

# Notion: From Note-Taking to Identity Infrastructure

How a PLG product engineered compounding behavioral investment —  
and where the architecture fractures between product and lifecycle.

Isaac Olukoya | Perennus

August 2026 | Case Study

---

# Summary

This case study examines the behavioral revenue architecture behind Notion's reported ~\$700M ARR. It is based on two primary product audits (February and June 2026), published behavioral psychology research, SaaS industry benchmark data, and a post-publication conversation with Notion's Head of EMEA Marketing, who engaged with the original research and shared it internally.

## Central Thesis

Notion's growth is best understood as a compounding behavioral sequence: identity priming → competence formation → habit engineering → data gravity → collaboration → champion emergence → community lock-in. Each stage raises the switching cost by one layer. The revenue is a lagging indicator of these decisions.

## Primary Observed Gap

Notion's onboarding collects four behavioral signals (name, use case, context, and specific intent items). The in-product experience deploys all four. The lifecycle email system appears to deploy none of them. This gap widened between February and June 2026.

## Primary Hypothesis

Notion's AI transition may produce faster early activation but shallower long-term behavioral investment. This is a hypothesis requiring internal cohort data to test, not a confirmed finding.

## Methodology

Observe → Diagnose → Hypothesize → Design state architecture → Build implementation → Define experiments → Connect to revenue.

# The Compounding Behavioral System

Notion did not build a note-taking tool. It built a behavioral system where each layer of user investment makes the next more likely and the exit more costly.

| STAGE                   | MECHANISM                                                      | WHAT IT PRODUCES                                                         |
|-------------------------|----------------------------------------------------------------|--------------------------------------------------------------------------|
| 01 Identity Priming     | Name + use case declaration before any feature                 | Commitment bias. User tells the product who they want to become.         |
| 02 Risk Elimination     | Templates transfer competence; freemium removes financial risk | Time-to-value collapses. Blank page fear neutralized.                    |
| 03 Competence Formation | Slash command creates feeling of command, not navigability     | Self-efficacy. Users return to activities where they feel capable.       |
| 04 Data Gravity         | Every page, database, linked task increases cost of leaving    | Switching cost grows automatically with usage.                           |
| 05 Habit Formation      | Multiple use cases drive multiple daily opens                  | Products opened daily become infrastructure.                             |
| 06 Multiplayer          | @mention or page share makes leaving disrupt someone else      | Interpersonal switching cost — categorically stronger than product lock. |
| 07 Champion             | Power user builds workspace their team depends on              | Unpaid internal advocate. Enterprise deal starts from 'yes.'             |
| 08 Community Lock-In    | Identity partially constructed around the product              | Leaving = self-erasure. Social switching cost no import tool addresses.  |




Figure 1 — Behavioral Revenue Architecture

Retention is not held by any single stage. It is held by the sequence. A user at Stage 4 is retained through data gravity. A user at Stage 6 is retained through social accountability. A user at Stage 8 is retained through identity — and identity is not migrated by any import tool.


Products that sell features compete on capability. Products that sell identities compete on belonging.

# Product Intelligence → Lifecycle Intelligence

The strongest observed finding is a structural disconnect between the behavioral intelligence the product collects and the behavioral intelligence reflected in the lifecycle experience.

## What the Product Collects

During onboarding, Notion asks: ‘How do you want to use Notion?’ with three declared use cases — For work, For personal life, For school. This is the first segmentation layer. As of June 2026, the product then collects specific intent items within each use case and the user’s name.

A screenshot of the Notion profile personalization interface. At the top, there is a circular profile picture placeholder with a large letter 'I' and a small checkmark icon. Below this, the heading 'Customize your profile' is displayed, followed by the subtext 'This is how you will appear in Notion'. A form field labeled 'Your name' contains the text 'isaac'. Below the form field is a blue button labeled 'Continue →'. At the bottom, there is a checkbox that is checked, with the text 'I agree to Notion marketing communications'. Below this, a small disclaimer states: 'You may unsubscribe from receiving marketing communications at any time. Notion's websites and communications are subject to our Privacy Policy.'

