OPENAI

STRATEGIC REVIEW AND ACTION LIST

Built from the intelligence brief, the five-angle evaluation, the marketing analysis and growth plan, the three one-page summaries, and the reconvened hundred-person focus group with its hundred marketing suggestions. August 10, 2026.

ANALYSES BASED ON PUBLIC INFORMATION. NOTHING HERE IS INVESTMENT ADVICE. NO POSITION / NO COMPENSATION / NO MNPI.

HOW TO READ THIS

Ten sections. One hundred points on marketing and websites, arranged in ten subsections of ten. Between twelve and sixteen points under each of the remaining nine headings. In every point, the one to four words that carry the main idea are bold, underlined and set one point larger than the surrounding text, so the document can be skimmed in about four minutes and read in about forty.

Where an item takes real time, the approximate elapsed time appears in parentheses at the end of the point. Those estimates assume a funded owner and normal internal review; they are elapsed time, not effort. Items marked immediate are decisions rather than projects.

OpenAI is a private company that publishes no audited financials. Every figure here is drawn from public reporting or from leaked material verified by major outlets, and should be treated as an estimate. Where sources conflict, the conflict is shown rather than resolved.

1. MARKETING AND WEBSITE IDEAS

One hundred points in ten subsections. Ordering within each subsection is rough priority order.

1.1 The pricing page and the commercial architecture

1. **One public pricing page** on openai.com listing Free, Go, Plus, Pro, Business, Enterprise and Education side by side, in plain English, with no login and no sales call. Five separate rooms across two studies named this independently and it is still not shipped. (Two weeks of design and copy, one week of legal review.)

2. **Put API pricing** on the same page, or one obvious click away, instead of on a separate developer domain. A buyer comparing total cost cannot hold two domains in their head, and the comparison is being made on a rival's page instead. (One week.)

3. Add a **business-versus-individual calculator** with five questions that tells a small-business owner whether the team tier beats buying five personal seats. The tamale-shop owner defaulted to five personal seats, which is worse for both sides. (Three weeks.)

4. Publish an **education discount** and an annual, invoiced, per-seat option without requiring a call. A university procurement officer who cannot find this quietly recommends a competitor and never tells you. (Two weeks, mostly finance sign-off.)

5. Show the **ad-free path** prominently inside pricing and keep it cheap. Opting out has to be genuinely available rather than nominally available, and pricing is where that promise is either made or broken. (One week.)

6. Publish a **country payment matrix** showing accepted local payment methods and currency per market. The $8 Go tier is under-performing in price-sensitive markets on rails, not on price. (Four weeks including payments work.)

7. Give every tier a **who this is for** line written by a copywriter rather than a product manager. Feature lists tell a reader what they get; they do not tell a reader whether they are the intended buyer. (Three days.)

8. **Date and version** the pricing page publicly so encyclopedias and journalists cite a live primary source rather than an eighteen-month-old article. This is the single cheapest way to fix the accuracy of everything written about you. (One day.)

9. Run a **recall test quarterly**: put five ordinary people in a room and ask them to write the lineup from memory. Iterate the page until four of five get it right. (One day per quarter.)

10. Add **transparent usage limits** per tier in numbers rather than adjectives. “Generous” and “extended” are the two words that generate the most support tickets and the most negative comparison content. (One week.)

1.2 Site architecture and answer-engine optimization

11. Appoint a **named GEO owner** with a budget, headcount and 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. (Hire in Q4 2026; scorecard live in four weeks after start.)

12. Instrument **citation share monthly** inside Gemini, Claude, Copilot, Perplexity and Grok for the two hundred highest-intent queries in the category, and report it to the executive team beside revenue. (Six weeks to build, then ongoing.)

13. Write **fifty canonical answers** to the highest-intent questions in the category - pricing, limits, data handling, model choice, migration - in extractable question-and-answer blocks on-domain. (Six to eight weeks.)

14. Add **structured data everywhere**: organization, product, offer, FAQ and how-to schema on every product and pricing page so machines can quote you without guessing. (Three weeks of engineering.)

15. **Flatten the hierarchy** so no commercially important page sits more than two clicks from the home page. Launch posts currently outrank pricing in internal link authority, which is backwards. (Four weeks.)

16. Rebuild the **internal linking model** to concentrate authority on pricing, comparison, security and enterprise pages instead of on the newsroom archive. (Two weeks.)

17. Publish **dated citable facts** - headcount, user numbers, model availability, pricing changes - as timestamped primary sources so third parties quote you accurately and encyclopedia editors have something to cite. (Ongoing; two weeks to set up.)

18. Run a **quarterly link audit** for broken links and orphan pages. The newsroom archive is already fraying and every dead link is authority thrown away. (Two days per quarter.)

19. Make **titles query-shaped** rather than campaign-shaped. Internal product naming in a page title is a decision to lose the search that page exists to win. (One week of retitling.)

20. Ship a **public model-selection guide** that tells a developer honestly when not to use the largest model. Honest downsell is the highest-trust content a platform can publish. (Two weeks.)

1.3 Comparison, migration and switching content

21. Build a **first-party comparison hub** covering Gemini, Claude and Copilot, updated monthly, honest about where you lose. Comparison queries are the highest-intent queries in any category and you have conceded all of them to affiliates. (Six weeks to launch, monthly upkeep.)

22. Publish a **migration guide** from each named rival, at both the indie-developer and the enterprise-architect level. An indie developer in the focus group moved a workload for exactly this reason and the model was not the deciding factor. (Three weeks per guide.)

23. Ship a **live cost calculator** that prices realistic workloads across your tiers and against published rival pricing. Rivals publish this. You publish launch posts. (Four weeks.)

