

AI Document–Report Businesses

Which one to build first, from a screened database of 335 candidates

THE QUESTION

Which business should a solo operator launch with **under \$5,000**, 30 days to first revenue, strong in-house design and paid-ads capability, and no licensed credentials? Every candidate shares one shape: **upload documents, get a valuable report back, pay once with a guarantee, subscription optional**.

335

CANDIDATES BUILT

18

INDUSTRIES SWEEPED

263

KILLED ON EVIDENCE

3,350

SCORED JUDGMENTS

100%

EVIDENCE-VERIFIED

7 wks

OF ANALYST TIME

HOW CANDIDATES WERE JUDGED

Each idea scores 1–5 on ten weighted criteria — profitability, upfront cost, order value, demand growth, underservice, social proof, ad reachability, word-of-mouth, subscription fit and speed of cheap invalidation — then takes a penalty for regulatory risk. Before scoring, every idea faces five kill-filters. **The most lethal by far**: is this deliverable already given away free as a sales lead-magnet? That one filter killed 26 ideas, four of them previously top-ten. You cannot charge for what competitors use as bait.

This runs in cycles, and that is the point. An automated loop rebuilds the database weekly from regulatory calendars and enforcement actions, re-screening what it already holds. Two cycles in, the honest result is not a higher ceiling — cycle one still holds the best idea. What cycling buys is **correction**: cycle two found a funded competitor in a gap cycle one called empty, retracted a regulatory deadline that never arrived, and killed 9 of the previous top 20. Desk research decays; a static report would still be repeating both errors.

How every business was graded

Filled squares = 1 to 5

THE TEN QUESTIONS ASKED OF ALL 335 CANDIDATES	WEIGHT	CONSTRUCTION TAKEOFF	HOA REVIEW	GRANT SCORING
Will it make money quickly?	15%			
Can it start cheaply?	10%			
Is each sale worth enough?	13%			
Is demand big and growing?	12%			
Is the need underserved?	12%			
Do buyers complain publicly?	7%			
Can we reach buyers with ads?	11%			
Will customers refer others?	9%			
Does a subscription fit?	5%			
Can we test it cheaply?	6%			
Legal or licensing risk (lower is better)	penalty	1 of 5	2 of 5	1 of 5
Final score out of 100		91.8	84.5	83.6

How to read this: every candidate was scored 1 to 5 on all ten questions, combined by weight, then penalised for legal risk. **Notice the pattern:** construction wins on money and reach but is weakest on underservice — a well-funded competitor already exists. Grant scoring is the reverse: the strongest untouched gap, but slower to get paid.

THE THREE TO BUILD



Construction Takeoff & Estimate

91.8
SCORE

Buyer: specialty trade contractors & small GCs · **\$249–\$1,499** per report · \$299–\$799/mo · regulatory risk 1/5

Upload a plan set, get a priced quantity takeoff back within hours. **92% of US construction firms cannot find enough workers, and estimators rank among the scarcest roles.** Margin pressure forces contractors to bid *more* jobs for the same volume, so demand rises as the industry gets harder. Purchase repeats weekly.

The wedge: Beam AI sells guaranteed done-for-you takeoffs at roughly \$8,000/yr per trade — priced for constant bidders. But 59% of specialty trades are under \$1M revenue and bid occasionally. Their only options are annual enterprise commitments or offshore shops. No credible per-report option exists.

Main risk: accuracy is the entire product, and the operator has no construction domain knowledge. Requires a retired estimator on QA from day one.



HOA / Condo Document Review

84.5
SCORE

Buyer: condo buyer under contract · **\$149–\$299** per review · no subscription · regulatory risk 2/5

Upload the HOA package the seller is obliged to hand over — budget, reserve study, minutes, bylaws, litigation disclosures — and get back, inside 24 hours, the special-assessment risk, the deferred maintenance, the rules that bite, and the questions to ask before the review window closes.

The wedge: the buyer has days, not weeks, and the downside is a \$5,000–\$100,000+ assessment. Agents will not read the package and attorneys charge by the hour to. The trigger is sharp, the price is small against the purchase, and every agent who has seen one saved deal refers the next buyer. Deliberately the opposite bet to pick #1: consumer, high volume, low ticket.

