
Module 8: Product Launch, Growth, and Monetization

Plan launches, design growth loops, run growth experiments, price and package, reduce churn and scale responsibly.

Product Management Master Handbook
v0.1 · 23 October 2026

Built from the 55-lecture course study guide and extended to the 12-module Product Core curriculum.
Hypothetical examples are labelled; no real company results are invented.

Contents

8.1 Go-to-market planning and launch objectives	3
8.2 Positioning and market segmentation at launch	5
8.3 Beta testing and pilot launches	5
8.4 Product adoption	6
8.5 Growth funnels and growth loops	7
8.6 Growth experimentation	9
8.7 User acquisition	10
8.8 Lifecycle marketing	10
8.9 Product-led growth (PLG)	10
8.10 Retention strategies and referral mechanisms	11
8.11 Pricing, subscriptions and freemium	11
8.12 Conversion optimisation, cross-selling and upselling	12
8.13 Monetisation experiments	12
8.14 Churn reduction	13
8.15 Growth metrics	14
8.16 Case study from your course: ride-hailing surge pricing	14
8.17 Launch retrospectives and scaling	15
8.18 Practical launch plans	16
8.19 Exercises with model answers	18
8.20 Interview perspective	19
8.21 Key takeaways	20
8.22 Check your understanding	21
8.23 Recommended resources	22

MODULE 8

Product Launch, Growth, and Monetization

Plan launches, design growth loops, run growth experiments, price and package, reduce churn and scale responsibly.

≈ 10 h study

About the numbers

All growth, pricing and churn figures are **hypothetical** and chosen for clear arithmetic. They are not benchmarks for any industry or company.

8.1 Go-to-market planning and launch objectives

Go-to-market (GTM) coordinates how the right users discover, adopt and get support for a product: audience, positioning, pricing, channels, sales and support readiness, timing and measurement (Module 2 covers the strategy). A **launch** is one event within GTM — *not the end of product work*.

From your course notes (Lecture 23; Extra topic 10)

Start with the **business case and timing**, define the **audience and proposition**, choose **pricing and explain the rationale**, plan **distribution and launch support**, set measurable success metrics and **adjust the roadmap using feedback after launch**. Pilot an early feature with a small group, train support, prepare a help article, monitor completion, confusion and complaints; **if accuracy fails, pause or roll back; expand only after the team reviews results**. Assign owners for product, engineering, support, analytics, operations and communications; distinguish an internal pilot from a public launch; name a rollback owner. In the cricket-equipment exercise, the instructor pushed learners to consider **partner incentives, sponsorship costs and likely return — not just list channels**.

Launch objectives should be outcome-based and time-bound:

Weak objective	Strong objective
"Launch the subscription"	"Reach 2,000 paying subscribers in 60 days at CAC ≤ ₹600 with refund rate < 3%"
"Get press coverage"	"Drive 15,000 qualified visits from launch content; 8% start a trial"
"Make users happy"	"Activation (first plan created) ≥ 45% within 7 days; CSAT ≥ 4.2/5"

8.1.1 Launch tiers

Not every release deserves a big launch. Tier them:

Tier	What	Example	Effort
Tier 1	New product or major capability; market-facing	New credit card; paid plan	Full GTM: PR, sales enablement, campaigns, support training
Tier 2	Significant feature for existing users	Early fee display; new integrations	In-app announcement, help docs, email to relevant segment
Tier 3	Improvements and fixes	UI polish; performance	Release notes

8.1.2 Product launch checklist

PRODUCT & QUALITY • Acceptance criteria met; QA/UAT signed off • Accessibility check • Performance and load test • Feature flag + kill switch • Rollback plan and owner	DATA & MEASUREMENT • Success metrics and guardrails defined • Events tracked and validated in staging • Dashboard ready; alert thresholds set • Experiment plan (if A/B)	RISK & COMPLIANCE • Security review • Privacy / consent review • Legal review of claims and terms • Fraud and abuse checks
GO-TO-MARKET • Positioning and messaging • Pricing page / plan changes • Launch channels and calendar • Sales enablement deck & FAQs	SUPPORT & OPERATIONS • Help articles • Support macros & training • Escalation path • Ops capacity (e.g., verification queue)	AFTER LAUNCH • Daily monitoring during ramp • Feedback collection • Retrospective scheduled • Decision date: expand / iterate / roll back

Launch checklist (reusable). Adapt by launch tier. The template version is in the Templates library.