24. Publish a **reliability benchmark** on real repositories and real tasks, and be willing to lose on it publicly some quarters. Credibility earned by an honest loss outlasts any benchmark win. (Eight weeks to design, quarterly thereafter.)

25. Write the **switching-cost story** deliberately: what you keep, what moves, what breaks. Switching costs in this category are approaching zero, so the winner is whoever removes friction from the switch. (Two weeks.)

26. Add a **why-not-us section** to the comparison hub naming the two or three cases where a rival is the better choice. Nothing raises conversion on the other twenty cases like it. (One week.)

27. Publish a **deprecation calendar** twelve months forward and never move it. Version churn is the single most-cited reason regulated buyers refuse to standardize on a vendor. (Two weeks, then permanent.)

28. Create a **bake-off kit** - evaluation harness, scoring rubric, sample tasks - that a prospect can run themselves. If you are winning evaluations you want more of them run; if you are losing, you want to know why before procurement does. (Six weeks.)

29. Buy and win the **alternative queries** - “alternatives to,” “vs,” “switch from” - with owned pages before affiliates finish consolidating them. The consolidation window closes roughly eighteen months after a category matures. (Ongoing.)

30. Publish a **model changelog** written for customers, not for researchers, so a buyer can see what changed under a workflow they validated. (One week to start, ongoing.)

1.4 Enterprise and trust pages

31. Publish the **boring artifacts publicly**: security architecture, data residency by region, non-training guarantees, indemnification terms, admin controls, audit trails and retention policy. A CISO in the focus group said the word intelligence had not come up in his last four procurement cycles. (Four weeks.)

32. Publish a **version-pinning policy** in plain language. A regulated buyer cannot validate a workflow against a model that changes underneath it, and this is a documented reason you are losing head-to-head evaluations. (Two weeks.)

33. Publish **measured error rates** by task category with methodology. Nobody in the category does this; the first to do it defines how buyers ask the question. (Eight weeks.)

34. Replace the **generic contact form** with named-account programs, a named human per account and a published response-time commitment. The form is currently the entirety of the remembered enterprise experience. (Six weeks to stand up.)

35. Assemble a downloadable **procurement kit**: standard security questionnaire answers, data processing addendum, sub-processor list, penetration-test summary, insurance certificate. No login. (Three weeks.)

36. Ship **one case study weekly**, per vertical, with a measurable number in the headline. Fifty-two named proofs a year is a different company from four beautiful films a year. (Ongoing; two weeks to build the machine.)

37. Market **Presence to buyers** of contact-center and business-process software in their vocabulary - containment rate, handle time, escalation policy - rather than in launch-post vocabulary. (Four weeks.)

38. Explain the **shareholder-competitor relationship** plainly on your own domain before a journalist explains it for you. It is a communications problem before it is a legal one. (One week.)

39. Publish a **public status history** with incident post-mortems and link it from the developer home page. Uptime honesty is enterprise marketing. (Three weeks.)

40. Commission one **serious print artifact** per year for the executive desk in named accounts. A well-made annual publication outperforms a quarter of programmatic in enterprise brand terms. (Twelve weeks, annual.)

1.5 Brand, promise and naming

41. Write the **one-sentence promise** and make every piece of creative defensible against it. Five people in one room produced five different sentences: intelligence, a helper, the smart one, the one with ads now, AGI. (Three weeks with the executive team.)

42. **Collapse the naming** to at most three consumer-facing names and one versioned capability ladder. Sol, Terra, Luna, Red, Rosalind, Codex, Presence, Daybreak, Atlas, Go, Plus, Pro, Business, Enterprise is a portfolio nobody can sort by capability. (One quarter including migration.)

43. **Retire mythological naming** for anything a consumer touches. Mythological model names are the naming pattern of a company selling identity rather than capability, and that transition is very hard to reverse once made. (Part of the naming project.)

44. **Separate the brands**: ChatGPT is the consumer brand, OpenAI is the corporate brand. Different sites, voices, audiences, metrics and owners. Today the enterprise sale starts by climbing out of the hole the consumer brand dug. (One quarter.)

45. Fix the **category-noun trap**. Being a verb is the rarest asset in marketing and it is habit, not preference. Habit is exactly what a rival with a default placement on a billion phones can take. (Continuous.)

46. Write **permission not persuasion**. The clergyman's finding is that people do not ask whether it works, they ask whether it is all right to use it, and no advertising in the category answers that. (One campaign cycle, roughly one quarter.)

47. Address the **shame problem** directly. People use the product constantly and mention it reluctantly; the art student hides it from her cohort. Media weight does not fix shame. Permission does. (One campaign cycle.)

48. Keep **ten percent frightened**: roughly half of American adults tell pollsters they are more concerned than excited about AI, so reserve a tenth of the creative budget for reassurance and spend the other ninety percent on ordinary and specific. (Budget decision, immediate.)

49. Define the **brand personality deliberately** rather than letting the model's default tone define it. Right now the personality is whatever the last system prompt did, which is not a brand decision. (Four weeks.)

50. Reconcile the **purpose and advertising**. The stated mission and the published advertising do not appear to be about the same thing, and ordinary people in the focus group could see the gap without prompting. (One quarter.)

1.6 Paid media and performance

51. **Defend the branded term** on every major engine. Rivals and affiliates bid on your product name, your quality score on your own name should be near-perfect, and the click costs pennies. This is a solved problem somebody has not solved. (One week.)

52. Shift budget from category awareness to **developer and migration intent** - the queries nobody is bidding on because everybody is chasing the brand term. The highest-value paid audience is not consumers. (Two weeks to restructure.)

53. Run **matched-market geo holdouts** before renewing any eight-figure broadcast commitment. In a category growing this fast, any measurement without a holdout will show that everything worked. (Eight to twelve weeks per test.)

