

Dossier

verified decision intelligence

Decision-grade intelligence, in days.

The AI strategic analyst for any business - fan-out research, adversarial verification, cited synthesis. Verified before it reaches the page.

a DropFly product · June 2026

Dossier · verified decision intelligence, in days

THE PROBLEM

Strategic decisions are bet-the-company. *The research behind them isn't.*

Leaders choose between two bad options: a consulting firm that takes six weeks and six figures, or a chatbot that answers in seconds but can't be trusted. One is too slow to afford; the other is too shallow to bet on.

2

\$80k+

a single market-entry assessment from a top firm

6 weeks

before you get the answer you needed today

4

0

verification behind a raw chatbot answer

THE DIFFERENCE

A chatbot gives an answer. *Dossier gives a verdict you can act on.*

Dossier fans out across dozens of sources, then tears its own findings apart - every claim faces independent skeptics and survives a vote or it's killed. What reaches you has already been cross-examined.

Verified, not generated. That is the entire product.

4,5

HOW IT WORKS

Six moves, every time.

01 Scope

Turn the decision into 4-6 falsifiable sub-questions.

02 Fan out

Dozens of sources per question - not one search.

03 Verify

Every claim faces skeptics. Survives a vote, or dies. The moat.

04 Synthesize

A cited report: bottom line first.

05 Human QA

A senior reviewer signs off before anything ships.

06 Deliver

Report + live readout: where to win, where not.

THE MOAT

Verification is the whole product.

Claims extracted

113

Put to the test

25

Confirmed

17

Killed before the page

8

In one real engagement, 8 of 25 tested claims were killed. A chatbot would have published all of them.

WHAT YOU GET

Every engagement is turnkey.

Interactive report site

A cinematic, navigable web report you can send as a link.

Comprehensive report

A full PDF - methodology, data, reasoning, every source.

Investor-ready deck

A cinematic pitch deck built from the same findings.



PROOF

One question, answered honestly.

The engine produced a full, cited market assessment of an AI logistics venture end to end - market sizing, the competitive teardown, the legal spine, the financial model, and the honest recommendation (don't build yet; validate). Every figure traces to a source.

It even told the client what NOT to do. That honesty is what a real report is for.

ANY SECTOR

Swap the question. Keep the rigor.

Energy

SaaS

Retail

Manufacturing

Healthcare

Logistics

Finance

Real estate

The engine is industry-agnostic. The rigor never changes.

THE VALUE

Same decision. A fraction of the cost.
In days, not weeks.

^{2,3}
\$12k

a Dossier assessment vs ~\$80k from a firm

days

to a decision-grade verdict vs six weeks

Priced to the value of the decision, not the hours.

Dossier · verified decision intelligence, in days

PRICING

Four ways to buy.

Brief

\$2,500

3 days

Assessment

\$12,000

~1 week

Deep

\$35,000

2-3 weeks

Strategy desk

\$7,500/mo

ongoing

A top firm bills ~\$80k over six weeks. Dossier delivers the same decision-grade output from \$2,500 - in days. Prices are fixed; broader scope is itemized in add-ons.

WHY NOW

The moment, and the method.

Frontier models can finally do the research

But raw model output is confident and often wrong - useless for a real decision.

The differentiator is the harness, not the model

Fan-out + adversarial verification + senior QA is what turns an answer into a verdict.

DropFly runs the engine and signs off the output

You get the speed of AI with the accountability of a human reviewer on every report.

BEGIN

Decide with evidence.

Know before you bet the company on it.

dropflyai@gmail.com

Dossier · verified decision intelligence · a DropFly product

Dossier · verified decision intelligence, in days

SOURCES

The market behind Dossier.

Every market fact in this deck traces to a source - click to open.

1. McKinsey - The economic potential of generative AI (\$2.6-4.4T/yr, concentrated in knowledge work)
<https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>
2. Yahoo Finance - Trends shaping the \$541B management-consulting market
<https://finance.yahoo.com/news/trends-shaping-541-billion-management-081300950.html>
3. AlphaSense - AlphaSense surpasses \$500M ARR - the category pays for decision-grade research
<https://www.alpha-sense.com/press/alphasense-surpasses-500m-in-arr/>
4. Perplexity - Perplexity Deep Research - free, but weak on investment-grade specifics (the free floor and its limits)
<https://www.perplexity.ai/hub/blog/introducing-perplexity-deep-research>
5. Bessemer (BVP) - AI pricing playbook - advice without the loop is soft-ROI; verified decisions command price
<https://www.bvp.com/atlas/the-ai-pricing-and-monetization-playbook>
6. BusinessWire - Gamma \$100M ARR / \$2.1B - lean AI-native scale is achievable
[https://www.businesswire.com/news/home/20251110805751/en/Gamma-Surpasses-\\$100M-ARR-Raises-at-\\$2.1B-Valuation-as-It-Replaces-PowerPoint-for-the-AI-Era](https://www.businesswire.com/news/home/20251110805751/en/Gamma-Surpasses-$100M-ARR-Raises-at-$2.1B-Valuation-as-It-Replaces-PowerPoint-for-the-AI-Era)