PURPLE KIWI ADVISORY

STRATEGIC BRIEF

Pull Products Don’t Compound

Converting crisis usage into habit — a retention architecture brief for founders

A STRATEGIC BRIEF · AUGUST 2026 · RETENTION ARCHITECTURE AND PRODUCT STRATEGY. NO IMPLEMENTATION CODE.

01 The pattern

There is a category of product that people genuinely love and almost never open.

Users describe it in interviews the way you'd describe a good doctor or a trusted friend — indispensable, hard to imagine being without. Then you look at the logs. Across the teams I have worked with the composite picture is the same: the median user has opened it a handful of times in a quarter. Support tickets are glowing. NPS is strong. Weekly actives are flat.

This is not a product problem, and it is usually not a churn problem. It is a **pull product**: one the user reaches for when something is wrong, and has no reason to touch when nothing is. Legal tech, incident response, HR case management, insurance claims, tax software, security tooling, and most of the interesting AI-native consumer products all share this shape. The value is real, it is acute, and it is episodic.

Episodic value is a fine business. It is a very hard *compounding* business. The purpose of this brief is to name why, and to lay out the specific architecture that converts pull into push without burning the trust that made the product work in the first place.

02 Your retention curve is lying to you

The standard retention curve assumes a return-visit model: a user who hasn't come back in thirty days has churned. In a pull product, that user may simply not have had a crisis yet.

The distinction matters because it changes the diagnosis entirely. **Latency is not churn.** A user with a sixty-day natural need interval will look identical to a churned user for fifty-nine days, and then transact. If your dashboard treats them the same, you will spend money reactivating people who were never gone and miss the people who quietly stopped trusting you.

Three corrections make the picture honest.

Set your churn window from observed need frequency, not convention.

Measure the distribution of gaps between sessions for users who *did* return. If the median gap is forty-five days, a thirty-day churn window is manufacturing a crisis that doesn't exist.

Separate crisis sessions from calm sessions.

Every session should be classified. In my experience most teams have never done this, and the result surprises them — in a true pull product, calm-state sessions are typically a small single-digit share of total volume.

Track interval compression, not visit count.

The goal of a pull-to-push strategy is not more visits in the abstract. It is a shrinking median gap between sessions. That single series tells you whether the architecture is working, months before it shows up in a cohort chart.

03 Is your product pull-based?

Five questions. Three or more yes answers and this brief is about you.

1. Can users articulate the moment they last used the product, but not the last week they used it?
2. Does your best marketing copy describe a problem rather than a routine?
3. Is your session length long and your session frequency low?
4. Would a user feel slightly odd opening the product when nothing was wrong?
5. Does your product get better the more it knows about the user — but only learn during crises?

Question five is the expensive one. If your product improves with accumulated context, then a pull pattern means your learning rate is bounded by how often the need arises. Context accumulates only on the days the user needs you most.

04 Why pull caps the business

Three ceilings, each independent.

The frequency ceiling.

Willingness to pay for episodic value is anchored to episode count. Subscription pricing on a four-episodes-a-year product creates a cancellation conversation at every renewal.

The memory ceiling.

Users forget you between episodes. Every crisis session begins with reacquisition — you are re-earning attention you already paid for. Word of mouth decays too, because nobody recommends a product they haven't thought about in two months.

The learning ceiling.

For any product whose quality depends on accumulated user context, low frequency means slow learning, which means the moat compounds at the rate the need recurs rather than at the rate the user engages. This is the ceiling founders discover last and care about most.

05 The Three Bridges

The move from pull to push is not one change. It is three, and the order matters, because each one funds the next.

