applied where problems are not well understood, it can add complexity without relieving pressure. If it is applied where friction is already visible, it has a better chance of supporting meaningful improvement rather than creating new work.

“Which parts of the learning operation are currently under the most pressure, and are those the areas this initiative is addressing? If executive expectations around AI disappeared tomorrow, would this still be the place we would focus our effort? What existing problems are we assuming AI will resolve without having named them clearly first?”

Would this change still matter if A.I. were not the proposed solution?

AI often enters the conversation before there is agreement on what actually needs to change. The discussion moves quickly to tools or capabilities, even though the underlying issue might be something L&D has been trying to fix for years.

Take content production. Wanting to create learning faster is not new. The real question is where time is currently being lost, and whether AI helps with that part of the process or just one visible step. If most of the delay sits in reviews, approvals, or coordination, speeding up drafting alone may not change very much.

This question helps separate the value of the change from the appeal of the technology. It makes it easier to see whether AI is being used where it genuinely reduces effort, or simply where it is easiest to introduce and demonstrate.

“If AI were removed from the proposal, what change would still feel worth pursuing? Where in the process is time or effort actually being lost today, and does this address that point? Are we improving the part of the work that limits progress, or the part

that is easiest to accelerate?”

Are we introducing A.I. at a point where it can support existing operations, or where it will expose gaps we are not ready to address?

AI tends to make existing ways of working more visible. When processes are well understood and reasonably consistent, that can be helpful. When they are fragmented, manual, or heavily dependent on workarounds, AI can surface those issues very quickly.

In L&D, this often shows up in areas like scheduling, reporting, or instructor-led training administration. If information is spread across systems, owned by different teams, or handled differently from one program to the next, AI can struggle to produce reliable outputs. What looks like an AI problem is often a reflection of how the operation already works.

“Is AI being introduced into a stable enough part of the operation to be genuinely supportive? Or is it likely to highlight gaps that have been manageable until now, but become harder to ignore once AI is involved?”

What assumptions are being made about data, process, or ownership that have not been discussed explicitly?

There may be an assumption that data is consistent, that processes are followed in the same way across teams, or that responsibility for outcomes is already clear. When those assumptions are wrong, the impact usually shows up later, once AI outputs start to be relied on.

In L&D, this can be especially subtle. Data may live across multiple systems, be updated inconsistently, or mean different things to different teams. Processes may exist in principle, but be handled differently in practice. Ownership may be shared, informal, or assumed rather than defined. AI does not

create these conditions, but it depends on them.

When assumptions remain unspoken, AI can appear unreliable or difficult to trust, even though the issue sits further upstream. Surfacing these assumptions early helps avoid situations where L&D is asked to explain outcomes that were shaped by decisions no one realized they were making.

“If this produces the wrong result, would it be clear why? Or would the cause be buried across systems, processes, or handoffs?”

Who is responsible if this creates confusion, errors, or additional work?

When AI is introduced, a lot of attention is paid to what the system will do and where it will be used. Much less attention is paid to who is expected to deal with the consequences when the output is wrong, incomplete, or creates extra work.

In practice, those consequences tend to show up quickly. Learners receive information that needs correcting. Teams add manual checks to avoid mistakes. Someone has to explain why the output cannot be relied on yet. Even if the decision to use AI was made collectively, the responsibility for managing these issues often ends up with L&D.

“If this produces output that people question, who decides whether it is good enough to use? If stakeholders ask why the result can’t be trusted yet, who is responsible for fixing and explaining that gap?”

How will we know when this is working well enough, and when it isn’t?

Once AI is in use, decisions are no longer about whether it should exist, but about how much it can be relied on. Outputs are produced, processes adjust around them, and expectations form about what should now be faster, easier, or more automated. At that point, the question of “is this working?” becomes harder to answer clearly.

What often gets lost is a shared view of what “good enough” actually means. Is it fewer manual steps, faster turnaround, fewer errors, or simply that something is running at all? Without agreement on that threshold, progress can look positive even when confidence in the results is mixed or uneven. Concerns tend to surface late, once reliance has already increased.

Clarity on how effectiveness will be judged makes it easier to intervene before reliance turns into obligation. It allows learning leaders to distinguish between acceptable limitations and signs that the approach needs to change, and to do so while the decision still feels open rather than fixed.

