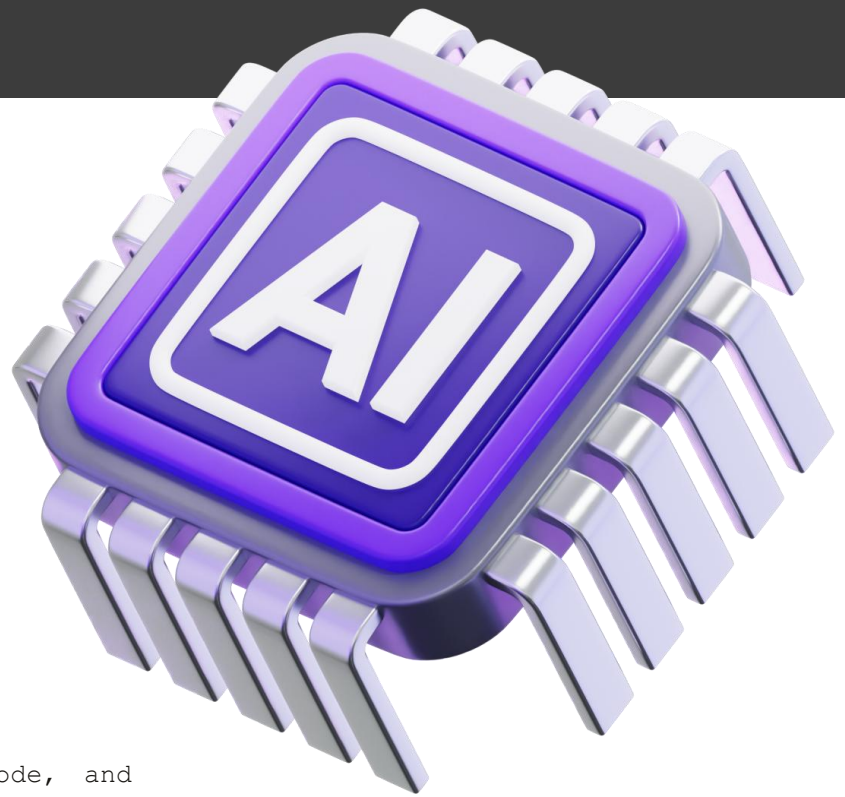
12 weeks of due diligence. Days. With an audit trail.



TWire reads the entire data room, computes every number in verified code, and returns findings a deal committee can sign. Sell-side memorandum, buy-side diligence and pitch analysis at launch, each in the client's own language.

ENGINEERED FOR CORRECTNESS. AUDITABLE BY CONSTRUCTION.

DOC CLASS · TECHNICAL BRIEF

SCOPE · PLATFORM CAPABILITIES & GUARANTEES

DISCLOSURE · PROPERTIES, NOT IMPLEMENTATION

REV · 4.0 · 2026

Twire Overview.

WHAT IS TWIRE?

AN AI-POWERED TRANSACTION INTELLIGENCE AND DUE-DILIGENCE PLATFORM FOR PROFESSIONAL DEAL TEAMS.

WHO USES IT?

M&A ADVISORS, ACCOUNTING/ADVISORY FIRMS, PE/VC FUNDS, FAMILY OFFICES, CORPORATE DEVELOPMENT TEAMS, INVESTMENT BANKS AND ACCELERATORS.

WHAT DOES IT DO?

IT ANALYZES COMPLETE TRANSACTION DATA ROOMS, FINANCIALS, CONTRACTS, MODELS AND COMMERCIAL DATA; IDENTIFIES RISKS AND INCONSISTENCIES; PERFORMS VERIFIED FINANCIAL CALCULATIONS; AND PRODUCES SOURCE-BACKED TRANSACTION DELIVERABLES.

WHY SHOULD I CARE?

FASTER DILIGENCE, LOWER ANALYST WORKLOAD, MORE CONSISTENT ANALYSIS, COMPLETE AUDIT TRAIL AND SUBSTANTIALLY FASTER SCREENING OF POTENTIAL TRANSACTIONS.