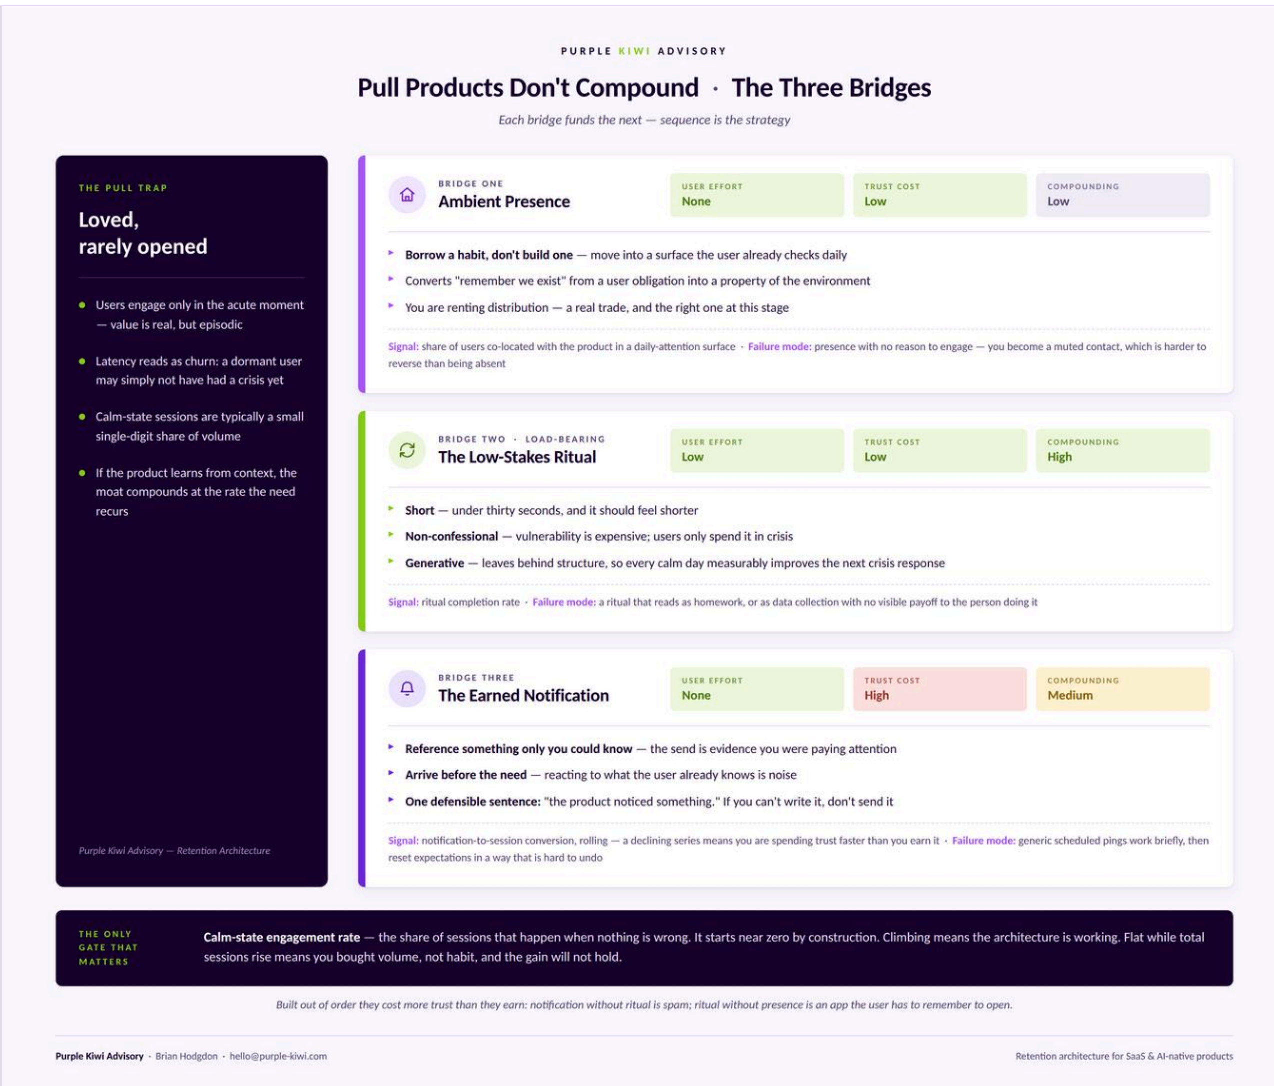

Exhibit 1. The Three Bridges. Each bridge funds the next; built out of order, they cost more trust than they earn.

Bridge 1 — Ambient Presence

Stop asking for a new habit and borrow an existing one. Move the product into a surface the user already checks: messaging, email, calendar, the tools they live in. This costs the user nothing and asks nothing, which is precisely the point. It converts "remember we exist" from a user obligation into a property of the environment.

The trade is platform dependency, and it is a real trade — you are renting distribution, and the landlord can change the terms. Rent it anyway at this stage. Owning a surface nobody opens is not an asset.

Signal to watch: the share of users co-located with the product in a daily-attention surface. **Failure mode:** presence with no reason to engage. You become a muted contact, which is worse than being absent, because it is harder to reverse.

Bridge 2 — The Low-Stakes Ritual

Give the user a reason to show up when nothing is wrong. This is the load-bearing bridge, and the one most teams skip.

A ritual that works has three properties. It is **short** — under thirty seconds, and it should feel shorter. It is **non-confessional** — it cannot require the user to be vulnerable, because vulnerability is expensive and users only spend it in crisis. And it is **generative** — it leaves behind structure the product can use, so that each calm interaction measurably improves the next crisis response.

That third property is where compounding lives. A ritual that only produces engagement is a treadmill. A ritual that produces *structured knowledge of the user* turns calm days into moat.

Signal to watch: ritual completion rate, and calm-state engagement rate. **Failure mode:** a ritual that reads as homework, or as data collection with no visible payoff to the person doing it.

Bridge 3 — The Earned Notification

Only now do you get to interrupt.

A notification is a withdrawal from a trust account that Bridge 2 funded. Three rules keep the account solvent. Reference something only you could know — the notification should be evidence that the product was paying attention. Arrive *before* the need, not after; a notification that responds to something the user already knows is noise. And make every send defensible in one sentence: *the product noticed something*. If you can't write that sentence, don't send it.