54. **Pre-register analysis plans** for major campaigns so results cannot be re-cut after the fact. This costs nothing and removes the single largest source of self-deception in the budget. (One week per campaign.)

55. Cap **frequency in saturated metros** and reallocate to the smaller Midwest, Mountain West and South markets already showing impression-share gains. (Two weeks.)

56. Use **document-format ads** on professional networks to deliver an actual artifact - a security overview, an implementation guide - rather than a slogan. That format is the one thing that platform does cheaply. (Three weeks.)

57. Build **named-account lists** rather than buying seniority-and-industry awareness. Company, title and seniority targeting is the platform's only real advantage and you are using it for the thing it is expensive at. (Four weeks.)

58. Buy **host-read audio** in shows the target already trusts. The purchase in this category is a trust purchase, and trust is transferable in a way owned media cannot replicate. (Six weeks to book.)

59. Publish **one incrementality result** publicly. Nobody in the category has published one; the first to do it owns the credibility and sets the standard rivals then have to meet. (One quarter after the first test completes.)

60. Give marketing science **veto over renewal** of any untested spend above a stated threshold. Without the veto the function is advisory and will be overruled every upfront. (Policy decision, immediate.)

1.7 Content, proof and programming

61. Turn the **small-business documentaries** into always-on weekly publishing with a fixed slot and a permanent budget. It is the best creative the company has made and it was made as a campaign. (Six weeks to convert, then ongoing.)

62. **Export the documentary format** to Japan, India, Germany and Brazil with local subjects and native crews, not translated American spots. It is the most exportable asset you own and nobody has exported it. (One quarter per market.)

63. Put **faces in the imagery**. The visual system is abstract and machine-made; faces convert, and every photographer in the study said so without being asked. (Immediate on the next shoot.)

64. **Repackage every post** into at least three formats - short video, audio, and a one-page artifact. You currently pay for content once and use it once. (Three weeks to build the workflow.)

65. Rebuild **video around search**: task tutorials titled for intent, not launch livestreams nobody searches for. The majority of the category's search volume is sitting on a platform you treat as a broadcast tower. (One quarter.)

66. Hire a **plain-language editor** for the newsroom whose only job is removing institutional voice. The documentation is currently the only honest writing on the site. (Hire in Q4 2026.)

67. Publish one **genuinely linkable resource** per quarter - a dataset, a benchmark, a tool - that other sites want to link to on its own merits rather than as coverage. (Six weeks per resource.)

68. Show **ordinary non-office work**: home health aides, two-person wildlife charities, tradespeople, teachers. The home health aide in the study assumed the product was for office workers because the advertising told her so. (Next campaign cycle.)

69. Publish a **customer-facing roadmap** written for buyers rather than developers, so an enterprise can plan against you instead of guessing. (Four weeks, then monthly.)

70. Explain **why things get killed**. Atlas, the video app and other efforts were sunset with no explanation to the people using them, who experienced a strategy as an abandonment. (One page, one week, every time.)

1.8 Community, earned media and employer brand

71. Build an **employee advocacy program** with a content library, posting guidance and measured employee-sourced reach. Thousands of employees will out-reach the brand handle on every professional network. (Eight weeks to launch.)

72. Put **named humans in communities** answering real questions in public. Brand-voice posting is fatal on those platforms; identified staff is the only thing that works. (Four weeks to write policy, then ongoing.)

73. **Never edit the encyclopedia**. Engage only on the talk page with disclosed affiliation, and publish citable primary sources so editors have accurate material. (Policy, one day.)

74. **Separate employer and product** feeds. Half the feed says "we are hiring" and half says "we are the future," neither has a call to action, and they are undercutting each other. (Two weeks.)

75. **Publish a headcount** and a hiring philosophy. A company that does not publish a headcount lets other people publish one for it, which is why 4,500 and 7,850 both circulate as fact. (One week.)

76. **Reconcile the contradiction** between "nearly doubling the team" and "doing more with fewer people." Candidates read both and draw the obvious conclusion about internal coherence. (One week.)

77. **Answer every applicant**. The careers experience is the only corporate-brand touchpoint most people will ever have, and a four-time applicant who never heard back is now telling that story in focus groups. (Four weeks to fix tooling.)

78. Establish a **rapid-response protocol** for viral product screenshots, with a named owner and a two-hour clock. This is a permanent condition of the category, not an incident. (Three weeks.)

79. **Celebrate builders weekly** by name. Ecosystem developers are the cheapest sales force available and none of them are being named. (Two weeks to start, ongoing.)

80. **End the feud content**. Public executive rivalry has handed reporters a reusable frame that will be deployed at the worst possible moment, most likely during the listing window. (Immediate.)

1.9 International and localization

81. Make **native creative locally** in Japan, India, Germany and Brazil. The user base is overwhelmingly outside the United States and roughly seventy-eight percent of the workforce is inside it, so the marketing is made where the employees are, not where the users are. (One quarter per market.)

82. Fix **local payment rails** before spending another dollar on acquisition in India and Southeast Asia. A cheap tier without local rails is a product decision wearing a marketing decision's clothes. (Six to twelve weeks per market.)

83. Make the **trust page the landing page** in Europe. In that market compliance documentation is the top-of-funnel asset, not a footer link. (Four weeks.)

84. Market the **Dublin build-out** - 88,000 square feet, 250 roles, roughly €105 million - as the European trust asset it is, landing exactly as EU AI Act enforcement powers activate. (Two weeks.)

85. Publish **local proof stories** per market: a Japanese small manufacturer, an Indian logistics operator, a German Mittelstand firm, a Brazilian retailer. (One per market per month.)

