



Most Graduate Programs Have No Playbook. Let's Fix That.

PLAY: [BRIDGE] The Bridge Matrix

Do You Have One Tool That Shows Where Your Learners, Your Curriculum, and Your Employers Actually Agree?

V1.0 Published July 22, 2026

By Rodney G. Alsup, D.B.A., CPA, CITP

Professor Emeritus

Part of the Graduate Program Intake Playbook — Part 0: Program Strategy & Foundation

Picture this: you're sitting across from your dean, and she asks a simple question — "How do we know our curriculum is actually preparing students for what employers want?"

You have an answer. It's a reasonable answer. It's built on instinct, a few advisory board comments, and last year's accreditation self-study. But you don't have anything you can put on the table and point to. I've had a version of this conversation more times than I can count, and it never gets easier to answer with a shrug dressed up as confidence.

Here's the problem underneath that moment: your program is running three separate conversations that never actually touch each other. One is about what your learners need to develop. One is about what your curriculum is designed to teach. One is about what employers require the day your graduates walk in the door. Each conversation has its own vocabulary, its own owner, and its own evidence. None of them talk to each other in a language everyone can read at the same time.

ISSUE

There is no shared framework connecting learner needs, curriculum design, faculty development, and employer expectations. Each stakeholder group describes "mastery" differently, so a program can look aligned in three separate reports and still be misaligned in practice.

WHY THIS MATTERS

Left unaddressed, the gap doesn't show up until it's expensive:

- Graduates test well on coursework but freeze in interviews that expect application, not recall.

- Faculty invest heavily in topics employers stopped asking for two years ago.
- Employer partners quietly lose confidence because the program can't show, in a single page, how a learning outcome maps to a job requirement.
- Accreditation self-studies scramble to reverse-engineer alignment evidence that should have existed all along.

BLOOM'S TAXONOMY, IN THIRTY SECONDS

You're going to hand this matrix to faculty, intake staff, and employer partners — most of whom have never heard the phrase "Bloom's Taxonomy" and have no reason to have. Here's the short version, so nobody has to take that on faith.

- Benjamin Bloom and a committee of collaborators built the original taxonomy in 1956 to classify educational objectives by complexity, from simple recall up through original creative work.
- In 2001, Lorin Anderson and David Krathwohl — Krathwohl was one of Bloom's original co-authors — revised it: they turned the static nouns into action verbs and swapped the order of the top two levels, putting Create above Evaluate.
- Nearly seventy years on, it's still the default framework most curriculum committees, teaching centers, and accreditors reach for when they write or review learning objectives. If your program has ever gone through an AACSB or regional accreditation self-study, you've already used it, whether anyone called it by name or not.

The six levels, from simplest to most complex: Remember, Understand, Apply, Analyze, Evaluate, Create. That's the entire vocabulary you need to read the matrix below. Nobody at the table needs a semester of instructional design theory — they need this paragraph.

WHY BLOOM'S TAXONOMY

That familiarity is exactly why it's useful here. Used as the spine of a shared matrix, Bloom's turns a vague claim like "students should be familiar with financial analysis" into a specific, comparable target: is "familiar" Remember-level recognition, or Analyze-level application under pressure? That distinction is the whole ballgame, and right now most programs never make it explicit.

WHAT TO DO: BUILD THE BRIDGE MATRIX

The matrix has two axes:

- Down the left side: the knowledge, skills, and capabilities (KSCs) your program is actually responsible for developing.
- Across the top: Bloom's six taxonomy levels.

Inside that grid, plot three overlays for every KSC:

1. Entry level (L) — where your learners actually start, from intake assessment or self-report.

2. Curriculum target (C) — the level your curriculum is designed to bring them to, pulled straight from your syllabi and learning objectives, not your marketing copy.
3. Employer expectation (E) — the level employers expect on day one, pulled from job postings, advisory board input, and alumni surveys.

A quick note on all three overlays: a survey checkbox will get you a number, but it won't always get you the truth. The richest data — especially for entry level (L) — comes from an actual conversation with the learner; how to structure that conversation is its own play, worth reading before you build your first cohort matrix. Curriculum target (C) benefits from the same treatment with faculty, since a syllabus rarely captures what an instructor actually emphasizes in the room. Employer expectation (E) is the one overlay where a well-built survey can be a legitimate shortcut rather than a compromise — employers are harder to get in a room repeatedly, and a tight, well-worded survey across a broader panel of employer partners often beats a handful of deep conversations you don't have time to schedule. Match the method to what each overlay actually needs, not to whichever one is easiest to run.

Read that table the way you'd read an X-ray. "Financial Statement Literacy" shows learners entering at Apply, the curriculum targeting Analyze, and employers expecting Evaluate — a design gap between what you teach and what the market wants. "Ethical Reasoning Under Ambiguity" shows a much wider spread: learners enter at Understand, the curriculum stops at Apply, employers expect Create. That's not a rounding error. That's a program telling on itself.

Here is a simplified example, showing where each stakeholder's mark lands on the same scale:

Knowledge / Skill / Capability	Remember	Understand	Apply	Analyze	Evaluate	Create
Financial Statement Literacy			L	C	E	
Cross-Functional Communication		L		C	E	
Data-Driven Decision Making			L		C	E
Ethical Reasoning Under Ambiguity		L	C			E

L = learner entry level C = curriculum target level E = employer expectation level

HOW TO BUILD IT

One caution before you start pulling KSCs together: the published literature and most program websites offer a nearly endless supply of competency lists — hundreds of knowledge, skills, and capabilities, most of them generic enough to apply to any program anywhere. That volume is a trap, not a resource; a list built from it will look rigorous and say nothing true about your program specifically. Resist starting there.

