

Product-Led Growth — Workbook

This workbook turns the course into reps you run on a real product. Each section mirrors one course module with hands-on exercises, fill-in worksheets, and checklists you apply to your own software product. Pick one product you own or work on and carry it through every section, and you will finish with a defined North Star and activation metric, an onboarding mapped to the aha moment, a chosen monetisation model, a measured viral loop, and a working product-qualified-lead definition with the math already done.

What Product-Led Growth Really Is

Decide honestly whether product-led fits your product, and map your own flywheel before changing anything.

Worksheet: PLG Fit Scorecard

Answer each question honestly for your product. The more you answer in favour of PLG, the stronger the case; if most point the other way, consider a hybrid or sales-led motion.

Can a brand-new user reach real, standalone value alone, in one session, without a human? (Yes/No + how)

Average contract value / typical price — is it low enough that you cannot afford a salesperson per deal? (amount + Yes/No)

Market shape — is it broad and horizontal with many self-serve buyers, or niche and top-down? (describe)

Cost to serve a free or trial user (compute + support) — can you afford it at scale? (rough cost + Yes/No)

Does normal use expose new people to the product (shared output, invites, links)? (Yes/No + the mechanism)

Verdict: product-led, hybrid (product-led sales), or sales-led — and the single biggest reason

Exercise: Map Your PLG Flywheel

Write the four flywheel stages for your product and the metric that proves each is working, then find the stage that is dragging hardest. That weakest stage becomes your focus for the rest of the workbook.

- Acquisition: how do users arrive today, and which channels are product-generated (shared links, invites, output)?

- Activation: what is the first real value, and roughly what percentage of new users reach it?

- Retention/adoption: what brings users back, and what does week-four retention look like?

- Expansion/advocacy: how do accounts grow (seats, usage, tiers) and how do users refer others?

- Circle the weakest stage — the worst conversion or biggest drop — and state why it is the priority.

Checklist: PLG Mindset Gut Check

- I can name a specific first value a user gets without talking to a human.
- I have an honest verdict (product-led / hybrid / sales-led) based on the fit scorecard, not on aspiration.
- I am thinking in a flywheel — referral and expansion feed acquisition — not a one-way funnel.
- I have identified the single weakest flywheel stage to focus on first.
- I treat onboarding and the invite flow as product I own and instrument, not as marketing add-ons.

Activation: Getting Users to Value Fast

Define your North Star and activation metric from data, then design an onboarding that drives every new user to the aha moment.

Worksheet: Define Your North Star and Activation Metric

Pin down the value-centred numbers your whole team can align on. Base the activation event on a real, value-delivering action, not a vanity count.

North Star metric (the one value-centred number, e.g. messages sent / nights booked / projects shared)

Why this captures delivered value rather than activity (1-2 sentences)

Activation event — the specific first action that signals real value (e.g. shared first file, invited a teammate)

Activation threshold if any (e.g. 3 projects, 2,000 messages) and the evidence behind it

Time-to-value target — how fast a new user should reach activation (minutes / one session / first day)

Exercise: Find Your Aha Moment in the Data

Use your product analytics (Amplitude, Mixpanel, PostHog, or Heap) to find the early action that best separates retained users from churned ones. This is empirical work, not a meeting decision.

- Define retention for your product (e.g. still active in week four) and the cohort you will analyse.
 - List the candidate early actions a user can take in their first session or week.
 - For each action, compare week-four retention of users who did it early vs those who did not — record the gap.
 - Identify the action with the biggest gap and a clear threshold, the way Slack found ~2,000 messages.
 - Sanity-check causation: does this action plausibly deliver value itself, or is it just what power users happen to do?
-

Worksheet: Onboarding Drop-off Map (Bowling Alley)

List the literal steps from sign-up to your activation event and the percentage reaching each one. The steepest gutter is the step to fix first; then add product and conversation bumpers.

Ordered steps from sign-up to activation, each with the percentage of users reaching it

The single steepest drop-off step (your first gutter to close)

Steps you can delete or defer because they are not required to reach first value

Product bumpers to add (empty-state help, template/sample data, progress checklist, tooltips)

Conversation bumpers to add (welcome email, stall-triggered nudge for users who sign up but do not activate)

Checklist: Activation Readiness Check

- My activation metric is a specific event tied to real value, found in data — not invented in a meeting.
- I have a time-to-value target and know my current activation rate.
- My onboarding ends exactly at the aha moment, not at a generic profile-completion step.
- The product is never blank: there is a helpful empty state and at least one template or sample.
- A stall nudge re-engages users who signed up but have not activated within a set window.

Monetisation: Freemium, Trials, and Pricing

Choose freemium, free trial, or a reverse trial; pick a value metric that scales with success; and design contextual upgrade prompts.

Worksheet: Choose Your Monetisation Model

Decide between freemium, free trial (opt-in or opt-out), or a reverse trial, using the benchmarks and your own economics. Make the trade-off explicit.

Model chosen: freemium / opt-in trial / opt-out (card-required) trial / reverse trial

Primary reason (viral exposure from free use / urgency from a clock / cost to serve / desired conversion vs volume)

Conversion rate you are targeting and the benchmark it is anchored to (e.g. freemium ~2-5%, opt-in trial ~15-25%, opt-out trial ~40-60%)

If a trial: length (7 / 14 / 30 days) and whether a credit card is required up front

Cost-to-serve a free/trial user and why your model is affordable at scale

Exercise: Pick a Value Metric and Package the Tiers

Choose the unit you charge for so revenue grows as the customer succeeds, then decide what is free versus paid. Keep the free tier genuinely valuable but bounded.

