FedLearn

WHITE PAPER

Illuminating the "Black Hole" of Online Learning
Measuring the Readiness Return on Training Investment

A Readiness-Economics Framework for U.S. Department of War Training Investment Decisions

Date: November 2026

FedLearn

FOREWORD

Seeing the Unseen: The "Black Hole" of Online Learning

Black holes are among the universe's greatest mysteries. No telescope can reveal what is happening inside. The interior remains hidden from view.

Yet astronomers do not conclude that black holes are unknowable. Instead, they study their effects. By measuring these observable consequences rather than attempting what is currently impossible, scientists can determine a black hole's location, mass and behavior with remarkable precision—even though its interior remains invisible.

Online learning poses a comparable measurement challenge. The process that matters most—the cognitive work of learning—cannot be directly observed. Instead, online learning platforms report only what is easy to capture, for instance, course completions or other simple metrics, which reveal that learning activities occurred, but they do not reveal what happened inside the learner. *They measure participation—not capability.*

If organizations mistake these observable signals for evidence of learning, they risk drawing the wrong conclusions about training effectiveness and return on investment.

The U.S. Department of War (DoW) invests heavily in online training. Yet the factors that decide whether that investment mattered—whether a learner truly engaged with the content, whether the knowledge and skills took hold and whether capability improved—are not directly visible in any online learning platform. Instead, these platforms measure participation: enrollments, minutes in a course, clicks and course completions. *These simple metrics are merely surface indicators of learning. Mistaken for learning itself, they mislead.*

The failure is *not* that these results go unmeasured. *The failure is that they are calculated and then asked a question they cannot answer.* For instance, course completion rate accurately measures whether learners finished a course but tells us little about whether they can apply what they learned on the job. *Completion rate tells us little about actual workforce readiness*.

This white paper makes *three arguments*. First, the metrics driving training investment decisions—course completion, time in training, learner satisfaction and seat license utilization—are poor predictors of workforce capability. This argument is supported by decades of research and reinforced by FedLearn data. Second, learner behavior provides measurable evidence of workforce capability, and the DoW has already mandated the data standards required to assess readiness. Third, once capability can be measured, training investment can be expressed in the currency that matters to a resource sponsor: mission-ready people produced per dollar spent.

The means to measure readiness already exist. What follows is a framework for doing so.

FedLearn

EXECUTIVE SUMMARY

The DoW spends heavily on online learning seat licenses and judges the return on training investment almost entirely by four measurements: the number of personnel who completed a course, time in training, learner satisfaction and the share of purchased licenses used. *None of these metrics indicates whether personnel became more capable.* All four should be retired as training decision metrics.

The evidence is not marginal. Much of it was produced by the Federal government itself, and FedLearn reproduced it directly across thousands of scored enrollments and learners.

- **Course completion provides little insight into workforce capability.** Across 221 online courses studied, the median completion rate was 12.6 percent.[i] In mandatory compliance training, it routinely tops 90 percent.[ii] When the same measure reads 13 percent in one setting and 90 percent in another, and the *difference comes down to whether attendance was ordered*, it cannot support a funding decision.
- **Time provides no evidence of learning, and FedLearn's data confirms this.** Research that investigated what learners actually did online found that time on task accounted for less than 2 percent of the difference in learning results.[iii] Across FedLearn's full course catalog, minutes spent in a course and demonstrated readiness essentially have no relationship at all.[iv]
- **Learner satisfaction is not a stand-in for learning.** Whether individuals liked a course has almost no bearing on what they learned and even less on whether they can use their new knowledge and skills on the job.[v]
- **Seat utilization is a financial signal—not a learning one.** Roughly half of all purchased enterprise licenses go unused.[vi] This matters for managing a budget, but it says nothing about whether the personnel who used their seat licenses became more capable.

The U.S. Government Accountability Office (GAO) documented the learning analytics gap inside the Department. Reviewing U.S. Army and U.S. Marine Corps training, GAO found the services had fielded online courses and automated systems but had no way to measure whether they were effective and recommended results-oriented metrics.[vii]

What has been missing is the means to fill that gap.

The FedLearn Solution

The FedLearn online learning platform directly measures workforce readiness. *Our Z1 behavioral engine—enabled by artificial intelligence (AI)—monitors how a learner works through a course and estimates continuously how much of the material that learner actually understood.* **That estimate becomes the Mission Readiness Knowledge Index (MRKI)—a 0 to 100 readiness score with a mission-ready floor at 61.** *The section of this white paper, "Measuring Readiness: The Z1 Engine and the MRKI," explains how the engine works and what the score means.*

Once readiness can be measured for an individual, it can be measured for a contract. Two metrics are the basis for the FedLearn solution:

- **Readiness yield.** Of the seat licenses a component paid for, what share produced someone who is mission ready? One metric replaces both course completion and license utilization, and it splits cleanly into the two outcomes that matter most: licenses being used and learners reaching the mission-ready floor. That split turns a score into a diagnosis because the two outcomes have different owners and require different solutions.
- **Cost per mission-ready learner.** What the component actually paid for each mission-ready individual. Contract size cancels out of the equation—allowing a 500-seat pilot and a 50,000-seat enterprise agreement to be compared directly—without normalization, indexing or weighting.

