

SAMPLE REPORT — prepared on a public company to demonstrate format. Real client reports are confidential.

# Due Diligence Report: Perplexity AI

**Prepared by:** Verdict — a Lucid AI Labs product **Report date:** June 25, 2026 **Subject company:** Perplexity AI, Inc. (perplexity.ai) **Stage:** Late-stage private (Series E / pre-IPO) **Decision context:** Invest / secondary share purchase

## 1. Executive Summary + Verdict

### VERDICT: PROCEED WITH CONDITIONS

Perplexity AI is a genuinely impressive growth story — ARR of ~\$450M [1] in under four years, 100M+ monthly active users, and a product that demonstrably changed how a segment of knowledge workers use search. The Comet browser move is strategically sound: it shifts Perplexity from a query endpoint dependent on Google's index to an ambient computing layer with first-party behavioral data.

The three reasons that drive the verdict:

- Growth is real and accelerating.** ARR compounded from ~\$16M in 2023 to ~\$450M by March 2026 — roughly a 10x in 24 months. The February–March 2026 pricing shift drove a verified 50% single-month jump [1]. This is not projection; it is reported actuals.
- The litigation overhang is existential, not cosmetic.** The New York Times, News Corp (WSJ, Barron's, NY Post), Merriam-Webster, Encyclopedia Britannica, Nikkei, Asahi Shimbun, Reddit, and others have active suits [2][3]. The core allegation — systematic circumvention of robots.txt and "stealth" crawlers — was confirmed by Cloudflare's own research [4]. If courts mandate content licensing at scale, the cost base could restructure materially.
- Valuation premium demands execution Perplexity has not yet proven.** At ~\$22B valuation [5] against ~\$450M ARR, the revenue multiple is ~49x. Comparable SaaS businesses at IPO trade at 10–20x ARR. The gap is paid for with a growth and margin expansion thesis that depends on: (a) winning the browser war, (b) building a durable proprietary search index, and (c) monetizing without killing the product that made it.

**Conditions for proceeding:** Confirm litigation reserve disclosures in the data room; validate whether publisher program revenue-sharing (\$42.5M pool) is accounted above or below the ARR line; stress-test the valuation multiple at 25x ARR for a downside scenario.

## 2. Company Overview

---

**What they do:** Perplexity AI operates an "answer engine" — a search interface that returns cited, synthesized responses to natural-language queries rather than a list of links. The product runs real-time web retrieval against multiple large language models (currently its proprietary Sonar models, plus GPT-4o, Claude, Gemini as options on paid tiers).

**Founding & stage:** Founded August 2022 by Aravind Srinivas, Denis Yarats, Johnny Ho, and Andy Konwinski. Launched publicly December 7, 2022. As of June 2026, late-stage private, tracking toward a stated 2028 IPO target per CEO commentary.

### Funding history:

Round	Date	Amount	Post-Money Valuation
Series A	April 2023	\$26M	~\$150M (est.) [6]
Series B	January 2024	\$73.6M	\$520M [6]
Series C	April 2024	\$165M (total raised)	\$1B+ [6]
Series D	June 2025	\$500M	\$14B [6][5]
Series D ext.	July–Aug 2025	\$100M	\$18B [5]
Series E	September 2025	\$200M	\$20B [5]
Series E-6	Early 2026	undisclosed	\$21.2–22.6B [5][7]

**Total funding raised:** ~\$1.72B across 11 rounds from 62 reported investors [7]. Notable strategic investors: Jeff Bezos (personal), Nvidia, Databricks, Accel, 1789 Capital.

**Headquarters:** San Francisco, CA. Employee count: ~350 (inference [8]; not publicly disclosed).

---

## 3. Market & TAM

---

**Market definition:** AI-augmented search and knowledge retrieval. The SAM is the intersection of: (a) the global search advertising market (~\$300B+ annually, dominated by Google at ~90% share) and (b) the emerging AI assistant / productivity subscription market.

**TAM range:** Perplexity addresses a market that is actively being redefined. Traditional search TAM is ~\$300B; AI assistant subscriptions are a newer line currently at a few billion dollars in disclosed revenue across providers. Industry analysts (cited in Perplexity's own pitch materials, per reporting — inference [9]) project AI search TAM at \$50–70B by 2030 as the query-to-subscription model displaces ad-per-click.

### Tailwinds:

- Google's antitrust judgment (August 2025) weakened the "default search" distribution advantage that historically locked competitors out [10]. This is a structural tailwind, not cyclical.
- Enterprise adoption of AI-augmented research tools growing; Perplexity's Enterprise Pro tier targets this segment.
- The Comet browser puts Perplexity in the ~\$20B browser market, currently a Chromium-based duopoly (Chrome / Edge).

