Rill Athletics | Case Study

An AI-native client intelligence platform for personal trainers

Dec 2023 – Jun 2025 • Lisbon, Portugal • Solo founder

The Problem

Personal trainers see their clients for 2-3 hours a week. The other 165 hours are invisible to them. Diet, sleep, stress, motivation, life events. All of it shapes whether a client progresses, and none of it shows up in the gym.

When a client stalls, the trainer has no diagnostic data. They guess. They ask questions during sessions, but session time is limited and clients default to vague answers ("doing okay"). Existing tools tracked diet logs and workout volumes. None of them captured the psychological and life-context inputs that actually determine whether a client follows the plan.

The result is predictable. Clients hit plateaus, trainers can't explain why, and churn follows. Trainers also lose the ability to spot upsell-ready clients, the ones who are low-stress, on-program, and primed for a higher-tier service.

The Insight

Trainers were already trying to gather this information. The friction wasn't the desire. It was the format.

Verbal probing during sessions takes time, requires the client to articulate vague psychological states on the spot, and depends on the trainer remembering details across 10-20 clients. The data trainers needed wasn't expensive to collect. It just needed a frictionless capture mechanism that worked between sessions.

Anonymized, structured surveys had the opposite friction profile. Clients shared more honestly when the format was simple (sliders, emoji scales, multiple choice) and the response felt low-stakes. A 60-second pre- and post-session check-in could capture mood, motivation, life-context, and habit adherence in a format clients would actually complete. An AI layer could then synthesize the responses into a longitudinal client model the trainer could read in 30 seconds before each session.

What I Built

Rill was a web app with two surfaces.

Client side. Two short surveys per session cycle. A pre-session check-in with 5 questions covering mood, motivation, weekly routine consistency, session intent, and an optional note to the trainer. A post-session check-in with 7 questions covering weekly habit adherence, perceived progress, motivation, and what felt hardest. Each survey took under 90 seconds. Questions were designed around behavioral psychology research to surface stress, adherence, and readiness signals indirectly, instead of asking clients to self-diagnose.

Trainer side. A dashboard that aggregated each client's survey responses into a running profile. Trainers got a quick pre-session summary ("here's where your client is mentally walking into today") and a longitudinal post-session view that tracked motivation, habit adherence, and life-context signals over time. Each client carried a color-coded status indicator (green, yellow, red) generated by the AI layer, so trainers could open the app and immediately see which clients needed attention instead of scrolling through every profile. The AI compiled survey responses into actionable, natural-language summaries. For example: *"Client motivation appears lower this week due to schedule instability and inconsistent sleep."* The product was built around one principle: surface insight, not data. The AI interpreted, the UI simplified, the trainer acted.

Stack. Built on Replit, with Supabase for backend and several Google Sheets API integrations for data flow. UI and frontend built from scratch as solo developer.

