Product Portfolio

Alexandr Petrov
B2C SaaS · Monetization & Growth
Email: serverwind.dev@gmail.com · Telegram: @lestrae · LinkedIn: linkedin.com/in/alexandr-petrov-sw · Portfolio: serverwind.dev

---

> This portfolio is not a list of achievements. It is a record of product decisions: what I saw, what I considered, what I chose, what I shipped, and what I learned.

Product context: an employer-owned B2C SaaS for seafarer certification. I managed the product; the company owns it. The numbers come from internal dashboards. Where data is missing, I say so.

Why these five cases: they show different PM muscles — business model, discovery, monetization, zero-to-one execution, and a product bet that failed. I did not choose them by the biggest number alone.

Case 1 — Turnaround: from $100 to ~$2,171/month

Context

A B2C SaaS for seafarers: test preparation and technical reference content in English. About 30,000 monthly visitors, a large question bank, and a small content team. The product was free and monetized only by context ads.

In January 2026, the owner was ready to shut it down. Revenue was about $100/month and did not cover costs. The team was at risk. I took over the product with no external PM guidance.

Problem

The product had an audience but no business model. Free content attracted traffic, but ads could not cover the cost of content and engineering. Users came to pass certification exams; the product gave them content, but not a reason to pay.

The question was not “which feature should we build?” It was: **can this product become a subscription business without destroying the organic traffic that keeps it alive?**

Evidence

- Revenue: ~$100/month from context ads.
- Audience: ~30,000 monthly visitors; no paid acquisition.
- Monetization: none — no paywall, no subscription, no paid archives.
- User behavior: support conversations and search behavior showed that users came for test questions and answers, not for ads or generic articles.
- Cost base: a small content team and hosting; the product was not covering its costs.

Options considered

1. **Keep ads and cut costs.** Reduce the team and accept a small ad-supported product.
2. **Sell one-off PDFs only.** Low-ticket, no recurring revenue.
3. **Put everything behind a paywall.** Maximum short-term revenue, but high risk to SEO and to users who had never paid.
4. **Freemium subscription plus paid PDF archives.** Free value first, paid depth at the moment of need.
5. **Paid acquisition.** Spend on ads to grow out of the problem.

Decision

I chose **freemium**: free content and a calibrated free tier, with premium subscriptions and paid PDF archives. No paid acquisition.

Why:

- Users do not pay for a pig in a poke. They need to see value first.
- The free tier had to be designed around exam behavior: **5 free questions** and **10 free searches per day**. That is enough to feel the product, but it runs out at the moment when a user needs more.
- A full paywall would have damaged the organic traffic that made the product valuable.
- Paid acquisition was not affordable and would not fix a broken business model.
- PDF archives gave a low-ticket entry point; subscriptions gave recurring revenue.

Execution

- Negotiated a two-month reprieve with the owner: the closure was postponed from March to May.
- Rebuilt the monetization: premium subscriptions, PDF archives, limited free search.
- Implemented an SEO-safe paywall with schema markup. The goal was to monetize without triggering a Google cloaking penalty; organic visibility and bounce rate stayed stable.
- Added payment rails for markets where PayPal was not enough: crypto for Turkey, Anypay for Russia.
- Handled unpaid orders personally: 1–2 of 5–7 daily unpaid orders were converted through direct follow-up; about 15 clients were brought to sale in the first program. Later trained a team member to run the same playbook.
- Led a small team: three editors at the start; later two editors and an SMM trainee. One editor left voluntarily; no layoffs during the turnaround.
- First sale: March 27, 2026, $10 via PayPal.

Sidebar: the card-payment outage

In August 2026, users were writing “I can’t pay by card.” The “Debit or Credit card” button opened a grey, unclickable form in every browser; PayPal worked. Support and the team initially assumed user error.

I checked the network requests: our side was clean. I disabled our custom scripts: not the cause. The root cause was an open bug in the WooCommerce PayPal Payments plugin (GitHub issue #4524, known since July 14, 2026). The plugin’s script was blocking the card form.

