RESEARCH BRIEF

The $103 Problem: Why Training Costs Are Rising and What AI-Native Platforms Change

Industry benchmark data shows cost per learning hour has surged 36% while delivery methods remain stuck in the past. This research brief examines how AI-native training platforms fundamentally restructure the economics of talent development.

April 2026 • 16 pages

hey LOOPY

Contents

About the data

Industry benchmarks referenced in this brief are drawn from publicly available findings in the Association for Talent Development's *State of the Industry* reports (2014-2022) and McKinsey & Company's *Reimagined: Development in the Future of Work* (2025 Perspective on Evolving Trends in L&D). All analysis, interpretation, and recommendations are original to HeyLoopy.

Executive Summary

The talent development industry has a cost problem that isn't visible in the headline numbers. Average spending per employee ($1,280) has barely moved in eight years. But underneath that stability, the **cost per learning hour used has surged to $103** - a 36% increase from the $76 average that held for most of the prior decade.

This isn't because organizations are buying more expensive programs. It's because the dominant delivery methods haven't changed despite a surface-level shift to technology. Most "technology-based" training is simply instructor-led sessions moved to video calls, carrying the same marginal cost per learner.

At the same time, the fastest-growing content category - managerial and supervisory training, now 20% of the average portfolio - is also the most expensive to deliver. And the learning method that organizations say they value most (on-the-job learning, with 62% commitment) is the one they measure least.

$103
Cost per learning hour used - up 36% from the prior decade average

5.0%
Of payroll spent on learning - the highest level ever recorded

33 hrs
Average learning hours per employee - declining since 2019

The result is a system under increasing pressure: flat budgets, rising unit costs, growing demand for the most expensive content types, and limited visibility into whether any of it is working.

McKinsey's 2025 research reinforces the urgency: while **72% of organizations** recognize the importance of becoming skills-based, only **11% report meaningful progress**. And 61% still plan their workforce strategy only one year out, leaving little room for the foresight that modern talent development demands.

This brief argues that AI-native training platforms represent a structural shift - not an incremental improvement - in the economics of talent development. By eliminating the marginal instructor cost, enabling content creation in minutes rather than months, and building measurement directly into the learning experience, these platforms address the root causes of the $103 problem rather than optimizing around them.

Section 1: The Spending Plateau

Between 2014 and 2021, average direct learning expenditure per employee moved from $1,229 to $1,280. That's a cumulative increase of 4.1% over eight years. During the same period, inflation in advanced economies compounded by over 15%.

In real terms, L&D budgets have been shrinking.

This matters because the demands on those budgets have not shrunk. Organizations are being asked to train more people, on more topics, across more locations, with more urgency - all on a budget that buys less each year. According to the World Economic Forum, **39% of existing skill sets will be transformed or become outdated between 2025 and 2030**. The training load is growing. The budget to handle it is not.

$1,229	$1,280	4.1%
Per employee spending, 2014	Per employee spending, 2021	Total growth in 8 years (vs. 15%+ inflation)

The size effect has reversed

For years, large employers (10,000+ employees) enjoyed cost advantages from scale - they could spread the same training programs across more people, driving down per-employee costs. In 2021, that dynamic reversed.

Large companies reported average expenditures of **$1,656 per employee** - 73% more than midsize organizations ($955) and 53% more than small companies ($1,083). The organizations that should benefit most from scale are now spending the most.

Several factors contribute to this reversal:

- **Complexity of distributed workforces.** Large organizations typically have more locations, more time zones, and more variation in role types - all of which increase delivery costs.
- **Return-to-office retraining.** Post-pandemic, large employers invested heavily in retraining workforces that had been reduced during 2020.
- **Compliance at scale.** Regulatory training requirements don't become cheaper with more employees - they become more expensive to coordinate.

> **What this means for L&D leaders**

> If your organization is large, you can no longer assume that scale works in your favor. The per-employee cost advantage that justified centralized training models has eroded. The question is whether your delivery methods can restore that advantage - or whether a fundamentally different approach is required.

