

# AI-Native Startups: How Investors Underwrite the AI Application Layer

*An evidence-led framework for global VC and growth investors and early-stage founders*

**Prepared by Matchpoint Partners**

Author attribution pending CK approval

August 2026

## ABSTRACT

**Background.** AI venture investment expanded sharply and became more concentrated in 2025 and the first half of 2026. **Objective.** This paper develops an evidence-led framework for underwriting AI-native application companies across product, technology, data, evaluation, revenue, margin, distribution, governance and capital. **Approach.** The analysis reviews 40 official, regulatory, provider, market, company-filing, standards and security sources available through 1 August 2026. **Findings.** The AI-native label has decision value only when machine inference materially changes a customer workflow and the company can evidence representative quality, retained production use, provider resilience, data rights and complete per-outcome economics. **Implications.** Investors and founders should use the quality-adjusted accepted customer outcome as the operating unit. Attributed Matchpoint or client revenue, cost reduction and loss reduction remain USD 0 until approved observed attribution exists.

## Highlights

**Unit.** A quality-adjusted accepted customer outcome connects model work to customer value and authority.

**Evidence.** Material AI, customer, productivity and financial claims receive a source, class, scope and owner.

**Economics.** Gross margin includes inference, retrieval, tools, failures, human review, support and reliability.

**Supply chain.** Model and cloud pricing, terms, lifecycle, concentration and migration are underwritten explicitly.

**Downside.** The investment case is tested against quality, provider, customer, sales-cycle and cost shocks.

**JEL Classification:** G24, G31, G32, L26, L86, M13, O32, O33

**Keywords:** AI-native startups, AI application layer, venture capital, growth equity, startup underwriting, model economics, inference cost, product evaluation, retention, data rights, AI governance, GCC venture ecosystem

*Research paper only. This document is not investment, legal, regulatory, accounting, tax, audit, privacy, intellectual-property, cybersecurity, valuation or technology advice. All worked inputs, score weights, shocks and thresholds are unverified illustrative management assumptions.*

## 1. INTRODUCTION

Artificial intelligence has become a large and concentrated category of venture investment. Stanford HAI reports that global corporate AI investment more than doubled in 2025, with generative AI capturing nearly half of private AI funding [1]. NVCA reports that US venture investment reached record levels in the first half of 2026 while most invested capital flowed to AI companies and financings of USD 100 million or more [2]. Carta records USD 119.5 billion raised by companies on its platform in 2025 across 4,859 rounds, the lowest annual round count in at least six years; its Series A data shows a 38% median valuation premium for AI companies [3]. OECD reports that AI-related investment exceeded half of overall venture investment in 2025 [6].


These figures describe capital markets, not investment quality. A label such as AI-native can refer to a product whose customer outcome depends on machine inference, a conventional software product that added a model interface, a services business using AI internally, a workflow company with a proprietary evaluation and data loop, or a founder narrative with limited production evidence. Each case has a different product, risk, margin, distribution and capital profile.

The underwriting problem is therefore evidential. An investor needs to establish what the application does, which customer outcome changes, what portion of the workflow is automated, which model and data dependencies matter, how quality is evaluated, how failures are controlled, how customers buy and retain the product, how revenue converts to cash, how cost to serve changes with usage and which rights allow the company to operate. A founder needs the same evidence to manage the company and prepare a credible financing process.

This paper develops a framework for **AI-Native Startups: How Investors Underwrite the AI Application Layer**. The primary ideal customer profile is A5, Global VC, Growth & Fintech Investors: international venture, growth and fintech-focused funds assessing UAE and GCC opportunities. The secondary profile is B5, Early-Stage & Growth Founders: founders raising from pre-seed through Series A and growth, particularly in AI and fintech across the UAE, GCC, UK and South Asia. These profiles, the title, hook, visual requirements, target length and Tech & AI service mapping were verified from Matchpoint Partners' Topic Tracker and ICP Legend on 1 August 2026.

The paper's central proposition is that the quality-adjusted accepted customer outcome is the correct operating unit. An accepted outcome is a completed customer task or decision that passes defined quality, authority and evidence criteria. Examples include a reviewed credit memo, a resolved customer request, an accepted invoice classification, a released software change, a completed compliance check or a customer-approved design. Tokens, calls, generations, agents and model benchmark scores are technical inputs. They become economically relevant when they contribute to accepted outcomes, retained customers and collected gross profit.

Seven underwriting principles follow. First, the AI-native label must be translated into an observable workflow and architecture. Second, the company should evidence model performance on representative customer tasks. Third, revenue quality requires contract, invoice, collection, cohort, concession and concentration evidence. Fourth, gross margin requires a complete per-outcome cost ledger, including human review and reliability costs. Fifth, model and cloud providers are part of the product supply chain. Sixth, data rights, privacy, intellectual property, security and regulatory role determine where and how the product can be sold. Seventh, the investment case should survive a model-price shock, a quality regression, a provider retirement, a customer loss and a slower sales cycle.



**Figure 1. Before and after: claim-led diligence versus evidence-led underwriting**

Source: Matchpoint underwriting framework.

The evidence cut-off is 1 August 2026. The analysis uses current official market reports, securities filings, provider documentation, regulator material, legislation, standards and security resources. No T17 company, founder, fund, customer, product, contract, cohort, data room, evaluation suite, cost ledger, cap table or investment memorandum was supplied. All worked volumes, prices, costs, conversion rates, productivity effects and valuation inputs are unverified illustrative management assumptions. Attributed Matchpoint or client revenue, cost reduction and loss reduction remain USD 0.

## 2. SCOPE, DEFINITIONS AND EVIDENCE BOUNDARIES

### 2.1 Application layer

The application layer converts models, data, software and human work into a product used for a specific customer workflow. It can include user interfaces, application logic, retrieval, memory, tools, model routing, prompts, fine-tuning, evaluation, access control, observability, billing and support. The application company may use proprietary, third-party or open-weight models. Its defensibility can arise from workflow integration, proprietary data rights, evaluation assets, distribution, switching costs, regulated permissions, customer trust or operational learning.

The model layer supplies inference or model weights. The infrastructure layer supplies compute, storage, orchestration and deployment. These boundaries can overlap. A company that fine-tunes or hosts a model still needs to show how the model changes the customer outcome and how the application captures value.

| Layer             | Principal components                               | Underwriting question                              | Minimum evidence                                          |
|-------------------|----------------------------------------------------|----------------------------------------------------|-----------------------------------------------------------|
| Customer workflow | User, task, decision, exception and approval       | Which costly or valuable job changes?              | Process map, baseline time, quality and authority         |
| Application       | Interface, logic, retrieval, tools and integration | What is the product customers buy?                 | Product demonstration, release record and customer use    |
| Evaluation        | Test cases, graders, thresholds and monitoring     | How is quality established and preserved?          | Representative suite, results, failures and release gates |
| Data              | Source, rights, lineage, quality and retention     | May the company use the data for this purpose?     | Rights register, samples, lineage and controls            |
| Model             | Provider, version, route, prompt and tuning        | Which capability is required and replaceable?      | Provider register, tests, costs and migration evidence    |
| Infrastructure    | Compute, storage, security and observability       | Can the service operate reliably and economically? | Architecture, invoices, logs, incidents and capacity plan |

| Layer      | Principal components                          | Underwriting question                        | Minimum evidence                     |
|------------|-----------------------------------------------|----------------------------------------------|--------------------------------------|
| Commercial | Contract, price, term, renewal and collection | Does use become retained, collected revenue? | Contract-to-cash and cohort evidence |

## 2.2 AI-native classification

An AI-native classification should be based on evidence. The classification does not require one technical architecture. It requires material dependence on machine inference within the customer value chain, a product designed around that capability and an operating system that measures its quality and economics.

| Test             | Evidence of a substantive AI-native application                 | Weak evidence                                 |
|------------------|-----------------------------------------------------------------|-----------------------------------------------|
| Customer outcome | Defined task with baseline, acceptance and consequence          | Generic productivity statement                |
| Product design   | Workflow built around inference, review and exception handling  | Chat interface added to existing software     |
| Evaluation       | Representative task suite and production monitoring             | One vendor benchmark or founder demonstration |
| Data             | Governed source, rights, lineage and feedback loop              | Unspecified proprietary data claim            |
| Economics        | Per-outcome price, cost and gross profit                        | Token cost quoted without workflow cost       |
| Operations       | Release, incident, rollback and provider migration process      | Best-effort manual troubleshooting            |
| Commercial       | Retained customer use tied to contract and collection           | Pilot logos or letters of intent only         |
| Defensibility    | Evidence loop, workflow integration, distribution or permission | Model access available to every competitor    |

The result can be recorded as AI-native, AI-enabled, AI-assisted service or unverified. The label is a diligence conclusion, not a valuation conclusion.