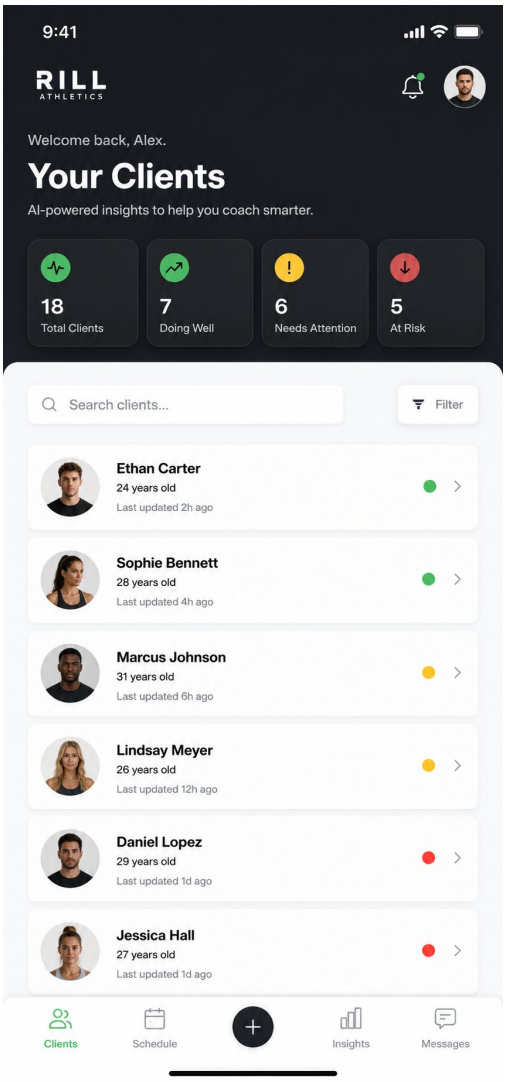

Trainer dashboard. Client roster with status indicators

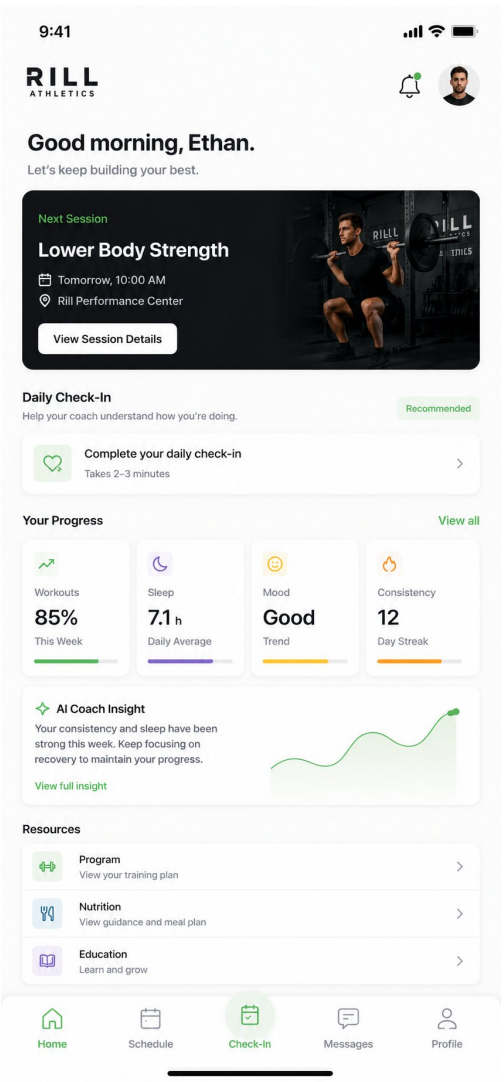

Client home. Daily check-in and progress view

Distribution & Revenue

Acquired the first paying trainers through a mix of in-person outreach at Mr. Big Evolution (one of Lisbon's top bodybuilding gyms), my own network, fitness expos, and cold calls to independent trainers in the Lisbon area. Qualifying questions surfaced trainers who already felt the information gap:

- "How do you figure out how a client is doing outside the gym?"
- "Would you benefit from knowing a client's stress levels and where the stress is coming from?"

Talked to roughly 75 trainers before landing the first 7 paying customers. A 25-45% conversion rate at the qualification stage. Pricing was $35/month per seat. Seven paying trainers at that price point validated that independent trainers were willing to put real money on a structured client-intelligence product.

The Outcomes That Validated the Thesis

Two pieces of evidence anchored the core thesis.

Documented training validation. Ran a structured four-month bulk protocol against a controlled sample, with full Rill data capture throughout. Pre and post-session surveys, readiness scoring, recovery and adherence tracking. Results showed roughly 50% improvement in tracked compound lift performance (squat, bench, deadlift) across the period. The protocol was designed to test whether Rill's data could actually inform meaningful training adjustments. It could.

An attributable customer upsell. One trainer noticed through the dashboard that a client was consistently on-program, low-stress, and reporting high motivation across multiple weeks. The data signaled the client was ready for a higher-tier service. The trainer pitched a tailored training-plus-diet package and converted the client from €100/month training-only to €150/month full-service. A 50% revenue lift on a single account, attributable directly to insight surfaced by Rill.

This was the case that proved the original thesis from the customer side. Closing the trainer's information gap doesn't just reduce churn. It unlocks revenue the trainer couldn't otherwise see.

What Didn't Work and What I Learned

Client survey adoption was the constraint that ultimately limited the product. Adoption rates were around 36%, meaning roughly 1 in 3 clients filled out surveys. For every client who did fill out the surveys, only about 50% of the time were both surveys completed. The system architecture required both pre- and post-session surveys to generate a complete client profile, and the pre-session survey was missed most often. Likely because clients showed up to the gym without thinking to open the app first. Without consistent data, the dashboard surfaced gaps instead of insights.

The deeper lesson was about who I designed the product for. I built Rill around the trainer's needs and treated client adoption as a downstream implementation detail. In retrospect, the client side should have been the primary design surface. The value proposition for the trainer collapses entirely if clients don't engage with the data capture. If I were rebuilding, I'd run dedicated discovery on the client experience first, redesign around a single post-session-only flow (which had materially higher completion rates), and make the trainer-facing value proposition something like: "you'll know more about your client than you do today, without asking them anything different in session."

The product wound down in mid-2025 when academic exams forced a pause in active outreach. Trainers needed consistent prompting to keep using the dashboard, and once the daily nudges stopped, usage decayed. The decision not to restart was deliberate. The friction between the product and the user was structural, not motivational. Continuing would have required rebuilding the engagement model from the ground up.

What I'd Take Forward

Three things from Rill that I bring to any product or operator role.

Customer discovery muscle. The 75 trainer conversations before landing the first paying customer taught me how to qualify quickly, which questions actually surface pain instead of agreement, and how to recognize when a prospect is signaling interest vs. politeness.

Designing for the bottleneck user. The hardest lesson from Rill was that the user you design around determines whether the product works. I built around trainers and lost on clients. In any future product or strategy role, the first question I ask is: which user has the highest leverage on whether this works, and am I designing primarily for them?

Compounding data products. Rill's underlying thesis was that small, frequent, structured data inputs compound into something irreplaceable over time. That's the same thesis underneath every winning AI product I've seen since. Getting the capture mechanism right is the whole game.