



Omnigraphs: workflow OS and trust infrastructure for transactional law

VERTICAL GRAPH RAG PURPOSE-BUILT FOR TRANSACTIONAL
LEGAL PRACTICE AREAS

Zero-hallucination • Real-time, single-source-of-truth matter management •
Domain-native, workflow-embedded legal AI

MAY 2026

FOUNDERS@FORMATIVELEGAL.COM

INTERACTIVE
DECK →



Eliminating hallucinations from LLMs is mathematically impossible. Every legal AI product on the market is built on top of that impossibility.

Xu et al. (2024) proved it using computability theory. OpenAI's own researchers confirmed it in September 2025. This is not a quality problem that better training will fix; it is a structural property of how probabilistic models work.

#1: WORKFLOW FRICTION

LOAD, UNLOAD, REPEAT

Dig for the right document, attach it, prompt engineer, pray the context window holds, wait for a conversational answer, repackage that answer into something actually usable. That is not a sustainable workflow — it is a workaround in the absence of anything better.

#2: TOKEN BURN

POSING DETERMINISTIC QUESTIONS TO PROBABILISTIC MODELS IS SENSELESS

LLMs re-read and reconstruct context every query, and subsequent responses hallucinate on prior hallucinations. That is a fundamental limitation of vector indexing, especially when applied to large volumes of unstructured data.

#3: <100% ACCURACY

WHERE'S WALDO?

Even in a world in which endless plugins, add-ons, skills and fine-tuning manage to improve LLMs to ~99.9% accuracy (which no one is currently able to claim)... if the user doesn't know where the 0.1% is, all 100% needs to be checked. Especially in legal, where one oversight can easily cascade and irreparably erode client trust.

Non-lawyers think Harvey, Legora and Claude Legal are specialist. The lawyers who use them daily know they are not.

Practitioners know that these models cannot be trusted to output work to the standard of a junior associate: nuanced, detailed, and domain-specific. Lawyers are bearing the brunt of "wrapper tax" — the three problems above all stem from the same root cause: that the AI, data shape and user interface are not purpose-built for the domain. And when the domain in question is transactional law with deal values in the millions/billions, retrofitting generalist frontier models through endless patching is simply too piecemeal and haphazard to be a genuinely usable solution.

Introducing the Omnigraph.

Our product platform, Matrix, structures and expresses **every moving part** in any given matter — whether a term, clause or email — in a **live-updated, single-source-of-truth** graph that all human and agentic queries are **first routed to**. The Omnigraph is **centralised** and **authoritative** as it pulls **live data** from across tools such as DMS (e.g. iManage), emails, DocuSign, etc.

The Omnigraph **deterministically resolves ~85%** of queries with provenance (hence the **zero-hallucination guarantee** assuming accurate ETL pipelines enforced by the human-in-the-loop); only queries it cannot answer are ever **re-routed to the probabilistic LLM layer, flagged** as such and **confidence-scored**. This graph-first infrastructure makes LLM outputs trustworthy.

KNOWLEDGE GRAPH (MASTER/REFERENCE DATA LAYER) + OPERATIONAL GRAPH (REAL-TIME UPDATES LAYER) = SINGLE-SOURCE-OF-TRUTH OMNIGRAPH

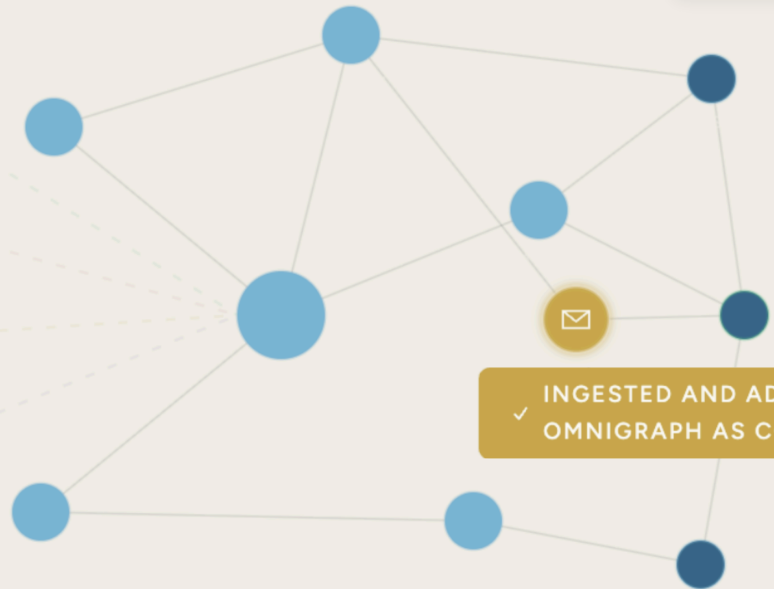
QUERY 1

QUERY 2

INBOUND EMAIL

NEGOTIATIONS

MATTER OMNIGRAPH



VIEW LIVE ON MOBILE
Scan to see the Omnigraph animation in your browser — best on desktop.

