The LMS as Platform in the AI Era

By Matthew Pittinsky, Ph.D. · **DRAFT** · *Last Revised May 8, 2026*

I was recently reminded of one of my most memorable moments from the early days of Blackboard. It was 1999 and across the table from me were Will Ethridge, then president of Pearson's Higher Education and Professional Publishing Group, Jim Bankey, and Pat Leonard. Blackboard was a small company; Pearson was the largest education publisher in the world. They wanted to license our software platform to wrap course management around offerings like MyMathLab. It was a classic "build versus buy" decision, with a wrinkle: institutions were increasingly standardizing on Blackboard as their campuswide digital learning platform, so if Pearson invested in integrating their content into Blackboard, faculty would already be fluent, and the same integrations could be deployed locally on each adopting campus. In other words, it was also a platform decision.

The negotiations took forever.[1] Pearson was worried -- reasonably -- about being disintermediated. About the LMS becoming the aggregator of faculty and institutions, with the publisher reduced to what today might be described as one tile in someone else's app store. The fact that we called third-party publisher offerings *course cartridges* did not help the case! We eventually agreed on a private-label deal announced in June 2000, and Pearson invested in Blackboard as part of the agreement, partly to realize some of the value created by anointing us and partly to be inside the tent of a potential disruptor.

1. Seriously, I was 27 with the maturity of a 18-year-old at best and I would have gnawed my arm off to escape... if the reason I was trapped was handcuffs, which it wasn't, except in the metaphorical golden sense given Pearson's market heft.

I share this story because as I think more about AI's impact on education, one of the implications I keep returning to is this: pedagogical value will move from the publishing layer to the delivery system. I alluded to this idea in passing in "The LMS at 30 Part 2." It deserves expanded consideration because its consequences are not confined to the LMS. The shape of the LMS as a platform reaches into the commercial education publishers (and OER), into the third-party EdTech ecosystem, and into how course materials are licensed.

Pedagogy and the Publisher Layer

Education publishing is more than just assembled, vetted content. Textbooks are a canonical example. The content is chunked into units, chapters, sections, subsections at sizes calibrated to working memory. Each chapter opens with learning objectives and a "before you read" frame and closes with a summary and a comprehensive set of review problems. New concepts are scaffolded with examples: a fully solved problem with every step explained, which is the design choice that cognitive-load research has shown most reliably helps a novice learner build a mental model. End-of-section practice runs a gradient from procedural to conceptual to applied. Key terms are bolded, defined in the margin, and gathered in a glossary. Diagrams are placed exactly where the prose needs them. Cumulative review at the end of the unit pulls forward what was learned earlier, a concept called *spaced retrieval*. And behind the visible book sit a study guide, a solutions manual, and a test bank for the instructor.

Each of those design elements represents pedagogical value. The content is necessary, but the content is not the only thing. The design is also the thing. A good textbook is the result of decades of authoring, peer review,

classroom field testing, and revision, layering all of that into a single bound (or licensed) artifact. *That* is what I mean by pedagogical value living in the publishing layer, at least in the analog era.

As publishers shifted from analog to digital, first with PC-based software and then Internet-based, the artifact of "courseware" embodied its own set of pedagogically driven design choices. Offerings like Pearson's MyLab series, McGraw Hill's Connect, MasteringChemistry, and ALEKS extended the textbook's pedagogical scaffolding with capabilities like adaptive practice. "Primitives" from the analog publishing era have evolved; for example, the worked example is now interactive. The test bank is automatically graded. These are genuine advances. But notice what did *not* change. The publisher remained the locus of pedagogical design. And, as a consequence, the pedagogical design remained per-publisher, per-course, maybe per program, with each student visiting different ones for different subjects. The LMS played a wrapper role -- managing access, accepting homework, and recording grades with a consistent user experience. But the pedagogical value stayed where it had always lived, with the publisher. So far so good, Will, Jim, and Pat.

