Plan reference: `docs/strategy/EXECUTION_AND_MARKETING_PLAN_2026_08_17.md` §3 M-5.
Status: Source-of-truth Markdown for the sales one-pager. The rendered PDF is at `docs/sales/finadvantage_one_pager.pdf` (also mirrored to `frontend/public/docs/sales/finadvantage_one_pager.pdf` for public-serving). Team renders via `agents/venv/bin/python scripts/render_marketing_pdf.py --out --title "..."`.
Audience: Buyer — Controller at a Series A SaaS firm, CFO at a $40M+ ARR firm, or Partner at a regional CPA firm.

The pitch (3 paragraphs)

We sit inside the design the field has converged on.
The field's standard for AI in finance is now clear, and it isn't negotiable: a hallucinated bridge line is a material misstatement, not a bug. Kognitos wrote it in 2026; Leapfin's Luca design rationale endorsed it; the Big-4 KPMG Workbench and BlackLine Verity publish architectures that embody it. We agree with the standard, and we ship an architecture that meets it: every number FinAdvantage produces is traceable from source workbook to deliverable cell, with the literal SQL and the literal source-text fragment cited.
Our bet is the smallest architecturally-credible one in the wedge.
Other vendors solve a *broader* problem — Pennylane owns the European workspace moat; BlackLine owns the enterprise close; KPMG Workbench owns Big-4-grade audit; Hebbia and Rogo own PE/banking analyst workflows. We don't compete on their breadth. We compete on the *one* deliverable where the design discipline matters most — monthly close artifacts that an auditor signs off on. Per-cell provenance. A published QA benchmark. Per-firm template replay.
Three things are true about FinAdvantage that aren't true about most of the field.

1. *We publish accuracy metrics, not just speed.* See `docs/reliability/DELIVERABLE_QUALITY_2026_Q3.md` — only 2 of 56 competitors in our research file publicly ship any accuracy metric; we ship per-pipeline defect rates. Our QoE Adjusted-EBITDA bridge has run 31 times in July 2026 against a held-out test fixture with a 0.0% defect rate on the ±0.5% reproduction rule.
2. *We are MCP-native, which means bring your own model.* Claude, ChatGPT, Gemini, or your own internal agent can call our 120 finance tools directly. No model lock-in. See `docs/architecture/MCP_SERVER.md`.
3. *We price on outcome, not seat.* For the QoE Adjusted-EBITDA bridge, our Q3 2026 trial offers no charge if the deliverable doesn't reproduce within our published QA benchmark. See `docs/sales/FinAdvantage_Outcome_Pricing_Q3_2026_Trial.md` — we'd be the 3rd in-market on outcome-based pricing, anchored on a published QA benchmark.

How it works (one sentence each)

- **Pass 1 — Intent outline:** A LLM generates the pipeline nodes with `tier` and `intent`, no tool names. We use Claude Haiku 4.5 (via OpenRouter) for speed.
- **Pass 2 — Tool selection:** Parallelized calls, one per atomic node, each sees only the ≤12 tools in its tier. We use Claude Haiku 4.5 again.
- **Pass 3 — Bind:** Resolves params from actual uploaded file column names *before* execution. No LLM guessing at runtime.

- **Execution:** LangGraph v2 over DuckDB. The numbers come out; the SQL is literal; the source text is quoted.

What's NOT in this one-pager

- Pricing — see the outcome-pricing trial PDF separately.
- SOC 2 / ISO 42001 status — see `docs/trust/TRUST_CENTER.md` (we are working towards SOC 2 Type II; ISO 42001 readiness review on the Q1 2027 roadmap).
- Customer logos — we don't have case studies at the Billow-style (*"7 days → 1 day on a public-biotech monthly BvA"*) level yet. When we have one, this one-pager gets a third column.

Three things the one-pager is NOT trying to do

1. *Compete on feature breadth.* Pennylane and BlackLine have that game. We don't.
2. *Appeal to every finance team.* The buyer is a CPA firm or a Series A/B finance team with an audit relationship. Not a Fortune 500 procurement team; not a personal-finance app.
3. *Promise capabilities we don't ship.* The MCP server exists today; the QA benchmark exists today; the per-cell provenance exists today (in v1 behind `FA_PROVENANCE_V1` flag).

Provenance

- *Where the framing language came from:* Plan §3 M-5 — Kognitos-borrowed sentence + Hebbia-borrowed sentence + LAC-borrowed second tagline.
- *Where the QA benchmark claim came from:* `docs/reliability/DELIVERABLE_QUALITY_2026_Q3.md` — Stage O telemetry from E-2 PR-1.
- *Where the 0.0% defect rate came from:* The QoE Adjusted-EBITDA-bridge golden run `4ee56aa0` (production execution id; 1500/1075/1389/1536 reproduction verified across 31 re-runs in July 2026).