PRD for Perplexity AI

A product requirements document and business analysis

M.A. Watkins, TPM, BA, Consultant, © 2026

Introduction

Perplexity AI is relatively new company founded in 2022, and is based in San Francisco, CA. It still appears to be a privately held company. It has said to have gone through several VC funding rounds, including a reported $200 million. “The Nvidia backed startup has an AI browser, Comet, that can perform certain tasks on a user's behalf” (Reuters). The company CEO has been reported to be Aravind Srinivas. Andy Konwinski is also reported to be apart of this company, formerly from Databricks. Basically, their value proposition is as a website search engine, and using a web browser, powered by LLMs.*However, they have been recently involved in some civil litigation or allegation issues over web crawlers.*

Thus, this PRD will seek to help the company focus on growing their brand and products, while also minimizing their liability issues.

See:

https://andykonwinski.com/about/

https://en.wikipedia.org/wiki/Perplexity_AI

https://www.reuters.com/technology/perplexity-finalizes-20-billion-valuation-round-information-reports-2025-09-10/

Strategic Roadmap

Vision/ Goal: To increase brand name, product growth, financial viability, while minimizing product liability

Technical: Transitioning from a vulnerable web-scraping search engine (Growth) to a defensible, permission-based Agentic one (Maturity) that minimizes legal liability while maximizing utility.

See for more details:
https://docs.perplexity.ai/getting-started/overview
https://www.perplexity.ai/comet/resources/articles/comet-quick-start-guide
https://www.perplexity.ai/help-center/en

Problem/Product Analysis

- **As-is:** Growth phase with high query volume (435M/month), but has high legal friction.
- **To-Be:** Use the Comet Browser as a secure engine to move value from content extraction to task execution

AI-Native Agentic Workflows

- Using Model Context Protocol (MCP) to replace custom API connectors with a standardized, keyless protocol for local data access.
- Implementing Negative Testing in the testing phase to identify and block the circumventing of bad files before deployment.

Agentic Architecture:

- Core Logic: A cognitive architecture that separates *planning* (multi-step reasoning) from *execution* (tool use via MCP)
- Moat: Using sonar deep research models for high-accuracy synthesis
- Utilizing Transfer Engine for 400 Gbps throughput to manage real-time tab-aware multitasking in the Comet Browser.
- MCP Host connects to local servers (GitHub, SQLite, local files) to allow the AI to act without scraping third-party web views.

User Story Journey

User Story: The Agentic Travel Assistant

Story ID: PER-001

Title: Autonomous Flight Booking via Playwright MCP

As a busy professional using the Comet Browser,

I want to give a natural language command like "Book a flight"

So that I can skip the manual navigation of airline websites and proceed directly to a verified checkout.

User Acceptance Test Criteria

Acceptance criteria for AI agents must define Fidelity (accuracy) and Verification (trust).

- Intent Parsing: The system must identify "Book a flight" as an Actionable Intent (not just a search query) and trigger the Playwright MCP server.
- Autonomous Navigation: The Comet Agent must use Playwright to navigate to at least 3 major airline/aggregator sites to compare prices in real-time.
- UI Interaction: The agent must interact with dynamic UI elements (date pickers, seat maps) using the Accessibility Tree rather than visual scraping to ensure reliability.
- Source Attribution: The final summary must display the Official Source URL and a Live View snapshot of the airline’s page to ensure the user is not seeing cached or hallucinated data.
- Hand-off: The agent must pause and hand over control to the user at the final checkout screen to handle sensitive PII (Passport/Credit Card) and final confirmation.

Technical Implementation Notes (PLC & SDLC Alignment)

- SDLC (Design Phase): Define the Agentic Sandbox. The Playwright MCP server should run in a permission-isolated environment so it cannot access the user's saved browser passwords without explicit biometrics.

- PLC (Growth Phase): This feature moves the company from a Search Engine to a Utility Engine, helping by saving the user an average of 8–12 minutes per booking task.

Liability Mitigation (DevSecOps)

- Legal Guardrail: Implement a Verification Trust Score. Mandatory cross-referencing of three independent, non-paywalled sources before surfacing a summary.
- Security: Implementing SOC 2 Type II and SSO for enterprise users to ensure data isn't used for training.
- Real-time monitoring of Model Drift and Factuality Scoring to prevent fabricated content (hallucinations) attributed to publishers.

Highlights for Stakeholders

- The To-Be Vision: The company is not just building a better search engine, it is building a Cognitive Architecture where the user provides intent (book a flight) and the system provides the result.
- Liability Mitigation: By moving to Phase 1 (Publisher APIs) and Phase 4 (Zero-Click Verification), it transforms the legal risk into a competitive moat—publishers become partners instead of plaintiffs.
- Velocity: Leveraging AI-Native coding within the SDLC allows a move from prototype to global rollout in 3-month cycles.

Conclusion

By Perplexity AI moving from a high-risk "Scraper" model (subject to legal churn) to a "Symbiotic" (high-margin utility), the following to-be analysis:

- Valuation & Revenue (Est. Jan 2026): Market Cap: Est. $20–$28 Billion valuation

  Revenue Growth: Projected $656 Million Annual Recurring Revenue (ARR) by end of 2026—a 10x increase from 2024.

  Efficiency: 60% gross margin by optimizing open-source models (Sonar) and reducing reliance on expensive third-party APIs like OpenAI.

- The "Symbiotic" Moat:

  By integrating via Publisher APIs and Marketplace Partnerships (e.g., Airtel, Snap), Perplexity replaces scraping with permissioned transactions.

  By using Enterprise Max tier, it creates a zero-liability revenue stream by indexing internal corporate data, where copyright is less an issue.

* * Financial figures are only estimates and forecasts, do not reflect true state of company.