## 2.3 Evidence classes

Every material claim should receive a class and owner. Observed evidence comes from a controlled system of record. Independently reproduced evidence comes from a test rerun by the investor, adviser or a qualified third party. Customer-attested evidence comes from a named authorised customer representative and should be reconciled to system data where possible. Management-estimated evidence is labelled and supported by a method. Unsupported claims remain open issues.

| Class                   | Definition                                                 | Permitted investment-memo use                          |
|-------------------------|------------------------------------------------------------|--------------------------------------------------------|
| O1 observed             | Dated source-system evidence with owner and reconciliation | Stated as observed within the defined period and scope |
| O2 reproduced           | Independently rerun method with retained inputs and result | Stated as reproduced within the tested environment     |
| O3 attested             | Named customer or counterparty confirmation                | Stated as attested; limitations disclosed              |
| M1 management estimated | Documented calculation using unverified inputs             | Labelled management estimate or illustrative scenario  |
| U unsupported           | No adequate evidence supplied                              | Excluded from base case; recorded as an open issue     |

## 2.4 Public comparables and market reports

Public filings demonstrate analytical methods and possible economic variation. Palantir reported USD 4.475 billion of 2025 revenue, an 82% gross margin and 954 customers [7]. UiPath reported USD 1.611 billion of fiscal 2026 revenue, USD 1.853 billion of annualised renewal run-rate and an 83% gross margin [8]. C3 AI reported USD 250.3 million of fiscal 2026 revenue and a 31% GAAP gross margin [9]. Duolingo disclosed that increased generative-AI and hosting costs reduced gross margin during 2025 and later described continued investment in AI infrastructure [10].

These companies differ in scale, product, deployment, customer mix, services, accounting and maturity. Their metrics do not determine a private startup's gross margin or valuation. They show why an investor should avoid a generic AI-software margin assumption.

Market reports also have boundaries. Carta's published data is aggregated from companies and securities holders on its platform and excludes companies that requested exclusion from anonymised studies [3-5]. NVCA's report covers the US venture ecosystem with PitchBook as its data provider [2]. OECD uses Preqin data in its AI venture brief [6]. Stanford HAI combines multiple public and proprietary sources [1]. The paper retains each source's scope.

## 2.5 Legal, investment and author boundaries

This paper does not determine whether a product is regulated, whether processing is lawful, whether intellectual-property rights are valid, whether a security is suitable, whether a valuation is fair or whether an investment should be made. Those decisions require the applicable contract, facts, jurisdiction and qualified authority.

Named-person author attribution remains pending CK approval. Matchpoint Partners is the organisational author. No founder performance, investor view or client outcome is attributed to a named person.

## 3. A5 INVESTOR AND B5 FOUNDER DECISION MAP

### 3.1 A5 investor decisions

A5 investors need a staged answer to five questions. Is the customer problem valuable and repeated? Does the product create an accepted outcome that customers will pay for? Can the company maintain quality and control through model and market change? Does retained revenue create attractive gross profit and cash characteristics? Is the financing price supported by the evidence and downside?

| A5 decision                    | Core evidence                                                    | Failure signal                                        | Authority                              |
|--------------------------------|------------------------------------------------------------------|-------------------------------------------------------|----------------------------------------|
| Screening                      | Customer problem, workflow, product and initial traction         | AI label without customer outcome                     | Deal team                              |
| Technical diligence            | Architecture, evaluations, data rights, security and portability | Demonstration cannot be reproduced                    | Technical reviewer and deal lead       |
| Commercial diligence           | Contract-to-cash, cohorts, references and pipeline               | Pilots counted as recurring production                | Deal team and commercial specialist    |
| Financial diligence            | Revenue bridge, gross margin, burn, runway and capital plan      | Costs omit inference, review or support               | Finance reviewer                       |
| Legal and regulatory diligence | Corporate, IP, data, licensing and product-role matrix           | Rights or permissions cannot be evidenced             | Qualified counsel and compliance owner |
| Valuation and structure        | Base, downside, dilution, milestones and reserves                | Premium driven only by AI category                    | Investment committee                   |
| Post-investment plan           | Evidence gates, hiring, product and financing milestones         | Plan measures output volume without accepted outcomes | Board and management                   |

### 3.2 B5 founder decisions

The same evidence supports founder control. A founder should know which customer outcome drives retention, which failure causes the most customer harm, which provider dependence constrains scale, which cohort creates gross profit and which milestone deserves the next unit of capital.

| B5 decision      | Required record                                                | Management use                                   |
|------------------|----------------------------------------------------------------|--------------------------------------------------|
| Product scope    | Accepted-outcome definition and exception boundary             | Prioritise the workflow customers value          |
| Model choice     | Representative eval, cost and latency by provider/version      | Route tasks and plan migrations                  |
| Data strategy    | Rights, quality, lineage and feedback register                 | Build lawful and useful learning loops           |
| Pricing          | Outcome value, usage, service intensity and willingness to pay | Select subscription, usage or outcome pricing    |
| Sales            | Funnel by segment, source, stage and time                      | Allocate founder and sales capacity              |
| Customer success | Cohort use, acceptance, support and renewal                    | Identify adoption and churn causes               |
| Hiring           | Bottleneck, capability and milestone                           | Link headcount to product or commercial evidence |
| Fundraising      | Use of funds, gates, runway and downside                       | Raise against evidence creation                  |

### 3.3 Shared evidence room

The investor and founder should work from the same reconciled evidence room. Different interpretation is legitimate. Different underlying numbers create avoidable risk.

| Evidence block             | Primary owner             | Reconciliation                                                |
|----------------------------|---------------------------|---------------------------------------------------------------|
| Customers and contracts    | Commercial lead           | CRM to contract repository to invoice and collection          |
| Product usage and outcomes | Product lead              | Event log to accepted-outcome definition                      |
| Evaluation and incidents   | Engineering or AI lead    | Release to evaluation report to production monitoring         |
| Cost and gross margin      | Finance lead              | Vendor invoice and payroll allocation to telemetry and ledger |
| Data and IP                | Legal or data owner       | Source, agreement, use, retention and output policy           |
| Security and privacy       | Security or privacy owner | Architecture, controls, tests, incidents and remediation      |
| Cap table and financing    | CFO or company secretary  | Legal register, option ledger and signed instruments          |
| Claims                     | CEO and functional owners | Claim, source, class, scope, date and approval                |

## 4. MARKET CONTEXT AND THE AI-CAPITAL CONCENTRATION PROBLEM

### 4.1 Premium and selection

AI capital growth expands the opportunity set and increases selection pressure. Carta's 2025 data shows more capital across fewer rounds and an AI valuation premium from Series A onward [3]. Its seed research reports leaner early teams, longer fundraising timelines and a widening gap between typical and top-tier valuations, particularly for AI companies [4]. Its H2 2025 compensation report describes AI-driven capital and talent effects while emphasising smaller teams at early stages [5].

The investor should therefore distinguish a category premium from company evidence. A valuation can reflect scarcity, competition, strategic option value, market size, growth, retention, team, technology, data rights or momentum. Each driver needs its own evidence and downside treatment.

### 4.2 Market-sizing discipline

An application-layer market should begin with customer units and workflow spend. Top-down AI totals can contextualise adoption. They do not establish obtainable revenue.

| Market layer         | Calculation                                                         | Evidence gate                                         |
|----------------------|---------------------------------------------------------------------|-------------------------------------------------------|
| Problem population   | Number of target organisations or users with the defined workflow   | Named segmentation source and exclusions              |
| Serviceable workflow | Population with required data, budget, jurisdiction and integration | Customer discovery and channel evidence               |
| Paying unit          | Contracted account, seat, workflow, transaction or accepted outcome | Actual pricing and buying process                     |
| Attainable market    | Qualified accounts reachable through available distribution         | Funnel, capacity, conversion and sales-cycle evidence |
| Base plan            | New and expanded collected revenue by cohort                        | Contract-to-cash model and downside                   |

For a vertical AI company, the market model should include the number of target firms, eligible departments, workflow frequency, current cost, required permission, integration burden, budget owner and plausible adoption. For a horizontal product, segmentation by use case and buyer prevents a superficial total-seat calculation.