TWire Works With Your Existing Deal Team

EXISTING DEAL TEAM

- Document ingestion and normalization
- Cross-document analysis
- Reconciliation
- Initial risk identification
- Financial computation
- Source tracing
- Drafting
- Q&A preparation

YOUR PROFESSIONALS RETAIN

- Professional judgment
- Materiality decisions
- Client interaction
- Assumption approval
- Transaction strategy
- Final sign-off

What TWire Can Analyze

One Platform Across the Transaction

01 · FINANCIAL

Financial statements, management accounts, budgets, forecasts, working capital, debt, EBITDA adjustments, cash flow, valuation.

02 · COMMERCIAL

CRM data, customer concentration, churn, pricing, sales pipeline, market position, competitors.

03 · LEGAL

Contracts, amendments, change-of-control, termination, customer/supplier agreements, corporate documents.

04 · OPERATIONAL

KPIs, productivity, capacity, suppliers, inventory, organizational data.

05 · TECHNICAL

Technical reports, product documentation, IP, technology architecture.

06 · ESG / REGULATORY

ESG documentation, compliance records and regulatory information.

WHY WE WIN

The system is the moat.

Not the model.

18-24 mo

FOR A FUNDED TEAM TO APPROXIMATE THE TECHNICAL SURFACE

+18-24 mo

STILL BEHIND ON VERIFICATION DATA, RELATIONSHIPS, REGULATORY POSITION

36-48 mo

TOTAL CATCH-UP WINDOW. THAT IS THE MOAT.

Every deal processed enlarges the golden corpus the next deal is verified against: accuracy compounds where competitors start from zero. Institutional trust is earned on transactions, not pitches. First movers set the EU AI Act audit standards later entrants retrofit against.

Shadow diligence: prove it on a deal you closed.

We take a transaction your team completed (same data room, under NDA) and run it end-to-end. You compare our red flags against the ones you found, our valuation against the one you booked, our audit trail against your file. Two weeks. Your data never leaves the EU.

ARRANGE A PILOT

More Deals. Less Repetitive Work.

Same Professional Control.

Metric	Traditional process	TWiRE-supported process
Analyst workload	160 analyst hours	1 analyst hour
Senior review	40 senior-review hours	1 senior-review hour
Timeline	8 weeks	**1 day**
Cost	€ internal cost	**€ cost**

RESULT: 50–70% reduction in repetitive analytical workload

THE PROBLEM, QUANTIFIED

A deal cannot be diligence one document at a time.

A demo summarizes a ten-page NDA. A live transaction is a different object: audited statements under local GAAP and IFRS, management accounts that reconcile to neither, an SPA on its ninth draft, and a CRM export holding the real revenue story. The answers that decide price live in the joints between documents, not in any excerpt that fits in a prompt.

15K-50K
DOCUMENTS IN A TYPICAL MID- MARKET DATA ROOM

1.2M+
PAGES OF UNSTRUCTURED CONTENT PER DEAL

~2,250
PAGES THAT FIT IN A 3M-TOKEN CONTEXT WINDOW

500×
GAP BETWEEN WHAT FITS AND WHAT IS REQUIRED

Any system whose unit of reasoning is “what fits in one prompt” has made a scoping decision about your deal that nobody signed off on.

Seven Independent Verification Layers.

EVERY MATERIAL OUTPUT PASSES THROUGH:

01	02	03	04	05	06	07
DOCUMENT VALIDATION	CONFIDENCE VALIDATION	UNIT VALIDATION	DETERMINISTIC CALCULATION	INDEPENDENT CROSS-CHECK	BENCHMARK VALIDATION	HUMAN APPROVAL

CONTROL CHAIN	Document validation → confidence validation → unit validation → deterministic calculation → independent cross-check → benchmark validation → human approval.
DUE DILIGENCE	Detailed technical architecture available during technical due diligence.

THE SUBSTRATE

Real Data Rooms Are Messy.

TWire Is Built for Them.

Anyone who has run a real process knows these are not edge cases. They are the median input, and the reason "upload it to a chatbot" stops working the day the data room opens.