GRAPH-ANSWERABLE PORTION

- 100% ACCURATE
- FULLY PROVENANCED
- ZERO TOKENS

LLM PORTION

- FLAGGED AS <100% ACCURATE
- CONFIDENCE SCORED

LLM FALLBACK LAYER

Four side-effects, each a fundable business in its own right.

The Omnigraph presents unexpected compounding upsides. The result: easier sales, multiple revenue streams and lower COGS.

AGGREGATE (BOTH NAMED AND ANONYMISED) INTELLIGENCE AS NEGOTIATION LEVERAGE


Thomson Reuters

...but fine-grained for each practice area, firm, historical counterparty. Every matter closed on Matrix gets aggregated to the firm's own intelligence layer and anonymised into the platform-wide market intelligence layer. Predict counterparty behaviour and market trends grounded in provenanced historical data. Even firms that never adopt the core product would pay 6 figures to access this real-time intelligence.

MULTI-PARTY COLLABORATION

The truth, the whole (or partial?) truth and nothing but the truth

The benefit of a single source of truth is that each stakeholder can view the slice of that truth that's pertinent to them, strictly permissioned to the highest InfoSec standards. Whether it's an investor counterparty wishing to DocuSign on Matrix, fund ops being handed over legal obligation monitoring key dates for the next 10 years or external counsel reviewing a redline, this is what lawyers asking for "better collaboration" meant all along.



AGENTIC READINESS

Domain-native harness engineering

Autonomous agents need a deterministic graph layer to query. Harvey shipped agents March 2026; Legora is betting on agentic head-first with "Legora aOS". Matrix firms get agent-ready infrastructure as a byproduct of running the core product.

MARGIN ESCAPE

40-60%

That is how much time-per-matter compresses with Matrix. Clients have been pressuring law firms for years to move off billable hours onto fixed/value-based pricing, and Matrix finally allows firms to capture productivity as margin. This is the upgrade path from labour-arbitrage to software-economics.

Whatever a firm bets on, we'll always be the correct parallel bet.


Firms are placing very different bets on AI: some have bought Harvey, some are building directly on Anthropic, some have already deployed and been disappointed, others are waiting and seeing. Whichever it is, Matrix is non-negotiable infrastructure alongside.

HARVEY/LEGORA

BUILD ON ANTHROPIC /
OPENAI

WAITING TO DECIDE

+



= Harvey or Legora's legal reasoning capabilities, constrained by what's confirmed in the matter Omnigraph. Defensible output, far less token spend.

= Matrix has already built the domain-specific substrate that in-house engineers would otherwise have to build themselves, brick by brick, by dragging fee earners from their desks to be interviewed.

= Deploy Matrix now and plug in whichever model wins, or even Matrix's own model offering. The need for trust infrastructure doesn't change.

MODEL-AGNOSTIC

No frontier model is a safe long-term bet: GPT-4 was the future two years ago. But law firms will always be liable for what their AI produces, and professional liability doesn't come with a confidence interval.

MATRIX BETS ON THE LIABILITY, NOT THE MODEL.

DIRECTION-AGNOSTIC

Every Harvey output lawyers produce without provenance is exposure firms can't defend. "We already have Harvey" isn't an objection to Matrix.


IT'S A REASON TO DEPLOY MATRIX YESTERDAY.

PRACTICE AREA-SPECIFIC

Horizontal AI tools optimise for breadth. Matrix optimises for transactional law's specific deliverables: LPAs, side letters, MFN matrices, credit agreements or SPAs.

THE VERTICALITY IS THE MOAT.

We're launching with fund formation, finance (fund and LevFin), and corporate M&A as proof of concept — the densest corners of a \$4.81tn global M&A market and \$570bn corporate legal segment, into a legal AI category growing 28% annually.