Pedagogy and the Delivery System

Why do I believe one consequence of AI is that pedagogical value will shift from the publisher layer to the delivery system? First, pedagogical design elements that previously were designed, "printed," and shipped – e.g., level-appropriate explanations, summaries, simulations – will increasingly be produced on demand by the LMS, drawing on the corpus of the course as raw material rather than as a finished pedagogical artifact.

Second, and more importantly, the algorithms that select the resources a particular learner needs at a particular moment necessarily need to operate at the platform level rather than at the publisher level if they are to benefit from understanding a student's full learning history, not just one chapter or one course. More on this in a moment. The LMS rebrand from Course Management to Learning Management, aspirational since the early 2000s, is finally about to be true.

I suspect this shift will occur slowly over a longer period than "AI Pilled" futurists predict. But it will occur, with implications for what commercial education publishers compete on, what platform-level affordances LMS providers expose to third-party EdTech tools to build against, and how course materials are licensed for any of this to work.

The LMS as Platform

As I described in "The LMS at 30," the LMS category had a strong dimension of platform to it from the beginning. The very first Blackboard business plan, in 1996, defined an instructional management system as "an operating environment in which teaching and learning can occur." In the abstract of Blackboard's mid-1997 business plan, I wrote:

> *Even a 19th century one-room schoolhouse had basic instruments for engaging in the learning process. The Blackboard was the operating environment. Chalk was the development tool. The pointer a navigation device. Ironically, here in the 20th century personal computers lack a comparable type of platform... Think of it as a learning platform. Think of it as an operating environment. Think of it as an instructional management system.*

The vocabulary was unsettled -- *learning platform, operating environment, instructional management system* -- but the conception was clear. The Educom Instructional Management Systems project, launched in 1996 with commercial partners and major research universities, was working in parallel toward the same idea. A project FAQ that autumn defined IMS this way:

> *An industry specification project aimed at creating an open architecture that vendors and I/T departments can base the development of their software [on] to ensure compatibility and simpler integration.*

In other words: middleware for discipline-specific, pedagogically rich tools, unified around the learner and the instructor. Five years later, in October 2001, Blackboard launched the Building Blocks program with an open SDK and a first cohort of partners ranging from Carnegie Mellon's Virtual Chemistry Lab to Centra, HarvestRoad, Questionmark, and WebEQ.[2]

IMS eventually became 1EdTech, which continues to champion LTI and much of the interoperability stack on which the modern EdTech ecosystem sits. Walk into any LMS conference exhibit hall and you can see the result: hundreds of vendors plugging into a handful of platforms through standards 1EdTech stewards. Back in the late 1990s, I expected each LMS course to reflect its discipline, its instructor's pedagogy, and its institution's approach to learning. Bounce between LMS courses at a sample of institutions today, and that level of differentiation is the exception not the rule.

2. I can't help but note that this type of app ecosystem approach was quite novel for on-premise "enterprise software," predating even Salesforce's AppExchange, which launched in 2005.

This Time is Different

So, at this point we are 0-2 on predictions or at least worries. The LMS did not disintermediate and commoditize education publishers as Pearson feared and, generally speaking, LMS courses did not meaningfully differentiate by discipline and pedagogy, at least at the LMS layer vs. what the LMS linked learners into. I don't think these predictions were wrong in substance, just in timing. For the two reasons introduced earlier, their fulfillment feels imminent: the affordances AI provides software to support learning, and the centrifugal forces of personalization and orchestration that move pedagogical value "down" into the LMS.

As noted earlier, a typical undergraduate might use one publisher's eText in one course, a different publisher's homework system in a second, an OER reader in a third, plus the professor's own materials. Often, no single publisher covers an entire section, let alone an institution, let alone the multi-year progression a degree represents. In the AI era, the pedagogical layer of education is a *many-to-many* problem -- many publishers, many tools, many courses, many learners -- and a publisher-bound "intelligence" layer answers a many-to-many problem by building yet another silo.