86. Localize the **pricing page first**, before the campaign. Price is the most-searched question in every market and translated marketing on an untranslated pricing page converts badly. (Three weeks.)

87. Staff **regional marketing leads** with real budget authority in the four priority markets rather than routing everything through San Francisco. (Hire in Q4 2026 and Q1 2027.)

88. Publish **data residency by region** as a public map. In several markets this single artifact is the difference between shortlisted and not. (Two weeks.)

89. Translate the **documentation estate** fully, not partially. Half-translated docs are worse than English-only because they signal abandonment mid-way. (One quarter.)

90. Sponsor **local developer communities** in the four markets rather than exporting a first-party conference. (Ongoing, one quarter to set up.)

1.10 Advertising inside the product - communications and measurement

91. Publish a **permanent trust page** on advertising, in plain language, before the next controversy rather than after it. The only durable defense against a trust attack is a document published before the attack. (Three weeks.)

92. Publish an **ad-load cap** as a stated percentage of responses and report against it quarterly. Third-party trackers reported penetration swinging from roughly twenty-six percent to near zero to roughly fifty-one percent in a single summer; that volatility is invisible to users and visible to researchers. (Four weeks.)

93. Maintain an **audited firewall** between ad ranking and answer generation and publish the auditor's letter. (One quarter to establish, annual thereafter.)

94. Commit in writing that **advertiser identity never enters** the model context for the organic answer. This is the specific claim critics will test, so make it testable. (Two weeks.)

95. **Restrict categories permanently** and publish the restricted list. Health, elections, financial products and children's categories are the ones that will define the coverage. (Three weeks.)

96. Ship **attribution before inventory**. Ads reached roughly $100 million of run rate in six weeks and then disappointed because advertisers cannot attribute return. Reach is not the gating problem. (One to two quarters.)

97. Label sponsored placements so **distinctly that screenshots** taken by critics make your case for you. Design the label for the hostile screenshot, not for the design review. (Three weeks.)

98. Publish **quarterly ad transparency** data the way a regulated exchange publishes volume: load, categories, complaint rates, removals. (One quarter to build.)

99. Never let **engagement time** become the ad business's optimization target, and publish the objective function you do use. An assistant people trust and an assistant optimized for time-in-conversation cannot both be the primary objective. (Immediate.)

100. Give subscribers a **written no-ads guarantee** they can point to. A paying teacher who watched free-tier ads appear in a friend's app has already stopped believing the informal version. (One week.)

2. HURDLES THAT MAY BE IN YOUR WAY

Sixteen. Ranked roughly by how much of the plan each one blocks.

1. The **flat run rate** is the hurdle every other hurdle sits on. Roughly $25 billion annualized since February 2026 after adding about $15 billion in the preceding eight months, with consumer assistant share sliding from about sixty-nine percent to the mid-forties. Marketing cannot fix a growth-rate problem it did not cause, but it will be asked to.

2. **Default dead economics**. A $20.92 billion operating loss on $13.07 billion of 2025 revenue, roughly negative one hundred twenty-two percent operating margin in Q1 2026, and $1.22 to $1.69 spent for every dollar earned. Every marketing ask competes against a compute bill.

3. **Single-vendor concentration**. $17.2 billion paid to one cloud vendor in 2025 - more than total revenue - is a structural fragility that also constrains what you can say publicly about that vendor, who is simultaneously a shareholder and a competitor.

4. **The registration window**. A confidential S-1 filed June 8, 2026 means every forward-looking marketing claim, executive post and celebratory metric is now a legal artifact. Marketing organizations are usually the last group told the rules changed. (Fix in two weeks by seating a securities reviewer in the workflow.)

5. **Empty leadership chairs**. Marketing and communications interim since early 2026, the applications chief moved to part-time advisory in July, eleven or more senior departures since April. You cannot run a repositioning without a permanent owner of the position.

6. **Distribution you rent**. Google, Meta and Apple own the surfaces; a rival sits behind the assistant button on more than a billion phones. Nothing in a marketing plan overcomes a default, and this one was not bought.

7. **Enterprise lost on paperwork**. Roughly thirty-two to forty percent rival share against twenty-five to twenty-seven percent, with reported seventy percent head-to-head win rates, decided on non-training guarantees, residency, version pinning, indemnity and published error rates. The hurdle is that none of those are marketing's to grant.

8. **Ads corrode the asset** you are trying to defend. Ad revenue at these margins will be the easiest money the company has ever made, and easy money reorients an organization faster than any strategy memo. The Yahoo pattern is the specific risk.

9. **Naming debt**. Fifteen live product and tier names, none sortable by capability by anyone in the study. Every campaign now spends a portion of its budget explaining the portfolio before it can sell anything.

10. **Your own news volume**. The company generates so much news that its marketing cannot be heard over it. A record round, an S-1, litigation, gigawatt infrastructure and a public rivalry all compete with the campaign for the same attention.

11. **Megawatt lead times**. Capacity is constrained by interconnection queues, transformers and turbines on multi-year lead times. Power secured in 2024 and 2025 cannot be bought back in 2027, so demand generation can outrun supply and turn a marketing win into a rate-limit complaint.

12. **Litigation overhang**. Publisher and training-data disputes are a marketing problem the marketing organization does not own, and they surface at moments chosen by opposing counsel rather than by you.

13. **Measurement immaturity**. Spending like a consumer packaged goods giant and measuring like a startup. Without holdouts, every campaign in a fast-growing category reports success, which means the budget cannot be defended when it is challenged.

14. **The shame problem**. Users use the product constantly and mention it reluctantly. Word of mouth is the growth mechanism you most need and the one most damaged by social discomfort about the category.