Direct expenditure as a percentage of payroll

In 2021, direct training expenditure reached **5.0% of total payroll** - the highest percentage ever recorded in the State of the Industry report. The prior peak was 4.25% in 2015.

This increase is partly mechanical: many organizations reduced headcount during 2020, shrinking the payroll denominator while training budgets held relatively steady. But it also reflects a real increase in the priority organizations place on workforce development, particularly in a tight labor market where training and development opportunities are a key retention lever.

The challenge is that this investment is not translating into proportionally more learning. Average learning hours per employee actually *declined* from 35.0 in 2019 to 32.9 in 2021. Organizations are spending a record share of payroll on training and delivering fewer hours of it.

Section 2: The Delivery Method Trap

The headline statistic looks encouraging: **67% of learning hours** are now delivered through technology-based methods. In 2019, that figure was 56%. The shift to technology-based delivery appears to be well underway.

But this number is misleading in a critical way.

The growth in technology-based delivery came overwhelmingly from one category: **instructor-led virtual classrooms (vILT)**. This is the same lecture-based format that dominated in-person training, moved to a video conferencing platform. It still requires:

- A live instructor for every session
- Scheduling coordination across participants
- Real-time facilitation and Q&A management
- A fixed ratio of instructor hours to learner hours

The marginal cost of adding one more learner to a vILT session is lower than in-person (no travel, no room), but the *structural* cost is unchanged. You still need an instructor. You still need a scheduled time block. You still have a ceiling on how many people can meaningfully participate.

What happened

- Classroom training moved to Zoom
- vILT replaced ILT as the primary method
- "Technology-based" percentage went up
- Cost structure remained instructor-dependent

What needed to happen

- Shift from instructor-dependent to self-paced
- Active learning replaced passive lectures
- Marginal cost per learner approached zero
- Quality scaled independently of instructor availability

Self-paced e-learning: still stuck at 30%

True self-paced e-learning - the delivery method that actually scales without marginal instructor cost - has remained at approximately **30% of learning hours used** since 2020. It didn't grow during the shift to technology. The technology shift bypassed it entirely.

This is the core of the cost problem. The one delivery method that could structurally lower cost per learning hour is the one that hasn't grown. And the method that did grow (vILT) simply moved the existing cost structure to a new medium.

> *"The platform shifted. The cost structure didn't. Organizations adopted technology as a channel for the same instructor-dependent model - like digitizing a paper form without changing the process behind it."*

Where the $103 comes from

Cost per learning hour used is calculated by dividing total direct learning expenditure by total learning hours consumed across the organization. At $103, this figure is 36% higher than the approximately $76 average that held from 2014 through 2018.

The drivers are compounding:

1. **Budgets held flat** while inflation eroded purchasing power
2. **Learning hours declined** from 35.0 to 32.9 per employee (a smaller denominator)
3. **Delivery methods remained instructor-dependent**, preventing cost-per-hour improvements from technology adoption
4. **Content mix shifted** toward more expensive categories (managerial training)

Each factor alone would pressure the cost per hour upward. Together, they produce a 36% spike that shows no sign of reversing under the current model.

Section 3: The Managerial Training Crisis

Managerial and supervisory training now represents **20% of the average learning portfolio** - up from 13% just two years prior. This is the largest single content category, surpassing even mandatory and compliance training.

The growth makes sense. The pandemic forced a rapid shift to remote and hybrid work, and managers who had relied on in-person proximity to lead their teams suddenly needed new skills: virtual facilitation, asynchronous communication, remote performance management, and digital-first coaching.

ATD's research found that 77% of managers had at least one direct report working partially or fully from another location by early 2021, compared with just 21% in 2019. Organizations responded by expanding virtual manager training programs - 56% offered them, with another 27% planning to add them.