Both are defined with worked examples in the "FedLearn Readiness Economics Model" section of this white paper. Together, they enable leaders to move beyond asking, "How many personnel were trained?"—and instead ask two more meaningful questions:

- **How many personnel are ready?**
- **What did it cost to train each individual?**

WHAT TODAY'S LEARNING METRICS REALLY TELL US

The common learning metrics—course completion, time in training, learner satisfaction and seat license utilization—have been studied against actual learning outcomes, and in every case the relationship is weak or nonexistent.

Course Completion

A study of 221 online courses found completion rates ranging from less than 1 to 52 percent with a median of 12.6 percent.[viii] Most of that spread comes from how completion is defined—*not* from how good the instruction was. A separate study of four courses revealed that the same course can report a completion rate of 30.0 percent, 43.1 percent or 48.1 percent depending only on which set of learners gets counted: everyone who enrolled, the active learners or those who said they intended to finish the course. *The researchers concluded that completion measured from initial enrollment is "a poor metric."*[ix]

When training is mandatory, the measurement fails in the opposite direction. Learning platform data across required compliance campaigns indicates completion rates of 97 to 99 percent and public-sector mandatory training in the 90 to 94 percent range.[x] When attendance is compulsory, nearly all personnel complete it. *The metric is reporting the existence of a mandate—not the effect of the instruction.*

Time in Training

Researchers who observed what learners actually did online—rather than asking the learners to report the information themselves—found that time on task accounted for less than 2 percent of the difference in learning scores and less than 1 percent on standardized tests. In eight of the 19 classrooms studied, more time on task created poorer learning outcomes—not better ones.[xi]

The sharpest criticism comes from the organization that invented the standard way of measuring student learning by tracking seat time: the "Carnegie Unit." The Carnegie Foundation for the Advancement of Teaching said the Carnegie Unit was "never intended to function as a measure of what students learned." It standardized only "exposure to subject material."[xii]

FedLearn Data Confirms the Research

Across thousands of scored enrollments, time spent in a course and demonstrated readiness have no measurable relationship at all. Seat time is also unreliable on its face: one record shows 429,833 minutes—298 days—inside a single course and 1.3 percent of records exceed 24 hours.

Learner Satisfaction

How much learners like a course turns out to have almost no bearing on what they learned and whether they use their new knowledge and skills at work. The most thorough review of the evidence, published as an Air Force Research Laboratory (AFRL) technical report and still available through the Defense Technical Information Center, found the relationship to be close to zero across every study it examined.[xiii]

Seat License Utilization

Seat license utilization is a real financial signal and should be tracked as a budget measure—not as a learning measure. Industry studies found that roughly half of purchased licenses are ever actually used and about half of all installed software is never opened. The average large organization also pays for roughly 15 overlapping online training applications.[xiv]

The distinction matters. Utilization tells a comptroller that money was spent without matching learner activity and cannot say whether the activity produced any capability. A program in which 95 percent of licenses are used, yet no one learns anything, would still receive a high utilization score—even though it provides no readiness diagnostic value.

Common Flaw

The four traditional metrics measure participation in training. None measure competency from training. A smart funding decision requires the latter.

THE MEASUREMENT GAP IS ALREADY FEDERAL POLICY

The measurement gap is a documented GAO finding and an unmet requirement of DoW policy.

Source	Finding
GAO-11-673 (2011)[xv]	Reviewing Army and Marine Corps training, GAO found the services had fielded online courses and automated systems but had no way to measure whether they worked. Recommended the DoW develop results-oriented performance metrics. The Department agreed.*
GAO-04-546G (2004)[xvi]	The standing federal framework for planning, designing, delivering and evaluating training. GAO deliberately declined to prescribe a specific set of measures and called the guide a starting point—leaving the actual measure to be supplied.
GAO-12-878 (2012)[xvii]	Found that chief human capital officers were not using practices that support cost-effective training investment decisions, including evaluating the benefits of the training they purchased.
DoD Manual 8140.03*[xviii]	Makes the work role the organizing principle for defense cyber workforce qualification, tied to the knowledge, skills, abilities and tasks (KSATs) the specific role requires, with training-based foundational qualification set at 70 percent coverage of core content

**Released prior to the Department's name change*

Practice has not caught up to policy. A survey of 340 organizations found that fewer than a quarter measure whether training achieved what it was meant to achieve.[xix] Fewer than half of learning professionals ever check whether learners applied the training on the job, and fewer than one in five look at whether it changed organizational results.[xx] The research literature has the same blind spot: studies measuring what learners know outnumber studies measuring organizational results by roughly nine to one.[xxi]

The Department has stated the requirement, mandated the data standard and named the vocabulary: measures of effectiveness and performance and key performance parameters.

