OpenAI: What I Would Actually Do

A two-page read pulled from the long strategic review. August 10, 2026. Public information only, no position, no compensation, no MNPI

I read everything in the file twice and the honest summary is that this is not a technology problem. The models are fine. The company is losing on paperwork, on defaults and on a pricing page that does not exist. Those are all cheap to fix, which is the good news and also the uncomfortable part.

What is actually wrong

1. Growth stopped. Roughly $25 billion of run rate since February, flat, after adding about $15 billion in the eight months before that. Consumer assistant share slid from around 69% to the mid-40s while the category kept growing. That is not a capacity story, and nobody in twenty focus-group rooms believed the capacity story.

2. The economics are default dead. A $20.9 billion operating loss on $13.07 billion of 2025 revenue, and $17.2 billion of it went to one cloud vendor who is also a shareholder and a rival. **The company runs on financing**, not on operations, and it has never published the date that changes.

3. Enterprise is being lost on contract terms, not on intelligence. Buyers want non-training guarantees, data residency, version pinning, indemnification and published error rates. A CISO in the study said the word "intelligence" had not come up in his last four procurement cycles.

4. Distribution is rented. A rival sits behind the assistant button on more than a billion phones. Nothing in a marketing plan beats a default.

The four things I would do this month

1. **Publish a plain pricing page** on openai.com. Every tier side by side, no login, no sales call, API pricing one click away. Five separate rooms named this independently across two studies. Two weeks of work. It is still not done.

2. Seat a permanent CMO and announce the name. The seat has been interim since early 2026, during a registration year, and an empty chair is a story competitors will tell for you.

3. Take the security, residency, indemnity and admin-control documents that already exist internally and put them on public pages. Three weeks. **This is the cheapest share** available anywhere in the business.

4. Name an owner for answer-engine citations, internally if you have to, and give them a monthly scorecard. You are the answer engine and you are under-cited inside every rival assistant. Nobody's objectives currently contain the word "citation."

Where the money actually is

1. Contracts, not capability. Terms are a marketing deliverable. Publish them, run named-account programs instead of a Contact Sales form, and take back the twelve or so points of enterprise API share that were lost on paperwork. One quarter.

2. Distribution as a board-level objective with a named owner and a real budget: OEM slots, carrier deals, an operating-system assistant default, the mobile share sheet and keyboard. Two to four quarters per deal, and every deal announced, because half the value of a distribution deal is the signal to the next partner.

3. Coding. Developers set the enterprise default about eighteen months out, they switch fastest, and they tell their friends fastest. **Win coding back on merit**, with migration guides, cost calculators and a published reliability benchmark you are willing to lose on occasionally.

What I would stop doing

1. Stop explaining the share loss with capacity constraints. It is not credible and it signals that nobody has diagnosed the problem.

2. Stop scaling advertising before attribution exists. It hit $100 million of run rate in six weeks and then disappointed, because advertisers cannot attribute return. Reach was never the gating problem.

3. Stop adding tiers and mythological model names. Fifteen live names and nobody in the study could sort them by capability. Collapse to three consumer names and one capability ladder.

4. Stop renewing eight-figure broadcast buys without a matched-market holdout. **You are measuring like a startup** and spending like Procter and Gamble.

Where the industry is going, and what it does to you

1. The category has moved from answers to finished deliverables - spreadsheets, slides, patches, resolved tickets. Marketing built around answer quality is marketing for the last product cycle.

2. Nobody has solved verification. Task-completion reliability is the next two years of competition, and **whoever publishes it first** sets the standard everyone else has to answer.

3. Megawatts are the constraint, not chips or capital. Power secured in 2024 and 2025 cannot be bought back in 2027, and the compute reservations are a genuine moat that nobody has bothered to market.

4. Switching costs are collapsing toward zero. Lock-in has to be built above the API - in workflow, data and evaluation - or it does not exist at all.

Three lines from the rooms worth keeping

1. "You are the answer engine. Explain to me why you are losing the answer."

2. "A hundred million people use you and there is no page that plainly says what it costs."

3. "You are spending like Procter and Gamble and measuring like a startup."

Who I would hire, in order

1. Permanent Chief Marketing Officer - now. 2. Enterprise product marketing lead - now. 3. Head of answer-engine optimization - Q4 2026. 4. VP of marketing science with veto power over renewals - Q4 2026, before the next upfront. 5. Four regional leads for Japan, India, Germany and Brazil - Q4 2026 into Q1 2027, because the users are outside the United States and roughly 78% of the staff is inside it.

The questions I cannot answer from public material

1. What is gross margin per token today, and what is the defensible curve to breakeven? 2. What are the real cancellation terms inside the $1.4 trillion of compute commitments? 3. Is advertising incremental revenue or transferred subscription revenue plus a permanent trust liability? 4. What would it cost to simply buy an assistant default, and why has that number never sat beside the marketing budget? 5. **Which two things** should this company be best in the world at? Every table converged on focus as the constraint, and that decision sits underneath all the others.

If only one thing gets done: the pricing page. It is two weeks of work, it has been the most-repeated finding across every room in every study, and the fact that it is still undone tells you more about the organization than any of the numbers do.