The cost problem within the cost problem

Managerial training is inherently expensive because it typically requires:

- **Small cohort sizes.** Leadership development works best in small groups with discussion and practice, not lecture halls of 200.
- **Senior facilitators.** You can't have a junior trainer teach leadership skills. The facilitator needs credibility and experience.
- **Extended programs.** A compliance module takes 30 minutes. A management development program takes weeks or months.
- **Role-play and practice.** Effective manager training requires simulated conversations, feedback practice, and coaching - all labor-intensive to facilitate.

When this category grows from 13% to 20% of the portfolio, it pulls the entire cost structure upward. And the demand isn't slowing - if anything, the shift to AI-augmented workforces is creating a new wave of management challenges that will require even more development.

> **The AI coaching opportunity**
>
> Managerial training is where AI-powered coaching platforms create the most dramatic cost advantage. An AI coach can simulate difficult conversations, provide instant feedback on communication style, and offer unlimited practice reps - capabilities that previously required a live facilitator for every session. The cost per practice session drops from hundreds of dollars to effectively zero.

Section 4: The On-the-Job Learning Gap

For six consecutive years, the ATD State of the Industry report has tracked organizations' commitment to on-the-job learning. And for six consecutive years, the finding has been remarkably consistent: **62% of organizations say they emphasize on-the-job learning to a "high or very high" extent.**

On-the-job learning includes knowledge sharing and coaching during work, job shadowing, rotational assignments, stretch assignments, and committee placements. It's the type of learning most likely to be relevant, timely, and directly applicable to performance.

It's also the type of learning that almost nobody measures.

The measurement gap

Formal learning is easy to track: hours completed, modules finished, certifications earned. These metrics fill dashboards and annual reports. On-the-job learning, by contrast, is inherently difficult to quantify because it happens in the flow of work, not in a separate "learning" context.

The result is a structural bias: organizations invest in what they can measure, and they can measure formal training. So formal training gets the budget, the technology, and the attention - even when everyone agrees that informal, on-the-job learning is where the real skill building happens.

62%
Say they're committed to on-the-job learning

~0%
Can effectively measure on-the-job learning outcomes

42%
Use technology for on-the-job knowledge sharing (up from 24% in 2018)

Technology-assisted on-the-job learning is growing

One bright spot: the use of technology for on-the-job knowledge sharing has grown steadily, from 24% of organizations in 2018 to 42% in 2021. This suggests that the infrastructure for measurable on-the-job learning is being built, even if outcomes tracking hasn't caught up.

McKinsey's 2025 learning perspective describes the future state clearly: "The goal is no longer to add learning into the flow of work - it's to merge work and development. Daily work is now designed as a developmental engine." AI copilots act as "real-time mentors, adjusting their support based on performance, stress, and cognitive load." Learning becomes "continuous, highly personal, and largely invisible."

AI-native training platforms sit squarely in this gap. They deliver training in the flow of work (short, daily interactions rather than pulled-out sessions), they draw content from the actual materials people use on the job (SOPs, playbooks, product docs), and they generate competency data as a byproduct of the training itself.

This is what "closing the on-the-job learning gap" actually looks like: not asking managers to log informal coaching sessions, but embedding training into daily work in a way that automatically produces the measurement data L&D teams need.

Section 5: How AI-Native Platforms Change the Math

The $103 problem has four root causes: flat budgets, instructor-dependent delivery, expensive content categories, and unmeasured on-the-job learning. AI-native training platforms address all four simultaneously.

1. Zero marginal instructor cost

In a traditional model, every learning hour requires an instructor hour (or a fraction of one, depending on class size). In an AI-native model, the AI coach delivers personalized training to each learner independently. The hundredth learner costs the same as the first.

This doesn't eliminate the need for human instructors entirely - complex leadership development and culture-building still benefit from human facilitation. But for the 70-80% of training content that is knowledge transfer, process training, product education, and compliance, AI delivery eliminates the marginal cost constraint.

> **Cost modeling example**
>
> A 500-person organization running 8 hours of product training per quarter currently needs approximately 160 instructor-hours annually (assuming 25-person sessions). At $150/hr for a skilled facilitator, that's $24,000 in instructor cost alone - before content development, scheduling overhead, or platform fees. An AI-native platform delivers the same knowledge transfer at a flat subscription cost, regardless of how many sessions or learners are involved.