*Figure 2 — Notion Profile Personalization*




Figure 3 — Declared Use Case Selection

| SIGNAL                                 | WHAT IT REVEALS                                             | POTENTIAL LIFECYCLE USE                           |
|----------------------------------------|-------------------------------------------------------------|---------------------------------------------------|
| 01 Name                                | Identity anchor                                             | Personalization across all communications         |
| 02 Use case (Work / Personal / School) | Intent category — first segmentation layer                  | Email sequence routing, activation path selection |
| 03 Context                             | Environment signal                                          | Content framing, feature prioritization           |
| 04 Specific intent items (June 2026)   | Granular declaration: spending, journal, projects, research | Micro-activation paths, template recommendations  |

## What the In-Product Experience Does

The product uses these signals. The June 2026 Welcome screen addresses the user by name, references their use case, and presents intent-specific options. The behavioral data collected during onboarding is actively deployed to personalize the first session.

## What the Lifecycle Email Does

The lifecycle email appears to use none of these signals. In both audits, the first email promoted the AI Agent to a user who had not established basic activation. No name. No reference to declared use case. A user who selected ‘For personal life’ received the same email as a product manager who selected ‘For work.’




Figure 4 — Lifecycle Email #1: AI Agent Introduction

## The Four-Email Forensic Timeline

| TOUCHPOINT                | CONTENT                                                | OBSERVED ISSUE                                        |
|---------------------------|--------------------------------------------------------|-------------------------------------------------------|
| Email #1 (Day 0, +15 min) | AI Agent: ‘You assign the tasks. Agent does the work.’ | Presents ceiling before floor. No personalization.    |
| Email #2 (Day 2)          | ‘Confused and need some help?’                         | Identity contradiction. Product said user is capable. |

|                            |                                             |                                                                       |
|----------------------------|---------------------------------------------|-----------------------------------------------------------------------|
| Email #3 (Day 9)           | CEO newsletter: Custom Agents deep-dive     | Context mismatch. Power-user content to potentially unactivated user. |
| Email #4 (Day 9, same day) | 'Don't miss your chance for expert support' | Orchestration collision. Two emails same day. Deliverability risk.    |




Figure 5 — Lifecycle Email #2: Expert Support Offer

Four emails. Four different assumptions about who the user is. Zero continuity from the behavioral signals collected during onboarding.

## Artifact 02 — Product to Lifecycle Intelligence Gap





Figure 6 — Product → Lifecycle Intelligence Gap

The argument is not 'Notion writes bad emails.' The argument is: there appears to be a structural disconnect between the behavioral intelligence the product collects and the behavioral intelligence reflected in the lifecycle experience.

## June 2026: The Gap Widened

Four months after the original analysis, a complete re-audit was conducted. The product had changed significantly. The lifecycle system had not changed at all.


A screenshot of the Notion sign-up page for AI workspace. At the top is the Notion logo (a blue 'N' in a square). Below it, the text reads "Notion: your AI workspace. Sign up with your work email". There is a text input field labeled "Work email" containing the email address "isaac@isaacolukoya.com.ng". Below the input field is a tip box that says "Tip: Use your work email for a free trial of Notion's AI Agent, Meeting Notes, and more". A large blue button labeled "Continue" is positioned below the tip. Underneath the button is a horizontal line with the text "or continue with" in the center. Below this line are two buttons: one for "Google" with its logo and one for "Microsoft" with its logo. At the bottom, there is a link that says "Existing user? Log in".

*Figure 7 — Notion's AI Workspace Positioning*

The signup page now positions Notion as 'your AI workspace' — the Commander identity, deployed as the primary acquisition framing. The profile screen remains structurally identical. Identity priming is intact; the identity being primed has shifted.

### The Fourth Signal

The most significant change: the product now presents granular intent items mapped to the user's declared use case. A 'personal life' user sees: 'Plan a trip,' 'Track my spending,' 'Set personal goals & habits,' 'Keep a journal.' A 'work' user would see project and collaboration options. A 'school' user would see research and note-taking options.




*Figure 8 — Personalized Welcome Experience*

Four behavioral signals now collected during onboarding. The product team added more personalization intelligence between February and June 2026.