15. **Governance complexity**. A nonprofit Foundation holding roughly twenty-six percent while appointing the entire PBC board is genuinely unusual, will be priced by public investors as a discount before it is priced as a premium, and is difficult to explain in a thirty-second spot.

16. **Talent market signaling**. Reported median total compensation near $608,182 and roughly $1.5 million average annual equity grants reach candidates through leveling sites and press. The loudest recruiting message you send is one you do not control.

3. THINGS YOU COULD BE DOING RIGHT NOW

Fifteen items that need a decision rather than a strategy. Most are shippable inside a quarter and several inside a week.

1. **Ship the pricing page**. It has been the most-repeated finding across two studies and twenty-plus rooms. It requires no strategy work, only a decision. (Two weeks.)

2. **Seat a permanent CMO** and announce the hire by name. An interim seat during a registration window is a story rivals will tell for you. (Search under way now; announce within the quarter.)

3. **Defend the branded search term** on every engine today. Your quality score on your own name should be near-perfect and the click costs pennies. (One week.)

4. **Name the GEO owner** even before you hire externally. An internal person with a scorecard beats an unfilled requisition. (One day to name; four weeks to first scorecard.)

5. **Publish the security pages** you already have internally: residency, non-training guarantees, sub-processors, retention, admin controls. These documents exist; they are behind a sales call. (Three weeks.)

6. **Put staff into communities** with names attached, starting with the developer forums and the largest public communities. Write the policy in an afternoon and start the same week. (One week.)

7. **Start the weekly case study** machine. One named customer per vertical per week with a number in the headline beats one beautiful film per quarter. (Two weeks to build, then ongoing.)

8. **Book the first geo holdout** before the next upfront so the next eight figures are measured rather than assumed. (Eight to twelve weeks; start now to have a result before renewal.)

9. **Publish the ad trust page** while ads are at roughly $100 million of run rate. Doing this now is a policy choice; doing it at $5 billion is a concession extracted by a regulator. (Three weeks.)

10. **Put a securities reviewer** inside the marketing approval workflow this week. The cost is one person's partial time and the downside it prevents is unbounded. (One week.)

11. **Retitle the video library** for search intent. This is an afternoon of work per playlist and recovers search volume you already paid to create. (Two weeks.)

12. **Answer the applicant backlog** and fix the careers auto-response. Cheap, visible, and currently generating the ugliest anecdote in every focus group. (Four weeks.)

13. **Publish the headcount** and the hiring philosophy in one short post, ending the 4,500-versus-7,850 confusion that third parties are resolving for you. (One week.)

14. **Date the primary sources**. Timestamp pricing, availability and user figures on-domain so encyclopedias and journalists stop citing stale articles. (One day.)

15. **Kill the feud content** by decision today. It costs nothing to stop and removes a frame that will otherwise be used against you during the listing. (Immediate.)

4. THINGS WORTH AVOIDING

Fifteen. Each one is either currently happening or currently being proposed somewhere in the organization.

1. **Do not explain** the share loss with capacity constraints. A rate limit does not explain why a user who already has your app installed opens a different one, and the explanation itself signals that nobody has diagnosed the problem.

2. **Avoid silver bullets**. Advertising, hardware and the IPO are all attractive precisely because they avoid the actual reason customers are leaving. The lead-bullets framing applies directly: win coding and enterprise contracts on merit first.

3. **Do not scale ads** before attribution exists. Selling more of an unmeasured product to advertisers who cannot attribute return produces churn, not revenue, and spends trust to do it.

4. **Never edit the encyclopedia** entry, and never let an agency do it on your behalf. The reputational cost of being caught vastly exceeds the value of any edit.

5. **Avoid abstract launch films** as the majority of the creative. They buy permission from the frightened, which is real work, but the regional documentary spots did more per dollar and the split should reflect that.

6. **Do not add tiers**. Every new tier was launched to answer a revenue question and nobody owns the story that connects them. The next one makes the lineup less legible, not more profitable.

7. **Stop counting sign-ups** as the headline marketing metric. In a category where the free tier is the product, sign-ups measure curiosity. Retained weekly usage measures the business.

8. **Avoid consumer-style launches** in health and education. A launch video actively reduces a hospital CIO's confidence, and the school district buys through procurement and references that brand advertising cannot reach.

9. **Do not bury governance** in the prospectus. If the Foundation-controls-the-PBC structure is a genuine long-horizon advantage, say so plainly and accept the initial discount as the price of not being managed by the median shareholder.

10. **Avoid translated creative** in the four priority international markets. Translated American work reads as American work, which is a specific and avoidable failure that local teams will name immediately.

11. **Do not renew broadcast** above eight figures without a completed holdout test. Renewal without measurement is the single largest recurring waste in the plan.

12. **Avoid engagement optimization** anywhere near the assistant. An assistant trusted because it tells the truth and an assistant optimized for time-in-conversation are different products, and only one can be the primary objective.

13. **Do not announce structure** - acquisitions, partnerships, corporate changes - without a communications plan attached. You have repeatedly let reporters set the frame on your own news.

14. **Avoid brand-voice posting** in communities. It is fatal on those platforms and the failure is public and permanent.

15. **Do not promise AGI** in marketing without an operational definition. Identity language cannot be reasoned about, and in a registration window it becomes a forward-looking statement with legal weight.

5. HOW SPECIALIZED VENTURE CAPITAL MIGHT VIEW YOUR PLANS

Fifteen. Written as growth-stage and applied-AI specialists would actually frame it in a partner meeting, not as they would say it to management.

1. Specialized venture capital will start with the **growth rate question** and nothing else. A flat run rate at a company still spending like a hypergrowth company is read as a change of species, not a plateau, and the diagnosis they want is whether the market stopped growing or you stopped being the thing people tell their friends about.