I did not replace the plugin — that would have been a large, risky migration for a problem that could be isolated. I shipped a CSS hotfix that unblocked the card fields.

After the fix, card payments became a major payment path: roughly **40–50% card, ~40% PayPal, the rest crypto/Anypay**. The business effect was not “a CSS fix”; it was a restored revenue path.

Result

Metric	Before	After
Monthly revenue	~$100	~$967 by July 2026; ~$2,171 total revenue by October 2026
Subscription revenue	$0	$879/month (October 2026)
Gross margin	Not meaningful	54.3%
Profit	Negative / not covering costs	~$1,179/month (October 2026)
Paid acquisition	None	None
Monthly traffic	~30,000 visitors (baseline)	~43,804 visits in October 2026; 52.5% organic
Team	At risk of closure	Retained; one voluntary departure

Total revenue includes subscriptions, PDF archives, lifetime sales, and advertising/partner placements. It is not MRR. Subscription revenue was $879/month; the rest was one-off and partner revenue.

The product moved from “shutdown planned” to a sustainable subscription business. The owner stopped talking about closing it.

What I did not have: repeat-purchase tracking, cohort retention, and LTV. I instrumented what I could and made decisions with the data I had. Building a proper metrics layer was the next step, not something I pretend I already had.

What I learned

- A business model is a product decision, not a finance afterthought.
- The free tier is the product. If it is not designed around the moment of willingness to pay, the paywall is just a wall.
- SEO and monetization can coexist, but only if the paywall is implemented with the search crawler in mind.
- The hardest part was not building the paywall. It was deciding what **not** to give away.
- What I would do differently: instrument repeat purchases and cohorts earlier; run pricing experiments before scaling; set kill criteria for new products before they start.

What this shows

P&L ownership · monetization · pricing · execution · technical fluency.

Case 2 — The 31-account study: from a killed hypothesis to +79% premium conversion

Context

By September 2026, premium users were buying the product but, from the outside, seemed not to use it. The team had a comfortable explanation: maybe users buy and forget, or maybe something is technically broken.

I had a stronger belief: **global search across questions and answers was the core value of the product.** The roadmap was being shaped around that belief.

Problem

If premium users do not use the product, retention and revenue are at risk. But the bigger risk was different: we might be investing in the wrong value proposition.

The question was: **what do premium users actually use — and what does that mean for the roadmap?**

Evidence

I audited 31 premium accounts with subscriptions expiring in less than 15 days. I checked what each account actually did.

Behavior	Accounts
Used global search more than 10 times	9
Did not use global search more than 10 times, but actively used test pages	19
Bought premium and never logged in	2
Bought premium and logged in once	1

Key findings:

- **61% (19/31)** did not use global search more than 10 times.
- **28 of 31** were actively using the product. The “technical problem” hypothesis was wrong.
- Users were not rejecting the product. They were using **test pages**, not search.
- Only **2 of 31 (6%)** bought and never logged in — a small edge case, not a trend.

Options considered

1. **Keep investing in global search and AI.** It was our belief; maybe users just needed better onboarding.
2. **Improve onboarding and communication.** Explain the search feature better.
3. **Redesign the test pages.** Users already preferred them; the design was old.
4. **Contact the non-users and ask why.** A small qualitative check.
5. **Do nothing.** The product was working; leave it alone.

Decision

I stopped further development of global search and made **test-page redesign the #1 priority**.

I also kept global search as a supporting feature, not the core value. I asked the team to contact the two non-users to understand whether they needed help or a refund.

Why:

- The evidence was behavioral, not opinion-based.
- 19 of 31 users were actively using test pages; that was the real product.
- Continuing to invest in search would have been investing in my own belief, not in user value.
- The test pages had an old design, no upsell, and no modern results screen. That was a clear product opportunity.

Execution

- Shared the findings with the owner and team. The owner trusted the analysis; there was no resistance.
- Paused further search development and deprioritized other initiatives, including Marlins expansion and AI search.
- Moved the team to the test-page redesign:
  - test start button on the first screen;
  - clearer test structure;
  - upsell to buy the test as a PDF;
  - certificate after passing (premium required);
  - better results screen.
- The redesign shipped in October 2026. I tracked bounce rate as a diagnostic (baseline 9%), but the commercial metrics were test starts, premium conversion, and PDF upsell.