## The Commander Transition — Now Observable

The Welcome page features a persistent AI prompt bar: ‘Describe your own.’ The Architect-to-Commander transition predicted in the original analysis is now visible in the production interface.



*Figure 9 — The Commander Interface*

## What Did Not Change

Fifteen minutes after completing the re-audit signup, the first lifecycle email arrived. Subject line: identical to February. Content: identical. No name. No reference to declared use case. No reference to specific intent items. A user who explicitly said ‘I want to keep a journal’ received an email promoting enterprise AI delegation workflows.

| DIMENSION                      | IN-PRODUCT (JUNE 2026)              | LIFECYCLE EMAIL (JUNE 2026)  |
|--------------------------------|-------------------------------------|------------------------------|
| User’s name                    | Used: ‘Welcome to Notion, isaac!’   | Not used                     |
| Use case signal                | Mapped to personalized intent items | Ignored                      |
| Specific intent items          | Collected: trips, spending, journal | Not referenced               |
| Identity framing               | Personal life, consumer context     | Enterprise AI delegation     |
| Signals collected vs. deployed | 4 collected, 4 used in-product      | 4 collected, 0 used in email |

The product team built a personalized, context-aware first session. The observed lifecycle sequence remained the same un-segmented template from February.

# Identity Mismatch & the AI Question

## The Identity Mismatch Problem

Notion sells the Architect — someone who builds systems. That identity is aspirational. The hypothesis: Notion’s activation architecture cannot distinguish between the user who will become the Architect and the user who merely wants to become one. These populations have categorically different lifetime values. They look identical at signup. They diverge by day fourteen.

The behavioral signals are readable: templates installed but never customized, single-page depth, fewer than three return visits, no database creation. Users who install templates without modifying them, remain at single-page depth, return infrequently, and never create structured data are showing a measurable mismatch between declared intent and behavioral investment.

## The AI Hypothesis: Architect → Commander

Notion’s AI transition reportedly improved early activation metrics. The Commander identity (‘direct the AI’) requires less manual investment than the Architect identity (‘build the system’).

**Hypothesis: If Commander-identity users build less manual structure, they accumulate less data gravity and lower switching costs — creating a potential divergence between early retention and long-term retention.**

This is explicitly a hypothesis, not a finding. The data required to answer it is inside Notion’s cohort analytics.




Figure 10 — AI-First vs. Manual-First Cohort Model

## Fast vs. Deep vs. Durable Activation

| LAYER              | WHAT IT MEASURES                  | AI-FIRST                                               | MANUAL-FIRST                                      |
|--------------------|-----------------------------------|--------------------------------------------------------|---------------------------------------------------|
| Fast Activation    | First value moment                | Minutes (strong)                                       | Weeks (slower)                                    |
| Deep Activation    | Structural investment             | Unknown — AI generates structure user may not maintain | Dense — manual construction = personal investment |
| Durable Activation | Compounding behavioral investment | Hypothesis: potentially lower switching cost           | Very high after day-30 habit threshold            |

AI dramatically improves Fast Activation. The open question is whether it produces equivalent Deep and Durable Activation. We don't know the answer. That's the experiment.

# The Behavioral Lifecycle Decision Model

The observation identified the fracture. The next question: what would a lifecycle system look like if it read the behavioral signals the product already collects?

The proposed model operates across five layers. Critically, it applies to all three declared use cases — not just personal users:

| LAYER                     | PERSONAL PATH                        | WORK PATH                             | SCHOOL PATH                         |
|---------------------------|--------------------------------------|---------------------------------------|-------------------------------------|
| 1. Declared Use Case      | For personal life                    | For work                              | For school                          |
| 2. Specific Intent        | Journal / Spending / Trips           | Projects / Goals / Meeting notes      | Research / Notes / Tasks            |
| 3. Observed Behavior      | Created page, no return              | Workspace created, no team invited    | Notes created, unorganized          |
| 4. Behavioral State       | AT_RISK_MISMATCH                     | ACTIVATING                            | ACTIVATING                          |
| 5. Lifecycle Intervention | Path reset: single-use-case template | Collaboration prompt: invite teammate | Study system: organization template |

## Artifact 03 — Behavioral Lifecycle Decision Model




Figure 11 — Behavioral Lifecycle Decision Model

Each user's state is determined by the intersection of their declared intent (from onboarding) and their observed behavior (from product analytics). The lifecycle system should read both signals continuously — not just at the moment of signup.

