

Round 4 & 5 — Scenarios and Board Red-Team

AI session transcript · Prompts P4 (scenarios) then P5 (red-team)

You

Create base, upside, and downside scenarios by varying **ONLY** these four drivers: new-logo quota per AE, logo churn rate, documents per firm per month, and inference cost per document. For each scenario, give the driver values and a one-sentence rationale grounded in the case facts — no new facts. Present as a table I can paste into my assumptions sheet. Then state which single driver you believe moves December cash the most; I will test your claim with a data table and report back.

Assistant

Here are three internally coherent scenarios that touch only the four drivers you named. Everything else — pricing, hiring schedule, prepay mix, headcount plans — stays at the case values.

Scenario driver table

Driver	Downside	Base	Upside
New-logo quota (logos / ramped AE / month)	2	3	4
Logo churn (seasoned, per month)	2.5%	1.5%	1.0%
Documents / firm / month (vs. base tenure curve)	−10%	curve as stated	+10%
Inference cost / document (Q1 start)	\$0.20	\$0.18	\$0.16

One-sentence rationales (grounded in the case)

- **Quota — Downside 2 / Base 3 / Upside 4.** The case sets quota at 3 for a ramped AE; a considered legal-tech sale could slip to 2 if the one-month ramp proves optimistic, or reach 4 if the ex-DocuSigner network shortens the sales cycle.
- **Churn — Downside 2.5% / Base 1.5% / Upside 1.0%.** The case fixes seasoned churn at 1.5%; weaker retention pushes it toward the higher early-life risk the grace window hides, while a stickier contract-review workflow could hold firms closer to 1.0%.
- **Documents/firm — ±10% around the curve.** The case's three-step tenure curve is the expansion engine; the scenarios flex how quickly firms climb it without changing the step structure itself.
- **Inference cost — Downside \$0.20 / Base \$0.18 / Upside \$0.16.** The case assumes \$0.18 declining 5% per quarter; the scenarios test whether model-optimization runs ahead of or behind that path, which is the central AI-margin question.

Most-sensitive driver (my claim to be tested)

My prior is that **inference cost per document moves December cash the most**. It scales directly with documents — the fastest-growing quantity in the model — and it hits gross margin exactly when Q4 volume peaks, so a swing from \$0.16 to \$0.20 compounds across both a larger book and a higher usage curve. Run a one-variable data table on December cash for each of the four drivers and see whether the churn and quota lines actually out-swing it before you accept this ranking.

You

You are a board director with a venture background reviewing a Series-A company's annual operating plan. Here are our forecast outputs: [pasted one-pager — active-firm trajectory, monthly revenue and gross margin, headcount plan, cash walk, deferred-revenue roll-forward]. Ask the five hardest questions you would ask in the board meeting. Then identify any internal inconsistencies — revenue growth vs. sales headcount, CAC vs. marketing spend, support ratio vs. customer count, commission expense vs. bookings. Be direct; do not soften.

Assistant

I will give you the five questions first, then the cross-checks. One of the cross-checks is a real break in this plan, not a nuance — fix it before this goes in front of the full board.

Five questions I would actually ask

1. **Runway and the next raise.** You are burning cash every month of this plan and never turn cash-flow positive inside FY2027. At this trajectory, in which month do you cross a six-month cash floor, and is the raise timed to close *before* that month — or are you planning to negotiate from weakness?
2. **The quality of NRR.** Your net revenue retention prints above 100%, but almost all of it comes from the assumed three-step document curve, not from price or seat expansion you have observed. How much of that expansion is in signed usage data today versus modeled, and what does NRR do if firms plateau one step below where the curve puts them?
3. **CAC and the one-month ramp.** Demand-gen more than doubles across the year while quota per AE is held flat, and you assume a brand-new AE reaches full quota in a single month. For a trust-driven legal buyer, what is your blended CAC by Q4, what is the payback in months, and where is the evidence a one-month ramp is real rather than convenient?
4. **Gross-margin durability.** The entire margin story rests on inference cost falling 5% per quarter. If that decline stalls — or per-document cost rises as heavy firms cross onto the top usage step — what is Q4 gross margin, and does the \$2 overage actually cover the marginal document?
5. **Cash quality behind the prepay.** Forty percent of new firms prepay a year, which flatters in-year cash and builds the deferred-revenue balance. What is your exposure if a prepaid cohort sours, how much of your “cash” is really unearned obligation, and what would a refund or dispute do to the runway math in questions 1 and 4?

Internal-inconsistency check

- **Revenue growth vs. sales headcount — consistent.** New-logo additions step only in your AE-hiring months, and revenue accelerates with the book. No obvious disconnect.
- **CAC vs. marketing spend — watch it.** Demand-gen grows faster than new logos in the back half, so CAC is drifting up even as the plan reads as “efficient growth.” Not a break, but it undercuts the efficiency narrative in question 3.

- **Support ratio vs. customer count — this is broken.** Your stated policy is one CSM per 15 active firms, hired one month early. But your customer-success headcount goes flat from mid-year: it holds at 10 CSMs from June onward while the active-firm book keeps climbing past 260 by December. That is roughly **one CSM to 26–27 firms** at year-end — nearly double your own 1:15 ceiling. Either the plan is silently under-staffing support (which understates cost of revenue and overstates gross margin and cash), or the headcount line stopped tracking the firm count after June. Restore the 1:15 rule with the hire-one-month-early timing and your CSM count should keep rising through December; expect cost of revenue to go up and year-end cash to come down once you do. This is the first thing I would make you fix.
- **Commission expense vs. bookings — consistent.** Commission tracks 8% of first-year platform ACV on new logos and lands in the quarter-end months, which matches your bookings pattern.

Bottom line: the numbers are internally tidy *except* for the support ratio, and that one is a policy break hiding inside otherwise-clean arithmetic. Fix the CSM line, re-run cash, and bring the corrected runway to the next meeting.