formative
LEGAL

INTERACTIVE DECK →

MATRIX: MATTER OMNIGRAPHS

Context + trust infrastructure for deal-shaped lawyers

Real-time single-source-of-truth Matter Omnigraphs, ready to **plug-and-play** into Harvey, Legora, Claude or any other **existing legal LLMs** law firms already use. **Category-defining deterministic state machine** for every deal with **built-in provenance, auditability** and **agent-ready permissioning** for lawyers running **cross-border transactions**.

PURPOSE-BUILT FOR $100M - $10BN+ DEALS

MODEL-AGNOSTIC CONTEXT INFRASTRUCTURE LAYER

JUNE 2026

INTERACTIVE DECK →

You can't make lawyers run billion-pound deals within a chatbot.

Every legal AI on the market asks **a probabilistic model** to **hold the state of a live matter** — every term, threshold, election and side-letter concession in flux across months, with consequences cascading into years down the line. **A context window is the wrong tool for holding state. State needs a substrate to live in**; a context window isn't one. The three problems below are symptoms of asking **a reasoning layer** to behave like **an authoritative matter record**.

#1: WORKFLOW FRICTION

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**.

LOAD, UNLOAD, REPEAT: THE FRICTION CYCLE MEANS HIGH LICENCE COUNTS, LOW DAILY USE

#2: LOSE-LOSE-LOSE

Posing deterministic questions to probabilistic models is senseless: you pay for tokens and wait for half-hours to play "Where's Waldo?". Even at 99.9% accuracy, if you don't know where the 0.1% is, **all 100% must be checked**. That is a **fundamental limitation of vector indexing**, especially when applied to large volumes of unstructured data.

BURN TOKENS, BUY HALLUCINATION RISK & UNAUDITABILITY IN RETURN

#3: INSUFFICIENT DOMAIN DEPTH

By the time a probabilistic model finishes burning tokens on what should've been deterministic all along, it's **run out of context window** to do **the real work that requires nuanced judgment**. This is why practitioners know that currently available products **cannot deliver redlines or insights to the level of a junior associate**.

"VERTICAL" AI FOR LAWYERS HAS PROVEN TO BE LESS VERTICAL THAN ADVERTISED

Buy vs build... or a new third way: the "build-on"?

The jury is filing back into court and the verdict has been reached: Kirkland & Ellis is investing **$500m** into building custom tooling in-house. Freshfields is **building directly with Anthropic**. Although there are only a handful of firms globally that have the revenue to justify this spend, these market signals point strongly in the same direction: that currently available solutions are simply **not deep enough into the vertical** to create **genuinely usable outputs** for practice area specialists who do this work day in, day out. As the elite powerhouse firms with generous innovation budgets are racing to build from scratch and the mid-market firms feel forced to buy generic tooling dressed as "vertical" off-the-shelf, we've been quietly building **a third way that serves both buyer personas**.

INTERACTIVE DECK →

Introducing plug-and-play Matter Omnigraphs

AND WHY IT'S A BIG DEAL FOR BIG DEALS.

Matrix gives the legal AI stack a **live, permissioned matter-state layer**. Firms keep the AI interfaces lawyers already use — Harvey, Legora, OpenAI, Claude and internal agents — while Matrix becomes the **authoritative graph** those tools call **before** they reason, draft, compare, escalate or act.

The Omnigraph resolves known facts **deterministically**, with source, permission, version and decision lineage attached. Only the remainder routes to the probabilistic model layer, **bounded by current matter context** instead of loose documents. **MCP makes the connector lightweight**; the product depth lives in deterministic state, permissioned retrieval and reusable matter context.

VIEW LIVE ON MOBILE

Scan to see the Omnigraph animation in your browser — best on desktop.