What has been missing is an effective and legitimate way to satisfy that need. The remainder of this white paper proposes one.

MEASURING READINESS: THE Z1 ENGINE AND THE MRKI

Measuring workforce readiness at scale has long been the challenge that leads organizations to rely on proxy metrics instead. **FedLearn solves that challenge in two steps: an AI-powered behavioral engine analyzes how learners actually perform online and converts these observations into a readiness score that leaders can use to make informed decisions.**

The Z1 Behavioral Engine

The Z1 behavioral engine was developed from peer-reviewed research at the Princeton University EDGE Lab. **Z1 monitors how a learner works through online course material**—what the individual focuses on, in what order, where the learner slows down, where the learner doubles back and what the learner gets right and wrong—**and estimates from that behavior how much of the content the learner genuinely absorbed.**

Three attributes make the Z1 estimate useful rather than merely interesting:

- **Runs continuously.** The estimate updates are ongoing, so a learner who is falling behind can be helped in week two rather than written up as a failure in week six.
- **Not tied to FedLearn content.** Z1 analyzes a variety of online content (e.g., videos, e-learning packages, documents and podcasts), so it can measure learning wherever it happens—not just in FedLearn courses.
- **Reports only when it is confident.** Z1 works to a confidence level of 90 percent or higher. When the evidence is insufficient, it says so rather than guessing.

What the Z1 Score Communicates

Z1 is a score from 0 to 100 that measures the level of learning occurring—in real time—while an individual is immersed in an online course. It is not an assessment grade or a course completion percentage. A learner who rushes through every screen in a course and retains little of the material would receive a low score. Conversely, a learner who completes only two-thirds of the course but genuinely absorbs the material would receive a higher score.

From Z1 to MRKI

MRKI puts the Z1 score on a workforce readiness scale. It is the Z1 score alone and is never blended with other metrics, like learner engagement, course completion or time spent, because the Z1 behavioral engine already uses them as inputs.

MRKI places learners into five bands. *Each band answers one question: Is this person mission ready and able to do the job?*

The MRKI bands are described in the table on the next page of this white paper.

Band	MRKI Score	Meaning for Readiness
At Risk	0–40	Not developing readiness. Will not get there without intervention
Developing	41–60	Building toward readiness but not there yet. Cannot be counted as ready
Competent	**61–80**	**Mission-ready floor. Can utilize the knowledge or skills learned on the job**
Proficient	81–90	Ready and consistent. Performs reliably without support
Expert	91–100	Ready to teach. Can instruct or lead others in the subject

Mission ready is achieving an MRKI score of 61 or higher. *Every readiness metric in this white paper evaluates learners against that floor*. A readiness measure must have a threshold because an average obscures the only thing a commander needs from a training report: a headcount of the personnel ready to do the job.

Why Measure Behavior Instead of a Test?

There are three reasons why measuring learner behavior is more valuable than a graded assessment:

- **Continuous—not final.** Readiness is estimated throughout the course, so a struggling learner can be helped (via "personal remediation" or providing snippets of information about the challenging content for the learner to review; *a FedLearn platform exclusive*) while in the course rather than written up as a failure later.
- **Does not rely on learners reporting on themselves.** Research shows that self-reported motivation is a much weaker predictor of performance in online learning rather than classroom-based training.[xxii]
- **Harder to game than completion.** Clicking through content advances a completion counter—not a readiness estimate built from behavior.

FEDLEARN READINESS ECONOMICS MODEL

Once readiness can be measured, the investment question can be answered in the Department's own currency.

Readiness Yield

Readiness Yield = Learners Reaching MRKI ≥ 61 ÷ Seat Licenses Purchased

The component to the right of the formula is the critical design choice: seat licenses purchased—not learners enrolled or who completed a course. That one decision lets a single measure do the work of both completion and utilization. A seat license that goes unused obviously does not produce a mission-ready learner and lowers yield. A learner who uses a license but stops partway does not become a mission-ready learner and lowers yield—unless that individual demonstrates readiness despite not finishing a course, which the model credits. And a learner who completes a course with an MRKI score of 45 does not become a mission-ready learner and lowers yield.