- List candidate value metrics (seats, projects, contacts, messages, storage, usage volume).
 - For each, does it rise as the customer gets more value? Pick the one that best aligns pay with value.
 - Set free-tier limits on that metric: generous enough to reach the aha moment, tight enough that serious use bumps the ceiling.
 - Define 2-4 paid tiers and what each unlocks across scale, collaboration, capability, and control.
 - Name the two or three natural upgrade triggers users will hit (limit reached, feature needed, team growing).
-

Worksheet: Design Contextual Upgrade Prompts

Specify in-product prompts that ask for money at the moment the user feels value and hits a limit. Frame each around what they unlock, not what they owe.

Upgrade trigger 1 (e.g. hit free project cap) — where it shows and the value-framed message

Upgrade trigger 2 (e.g. inviting the Nth teammate / locked premium feature) — placement and message

Value-recap nudge near end of trial — what accomplishments it summarises

Expansion prompt — when an admin is invited to add seats or move up a tier

Target net revenue retention (NRR) for the account base and how you will track it by cohort

Checklist: Monetisation Sanity Check

- My free experience reaches the aha moment and builds a habit — it is a front door, not a crippled demo.
- My value metric goes up on its own as the customer gets more value (built-in expansion).
- Upgrade prompts fire at the moment of need and are framed around value unlocked.
- Inviting teammates and adding seats is frictionless so accounts can expand themselves.
- I track NRR by cohort and have set a target above 100%.

Loops, PQLs, and Scaling the Engine

Build a measurable viral loop, define and score product-qualified leads, and run a disciplined experiment programme that compounds.

Exercise: Design a Viral Loop and Calculate K-Factor

Choose a viral loop that fits your product, then measure it. Compute K-factor and cycle time so you know what the loop is actually worth.

- Pick the loop type (inherent/collaborative, word-of-mouth, incentivised referral, embed/artefact) and describe the exact mechanic.

- Estimate invites or exposures per active user (i) — e.g. 4 shared links.

- Estimate the conversion rate of those invites to new active users (c) — e.g. 25%.

- Calculate $K = i \times c$, and interpret it (below 1 amplifies other channels; at/above 1 is self-sustaining).

- Estimate cycle time (days from joining to invites converting) and one way to shorten it.

Worksheet: Define Your Product-Qualified Lead (PQL)

Combine activation, engagement, and intent/limit signals into a PQL definition precise enough to act on automatically and predictive enough to justify human follow-up.

Activation signal(s) that must be true (e.g. completed the aha moment, set up the core workflow)

Engagement signal(s) (e.g. returned 3+ days, used across multiple sessions/features)

Expansion/team signal(s) (e.g. invited teammates, multiple active users, usage near a limit)

Intent signal(s) (e.g. viewed pricing, hit a paywall, started an upgrade flow)

PQL score thresholds and routing rule — low-value PQL to automated flow, high-value PQL/account to a human

Exercise: Plan Your First PLG Experiment

Run one disciplined A/B test on your weakest flywheel stage. Size it and set a stopping rule before launch so the result is honest.

- Which flywheel stage are you testing (the weakest one) and the target metric it should move?
 - Hypothesis in if-then-because form, changing only one thing.
 - Sample size and runtime from a calculator (Optimizely / VWO / Evan Miller), and your pre-committed stopping rule.
 - Tool you will run it in (PostHog / GrowthBook / Statsig / Optimizely / VWO).
 - How you will read it: confidence interval (not just headline lift) and the segments you will break it down by.
-

Checklist: Scaling Engine Check

- I have a deliberate viral loop with a measured K-factor and an estimated cycle time.
- My PQL definition combines activation, engagement, and intent/limit signals from real upgrade behaviour.
- Low-value PQLs route to automated flows; high-value PQLs/accounts route to a human (product-led sales).
- I prioritise experiments with ICE and focus on the weakest flywheel stage first.
- I size every test and set a stopping rule before launch, read the confidence interval, and log the insight win or lose.

Your Action Plan

1. Choose one product you own or work on to carry through the whole plan, and write an honest product-led / hybrid / sales-led verdict using the fit scorecard.
2. Map your PLG flywheel (acquisition, activation, retention, expansion) and mark the single weakest stage as your focus.
3. Define a value-centred North Star metric and a specific activation event with a time-to-value target.
4. Use product analytics to find your aha moment empirically — the early action that best separates retained from churned users.
5. Map the steps from sign-up to activation, delete or defer the non-essential ones, and add product and conversation bumpers (Bowling Alley) to close the worst gutter.
6. Choose a monetisation model (freemium / free trial / reverse trial) against the benchmarks and your cost to serve.
7. Pick a value metric that scales with customer success, set free-tier limits, and design 2-4 paid tiers.
8. Design contextual, value-framed upgrade prompts at real limit and feature moments, and set a target NRR above 100%.
9. Design one viral loop, calculate its K-factor and cycle time, and define + score your product-qualified lead with a routing rule.
10. Run one disciplined A/B test on your weakest flywheel stage — sized, with a stopping rule set before launch — read the confidence interval and segments, and log the insight either way.