8.2 Positioning and market segmentation at launch


Re-use Module 2's positioning statement and choose a **beachhead segment**: a narrow group with acute need, reachable cheaply, whose success creates references for the next segment. Messaging should lead with the job and outcome ("Know your total before you pay"), not the feature ("Fee breakdown component").

8.3 Beta testing and pilot launches

- **Alpha**: internal or very small trusted group; find major bugs.
- **Closed beta**: invited external users; usability, value and reliability; often with feedback channels and NDAs.
- **Open beta**: anyone can join; scale and edge cases.
- **Pilot**: a real launch in a constrained scope (one city, one customer, 5% of traffic) with success criteria and a go/no-go decision.
- **Staged rollout / canary**: gradually increase exposure (1% → 5% → 25% → 100%) while watching guardrails.

From your course notes (Lecture 27)

A/B testing is different from alpha and beta release testing: alpha/beta reveal problems and gather feedback; A/B testing compares variants against a defined outcome with random assignment.




Staged rollout with gates (hypothetical). Each arrow is a gate with explicit criteria (e.g., error rate < 0.5%, no rise in complaints). Anyone on the team can pull the kill switch.

8.4 Product adoption

Everett Rogers' **diffusion of innovations** model segments adopters by timing: **innovators (2.5%), early adopters (13.5%), early majority (34%), late majority (34%), laggards (16%).**

Geoffrey Moore's *Crossing the Chasm* argues the jump from early adopters to the early majority is the hardest because the majority wants proven, complete solutions and references.




Diffusion of innovations — adopter categories (Rogers' model). Rogers' model proportions (based on a normal distribution of adoption time), not measured data for any product. Use it to plan messaging and proof needed at each stage.


Adoption metrics: activation rate, time-to-value, feature adoption rate (users of feature ÷ eligible users), depth (frequency of use), breadth (number of features used), accounts with ≥ N active seats (B2B).

8.5 Growth funnels and growth loops


A **funnel** is linear (acquire → activate → retain → monetise → refer). A **growth loop** is circular: the output of one cycle becomes input to the next, so growth compounds.



Viral (referral) loop. The loop compounds only if enough invited people sign up and activate. Quality of value at step 2 drives everything after it.



Content / SEO loop (user-generated content). Typical of review sites and community Q&A. Content quality and moderation are the guardrails — spam breaks the loop and trust.



Paid acquisition loop. Works only if contribution margin per customer exceeds CAC within an acceptable payback period (Module 7).

8.5.1 Viral coefficient (K-factor)

$K = \text{invites sent per user} \times \text{conversion rate of invites}$
Total users from an initial cohort N ($K < 1$) $\approx N \div (1 - K)$

Worked example

Each new user sends 3 invites; 10% of invites convert $\rightarrow K = 0.3$. From 1,000 seed users: 1,000 + 300 + 90 + 27 + ... $\approx 1,000 \div (1 - 0.3) = \approx 1,429$ users. Virality *amplifies* other channels rather than replacing them unless $K > 1$, which is rare and usually temporary. **Cycle time** matters as much as K : a loop that turns weekly compounds far faster than one that turns quarterly.

8.6 Growth experimentation

Growth teams run a **pipeline** of experiments: generate ideas from funnel data and research $\rightarrow$ prioritise (ICE) $\rightarrow$ design (hypothesis, metric, guardrails, sample size) $\rightarrow$ run $\rightarrow$ analyse $\rightarrow$ document $\rightarrow$ next. Keep a public **experiment log** with results, including failures.