#### **Headwinds:**

- Google launched AI Overviews and continues to integrate Gemini into Search — the incumbent is not standing still.
  - OpenAI's ChatGPT with search, Microsoft Copilot, and Anthropic Claude all run web search features that directly compete.
  - AI search queries cost approximately 10x traditional search queries to serve [11] — margin expansion is non-trivial.
- 

## **4. Product & Technology**

---

**Core product:** The Perplexity answer engine — real-time RAG over the web, multi-LLM, with inline citations. Available on web, iOS, Android, and via API (Sonar API).

#### **Key product milestones:**

- **May 2024:** Pages — structured multi-page report generation from queries.
- **January 2025:** Perplexity Assistant — multi-app task handling (email, calendar, actions beyond search).
- **July 2025:** Comet browser launched (desktop, Chromium-based), initially Max-tier only at \$200/month [12].
- **October 2025:** Comet made free globally; hit #3 on the US App Store within 48 hours of iOS launch [12].
- **March 2026:** Comet iOS launch.
- **Feb–Mar 2026:** "Computer" agentic tool launched; pricing restructured to usage-based, triggering 50% ARR jump [1].

**Technology architecture:** Perplexity runs its own proprietary Sonar models (built on Meta's LLaMA foundation — inference [4][11]) alongside licensed access to frontier models. It maintains a partial proprietary search index (scope not publicly disclosed) and runs web crawlers for real-time retrieval. The crawler practices have been the subject of significant controversy (see Section 8).

**Defensibility assessment:** Mixed. The product experience — citations, synthesis, clean UX — is genuinely differentiated from raw LLM chat. However, the underlying capability stack (web search + LLM generation) is replicable by any frontier lab with a distribution channel.

Perplexity's defensibility lies in: (1) brand/habit among its 100M MAU base, (2) the Comet browser as a first-party data moat (if it scales), and (3) enterprise workflow integrations. None of these is a hard technical moat; all are network/distribution advantages that require continued execution.

---

## 5. Team & Founders

---

**Aravind Srinivas (CEO & Co-founder):** PhD in Computer Science, UC Berkeley; prior roles at OpenAI (research) and Google DeepMind. Strong technical credentials; has grown into a credible public-facing CEO. Known for transparent, media-forward communication style — both an asset (press, recruiting) and a risk (has made statements about scraping practices that attracted regulatory attention).

**Denis Yarats (CTO & Co-founder):** Former AI researcher at Meta (FAIR). Research background in deep RL and sequence modeling. Low public profile, which is appropriate for a CTO in this stage.

**Johnny Ho (CSO & Co-founder):** Former Quora engineer with search and ranking background. The search-specific expertise on the founding team is a genuine differentiator vs. pure LLM labs entering the space.

**Andy Konwinski (Co-founder):** Co-founder of Databricks. Brings enterprise software scaling credibility; brings the Databricks investor network. His involvement signals the enterprise channel ambition is more than marketing.

**Key-person risk:** High concentration in Srinivas. He is the face, the narrative, and the public strategy. A departure or reputational event (not hypothetical given ongoing litigation) would be material.

**Team gaps:** Public information suggests limited depth in legal/compliance leadership, which is a concern given the litigation surface area. Enterprise sales motion appears nascent relative to the ARR growth pace (inference from public hiring signals [8]).

---

## 6. Competitive Landscape

---

### Primary competitors:

Competitor	Moat	Threat level
Google Search + AI Overviews	Monopoly index, distribution via Android/Chrome default	High — but antitrust ruling created opening
ChatGPT (OpenAI)	500M+ users, top brand, browser extension, operator network	High — direct feature overlap

Competitor	Moat	Threat level
Microsoft Copilot / Bing	Enterprise O365 integration, Bing index, enterprise contracts	Medium — weak consumer brand but dangerous in enterprise
Anthropic Claude	Strong on professional research; growing web access	Medium — positioning more complementary currently
You.com, Phind, Kagi	Niche AI search alternatives	Low — smaller scale

**Positioning:** Perplexity competes on UX clarity and citation quality. It is not trying to win on model capability (it licenses frontier models) or on index comprehensiveness (Google's index remains superior). The bet is that a purpose-built "answer interface" wins mindshare with a segment willing to pay.

**Moat assessment:** Perplexity has a distribution and brand moat — not a technical one. 100M MAU is real. The Comet browser, if it retains users, converts that into a behavioral data moat over time. The core risk: Google, OpenAI, and Anthropic all have stronger models and larger user bases; if any of them make search-quality UX a priority, Perplexity's differentiation narrows. The strategic logic of the Comet pivot is exactly this: escape the "search endpoint" category before the giants close the gap.