### 4.3 Fundraising probability and runway

The market evidence supports a concentrated financing environment. It does not establish that a given company can raise. The financing plan should include a no-raise case, a delayed-raise case and a lower-price case. It should show

monthly cash, committed cost, discretionary cost, hiring triggers, provider commitments, working capital and minimum operating runway.

| Financing case | Unverified illustrative assumption   | Required management response                                |
|----------------|--------------------------------------|-------------------------------------------------------------|
| Base           | Next round closes in month 12        | Hire only against product and commercial gates              |
| Delay          | Next round closes in month 18        | Preserve cash and defer non-critical commitments            |
| Lower price    | Pre-money valuation 30% below base   | Model dilution, structure and option-pool effect            |
| No raise       | No external capital within 24 months | Define break-even, strategic sale or orderly wind-down path |

These cases are unverified illustrative management assumptions. They demonstrate planning mechanics. They do not predict financing availability, valuation or return.

## 5. THE AI-NATIVE APPLICATION-LAYER TEST

### 5.1 Customer outcome before model capability


The first diligence session should reconstruct one customer workflow without the product. The team records trigger, input, operator, decision, output, review, exception, elapsed time, error cost and authority. It then reconstructs the same workflow with the product. The comparison identifies which steps disappear, change, accelerate or require new controls.

The exercise exposes category errors. A faster draft can increase review effort. A lower handling time can increase correction or escalation. A higher automation rate can hide lower customer acceptance. A product that creates more output can reduce value if users cannot verify it within their decision window.

| Workflow field | Baseline record                     | Assisted record                         | Acceptance evidence           |
|----------------|-------------------------------------|-----------------------------------------|-------------------------------|
| Trigger        | Event that starts the task          | Same or changed event                   | Source-system timestamp       |
| Input          | Documents, data and context         | Data passed to application and model    | Lineage and rights            |
| Work           | Human and software steps            | Automated, assisted and retained steps  | Process log and observation   |
| Output         | Decision, content or action         | Generated candidate and final output    | Version and reviewer record   |
| Quality        | Error, completeness and policy test | Same tests plus model-specific failures | Representative evaluation     |
| Authority      | Person allowed to release action    | Human or system release boundary        | Role and approval record      |
| Consequence    | Revenue, cost, risk or experience   | Observed change                         | Customer and finance evidence |

### 5.2 Accepted customer outcome

An accepted outcome has five elements. The task is defined. The source evidence is sufficient. The output passes quality and policy thresholds. The authorised person or system accepts it. The outcome reaches the customer workflow. A generated answer that is discarded, a draft that requires full recreation and an automated decision reversed by review are produced outputs, not accepted outcomes.



**Figure 2. Accepted customer outcome evidence chain**

Source: Matchpoint operating framework.

The company should measure accepted outcomes per customer and cohort. Useful measures include first-pass acceptance, accepted outcome per active user, review minutes per accepted outcome, exception rate, correction rate, incident rate, customer escalation and time to acceptance. The definition should be stable enough for cohort analysis and specific enough to prevent volume inflation.

### 5.3 Claim register

SEC enforcement concerning AI statements shows the importance of accurate representations [12]. A private financing also depends on reliable claims. The claim register records each material assertion used in the pitch, data room, customer proposal and investment memo.

| Claim                     | Class                 | Evidence                                | Scope and date                | Owner            | Memo treatment                         |
|---------------------------|-----------------------|-----------------------------------------|-------------------------------|------------------|----------------------------------------|
| "Cuts review time by 60%" | M1 until observed     | Time-study method and sample            | Named workflow and period     | Product owner    | Management estimate or observed result |
| "Proprietary dataset"     | U until rights proven | Dataset inventory and agreements        | Data fields and jurisdictions | Legal/data owner | Open issue                             |
| "Enterprise production"   | O1                    | Contract, environment and usage         | Customer, product and period  | Commercial lead  | Observed with defined scope            |
| "Best-in-class accuracy"  | U                     | Representative benchmark and comparator | Task, sample and version      | AI lead          | Exclude until reproduced               |
| "90% gross margin"        | O1 or M1              | Ledger, invoices and allocation         | Product and period            | CFO              | Reconciled actual or management plan   |

The claim register also helps founders. Unsupported claims can be corrected before a financing process. Management estimates can be converted into observed evidence through structured tests.

### 5.4 Evidence flywheel

A defensible application can build a lawful and useful evidence flywheel. Customer tasks create consented or contractually permitted usage data. Review and outcome signals improve evaluation. Evaluations improve routing, workflow and release decisions. Better accepted outcomes improve adoption and retention. Retained use produces more representative evidence. Each link needs rights, quality and economic value.

The investor should reject a circular data-moat claim. More data has value only when it is permitted, relevant, sufficiently labelled, incorporated into a controlled improvement process and difficult for a competitor to reproduce.



## 6. PRODUCT VALUE, CUSTOMER WORKFLOW AND PRODUCT-MARKET EVIDENCE

### 6.1 Baseline and counterfactual

Product value needs a baseline. The baseline can be current human work, rules, incumbent software, outsourced service or no action. The counterfactual should reflect what the customer would actually do without the product. Comparing a product with an unrealistic manual process overstates value.

The strongest early evidence is a representative within-customer comparison. It uses the same task family, service requirements and decision window. It records input differences, selection effects and review. Randomised tests can be appropriate for product features. Operational and regulated workflows may require matched periods, staggered rollout or controlled shadow use.

### 6.2 Productivity and revenue bridge

The Topic Tracker hook asks how technology multiplies productivity and revenue. The paper separates gross activity from verified enterprise benefit.

| Bridge line          | Calculation                                                             | Evidence requirement                         |
|----------------------|-------------------------------------------------------------------------|----------------------------------------------|
| Baseline capacity    | Accepted outcomes divided by full labour hours                          | Time study and stable outcome definition     |
| Assisted capacity    | Accepted outcomes divided by assisted labour and exception hours        | Same task mix and quality                    |
| Quality adjustment   | Accepted outcomes passing all thresholds                                | Evaluation and customer acceptance           |
| Customer value       | Avoided cost, released capacity, lower loss or higher collected revenue | Approved counterfactual and finance evidence |
| Product price        | Collected revenue net of credits and concessions                        | Contract, invoice and bank/ledger record     |
| Product gross profit | Collected or recognised revenue less complete cost to serve             | Reconciled ledger and allocation policy      |

Productivity becomes a customer value proposition when the customer can redeploy capacity, reduce actual cost, serve more demand, improve time-to-revenue or reduce measured loss. Revenue becomes company value when contracted use converts to collected and retained gross profit.

### 6.3 Customer evidence ladder

| Level               | Evidence                                           | What it supports                      |
|---------------------|----------------------------------------------------|---------------------------------------|
| C0 discovery        | Interview and observed pain                        | Problem hypothesis                    |
| C1 design partner   | Access, feedback and workflow                      | Product design evidence               |
| C2 controlled pilot | Representative tasks and agreed success criteria   | Feasibility and user evidence         |
| C3 paid pilot       | Signed fee and collection                          | Willingness to pay within pilot scope |
| C4 production       | Repeated accepted outcomes in operating workflow   | Product use and reliability           |
| C5 renewal          | Continued paid contract after an informed decision | Retention evidence                    |
| C6 expansion        | Additional product, volume, team or geography      | Broader value evidence                |

Letters of intent, memorandum of understanding, unpaid pilot, paid pilot, production contract and collected renewal should be reported separately. Logos do not substitute for this classification.

### 6.4 Customer reference protocol

A customer reference should be authorised and structured. The interviewer confirms problem, prior workflow, buying process, deployment, time to value, quality, exceptions, support, use frequency, user breadth, price, renewal intent and alternatives. The investor records the relationship of the interviewee to the purchase and daily use.

The reference is customer-attested evidence. It is stronger when reconciled to contract, product and finance records. A founder-selected champion can provide valuable detail and may not represent procurement, security, finance or end-user views.

## 6.5 Distribution and sales motion


AI applications can use founder-led sales, product-led adoption, direct enterprise sales, marketplaces, cloud partners, systems integrators, channel partners and embedded distribution. Each motion has distinct cost, control, time and concentration.

| Distribution measure  | Definition                                                           | Common distortion                                |
|-----------------------|----------------------------------------------------------------------|--------------------------------------------------|
| Qualified opportunity | Account meets explicit customer, need, authority and timing criteria | Demo request counted as pipeline                 |
| Sales cycle           | First qualified contact to signed contract                           | Excludes procurement or security delay           |
| CAC                   | Fully loaded acquisition cost for the relevant cohort                | Omits founder, partner or implementation cost    |
| Payback               | CAC divided by cohort gross-profit contribution                      | Uses revenue or target margin                    |
| Win rate              | Won opportunities divided by comparable resolved opportunities       | Removes losses or stale deals                    |
| Expansion             | Additional collected recurring value from existing customer          | One-off services included                        |
| Channel concentration | Share of qualified pipeline, revenue and collection by source        | Partner-sourced and partner-influenced conflated |

## 7. TECHNICAL ARCHITECTURE, MODEL SUPPLY CHAIN AND PORTABILITY

### 7.1 Architecture evidence

The architecture should show the complete path from user and source system to accepted outcome. It includes identity, permissions, data movement, retrieval, prompts, model routes, tools, code execution, guardrails, human review, storage, observability and billing. The diagram should distinguish customer environments, company systems and third parties.