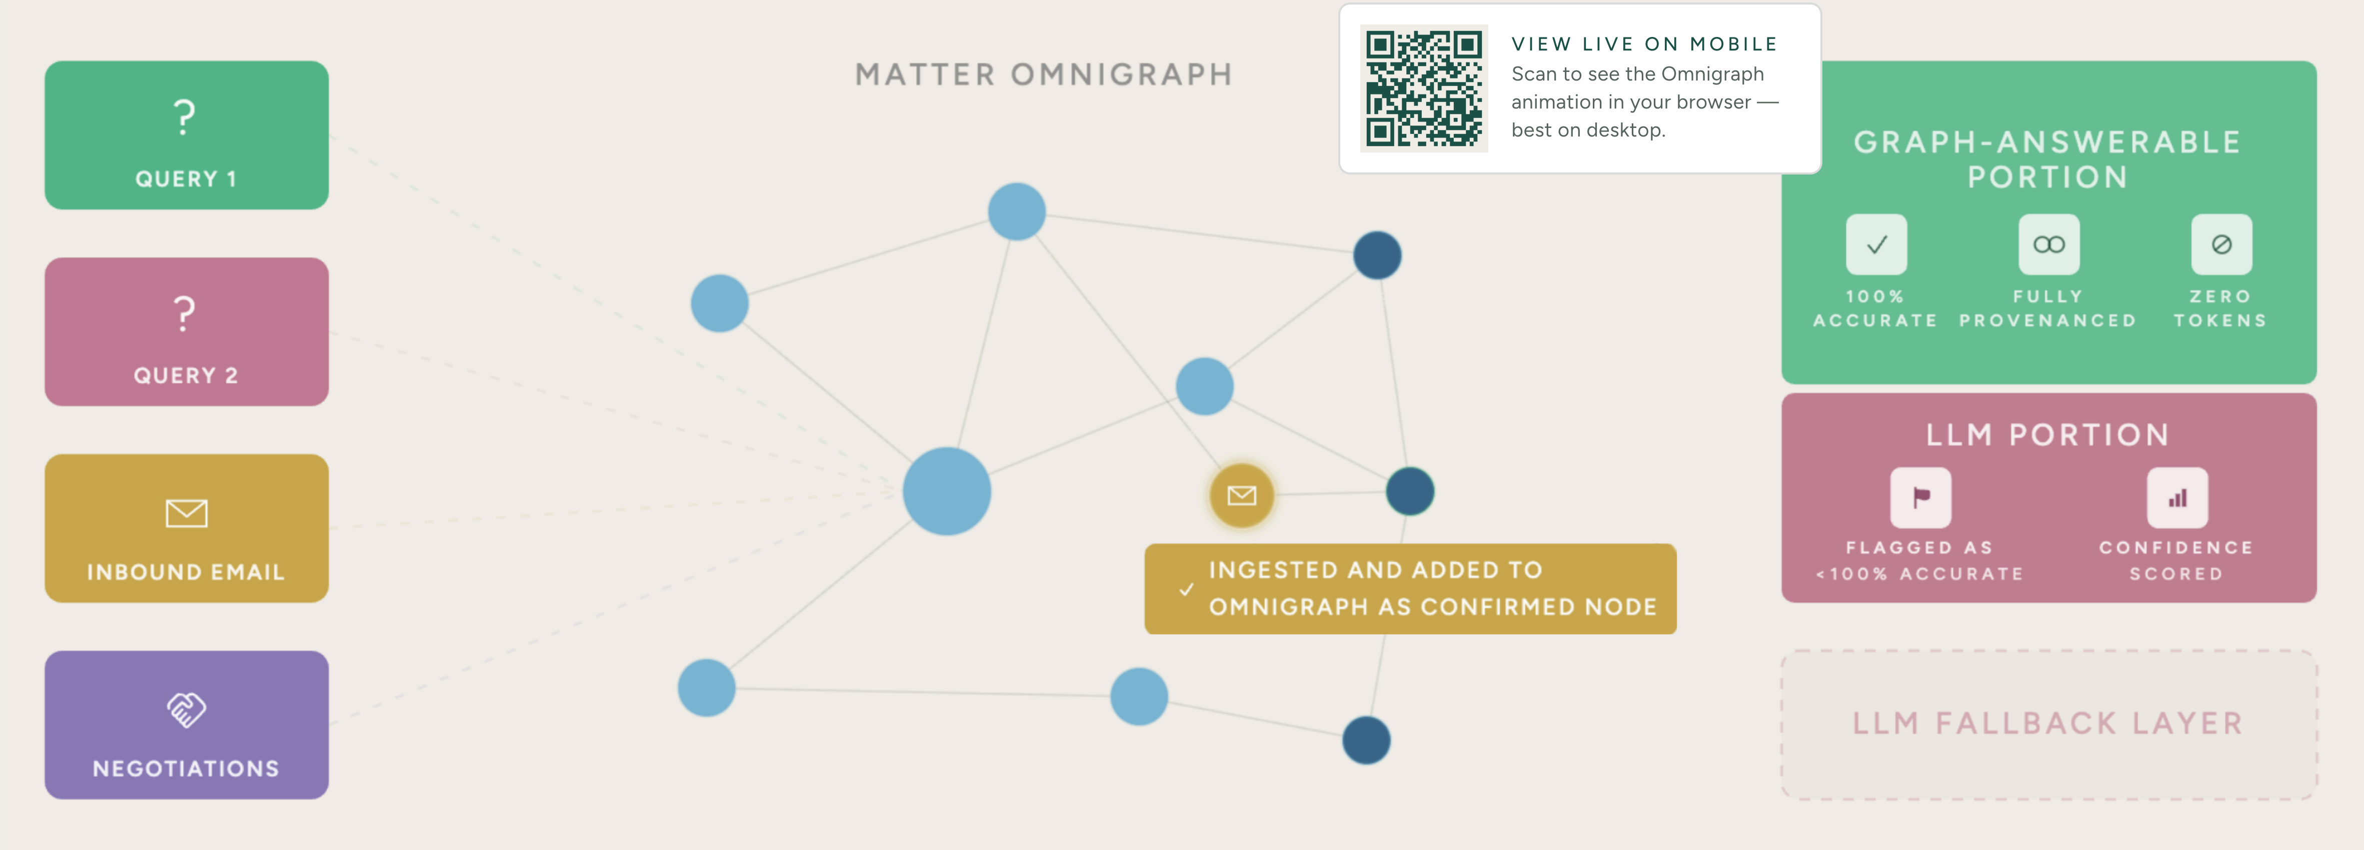

3. PRODUCT ARC: V1 TO V3

INTERACTIVE DECK →

Our v1 MCP server is just the wedge. Our v2 Deal Engine & v3 market benchmarking capabilities are the company.