## 7. Traction & Financial Signals

**Revenue / ARR (all figures from reported sources; flag: private company, no audited financials):**

- ARR ~\$16M: estimated 2023 [1]
- ARR \$100M: March 2025 [1]
- ARR ~\$150–200M: August–September 2025 [5][13]
- ARR ~\$305M: pre-pricing-shift February 2026 [1]
- ARR \$450M+: March 2026, post-pricing restructure [1][13]
- ARR target \$656M: internal 2026 target, per FT reporting [1] — *inference; not confirmed*

**User metrics:**

- Monthly active users: 100M+ as of early 2026 [1] (up from ~22M at start of 2025)
- Query volume: ~780M queries/month as of May 2025, up from 230M/month in mid-2024 [14]

**Business model:**

- Freemium: free tier (limited queries), Pro (\$20/month), Max (\$200/month), Comet Plus (\$5/month browser-first tier)

- As of February 2026: shifted to subscription-first model, de-emphasizing ad-supported free tier [6]
- Publisher revenue sharing: \$42.5M pool announced, distributing 80% of subscription revenue to participating publishers (Blavity, Der Spiegel, Fortune, Gannett, Time, others) [15]
- Advertising: announced but not yet launched as of report date; "related questions" ad units in development

**Burn / cost signals:** Not disclosed. AI search query costs estimated at ~10x traditional search [11]. At \$450M ARR and significant inference costs, gross margin is likely 40–60% — below SaaS benchmarks of 70–80% — though no public data to confirm. *Flag: treat any margin estimate as inference until confirmed.*

**Hiring signal:** LinkedIn shows continued aggressive hiring in engineering, enterprise sales, and legal (inference [8]). The legal hiring acceleration post-December 2025 lawsuits is notable.

---

## 8. Risk Register / Red Flags

---

*Ranked by severity (1 = highest):*

---

● **RED FLAG #1 — Content-licensing litigation: systemic and escalating** *Severity: Critical*

The scope of the lawsuit wave is broader than most coverage suggests. Active litigants as of June 2026 include: The New York Times, News Corp (WSJ/Barron's/NY Post), Encyclopedia Britannica, Merriam-Webster, Nikkei, Asahi Shimbun, Reddit, and others [2][3]. The core allegation is not a technical gray area — Cloudflare's independent research confirmed Perplexity used "stealth" crawlers that actively bypassed robots.txt blocks and hard-blocks implemented by publishers, with Cloudflare's CEO describing the behavior as resembling "North Korean hackers" rather than a legitimate AI company [4].

The non-obvious risk: if even one major case results in a precedent requiring retroactive licensing fees or injunctive relief against crawling practices, the entire retrieval architecture would require rebuilding or expensive licensing at scale. The \$42.5M publisher program appears structured partially as litigation mitigation — a fraction of the potential liability.

**Investor question:** What is the estimated litigation reserve in the current balance sheet? Are any cases approaching trial dates?

---

● **RED FLAG #2 — Valuation-to-revenue multiple demands flawless execution** *Severity: High*

At ~\$22.6B valuation against ~\$450M ARR, the implied multiple is ~50x ARR. For context: high-growth public SaaS at IPO typically commands 10–20x ARR; the most premium AI

companies (post-correction) trade at 20–30x with clear paths to profitability. Perplexity's multiple is priced for a world where it: (a) sustains 50%+ ARR growth, (b) expands gross margins as LLM inference costs decline, and (c) converts the browser into a durable monetization layer.

The non-obvious risk: the February–March 2026 50% ARR jump came from a pricing restructure — a one-time re-rating of the subscriber base, not necessarily ongoing organic growth acceleration. If the cadence reverts to 10–15% quarterly, the multiple compresses rapidly at IPO.

---

● **RED FLAG #3 — Technology moat is thinner than the narrative suggests** *Severity: High*

Perplexity is explicitly not a frontier model lab. It runs Sonar (a fine-tuned derivative of Meta's LLaMA [4][11]) and licenses GPT-4o, Claude, and Gemini as Pro options. This means its primary suppliers (OpenAI, Anthropic, Google, Meta) are also its primary competitors. The risk is not theoretical: ChatGPT's web-search feature, Google's AI Overviews, and Anthropic's web-browsing capability all compress the product surface that Perplexity defined.

The Comet browser is a smart hedge against this — it creates a distribution layer Perplexity owns. But browser market share is notoriously sticky (Chrome held 65%+ for a decade); Comet reaching 5% market share in 3 years would be exceptional.