**Figure 3. AI application-layer architecture and evidence stack**

Source: Matchpoint architecture; provider, risk and security context [11,13-30].

| Component            | Underwriting evidence                              | Key dependency                 |
|----------------------|----------------------------------------------------|--------------------------------|
| User and identity    | Roles, authentication and authorisation tests      | Customer identity provider     |
| Source integration   | Data contract, lineage, failure and reconciliation | Customer API and data quality  |
| Retrieval and memory | Corpus, freshness, access and deletion             | Vector store and source rights |

| Component     | Underwriting evidence                               | Key dependency                           |
|---------------|-----------------------------------------------------|------------------------------------------|
| Orchestration | Workflow version, retries, timeout and idempotency  | Application code and tool APIs           |
| Model route   | Provider, model/version, fallback and thresholds    | Price, capacity, lifecycle and terms     |
| Tools         | Permission, validation and transaction boundary     | External service and credential security |
| Human review  | Reviewer, queue, authority and escalation           | Skilled capacity and operating hours     |
| Monitoring    | Quality, cost, latency, incident and drift measures | Event completeness                       |

## 7.2 Provider register


OpenAI, Anthropic, Google and AWS publish different pricing structures for models, caching, batching, tools, context and deployment [13-16]. Google also publishes model and endpoint lifecycle information, including retirements [17]. The register records the exact service used, not the provider brand alone.

| Field            | Required content                                                         |
|------------------|--------------------------------------------------------------------------|
| Service identity | Provider, product, model, version, region and account                    |
| Commercial       | Price unit, tier, committed spend, discount, currency and renewal        |
| Data             | Input/output use, retention, logging, training setting and subprocessors |
| Performance      | Representative quality, latency, availability and capacity               |
| Lifecycle        | Release, deprecation, retirement and migration notice                    |
| Control          | Version pinning, fallback, rate limit and spend limit                    |
| Exit             | Alternative provider/model, migration work, test status and data export  |

List prices can change and realised spend depends on input, output, retries, caching, tool use, context, modality and volume. Diligence should reconcile invoice line items to product telemetry for a representative period.

## 7.3 Concentration and switching

The FTC's study of selected cloud-provider and AI-developer partnerships records cloud-spend commitments, exclusivity features, access to sensitive information and potential switching costs [11]. The report concerns those partnerships and does not establish a private application company's terms. It provides a useful diligence lens.



**Figure 4. Provider concentration, portability and migration map**

Source: Matchpoint framework; provider and lifecycle context [11,13-22,40].

Concentration is measured across spend, accepted outcomes, critical features, regions and customers. A company can appear multi-model while one provider handles every revenue-critical task. A fallback that passes a synthetic prompt can still fail on customer tools, long context, safety behaviour, structured output or latency.

#### 7.4 Portability test

A portability test runs a representative evaluation suite against an alternative model or deployment. It records code changes, prompt changes, output-schema differences, quality, review burden, latency, cost, safety, data terms and migration time. The result can be portable, partially portable or concentrated.

The investor should model at least three shocks: a 50% increase in model unit cost, withdrawal of the primary model within 90 days and capacity throttling during peak use. These percentages and periods are unverified illustrative scenario inputs. Management should replace them with its contractual and operating evidence.

#### 7.5 Open-weight and licence dependency

Meta's Llama 4 Community License grants specified rights and includes redistribution, attribution, acceptable-use and scale-related terms [22]. Other open-weight models have different licences. Open-weight does not mean unrestricted. The company should retain the exact licence version, model card, download source, modifications, notices, acceptable-use review and redistribution analysis.

### 8. EVALUATION, RELIABILITY, SECURITY AND OPERATING CONTROLS

#### 8.1 Evaluation contract

NIST's AI RMF uses govern, map, measure and manage functions [23]. Its Generative AI Profile identifies generative-AI risks and recommended actions [24]. OpenAI's current model guidance recommends representative task evaluation before migration or optimisation [29]. Anthropic's evaluation guidance begins with specific, measurable and relevant success criteria [30]. These sources support a practical release discipline.

An evaluation contract defines task population, sample, ground truth, grader, dimensions, threshold, severity, subgroup, model/version, prompt, tool environment and release decision. The test set should include normal, difficult, adversarial, incomplete and out-of-scope cases.

| Evaluation dimension | Example measure                          | Release treatment                              |
|----------------------|------------------------------------------|------------------------------------------------|
| Task success         | Accepted answer or completed tool action | Minimum weighted pass rate                     |
| Factual support      | Claims supported by permitted evidence   | Critical failure if material claim unsupported |

| Evaluation dimension | Example measure                                            | Release treatment                         |
|----------------------|------------------------------------------------------------|-------------------------------------------|
| Completeness         | Required fields and decisions present                      | Schema and reviewer check                 |
| Safety and policy    | Prohibited, harmful or unauthorised action                 | Zero tolerance for defined critical cases |
| Privacy              | Personal or confidential data handling                     | Data-flow and output test                 |
| Robustness           | Performance under noise, prompt variation and missing data | Threshold by customer segment             |
| Latency              | Time to accepted outcome                                   | Customer decision-window threshold        |
| Cost                 | Full cost per accepted outcome                             | Product and margin threshold              |

## 8.2 Release and monitoring

One benchmark is insufficient. Releases can change model version, prompt, tools, retrieval, policies, code, data and customer configuration. The company should preserve a release manifest and rerun the required suite. Production monitoring then compares acceptance, failure, latency and cost with release expectations.

| Release record | Minimum evidence                                                |
|----------------|-----------------------------------------------------------------|
| Identity       | Application, model, prompt, tool, data and policy versions      |
| Change         | Description, owner, reason and affected customers               |
| Evaluation     | Dataset version, results, failures and comparison               |
| Approval       | Product, engineering, security and domain authority as required |
| Rollout        | Cohort, feature flag, limits, monitoring and rollback           |
| Disposition    | Release, restricted release, hold or retire                     |

## 8.3 Reliability and incident evidence

Reliability is measured at the accepted-outcome boundary. Availability can be high while a critical tool silently fails. A model can respond while providing an invalid schema. A workflow can complete after the customer's decision window. The incident register should include provider failures, degraded quality, retrieval errors, data staleness, tool errors, security events, spend anomalies and human-review backlogs.

| Incident field                      | Purpose                         |
|-------------------------------------|---------------------------------|
| Start, detection and recovery       | Measure exposure and response   |
| Affected task, customer and outcome | Establish operating consequence |
| Model, data, tool and code versions | Reproduce the condition         |
| Severity and authority              | Trigger required escalation     |
| Customer communication              | Evidence contractual response   |
| Root cause and corrective action    | Reduce recurrence               |
| Cost, credit and concession         | Reconcile economic impact       |

## 8.4 Security

NIST's SSDF community profile adds generative-AI practices to secure software development [25]. CISA and NCSC address secure design, development, deployment and operation for AI systems [26]. OWASP identifies risks including prompt injection, sensitive-information disclosure, supply-chain vulnerabilities and unbounded consumption [27]. MITRE ATLAS provides an evolving threat knowledge base for predictive, generative and agentic AI systems [28].

| Threat                      | Application exposure                              | Control evidence                                                |
|-----------------------------|---------------------------------------------------|-----------------------------------------------------------------|
| Prompt or context injection | Untrusted content alters model behaviour          | Trust boundaries, sanitisation, instruction hierarchy and tests |
| Excessive agency            | Model invokes consequential tool beyond authority | Least privilege, allowlists, approval and transaction limits    |
| Sensitive disclosure        | Prompt, retrieval, log or output reveals data     | Access, minimisation, redaction, encryption and monitoring      |
| Supply-chain compromise     | Model, package, embedding or tool is malicious    | Provenance, scanning, pinning and vendor review                 |
| Unbounded consumption       | Requests create service degradation or cost shock | Rate, token, tool, recursion and spend controls                 |
| Model theft or extraction   | Repeated queries reproduce proprietary behaviour  | Abuse detection, limits and response                            |
| Data or RAG poisoning       | Source content corrupts output                    | Source authority, ingestion control and anomaly tests           |

Security certifications and penetration tests are evidence within their scope and date. They do not replace architecture review, incident history or customer-specific controls.