- AUDITED FINANCIALS AS SCANNED PDFS AT 200 DPI, ROTATED PAGES, NOTARY STAMPS AND WATERMARKS BLEEDING INTO THE TABLE CELLS
- LOCAL-LANGUAGE SHAREHOLDER RESOLUTIONS INTERLEAVED WITH ENGLISH BOARD PACKS, ONE ENTITY UNDER FIVE NAME VARIANTS
- SIGNED CONTRACTS WITH HAND-ANNOTATED AMENDMENTS THAT WERE NEVER TYPED UP, AND GOVERN ANYWAY
- THE BANKER'S MODEL WITH CIRCULAR REFERENCES, HIDDEN TABS, AND EXTERNAL LINKS TO A LAPTOP THAT LEFT THE FIRM IN 2023
- THREE VERSIONS OF THE SAME Q3 REPORT, NONE DATED, ONLY ONE OF WHICH THE SELLER INTENDS THE BUYER TO RELY ON
- MSG EMAIL EXPORTS WHOSE EMBEDDED ATTACHMENTS HOLD THE ONLY COPY OF THE SIDE LETTER THAT CHANGES THE EARN-OUT

A system that cannot ingest, normalize, and cross-reference this material is not a diligence platform. It is a demo. The first hundred engineering hours of TWwire were spent on parsing, not prompts.

FAILURE MODES OF ASSISTANT-GRADE AI

Why Generic AI Fails on Live Transactions.

F–01	**Retrieval-relevance mismatch**	Top-k retrieval surfaces fifty mentions of working capital and misses the seasonal-build clause that moves the peg by millions. Semantic similarity is not materiality.
F–02	**No cross-document closure**	A change-of-control clause in the top customer's MSA. 38% revenue concentration in the CRM export. A projection that assumes both away. Three documents, one deal-breaker, invisible to flat retrieval.
F–03	**Numeric transcription drift**	A model that reads EUR 12.4M and writes EUR 1.24M has not malfunctioned: that is what statistical text generation is. At an 8× multiple, the difference is nine figures of mispricing.
F–04	**Temporal incoherence**	The data room updates nightly through Q&A. A system without version lineage cites the seller's superseded figures with full confidence, inside your investment memo.
F–05	**Non-reconstructible reasoning**	When the committee asks "why this number?", "the model said so" ends the meeting. Under the EU AI Act, it also ends the audit.

Each failure mode is individually solvable. Solving all five simultaneously, in one system that also ingests scanned paper and emits committee-grade deliverables, is the engineering problem.

REQUIREMENTS · ASSISTANT VS INSTITUTION

Your analysts already use AI.

A committed deal sets a higher bar.

		GENERAL ASSISTANT	SELF-BUILT AGENT STACK	Twire
R1	Complete corpus coverage	One context window at a time	Whatever each script happens to fetch	Every document parsed, indexed, accounted for, incl. the ones nobody opened
R2	Deterministic computation	Numbers generated as text	Numbers generated as text	All arithmetic in versioned, tested code
R3	Cell-level provenance	None	Logs, if someone builds and maintains them	Document, page, cell, confidence, formula, timestamp, on every figure
R4	Temporal consistency	None	None	Nightly Q&A updates reconciled; superseded versions excluded
R5	Reproducible audit	A chat transcript	Fragile, unversioned	Any conclusion replayable from logged inputs and model versions
R6	Residency & isolation	Consumer terms, US cloud	Your team's ongoing liability	EU processing, single-tenant option, no training on client data
R7	Durable multi-week execution	The session forgets	State lives in scripts that break	Stateful six-week process; state survives restarts and deploys
R8	Named human accountability	None	None	Identity-logged review gates before anything reaches a client

Chat assistants are excellent at fifty-page tasks; your teams should keep using them there. The bar was set by deal practice and regulation, not by us.

THE PRODUCT

One command center. Deliverables a committee can sign.