2. They will treat **gross margin** as the only number that matters and operating margin as a distraction. Cost per token, serving efficiency, routing and custom silicon are the levers; a modeled base case reaching operating breakeven around 2029 on roughly $88 billion of revenue is credible only if the margin curve is published and defended.

3. They will ask for the **default-alive date**. A default-dead company that knows its date behaves completely differently from one that does not, and refusing to name one is read as not having done the work rather than as prudence.

4. The **$852 billion mark** is the problem in the room. A modeled base case supports roughly $700 billion to $1.1 trillion, meaning the last round already priced the story. Growth investors will want two consecutive quarters of reaccelerating revenue and improving gross margin before they underwrite it again.

5. They will read **$1.4 trillion of commitments** against roughly $73 billion of cash as a leveraged bet on inference costs falling fast. Sophisticated investors will discount the nameplate figures for option value and cancellation terms, then ask to see those terms, and the answer determines the discount.

6. **Enterprise share loss** will be the cheapest thing they point at. Losing on contract terms rather than on capability is, to a venture investor, the most fixable problem in the file and therefore the most damning that it is unfixed.

7. They will price **distribution as destiny**. Renting the surfaces from Google, Meta and Apple while a rival owns the assistant default on a billion phones is the structural fact that most constrains the terminal value in their model.

8. The **ads business** will divide them. Some will value it as high-margin incremental revenue on existing traffic; others will read it as evidence the core business could not fund itself and as a permanent constraint on trust. Both will want the ad-load cap and firewall in writing.

9. They will treat **governance as a discount** first. A nonprofit appointing the entire board of an $852 billion PBC is unpriceable by comparables, and the Foundation-PBC board overlap is the specific item they will want unwound before a listing.

10. **Empty executive seats** during a registration year are read as an execution risk multiplier. An interim CMO, a part-time applications chief and eleven-plus senior departures since April are, in diligence, treated as a single signal rather than as separate events.

11. They will want a **benefit metric** more than most operators expect. Mission-aligned capital in this category increasingly asks whether users are measurably better off, because that is the metric that survives a regulatory cycle. Nobody publishes it, which makes it cheap to own.

12. They will ask what the **top idea** in the chief executive's mind has been for eighteen months. Between a record round, an S-1, litigation and gigawatt-scale negotiation, the honest answer predicts the next two years better than any projection.

13. Specialists in **applied AI infrastructure** will see the compute reservation - ten gigawatts of custom accelerators, Stargate, up to nine hundred thousand DRAM wafers a month - as a genuine moat inside a semiconductor supply chain that no software company has ever built before, and will be surprised nobody markets it.

14. They will value **Codex and agents** above the consumer app on a five-year view, because developers set the enterprise default roughly eighteen months out and agentic task completion is where the durable margin sits.

15. They will assume **marketing spend is unmeasured** until shown otherwise. Publishing one incrementality result would change diligence conversations more than any deck, because no company in the category has published one.

6. COMPETITOR ANALYSIS

Eighteen points across what is working for them, what is not, and what it means for you.

5.1 What is working for them

1. **Google owns the default**. Gemini reaches the assistant button on more than a billion phones through platform placement, runs its own silicon, and has cut prices aggressively. Distribution plus vertical integration is the strongest position in the category and it does not require winning a benchmark.

2. **Anthropic won the contract**. Enterprise-first positioning, an ad-free promise, non-training guarantees and reliability messaging have produced roughly thirty-two to forty percent of enterprise API share and reported seventy percent head-to-head win rates. They are not winning on intelligence; they are winning on procurement.

3. **Microsoft bundles Copilot** into an existing enterprise agreement, which means the purchase decision is a renewal rather than an evaluation. Being simultaneously your largest shareholder and your rival is a position no marketing plan can neutralize.

4. **Meta uses distribution** the same way, embedding assistants inside social surfaces that already have daily habit, and modeling conversational advertising economics that the whole category will eventually copy.

5. **Open weights set the floor**. DeepSeek, Mistral, Alibaba and Meta's open models make price a race to a floor you cannot control, and routers like OpenRouter make switching a configuration change rather than a migration.

6. **Perplexity owns the answer** framing in a way you should have. It is a much smaller company that has made citation share a measured discipline with a named owner, which is precisely the gap in your organization.

5.2 What is not working for them

7. **Their consumer brands lag**. None of them is a verb. The category noun is still yours, and no competitor has produced a consumer brand asset comparable to it despite spending heavily to try.

8. **Bundling breeds indifference**. Copilot seats bought in a renewal are frequently unused, and low activation inside enterprise agreements is a soft underbelly that a well-run competitive campaign could attack with usage data.

9. **Enterprise-first costs reach**. The rival winning contracts has a much smaller consumer footprint, which means slower pilot approval and no bottom-up adoption pulling procurement along. Your consumer familiarity gets a pilot approved in a week rather than a quarter.

10. **Price cuts compress everyone**. Aggressive discounting to buy share damages the discounter's own margin story and creates an industry-wide expectation of falling prices that will eventually constrain their ability to fund frontier training.

11. **Open models lack accountability**. No indemnity, no named human, no incident process. Regulated buyers cannot use them for anything material, which leaves a defensible commercial band that you are currently not defending on paperwork.

12. **Nobody has solved verification**. Every competitor claims agentic task completion and none can prove reliability on real work. The first company to publish credible task-completion reliability data defines the metric for the category.

5.3 What it means for you

13. **Compete on contracts** immediately, because that is where the enterprise loss is occurring and it is the cheapest share available anywhere in the business. Terms are a marketing deliverable, not a legal one. (One quarter to publish the full set.)