BEACHHEAD SAM • 3 PRACTICE AREAS AT LAUNCH

\$3 billion

Fund formation, finance, and corporate M&A at the top 200 transactional firms globally. **Underlying market:** ~\$200bn in annual legal fees across these three PAs (AmLaw 100 alone: \$158bn revenue, ~50% transactional). **Tech budget addressable:** 3–5% of fees = \$6–10bn current, with matter intelligence ~25% subset.

Sources: AmLaw 100 2025, Mergermarket FY2025, Chambers UK 2025, Legal 500 UK 2026

TAM • ALL TRANSACTIONAL LEGAL PRACTICE AREAS

\$8 billion+

All transactional commercial law globally — real estate, project finance, capital markets, restructuring. **Corporate legal segment: \$570bn globally**, growing 6.4% CAGR — the highest of any segment. Transactional ~60% = \$340bn in annual fees. Tech budget addressable: \$8–12B+ by 2030 at legal AI's 28% CAGR.

Sources: Grand View Research 2024, Mordor Intelligence 2025, McKinsey Global Private Markets Report 2026

↗ **28.3%**

LEGAL AI SOFTWARE CAGR
2025–2030

MarketsandMarkets, 2025

↗ **31.8%**

CONTRACT DRAFTING + REVIEW
AI CAGR

MarketsandMarkets, 2025

↗ **+41%**

GLOBAL M&A DEAL VALUE
2025 VS 2024 (\$4.81TN)

Mergermarket FY2025

↗ **11% → 30%**

AI ADOPTION IN LEGAL
2023 TO 2024 (ONE YEAR)

Bloomberg Law / Industry Research

Centari (deal intelligence, M&A, \$14M Sept 2025) and **Ontra** (PE legal, \$310M raised) validate that the transactional legal market pays for structured intelligence. Neither is building trust infrastructure. Neither is graph-first. **Harvey** (now ~\$11B valuation) validates the broader legal AI category.

The window is open: whoever ships graph-native first compounds away from everyone else.

THE WINDOW IS OPEN...

VALIDATION

The buyer is no longer hypothetical.

Harvey ~\$11B. Legora \$100M+ ARR.
Centari \$14M. Ontra \$310M+. **Willingness to pay is settled; architecture isn't.**

GAP

No trust layer exists.

Every funded player is LLM- or document-first. **None graph-native.**
Harvey's \$200M Series D went to agents, not foundations.

FUTURE

Agentic in legal is < 18 months out.

Harvey shipped agents **March 2026.**
Deploying agents without a deterministic trust layer is liability waiting to happen.

REGULATION

Audit trail is becoming a regulatory baseline.

EU AI Act high-risk obligations live **2 Aug 2026.** SRA demands explainability across 200,000 E&W solicitors.

...AND THE MOAT COMPOUNDS FROM DAY 1.

IMMEDIATE

Every matter feeds the next.

Every matter produces an Omnigraph — **a training signal for the next matter, firm, practice area.** Data accumulates inside the architecture, not on top of it.

WITH EVERY EDGE CASE

The ontology deepens.

Every fund structure, deal type, side letter extends the ontology. **A 2027 entrant will need years of edge cases to catch up.**

~0 AT M1 • VERY HIGH AT Y3

Switching costs compound.

1,000 matters = 1,000 living omnigraphs feeding the precedent library. **Migrating off means walking away from years of institutional knowledge.**

WITH EVERY COUNTERPARTY

Counsel network propagates.

Every cross-firm transaction is a **Matrix-to-Matrix handshake.** The network spreads sideways through deals, not top-down through sales.

The real moat isn't GraphRAG as an architectural decision...


anyone with a whiteboard and sufficient domain knowledge can sketch the graph schema specific to each practice area. It's the **structured data crystallising inside that architecture, matter by matter** — institutional memory competitors would need years of edge cases to replicate. **Whoever ships graph-native first, owns the compounding.**



Workflow-critical by design. Recurring revenue by nature. Intelligence features compound cost of switching.