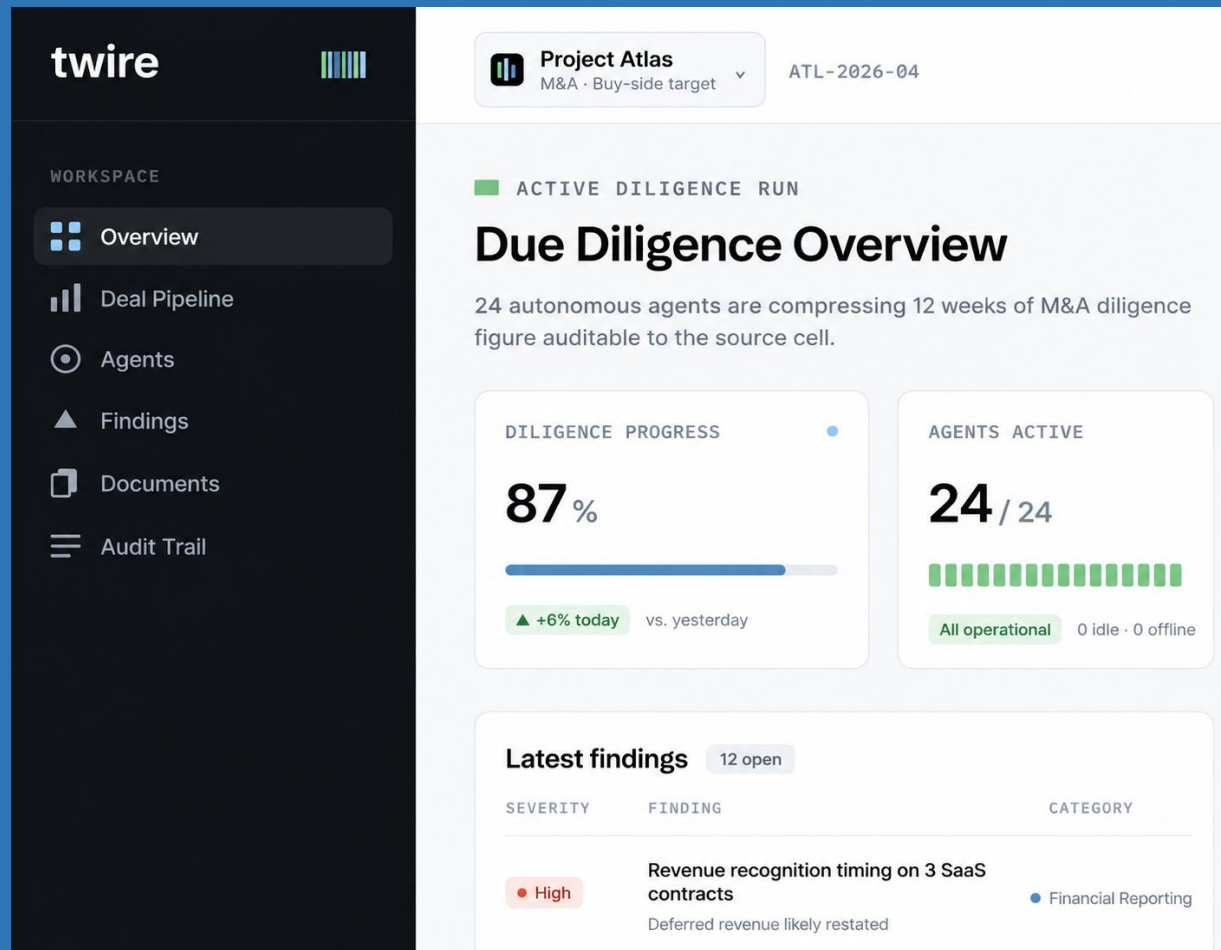
Live engagement state on one screen, and at the end, the artifacts your process already expects: red-flag report by workstream, QoE-style databook, valuation with sensitivities, CIM drafts and a citation-complete Q&A log. Each in the language the client asks for.

Live engagement state

Progress, workstreams, and open risk flags move in real time.

Findings with provenance

Each finding links to the exact source cell, e.g. FIN-REV!D42.



NUMERIC INTEGRITY

AI never does arithmetic.

Not one number, not once.