14. **Buy or build distribution** as a board-level objective with a named owner and a budget: OEM slots, carrier deals, an operating-system assistant default, the mobile share sheet and keyboard, and a real reseller channel. (Two to four quarters per major deal.)

15. **Protect the category noun** deliberately. It is the one asset no competitor has, it was won rather than bought, and it is currently held together by habit that a default can take.

16. **Win coding back** on merit. Developers set the enterprise default roughly eighteen months out, they switch fastest, and they tell their friends fastest. Codex at roughly five million weekly users is a strong base to defend rather than a consolation prize.

17. **Attack unused seats** where competitors bundled rather than sold. Activation data is the most persuasive competitive material in enterprise software and nobody in this category is using it.

18. **Publish reliability first**. Whoever publishes credible task-completion reliability data first sets the standard everyone else must answer, and you have more real-world usage to draw on than anyone.

7. TRENDS AND DIRECTION FOR YOUR INDUSTRY

Fourteen structural shifts, each with a marketing consequence.

1. The category is moving **from answers to deliverables**. ChatGPT Work, Presence, Codex and Daybreak all point the same way: the output is a finished spreadsheet, slide deck, patch or resolved ticket rather than a paragraph. Marketing built around answer quality is marketing for the last product cycle.

2. **Verification is the frontier**. Nobody has solved proving that an agent did the work correctly. Task-completion reliability, not raw capability, is the competitive battleground for the next two years and the first credible public benchmark will be a positioning event.

3. **Price-performance beat frontier supremacy** this cycle. The July 2026 release positioned explicitly on efficiency rather than capability, landing within about a point of the leading rival on aggregate intelligence indices at roughly a third of the price. Efficiency is now the story the whole category will tell.

4. **Megawatts became the constraint**. AI data-center capacity near 29.6 gigawatts - roughly New York State at peak - with multi-year interconnection and transformer lead times. Power procured today determines product capability in 2028.

5. **Silicon supply is strategy**. Ten gigawatts of custom accelerators, a five-hundred-billion-dollar joint venture, and DRAM reservations approaching forty percent of global output mean the moat is now inside the semiconductor supply chain. No software company has ever competed there before.

6. **Switching costs are collapsing**. OpenAI-compatible APIs are published as a courtesy, routers make substitution a configuration change, and open weights set a price floor. Lock-in must be built above the API, in workflow, data and evaluation, or it does not exist.

7. **Enterprise buys governance**. The purchase criteria have shifted decisively toward data handling, residency, version pinning, indemnity, auditability and published error rates. This is a permanent structural shift, not a procurement fashion.

8. **Advertising is arriving** across the whole category, and the terms on which it arrives are being set in 2026. Whoever publishes constraints first defines what regulators later require of everyone.

9. **Regulation is localizing**. EU AI Act enforcement powers, state-level rules in the United States and sectoral requirements in health and education mean compliance artifacts are becoming top-of-funnel marketing content in some markets.

10. **Vertical models are proliferating** - life sciences, security, health records. Depth in a regulated vertical is now a more defensible position than breadth, and it is bought with evidence rather than with advertising.

11. **Consumer share is fragmenting**. Assistant usage is splitting across defaults, social surfaces and specialists rather than consolidating. Planning on a single winner-take-all consumer outcome is no longer supportable by the data.

12. **Trust is the scarce input**. As capability converges, the differentiator becomes whether a user believes the answer. Every monetization decision spends some of that balance, and the balance is not replenished by advertising.

13. **Talent costs are equity**, not salary. Median compensation near $608,182 with roughly $1.5 million average annual grants means labor is nearly irrelevant against silicon and electricity in the cost structure - an inversion with no precedent in software.

14. **Capital stopped picking sides**. The same investors now fund direct rivals; Amazon leads a round here while remaining the largest backer of the closest competitor. Exclusivity as a strategy is over across the entire supply chain.

8. WHERE AI MIGHT FIT INTO YOUR BUSINESS

Thirteen. Specifically, where your own technology should be used on your own marketing and go-to-market problems, which is the least exploited surface in the company.

1. **AI writes the canonical answers** at scale. The fifty highest-intent questions in the category, localized into four markets, is exactly the kind of volume work a model does well under human editorial control. (Six weeks with two editors instead of six months with an agency.)

2. **Measure citation share automatically**. A scheduled agent querying rival assistants across two hundred prompts weekly, logging which sources are cited, is a build of a few weeks and would give you the scorecard you currently lack. (Four to six weeks.)

3. **Generate the comparison hub** and keep it current. Model-versus-model comparison content goes stale in weeks; an agent that re-runs the evaluations monthly and drafts the updates makes an unmaintainable asset maintainable. (Six weeks, then automatic.)

4. **Triage the community**. An agent that reads developer forums and community platforms, clusters recurring complaints, and routes them to named humans turns an unread firehose into a weekly product input. (Four weeks.)

5. **Draft the case studies**. Fifty-two named customer stories a year is a machine problem: agent-drafted from an interview transcript and usage data, human-edited, legal-reviewed. This is the only realistic way to hit weekly cadence. (Three weeks to build the pipeline.)

6. **Localize with native review**. Model translation plus a native creative director reviewing is not the same as translated American creative; the failure mode is skipping the second half. (Ongoing, one editor per market.)

7. **Automate the procurement kit**. Security questionnaires are the highest-volume, lowest-creativity work in enterprise sales, and an agent grounded in your own approved documents can answer eighty percent of them for human sign-off. (Eight weeks.)

8. **Run pre-registered analysis** pipelines. An agent that enforces the pre-registration, runs the geo-holdout analysis as specified and refuses to re-cut it afterwards removes the largest source of self-deception in media measurement. (Six weeks.)

9. **Personalize the onboarding** path by declared use case rather than by demographic. The single largest retention lever in the product is showing a new user the one thing their job actually needs in the first session. (One quarter.)