Good growth experiments are tied to a funnel step, have a single primary metric, enough traffic for the effect size (Module 7), and guardrails that protect users.

8.7 User acquisition

Channel	Strengths	Weaknesses	Watch metric
Organic search / content	Compounding, low marginal cost	Slow to build	Organic sign-ups, rankings
App-store optimisation	Captures intent	Competitive; algorithms change	Store conversion rate
Paid social / search	Fast, targetable	Costly; can attract low-intent users	CAC, payback, retention by channel
Referral	High trust, low CAC	Depends on product love	K-factor, referred retention
Partnerships	Access to trusted audiences	Incentive alignment, dependency	Partner-sourced activation
Sales (B2B)	Large contracts	Expensive, slow	Win rate, sales cycle, ACV
Community / events	Loyalty, feedback	Hard to scale	Engagement, referrals

Compare channels on quality, not volume: activation, 30/90-day retention, contribution and payback — not just installs.


8.8 Lifecycle marketing

Messages (email, push, in-app, SMS) timed to the user's stage: **onboarding** (help reach first value), **engagement** (habit-building tips, relevant content), **expansion** (upgrade when they hit limits), **win-back** (lapsed users with a reason to return). Respect consent and frequency caps; measure incrementality with holdout groups — many “opens” would have happened anyway.

8.9 Product-led growth (PLG)

PLG makes the product the main driver of acquisition, conversion and expansion: self-serve sign-up, fast time-to-value, a free tier or trial, in-product upgrade prompts at **moments of value or limits**, and natural virality (sharing, collaboration). The PM's levers: **activation design**,

usage limits that align with value, **pricing page clarity**, **product-qualified leads (PQLs)** for sales-assist.



8.10 Retention strategies and referral mechanisms

Retention strategies: improve activation (the biggest lever for most products); build habits around a real recurring need; deliver ongoing value (new content, personalisation); reduce friction (saved preferences, reorder); proactive support; win-back flows; loyalty programmes that reward genuine value (not just discounts).

Referral mechanisms: double-sided rewards (both get value), referral at moments of delight (after a successful delivery), easy sharing, fraud controls (self-referral, fake accounts), and measuring *referred users' retention*, not just referral counts.

8.11 Pricing, subscriptions and freemium

From your course notes (Lecture 26; Extra topic 5)

Pricing is how much and on what terms; monetisation is the broader way a product earns revenue. Consider value, costs, competition and willingness to pay — not cost alone. Models include subscription, freemium, advertising, affiliate, usage-based, licensing and transaction fees, and they can be combined. **A paid user is not automatically a happy user** — watch activation, repeat use, revenue, refund requests and fairness together. A free trial should disclose when payment begins and how to cancel.

8.11.1 Freemium vs premium (economics)

Worked example — is the free tier paying for itself?

100,000 free monthly active users; 3% convert to paid → **3,000 subscribers** × ₹299 = **₹8,97,000/month** revenue. Serving costs: free users ₹2/month each → ₹2,00,000; paid users ₹20/month each → ₹60,000. **Contribution before acquisition:** ₹8,97,000 – ₹2,60,000 = **₹6,37,000/month**. If conversion fell to 1.5%, revenue halves to ₹4,48,500 and contribution drops to ₹4,48,500 – (₹2,00,000 + ₹30,000) = **₹2,18,500**. The PM's job: find the free/paid boundary where free users experience real value *and* paying users clearly get more.

8.11.2 Pricing psychology and packaging (use ethically)

- **Good–better–best tiers** with a clear “most popular” anchor.
- **Value metric:** charge on the unit that tracks customer value (seats, projects, transactions).
- **Annual discounts** improve cash flow and retention.
- **Avoid dark patterns:** hidden fees, forced continuity without notice, hard-to-cancel flows — they damage trust and may be illegal in some places.