Result

The result was a **roadmap change and a measurable lift**:

- Global search was no longer the center of the roadmap.
- The team redesigned the test experience around what users actually used.
- After the redesign, organic-to-premium conversion rose **79%**, organic-to-PDF conversion rose **47%**, and sales rose **23%**.
- The lift held for two weeks. This is a before/after result, not a controlled A/B, so I do not claim pure causality; seasonality and other changes were not isolated.

The important result is that the decision changed before more money was spent on the wrong feature — and the commercial metrics moved in the right direction after the change.

What I learned

- The most dangerous assumption is the one you do not test.
- Users' behavior is more honest than their words — and more honest than my own beliefs.
- A PM's job is to kill their own ideas with evidence, not to defend them.
- A small sample can be a signal, not statistics. I treated it as a signal and kept the decision reversible.

- A before/after lift is enough to act, but not enough to claim causality. The next step would be a controlled test.
- What I would do differently: instrument behavioral analytics earlier and define the success metric for the redesign before launch, not after.

What this shows

Discovery · data literacy · prioritization · intellectual honesty.

Case 3 — Pricing experiments: the $1/day failure and the 5→3 reset

Context

The premium product had five plans, including 1 month, 3 months, 6 months, 12 months, and lifetime. PDF archives were priced by intuition rather than by a pricing model.

The team wanted to test a low-ticket entry point to reduce the barrier to first purchase.

Problem

Conversion was limited, and the pricing page was confusing. But a cheap plan could also cannibalize the main subscription. The question was: **can a micro-plan grow revenue, or does it just move existing demand to a cheaper product?**

Evidence

- 6- and 12-month plans produced almost no sales: **one sale in August**.
- Most sales came from the 1-month and 3-month plans; lifetime was occasional.
- The $1/day experiment:
  - 10 purchases;
  - PayPal took roughly 30% on each micro-payment, leaving about $0.60 net per sale;
  - monthly subscription sales fell to **1** during the experiment, versus the usual **3–4**;
  - net revenue was lower than before the experiment.
- AI search was built, but it fired rarely and users barely saw it. Its cost was under 2% of revenue.

Options considered

1. **Keep all five plans.** Avoid change; accept complexity.
2. **Create a separate Premium+ plan for AI.** A new SKU with a higher price.
3. **Put AI only in 3+ month plans.** Use it as an upsell to longer commitments.
4. **Three plans, AI included, with limits.** Simplify the choice and let AI support conversion.
5. **Keep the $1/day plan.** Use it as a low-ticket entry.

Decision

- **Killed the $1/day plan.** The unit economics did not work, and it cannibalized the monthly subscription.
- **Reduced five plans to three:** 1 month ($10.99), 3 months ($25.99), lifetime ($99.99).

- **Removed 6- and 12-month plans.** They did not sell and added confusion.
- **Included AI in all paid plans** instead of creating a separate Premium+ SKU. AI was an amplifier, not a product: it fired rarely and cost little.
- **Did not raise the monthly price yet.** The team needed conversion data after the plan change before adding another variable.

Execution

- Rolled back the $1/day experiment.
- Rebuilt the pricing page with three plans, a clearer value description, and price-per-month framing.
- Added AI limits to the 1-month plan (20 answers/day) and kept it unlimited on 3-month and lifetime.
- Planned to measure conversion, plan mix, and AI usage in October before changing the price again.

Result

- The $1/day hypothesis was rejected: fees and cannibalization made it worse than the standard subscription.
- The 5→3 simplification shipped. October data was not final at the time of writing.
- AI was launched as a feature inside paid plans, not as a separate product. Usage and cost were being measured.

This case is included because the **decision quality** matters more than a clean win. The failed experiment produced a clear rule: do not launch a micro-plan before calculating payment fees and cannibalization.