## 9. DATA RIGHTS, PRIVACY, REGULATION AND INTELLECTUAL PROPERTY

### 9.1 Exact service terms

OpenAI states that business and API inputs and outputs are excluded from training by default [18]. Its services agreement sets the business relationship and relevant content rights [19]. Anthropic publishes commercial-product data-use information [20]. Google's Gemini API terms distinguish unpaid and paid services, including use of submitted content and human review for unpaid services and different treatment for paid services [21]. Google's zero-data-retention guidance identifies product settings and features relevant to retention [40].

These terms can change. The company should retain dated contract documents, account settings and data-flow evidence. A privacy answer based only on a provider's public brand-level statement is incomplete.

### 9.2 Data-rights register

| Data class        | Source                 | Right and purpose                        | Training or evaluation use | Retention and deletion      | Transfer and processor      |
|-------------------|------------------------|------------------------------------------|----------------------------|-----------------------------|-----------------------------|
| Customer input    | Customer system        | Contract and applicable lawful basis     | Defined by agreement       | Customer and company policy | Named providers and regions |
| Customer feedback | User review            | Product operation and improvement term   | Scope-specific             | Defined period              | Approved processors         |
| Public data       | Identified publication | Copyright, database and use analysis     | Recorded                   | Provenance retained         | Provider terms checked      |
| Licensed data     | Vendor or partner      | Licence scope, field, geography and term | Express permission         | Contract term               | Transfer restriction        |
| Synthetic data    | Company generation     | Source and generation rights             | Documented                 | Versioned                   | Model terms checked         |
| Employee content  | Staff or contractor    | Assignment, confidentiality and policy   | Restricted purpose         | Employment and legal policy | Approved tools              |

The register should record data quality and usefulness separately from legal rights. Lawful data can be irrelevant. Useful data can be unusable because the required right is absent.

### 9.3 UAE, DIFC, ADGM, Saudi and UK data protection

The UAE Federal Decree-Law No. 45 of 2021 establishes a federal personal-data framework within its scope [33]. DIFC Regulation 10 addresses personal-data processing through autonomous and semi-autonomous systems [34]. ADGM guidance addresses automated individual decision-making, including the need for meaningful human review in relevant cases [35]. Saudi Arabia's PDPL and implementing framework establish controller obligations within scope [36]. The UK Information Commissioner's Office publishes AI guidance covering lawfulness, fairness, transparency, purpose limitation, minimisation, accuracy, storage, security and accountability [37].

| Jurisdiction or regime | Product-role question                                                           | Diligence evidence                                                |
|------------------------|---------------------------------------------------------------------------------|-------------------------------------------------------------------|
| UAE federal            | Is personal data processed within territorial and material scope?               | Processing record, basis, notice, processor and transfer map      |
| DIFC                   | Does an autonomous or semi-autonomous system process personal data?             | Regulation 10 assessment and control record                       |
| ADGM                   | Does solely automated processing create legal or similarly significant effects? | Decision map, exception, meaningful human review and appeal       |
| Saudi Arabia           | Is the company a controller or processor under the PDPL framework?              | Basis, purpose, notice, security, transfer and retention evidence |
| United Kingdom         | How do UK data-protection principles apply across the AI lifecycle?             | DPIA, fairness, transparency, rights and security evidence        |

This table is a diligence index. It is not a legal conclusion.

### 9.4 EU AI Act roles and transparency

The EU AI Act distinguishes model providers, system providers, deployers and other operators and applies obligations by role and risk [31]. The European Commission states that Article 50 transparency obligations apply from 2 August 2026, with specific transition treatment for certain systems placed on the market earlier [32]. A GCC company can be within scope through product placement, deployment or affected persons in the Union, depending on the facts.

The product-role matrix should be completed for each use case and jurisdiction. A company can be a provider for one feature, a deployer for another and a downstream integrator of a third-party general-purpose model. Contract language should align with actual architecture and control.

### 9.5 Copyright, patent and code provenance

The US Copyright Office has published reports addressing digital replicas, copyrightability of generative-AI outputs and AI training [38]. USPTO's revised November 2025 guidance states that ordinary inventorship standards apply to AI-assisted inventions and only natural persons can be inventors [39]. Provider agreements and open-model licences allocate and reserve different rights [19,21,22].

| IP asset             | Required record                                            | Diligence issue                            |
|----------------------|------------------------------------------------------------|--------------------------------------------|
| Source code          | Repository, author, assignment, licence and dependency     | Ownership and open-source obligations      |
| Prompt and workflow  | Version, author, customer restriction and secret treatment | Copyright, confidentiality and portability |
| Dataset              | Source, licence, collection and permitted purpose          | Copyright, database, privacy and contract  |
| Fine-tune or adapter | Base licence, training data, method and deployment         | Rights and provider conditions             |
| Output               | Human contribution, review and customer terms              | Protectability and infringement risk       |
| Invention            | Human conception record and assignment                     | Inventorship and patent ownership          |
| Brand and domain     | Registration, ownership and use                            | Freedom to operate and customer confusion  |

Generated code should pass the same review, test, security, provenance and licence controls as human-written code. Diligence should sample commits and dependencies rather than rely only on a policy document.

## 10. REVENUE QUALITY, RETENTION AND DISTRIBUTION

### 10.1 Contract-to-cash

Revenue underwriting begins with a customer-level bridge. The bridge links opportunity, signed contract, performance obligation, product access, usage, accepted outcomes, invoice, collection, credit, concession, renewal and expansion. It also identifies reseller, related-party, founder-network and strategic-partner relationships.

| Stage       | Evidence                                 | Investor test                        |
|-------------|------------------------------------------|--------------------------------------|
| Opportunity | Qualified need, buyer, budget and timing | Definition and ageing are consistent |

| Stage                | Evidence                                               | Investor test                                           |
|----------------------|--------------------------------------------------------|---------------------------------------------------------|
| Contract             | Executed agreement, term, scope, price and termination | Signed by authorised parties; side letters captured     |
| Deployment           | Environment, user and integration record               | Production, pilot and test are separated                |
| Performance          | Accepted outcomes and service evidence                 | Contracted service was delivered                        |
| Invoice              | Invoice, date, amount, tax and due date                | Reconciles to contract and ledger                       |
| Collection           | Bank or payment record and allocation                  | Cash is received and not reversed                       |
| Credit or concession | Credit note, discount, free period or service recovery | Net revenue and customer issue captured                 |
| Renewal              | Informed continuation and new term                     | Auto-renewal, negotiated renewal and holdover separated |

Annual recurring revenue should have a written policy. The investor should test whether pilots, consumption commitments, implementation, professional services, non-cancellable contracts, usage minimums, one-time fees and overdue invoices are included. Reported ARR should reconcile to a customer schedule and general ledger within the policy.

## 10.2 Cohorts

A cohort groups comparable customers by start period, segment, product, contract or acquisition source. Revenue retention should be shown on both gross and net bases. Logo retention should sit beside revenue retention because one expanding customer can conceal broad customer loss.

| Cohort measure             | Numerator                                                           | Denominator                       | Interpretation                |
|----------------------------|---------------------------------------------------------------------|-----------------------------------|-------------------------------|
| Gross revenue retention    | Opening recurring revenue less churn and contraction                | Opening recurring revenue         | Retention before expansion    |
| Net revenue retention      | Opening recurring revenue less churn and contraction plus expansion | Opening recurring revenue         | Retention including expansion |
| Logo retention             | Opening customers retained                                          | Opening customers                 | Breadth of customer retention |
| Accepted-outcome retention | Outcomes from retained comparable workflows                         | Opening comparable outcomes       | Continued product use         |
| Collection retention       | Collected recurring value from retained customers                   | Opening collected recurring value | Cash quality                  |

The company should disclose cohort size and age. A young cohort can show high net retention before the first meaningful renewal. A customer with a large initial rollout can create contraction without dissatisfaction. The evidence should explain the movement.

## 10.3 Concentration

Concentration is measured by revenue, gross profit, collection, accepted outcomes, pipeline and critical data rights. A high-revenue customer can be low margin due to dedicated infrastructure, services, support or model consumption. A low-revenue design partner can be operationally critical because the product depends on its data or domain expertise.

## 10.4 Sales efficiency

Customer-acquisition cost should include sales compensation, marketing, events, founder time, channel fees, solution engineering, security diligence, pilots and implementation where they are required to win the cohort. Payback uses cohort gross profit rather than revenue.

An early-stage company can report the underlying activity even when a stable CAC is unavailable. Useful evidence includes founder hours per stage, number of security reviews, pilot duration, implementation effort, win/loss reasons and time from signature to production acceptance.