Instead, anchor your pilot in a current intake cohort — the group you already have in the building right now. Run the KSC list and the three overlays against real learners, real syllabi, and real employer conversations tied to that specific group, rather than a hypothetical. This does two things at once: it proves the gaps already exist using your own data instead of a framework, and it gives that cohort an early, substantive reason to feel like the program invested in them specifically — the same relationship you're ultimately hoping shows up again at graduation and well beyond it.

1. Build the KSC list first. This is the hard part, and it's tempting to skip straight to the taxonomy. Pull candidates from three sources: your curriculum's stated learning objectives, a scan of employer job postings and advisory board notes, and your accreditation standards. Where two of the three sources agree a capability matters, it earns a row.
2. Score each source separately and resist the urge to average them in your head as you go. The gap only shows up if you score them independently first.
3. Overlay the three scores and look for three kinds of gaps: a readiness gap (entry vs. curriculum target — are you teaching to where students actually are?), a design gap (curriculum target vs. employer expectation — are you teaching to the wrong altitude?), and a perception gap (what your program believes it delivers vs. what alumni and employers report).
4. Build one matrix per learner at intake, then aggregate them into a single cohort-level matrix. The individual profiles tell you who needs what. The aggregate tells you what to prioritize in your next curriculum revision and where faculty development dollars should actually go.

A note on capacity: this doesn't have to stay a one-person job forever but sequence the help carefully. Pilot the process yourself first, on that one current cohort, until you trust the method. Only then is it a strong candidate for faculty researchers or class projects — a marketing research methods course building out the interview protocol, or an education doctoral student treating KSC validation as a thesis project. A class project can execute a known method well. It's the wrong venue for working out the kinks in a brand-new one. (More on these campus partnerships in the SILO cluster of plays.)

WHERE TO START THE BALANCE

You can build this matrix from any starting point, but I'd argue for one approach over the others: start with the learner and the employer and let the curriculum follow. Curriculum is the piece you have the most direct control over, which makes it tempting to lead with it. But curriculum should be the answer to a question, not the question itself. Anchor the matrix in what learners actually bring and what employers actually need, then let the gaps tell you what the curriculum has to become — not the other way around.

WHAT TO DO WITH THIS TODAY

You don't need the full matrix built by Friday. Pick three KSCs your program leans on most heavily in its marketing. Score them from all three angles — learner entry, curriculum target, employer expectation — using nothing more sophisticated than your

own best judgment and one conversation with a faculty member and one with an employer partner. See if the gaps surprise you. In most cases, they will.

RUN IT TWICE: THE MATRIX AS BOTH A PLAN AND A PROOF

There's a second use for this same matrix, and it might be the more persuasive one. Build it once at the start of an intake cycle to plan. Build it again — same KSCs, same six columns — near the end, and plot where each learner actually landed against where they started.

That second pass turns the Bridge Matrix from a planning tool into evidence, and three audiences care about that evidence in different ways:

- Accreditors — and your own cycle review process — want to see whether the design gap actually closed, not just whether you identified it. This is exactly the before-and-after comparison the Cycle Review and Reset process is built to use.
- Employers who supplied the original E-column input get to see whether the program delivered. Closing that loop with them — showing them their voice changed something concrete — does more for the relationship than any amount of relationship management without evidence behind it.
- Prospective students and other constituents respond to specific, honest proof of movement far more than another claim of excellence. "Our last cohort entered at Understand and left at Analyze" is a fundamentally different pitch than "our graduates are exceptionally well prepared."

Turning that second matrix into a finished report, an employer conversation, or marketing material is real work in its own right — worth its own play down the road. For now, the discipline that matters is simpler: build the matrix twice, not once, and keep the raw comparison somewhere you can find it.

FOR FURTHER READING

Bloom, B. S., Engelhart, M. D., Furst, E. J., Hill, W. H., & Krathwohl, D. R. Taxonomy of Educational Objectives: The Classification of Educational Goals. Handbook I: Cognitive Domain. David McKay Company, 1956. The original.

Anderson, L. W., & Krathwohl, D. R. (Eds.). A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives. Longman, 2001. The revision most programs actually use today.

This Playbook only gets built by hearing what's missing from it. If this play left a gap for your situation — or if you already know the play you wish existed but haven't seen written anywhere — send it my way. That's exactly how the next one gets written.

Need support? Three Ways to Engage – All Free

I built the Playbook to be useful, not to be sold. That means no paywalls, no retainers, and nothing to register for just to read the work.

Read the Plays. Every Play is free at <https://rgalsup.com/grad-program-intake-playbook/>. Start wherever your program needs the most help — each one stands on its own.

Join a free webinar. I host free monthly webinars on specific Playbook topics, built around real conversations with people working on the same problems you are. One-time registration is required. The meetings occur monthly, on the third Wednesday at 11:00 AM EST. [Click here to register](#).

Talk to me directly. If something here is hitting close to home for your program — faculty enablement, hidden data in your platforms, evidence that nobody is using — I would genuinely like to hear about it. Free one-on-one conversations are available via Calendly. Click here to schedule a [One-on-one with Rodney](#). Please book any available slot that works for you, as my schedule fills up fast.

The offer is open. Let's make a difference one play at a time.

Subscribe to the Grad Program Intake Playbook newsletter via LinkedIn so you do not miss a Play. View all published Plays at: <https://rgalsup.com/grad-program-intake-playbook/>.