8.12 Conversion optimisation, cross-selling and upselling

- **Conversion optimisation (CRO):** identify the leakiest step (Module 7 funnels), research why, test fixes. Typical checkout fixes: show total early, fewer fields, guest checkout, trust signals, saved payment methods, clear errors.
- **Upsell:** move a customer to a higher tier (more value). **Cross-sell:** add a related product (e.g., add-on dessert, insurance with a card). Make offers *relevant* and *timely*; measure attach rate, AOV, refunds and complaints.

From your course notes (Lecture 31)

In the quick-commerce AOV case, ideas included cart recommendations, transparent discounts on near-expiry items, bundles, thresholds and loyalty rewards — **assess whether each genuinely helps customers rather than encouraging unwanted spending**. Higher AOV alone does not prove better profitability or experience.

8.13 Monetisation experiments

Price tests, packaging tests, paywall placement, trial length, discount offers. **Special care:** price tests can feel unfair if customers see different prices — consider testing on *new* users or by


geography, or test packaging rather than price; check local laws; measure long-term retention and refunds, not just conversion.

Example (hypothetical) — trial length

7-day vs 14-day trial for a SaaS tool, randomised by new sign-up. Primary metric: paid conversion at day 30; guardrails: refund rate, day-60 retention. Result interpretation must wait until day 60 for the 14-day arm — define the analysis window in advance.

8.14 Churn reduction

- **Voluntary churn:** customer chooses to leave (value, price, competitor, need ended). Tools: exit surveys + interviews, save offers that address the reason (pause, downgrade), product fixes for top reasons.
- **Involuntary churn:** payment failures (expired cards, insufficient funds). Tools: **dunning** (smart retries, reminders, card-updater services), grace periods.



Reasons for cancellation (hypothetical exit-survey, 200 responses). Hypothetical. "Payment failed" is involuntary churn — often the cheapest to fix with dunning. "Not using it enough" points back to activation and habit.

8.15 Growth metrics

Net revenue retention (NRR) = (Start MRR + expansion - contraction - churned MRR) ÷ Start MRR

Gross revenue retention (GRR) = (Start MRR - contraction - churned MRR) ÷ Start MRR

SaaS quick ratio = (new MRR + expansion MRR + reactivation MRR) ÷ (churned MRR + contraction MRR)


Worked example

Start MRR ₹100 lakh; expansion ₹12 lakh; contraction ₹3 lakh; churned ₹5 lakh; new ₹15 lakh; reactivation ₹1 lakh. **NRR** = $(100 + 12 - 3 - 5) \div 100 = 104\%$ — existing customers grew revenue even after churn. **GRR** = $(100 - 3 - 5) \div 100 = 92\%$. **Quick ratio** = $(15 + 12 + 1) \div (5 + 3) = 28 \div 8 = 3.5$ — revenue added is 3.5× revenue lost.

8.16 Case study from your course: ride-hailing surge pricing

From your course notes (Lecture 15)

Design hypothetical surge-pricing logic with an AI component. A useful goal balances **availability, driver earnings, affordability and trust**. Groups considered drivers, requests, location, time, weather, festivals and events; distinguish the ordinary fare from a temporary surcharge; make assumptions explicit. Ideas: use historical demand and weather to anticipate busy periods, alert customers and help drivers move toward demand. **Test price changes with guardrails**: measure ride completion, cancellations, driver earnings, affordability and trust; price caps and local rules may matter. *A participant's idea of charging more based on low phone battery or repeat booking behaviour is an unverified proposal — exploiting urgency or personal data undermines trust and raises privacy, fairness and legal concerns.*



Surge pricing as a balancing loop (conceptual). Conceptual. The loop rebalances supply and demand, but each turn has a trust cost. Guardrails: cap multipliers, show the fare before booking, monitor cancellations and complaints.