Main risk: low order value leaves little room for an expensive click — it works only if the ads stay cheap, which is what the Stage 1 test measures first.

Grant Readiness

Scoring

83.6
SCORE

Buyer: small nonprofits · **\$299–\$1,499** per review · \$199/mo · regulatory risk 1/5

Upload a draft application and the funder's RFP; get back a simulated reviewer score against that funder's own rubric, every weak section flagged — **before** submitting.

The wedge: grant writers charge \$1,500–\$5,000 per foundation proposal, \$5,000–\$15,000 federal — so a \$399 pre-submission check is easily justified. Existing tools either *write* narrative (Grantable) or *find* funders (Instrumentl). Downgraded in R003: a consultancy gives away a mock-review toolkit as a lead magnet, and universities and SBIR programmes score drafts free for their own people. What is left is the nonprofit with neither behind it.

Main risk: nonprofits pay slowly and the product's value is unprovable for 6–9 months, so no results-based case study is possible. Deferred from the first test round on cycle time, not on quality.

CONFIDENCE — READ HONESTLY

- **Every row has been screened against the kill-filters.** 38 carry sourced competitor pricing and a demand statistic; the rest are screened priors.
- **No idea has full verification.** None has real cost-per-click or search-volume data attached.
- **Score gaps under ~5 points are noise.** Verification has moved ideas by as much as 14.6 points.
- **Zero customer contact so far.** This is entirely desk research.

Next steps, in priority order: (1) launch the \$150 Stage 1 signal test on Construction Takeoff — the Status page names it as the single next action; (2) run HOA, CAM and Supplement Substantiation alongside it, the four being deliberately different kinds of bet; (3) advance whatever clears the thresholds to a \$350 Stage 2 with a real checkout, measuring the only signal that really matters — whether anyone buys, and then buys twice. The FTC Safeguards WISP that held this page's third slot was killed on 26 August 2026 when its enforcement check failed.

The Next 7

Ideas that survived every kill-filter but lost the top three on score

These are **not** rejects — every one cleared all five kill-filters. They sit just below the top three, and the gap is often smaller than the margin of error in the scoring. One is the natural second product for pick #1. The rest of the active set sits below these in the Leaderboard.

Score = weighted rubric result out of 100 after the risk penalty. **Risk 1** = no licensure exposure, 5 = effectively requires a licence. **Tier** = evidence level: **T0** | unverified prior | **T1** | screened against the kill-filters | **T2** | competitor pricing and demand statistic sourced | **T3** | adds live ad-cost data — nothing has reached T3 yet. Excluded here: every candidate killed on evidence. Nothing is unscreened. Full database on the site.

Construction & trades — the same buyer as pick #1

Bid Leveling & Comparison Report 87.6

GC or building owner · \$199–\$899 · \$299/mo · risk 1/5 **T2**

Owners comparing 4–8 subcontractor bids have no cheap way to normalise scope. Same upload motion as pick #1 — this is SKU 2, not a second company.

Regulated & compliance-driven — mandated documents

Supplement & Health Claim Substantiation Dossier 82.8

Supplement / wellness brand · \$1500–\$7500 · \$399/mo · risk 3/5 **T2**

Published price band of \$1,500–\$7,500 per claim set, FTC penalties above \$53,000 per violation, and ad lawyers dislike document assembly.

TTB COLA Label Pre-Submission Rejection Risk Review 79.5

Craft brewery / distillery / winery · \$300–\$900 · \$149/mo · risk 2/5 **T2**

Small producers submit dozens of labels a year and every change needs a fresh COLA, while beverage lawyers bill \$350–\$600/hr. Rejections delay launches by weeks. Nobody serves that volume at a per-label price.

Provider Enrollment Revalidation & Credentialing Application Audit 79.2

Small practice / new provider · \$400–\$1500 · \$149/mo · risk 1/5 **T2**

Missing a Medicare revalidation deadline deactivates billing privileges and stops payment outright. Platforms are priced for volume and health systems do it free for employed providers — the independent practice with an episodic need is the gap.

