

STONEBOUND

Learning Should Be Experienced, Not Just Taught

DESIGNING THE INTERVENTION FOR WHEN LEARNING STALLS · A MULTI-GRADE INSTRUCTIONAL DESIGN CASE STUDY

Tazz Kelly

Sole designer, curriculum architect, and Story Watcher (GM) · tavariskelly@gmail.com

The Problem

Engagement is the variable learning systems control worst. Content can be excellent, platforms can work exactly as built, and progress still stops — because nothing in the experience gives the learner a reason to care. The default fix is intensity: more practice, more reminders, more supervision. I've watched that fail often enough that I designed around the failure mode instead. When learning stalls, the content usually isn't broken. The experience around it is.

Stonebound is the intervention I designed for that failure mode — not a replacement curriculum, but an experience layer engineered on one premise: the way a challenge feels is as important as what it teaches.

I wrote myself a deliberately hard brief: three learners at three different levels (rising 1st, 3rd, and 5th grade), real academic work held at real grade-level standards, zero content dilution — and an experience they would ask to return to. Everything below is the set of design decisions that brief forced.

The Structure: A Week With a Plot

My first structural decision was to organize the week by narrative tension instead of by subject. Each week runs on a three-day rhythm, and I gave each day exactly one job:

- Tuesday — Campaign Day. A collaborative challenge opens the week's region of the story world, individual missions follow, and the day ends on a tease: something is coming Thursday.
- Wednesday — Training Day. Independent academic work on IXL and Miacademy, framed as two mini-bosses per child. Clearing them earns physical Advantages that change how Thursday's fight runs.
- Thursday — Boss Day. The collaborative boss fight lands everything Tuesday built and Wednesday earned. A region of the map is restored. A real-world field trip unlocks.

The sequencing is the design. Tuesday manufactures the want. Wednesday positions the academic work as the only path to it. Thursday pays it off publicly. I deliberately placed the academic platforms — the point where engagement most often dies — at the center of the week holding its highest-stakes currency, instead of at the end of a to-do list holding none. That placement decision does more work than any single mechanic.

Mechanics That Embody the Concept

My rule for every mechanic in the campaign: the game may never be decoration on top of the content. If the mechanic doesn't physically enact the academic concept, I cut it and design another. Two examples of that rule applied, from the live campaign:

The Number Dwarf — pattern under pressure

For a week targeting states of matter and math fluency, I needed classification to happen at speed, not on paper. So I built the boss as a pressure system: the Dwarf calls a state — solid, gas, or liquid — and the team answers with the right colored balls inside a timed window, draining 50 HP. I opened the design at ten seconds and widened it to thirty in play, tuned against real reaction times rather than the ones I imagined. I tuned the register to deliberate chaos on purpose — the Dwarf feeds on confusion, and the only counter I gave the team is pattern. No one is told that classification and quick recall are the lesson. The mechanic makes them the only way to win.

The Bridge Troll — sequencing you can step on

For sequencing — a skill that stays abstract on worksheets — I wanted physical commitment. The boss is a tile bridge: each child crosses a lane of six tiles by naming what comes next in the sequence before stepping. I made a wrong step cost the whole team their progress, deliberately: committing your body weight to an answer, in front of teammates who depend on it, converts a private academic guess into a public, felt decision. That conversion is the design.

One Activity, Three Grade Levels

The hardest constraint in my brief was multi-grade design: a rising 1st grader and a rising 5th grader in the same activity, both legitimately challenged, neither bored nor drowning. The standard solution is separate activities — which quietly abandons the shared experience. I refused that trade. My solution was to differentiate inside a single shared mechanic: difficulty is assigned to the lane, not to the child.

Lane	Sequence content	Academic level
Whale (rising 1st)	Color runs and simple repeating patterns	Pattern recognition, sequencing
Owl (rising 3rd)	Nature patterns and symbol chains	Abstract patterns, categorization
Elephant (rising 5th)	Logic, doubling, and multiplication sequences	Multiplicative reasoning

All three lanes run simultaneously, and I set the victory condition to require every Wielder arriving together. When the campaign later doubled to eight players, I restructured the lanes to pairs — two Wielders per path, never sharing a row — without touching the underlying mechanic; the differentiation held because it lives in the lane, not the roster. That choice makes the 1st grader's success and the 5th grader's success mechanically interdependent — converting age difference from the management problem it usually is into the engine of the teamwork itself. I also imposed a pairing rule on every collaborative challenge — each one combines two academic subjects (math with science, science with language arts, language arts with social studies) — so every shared activity carries double curricular load by design.

The Hidden Layer

I architected the system in two layers on purpose. Underneath the story sits an unglamorous, fully conventional academic map: grade-level standards, real IXL and Miacademy task lists, weekly completion targets. That layer is mine — the operator's layer — and the kids never see it. What I expose to them is a mission card.



Print on card stock

A Wednesday mission card. The narrative layer ('The Thing that didn't MATTER' — that week's states-of-matter theme) sits directly on top of the real assignment: complete four IXL sessions on the child's standards-aligned recommendation path.

The split is the principle: the learner-facing layer carries identity and momentum; the operator layer holds precision and standards. 'Complete IXL x4' is the identical assignment it would be on a chore chart — I changed everything around it instead: what it's called, what it earns, and what it means on Thursday. That's the intervention. Not new content — a redesigned experience wrapped around unchanged content.

Discovery Over Instruction

I imposed one rule on every printed material: mission cards withhold explanations. Rules are discovered through play, never read from the page. Players hypothesize, test, and correct — so the campaign's way of teaching its own rules rehearses the exact inquiry behavior I want it to build. Decoding a new boss mechanic is the first puzzle. That's not sparse writing; it's a documentation standard I held across eight weeks of materials.

I applied the same philosophy to assessment. Instead of opening with a placement test, I designed five short Trials — Logic, Observation, Insight, Craft, Focus — each built to measure what a learner can confidently do, not where they fall short. I surfaced results as starting stats and strengths, never as scores. I built the assessment to do one job: assign the right challenges. Labeling the learner isn't a job I gave it.

The Payoff Loop



The campaign map. Each Thursday victory physically clears fog from one region — restoration the kids can watch accumulate across the summer.

Every week closes the same way, by design: after the boss falls, a Wielder — never the adult — clears the fog from that week's region. Then the field trip unlocks: a real-world outing themed to the restored region, sourced from a single family activity pass so the whole reward layer runs on a fixed budget — a constraint I built the region themes around, not against. The loop closes as architected: story creates desire, academics earn advantage, collaboration wins the day, and the world — fictional and real — visibly opens.

Status & What This Demonstrates

I designed and locked the full eight-week campaign before summer began: sixteen session runsheets, eight boss mechanics with per-week advantage mappings, printed mission cards and game components for every week, and a companion web app I built alongside it. The campaign is live and mid-run, and has absorbed growth from three participants to eight without structural changes — because scalability was a design assumption, not an afterthought.

As a work sample, this is what it demonstrates about how I design: backward from engagement rather than forward from content; differentiation engineered inside shared activities rather than around them; real standards held under a motivating identity layer; assessment built to assign challenges rather than assign labels; and complete, produced, facilitator-ready delivery — not a concept deck.

A companion case study, 'Designing a Cross-Platform Learning Economy,' documents the motivation and reward architecture of the same campaign.