We enter where firms already have approved AI interfaces. The same Matter Omnigraph architecture then compounds into workflow execution, and ultimately into firm memory no model or point solution can recreate.

Architectural throughlines

- **SINGLE SOURCE OF TRUTH**
  Answers, workflows and firm memory all call on the same immutable record.
- **BUILT-IN AUDITABILITY**
  Source trails are attached from ingestion, not reconstructed later.
- **STRICTLY PERMISSIONED**
  Matter data stays sealed, region-hosted and scoped per user.
- **AGENT HARNESSING**
  Retrieval becomes workflow execution, then institutional memory.
- **CROSS-MATTER INTELLIGENCE**
  The same schema unlocks client and practice-area insight.

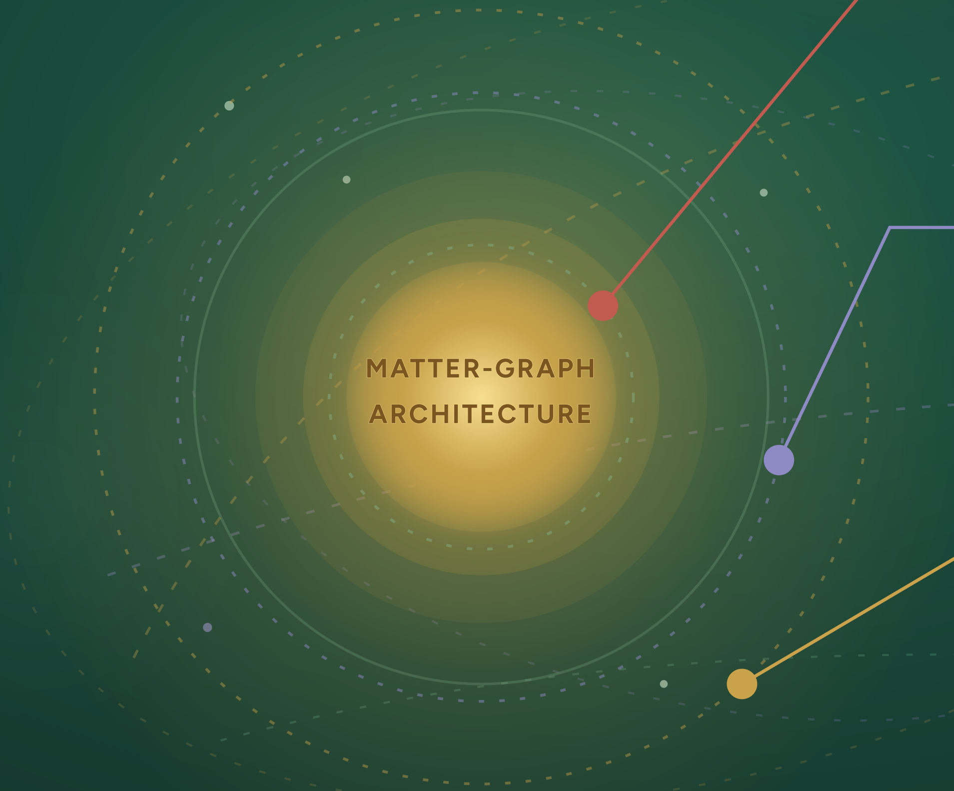

V1: MCP SERVER • WEDGE

Existing AI gets the matter graph: live state, permissions, provenance and graph-grounded answers.

V2: DEAL ENGINE • WORKFLOW UI

The same graph runs workflows: drafting, trackers, eSigning, agents and practice-specific deal engines.

V3: FIRM BRAIN • CROSS-PRACTICE AREA QUERYABLE INTRANET

Structured knowledge becomes searchable across matters, clients and practice areas.

INTERACTIVE DECK →

Harvey and Legora unlocked the distribution for Matrix to ride on. We are the infrastructure layer they have to call into.

The legal AI budget has already been unlocked by **horizontal reasoning platforms**. That is not a threat to Matrix; it is **the insertion point**. As firms standardise on approved AI interfaces, agents need a deterministic context infrastructure layer underneath them: live matter state, permissions, provenance and auditability before they draft, compare, escalate or act.