## 11. UNIT ECONOMICS, GROSS MARGIN AND CAPITAL PLAN


### 11.1 Per-outcome cost ledger

Provider token pricing is only one cost line [13-16]. The full cost of an accepted outcome can include source-system access, document processing, embeddings, retrieval, context storage, model input and output, caching, grounding, tools, code execution, retries, failed calls, orchestration, observability, security, human review, support and customer-specific infrastructure.

| Cost pool                  | Allocation driver                                           | Evidence                                |
|----------------------------|-------------------------------------------------------------|-----------------------------------------|
| Model inference            | Actual tokens, images, audio, calls or provisioned capacity | Telemetry to invoice                    |
| Retrieval and storage      | Indexed volume, query, storage time and egress              | Cloud records and invoice               |
| Tools and data services    | Invocation, search, transaction or licence                  | Tool log and invoice                    |
| Human review               | Minutes by role and outcome class                           | Workflow time record and payroll rate   |
| Support and operations     | Tickets, incidents, on-call and customer environment        | Ticket and staffing record              |
| Infrastructure             | Compute, database, network, observability and security      | Tagged cloud cost and allocation policy |
| Customer-specific delivery | Integration, configuration and dedicated environment        | Project record and contract             |

### 11.2 Illustrative gross-profit bridge

Assume, solely for illustration, that one customer purchases 10,000 accepted outcomes per month at USD 1.20 each. Monthly revenue is USD 12,000. Model and retrieval cost is USD 1,650; tools and infrastructure cost USD 900; human review costs USD 1,800; support and reliability allocation costs USD 750. Illustrative gross profit is USD 6,900 and gross margin is 57.5%.



**Figure 5. Illustrative monthly gross-profit bridge per accepted outcome**

Source: Unverified illustrative management assumptions; Matchpoint calculation.

| Line                            | Unverified illustrative assumption | Per accepted outcome |
|---------------------------------|------------------------------------|----------------------|
| Collected or recognised revenue | USD 12,000                         | USD 1.20             |
| Model and retrieval             | USD 1,650                          | USD 0.165            |
| Tools and infrastructure        | USD 900                            | USD 0.090            |
| Human review                    | USD 1,800                          | USD 0.180            |

| Line                      | Unverified illustrative assumption | Per accepted outcome |
|---------------------------|------------------------------------|----------------------|
| Support and reliability   | USD 750                            | USD 0.075            |
| Illustrative gross profit | USD 6,900                          | USD 0.690            |
| Illustrative gross margin | 57.5%                              | 57.5%                |

Every input in this example is an unverified illustrative management assumption. It proves no company or customer result. Attributed revenue, cost reduction and loss reduction remain USD 0.

### 11.3 Margin decomposition

Margin movement should be decomposed by price, volume, mix, model route, prompt/context size, caching, batch use, review rate, failure rate, customer-specific service and accounting classification. Duolingo's 2025 disclosure connecting generative-AI and hosting costs to gross-margin movement illustrates why direct AI cost belongs in the bridge [10]. The different gross margins reported by Palantir, UiPath and C3 AI illustrate product and business-model variation [7-9].

A declining provider price can improve margin. It can also invite price competition, enable new entrants or support a higher-quality model route that preserves unit cost. Underwriting should model the company response rather than assume all provider savings remain with shareholders.

### 11.4 Capital intensity

Application companies can be capital-light or capital-intensive. Capital requirements rise with long enterprise sales cycles, dedicated deployments, data acquisition, model training, guaranteed capacity, regulated permissions, security requirements, human operations and working capital. Committed cloud spend can create a fixed obligation before customer volume arrives.

| Capital line             | Base evidence                                       | Downside question                               |
|--------------------------|-----------------------------------------------------|-------------------------------------------------|
| Product and engineering  | Headcount, roadmap and release gates                | Which work is required for retention?           |
| AI and data              | Provider, training, evaluation and data commitments | Which cost scales ahead of revenue?             |
| Sales and implementation | Funnel, cycle and delivery capacity                 | What happens when deals delay?                  |
| Security and compliance  | Customer and jurisdiction requirements              | Which permission blocks revenue?                |
| Working capital          | Billing, collection and vendor payment terms        | Does cash use grow with revenue?                |
| Financing                | Cash, debt, equity and restricted funds             | Which covenants or milestones constrain action? |

### 11.5 Runway bridge


The runway model should start with bank-verified cash and reconcile monthly receipts and payments. Management then applies hiring, provider, sales and regulatory gates. A board should see base, delay, lower-revenue, cost-shock and no-raise cases.

The plan should distinguish committed, contracted, approved, discretionary and unapproved spend. Use of funds becomes a list of evidence milestones: production acceptance, representative evaluation, renewal cohort, gross-margin threshold, provider migration, security requirement or licensing step.

## 12. UNDERWRITING SCORECARD, INVESTMENT MEMO AND DOWNSIDE

### 12.1 Scorecard design

The scorecard records evidence sufficiency, not founder charisma or category excitement. Each domain receives a status: verified, partially verified, management estimated, unsupported or adverse. A numerical score can support consistency but should not conceal a critical issue.

**Figure 6. Underwriting scorecard and gating logic**

Source: Matchpoint underwriting framework.

| Domain                       | Weight | Evidence gate                                   | Critical hold                         |
|------------------------------|--------|-------------------------------------------------|---------------------------------------|
| Customer problem and outcome | 12     | Observed workflow and accepted outcome          | No valuable repeated task             |
| Product and usage            | 10     | Production use and adoption                     | Demonstration only                    |
| Evaluation and reliability   | 12     | Representative release and production evidence  | Material failures uncontrolled        |
| Architecture and portability | 8      | Dependency map and tested fallback              | Unmanageable provider concentration   |
| Data, privacy and IP         | 12     | Rights and role evidence                        | Material rights absent                |
| Revenue and retention        | 14     | Contract-to-cash and cohorts                    | Revenue materially misstated          |
| Unit economics               | 12     | Reconciled cost per accepted outcome            | Cost ledger incomplete                |
| Distribution                 | 8      | Qualified funnel and repeatable buying evidence | Pipeline unsupported                  |
| Team and governance          | 6      | Named owners and control record                 | Authority or integrity concern        |
| Capital and terms            | 6      | Runway, downside and financing plan             | Insolvency or financing gap unmanaged |

The weights are unverified illustrative management assumptions. The investment committee should approve its own weights and critical holds.

## 12.2 Investment-memo evidence packet

| Memo section | Required exhibit                                                         |
|--------------|--------------------------------------------------------------------------|
| Thesis       | Claim register and disconfirming evidence                                |
| Market       | Bottom-up workflow market and source limitations                         |
| Product      | Before/after workflow, accepted outcome and product demonstration record |
| Technology   | Architecture, provider register, evaluation and portability test         |
| Customers    | Contract-to-cash sample, cohorts and structured references               |
| Economics    | Per-outcome ledger, gross-margin bridge and sales efficiency             |

| Memo section         | Required exhibit                                                  |
|----------------------|-------------------------------------------------------------------|
| Legal and regulatory | Corporate, IP, data and product-role matrices                     |
| Team                 | Roles, references, ownership, incentives and key-person plan      |
| Financing            | Cap table, instrument terms, use of funds, runway and downside    |
| Decision             | Conditions, reserves, governance rights and post-investment gates |

### 12.3 Downside scenarios

| Scenario   | Illustrative shock                                   | Evidence to model                            | Management response                       |
|------------|------------------------------------------------------|----------------------------------------------|-------------------------------------------|
| Model cost | Primary model unit cost rises 50%                    | Current invoice and workload distribution    | Routing, price, context and fallback plan |
| Quality    | First-pass acceptance falls 15 percentage points     | Eval and production acceptance history       | Rollback, review and customer restriction |
| Provider   | Primary version retires in 90 days                   | Lifecycle notice and migration test          | Alternative model and release plan        |
| Customer   | Largest customer churns                              | Revenue, gross profit and cost concentration | Cost, cash and pipeline response          |
| Sales      | Median enterprise cycle extends six months           | Stage history and procurement evidence       | Hiring and runway adjustment              |
| Regulation | Product requires additional assessment or permission | Role matrix and counsel analysis             | Scope, geography or control change        |

All shocks are unverified illustrative scenario inputs. The diligence team should replace them with company-specific evidence and committee-approved severities.

### 12.4 Valuation and structure

The valuation analysis should separate current evidence from option value. Current value can use observed revenue, retention, gross profit, growth and cash characteristics. Option value can reflect new products, geographies, datasets or model capability and should carry probability, capital and time.

Structure can address evidence gaps through staged funding, milestones, reserves, information rights, consent rights, board rights, founder vesting, option pools and warranties. Legal counsel should draft and assess the transaction. Milestones should be within management influence and defined by auditable evidence.