Completion Is Not on the Readiness Path

Because MRKI is *estimated continuously*, a learner can reach mission-ready status without finishing a course. Requiring completion would squander those results—and in FedLearn's platform data, that would mean disposing 21.8 percent of all mission-ready course results. Completion is a diagnostic that helps explain results—not a gate.

Cost per Mission-Ready Learner

There are two ways to calculate this measure, and they always give the same answer.

The first is straightforward accounting. Divide the total contract cost by the number of mission-ready personnel produced. A comptroller can calculate the measure directly from an invoice and a readiness report.

Cost per Mission-Ready Learner = Total License Investment ÷ Learners Reaching MRKI
≥ 61
which reduces exactly to
Seat License Price ÷ Readiness Yield

The second formulation eliminates contract size from the equation altogether: seat license price divided by readiness yield. This works because the number of purchased seat licenses appears in both the numerator and denominator of the original formulation and therefore cancels out.

Consider a contract for 1,000 seat licenses at $61 each, for a total investment of $61,000. If 50 percent of those licenses produce mission-ready learners, the cost is $122 per mission-ready individual. Now consider a contract for 10,000 seat licenses at the same price and the same readiness yield. The total investment increases to $610,000 and the number of mission-ready learners rises to 5,000—but the cost remains $122 per mission-ready individual. *The contract is ten times larger, yet the measure remains unchanged because it is independent of contract size.*

That is what makes the second version useful to a buyer. *A 500-seat pilot and a 50,000-seat enterprise agreement can be set side by side with no adjusting for size, no index and no weighting scheme to argue about.* It also settles what a seat license really costs because the quoted price only holds if every license produces a mission-ready individual. *The table below shows what a $61 license actually costs at each level of readiness yield.*

Readiness Yield	Cost at $61	Effective Multiple
90%	$67.78	1.1×
75%	$81.33	1.3×
60%	$101.67	1.7×
50%	**$122.00**	**2.0×**
40%	$152.50	2.5×
20%	$305.00	5.0×
10%	$610.00	10.0×

The table demonstrates that a FedLearn seat license priced at $61 for Federal customers only costs $61 if every license produces a mission-ready learner. If only 40 percent of them do, each ready individual costs $152.50 because the other three-fifths produced no competency at all. Raise readiness yield to 80 percent and the same license costs $76.25 per mission-ready learner. Same price, same contract and half the cost per ready operator. *This data is what a DoW leader must have to understand the true cost of readiness.*

What Counts as One Mission-Ready Individual

One mission-ready learner is one individual who cleared the MRKI floor of 61 in at least one course. This is the definition used throughout this white paper, and it is exactly what readiness yield measures, so the two never disagree.[xxiii]

Saying so plainly matters because the same investment can be made to appear more economical by counting course completions rather than personnel. A single learner may take and pass several courses. FedLearn counts learners because a leader needs a headcount—not just a course count. The number of courses it takes to make someone mission-ready for a specific job is a different question.

Appendix A of this white paper applies these measures to two real contracts held by the same Federal customer at effectively the same price per seat and explains what the comparison reveals when one platform measures readiness and the other does not.

Readiness for a Person Requires a Work-Role Mapping

DoDM 8140.03 defines qualification against the tasks and KSATs a work role requires, with training-based foundational qualification set at 70 percent coverage of core content. That is a ready-made, policy-grounded definition of role readiness, so FedLearn does not need to invent one. Mapping courses to work-role KSATs is exactly what the Total Learning Architecture's Reusable Competency Definitions are meant to carry. Until that mapping exists for a given customer, one mission-ready learner means one person who cleared the floor in at least one course, and it should be described that way.

WORKED EXAMPLE: READING A CONTRACT

Take the example of an enterprise agreement for 2,000 seat licenses at the $61 Federal price. Total investment: $122,000. The table below explains what that contract actually produced.

Stage	Rate	Result
Seat licenses purchased	—	2,000
Stage 1: Activation (licenses opened)	62%	1,240 learners opened one
Stage 2: Readiness (reaching MRKI ≥ 61)	84.6%	1,049 mission-ready learners
Readiness yield	**52.4%**	**1,049 of 2,000 licenses produced a ready person**
Cost per mission-ready learner	**$116.30**	**$122,000 ÷ 1,049, or $61 ÷ 0.524**

As the table shows, a total of 2,000 seat licenses were purchased, and about 62 percent of the licenses were ever opened. This leaves 1,240 individuals who started a course, and of those personnel, 84.6 percent attained the mission-ready threshold, or a total of 1,049 mission-ready learners.

In sum, the contract purchased 2,000 seat licenses and produced 1,049 ready personnel—representing a readiness yield of 52.4 percent and meaning that each ready learner actually cost $116.30 rather than the $61 quoted in the contract.