SPRAWLING TOOLS ARE SYMPTOMS

point tools • one task, one workaround

Contract review
Add-on
Drafting
Word plug-in
Chatbot
Wrapper
Legal research
Summariser
Due diligence
Negotiation assistant
Legal analytics
DMS
eSignature
Document automation

COLLAPSE INTO
get absorbed

MATTER OMNIGRAPHS

Matrix is the substrate

Every workaround on the left exists because no one owns matter state. Matrix is **not another layer on the stack; it is the ground the stack stands on**: auditable, permissioned, workflow-embedded, and purpose-built for transactional practice.

CALL INTO
query, not replace

THE PLATFORMS BECOME OUR TENANTS

funded, horizontal reasoning layers

Harvey
HORIZONTAL LLM

Broad productivity and reasoning, sold practice-wide. Powerful distribution, but **no record of what has actually been agreed in the matter**.

Legora
TABULAR • AGENTIC

Agentic workflows surface complex relationships; they still need a **liability-grade deterministic layer** for agents to query.

Frontier models
ANTHROPIC • OPENAI

Reasoning keeps improving. Domain state does not appear by magic. **Build the substrate yourself, or call Matrix.**

MCP gives Matrix a lightweight insertion point. Distribution first, substrate ownership second, workflow expansion third.

INTERACTIVE DECK →

No amount of fine-tuning models can make up for lack of structured, deterministic, real-time deal context.

The wedge is intentionally thin: plug into Harvey, Legora, Claude, OpenAI and firm-built agents as the matter-state and provenance layer. The moat compounds as each deployment creates **reusable transaction memory**, **workflow precedent** and **benchmark intelligence**. MCP is the connector; the durable asset is the graph, audit trail, ontology and cross-matter memory.

PILLAR #1

System of record adjacency

Matrix sits next to DMS, billing, email and deal-room systems, so it captures **live matter state** rather than static prompts or copied documents.

PILLAR #2

Audit-grade provenance

Every answer can carry source, permission, version and decision lineage. That makes the layer useful for **risk, privilege and client accountability**.

MATRIX

Trusted context layer

One firm-controlled layer normalises the messy context that models need: who can see what, what changed, which precedent matters, and where the deal actually stands.

KNOWLEDGE GRAPH	entities, documents, clauses, people and obligations linked to each matter.
OPERATIONAL GRAPH	tasks, approvals, status, deadlines and handoffs captured as work happens.
TRUST GRAPH	permissions, provenance, audit trails and client-specific guardrails enforced before retrieval.

COMPOUNDING FLYWHEEL

MATTERS LAND > MEMORY IMPROVES > BENCHMARKS COMPOUND

PILLAR #3

Workflow precedent

Each transaction teaches playbooks, approval paths and exception patterns. The product gets better at **how the firm actually executes**.

PILLAR #4

Model-agnostic distribution

MCP makes Matrix callable from the tools firms already adopt. When models churn, the **context, ontology and trust layer stay**.

INTERACTIVE DECK →

No, they won't: why we're the only one with the full hand.

	ONTOLOGY	ARCHITECTURE	CROSS-FIRM INTELLIGENCE	INCENTIVE		WHY THEY WON'T
Powerhouse firms, building in-house ONLY A HANDFUL (THE HIGHEST-GROSSING FIRMS GLOBALLY): KIRKLAND • FRESHFIELDS	✓ Deeply domain-expert lawyers	✓ Budget + talent to build	✗ Data never leaves firm's walls	✗ Off-mission; billers off deals	=	IT'S OFF-MISSION. Powerhouse firms can build internal systems, but they cannot create a cross-market context layer. Data stays inside the firm's walls, and fee earners get dragged off live matters to define infrastructure instead of doing deals. Matrix gives every other firm a foundation to build on.
LLM wrappers, fine-tuned for lawyers HARVEY • LEGORA	◐ Legal-general, not transactional-deep	◐ Retrofit data lakes into schema?	◐ Unstructured + siloed volume	✗ Off-thesis to rebuild from first principles	=	IT'S OFF-THESIS. Their forward story is agents and reasoning, not rebuilding deterministic foundations underneath every live matter. "We already have Harvey" is reason to deploy Matrix, yesterday: the reasoning layer needs an authoritative matter-state layer to call.
Frontier models ANTHROPIC / CLAUDE FOR LEGAL • OPENAI	◐ "Legal" as monolith is the vertical	◐ Retrofit data lakes into schema?	◐ Unstructured + siloed volume	✗ Diametrically antithetical	=	IT'S ORTHOGONAL. Frontier models make reasoning cheaper and more capable. They do not create live matter state, permission lineage or practice-area ontology by magic. Build the substrate yourself, or call Matrix.
Matrix GRAPH-NATIVE • TRANSACTIONAL-SPECIALIST	✓	✓	✓	✓	=	WE WILL — AND WE'RE ALL IN. The only player whose **entire business is the trust layer itself** — and **the parallel bet alongside whatever a firm has already bought**.