£4.6m Y1 ARR*

3 PRACTICE AREAS TARGETED FOR LAUNCH IN Y1 • 12 ANCHOR FIRMS




MONTHLY LICENCE FEE
£1,200 per user/month
 flat across boutique, mid-market and elite/mega firms

ONE-OFF SETUP FEE
£320 – 430K
 total on the 12-firm Y1 book at signing
 £15-20k (Boutique) • £20-30k (Mid-market) • £50-60k (Elite/Mega)

Y1 ANCHOR FIRM PROJECTIONS

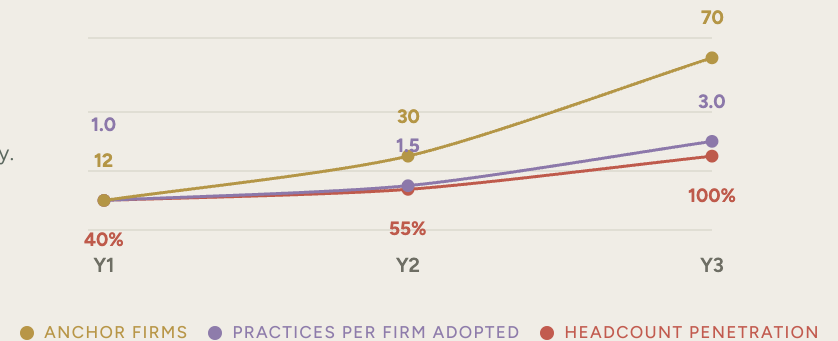
	×2	×3	×1
FUND FORMATION	£115K	£432K	£403K
FINANCE	—	£461K	£1.04m
M&A	—	£691K	£1.44m



GO-TO-MARKET STRATEGY

- FOUNDER-LED** Both founders ex-Mid-Market/Powerhouse • warm partner-level intros across the 30-firm Y1 funnel.
- DESIGN PARTNERS** Pilots refining the product • pilot firms converted to anchor M6 onwards.
- ADVISOR NETWORK** 3–5 senior partners • ~3 warm doors each, plus reference credibility.
- INBOUND ENGINE** Conference circuit + category-creation thought leadership.

GROWTH LEVERS




Each lever is multiplicative — growth is back-loaded.

*Ultra-conservative — does not include ARR from firms purchasing the market intelligence layer as a standalone product (see Slide 4, Upside #1), revenue streams from multi-party usage (e.g. fund ops, see Slide 4, Upside #2), or regional/independent firms with smaller transactional practices.

Matrix was always going to exist. It just needed repeat co-founders with a 7+ year track record of building together — and their exact combined backgrounds to bring it to life.

Two BAME, female, socially mobile lawyers. One bedroom startup founded 7 years ago that became the multi-award-winning social mobility charity, STRIVE Talent. Now building the intelligence layer that transactional lawyers have needed for decades. Between them: the transactional legal experience and the technical vision to build what nobody else has thought to build in quite this way.



Sana Shafi

CEO • PRACTISING FUNDS LAWYER (4 PQE, EX-KIRKLAND)

Six years at top fund practices in the City; formation of funds with AUM in the billions is her bread and butter. She has done, manually, everything Matrix is built to replace. Sana does not have a theory about how transactional lawyers work: she is one. She has sat in every closing, negotiated every side letter, and managed every MFN process that Matrix handles. When she says the current workflow is broken, she means she lived it — just last Thursday. She is the domain expert on call who shapes design and schema decisions, tightening the feedback loop from months to minutes.



Bertilla Chow

CHIEF PRODUCT & TECHNICAL OFFICER • SOLICITOR & BARRISTER TURNED SOFTWARE ENGINEER

Chancery-trained barrister. Silver-circle trained solicitor. 42-trained software engineer with systems/architecture depth and product/design instinct. She has acted for Deutsche Bank, Goldman Sachs and Barclays on transactions up to £350M, then as Chief of Staff at a fast-scaling professional services firm drove 40% revenue growth while digitising its entire client base. Her capability as a high-resolution translator between Sana's domain expertise and her varied skills stack means we can execute at breakneck velocity, with crystal clarity of vision, and stay leaner for longer.


£1 million

PRE-SEED • TARGETING ROUND CLOSE END OF Q3 2026 • MINIMUM £800K

OPERATIONS & RUNWAY

25% • £250K

18-month runway plus the infosec stack institutional firms demand to sign — SOC2 Type I → II, ISO 27001 by M18. Founder salaries, legal, finance, insurance included.



PRODUCT & ENGINEERING

45% • £450K

Graph schema and UI build for the 3 target practice areas for launch. Lead Engineer secured and ready to start from Day 1, working alongside domain experts in fund finance, leveraged finance and M&A. Funds, Finance and M&A MVP leading into v2.0 iteration.

GO-TO-MARKET & FIRST CLIENTS

30% • £300K

GTM Lead hire plus the pilot-to-paying motion. Converts 12 anchor firms across fund formation, finance and M&A into £4.6m ARR by Y1 close.