What Happened to the Other Seat Licenses Purchased?

A total of 951 seat licenses were purchased and produced no mission-ready personnel for two reasons.

- 760 seat licenses were never used—never opened.
- 191 licenses were opened by a learner who did not reach the mission-ready floor.

Both outcomes are typically the responsibility of the purchasing organization and often stem from the same underlying cause. Personnel are far more likely to engage seriously with training when organizational expectations make its importance clear: seat licenses are assigned to named individuals rather than maintained in a shared pool, courses are tied to duty requirements or billet qualifications and supervisors monitor learner progress. Where these conditions are absent, some personnel never begin the training, while others participate without fully engaging with the material. This conclusion is not speculative.

Decades of research consistently show that persistence, engagement and knowledge transfer are driven primarily by organizational support and perceived job relevance—not by learner motivation alone.

One exception does exist. If learners are actively engaging with a course yet consistently fail to reach the mission-ready threshold, the problem likely lies with the course itself rather than the organization. Because MRKI is reported at the course level, ineffective courses can be identified, revised and reevaluated to improve learning outcomes.

Improving Readiness Without Increasing Costs

The preceding table illustrates the **central argument of this white paper: the true return on training investment becomes visible only when workforce readiness can be measured directly.** The quoted contract price reflects the cost of a seat license. The readiness cost reflects the cost of producing one mission-ready individual. The difference between the two is determined not by contract negotiations, but by execution.

That distinction is encouraging because execution is the least expensive variable to improve. In this example, seat license activation is 62 percent. Increasing activation to 85 percent requires no additional funding and no new course content. It does require better execution: assigning seat licenses to named personnel, linking training to duty requirements and having supervisors monitor learner progress.

Readiness is already 84.6 percent—consistent with the levels FedLearn has observed across its platform. With relatively little room for improvement, further gains will most likely require course redesign, which will demand additional funding and time.

The greatest opportunity, therefore, lies in improving seat license activation. The table below holds the total contract investment of $122,000 constant. No additional funding is assumed and no new course development is required. It shows only the reduction in cost per mission-ready individual that can be achieved through better execution.

Intervention	New Yield	Learners Gained	New Cost per Learner
Activation 62% → 85%	**71.9%**	**+389 learners**	**$84.84**
Readiness 84.6% → 92%	57.0%	+92 learners	$106.92
Both improved	**78.2%**	**+515 learners**	**$78.01**

The same $122,000 produces either 1,049 mission-ready learners or 1,564—depending entirely on execution. That is a 33 percent cut in cost per mission-ready learner with no change in price, scope or contract value.

DIAGNOSTIC: READING SCORES AND ACTING ON THEM

Readiness yield pinpoints a problem and does not by itself explain what is causing the challenge. Because yield is two sequential rates multiplied together, each shortfall can be isolated, assigned to the program office that owns the readiness yield issue and matched to the solution supported by evidence.

Readiness Yield = Number of Mission-Ready Personnel ÷ Total Seat Licenses Purchased

Course completion in this section of the white paper is not a gate but a diagnostic that helps explain which of the two stages (described below and on the next two pages of this white paper) is failing and why.

Stage 1: Activation—Are the Seat Licenses Being Used?

Definition: Share of purchased seat licenses that learners ever use

Owner: Command, program office or training officer responsible for distributing the licenses. Represents a distribution and accountability problem—not a learning problem. This is why improving a course will never address the challenge of unused licenses.

The table on the next page of this white paper identifies four levels of seat license activation and recommends actions to improve utilization and workforce readiness.

Band	Activation	Recommended Actions
Critical	Below 25%	▪ Stop buying seat licenses until distribution issue is addressed. More licenses purchased at this rate produce almost no readiness. ▪ Check whether licenses reach named individuals or remain unassigned in a pool. Unassigned pools are the most common cause of activation collapse. ▪ Tie the training to a duty requirement, billet qualification or evaluation input. Whether adults stick with online learning depends on organizational support and whether the training feels relevant—not on the individuals themselves.[xxiv] ▪ Brief supervisors directly. Supervisor support is one of the strongest predictors of whether training gets used and is the cheapest lever available.[xxv]
Low	25–49%	▪ Identify who has not used a seat license by organization and role ▪ Activation failure is usually concentrated—not spread evenly. ▪ Institute a 30-day activation window with command-level reporting on assigned but unstarted licenses ▪ Reallocate unused licenses to other personnel with waiting demand rather than allowing them to expire
Moderate	50–74%	▪ This band already beats commercial norms. Typical enterprise software utilization sits near 49 percent.[xxvi] ▪ Target the residual through scheduled cohort starts rather than open enrollment ▪ Shift optimization effort to stage two where the remaining headroom is larger
Strong	75%+	▪ Maintain current distribution practice and document it as a repeatable process for other components ▪ Consider expanding the seat license count. At this activation rate, marginal licenses convert efficiently into readiness.