What I learned

- Micro-payments can die on payment fees before they ever reach product-market fit.
- A cheap plan can cannibalize the core subscription even when it looks like “growth.”
- Fewer plans reduce decision fatigue; pricing is part of the product experience.
- Do not create a new SKU before you know the feature is visible and valuable.
- What I would do differently: calculate unit economics before the experiment, define a kill criterion, and test with a small cohort instead of launching to everyone.

What this shows

Monetization · unit economics · experimentation · willingness to kill.

Case 4 — App 0→1 in three weeks

Context

Users asked for a mobile app. The website was built on WordPress and was heavy with content; a PWA was not a realistic path. The product’s core value — tests and answers — was a good fit for a focused mobile experience.

There was also a distribution hope: Google Play could bring a new acquisition channel.

Problem

Mobile users needed a focused way to practice, without the weight of the full website. The team was small, with no native mobile team. The question was: **what is the smallest app that can validate the value, and how do we ship it fast without breaking the existing product?**

Evidence

- Users explicitly asked for an app.
- Competitor research showed the standard feature set and monetization patterns.
- The site was overloaded with content; the app could be focused on tests.
- The existing premium subscription could be reused through access keys.
- Beta results: 7 sign-ups, 4 active (57%), 215 sessions, D3 retention 50%, 0 crashes, 10 bug reports — mostly content, not app stability.
- Legacy premium keys: only 10 of 100 active premium users activated app access. Older buyers were less likely to activate; many had already passed their exams and were at sea.

Options considered

1. **Build a PWA.** Reuse the web stack, but the experience would still be heavy.
2. **Build a native app.** Best experience, but slower and more expensive.
3. **Wait for Google Play.** Avoid alternative distribution, but delay learning.
4. **Build a focused MVP and release on alternative Android stores first.** Ship fast, validate stability, handle Google Play later.
5. **Build a full-featured app.** More scope, more risk, slower feedback.

Decision

I chose the focused MVP:

- tests, answers, and premium access via key;
- simple, fast UI;
- no extra content or social features;
- release on alternative Android stores first;
- use the beta to validate stability and content quality, not growth;
- keep monetization through the existing premium subscription, not a separate app subscription.

Execution

- Sep 5: competitor research and MVP feature list.
- Sep 7: development start.
- Sep 14: content and SMM preparation; editor tasks for SEO/AEO; beta preparation.
- Sep 21–24: beta with 7 users.
- Sep 25: release candidate.
- Sep 28: landing page and release.

- Sep 30: added to other Android stores; Google Play still pending.
- Generated keys manually for legacy premium users and handed them to SMM for distribution.

Result

- MVP shipped in about three weeks.
- Beta: 0 crashes, D3 retention 50% (n=4), positive feedback on speed and simplicity.
- 10 of 100 active premium users activated app access.
- Adoption was concentrated among recent buyers; older buyers often did not need the app because they had already passed their exams.
- The main constraint was **distribution**, not the app itself: without Google Play and with .apk friction, the app could not become a growth channel yet.

What I learned

- Shipping the app was the easy part. Distribution and activation are the product.
- A small beta is a signal, not statistics. I used it to validate stability, not to claim product-market fit.
- Legacy users may not need the app at the moment you launch it; timing matters.
- What I would do differently: solve Google Play distribution before building; instrument activation and retention from day one; define the app's job-to-be-done (practice on the go) and measure it explicitly.

What this shows

Zero-to-one execution · MVP scoping · beta · cross-functional work · honesty about distribution.

Case 5 — PSC Guide: a product bet that sold zero

Context

The owner wanted a PSC Inspection Guide for gas carriers. The author provided 143 inspection items across 6 modules. The target audience was narrow: officers on board during PSC inspections.

I told the owner upfront that I did not believe the niche was large enough. He wanted it built anyway. The go/no-go decision was his, not mine.

Problem

The product needed 10–15 sales to recoup development costs. The question was not only "can we build it?" but **"is there enough demand, and what is the right form for this product?"**

A plain text manual would be a commodity. An interactive tool could be more useful — but only if the audience trusted it and could find it.