8.17 Launch retrospectives and scaling

Retrospective (blameless): goals vs results (with data), what went well, what didn't, surprises, decisions to keep/change, owners and dates. Hold it 2–4 weeks after launch and publish it.

Scaling a product: (1) prove retention in the beachhead before expanding; (2) industrialise operations (support, onboarding, fraud); (3) harden reliability and observability (Module 9); (4) localise (language, payments, regulation); (5) expand segments and channels one at a time; (6) protect unit economics as you grow.

From your course notes (Extra topic 13)

An incident is an event that disrupts a service or harms users. Define alert thresholds and owners before release; track failed orders and recovery; maintain an incident timeline; follow up with actionable prevention items. **Never assume an incident is over merely because the dashboard looks better for a few minutes.**

8.18 Practical launch plans

8.18.1 Fintech: launching a digital credit card application (hypothetical)

Phase	Key activities	Gate criteria
Readiness	Regulatory and compliance approvals; KYC vendor integration; credit-risk policy; fraud rules; disclosures and consent copy reviewed	Compliance sign-off; security and privacy review complete
Employee pilot	Staff apply with real KYC; manual review of every decision	No critical defects; approval decisions match policy
Limited pilot	2% of eligible traffic; staged by region; human review of edge cases	Completion $\geq$ target; fraud rate within limit; complaint rate stable
Scale	Increase traffic weekly; marketing campaigns; partner channels	Unit economics within plan; ops queue SLA met
Optimise	A/B test document-upload UX, pre-fill, status tracking	Experiments show completion lift without risk increase
Metrics: application start → submit → approve → activate (first transaction); time to decision; drop-off by step; fraud rate; cost per approved account; complaints. Guardrails: credit losses, fraud, regulatory breaches.		

8.18.2 B2B SaaS: launching a paid tier (hypothetical)

Phase	Activities
Beta	20 design-partner accounts; usage telemetry; weekly feedback calls; pricing interviews
Pricing	Value metric choice (seats vs projects); good–better–best tiers; pricing page test
GA launch	Docs and help centre; in-app upgrade prompts at limits; sales enablement for large accounts; webinar
Post-launch	Conversion funnel by plan; expansion; churn reasons; NRR tracking
Metrics: trial → paid conversion, time-to-value, PQL → won rate, expansion MRR, NRR, support tickets per account.	

8.18.3 Consumer app: launching in a new city (hypothetical)

Phase	Activities
Pre-launch	Supply onboarding (restaurants, couriers) to minimum density; localisation; payment methods
Soft launch	Invite-only waitlist; referral codes; monitor delivery times and cancellations
Public launch	Local partnerships, app-store optimisation, targeted ads with CAC caps
Scale	Expand zones only where wait times and contribution meet targets
Metrics: first-order conversion, repeat within 14 days, delivery time, cancellation rate, contribution per order, CAC payback. Guardrails: rider/courier earnings, complaints, safety incidents.	

8.19 Exercises with model answers

Exercise 1 · K-factor

Users send 2 invites each; 25% convert. Seed = 5,000. Estimate total users after the loop exhausts. How would halving cycle time change things?

Model answer

$K = 2 \times 0.25 = 0.5 \rightarrow \text{total} \approx 5,000 \div (1 - 0.5) = \mathbf{10,000}$. Halving cycle time doesn't change the eventual total when $K < 1$, but you reach it twice as fast — which matters for compounding with other channels and for time-bound goals.

Exercise 2 · NRR and quick ratio

Start MRR \$200k; new \$30k; expansion \$20k; contraction \$10k; churn \$14k; reactivation \$4k. Compute NRR, GRR, quick ratio.

Model answer

$NRR = (200 + 20 - 10 - 14) \div 200 = 196 \div 200 = \mathbf{98\%}$. $GRR = (200 - 10 - 14) \div 200 = \mathbf{88\%}$. Quick ratio = $(30 + 20 + 4) \div (14 + 10) = 54 \div 24 = \mathbf{2.25}$. Growth relies on new logos; existing customers shrink slightly — focus on churn reasons and expansion paths.