ATF FFL Bound Book & 4473 Discrepancy Audit 78.3

Gun shop / FFL holder · \$900–\$3000 · \$249/mo · risk 2/5 **T2**

ATF revokes licences for wilful recordkeeping violations, so the bound book must be inspection-ready at all times. NSSF subsidises on-site visits for Premium members, but most small FFLs are not members, and a remote document audit is a different product from a site visit.

Money the buyer is already owed

CAM / Operating Expense Reconciliation Audit 83.1

Commercial tenant · \$499–\$2499 · \$299/mo · risk 2/5 **T2**

~40% of reconciliations contain material errors and recovery averages 15–20% of CAM billed. Incumbents take 30–50% contingency; a flat fee where the tenant keeps everything is the wedge.

Other shapes worth a look — different buyers, different rhythms

Tariff Exposure & Mitigation Report 79.6

Importer / manufacturer · \$999–\$4999 · \$499/mo · risk 3/5 **T2**

Tariffs are a dominant 2026 operating issue for importers and manufacturers, and contractors are already re-bidding around the volatility. Urgent, expensive and poorly served at the SMB importer level — but the riskiest card here at 3/5.

\$ What each would have to do to matter		Not a forecast — a break-even model	
PER ORDER	CONSTRUCTION TAKEOFF	HOA REVIEW	GRANT READINESS
Blended price	\$499	\$189	\$399
AI + payment fees	–\$23	–\$11	–\$22
Gross per order	\$476	\$178	\$377
ORDERS PER MONTH NEEDED TO NET \$10,000			
If CAC is low	29	80	35
If CAC is mid	38	108	42
If CAC is high	135	never	94
CAPACITY AND SHAPE			
Minutes to fulfil one	30	25	25
Solo ceiling / month	240	288	288
Net at mid CAC, if volume is reached	\$123k/yr	\$120k/yr	\$120k/yr
Read the middle block, not the bottom one. All three reach roughly the same income <i>if</i> they reach volume — so the real question is which volume is believable. 38 takeoffs a month from small subcontractors is a plausible year-one business. 108 condo reviews or 42 grant reviews a month is a different job. None of the three has the high-ticket cushion the killed WISP offered: all three go permanently underwater if acquisition proves expensive, which is why the high-ticket bet now sits with CAM Reconciliation and Supplement Substantiation in the test order (see the Proforma). CAC is the single unknown that decides all three, and no test has been run yet. One point favours all three: payment is taken at or before delivery, so there are no receivables and no working capital tied up — which is why \$3,000 of testing is enough to learn something. Two things this table cannot score. Whether frontier AI commoditises the deliverable — defensible, but only through choices made on day one, which the playbook covers. And what the work feels like: 38 contractor conversations a month is a different life from 108 anxious condo buyers.			

🕒 What this took — and what it would have taken by hand		Estimated, assumptions shown	
WORK PERFORMED	VOLUME	MANUAL RATE	ANALYST HOURS
Specify each candidate — buyer, input documents, output report, price band	335 ideas	10 min ea.	56
Score every candidate against 10 weighted criteria plus a risk rating	3,350 judgments	8 min/idea	45
Research threads run — search, read sources, extract and record findings	111 threads	25 min ea.	46
Deep verification — named competitors, live pricing, demand statistics	38 candidates	60 min ea.	38
Kill decisions documented with reason and source, plus synthesis, ranking and full write-up	263 kills / 8,800 wds	—	104
Total human effort to reproduce this database			~288 hrs
7 weeks BY HAND	36 days FULL-TIME EQUIVALENT	\$43,300 AT \$150/HR ANALYST RATE	~6 hrs ACTUAL ELAPSED
			48x COMPRESSION
<i>Assumptions: manual rates assume an experienced analyst working from a brief without prior domain knowledge across 18 industries. Volumes are counted from the database and research log, not estimated. \$150/hr is a mid-market consulting benchmark. The database refreshes weekly at no additional analyst cost — a recurring saving this one-off comparison does not capture.</i>			