---

● **RED FLAG #4 — Key-person concentration: Srinivas** *Severity: Elevated*

Srinivas is the product vision, the public narrative, and the investor relationship simultaneously. His public statements about crawler behavior (in particular, not committing to ceasing unauthorized scraping despite direct questioning) have been cited in litigation [2]. A reputational event, regulatory action targeting him personally, or departure would be disproportionately destabilizing.

---

● **RED FLAG #5 — Google antitrust remedy creates opportunity AND a timing risk** *Severity: Elevated*

The antitrust ruling (August 2025) against Google's default-search distribution agreements is a tailwind — but the remedy implementation timeline is unknown. If Google defaults remain in place through extended appeals for 2–3 years, Perplexity's distribution advantage from the ruling is deferred. Meanwhile, Google is accelerating AI Overviews and its own AI Mode search product with an urgency the antitrust pressure is producing.

---

● **RED FLAG #6 — Monetization model is still unproven at scale** *Severity: Moderate*

The shift to subscription-first (February 2026) improves revenue predictability but reduces top-of-funnel reach. The ad product (announced, not yet launched) is navigating a minefield: Perplexity's value proposition is ad-free, high-quality answers; introducing advertising risks exactly the user trust that drove paid conversion. The publisher revenue-sharing model (80% to publishers) is structurally generous and may be unsustainable if total subscription revenue doesn't compound to offset the content liability.

---

## 9. Key Questions for Management

---




Bring these into any diligence call:













1. **Litigation reserve:** What is the current litigation reserve on the balance sheet? Have any of the active lawsuits entered discovery or received trial dates? Is there settlement negotiation ongoing with any major plaintiff?
  2. **Gross margin:** What is actual gross margin today — accounting for inference costs (LLM API fees + own compute), crawler infrastructure, and publisher revenue share? How does it trend if LLM inference costs decline 50% over 18 months?
  3. **ARR composition:** What percentage of the \$450M ARR is subscription (recurring), vs. enterprise contract, vs. API (Sonar API), vs. publisher-program fees? Was the March 2026 50% jump primarily from pricing tier migration or from net new subscribers?
  4. **Comet retention:** What is 60-day and 90-day retention for Comet users? Browser switching has historically been a one-time event — what does re-engagement data show?
  5. **Model strategy:** What is the 3-year plan for the Sonar model? Will Perplexity build and train its own frontier model, or remain dependent on LLaMA-derived fine-tuning + licensed frontier models? What is the trigger that would cause a shift in strategy?
  6. **Antitrust upside:** Has the company quantified the distribution impact of the Google antitrust remedy? What deals are in negotiation for OEM or carrier defaults that the remedy may unlock?
  7. **IPO readiness:** The 2028 IPO target is public. What are the internal financial gate conditions (ARR, gross margin, EBITDA) that trigger the process? Is there a dual-track (IPO + acquisition) analysis happening?
  8. **Publisher program sustainability:** The \$42.5M publisher revenue-sharing pool represents 80% of pooled subscription revenue going to publishers. Is this structure compatible with the \$656M ARR target and eventual profitability? What is the cap or adjustment mechanism?
- 

## 10. Sources & Confidence Notes

---

**Confidence framework:**

-  **Confirmed** — directly reported by named publications with attributed sources
-  **Inference** — derived from multiple signals; not directly disclosed; flagged in text
-  **Unverified** — single-source or unattributed; treat with additional skepticism

