RED BEAR: ALPHA ENGINE · DATA SHEET · VALIDATE · INTELLIGENT DUE DILIGENCE

Buy only the answers you do not already have

Work out what the firm genuinely does not know, rank it by whether it would change the decision, brief advisers on that alone, and bring the findings back as evidence rather than a PDF.

KEY POINTS

- Everything known about the target is classified established, assumed, uncertain or absent, and every classification traces to the evidence that produced it
- The thesis is decomposed into its assumptions, and each is ranked by how much of the predicted uplift rests on it; that is what gets validated
- A structured work package, not a letter: questions, hypotheses, evidence required, outputs expected, confidence needed, scope boundaries; issued anonymised if the market must not learn what you are looking at
- Advisers bid against the same package, so cost and coverage are comparable for the first time; coverage is computed against the package, not self-reported
- Findings return as structured evidence, every prior assumption gets a verdict, and the prediction can be re-run; diligence never ends inside diligence

What it is

Before a firm commits, it hires external advisers to verify the investment case, and preliminary diligence is where it first spends real money on a deal. Intelligent Due Diligence audits the firm's own evidence base before the brief is written. It classifies everything known about the target, works out which of the uncertain things would actually change the decision if they turned out differently, and builds the work package from those alone, so advisers are pointed at what is genuinely unknown rather than re-establishing what the firm already holds.

The output is a structured work package rather than a letter, which several advisers bid against like for like. What comes back enters the evidence ledger with its source and status, is reconciled against what the firm believed going in, and re-runs the prediction from the previous stage.

The pain point

Diligence reads everything and checks nothing in particular. The brief is written generically because writing a precise one would mean auditing the firm's own evidence first, and nobody has time; so adviser spend goes on re-deriving what the firm already has, the standard checklist samples evenly across what matters and what does not, and the late surprise is the uncertainty nobody ranked. Proposals arrive in free form, so cost and coverage cannot be compared. And the report comes back as a PDF: the case that goes to committee is largely the one that went into diligence, because the findings never reached the model.

The costs are real. Diligence often runs to hundreds of thousands of dollars, and a deal that will not survive should be killed before the fee is committed, not after.

How it works

Know what you know. For a target, every material attribute is classified established, assumed, uncertain or absent, from evidence status, source count, contested flag and staleness across the ledger, model outputs, data-room findings and recorded management claims; the classification is recomputed on demand and traceable to its evidence (FR-VAL-001). The thesis is decomposed into its assumptions, each linked to the attributes and model outputs it depends on, and ranked by the sensitivity of the predicted uplift to it. An assumption carrying 40% of the uplift ranks first. The analyst can re-order the ranking, with a rationale (FR-VAL-002).

Scope it, shape it, issue it. A work package per workstream (commercial, financial, operational, technology, legal, ESG) specifies the questions requiring answers, the hypotheses to validate, the evidence and datasets required, the structured outputs expected, materiality, the confidence required, and explicit scope boundaries including what is out of scope; every question traces to a material gap or assumption, and packages export in a machine-readable form (FR-VAL-003). A package can be issued with the target's identity and any identifying attributes suppressed at generation, verified before issue, with a record of which adviser received which version (FR-VAL-004). Responses are laid side by side on cost, coverage of the specified questions, method, team and timeline; unanswered questions show as coverage gaps and cost is normalised per question (FR-VAL-005). The investment director chooses the adviser; the platform only makes the choice comparable.

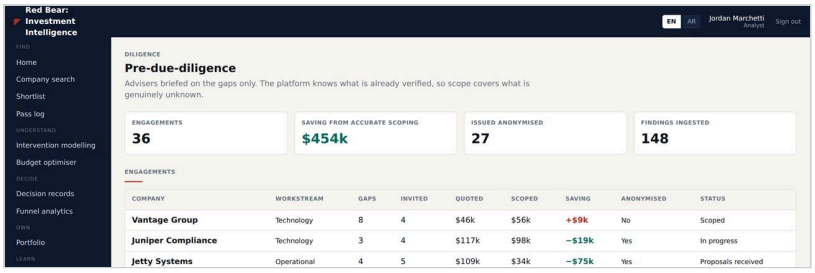
Bring it back as evidence. Advisers submit findings, assumptions tested, evidence, confidence, risks and recommendations in a defined schema alongside the conventional report, which is kept and linked; submitted values enter the ledger as *externally sourced*, attributed to the adviser, until verified (FR-VAL-006). Each prior assumption is adjudicated confirmed, disproven, conflicting or unresolved; conflicts are surfaced, never auto-resolved, and an adjudication can be overridden with a reason (FR-VAL-007). Reconciled evidence marks affected estimates stale and offers, never forces, a re-run; both cases are kept and the confidence change reported (FR-VAL-008). Quoted, scoped and final cost are

recorded per workstream, with the difference between a generic-scope estimate and the evidence-driven scope reported as an estimate and labelled so (FR-VAL-009). Every question, engagement, submission, adjudication and override is recorded immutably and exportable for an LP or regulator (FR-VAL-010).

What you get

Capability	What it does	Requirements
Knowledge state and thesis dependency	Every attribute classified; assumptions ranked by their weight in the predicted uplift	FR-VAL-001, FR-VAL-002
Work package generation	Structured, per-workstream, exportable; every question traces to a material gap	FR-VAL-003
Anonymised issue	Target identity suppressed at generation, verified before issue	FR-VAL-004
Proposal comparison	Cost, computed coverage, method, team and timeline side by side	FR-VAL-005
Structured evidence submission and reconciliation	Findings in a schema, attributed to the adviser; every prior assumption gets a verdict	FR-VAL-006, FR-VAL-007
Feedback to prediction	Stale estimates re-run on offer, both cases kept, confidence delta reported	FR-VAL-008
Cost tracking and audit record	Quoted, scoped and final cost per workstream; an immutable, exportable diligence record	FR-VAL-009, FR-VAL-010

Screens



S-13 Pre-due-diligence: Pre-due-diligence: the knowledge state of the target, the assumptions ranked by materiality, and the work package built from the gaps. Prototype render on synthetic data for fictional entities.

RETURN ON INVESTMENT

Lower diligence spend. Scope covers what is genuinely unknown rather than everything that could be asked; the product records the difference between the generic-scope estimate and the evidence-driven scope per engagement, labelled an estimate, and that is the measure this pillar is judged on (FR-VAL-009). **Genuine competitive tension.** Advisers price against a common structured brief, so cost and coverage are comparable in a way free-form proposals never are. **Fewer late surprises.** Scoping follows materiality rather than a checklist that samples evenly. **A case that improves.** Structured findings enter the model and re-run the prediction, so the paper the committee sees is better evidenced than the one that went into diligence. Ask the data room directly during the engagement: a plain-language question returns the supporting passage or "not found" (FR-FIND-006).

Platform engines

Intelligent Due Diligence consumes the Research engine for cited, refusable answers from the data room (RBP FR-RS-02, RBP FR-RS-04), the Organisation Information Gathering engine's rule that AI extraction writes proposals a person confirms (RBP FR-OG-02), and the Provenance service's staleness and confidence model (RBP FR-PV-01, RBP FR-PV-03). Platform behaviour is specified in the platform functional specification, not here.