2. Content creation in minutes, not months

Traditional e-learning development follows a well-documented timeline: 40-200 hours of development time per finished hour of training content, depending on complexity. For a custom module with interactions and assessments, 100:1 is a common ratio.

AI-native platforms invert this ratio. You upload existing documents - SOPs, playbooks, product manuals, process guides - and the platform generates interactive training from them automatically. The content your organization already has *becomes* the training.

This eliminates the development bottleneck that forces L&D teams to ration which topics get proper training and which get a shared drive link and a prayer.

3. Active learning replaces passive consumption

The dominant delivery model - whether in-person, vILT, or self-paced e-learning - is fundamentally passive. The learner watches, listens, or reads. Maybe they answer a few multiple-choice questions at the end.

Cognitive science has established for decades that active recall (retrieving information from memory) and spaced repetition (revisiting material at increasing intervals) produce dramatically better retention than passive exposure. Studies consistently show 200%+ improvement in long-term retention with spaced repetition versus massed practice.

AI-native platforms build these principles into every interaction. The learner doesn't just consume content - they answer questions, work through scenarios, and demonstrate understanding. The AI adapts difficulty based on performance, spending more time where the learner is weak and less where they're strong.

4. Measurement as a byproduct, not a project

In the traditional model, measuring training effectiveness is a separate initiative - surveys, assessments weeks later, correlation studies with performance data. It's expensive and unreliable.

In an AI-native model, every training interaction generates competency data. The platform knows which topics each person has mastered, which they're still struggling with, and how their knowledge changes over time. This data exists as a natural byproduct of the training itself, not as a separate measurement project.

For L&D leaders who have struggled to demonstrate training ROI, this is transformational. The question shifts from "can we prove training works?" to "which specific skills need more attention?"

Section 6: Benchmarks That Actually Matter

The training industry has historically tracked metrics that measure activity rather than impact. Hours delivered, modules completed, satisfaction scores - these tell you that training happened, not that it worked.

McKinsey's research underscores this shift: people development must move "from a support function to a strategic business driver," using "purposeful, actionable data - not just tracking course completions or attendance, but leveraging analytics to understand impact, track growth, and identify skill gaps." Organizations need "predictive modelling and impact visualization" to tie learning directly to business outcomes.

Metrics to retire

Traditional Metric	Why It Misleads
Completion rate	Measures clicks, not learning. A 95% completion rate says nothing about whether anyone retained or can apply the material.
Seat time	Incentivizes longer training, not better training. The best training is often the shortest.
Satisfaction scores	Learners rate easy, entertaining training highest. Challenging training that produces growth often scores lower.
Hours per employee	More hours is not better. Efficient delivery that achieves competency in fewer hours is the goal.

Metrics to adopt

Modern Metric	What It Tells You
Competency rate	Percentage of employees who have demonstrated proficiency in required skills - not just completed a module.
Time-to-proficiency	How quickly new hires reach defined competency thresholds. Directly correlates with business impact.
Knowledge retention at 30/60/90 days	Whether training sticks. Spaced repetition platforms can measure this automatically.

Modern Metric	What It Tells You
Skill gap coverage	What percentage of identified skill gaps have active training addressing them.
Cost per competency gained	Total training cost divided by the number of verified competency improvements. The true efficiency metric.

> **The shift in one sentence**
>
> Move from measuring *training activity* (did it happen?) to measuring *workforce capability* (can they do the thing?). AI-native platforms make this shift practical because competency data is generated continuously, not sampled periodically.

Section 7: Cost Modeling Framework

Use this framework to estimate the potential impact of shifting from instructor-dependent to AI-native delivery for your organization.