The first-mover position in legal trust infrastructure is ripe for occupation by lawyers who've lived the problem and are ready to build an all-in-one solution. Join us before the window closes.

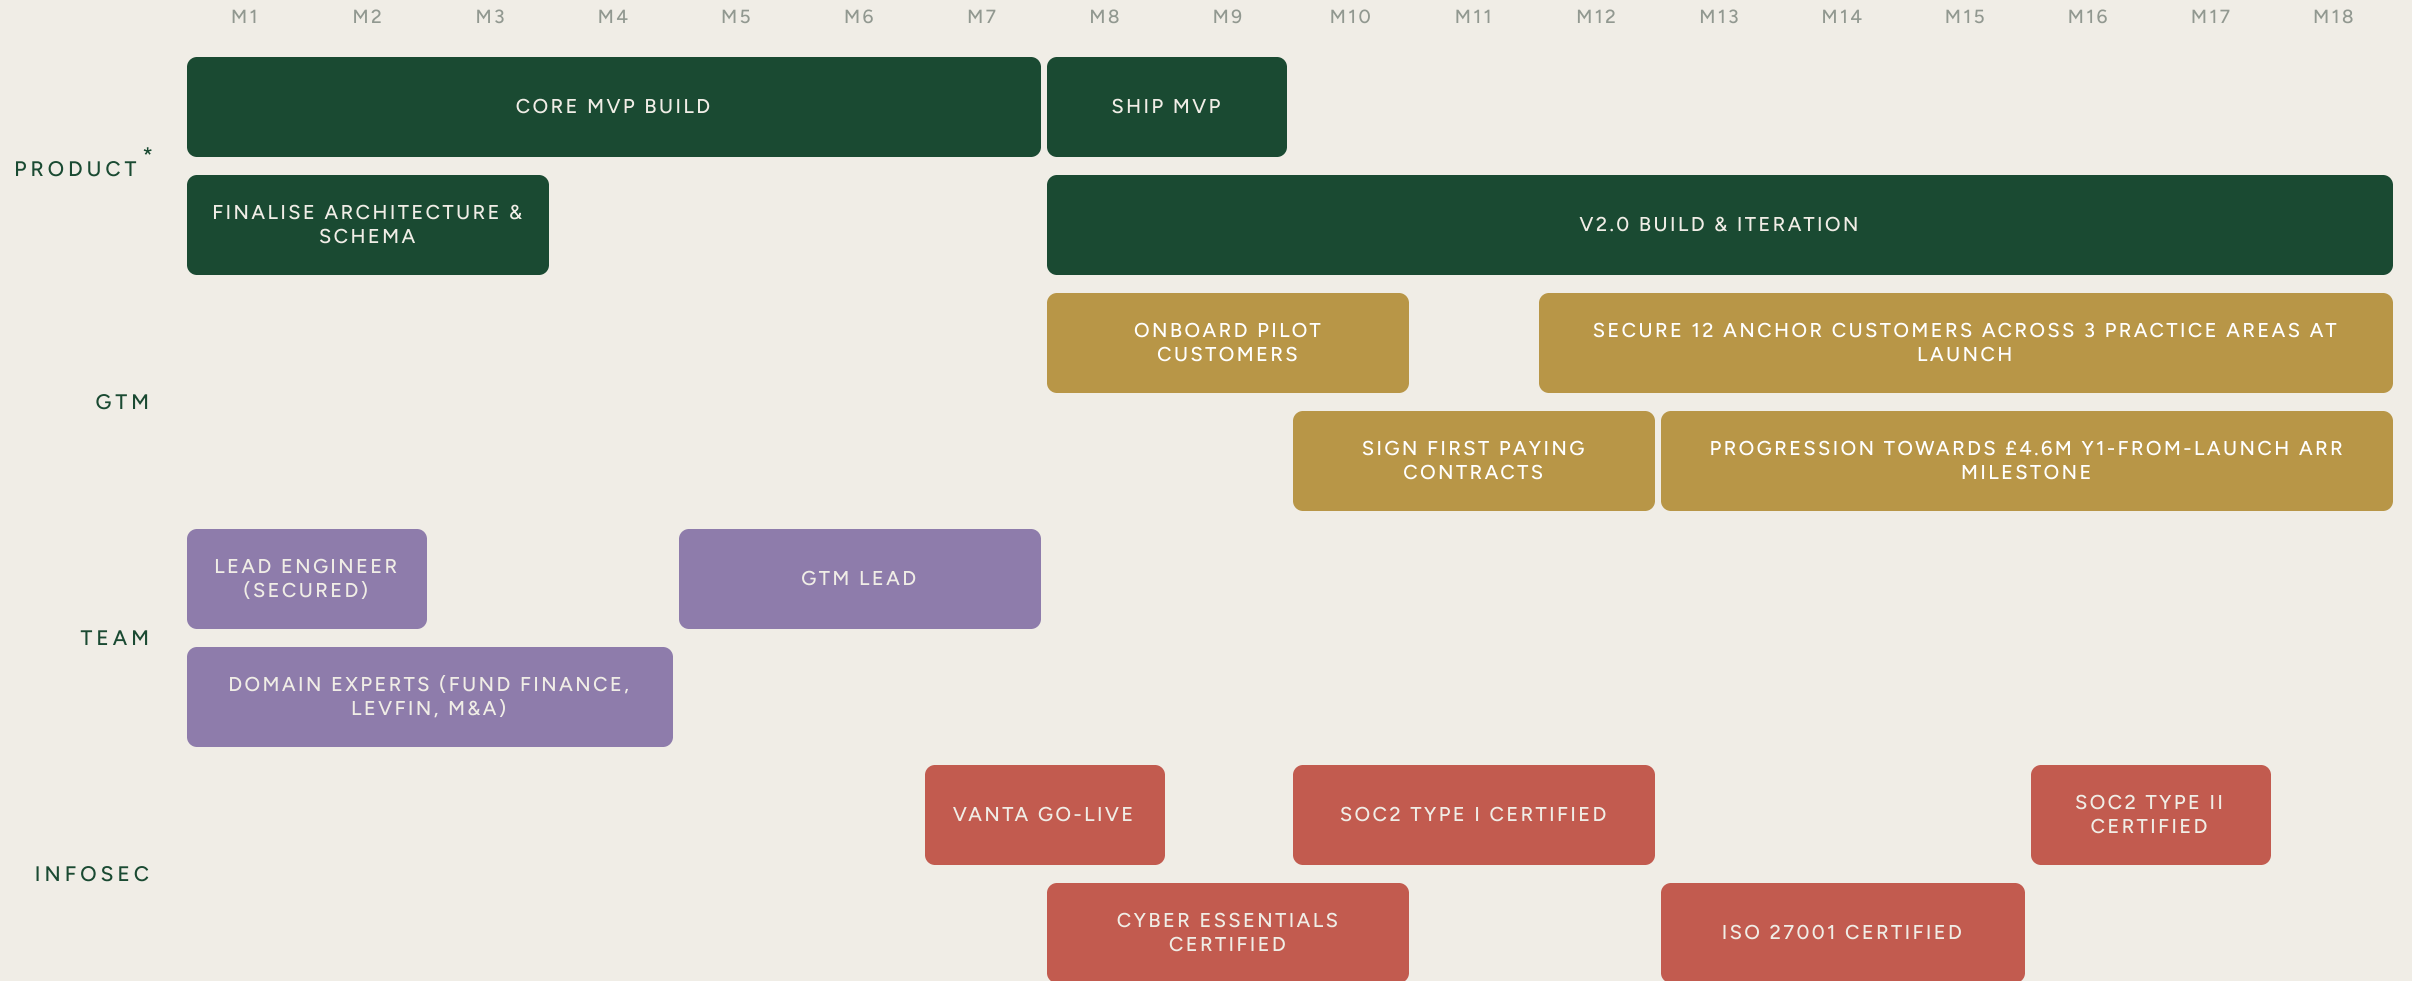
FOUNDERS@FORMATIVELEGAL.COM • BOOK A CALL