INTERACTIVE DECK →

We're not in the AI race. We're **the toll on the road** charging **almost nothing, on everything.**

≈1% of the firm's fee per mandate, across three revenue streams — a **~$5bn TAM**. £3.2m today → £54m by Y3, ~9% of the $800m beachhead.

Y1 lights a thin first layer; the headroom is the adjacent practices ($1.3bn SAM) and the full firm universe ($5bn TAM).

~$3–7bn
INTELLIGENCE LAYER • A 2ND MARKET

~$800m
SOM • 3-PA BEACHHEAD
3 streams • core ~$0.56bn (≈1% of the $56bn fee pool)

~$1.3bn
SAM • YOUR FIRMS
3 streams • core ~$0.9bn (≈1% of the $90bn fee pool)

5–10%
MULTI-PARTY FAN-OUT • LPs + COUNSEL

~$5bn
TAM • ALL FIRMS • ALL TRANSACTIONAL PAs
3 revenue streams (see next slide):
• Core ~$3.4bn — ≈1% of the $336bn fee pool
• Multi-party fan-out • Market Intelligence

LAUNCH — YEAR 1

- **FUND FORMATION** — LPAs • side letters • MFN
- **FINANCE** — Leveraged + fund finance
- **CORPORATE M&A** — Private + public deals

ADJACENT EXPANSION • ROADMAP

- **REAL ESTATE** — Acquisition + finance
- **CAPITAL MARKETS** — ECM • DCM • securitisation
- **PROJECT FINANCE** — Infra • energy
- **RESTRUCTURING** — Multi-creditor

Legal services $1.05tn × transactional ~31% ≈ $336bn (Grand View / Mordor / GM Insights 2025–26); legal analytics ~$3.15–7.7bn (Mordor / ResearchNester). Take rate derived from price card; expansion shown as upside, held at zero in base case.

8. BUSINESS MODEL & FINANCIAL PROJECTION

Every practice area unlocks three revenue streams.

INTERACTIVE DECK →

ARR TRAJECTORY (CONSERVATIVE BASE) • £M

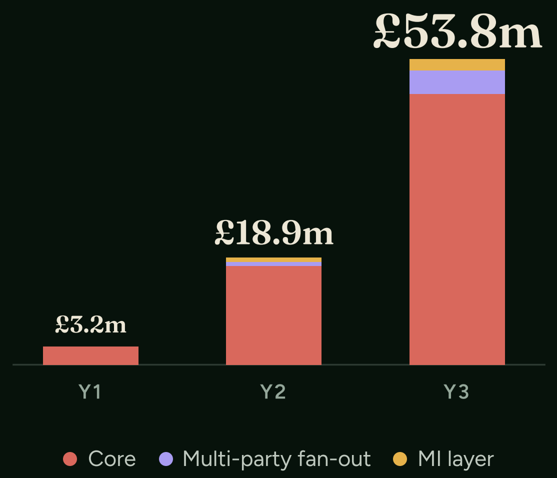

GROWTH = DEPTH × BREADTH

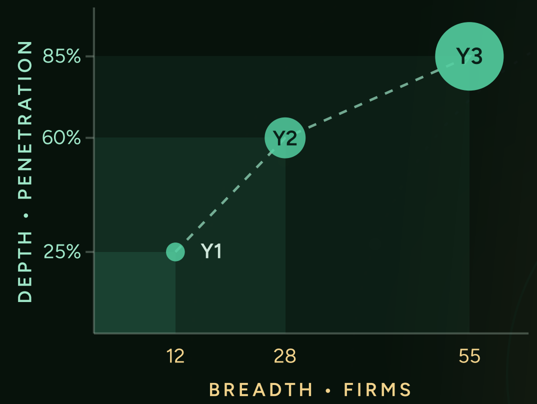

Bubble = ARR (£3.2m → £54m). Both levers compound — growth is back-loaded.

CORE • law firms

Value-based (per-matter) pricing, not per-seat. Blended **~1% take-rate** on every mandate onboarded onto Matrix.

£47.7m Y3

FAN-OUT • counterparties

Cross-sell to the counterparty universe within **each multilateral practice area**.

High fan-out potential: fund formation (many LPs), project finance (sponsor + lender syndicate + ECAs + gov + contractors), restructuring (multi-creditor), syndicated/LevFin.

Medium: capital markets, real-estate JVs.

Low: bilateral M&A and real estate.

£4.1m Y3

PRICING — VALUE-BASED, CAPPED AT WILLINGNESS-TO-PAY (WTP) CEILING

PER MATTER, NOT PER SEAT	
Small	£5k
Medium	£10k
Large	£20k
XL	£25k

MATURE ACV CAP • PER FIRM	
Boutique	£250k/yr
Mid-market	£700k/yr
Elite	£2m/yr

≈1% of the firm's fee — the WTP ceiling

Plus a **one-off setup fee to cover ingestion and onboarding costs** — excluded from ARR above

MARKET INTELLIGENCE LAYER • principal & intermediary universe

Market intelligence & benchmarking as a standalone offering • blended ~£50k per subscription. LPs, GPs, banks & corporates already spending on PitchBook, Preqin and Bloomberg terminals.