Financial calculations cannot live inside a language model; that is an audit requirement, not an opinion. The model's role is bounded and inspectable; the numbers are produced by the same class of code an auditor would review.

COMPUTED IN VERIFIED, VERSIONED CODE

- DCF with terminal value, WACC build-up, sensitivity grids
- Trading & transaction comparables with documented selection
- QoE-style EBITDA bridges and add-back schedules
- Net-working-capital analysis and peg mechanics
- DSCR & covenant-headroom modeling under stress scenarios
- Monte Carlo valuation distributions, assumptions stated

THE MODEL'S PERMITTED ROLE

- Select which computation engine a question requires
- Propose inputs, each traceable to a source document and page
- Narrate outputs in natural language, numbers passed through unmodified
- Flag deviations from benchmarks for human review

PRINCIPLE A number either traces to a specific page, cell, extraction confidence, and formula, or it does not ship. This is what “does not hallucinate” means when it is engineered rather than promised.

PROVENANCE · ONE NUMBER, END TO END

Anatomy of a shipped number.

Every figure carries this record.

The adjusted-EBITDA figure in a CIM is not a string the model produced. It is the visible end of a six-link chain, generated automatically and queryable forever.

01 · SOURCE

Q3_Mgmt_Accounts.xlsx

Sheet "P&L_2025", cell D42. Document checksummed and version-lineaged at intake.

doc a1f4…9c02

02 · EXTRACT

EUR 12.4M

Read under structural validation (cross-foot ok), calibrated confidence 0.973, above gate threshold.

ext 7b21…e4d8

03 · NORMALIZE

Canonical form

EUR thousands → EUR; fiscal period aligned to FY2025; entity resolved to TARGET_01.

nrm c583…1a77

04 · COMPUTE

EUR 14.1M adjusted

EBITDA_adj = EBITDA + Σ add- backs; 11 add-backs, each sourced; engine version pinned; sensitivity attached.

cmp f0ab…62c4

05 · NARRATE

CIM §4.2

Model drafts the prose around the figure. The number itself is passed through unmodified; narration cannot alter it.

out 3d9e…b510

06 · SIGN-OFF

Approved

Reviewer identity and timestamp bound to the record at Gate 3. Without this link, the deliverable does not ship.

rev analyst…t+0

AUDIT QUERY "Where does this number come from?" resolves in one lookup: source, confidence, transformations, formula, engine version, reviewer. For every number, on every page, of every deliverable.

VERIFICATION & EVALUATION

Correctness is not claimed.

It is measured.

Four mechanisms make accuracy a property that is tested, tracked, and reported: per component, per deal type, per release.

Contract-tested components

Every analysis component ships with a closed specification of admissible outputs and is regression-tested against it on every release. A component without a testable contract does not enter the pipeline.

Golden-deal benchmark corpus

A curated set of completed transactions with known ground truth: SaaS, industrial, carve- out, portfolio. Accuracy is measured against it continuously; regressions block release.

Independent recomputation

High-stakes outputs (valuation ranges, risk scores, legal flags) are derived through multiple independent paths and reconciled. Unresolved disagreement escalates to a human, never averages away.