Appendices

- A** The next 18 months: projected milestones
- B** Competitive landscape: Harvey vs Legora vs Matrix
- C** Standard RAG vs GraphRAG: the change-of-control worked example
- D** Information security, compliance & data residency

The next 18 months: projected milestones

Product, GTM, team and infosec workstreams to first paying contracts, £4.6m ARR and SOC2 Type II.



* Parallel across all three practices targeted for launch: fund formation, finance (fund finance and LevFin), and M&A.

Same buyers, different layer.

"How do you compete with Harvey/Legora?"

We don't — they're competing with each other. We're in the deal workflow that both of them eventually need to call into.

Harvey¹

Horizontal LLM • Founded 2022 • \$11B valuation

ARR **\$190m**

CUSTOMERS **1,000+ law firms[†]**

~50 AmLaw 100 + 400+ mid-sized + ~550 smaller/regional

Plus ~500+ in-house corporate teams — separate count

BUYER **Innovation team**

WEDGE

Firm-wide productivity layer. Sold per seat across every practice — contract review, drafting, research, due diligence.

Legora²

Horizontal + differentiating through agentic focus • Launched 2024 • \$5.55B valuation

ARR **\$100m+**

CUSTOMERS **1,000+ firms + in-house[†]**

50 markets • firm/in-house mix not publicly broken out

Named BigLaw: White & Case, Linklaters, Cleary, Goodwin

BUYER **Innovation team**

WEDGE

Agentic workflows on top of foundation models. Per-seat. Wide coverage of legal tasks, shallow per workflow.

Matrix

Graph-native, domain-specialist, workflow-embedded trust infrastructure for transactional practice areas

ARR **£75m+** Y3 TARGET

CUSTOMERS **70 enterprise firms** Y3 TARGET

Enterprise tier • top 100 by revenue + Magic Circle + boutique + regional firms with transactional deal flow

BUYER **Transactional partners**

WEDGE

Workflow-critical depth inside specific transactional practices. Per-seat licence sold to the transactional practice group, not the firm-wide AI budget. Adopted alongside Harvey/Legora, not instead.

SAM The universe of target commercial law firms — top 100 firms by revenue globally — totals roughly **~120 firms**. Harvey and Legora's 1,000+ counts inflate via mid-sized firms (Harvey: 400+) and in-house teams; their actual enterprise base sits closer to **50–100 firms each**. Matrix's Y3 target of rollout across **70 client firms** is a meaningful share representing the transactional-heavy segment of this finite, well-defined market — plus selective boutique + regional firms with transactional deal flow, counted on a separate line and additive to the 70.

[†] Counting methodology not publicly disclosed. **Working assumption:** at BigLaw firms, office budgets are set locally, not centrally — even when enterprise licensing is centralised firm-wide. The 1,000+ headline likely registers each office as a separate customer; one firm with offices in London, New York and Singapore reads as three. The most plausible reconciliation against a finite global BigLaw universe.

¹ Harvey, "Helping Law Firms and Companies Collaborate at Scale," 13 Mar 2026; Harvey, "How Harvey Helps Mid-Sized Law Firms Scale Legal Work," 24 Oct 2025; CNBC, 25 Mar 2026 (\$190m ARR, \$11B valuation).

² Legora, "Legal teams' adoption of AI propels Legora past \$100M in ARR," 2 Apr 2026; Bessemer Atlas, "Legora: The fastest enterprise business to reach \$100M ARR" (named BigLaw customers); TechFundingNews on \$5.55B Series D.

Standard RAG vs GraphRAG: the change-of-control worked example

A worked example showing why semantic similarity alone is insufficient for legal drafting — and how Matrix's hybridised signal closes the gap.

No Limited Partner may Transfer all or any part of its Interest without the prior written consent of the General Partner, which may be withheld in its sole discretion, if following such Transfer the transferee would, together with its Affiliates, hold an aggregate Interest representing more than ~~5%~~ **10%** of the total Commitments of the Fund.

Every word is identical except one number. Standard RAG computes a similarity score of **~99%** because the surrounding language — Transfer, Limited Partner, Interest, Affiliates, prior written consent, sole discretion — is so dominant that the 5% vs 10% difference barely registers.

Matrix's GraphRAG architecture is structured differently. It computes **two scores**:

1. **Semantic similarity** (vector-based): **~99%** — the vector index, co-located alongside the graph index, identifies these clauses as near-identical.
2. **Threshold equivalence** (graph-based): **FALSE** — a clause where **CoC_threshold = 0.05** is plainly not equal to a clause where **CoC_threshold = 0.10**.

STANDARD RAG SEES ONLY THE FIRST SIGNAL. MATRIX SEES BOTH.

This hybridised signal is what gives counsel the nuance they need in negotiations and drafting.