Stage 2: Readiness—Are Learners Reaching Mission-Ready Status?

Definition: Share of learners who used a seat license and attained an MRKI score of 61 or higher

Owner: Whoever designs the content and the assessments. This is the stage that no conventional learning platform can measure and where the difference between attendance and readiness is decided. The FedLearn platform benchmark is 84.6 percent.[xxvii]

The table on the next page of this white paper details sets of recommended actions that are used to address the MRKI scores generated from FedLearn's AI.

Band	MRKI ≥ 61	Recommended Actions
Critical	Below 61	▪ Treat this as a content problem before treating it as a learner problem. When most participants fail to reach readiness, the course is not teaching the skill being measured. ▪ Add retrieval practice throughout. Having learners recall what they learned outperforms having them re-review it—consistently and by a wide margin[xxviii] ▪ Insert worked examples that show reasoning step by step rather than stating conclusions ▪ Add adaptive remediation. Intelligent tutoring systems produce large, well-documented gains over standard instruction.[xxix]
Low	61–80	▪ Introduce spaced review. Spaced practice outperforms cramming it together, and the best gap widens the longer the material has to be retained[xxx] ▪ Look at MRKI by individual competency rather than in total to determine the specific skills that are not being developed ▪ Add scenario-based practice at increasing difficulty—consistent with AFRL guidance on building proficiency faster[xxxi]
At Benchmark	81–90	▪ Target the "Developing" band (MRKI 41–60) specifically. These learners are closest to the floor require the least expensive investment to cross it. ▪ Add follow-up microlearning reinforcement after course completion rather than extending the course itself
Strong	91+	▪ The instruction is performing above the platform benchmark. Shift the focus to retention (see "Sustainment: Readiness Decays") ▪ Consider moving the reporting threshold up to "Proficient" (81–90) for roles that require teaching or leading others

Overall: Cost per Mission-Ready Learner

The two-stage diagnostics describe performance. Cost per mission-ready learner sums it up for the funding decision.

The levels described in the table on the next page of this white paper are expressed as a multiple of seat license price, so they apply at any price and any contract size.

Assessment	Cost vs. Seat License Price	Recommended Posture
Efficient	Under 1.5× (yield above 67%)	Renew and expand. Each additional seat license converts efficiently. Document the distribution and teaching practices producing this result
Acceptable	1.5× to 2.5× (yield 40–67%)	Renew at current scale. Fix the weaker stage before buying more licenses because added licenses inherit the current yield
Inefficient	2.5× to 5× (yield 20–40%)	Hold scale. Do not expand until yield improves. Diagnose using the two stages above and remeasure after a defined period of time
Critical	Above 5× (yield under 20%)	Suspend further purchase. The main cause is nearly always activation and it can be fixed without renegotiating price or changing content.

Sustainment: Readiness Decays

A learner who reaches a particular MRKI score does not remain at that result forever. Skill decay is well documented in defense training research, and AFRL guidance recommends refresher training aimed specifically at the skills that fade fastest—rather than blanket annual retraining.[xxxii]

Measuring MRKI again at 90 and 180 days transforms a one-time score into an ongoing readiness picture and relays which knowledge and skills need a refresh and which hold. For a resource sponsor, this changes the question from "How many personnel did we train last year?" to "How many personnel are mission ready today?" *The second question is the one that should have been asked all along.*

CONCLUSION: MAKING THE INVISIBLE VISIBLE

Astrophysicists did not solve the black hole problem by trying to observe the unobservable. They accepted that the interior would remain hidden and developed instruments capable of measuring its observable effects with remarkable precision. Workforce readiness presents the same measurement challenge. *The learning process itself cannot be observed directly, but its effects can.*

GAO documented the DoW's measurement gap. Department policy established the data standards needed to close it. The FedLearn platform provides the evidence: *across thousands of scored enrollments, time spent in a course has no measurable relationship to workforce readiness.*

What changes when readiness is measured from learner behavior is not that training becomes visible—workforce readiness becomes measurable. A seat license is no longer simply a purchased expense. It becomes a measurable investment that produces a specific number of mission-ready personnel at a quantifiable cost per individual.

That outcome can now be measured. It is calculated from data the DoW has already directed its components to collect—expressed in the same economic terms resource sponsors use to evaluate investments and designed to identify where corrective action is needed when results fall short. *Leaders no longer have to infer whether training worked from activity-based metrics.* **They can quantify what training produced—mission-ready personnel—identify where improvement is needed and understand the readiness return on every training dollar invested.**

FedLearn

APPENDIX A

Two Contracts at the Same Price

The figures in this appendix are real. They describe two online learning contracts held by the same DoW organization and cover the same period and workforce. One contract was with FedLearn. The other was with an unnamed commercial learning platform.