Step 1: Calculate your current cost per learning hour

> **Formula**
>
> **Total direct learning expenditure / (Number of employees x Average learning hours per employee) = Cost per learning hour**

Compare your result to the industry benchmarks:

Organization Size	Avg. Spend/Employee	Avg. Hours/Employee	Implied Cost/Hour
Small (1-499)	$1,083	32	$34
Midsize (500-9,999)	$955	22	$43
Large (10,000+)	$1,656	41	$40
Consolidated	**$1,280**	**33**	**$39**

Note: The $103 cost per learning hour referenced throughout this brief is calculated differently by ATD - using total expenditure divided by total learning hours (not averaged across size categories). The per-category figures above show how the math varies by organization size.

Step 2: Identify your shiftable training volume

Not all training can move to AI-native delivery. Estimate the percentage of your learning hours that fall into these categories:

Content Type	AI-Native Suitability	Typical % of Portfolio
Product/service knowledge	High	10-15%
Process and procedure training	High	10-15%
Mandatory/compliance	High	10%

Content Type	AI-Native Suitability	Typical % of Portfolio
Onboarding	High	5-10%
Interpersonal/soft skills	Medium-High	12%
Managerial/supervisory	Medium	20%
Industry-specific technical	Medium	10%
Executive leadership	Low	3-5%

For most organizations, **60-75% of training volume** is suitable for AI-native delivery - either fully or in a blended model where AI handles knowledge transfer and practice while humans facilitate discussion and advanced application.

Step 3: Model the cost reduction

AI-native platforms reduce cost per learning hour through three mechanisms:

1. **Instructor cost elimination** (typically 40-60% of delivery cost for the shiftable portion)
2. **Content development acceleration** (reducing the 100:1 development ratio to near 1:1)
3. **Administrative overhead reduction** (scheduling, tracking, reporting automated)

Conservative estimates suggest a **40-60% reduction in cost per learning hour** for content shifted to AI-native delivery, with the total organizational impact depending on how much volume is shifted.

Section 8: Building the Business Case

The most common mistake in proposing training infrastructure changes is framing them as a *new investment*. The $103 problem gives you a different frame: this is a **cost reduction initiative** that simultaneously improves outcomes.

Three arguments that resonate with leadership

1. "We're spending more per hour than ever and getting less." The cost per learning hour data makes the current trajectory unsustainable clear. This isn't about adding budget - it's about restructuring how the existing budget is spent.

2. "Our 'technology transformation' didn't actually transform costs." The shift to 67% technology-based delivery sounds like modernization, but it moved the same model to a new channel. Actual cost structure transformation requires a different approach.

3. "We can measure what was previously unmeasurable." The on-the-job learning gap (62% commitment, near-zero measurement) is a known frustration for leadership. AI-native platforms close this gap by embedding measurement into training delivery.

4. "The skills gap is accelerating and our current approach can't keep up." McKinsey's 2025 research finds that 83% of leaders believe leadership is key to the skills-based transition, but only 28% of employees feel the strategy is being clearly communicated. Meanwhile, 85% of people believe businesses are obligated to "train or reskill" their employees. The expectation is there. The delivery mechanism is what's missing.

Positioning the pilot

Don't propose a full replacement. Propose a **90-day pilot** with one content area (product knowledge and onboarding are typically the easiest starting points) and one team. Define success metrics in advance:

- Time-to-proficiency for new hires (compared to historical average)
- Knowledge retention at 30 and 60 days (measured by the platform)
- Manager time saved on repetitive training (self-reported)
- Cost per learning hour for the pilot group vs. the organizational average

A well-run pilot generates the data to justify broader adoption. And because AI-native platforms work with existing content (uploaded documents become training immediately), the pilot can launch in days, not months.

The bottom line

The $103 problem isn't going to solve itself. The four forces driving it - flat budgets, instructor-dependent delivery, expensive content growth, and unmeasured informal learning - are structural, not cyclical. Addressing them requires a structural change in how training is delivered. AI-native platforms provide that structural change while reducing cost, improving retention, and finally making workforce competency measurable at scale.

HeyLoopy - Build Expertise. Unleash Potential.

heyloopy.com