## 13. FOUNDER AND INVESTOR PLAYBOOKS

### 13.1 Founder diligence-room index

| Folder                  | Minimum content                                                      | Update frequency      |
|-------------------------|----------------------------------------------------------------------|-----------------------|
| 01 Corporate            | Incorporation, registers, group, licences and board materials        | On change             |
| 02 Capital              | Cap table, instruments, options, debt and waterfall                  | On transaction        |
| 03 Customers            | Contract schedule, signed samples, invoices, collections and credits | Monthly               |
| 04 Product              | Roadmap, releases, usage and accepted outcomes                       | Each release/monthly  |
| 05 AI and evaluation    | Models, prompts, suites, results, incidents and migrations           | Each release          |
| 06 Data and IP          | Rights, provenance, assignments, licences and inventions             | On change/quarterly   |
| 07 Security and privacy | Architecture, policies, tests, incidents, DPIAs and processors       | Quarterly/on incident |
| 08 Finance              | Management accounts, bank, budget, margin and runway                 | Monthly               |
| 09 Commercial           | Pipeline, funnel, pricing, losses and partners                       | Weekly/monthly        |
| 10 Team                 | Organisation, contracts, compensation, references and hiring plan    | Monthly               |

| Folder    | Minimum content               | Update frequency          |
|-----------|-------------------------------|---------------------------|
| 11 Claims | Pitch and memo claim register | Before every external use |

The founder should preserve original records and provide reconciliations rather than a curated collection of favourable screenshots. Customer-identifiable and personal data should be shared only under appropriate authority and controls.

### 13.2 Investor fieldwork

The investor begins with a request list and evidence map. It samples customers across retained, expanded, pilot, churned and lost-opportunity groups. It reproduces a bounded evaluation. It traces selected contracts to usage, invoices and cash. It reconciles provider invoices to telemetry. It samples source-code ownership and data rights. It tests one provider migration. It reviews incidents and board reporting.

| Fieldwork                   | Sample objective                     | Evidence retained                        |
|-----------------------------|--------------------------------------|------------------------------------------|
| Workflow observation        | Confirm task, review and acceptance  | Dated process map and limitations        |
| Evaluation reproduction     | Test representative quality          | Inputs, versions, output and disposition |
| Contract-to-cash tracing    | Test revenue existence and quality   | Contract, invoice, collection and credit |
| Cohort recreation           | Test retention calculations          | Customer schedule and formula            |
| Cost reconciliation         | Test full cost per accepted outcome  | Invoice, telemetry and allocation        |
| Provider migration          | Test portability and hidden work     | Change log, eval, cost and time          |
| Security/privacy review     | Test material controls and incidents | Architecture, test and remediation       |
| Founder and team references | Test history, roles and integrity    | Source, relationship and factual notes   |

### 13.3 Management questions

1. What exact customer outcome is accepted, by whom and under which quality threshold?
2. Which production task fails most often, and what is its customer consequence?
3. Which reported productivity result has an observed baseline and counterfactual?
4. Which provider, model or tool is critical to retained revenue?
5. How long would a representative migration take, and when was it last tested?
6. Which data claim has the weakest rights evidence?
7. Which customer cohort produces the highest collected gross profit?
8. Which customer requires the most human review or support?
9. Which material claim in the pitch deck remains a management estimate?
10. What result would cause management to restrict, redesign or stop the product?
11. What cash plan applies if the next financing is delayed by six months?
12. Which milestone converts the largest remaining uncertainty into evidence?

## 14. GATED 12-18 MONTH ROADMAP

The roadmap begins with definition and reconciliation. It advances through representative evaluation, production evidence, retained economics and financing readiness. Calendar timing depends on company state, customer process and regulatory requirements.

**Figure 7. Gated 12-18 month founder and investor roadmap**

Source: Matchpoint adoption framework.

| Phase              | Target period | Deliverable                                                   | Exit gate                                    |
|--------------------|---------------|---------------------------------------------------------------|----------------------------------------------|
| 0 Definition       | Month 0-1     | Workflow, accepted outcome and claim register                 | Named owners and stable definitions          |
| 1 Evidence room    | Month 1-2     | Contract-to-cash, architecture, provider and rights registers | Core records reconcile                       |
| 2 Evaluation       | Month 2-4     | Representative suite, release gate and failure taxonomy       | Critical thresholds pass                     |
| 3 Economics        | Month 3-6     | Per-outcome cost and cohort gross-margin bridge               | Invoice, telemetry and ledger reconcile      |
| 4 Portability      | Month 4-7     | Alternative model/provider migration test                     | Defined quality, time and cost threshold     |
| 5 Commercial proof | Month 4-10    | Production adoption, reference and first renewal evidence     | Customer and collection gates pass           |
| 6 Scale control    | Month 8-14    | Monitoring, security, support and unit-economics controls     | Growth does not degrade outcome or margin    |
| 7 Financing        | Month 10-18   | Evidence-led memo, downside, valuation and use of funds       | Committee or board-approved transaction path |

The roadmap is a sequence of evidence gates. A company can move faster when reliable evidence already exists. It can remain at a phase when a gate fails. The disposition can be advance, restrict, redesign, retain as research or stop.

## 15. LIMITATIONS, RESEARCH AGENDA AND CONCLUSION

This paper presents a diligence framework based on public evidence. It does not report an investment, customer cohort, production evaluation or financial return. Market data describes different datasets and periods. Public-company filings describe mature businesses with economic profiles that may not resemble a private application startup. Provider pricing and terms can change after the evidence cut-off.

The framework requires company evidence for implementation. Future research should analyse anonymised contract-to-cash cohorts, accepted-outcome economics, model migration events, incident costs and valuation outcomes across application categories. Comparative studies should report complete workload, model, prompt, tool, review, latency, failure and cost boundaries. GCC research should distinguish federal, DIFC, ADGM and Saudi legal contexts and use qualified legal analysis.

The investment conclusion is disciplined. AI-native application underwriting starts with one valuable customer workflow and one quality-adjusted accepted outcome. It then follows the evidence through architecture, model

supply chain, evaluation, data rights, production use, contract, collection, cohort, cost and control. The company earns an AI-native classification when this chain is material and reproducible. It earns an investment case when retained customer value, gross profit, governance and financing price remain credible under downside.

For A5 investors, the framework converts category excitement into a claim register, fieldwork plan, scorecard and evidence-led memorandum. For B5 founders, it creates a management system for product quality, customer value, provider dependence, data rights, gross margin and fundraising readiness. Both parties benefit from one reconciled evidence room and explicit treatment of uncertainty.

## DECLARATIONS

Author. Matchpoint Partners. Named-person author attribution remains pending CK approval.

Evidence disclosure. The paper uses public market reports, company filings, provider documentation, laws, regulator material, standards and security resources available through 1 August 2026. It contains no T17 company, customer, fund, product, data-room or investment-committee evidence.

Scenario disclosure. All worked prices, costs, productivity effects, conversion rates, score weights, shocks, timing and valuation inputs are unverified illustrative management assumptions. Attributed Matchpoint or client revenue, cost reduction and loss reduction remain USD 0.

Advice boundary. This paper is not investment, legal, regulatory, accounting, tax, audit, privacy, intellectual-property, cybersecurity, valuation or technology advice. It is not an offer, solicitation, recommendation, warranty or promise of results.