Exercise 3 · Launch plan

Draft a one-page launch plan for “early fee display” (Tier 2).

Model answer outline

Objective: +2 pts checkout completion, complaints $\leq 1.3\%$. Audience: all app users (A/B 50/50, then ramp). Readiness: flag, events validated, help article, support macros, fee-service fallback tested. Rollout: 1% internal → 5% canary → A/B 14 days → 100% with 5% holdout. Comms: in-app tooltip; no press. Monitoring: completion, payment errors, complaints daily. Rollback owner: Eng lead. Retro: 3 weeks after ramp.

8.20 Interview perspective

“How would you grow a product's revenue by 20%?” — structure

Decompose revenue: users $\times$ conversion $\times$ ARPPU (or orders $\times$ AOV $\times$ take rate). Find the biggest, most fixable lever with data (funnel, cohorts, segments). Generate options per lever (acquisition, activation, retention, pricing/packaging, upsell/cross-sell). Prioritise by impact, confidence, effort and risk. Set metrics and guardrails (refunds, complaints, churn). Sequence experiments. Mention unit economics: growth that loses money per unit is not success.

“Our churn went up. What do you do?” — 60-second answer

“First, confirm it's real and which churn — voluntary or payment failures — and for which cohorts and plans. Payment-failure churn can often be cut quickly with smarter retries and reminders. For voluntary churn, I'd segment by reason using exit surveys and talk to recently churned customers, then check whether activation or a recent change explains it. I'd pick the biggest addressable reason, test a fix — product change, better onboarding, or a pause/downgrade option — and track churn by cohort plus guardrails like refunds and complaints.”

8.21 Key takeaways

- Launch = one step in GTM; **tier** launches and **gate** rollouts with criteria and a kill switch.
- **Growth loops compound**; funnels don't. K-factor < 1 amplifies, not replaces, other channels.
- Compare channels on **quality and payback**, not volume.
- **PLG** depends on time-to-value and upgrade prompts at moments of value.
- Model **freemium economics** explicitly; free users cost money.
- Distinguish **voluntary vs involuntary churn**; track **NRR, GRR and quick ratio**.
- Monetise **fairly** — avoid dark patterns; watch refunds and trust.

8.22 Check your understanding

CHECK YOUR UNDERSTANDING

Users send 4 invites each and 5% convert. What is K?

- A. 0.2
- B. 0.8
- C. 4.05
- D. 20

Why: $K = \text{invites per user} \times \text{conversion} = 4 \times 0.05 = 0.2$.

NRR above 100% means...

- A. No customers churned
- B. Revenue from existing customers grew after expansion, contraction and churn
- C. New customers grew fast
- D. Prices were raised

Why: NRR measures revenue from the starting cohort including expansion, contraction and churn; it excludes new customers.

Which is involuntary churn?

- A. Customer switches to a competitor
- B. Card expires and renewal payment fails
- C. Customer finds it too expensive
- D. Customer no longer needs it

Why: Involuntary churn comes from payment failures rather than a choice to leave.

A staged rollout primarily helps you...

- A. Avoid measuring results
- B. Limit blast radius and check guardrails before full exposure
- C. Skip QA
- D. Run a price test

Why: Gradual exposure with gates and a kill switch limits damage from unexpected problems.

In PLG, the best moment for an upgrade prompt is...

- A. Immediately at sign-up
- B. At a random time
- C. When the user hits a limit after getting real value
- D. Never

Why: Upgrade prompts convert best at moments of demonstrated value or need.

8.23 Recommended resources

- [Reforge: Growth loops are the new funnels](#).
- [OpenView: Product-led growth resources](#).
- Geoffrey Moore, *Crossing the Chasm* — [overview \(Wikipedia\)](#).
- [Stripe: Reduce involuntary churn \(Billing docs\)](#).
- [Lenny's Newsletter](#) — growth and pricing essays.