The costs of doing it that way are concrete. The learner data model fragments across publishers, so no single layer sees the full picture. Faculty and institutions lose transparency and accountability over what is being recommended to students, on what basis, and with what guardrails. Personalization gets harder, not easier, because each engine sees only its own course's content and not the rest of the student's week and work.

From the outside it seems the publishers see the centrifugal forces coming and are racing to keep the pedagogical layer in their hands. McGraw Hill now describes itself not as a publisher but as an "education science company," investing in "data-driven personalized learning experiences, interactive simulations, 3-D models, experiential learning activities, and generative AI tools." Pearson describes itself as "the world's lifelong learning company," with product positioning to match: a Pearson+ Channels feature it calls an "AI Tutor" that ingests a student's syllabus to sequence personalized study; an AI Reader inside McGraw Hill Connect that generates explanations and quizzes from any passage; ALEKS expanding into Calculus and K–3 on the strength of a quarter century of student-skill data.

The publishers have deep content, decades of learning-outcomes research, and real adaptivity engines. The AI-augmented versions of those products will, in many cases, work well for students. That said, the logic of the platform taking on more of what these publishers are building seems overwhelming.

Pedagogy and the Delivery System, Continued

Justin Ménard's recent ListEdTech analysis of Canvas is illustrative: he argues the most widely adopted LMS in higher education is pursuing a strategy in which the LMS moves from "system of record" to "system of interaction." Ménard's analysis points to an LMS that takes on more of the active instructional role that the publishers used to monopolize.

Consider four broad capabilities that AI enables. First, a *proactive academic planner* that holds the program of study, tracks deadlines, flags missing assignments, and models GPA scenarios. Second, a *GPA fitness*

app that treats academic performance the way Strava treats running -- streaks, nudges, the accountability of a friend keeping score with you. Third, a *Hogwarts Book of Magical Study Aids* -- software that generates flashcards, practice sets, simulations, and Socratic dialogues from any course corpus on demand. And fourth, a *deep course-aware context*, guided by the instructor, that indexes the syllabus, the assessments, and the rubric and helps a student work through them -- software, alongside the human TA, that scales the kind of help every student should have access to.

None of those four are publisher-specific applications. All four require an "intelligence" layer that crosses courses and operates on the full content of every course under whatever license the institution has secured. Compared to today's scope of LMS interoperability -- gradebook passback, identity, rostering -- what is being asked of the platform is an order of magnitude wider, and an order of magnitude more valuable.

From SaaS to PaaS -- Vertical PaaS for Education

An extreme form of this argument is that the LMS will shift from software-as-a-service (SaaS) to platform-as-a-service (PaaS), something ASU CIO Lev Gonick has written about in his TechEd series. As the delivery system gets more capable, the LMS becomes less bundled software and more programmable platform. Curriculum providers, third-party tool builders, and the institutions themselves build on a shared layer of identity, learner data, compliance, course context, and software orchestration, instead of each one rebuilding that infrastructure from scratch. The most ambitious mastery- and competency-based providers I know are already constructing this kind of infrastructure for their own use; the next move is for the field to stop reinventing it

privately and to make it shared infrastructure inside the system institutions already govern. *Vertical PaaS for education* is a mouthful and also, notably, the conception IMS and the 1997 Blackboard abstract had in mind when it asked the reader to think of the LMS as platform and operating environment.

The “Power Grid” Objection

An objection I’ve heard when discussing these ideas is that the LMS should remain trusted infrastructure -- closer to a power grid than to a pedagogical actor. In that view, the LMS's job is to deliver the aforementioned access, identity, enrollment, workflow, grades, and institutional governance reliably, predictably, and interoperably. Pedagogical value belongs in discipline-specific tools built by expert authors with domain-specific learning models. Making the LMS "smart" risks turning neutral infrastructure into an opaque instructional actor -- harder to manage, more invasive to trust, more expensive to maintain, and less pedagogically precise than the publisher and courseware products that have been refining their craft for decades.