This is not basic personalization. This is behavioral lifecycle architecture: the user tells us who they are trying to become, their behavior tells us whether they are becoming it, and the system reconciles those two things.

# From Behavioral Model to Executable Lifecycle

The proposed decision model translates into a Customer.io automation: onboarding context establishes the initial path, a short observation window allows behavioral signals to accumulate, and those signals determine the next intervention rather than elapsed time alone.




Figure 12 — Behavioral Lifecycle Implementation

The workflow branches on behavioral depth signals after a 3-day observation window.

Engaged users (`pages_created > 0`) continue to a deepening check: has the user customized a template? If yes, they reach ACTIVATED state. If not, they receive a deepening intervention. Users who show no meaningful engagement route to a Day-10 path reset.

The workflow does not blindly send the next email after three days. It observes behavior and branches. That is the difference between elapsed-time sequencing and behavioral lifecycle orchestration.

**This is a proposed implementation — not a claim about Notion’s internal automation architecture. It demonstrates how the behavioral model identified in this audit could be operationalized in a lifecycle platform.**

# What I Would Test

Three experiments. Each connects observation → hypothesis → intervention → measurement → revenue. Click-to-expand format in the digital version; full methodology below.

## EXPERIMENT 01: AI-FIRST vs MANUAL-FIRST COHORT DEPTH

**PHASE 1 (OBSERVATIONAL):** Do AI-first and manual-first users exhibit different long-term behavioral depth? Compare at D30, D90, D180, D365. Measure: retention, engagement depth, data gravity, collaboration, expansion, churn.

**PHASE 2 (IF DIVERGENCE FOUND):** Test whether interventions encouraging manual structure among AI-first users improve downstream retention.

**POPULATION:** All new users, segmented by primary activation path in first 7 days.

**CONTROL:** Current experience — both paths available, no segmentation in lifecycle.

**TREATMENT:** AI-first users receive prompts to convert AI output into manual structure.

**PRIMARY METRIC:** D180 retention rate, segmented by AI-first vs. manual-first cohort.

**GUARDRAILS:** Month-1 retention and activation rate must not degrade.

**COMMERCIAL OUTCOME:** Answers whether the AI transition is strengthening or diluting the retention architecture.

## EXPERIMENT 02: DAY-10 IDENTITY MISMATCH INTERVENTION

**OBSERVATION:** Users who declared an identity during onboarding but haven't developed corresponding behavioral depth by day 10.

**DETECTION SIGNALS:** Template installed but not modified + single-page depth + fewer than 3 return visits + no database creation.

**POPULATION:** Users matching **AT\_RISK\_MISMATCH** state by day 10.

**CONTROL:** Current lifecycle experience (generic re-engagement if any).

**TREATMENT:** Path reset: simplified single-use-case activation mapped to declared onboarding intent. One template. One completed action.

**PRIMARY METRIC:** D30 retention rate for treated **AT\_RISK** users vs. control.

**GUARDRAILS:** Unsubscribe rate should not increase > 0.5%.

**COMMERCIAL OUTCOME:** If 10% of at-risk users are recovered into the activation funnel, incremental paid conversion at scale is significant.

### EXPERIMENT 03: BEHAVIORAL LIFECYCLE ORCHESTRATION

**OBSERVATION:** Notion collects four onboarding signals. The lifecycle system appears to use none of them.

**POPULATION:** All new free-tier users. Segmented by declared use case (Work / Personal / School) and first-week behavioral state.

**CONTROL:** Current experience — identical email sequence regardless of declared intent or behavior.

**TREATMENT:** Behaviorally adaptive lifecycle: use case determines email sequence, name appears in every communication, feature used first determines next action, days without core action triggers intervention thresholds.