Evidence and assumptions

- Audience: officers on board during inspections — narrow, high-stakes, not mass-market.
- Author credibility: a chief officer's assistant; his video/explanation could be a trust signal.
- No demand validation before building.

- No paid acquisition planned.
- Break-even: 10–15 sales.
- The author did not provide the launch assets we asked for: a video address for the landing page and short explanatory clips.

Options considered

1. **Do not build it.** My preferred option: the niche was too narrow.
2. **Build a plain PDF/manual.** Cheapest, but a commodity.
3. **Build an interactive checklist tool.** Better product, more development.
4. **Use freemium instead of a one-off purchase.** Lower barrier, but no recurring value yet.
5. **Give it away free first to validate demand.** Learn before charging.

Decision

- Build the **interactive tool**, not a plain manual: checkboxes, progress, modules, a “Pass or Detain” test.
- Ship on time at the end of August.
- Use a one-off purchase model.
- Ask the author for video and launch assets to build trust.

The product decision was mine; the go/no-go was the owner’s. I disagreed with the bet but executed it properly.

Execution

- Designed the tool and its modules around the inspection workflow.
- Built checkboxes so an officer could walk the ship and tick each item.
- Added progress tracking and the final “Pass or Detain” test.
- Launched with a landing page.
- The author did not deliver the video/explainers; distribution and trust were weaker than planned.
- No paid acquisition.

Result

- **0 sales in September and October.**
- Break-even was 10–15 sales; it was not reached.
- The tool itself was good; the product did not find demand or distribution.
- The owner’s decision rights meant the project went ahead despite my objection.

What I learned

- A good tool without demand or distribution is a hobby, not a product.
- For a narrow, high-trust product, the author’s credibility and launch assets are part of the product.
- Saying no to the owner is not enough. You need a decision framework and a pre-agreed kill criterion.
- I should have tested demand cheaper before building: a landing page, a waitlist, a free pilot, or a short paid pre-order.

- What I would do differently: set a 60-day kill criterion, validate demand before development, and use freemium or a paid pilot instead of a one-off purchase.

What this shows

Stakeholder management · product judgment · failure learning · honesty.

Other experiments

These are smaller, but they show how I work.

Checkout revamp: when the fix did not move the metric

- 1,074 sessions on /checkout without a purchase.
- 67% high-intent users; 23.5% quick backs; 32.5% dead clicks; 58.6% scroll depth.
- Dead clicks were inflated by third-party payment iframes; the data had to be interpreted, not taken literally.
- I changed the express checkout to show only card first, with PayPal/crypto below.
- The metric barely moved.
- The lesson: the problem was not the form layout, but the moment users were pushed into checkout. I stopped investing in the form instead of forcing a fix.

Question-bank audit: trust is a product feature

- 250 questions sampled from CES Engine Management.
- 18 findings: 4 almost certainly wrong, 8 questionable, 6 missing images.
- AI found candidates; editors verified and fixed them manually.
- Target: reduce error reports from 15–20/month to 2–3/month.
- The principle: we must be a source of truth. Errors damage trust more than outdated content.

Registration as a lead magnet

- Users asked: “How do I buy premium? I don’t see registration.”
- Added open registration with a lead magnet: anonymous users get 5 searches/day, registered users 10, plus a 15% first-purchase coupon.
- 5 registrations on the first day; target was 20/month.
- Early signal, not a proven result.

What this portfolio shows

PM muscle	Case
Business model and P&L	Case 1
Discovery and prioritization	Case 2

PM muscle	Case
Pricing and unit economics	Case 3
Zero-to-one execution	Case 4
Stakeholder management and failure	Case 5
Technical fluency	Case 1 sidebar
Data literacy and stopping	Other experiments

What I am looking for

A B2C SaaS where the PM owns outcomes, not just a backlog — especially in subscription, monetization, or growth. Seed to Series A, remote.

I am not looking for a developer role, a delivery-manager role, or a “technical PM” role that is really project management. I am looking for a product role where technical depth is useful, customer closeness matters, and the PM is accountable for the result.

© 2026 Alexandr Petrov — Product Portfolio