Calibrated confidence gating

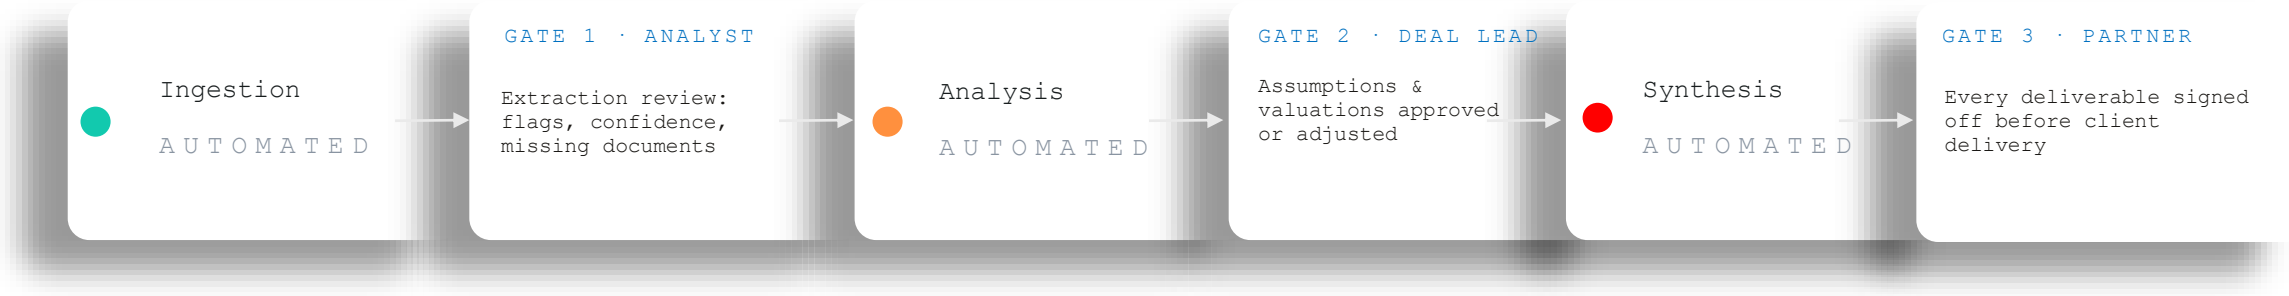
Every extracted value carries a calibrated confidence score. Below-threshold values are routed to a human verification queue before they can enter any downstream analysis.

Accuracy figures are meaningless without the methodology behind them. Ours ships with every engagement report.

HUMAN OVERSIGHT

Autonomous where it's safe.
Accountable where it matters.

The platform carries the volume; your deal team holds the pen. Processing pauses at three hard gates mapped to how deal teams already work, and nothing reaches a client without a named reviewer's sign-off.



Hard stops, not advisories

Processing does not continue past a gate. Flagged items are resolved or explicitly accepted; there is no third state.

Editable assumptions

Reviewers override any assumption (a peg, a discount rate, an add-back) and the engines recompute downstream results immediately, provenance intact.

Maker-checker by construction

Approvals are logged with the reviewer's identity: the four-eyes model regulated institutions already operate, applied to AI output.

Timed escalation

Unattended gates notify a team lead, then a partner. Engagements never silently stall.

This is the human-oversight posture Article 14 of the EU AI Act asks for, implemented as process design, not as a policy document.

OPERATING ENVELOPE

Institutional scale.
Deal speed.

Scale and latency are control features, not marketing numbers. A system fast enough to re-run is a system deal teams actually verify, and a process that answers in seconds changes how the process itself is run.

- **Red flags in hours, not week six**

  First findings earlier, questions answered in seconds, workstreams run concurrently, updates don't require restarting analysis.

- **Parallel workstreams**

  Financial, legal, commercial, technical, and ESG analyses proceed concurrently and reconcile continuously against the same corpus.

- **Q&A at deal tempo**

  A diligence question answered with citations in seconds, not in the next batch, three days from now.

- **Durable, resumable execution**

  Engagements span weeks. Any interrupted phase resumes exactly where it stopped, without recomputing completed work.

50K+

DOCUMENTS PER DEAL CORPUS, ALL FORMATS, ALL LANGUAGES

<200 ms

RETRIEVAL LATENCY AT FULL CORPUS SCALE

100%

OF FINDINGS CARRY A SOURCE TRACE

0

DEAL STATE LOST TO RESTARTS, DEPLOYS, OR FAILURES

SECURITY · RESIDENCY · REGULATION

Your Deal Data Stays Under Your Control.

Built for regulated clients from day one. Reconstructing any reasoning chain is a database query, not a forensic investigation.

EU residency, default and option

Processing in Frankfurt; no deal data transfers to a US cloud. Single-tenant and on- premise deployments available for institutions that require them.

Record-keeping

Every automated action is logged with input hash, model version, retrieval scope, and output hash; every human approval with the reviewer's identity. The record the EU AI Act's logging obligations anticipate.