**REQUIRED SIGNALS:** **use\_case\_selected**, name, **first\_page\_created**, **template\_installed**, **template\_customized**, **database\_created**, **ai\_interaction**, **collaboration\_started**, **days\_since\_last\_session**.

**SUPPRESSION:** Max one lifecycle email per 48 hours. Suppress generic product marketing for users in adaptive path.

**PRIMARY METRIC:** Time-to-activation (time from signup to first ACTIVATED-state event).

**GUARDRAILS:** Email deliverability and free-to-paid conversion must not decrease.

**COMMERCIAL OUTCOME:** Improving activation depth produces measurable downstream retention and expansion improvement.

# How Behavioral Change Becomes Revenue

Every intervention must connect to a commercial outcome. The measurement framework:

| BEHAVIORAL CHANGE               | PRIMARY METRIC                      | COMMERCIAL OUTCOME                   |
|---------------------------------|-------------------------------------|--------------------------------------|
| Faster time-to-activation       | Time-to-value, activation rate      | Free-to-paid conversion, CAC payback |
| Higher engagement depth         | DAU/MAU, content complexity         | Monthly churn, NRR                   |
| Reduced identity mismatch churn | D14 depth, D30 retention            | Gross churn, LTV                     |
| Improved lifecycle relevance    | Email engagement, 24hr return rate  | Week-1 churn, deliverability health  |
| AI-first retention parity       | D180 retention by cohort            | NRR, LTV:CAC ratio                   |
| More expansion triggers         | Page shares, @mentions, invitations | Expansion revenue, ACV, sales cycle  |




Figure 13 — Behavioral Change → Commercial Outcome

## What We Cannot Measure Externally

This analysis does not have access to Notion's internal cohort data, conversion rates, NRR figures, or lifecycle platform configuration. Where figures are cited, they are directional estimates from industry benchmarks and published commentary, not Notion-specific measurements. The framework defines what should be measured — not what the measurements would show.

# What This Means Beyond Notion

The structural findings in this case study are not unique to Notion. They describe patterns common across product-led SaaS.

## The Product → Lifecycle Fracture

Most PLG companies collect behavioral data during onboarding that their lifecycle systems ignore. The product team and the lifecycle team operate from different data architectures. This is the most common and most expensive structural leak in product-led SaaS. It is also the most solvable.

## The Identity Mismatch Problem

Any product that sells an aspirational identity faces a population of users who want to inhabit that identity but cannot sustain it. The intervention is detection and path reset, not re-engagement.

## The AI Depth Question

Any product adding AI-mediated workflows should ask: does AI deepen the user's relationship with the product, or does it allow the user to extract value without accumulating behavioral investment? The answer determines whether AI improves retention durability or merely front-loads it.

---

The thinking is the product. The company is the case study.

---

ABOUT THIS WORK

# About This Work

This case study is the first in a planned series examining how PLG companies engineer user behavior into revenue outcomes. The series includes Notion, Canva, Figma, Duolingo, and Slack, concluding with a cross-company synthesis.

The analysis is based on primary product observation (two complete signup audits across four months), published behavioral psychology research (Cialdini, Nunes & Dreze, Fogg), SaaS industry benchmark data (OpenView Partners, Amplitude), and a post-publication conversation with Notion's Head of EMEA Marketing. Where this document is uncertain, it says so.

## THE BACK HALF

### Your product already has the behavioral data.

The question is whether your lifecycle system knows what to do with it.

The Notion example demonstrates a broader PLG problem: the product can know what a user is trying to accomplish, what they actually do, and how deeply they are investing. The lifecycle layer often does not reflect that same intelligence. That's where activation leaks begin, identity mismatch goes undetected, and revenue compounds more slowly than the architecture allows.

#### WHAT PERENNUS DOES

- Detect the user's behavioral state from product signals
- Adapt lifecycle interventions to that state
- Build deeper product investment through behavioral orchestration
- Connect behavioral changes to activation, retention, expansion and revenue

**If you think your product knows more about your customers than your lifecycle system does, let's talk.**

Start a conversation: [perennus.com](https://perennus.com)

Isaac Olukoya | Founder, Perennus

[isaac@isaacolukoya.com](mailto:isaac@isaacolukoya.com)