The negotiated price per seat license on the two contracts differed by 11 cents—only the total number of seat licenses and contract value differed. Because the customer, the personnel and the seat license price are essentially identical, the difference in what each contract produced came down to execution and measurement—*not* what was purchased.

Contract Inputs

	FedLearn	Unnamed Vendor
Seat licenses purchased	600	16,000
Contract value	$31,752	$845,000
Price per license	**$52.92**	**$52.81**

As the table displays, there is a 0.2 percent difference in seat license price between the two vendors. *Whatever separates the two investments is not the negotiated price.*

Funnel Inputs

Readiness yield is two rates multiplied together: the share of purchased seat licenses that learners use and the share of those individuals who reach the mission-ready threshold. Both contracts can report the first metric. *Only one—FedLearn—can report the second.*

	FedLearn	Unnamed Vendor
Learners who used a seat license	94	720
Activation rate	**15.7%**	**4.5%**
Learners reaching MRKI ≥ 61	79	*Not measured*
Reached mission-ready status	**84.0%**	*Not measured*

The FedLearn activation rate is *more than three times* that of the unnamed vendor's rate. Among the learners who participated in the courses, 84.0 percent reached the mission-ready floor.

Readiness Economics

	FedLearn	Unnamed Vendor
Cost per learner who used a seat license	$337.79	$1,173.61
Readiness yield	**13.2%**	*Not computable*
Cost per mission-ready learner	**$401.92**	*Not computable*
Unused licenses	506	15,280
Spend on unused licenses	$26,778	$806,975

The blank cells in the table are telling.
The unnamed vendor's contract cannot be evaluated on workforce readiness—because the measurement does not exist in its learning platform. The vendor's contract value of $845,000—against FedLearn's $31,752 for the same workforce—can only be reported in terms of the personnel who logged in.

When a Learning Platform Does Not Measure Readiness

A comparison like this example invites the objection that the unnamed vendor's readiness results are simply unknown and might be excellent. That objection can be answered without knowing the results.

Assume—generously and with no evidence at all—that every single learner who used a seat license on the unnamed vendor's platform reached the mission-ready threshold. That assumption produces the lowest cost per mission-ready learner the contract could possibly achieve, as displayed in the table below.

Assumed Share Reaching Mission-Ready Status	Readiness Yield	Cost per Mission-Ready Learner
100% (best possible case)	**4.5%**	**$1,173.61**
85%	3.8%	$1,380.72
70%	3.1%	$1,676.59
50%	2.2%	$2,347.22

Even if every single learner reached mission-ready status, the unnamed vendor's contract would cost $1,173.61 per mission-ready learner—*2.9 times* the FedLearn price of $401.92.

That is a floor—not an estimate—and supports a stronger conclusion than a straight comparison would.

The contract cannot match FedLearn's cost effectiveness—not because the course content is inferior, but because *low seat license activation limits readiness yield before course quality can influence the outcome*.

What This Case Demonstrates

- **Price per seat is not price per outcome.** Two contracts with almost identical seat license prices delivered results that are not remotely comparable, and the difference is invisible in each procurement.
- **Activation caps everything downstream.** Excellent content cannot make up for seat licenses that are never used. Readiness yield is the number of mission-ready learners divided by the total number of licenses, and neither one can compensate for a collapse in the other.
- **Unmeasured is not the same as adequate.** The absence of a readiness figure is not a neutral gap in reporting. It means the unnamed vendor's contract—27 times larger in value—cannot be judged on an outcome it was purchased to produce.
- **The measure produces an instruction—not just a score.** It identifies whether activation or readiness failed, which office owns the issue and what addressing it is worth in dollars per mission-ready learner.

Appendix figures are drawn from FedLearn platform records for August 2025 through July 2026 and from contract values supplied by the customer. The unnamed vendor's seat count, contract value and utilization rate were also supplied by the customer. No readiness data exists for that platform. Cost per mission-ready learner is calculated at the learner level on both sides, so the two figures are directly comparable.

---

[i] Jordan, K. (2015). Massive open online course completion rates revisited: Assessment, length and attrition. International Review of Research in Open and Distributed Learning, 16(3), 341–358.

[ii] KnowBe4. (2023, January 26). What is a good completion percentage for security and compliance training?; Milton Keynes University Hospital NHS Foundation Trust. (2022, September 20). Freedom of Information Act 2000 request response: Statutory and mandatory training (Ref. 7667).