10. **Detect brand-term leakage** continuously. An agent monitoring paid search results for your brand terms across engines and markets, alerting when a rival or affiliate outranks you, is a two-week build against a permanent leak. (Two weeks.)

11. **Summarize competitor releases** into a weekly internal brief with implications flagged by function. The company generates and consumes so much news that the marketing organization is structurally under-informed about its own category. (Three weeks.)

12. **Draft the regulatory reads**. An agent that reads incoming rules in each market and drafts the marketing-claims implications for counsel to confirm shortens the slowest step in the review workflow. (Six weeks.)

13. **Instrument the benefit metric**. Measuring whether users are better off requires classifying outcomes at a scale only a model can reach; this is the cheapest item in the plan and the only way the mission can ever be shown to have succeeded. (One quarter to define, then continuous.)

9. BRIEF HIRING IDEAS

Twelve roles, with timing. The first three are the ones that unlock everything else in this document.

1. **Permanent Chief Marketing Officer** - hire now, not next quarter. An empty top marketing seat during a listing window is itself a market signal, and every recommendation in this document needs a permanent owner to survive contact with the organization.

2. **Head of Answer Engine Optimization** - hire in Q4 2026. You are the answer engine and you are under-cited inside everyone else's. This role needs a budget, engineering support and a monthly citation-share scorecard reported to the executive team.

3. **VP Marketing Science and Incrementality** - hire in Q4 2026, before the next television upfront, so the next eight figures are measured rather than assumed. The role only works with veto authority over renewal.

4. **Enterprise Product Marketing Lead** - hire now. Someone whose entire job is turning contract terms, residency, indemnity and error rates into public pages and named-account programs. This is where the cheapest share in the business sits.

5. **Developer Marketing Lead** - hire now. Docs, quickstarts, cost calculators, migration guides and model-selection content funded as marketing rather than as engineering overflow.

6. **Four Regional Marketing Leads** - Japan, India, Germany, Brazil, hired Q4 2026 through Q1 2027 with real budget authority. The users are outside the United States and the marketing is not.

7. **Plain-Language Newsroom Editor** - hire in Q4 2026. One person whose only job is removing institutional voice from everything the company publishes.

8. **Community Manager, Named Presence** - hire now. Identified staff answering real questions in public, with a written policy and measured participation, across developer forums and the largest public communities.

9. **Marketing Counsel, Embedded** - fill immediately, internally if necessary. A securities and regulatory reviewer inside the marketing approval workflow rather than at the end of it.

10. **Documentary Series Producer** - hire in Q4 2026 to convert the best creative the company has made from a campaign into an always-on weekly programme with a fixed slot.

11. **Ad Trust and Transparency Lead** - hire now, while the ad business is small enough to constrain. Owns the load cap, the firewall audit, the category restrictions and the quarterly transparency report.

12. **Employer Brand Manager** - hire in Q1 2027. Owns candidate experience, the headcount narrative, employee advocacy enablement, and the separation of the hiring feed from the product feed.

10. OPEN QUESTIONS WORTH EXPLORING

Thirteen questions that nothing in the public record answers, ordered by how much the answer would change the plan.

1. What is the **actual gross margin** per token today, and what is the defensible curve to 2029? Nothing in the public record answers this, and it is the single number that determines whether every other plan in this document is affordable.

2. Did the market stop growing, or did you **stop being told about**? The distinction between a demand problem and a love problem determines whether the answer is distribution or product, and the evidence points to the second.

3. What are the **actual cancellation terms** inside the $1.4 trillion of compute commitments? Option value changes the entire risk picture, and nobody outside has seen them.

4. Is **advertising incremental** revenue, or is it revenue transferred from subscriptions plus a permanent trust liability? No published analysis separates the two, and the answer should govern how fast ads scale.

5. What would it **cost to buy** a genuine assistant default - an OEM slot, a carrier deal, an operating-system placement - and why has that number never been put beside the marketing budget?

6. Can **task-completion reliability** be measured credibly enough to publish? If yes, it is the most valuable marketing asset available in the category. If no, that is itself the most important product finding of the year.

7. What is the **benefit metric**, concretely? Defining and publishing whether users are measurably better off is cheap, unowned by anyone in the category, and the only way the stated mission can ever be shown to have succeeded.

8. What does the company **do about Microsoft** being simultaneously a shareholder, the largest single vendor, and a direct competitor - and when is that explained to the public in plain language?

9. Should the **listing wait** for two consecutive quarters of reaccelerating revenue and improving gross margin, and what is the cost of waiting versus the cost of listing into a flat curve?

10. Which **two things** should the company be best in the world at? Every focus-group table converged on focus as the constraint. Naming the two, funding them past comfort and partnering or shutting the rest is the decision underneath all the others.

11. What is the **honest kill criterion** for the hardware programme? Roughly $6.5 billion spent and nothing shipped; the criterion must be set before launch, while it can still be applied.

12. How much of the **consumer share loss** is recoverable by product versus how much is structurally lost to defaults? The answer determines whether to spend the next billion on brand or on distribution.

13. What is the **top idea** in the chief executive's mind, honestly, over the last eighteen months? If it has not been the product, that fact predicts the next two years more reliably than any projection in this file.

CLOSING NOTE

Total points: 231. Four items recur in every source document and in nearly every focus-group room: publish a plain pricing page, name an owner for answer-engine citation share, publish the enterprise contract artifacts as public pages, and seat a permanent Chief Marketing Officer. None of the four requires a strategy debate, a budget approval above the noise floor, or a product change. That they remain undone is the most useful diagnostic in this file.

NO POSITION / NO COMPENSATION / NO MNPI.