

Adaptive Learning Readiness Checklist

Run this gate before building an adaptive learning path. If most boxes are unchecked, the honest first project is instrumentation and content granularity, not an engine.

A. Goal vs method

- We separated the goal (a path that fits each learner) from the method (automatic adjustment)
- We confirmed we need genuine adaptivity, not authored branching with good prerequisites
- We can name which levers will adapt: pace, path, difficulty, content, or support

B. The data the engine needs

- Player and assessment events are captured as fine-grained, skill-tagged statements
- Tracking uses xAPI 1.0.3 (or Caliper) writing to a Learning Record Store
- Statements carry per-skill detail: item, right/wrong, time, retries, segments, hints

C. Content and the learner model

- Content is granular and tagged to skills, not one undifferentiated unit
- There are enough items and units per skill to route a learner to
- We chose a learner model (rule-based, IRT/CAT, knowledge tracing, or hybrid LLM)
- Any LLM layer sits over a structured, auditable model for routing decisions

D. Decision rules and evaluation

- Decision rules are defined, tested, and explainable to an instructor
- We measure outcomes (completion, mastery, retention) vs a non-adapted cohort
- The same xAPI/Caliper signals feed analytics so we can prove the path helped