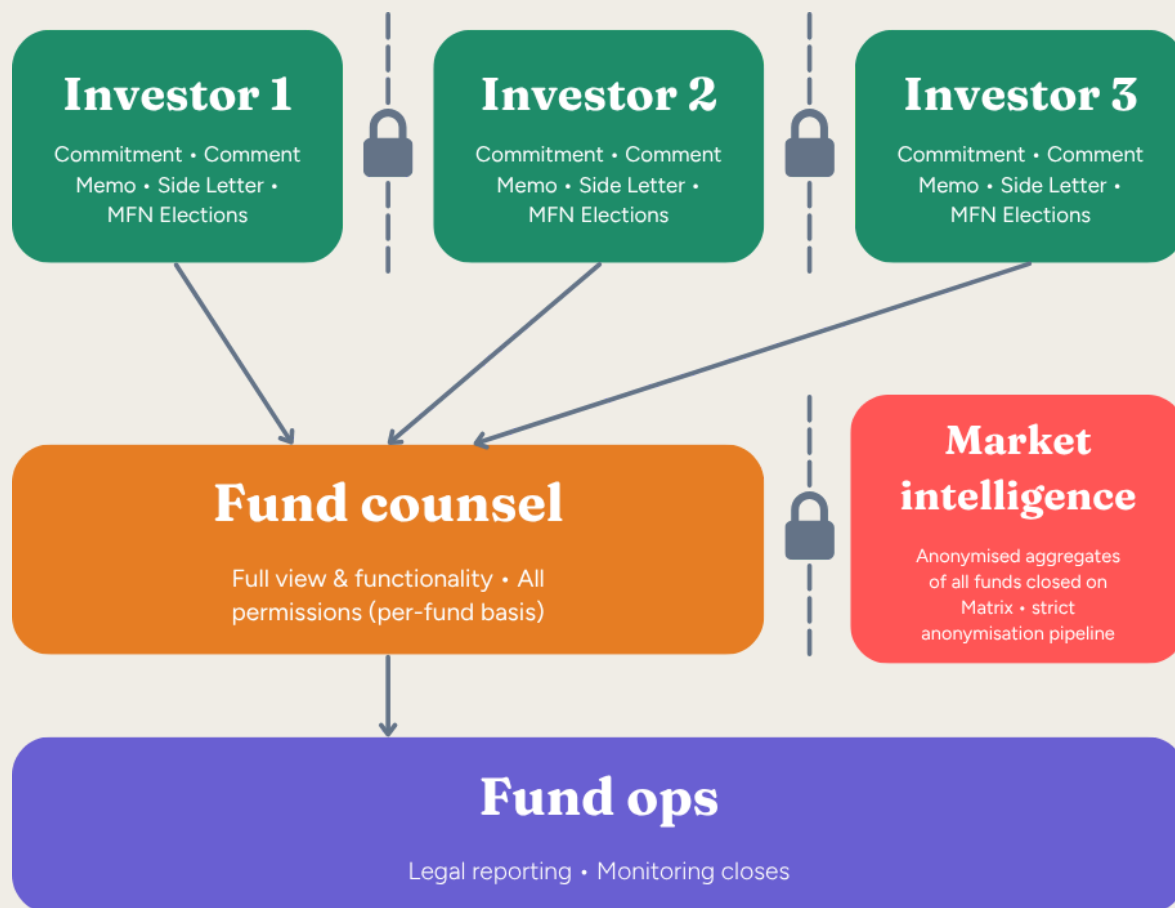
As counsel draft a change-of-control clause, Matrix surfaces semantically comparable provisions from the firm's own clause repository — ranked by how often each has been agreed in practice.

"Matrix recommends this formulation: you have accepted it in 95% of comparable mandates. Here are four alternatives if you want to consider them."

INSTITUTIONAL KNOWLEDGE BECOMES A LIVE DRAFTING INPUT, NOT AN ARCHAEOLOGICAL EXCAVATION.

Information security, compliance & data residency

Isolation enforced at the architecture layer — not just the UI — with role-based access and EU / UK / US data residency. Here is an example diagram of how InfoSec would look for fund formation, but the same principles apply across Finance and M&A.



INVESTOR-LEVEL DATA ISOLATION

Each investor's side-letter terms, negotiation history and concessions are stored in discrete, permissioned nodes. No investor can query, view, or infer another's position. Confidentiality is enforced at the architecture layer, not just the UI.

ROLE-BASED ACCESS CONTROL

Every user operates within a defined permission tier — Fund Counsel (full fund view), Fund Ops (operational data only), Investor Counterparty (own data only). Queries requiring graph traversals beyond each user's permissions will be denied. Access is logged, timestamped and auditable.

DATA RESIDENCY

Matrix's cloud infrastructure is deployable with EU-based data centres (AWS Frankfurt / Dublin or Azure Netherlands / Ireland) for GDPR-bound funds, UK regions for post-Brexit clients, and US-East / US-West for US clients. Region selection is configured at onboarding and documented in the Data Processing Agreement.