## REFERENCES

- [1] Stanford Institute for Human-Centered AI (2026). AI Index Report 2026, Chapter 4: Economy. [https://hai.stanford.edu/assets/files/ai\\_index\\_report\\_2026\\_chapter\\_4\\_economy.pdf](https://hai.stanford.edu/assets/files/ai_index_report_2026_chapter_4_economy.pdf)
- [2] National Venture Capital Association and PitchBook (2026). Q2 2026 PitchBook-NVCA Venture Monitor. <https://nvca.org/pitchbook-nvca-venture-monitor/>
- [3] Carta (2026). State of Private Markets: 2025 in Review. <https://carta.com/data/state-of-private-markets-q4-2025/>
- [4] Carta and a16z speedrun (2025). The State of Seed 2025. <https://carta.com/data/resources/state-of-seed-2025/>
- [5] Carta (2026). State of Startup Compensation: H2 2025. <https://carta.com/data/startup-compensation-h2-2025/>
- [6] OECD.AI (2026). Venture capital investments in artificial intelligence through 2025. <https://oecd.ai/en/ai-publications/venture-capital-investments-in-artificial-intelligence-through-2025>
- [7] Palantir Technologies Inc. (2026). Annual Report on Form 10-K for the year ended 31 December 2025. <https://www.sec.gov/Archives/edgar/data/1321655/000132165526000011/pltr-20251231.htm>
- [8] UiPath, Inc. (2026). Annual Report on Form 10-K for the year ended 31 January 2026. <https://www.sec.gov/Archives/edgar/data/1734722/000173472226000012/path-20260131.htm>
- [9] C3.ai, Inc. (2026). Fiscal Fourth Quarter and Full Fiscal Year 2026 Results. <https://ir.c3.ai/news-releases/news-release-details/c3-ai-announces-fiscal-fourth-quarter-and-full-fiscal-year-2026>
- [10] Duolingo, Inc. (2026). Annual Report on Form 10-K for the year ended 31 December 2025. <https://www.sec.gov/Archives/edgar/data/1562088/000162828026012494/duol-20251231.htm>
- [11] US Federal Trade Commission (2025). Partnerships Between Cloud Service Providers and AI Developers. [https://www.ftc.gov/system/files/ftc\\_gov/pdf/p246201\\_aipartnerships6breport\\_redacted.pdf](https://www.ftc.gov/system/files/ftc_gov/pdf/p246201_aipartnerships6breport_redacted.pdf)
- [12] US Securities and Exchange Commission (2024). SEC Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence. <https://www.sec.gov/newsroom/press-releases/2024-36>
- [13] OpenAI. API Pricing. <https://developers.openai.com/api/docs/pricing>
- [14] Anthropic. Pricing. <https://docs.anthropic.com/en/docs/about-claude/pricing>
- [15] Google. Gemini Developer API Pricing. <https://ai.google.dev/gemini-api/docs/pricing>
- [16] Amazon Web Services. Amazon Bedrock Pricing. <https://aws.amazon.com/bedrock/pricing/>
- [17] Google Cloud. Vertex AI Model Versions and Lifecycle. <https://cloud.google.com/vertex-ai/generative-ai/docs/learn/model-versions>
- [18] OpenAI. Business Data Privacy, Security, and Compliance. <https://openai.com/business-data/>
- [19] OpenAI (2026). OpenAI Services Agreement. <https://cdn.openai.com/osa/openai-services-agreement.pdf>
- [20] Anthropic (2026). How do you use personal data in model training? <https://privacy.claude.com/en/articles/7996885-how-do-you-use-personal-data-in-model-training>
- [21] Google (2026). Gemini API Additional Terms of Service. <https://ai.google.dev/gemini-api/terms>
- [22] Meta Platforms, Inc. (2025). Llama 4 Community License Agreement. <https://github.com/meta-llama/llama-models/blob/main/models/llama4/LICENSE>
- [23] NIST (2023). Artificial Intelligence Risk Management Framework 1.0. <https://doi.org/10.6028/NIST.AI.100-1>
- [24] NIST (2024). Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. <https://doi.org/10.6028/NIST.AI.600-1>
- [25] NIST (2024). Secure Software Development Practices for Generative AI and Dual-Use Foundation Models. <https://doi.org/10.6028/NIST.SP.800-218A>
- [26] CISA and NCSC (2023). Guidelines for Secure AI System Development. <https://www.cisa.gov/news-events/alerts/2023/11/26/cisa-and-uk-ncsc-unveil-joint-guidelines-secure-ai-system-development>
- [27] OWASP GenAI Security Project (2024). OWASP Top 10 for LLM Applications 2025. <https://genai.owasp.org/resource/owasp-top-10-for-llm-applications-2025/>
- [28] MITRE. Adversarial Threat Landscape for AI Systems. <https://atlas.mitre.org/>
- [29] OpenAI. Model Guidance and Representative Evaluations. <https://developers.openai.com/api/docs/guides/latest-model>
- [30] Anthropic. Define Your Success Criteria. <https://docs.anthropic.com/en/docs/test-and-evaluate/define-success>
- [31] European Union (2024). Regulation (EU) 2024/1689, Artificial Intelligence Act. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32024R1689>
- [32] European Commission (2026). Transparency obligations under Article 50 of the AI Act. <https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act>
- [33] United Arab Emirates (2021). Federal Decree-Law No. 45 of 2021 concerning the Protection of Personal Data. <https://uaelegislation.gov.ae/en/legislations/1972/download>
- [34] Dubai International Financial Centre. Regulation 10: Processing Personal Data through Autonomous and Semi-autonomous Systems. <https://www.difc.com/business/registrars-and-commissioners/commissioner-of-data-protection/regulation-10>
- [35] Abu Dhabi Global Market Office of Data Protection (2025). Automated Individual Decision-Making, Including Profiling. <https://www.adgm.com/documents/office-of-data-protection/resources/brochures-flyers/adgm-data-subject-rights-automated-individual-decision-making-brochure.pdf>
- [36] Saudi Data and Artificial Intelligence Authority. Personal Data Protection Law. <https://sdaia.gov.sa/en/SDAIA/about/Documents/Personal%20Data%20English%20V2-23April2023-%20Reviewed-.pdf>
- [37] UK Information Commissioner's Office. Guidance on AI and Data Protection. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/about-this-guidance/>
- [38] US Copyright Office (2024-2025). Copyright and Artificial Intelligence Reports. <https://www.copyright.gov/ai/>

[39] US Patent and Trademark Office (2025). Revised Inventorship Guidance for AI-Assisted Inventions. <https://www.uspto.gov/subscription-center/2025/revised-inventorship-guidance-ai-assisted-inventions>

[40] Google (2026). Zero Data Retention in the Gemini Developer API. <https://ai.google.dev/gemini-api/docs/zdr>

#### APPENDIX A. MINIMUM CLAIM REGISTER

| Field      | Content                                                                |
|------------|------------------------------------------------------------------------|
| Claim ID   | Stable identifier                                                      |
| Claim text | Exact words used externally or in the memo                             |
| Type       | Product, AI, customer, financial, legal, market or team                |
| Class      | O1 observed, O2 reproduced, O3 attested, M1 estimated or U unsupported |
| Source     | System, document, test or person                                       |
| Scope      | Product, customer, cohort, geography and period                        |
| Method     | Calculation or evaluation method                                       |
| Owner      | Named accountable person                                               |
| Approval   | Approved use and date                                                  |
| Expiry     | Date or event requiring refresh                                        |

#### APPENDIX B. MINIMUM EVALUATION RECORD

| Block        | Minimum content                                               |
|--------------|---------------------------------------------------------------|
| Identity     | Product, task, release, model, prompt, tools and owner        |
| Dataset      | Source, rights, sample, version, population and exclusions    |
| Ground truth | Method, reviewer, disagreement and uncertainty                |
| Criteria     | Quality, safety, privacy, latency and cost thresholds         |
| Results      | Aggregate, segment, severity, failure and confidence evidence |
| Comparison   | Prior release, baseline, alternative model and ablation       |
| Decision     | Release, restrict, hold, redesign or retire                   |
| Monitoring   | Production measures, limits, alert and rollback               |

#### APPENDIX C. CONTRACT-TO-CASH SAMPLE

For each sampled customer, retain the signed agreement, amendments and side letters; classify pilot or production; confirm product, term, price, termination and service; reconcile active users and accepted outcomes; trace invoice and credit notes; trace collection; identify concessions; and confirm renewal or churn. Record every exception and its memo treatment.

#### APPENDIX D. ILLUSTRATIVE UNIT-ECONOMICS FORMULAE

Accepted outcome rate equals accepted outcomes divided by generated candidate outcomes. First-pass acceptance equals accepted outcomes with no material correction divided by reviewed outcomes. Review minutes per accepted outcome equals complete human review and escalation minutes divided by accepted outcomes. Full variable cost per accepted outcome equals all attributable model, retrieval, tool, infrastructure, review and support cost divided by accepted outcomes. Cohort gross-profit contribution equals cohort revenue less complete cohort cost to serve.

All formula inputs require a defined period, product, customer population and accounting policy. Management-estimated inputs remain labelled until reconciled to observed records.

#### APPENDIX E. RED-FLAG REGISTER

1. AI capability cannot be reproduced on representative customer tasks.
2. Pilots, proofs of concept or memoranda are reported as production recurring revenue.
3. Revenue schedules do not reconcile to invoices, collections and credits.
4. Gross margin excludes model, cloud, human review, support or implementation required to serve customers.
5. A proprietary-data claim lacks source and rights evidence.
6. One provider supports nearly every accepted outcome and no migration test exists.
7. The product can take consequential actions without tested authority limits and rollback.
8. Evaluation uses synthetic happy-path cases without production failure classes.
9. Customer concentration is measured on revenue and omits gross profit, data and operating dependence.
10. Market size begins with aggregate AI spending and lacks a workflow-level paying unit.
11. The financing plan assumes a future round without a delayed or no-raise case.
12. Material pitch-deck claims lack a dated source, scope and accountable owner.