“What would tell us that this is no longer reliable enough to depend on? Where is the line between acceptable limitation and a real problem in this context? Who has the authority to say that line has been crossed?”

What all of these questions point to.

Taken together, these questions surface a pattern that often gets lost in day-to-day AI discussions.

When AI initiatives struggle in learning and development, it is rarely because the technology itself is incapable. More often, it is because AI is being asked to operate inside systems that are fragmented, inconsistently understood, or only held together by workarounds that exist in people's heads rather than in formal processes.

In those conditions, AI can reveal uncertainties rather than simplify work.

Ways of working that were once flexible or informal tend to become more defined. Variations that were previously absorbed through experience start to matter in new ways, simply because outputs are now being generated and shared more widely.

This shift brings a different set of considerations into focus. Questions about consistency, traceability, and decision-making naturally move earlier in the conversation. The emphasis moves away from what is possible in theory, and toward what can be relied on in practice.

In that context, pacing matters. Introducing AI successfully often depends less on speed, and more on ensuring that the underlying structures it draws on are clear enough to support it with confidence.

Why A.I. In L&D Must be Shaped by Operational Readiness, Not Models

AI tends to perform best when it has a clear operational context to work within. That context is created less by technical sophistication and more by how learning activity is organized day to day.

When data is coherent, processes are understood, and responsibility is clear, AI outputs are easier to interpret and trust. When those elements are loosely defined, AI still functions, but its value is harder to judge with confidence.

In learning environments, this distinction matters. Training operations often span multiple systems and teams, and evolve over time to meet changing needs. Introducing AI into that environment naturally places more emphasis on how those moving parts connect.

Seen this way, successful AI use in L&D is less about selecting advanced capabilities, and more about ensuring that the operational layer beneath them is legible and aligned.

What operational readiness looks like in real terms.

Operational readiness does not imply rigid standardization or complete uniformity. It is about having enough shared understanding for the system to behave predictably.

In practice, this means clarity around where data originates, how it is interpreted, and how it moves through learning workflows. It means processes that can be described without relying on individual

memory. It means knowing where flexibility is intentional, and where consistency matters.

When AI is introduced into an environment like this, it becomes easier to set expectations appropriately. Outputs can be used with a clear sense of their boundaries. Decisions about reliance feel deliberate rather than provisional. Conversations about improvement become more concrete.

For learning leaders, this kind of clarity supports confident decision-making and clear communication with stakeholders, without requiring technical depth or continuous justification.

From this perspective, progress with AI looks slightly different.

It is not defined solely by how quickly new capabilities are adopted, or how prominently AI features appear in the learning stack. It is defined by how well AI fits into existing operations, and how comfortably its outputs can be explained and relied on.

This reframing allows AI discussions to become more grounded. It creates space to talk about sequencing, scope, and readiness without those considerations being interpreted as hesitation. It also supports a more durable definition of success, one that values confidence and reliability alongside efficiency.

A stable foundation for learning operations and A.I. implementation.

As AI continues to enter learning and development, the most durable progress tends to come from systems that already reflect how training actually operates across schedules, instructors, resources, locations, dependencies, and the many small decisions that make delivery possible.

That operational layer is where clarity is built. It's where data becomes consistent enough to trust, processes become visible enough to explain, and coordination becomes possible without relying on individual memory or manual reconciliation. When that foundation is in place, new capabilities can be introduced sustainably, without destabilizing the work that already needs to happen.



Administrate is the foundation A.I. needs to deliver value.

AI performs best with coherent data, clear processes, and defined responsibility.

Administrate establishes all of these through a centralized training management platform

By organizing learning operations in one system (schedules, instructors, resources, locations, dependencies), Administrate becomes the operational layer that makes AI outputs trustworthy.

Administrate's structured approach ensures AI fits sustainably into existing workflows rather than destabilizing them.

Administrate ensures you have the consistent data, visible processes and coordination infrastructure that allow AI to function predictably.

Administrate brings operational A.I. to enterprise training.

API-first interoperability synchronizes every system, creating clean data for smarter decisions. Operational AI automates workflows at scale. Revenue functionality delivers results: our customers generate seven and eight-figure increases.

Book time to see
Administrate

Get a demo

 1 (800) 265-7163

 info@getadministrate.com

 getadministrate.com