£2.0m Y3

INTERACTIVE DECK →

As the substrate, every current market force becomes a tailwind for us.

Build the substrate, and it doesn't just collapse the point-solution sprawl — **it throws off value the sprawl never could**. The same Omnigraph that wins the core deal spins off four more, each pulled into existence by a market force already in motion.

FORCE #1: AGENTIC IN LEGAL IS < 18 MONTHS OUT

Harvey shipped agents March 2026; Legora is leaning heavily on their "aOS". **Agents need a deterministic layer to query** — or they're just **ticking liability time-bombs**.

→

"Vertical within a vertical"-level harnesses to constrain agentic workflows

Firms get **agent-ready infrastructure as a byproduct** of the core product — and every agent vendor becomes **a tenant of the graph, not a rival to it**. MCP makes that context layer callable from the AI front doors firms already bought.

FORCE #2: THE MARKET HAS NEVER BEEN MORE WILLING TO PAY, BUT ALSO NEVER MORE FATIGUED FROM ENDLESS LLM-ON-LLM PATCH SOLUTIONS

Legal-AI spend is compounding **~28% a year** — but it's scattered across ten overlapping point tools that don't hold. **Buyers are done paying for patches**; they want **one thing that works**, and rightly so.

→

Omnigraphs are accidentally the all-in-one solution

GitHub for lawyers, matter tracker, AI-assisted predictive drafting, agentic onward actions, market-wide Thomson Reuters fine-grained for each practice area, firm-wide negotiation analytics, or precedent library... **we happen to be all of the above. We bank the budget firms were fragmenting across patches** — the good reasoning tools still plug in; narrow patches and point-solutions **won't survive the consolidation**.

FORCE #3: THE REGULATORS ARE WATCHING — LIKE A HAWK

The SRA already holds **200,000 E&W solicitors to a duty of competence**, and "the model hallucinated and we don't know how" is **no defence to a negligence claim or referral**. The EU AI Act adds pressure (fines to **€15M / 3%**) — though its high-risk obligations, proposed for 2 Aug 2026, may slip to Dec 2027 and aren't settled.

→

Auditability as a by-product of the architecture right from ingestion, not reconstructed post-hoc by some skill or plugin

Append-only and immutable, the Omnigraph is the audit trail: every fact permanent and sourced the moment it's ingested and confirmed. The sooner a firm adopts, the longer the period covered by the auditable trail. **Nothing to rebuild**, no plugin burning tokens poring over reams of conversation history to certify itself five years later.

FORCE #4: CLIENTS HAVE BEEN FORCING FIRMS OFF THE BILLABLE HOUR FOR YEARS

Under the billable hour, **working faster just means billing less — efficiency is a pay cut**. Announcing its $500m in-house platform, **Kirkland's chair framed it as a shift off the billable hour** toward value-based pricing — the top-grossing firm making our case.

→

40-60% margin escape

Matrix is projected to **compress time-per-matter 40-60%**; firms bank it as margin instead of writing it off. **Labour arbitrage becomes software economics.**

INTERACTIVE DECK →

3 qualified lawyers and an ontology-obsessed engineer between 2 repeat co-founders: Matrix was always going to exist, it just needed the right team.

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 in practice 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**.

INTERACTIVE DECK →

£1.5 million

PRE-SEED • TARGETING ROUND CLOSE END OF Q3 2026

~18-month runway — we raise the seed at **~M12** from a position of strength, not necessity. Cash covers the full plan; the raise is a choice, with ~6 months of cushion behind it.

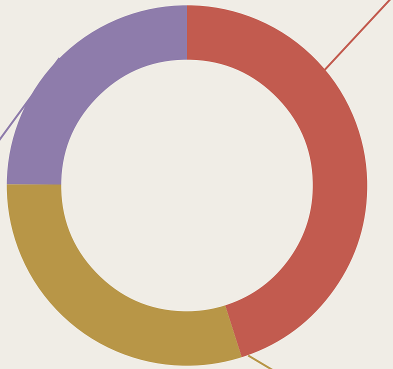

PRODUCT & ENGINEERING ~46% • £690K

Graph schema, MCP connectors and UI across the 3 launch practice areas. Lead Engineer (secured) plus additional engineering, working alongside fund finance, LevFin and M&A domain experts. Cloud, Neo4j and LLM inference. **MVP → v2.0.**

OPERATIONS & RUNWAY ~26% • £390K

The infosec stack institutional firms must see before they sign — SOC2 Type I → II, ISO 27001. Professional indemnity & cyber insurance, legal, finance, and contingency.