Signal to watch: notification-to-session conversion on a rolling basis. A declining series means you are spending trust faster than you are earning it, regardless of what the absolute number looks like. **Failure mode:** generic scheduled pings. They work briefly, then reset the user's expectation of what your alerts mean — a reset that is very hard to undo.

> **Sequence is the strategy. Built out of order, the bridges cost more trust than they earn. Notification without ritual is spam. Ritual without ambient presence is an app the user has to remember to open — which is the original problem, restated.**

06 The one metric that matters

Most of the instrumentation above is diagnostic. One number is the gate.

> THE GATE
>
> **Calm-state engagement rate** — the percentage of sessions that occur when the user is not in an acute moment.

It starts near zero by construction. If it is climbing, the architecture is working even when nothing else has moved yet. If it is flat while total sessions rise, you have bought volume rather than built habit, and the gain will not hold.

07 In practice: HeyLina

HeyLina is an AI relationship strategist for emotionally intelligent dating — *"dating feels like chaos, we built clarity."* Led by CEO Briana Longe and COO Angus Hally, the team has built something genuinely uncommon: a product designed for the moment the need is sharpest, and built to hold the kind of detail users would hesitate to tell a friend. Trust like that is among the hardest assets in consumer AI to acquire, and the team acquired it early.

That same strength produces a classic pull pattern. Users come to Lina after the confusing text, the sudden silence, the argument — moments of real emotional weight that arrive on their own schedule rather than a daily one. Product quality was never the constraint. The strategic question was different and more interesting: *how do you become present between the hard moments?*

The team's architecture maps cleanly onto the three bridges, and they were already building it. **Ambient presence** through WhatsApp integration, putting Lina inside a messaging surface users already live in rather than asking them to open a separate app during a stressful evening. **A low-stakes ritual** in the form of daily check-ins — short, light-touch, and generative, because each one enriches Lina's model of the user. And a **roster** of structured context on the people in a user's life, which turns calm-day input into an asset that makes the next crisis conversation materially better rather than starting cold.

What makes the design strong is the compounding link between the layers — and a decision about what *not* to build. HeyLina does not score its members. What accumulates instead is a non-pejorative profile of each member, and an emotional intelligence layer that grows more capable as those profiles deepen. Every calm-state interaction improves the product's response in the hard moments, and none of it requires reducing a person to a number.

That restraint is worth dwelling on, because it is Bridge 3's argument in a different register. Scores are legible, demoable, and easy to build an interface around — which is exactly why so many products in this category reach for them — and they convert a relationship into a judgment the instant a user sees it. Choosing the profile over the score protects the trust the product runs on, at the cost of the more visible artifact. That is the sharpest version of the trade this brief describes. Briana and Angus were building toward this before it had a name — usually the sign of a team whose product instincts are running ahead of the vocabulary the industry has for them.

08 Four ways this goes wrong

None of these are HeyLina-specific — they are the four ways I most often see this architecture fail.

Buying frequency with volume.

Notification cadence is the fastest lever and the hardest to reverse. Trust spends down quickly and rebuilds slowly.

Gamifying a serious product.

Streaks punish the user for having a good week. On any product with emotional or professional weight, streak mechanics invert the relationship: the user now owes *you* something. Reserve them for products where the user's goal genuinely is frequency.

Feature-stuffing to manufacture reasons to visit.

Adding surface area to create visit occasions dilutes the thing users came for and rarely moves calm-state engagement. One well-designed ritual beats six mediocre features.

Mistaking the ritual for the product.

The crisis moment is still where value is proven and where willingness to pay is set. The ritual exists to make that moment better and to keep you present until it arrives. If ritual quality starts to crowd out crisis-response quality, the strategy has eaten itself.

09 A 30-day diagnostic

WEEK 1 — **Baseline.** Classify every session as crisis or calm. Compute calm-state engagement rate and the median inter-session interval per cohort. Reset your churn window to observed need frequency. *Expect the calm-state number to be lower than anyone on the team guessed.*

WEEK 2 — **Habitat.** Identify where your users' attention already lives daily. Choose one surface. Do not build a second.

WEEK 3 — **Ritual.** Ship exactly one ritual. Thirty seconds or less, no vulnerability required, and it must produce structured data the product can use. Instrument completion.

WEEK 4 — **Gate.** Read calm-state engagement rate, not DAU.

> THE DECISION GATE
>
> If calm-state engagement has not moved meaningfully above baseline within thirty days of shipping the ritual, **the ritual is wrong — not the strategy.** Change the ritual. Do not reach for notifications to compensate; Bridge 3 built on an unfunded Bridge 2 will produce a short spike and a permanent cost.

Brian Hodgdon is a technology and product executive — previously VP of Engineering at CDK Global, with senior roles at Dropbox and Hewlett Packard Enterprise — and advises SaaS and AI-native companies on retention architecture, activation, and engagement through Purple Kiwi Advisory. This brief draws on advisory work with HeyLina and is published with the team's permission.