Data protection

Encryption in transit and at rest, strict tenant isolation, client data never used for model training, verifiable deletion at engagement close.

Regulatory posture

Designed to support institutional compliance requirements, including human oversight, traceability and record keeping.

"The system extracted these inputs from these pages, computed the result in auditable code, and our analyst approved the output." A normal Tuesday at a bank.

PITCH LENS · ADD-ON SERVICE

Every pitch, scored Against your thesis.

For funds, accelerators, and corporate scouts reading thousands of decks a year. Any pitch is evaluated against your stated thesis (stage, sector, geography, check size, the signals you actually weight) and returns a decision-ready match report with independent market analysis. Same provenance discipline: every claim in the report carries its source.



01 INPUT
Ingest the deck

02 ANALYSIS
Score against your thesis

03 RESEARCH
Independent market research

04 OUTPUT
Match verdict + report

EVERY REPORT INCLUDES

- Thesis-fit rating with sub-scores
- Market & competitor analysis with sources
- Risk flags, each cited
- One comparable scale across all decks

A hundred pitches land on one comparable scale, and the ones worth a partner meeting surface themselves.

PRECEDENT

The market prices the system, not the model.

≈ $1.3B	Databricks acquires MosaicML · 2023	Not a category-heading model: the engineering that made large-scale training accessible. The system and the team were the asset.
≈ $14B	Meta invests in Scale AI · 2025	A fourteen-billion-dollar position bought repeatable, auditable data infrastructure, not a model.
≈ $185M	Snowflake acquires Neeva · 2023	A failed consumer product concealing enterprise-grade retrieval engineering, folded straight into Cortex AI.
≈ $700M	Hebbia Series B · 2024	No foundation model: financial-document reasoning with traceability. Investors priced the system.
Multi-$B	Harvey AI · 2024-25	Runs on the same rented models as everyone else. Owns the workflow, the law-firm embed, and the trust layer. Capital follows the workflow.

None of these valuations was justified by the model layer. Each was justified by what surrounds it: ingestion, verification, provenance, regulatory trail. That is the layer TWire is built on, deliberately, for M&A.

Arrange Your TWire Pilot



TW CONSULTING & TRADE LTD.

TW Consulting & Trade Ltd. (Malta)

Wolfgang Tweraser
CEO/President

Tel EU: +356 2701 9052 (office)
Tel EU: +356 9942 9829 (mobile)
Tel UAE: +971 50 209 3055

wolfgang@TWCTMT.com
www.TWCTMT.com

Addendums

- SYSTEM GUARANTEES
- THE PROCESSING FABRIC · EXTERNAL VIEW
- DEFENSE IN DEPTH



SYSTEM GUARANTEES

Five invariants.

Not goals.Preconditions of shipping.

I-01	Coverage	Every artifact is parsed, classified, and indexed; unprocessable items are surfaced for human handling, never silently dropped. Deals die in the dark corners: the system is not allowed to have any-
I-02	Provenance	Every emitted figure carries its source document, page and cell, extraction confidence, and full transformation chain. It is the difference between an opinion and a finding.
I-03	Determinism	All arithmetic executes in versioned, unit-tested code. Language models select inputs and narrate results; they never compute. Re-run the deal, get the same numbers.
I-04	Temporal consistency	Version lineage is tracked at content level. Current-state answers exclude superseded material; the state of knowledge at signing is reconstructible at closing.
I-05	Reconstructibility	Any conclusion can be replayed end-to-end from logged inputs, model versions, and retrieval scopes. An audit is a query, not an investigation.