[iii] Godwin, K. E., Seltman, H., Almeda, M., Davis Skerbetz, M., Kai, S., Baker, R. S., & Fisher, A. V. (2021). The elusive relationship between time on-task and learning: Not simply an issue of measurement. Educational Psychology, 41(4).

[iv] FedLearn platform analysis. Figures are calculated from all FedLearn learning activity between January 1, 2025, and July 31, 2026, covering thousands of enrollments and learners across dozens of organizations. No individual learner identifiers were used; all figures are course, subject-area and organization totals.

[v] Alliger, G. M., Tannenbaum, S. I., Bennett, W., Jr., Traver, H., & Shotland, A. (1997). A meta-analysis of the relations among training criteria. Personnel Psychology, 50(2), 341–358. Also issued as AFRL technical report ADA368508, available through the Defense Technical Information Center.

[vi] Zylo. (2024). 2024 SaaS Management Index; Nexthink. (2023, February 6). Half of software licenses goes unused by employees; Productiv. (2021, September 15). Less than half of company SaaS applications are regularly used by employees. All three are vendor telemetry studies and none are specific to training licenses.

[vii] U.S. Government Accountability Office. (2011). Army and Marine Corps training: Metrics needed to assess initiatives on training management skills (GAO-11-673).

[viii] See note 1.

[ix] Celik, B., & Cagiltay, K. (2024). Uncovering MOOC completion: A comparative study of completion rates from different perspectives. Open Praxis, 16(3), 445–456.

[x] See note 2.

[xi] See note 3.

[xii] Silva, E., White, T., & Toch, T. (2015). The Carnegie Unit: A century-old standard in a changing education landscape. Carnegie Foundation for the Advancement of Teaching.

[xiii] See note 5.

[xiv] See note 6.

[xv] See note 7.

[xvi] U.S. Government Accountability Office. (2004). Human capital: A guide for assessing strategic training and development efforts in the federal government (GAO-04-546G).

[xvii] U.S. Government Accountability Office. (2012). Federal training investments: Office of Personnel Management and agencies can do more to ensure cost-effective decisions (GAO-12-878).

[xviii] U.S. Department of Defense. (2023). Cyberspace workforce qualification and management program (DoD Manual 8140.03).

[xix] Association for Talent Development. (2026). Highlights from ATD's 2026 State of the Industry report.

[xx] Kennedy, P. E., Chyung, S. Y., Winiecki, D. J., & Brinkerhoff, R. O. (2014). Training professionals' usage and understanding of Kirkpatrick's Level 3 and Level 4 evaluations. International Journal of Training and Development, 18(1), 1–21.

[xxi] Arthur, W., Jr., Bennett, W., Jr., Edens, P. S., & Bell, S. T. (2003). Effectiveness of training in organizations: A meta-analysis of design and evaluation features. Journal of Applied Psychology, 88(2), 234–245.

[xxii] Walker, A., Aguiar, N. R., Soicher, R. N., Kuo, Y.-C., & Resig, J. (2024). Exploring the relationship between motivation and academic performance among online and blended learners: A meta-analytic review. Online Learning, 28(4), 76–116.

[xxiii] Why the definition changes the number. On the FedLearn platform, 820 learners held a seat license at the $61 federal price, a $50,020 investment. Dividing that one investment three ways: $12.99 for each of the 3,852 courses passed at MRKI ≥ 61; $67.78 for each of the 738 learners who passed at least one course; $153.91 for each of the 325 who passed at least three. The same people account for all three figures. This white paper uses $67.78, the learner-level figure, because that is what readiness yield counts.

[xxiv] Park, J.-H., & Choi, H. J. (2009). Factors influencing adult learners' decision to drop out or persist in online learning. Educational Technology & Society, 12(4), 207–217.

[xxv] Blume, B. D., Ford, J. K., Baldwin, T. T., & Huang, J. L. (2010). Transfer of training: A meta-analytic review. Journal of Management, 36(4), 1065–1105.

[xxvi] See note 6.

[xxvii] See note 4.

[xxviii] Rowland, C. A. (2014). The effect of testing versus restudy on retention: A meta-analytic review of the testing effect. Psychological Bulletin, 140(6), 1432–1463.

[xxix] Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review. Review of Educational Research, 86(1), 42–78.

[xxx] Cepeda, N. J., Pashler, H., Vul, E., Wixted, J. T., & Rohrer, D. (2006). Distributed practice in verbal recall tasks: A review and quantitative synthesis. Psychological Bulletin, 132(3), 354–380.

[xxxi] Hoffman, R. R., Feltovich, P., Fiore, S., Klein, G., Missildine, W., & DiBello, L. (2010). Accelerated proficiency and facilitated retention (AFRL-RH-AZ-TR-2011-0001). AFRL.

[xxxii] See note 30.