I take that risk seriously. But I do not think it survives the compelling capabilities of generative AI. When software can produce worked examples, practice gradients, and explanations from any course corpus on demand, the question is no longer whether the platform will host pedagogy. It will. The question is whether the institution gets to govern what the platform does, on what data, and under whose policies. Keeping the LMS "dumb" will not preserve neutrality; it will cede the pedagogical layer to whatever AI sits closest to the learner. The instinct to keep the LMS predictable and well-bounded is the right one. The mistake is in thinking it requires keeping the LMS inert.

Implications for Course Materials Licensing

A delivery system can only do this work if it has rights to operate on the *full corpus of the course* -- every text, dataset, problem set, and learning application assigned to every enrolled student. That is incompatible with the historical model in which each student decides individually whether to buy or rent the assigned materials, with the result that on day one a fraction of the class has full access to course materials.

Two licensing models the industry has been quietly migrating toward for a decade fit this moment exactly. *Inclusive access* -- sometimes called *course* access -- gives every student in a course Day-One rights to the assigned materials at a negotiated price folded into tuition or course fees. *Equitable access* -- sometimes called *total access* -- extends the same idea to a campus-level bundle across publishers so that every enrolled student has rights to every assigned title across every course. The original case for both was the simple finding that students perform better when they all have the materials on day one. There are also fantastic win (learner), win (college), win (publisher), win (costs to both) dynamics.

I recognize that course and total access programs are also contested in higher education -- around opt-out transparency, student choice, price competition, and whether day-one access reliably improves academic outcomes. Those concerns are worth taking seriously on their own terms. That said, they do not eliminate the structural point: an AI-enabled delivery system will need a clear rights envelope to operate across the full corpus of a course. The form of that envelope may evolve. The need will not if the delivery system is to generate, personalize, and adapt across the full corpus the way the generative AI era will demand. Total and course access were already on the rise. The migration of

pedagogical value to the delivery system gives them an additional tailwind they did not have before. Content licensing and platform architecture become linked.

A Coda on Locks and Tolls and Standards

A worry to name before closing, followed by an alternative path. We are in a fraught moment around AI software acting on behalf of learners and undermining academic integrity. That worry is well-founded -- institutions need confidence that the entity engaging with the learning environment is the enrolled human, not software acting on her behalf. That said, these legitimate concerns could create cover for a major LMS player to lock down its APIs, charge increased tolls for partner status, and narrow who gets to integrate at all... in the name of AI safety. I am already hearing from early-stage companies that this is happening in practice, not just in policy documents. Some reconsideration of platform access is warranted. But we must guard against opportunism dressed in the language of integrity, and it is a particularly dangerous instinct at this moment when the balance of platform-level services and ecosystem-level services is shifting.

There is also, in principle, a different path that could satisfy the centrifugal forces in a different way. Open standards and protocols could enable a world where distributed courseware and learning applications interoperate at the granularity AI personalization requires (e.g., learner-level state, mastery signals, generative-content rights, cross-course context).

I think this scenario is less likely for two reasons. First, the competitive drivers point the wrong way. LMS providers have strong commercial

incentives to bring more of the learning capability into their platforms, not federate it through standards. Second, the history is sobering. IMS turned 1EdTech and LTI are meaningful achievements -- I am a big fan -- but the level of granularity and adoption a standards-only model would require to satisfy the needs of AI-era personalization is, so far, unprecedented in this sector. Perhaps a horizontal, general AI/agent protocol will do the trick, relegating me to 0-3 on predictions!

If not, twenty-six years after our negotiations, Will, Jim, and Pat's concerns are finally arriving as AI moves the pedagogical value into the LMS. If true, the LMS will become what it was always supposed to be, a platform for learning.

Matthew Pittinsky is the incoming CEO of Blackboard, which he co-founded in 1997 and previously helped lead as CEO and Executive Chairman. More at matthewpittinsky.com.