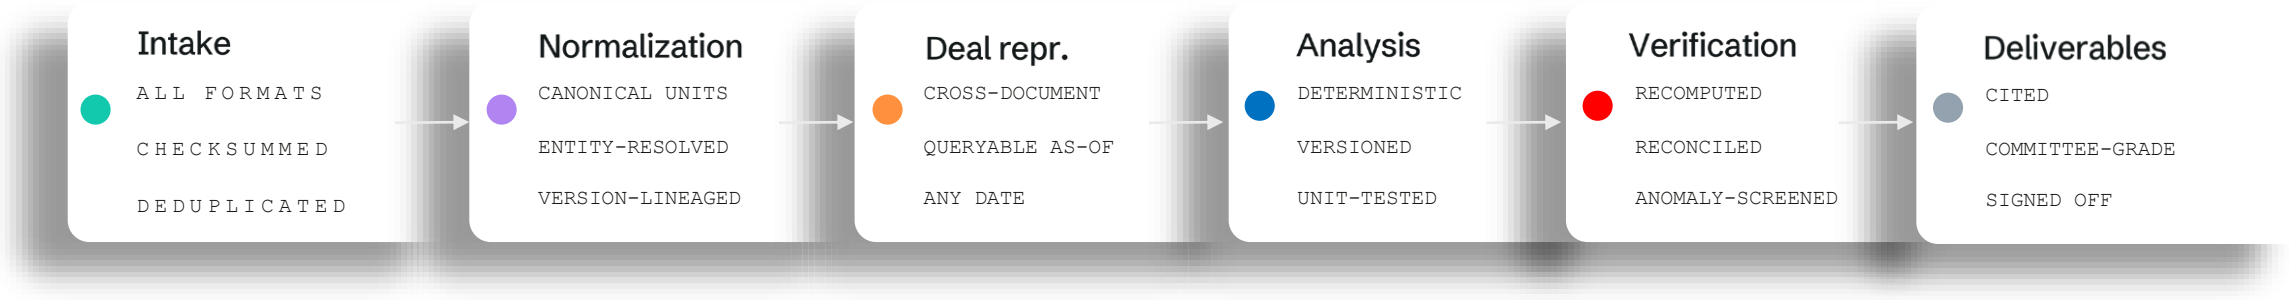
The invariants are enforced by the platform's test suite and release gates. A build that violates one does not deploy: correctness failures are build failures, not incident reports.

THE PROCESSING FABRIC · EXTERNAL VIEW

Six stages.Two control planes.

Visible only through their guarantees.

This is the platform's external contract: the stages a deal passes through and the planes that govern them. Each stage is independently contract-tested, idempotent, and resumable; the internal machinery is deliberately not part of this document.



PROVENANCE LEDGER Every artifact, transformation, and decision hash-chained end to end. Nothing enters a deliverable without a complete, queryable lineage record.

SPANS ALL SIX STAGES

HUMAN CONTROL PLANE Three identity-logged gates with hard stops and timed escalation. The pipeline cannot advance past a gate without a named reviewer's decision.

SPANS ALL SIX STAGES

This is the entire externally observable surface. Everything a client, auditor, or regulator interacts with reduces to these six stages and two planes, by design.

DEFENSE IN DEPTH

Seven layers between a raw document and a delivered number.

I-1	Structural parse validation	Extraction is checked against the document's own structure (totals, cross-foots, table geometry) before anything downstream sees it.	STRUCTURAL
I-2	Calibrated confidence gating	Every extracted value carries a calibrated confidence score; below-threshold values divert to human verification instead of entering the analysis.	STATISTICAL
I-3	Schema & unit conformance	Values must conform to the deal's typed schema (currency, unit, period, entity) or they are rejected. EUR thousands never silently meets EUR millions.	TYPE-LEVEL
I-4	Deterministic computation	All arithmetic in versioned, unit-tested engines. The model never computes; transcription drift cannot occur where there is no transcription.	ARITHMETIC
I-5	Independent recomputation	Valuation ranges, risk scores, and legal flags derived through independent paths and reconciled; disagreement escalates, never averages away.	REDUNDANT
I-6	Benchmark anomaly screening	Results are screened against sector benchmarks and the deal's own history. A 60% margin in a 30% industry is a question, not a finding.	CONTEXTUAL
I-7	Named human sign-off	A reviewer with identity and timestamp approves every deliverable. The last layer is a person, by design.	ACCOUNTABLE

The layers are uncorrelated by construction: structural, statistical, type-level, arithmetic, redundant, contextual, human. An erroneous number must defeat all seven independently to reach a client.