Salaries sit within the functions above — this line is pure overhead.

GO-TO-MARKET & FIRST CLIENTS ~28% • £420K

GTM Lead hire plus the founder-led partner motion. Design partners **converted to 12 paying anchor firms**. Conference circuit and category-creation thought leadership.

FOUNDERS@FORMATIVELEGAL.COM • BOOK A CALL

Appendices

A THE NEXT 18 MONTHS: PROJECTED MILESTONES

B COMPETITIVE LANDSCAPE: HARVEY VS LEGORA VS MATRIX

C INFORMATION SECURITY, COMPLIANCE & DATA RESIDENCY

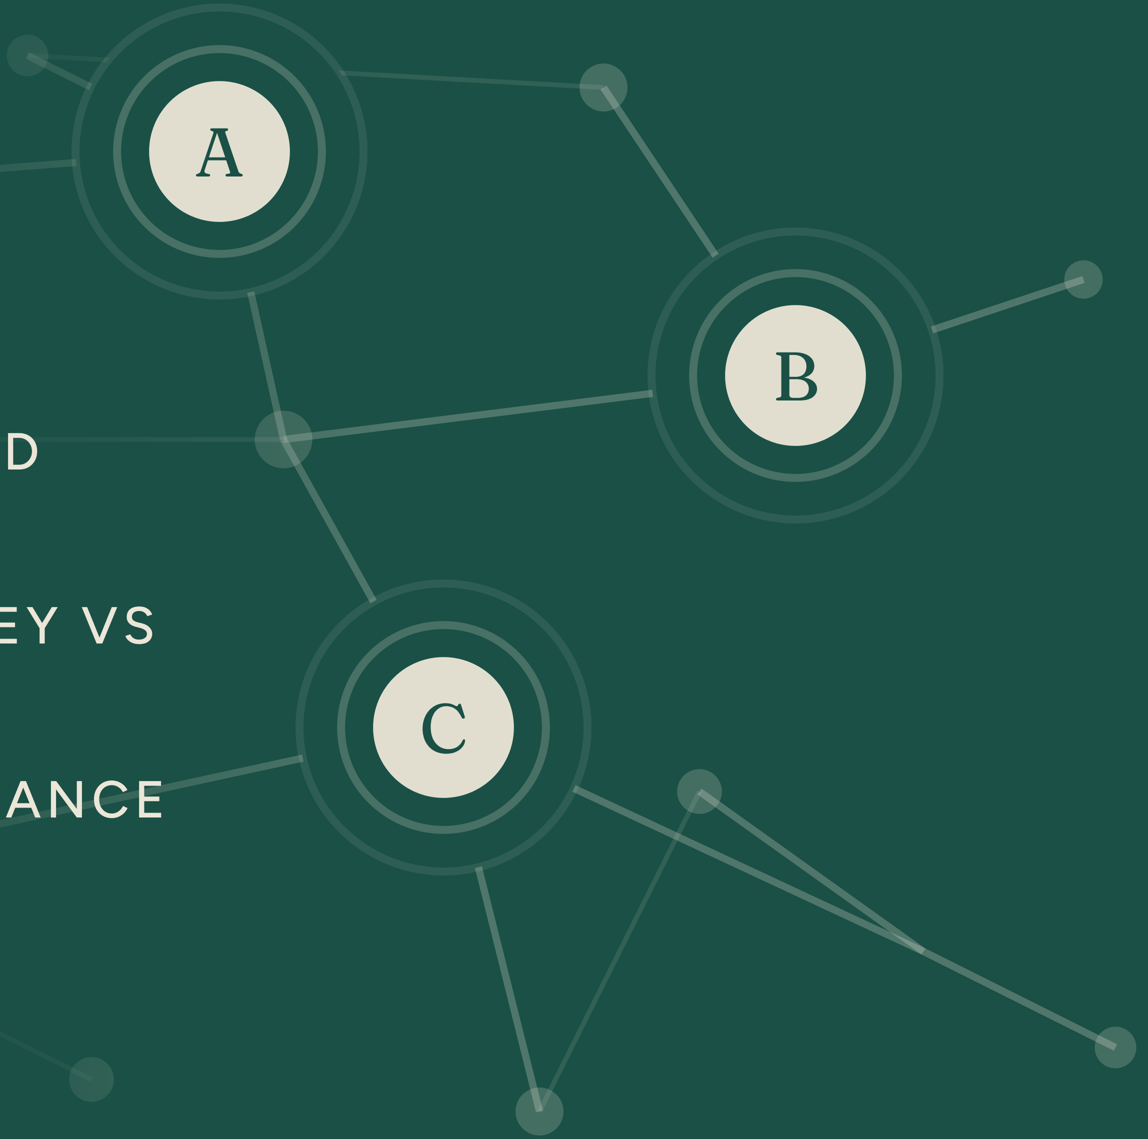

INTERACTIVE DECK →

INTERACTIVE DECK →

v1 (MCP server) proves the context layer as infrastructure. v2 (Deal Engine) is the UI layer on top: end-to-end OS built for all transactional workflows.