#	Source	Confidence	URL
[1]	FT / MSN: "Perplexity revenue soars 50%, ARR hits \$450M"		<a href="https://www.msn.com/en-us/money/other/ai-startup-perplexity-s-revenue-soars-50-arr-hits-450m-ft/ar-AA20nAY2">https://www.msn.com/en-us/money/other/ai-startup-perplexity-s-revenue-soars-50-arr-hits-450m-ft/ar-AA20nAY2</a>
[2]	TechCrunch: "New York Times sues Perplexity AI" (Dec 5, 2025)		<a href="https://techcrunch.com/2025/12/05/the-new-york-times-is-suing-perplexity-for-copyright-infringement/">https://techcrunch.com/2025/12/05/the-new-york-times-is-suing-perplexity-for-copyright-infringement/</a>
[3]	CNBC: "NYT sues Perplexity, alleging copyright infringement"		<a href="https://www.cnbc.com/2025/12/05/the-new-york-times-perplexity-copyright.html">https://www.cnbc.com/2025/12/05/the-new-york-times-perplexity-copyright.html</a>
[4]	Wikipedia: Perplexity AI — controversies, Cloudflare confirmation		<a href="https://en.wikipedia.org/wiki/Perplexity_AI">https://en.wikipedia.org/wiki/Perplexity_AI</a>
[5]	TechCrunch: "\$200M raise at \$20B valuation" (Sept 2025)		<a href="https://techcrunch.com/2025/09/10/perplexity-reportedly-raised-200m-at-20b-valuation/">https://techcrunch.com/2025/09/10/perplexity-reportedly-raised-200m-at-20b-valuation/</a>
[6]	Wikipedia: Perplexity AI — full funding history		<a href="https://en.wikipedia.org/wiki/Perplexity_AI">https://en.wikipedia.org/wiki/Perplexity_AI</a>
[7]	Tracxn / Texau: \$1.72B total, 62 investors	 Inference	<a href="https://tracxn.com/d/companies/perplexity/_V2BE-5ihMWJ1hNb2_u1W7Gry25JzPFCBg-iNWi94XI8">https://tracxn.com/d/companies/perplexity/_V2BE-5ihMWJ1hNb2_u1W7Gry25JzPFCBg-iNWi94XI8</a>
[8]	LinkedIn hiring signals — legal/enterprise team growth	 Inference	<a href="https://www.linkedin.com/in/aravind-srinivas-16051987/">https://www.linkedin.com/in/aravind-srinivas-16051987/</a>
[9]	AI Business Weekly: TAM projections	 Inference	<a href="https://aibusinessweekly.net/p/perplexity-ai-valuation-2026">https://aibusinessweekly.net/p/perplexity-ai-valuation-2026</a>
[10]	Bloomberg: "Google Antitrust Ruling a Mixed Outcome for OpenAI, Perplexity"		<a href="https://www.bloomberg.com/news/newletters/2025-09-05/google-antitrust-ruling-a-mixed-outcome-for-openai-perplexity">https://www.bloomberg.com/news/newletters/2025-09-05/google-antitrust-ruling-a-mixed-outcome-for-openai-perplexity</a>
[11]	Medium / cybernews: AI search query cost vs. traditional search	 Inference	<a href="https://cybernews.com/ai-news/what-startup-300-ai-insiders-short/">https://cybernews.com/ai-news/what-startup-300-ai-insiders-short/</a>
[12]	TechCrunch / CNBC: Comet browser launch and free		<a href="https://techcrunch.com/2025/10/02/perplexity-comet-ai-browser-now-free-">https://techcrunch.com/2025/10/02/perplexity-comet-ai-browser-now-free-</a>

#	Source	Confidence	URL
	availability		<a href="#">max-users-get-new-background-assistant/</a>
[13]	Yahoo Finance: "Perplexity ARR tops \$450M after pricing shift"	✓	<a href="https://finance.yahoo.com/sectors/technology/articles/perplexity-arr-tops-450m-pricing-132500539.html">https://finance.yahoo.com/sectors/technology/articles/perplexity-arr-tops-450m-pricing-132500539.html</a>
[14]	DemandSage / GetPanto: query volume and MAU statistics	⚠ Inference	<a href="https://www.getpanto.ai/blog/perplexity-ai-statistics">https://www.getpanto.ai/blog/perplexity-ai-statistics</a>
[15]	Digiday / Storyboard18: Publisher program \$42.5M / 80% structure	✓	<a href="https://www.storyboard18.com/how-it-works/perplexity-ai-announces-42-5m-revenue-sharing-plan-to-pay-publishers-79622.htm">https://www.storyboard18.com/how-it-works/perplexity-ai-announces-42-5m-revenue-sharing-plan-to-pay-publishers-79622.htm</a>

---

### Analyst notes on confidence gaps:

- **Gross margin / burn:** No public data. All margin commentary in this report is flagged as inference. Request audited financials or a CFO call before any investment decision.
- **Litigation outcomes:** Outcome of NYT and News Corp suits is genuinely unknown. The risk register reflects severity of outcome, not probability. The structural scraping behavior confirmed by Cloudflare makes the legal exposure hard to dismiss as speculative.
- **User count methodology:** "Monthly active users" figures from multiple third-party statistics aggregators. Perplexity has not publicly defined MAU methodology. Treat all user counts as directional, not precise.
- **ARR trajectory:** The March 2026 50% jump is real per FT sourcing, but the base effect of a pricing restructure means sequential growth in Q2–Q3 2026 will be a more meaningful signal. Request monthly ARR cohort data in diligence.

---

*This report was produced by Verdict — a Lucid AI Labs product. All material claims are sourced or explicitly flagged as inference. This document does not constitute investment advice. For professional investment decisions, supplement with primary management interviews, data room review, and legal counsel.*

---

**Word count:** ~3,300 words (body) | **Format:** PDF-ready markdown