Built **with design partners and live pilots from month one** — not in isolation. Product, GTM, team and infosec run in parallel to first revenue, 12 anchor customers, £3.2m ARR and SOC2 Type II. **The pre-seed funds the full 18-month plan; we open the Series Seed at ~M12, from a de-risked position.**

*

INTERACTIVE DECK →

"How do you compete with Harvey/Legora?"

We don't — they're competing with each other. **We're the trusted matter-state layer their agents eventually need to call into**.

Harvey[1]

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.

Legora[2]

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

ARR	**$100m+**[2]
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, MCP-callable trust infrastructure for transactional practice areas

ARR	**£53.8m** Y3 · BASE CASE
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-matter, value-based pricing — tiered by deal size, ACV-capped per firm — sold into 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's $100m+ ARR is **self-reported** (Legora, "...propels Legora past $100M in ARR," 2 Apr 2026); independent reporting put ARR at ~$23m as of Sept 2025 — the basis of the ~240× multiple on the $5.55B Series D (TechFundingNews). Harvey figures are externally corroborated (CNBC).

INTERACTIVE DECK →

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.

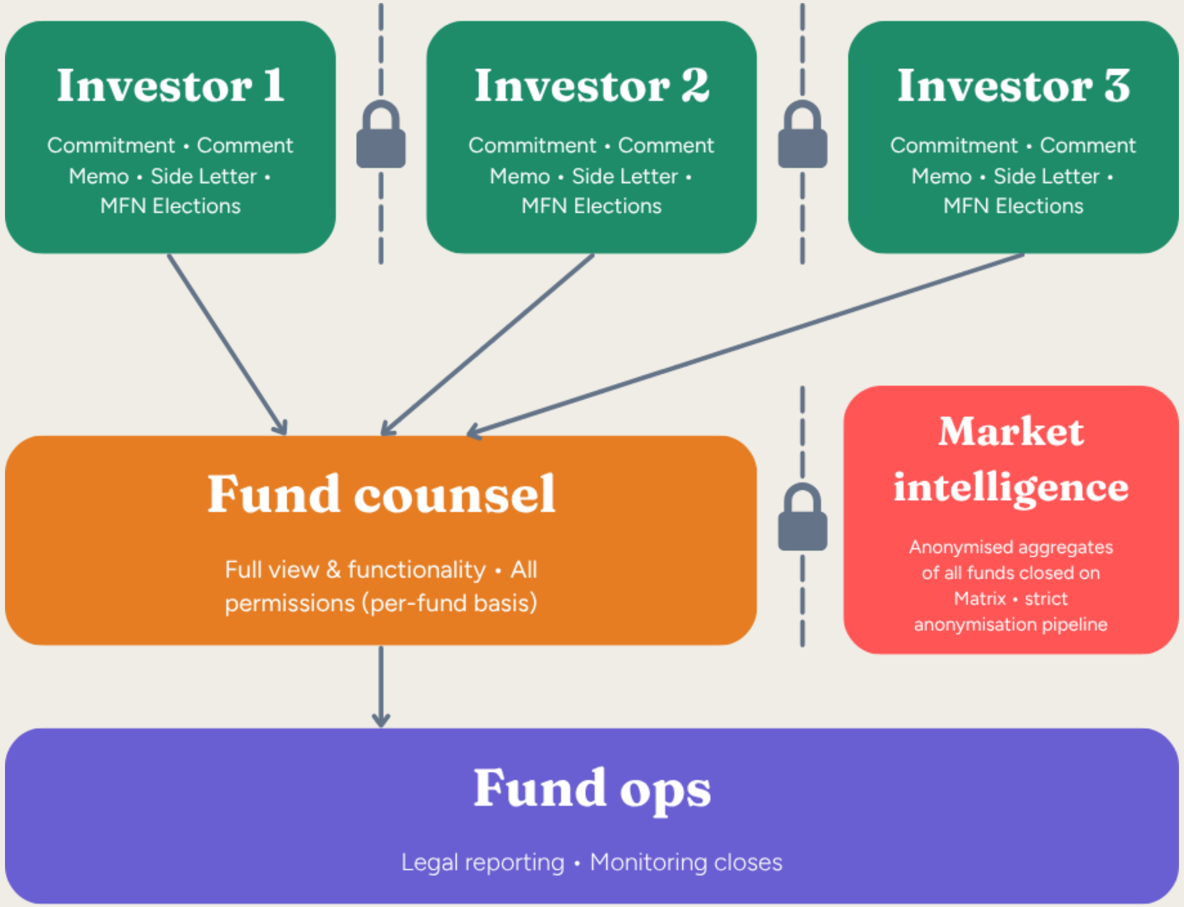

INVESTOR-LEVEL DATA ISOLATION

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.