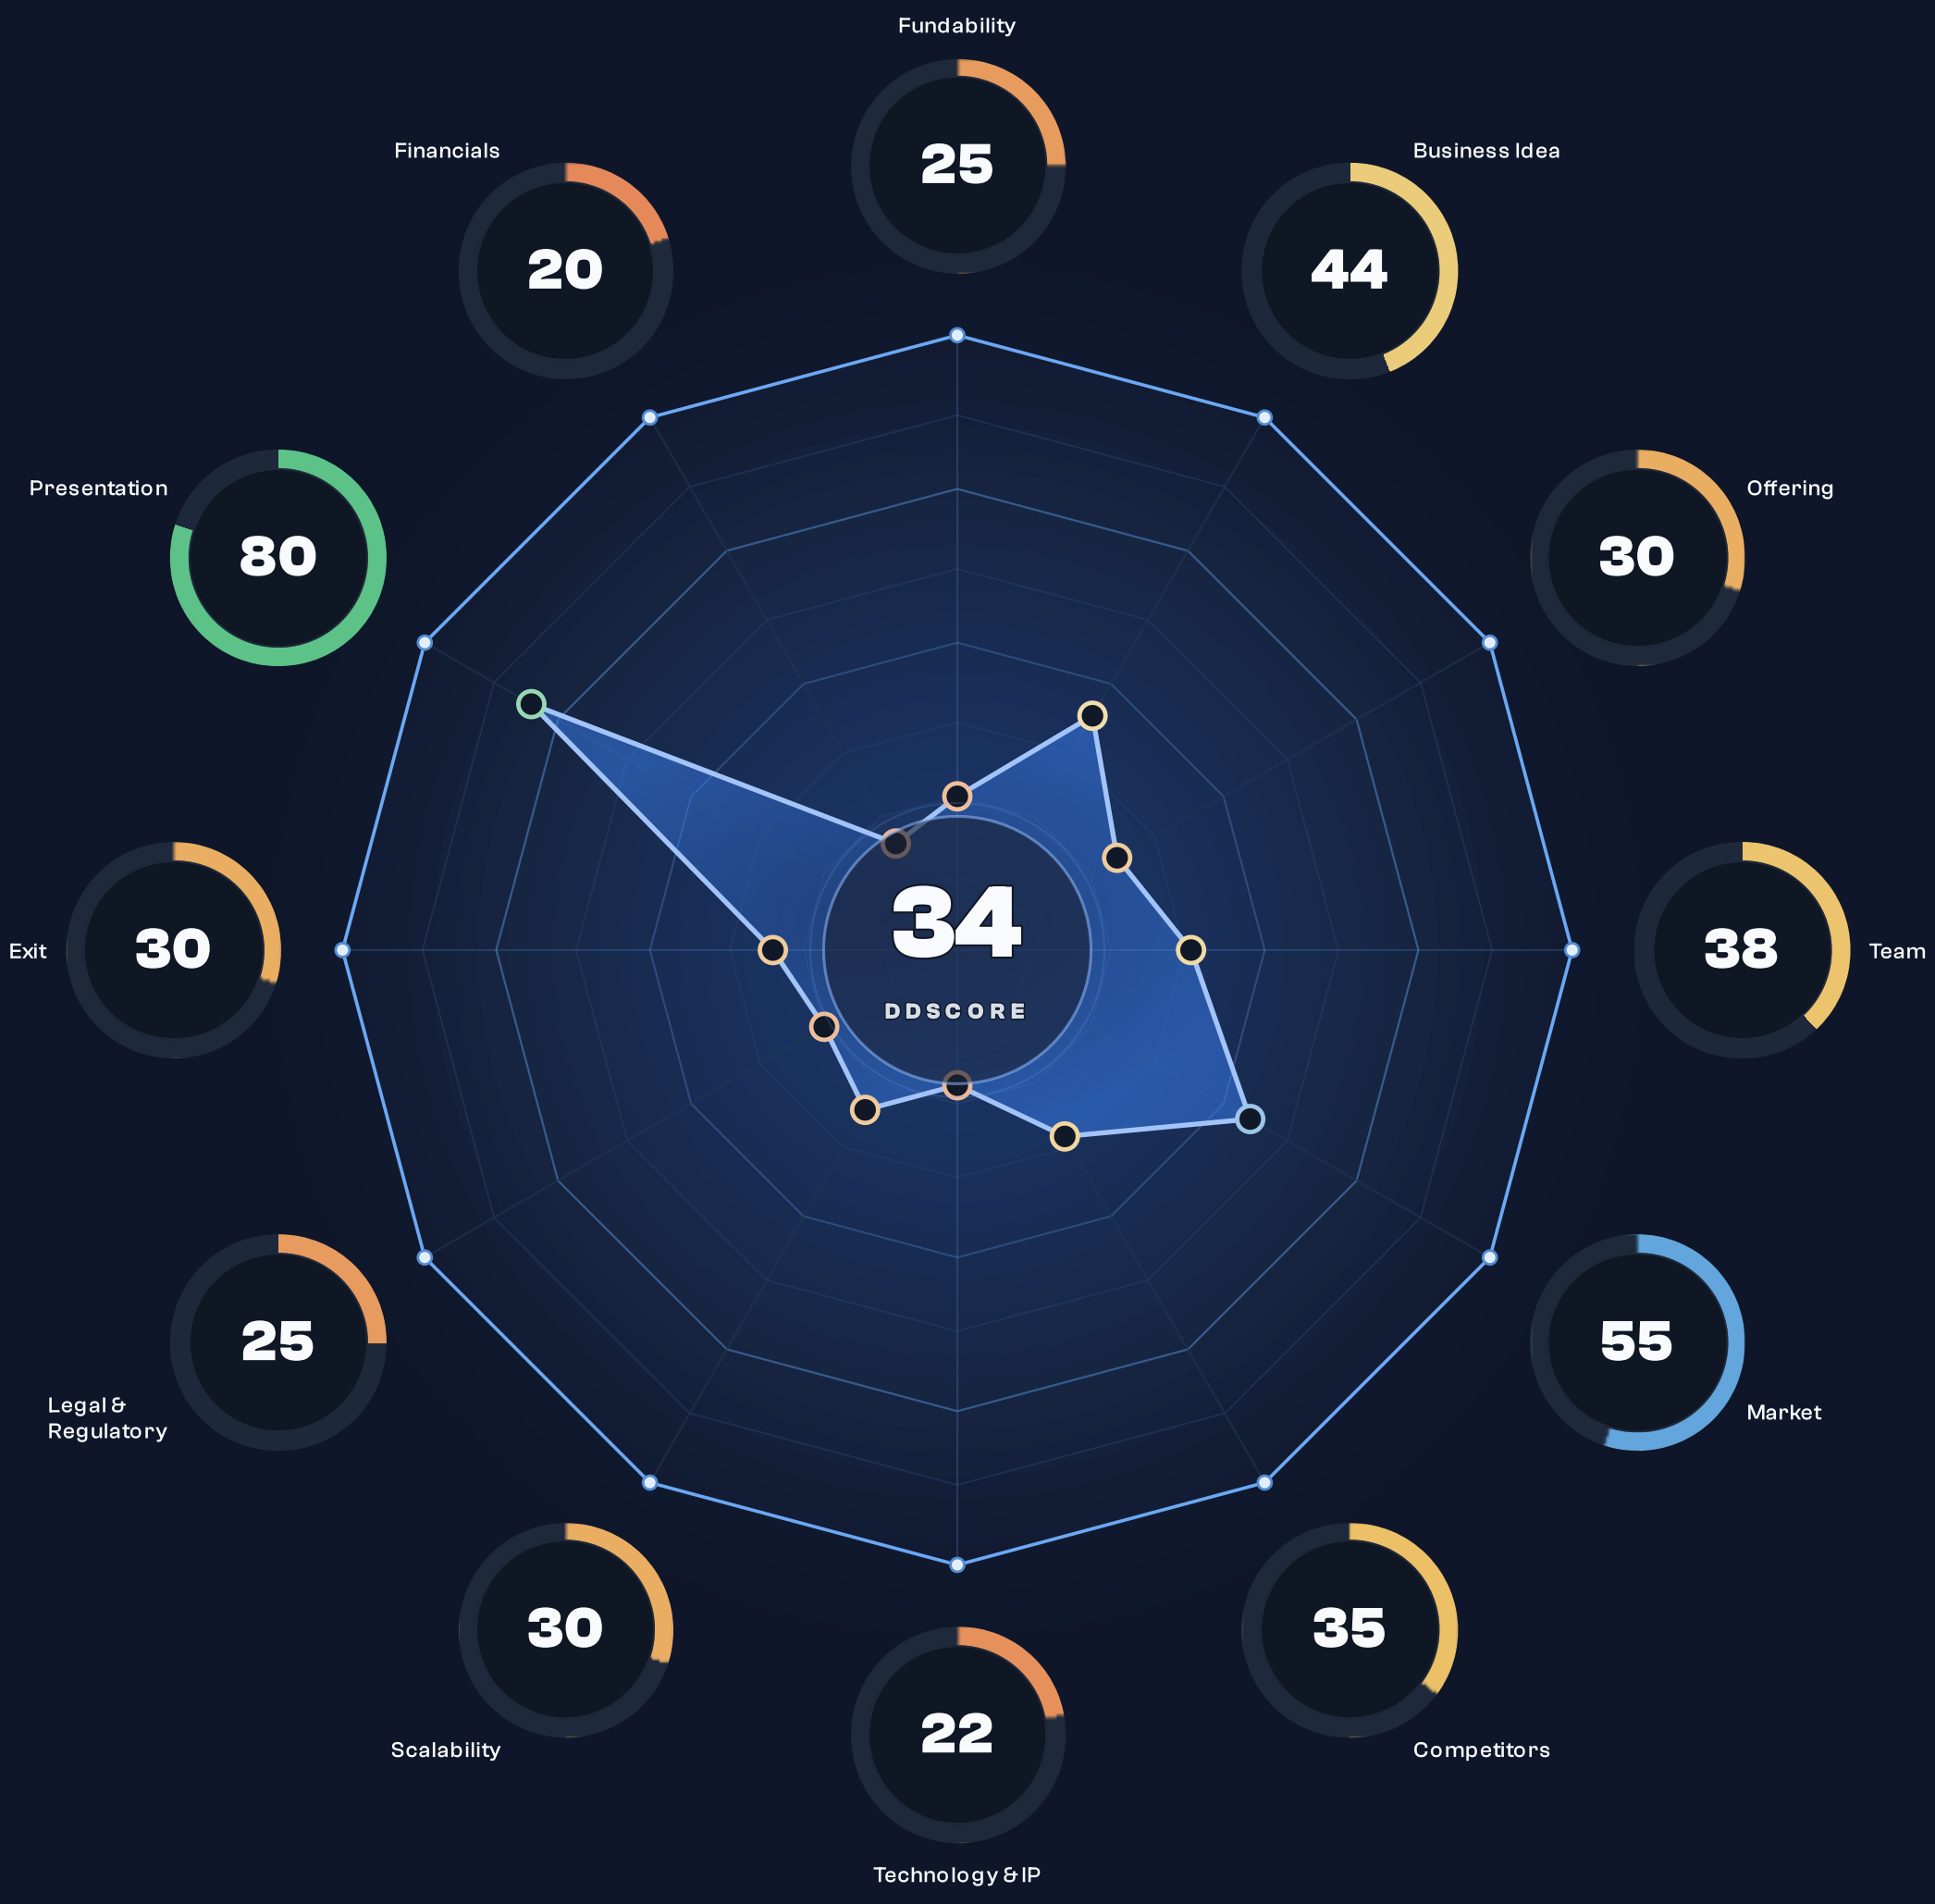
RND PlayfulPixels — Pasi Puikkonen

Investment analysis for ExampleForge Marketing AI Ltd

DDScore Report — October 5, 2026

Score Overview

DD SCORE.AI



Make Better Decisions

0422 2026-10-05

Scale (0-100%)

Analyzed by DDScore.ai — not investment advice

First time reading a DDScore report? Read first the interpretation guide (https://www.ddscore.ai/report-reading-guide): this report is based on a statistical probability model with company-specific weights for every dimension, built with outside sources.

Executive Summary

ExampleForge Marketing AI Ltd, a Seed-stage German AI recruitment-marketing prototype raising EUR 8.75M at EUR 35M pre-money, sits in deep Elevated Risk. The pain is genuinely documented: 36 discovery interviews, 6 to 10 hours of manual setup per recruitment campaign, and a EUR 420M German wedge budget pool sized from IAB Europe and OVK/BVDW data. The case for caution is stronger: zero paying customers, EUR 312k of one-off consulting sales, and an OpenAI serving cost of EUR 0.15 per impression against a EUR 0.012 to 0.019 price, so every served impression destroys value. Alexandra Hartmann (CEO) is the one standout, with EUR 52M of prior fundraising exposure across FunnelNova GmbH, RetailPilot AG and Northbridge Ventures, but she is commercially only, while Oskar Lehto, Mia Rosen and Daniel Weber are junior and no CTO, ML, security or DPO owner exists. The data room's cash bridge only balances after correcting a sign error that overstates ending cash by EUR 9,050k. What would most change the view: paid pilots at corrected pricing with positive contribution margin, a signed full-time technical lead, completed IP assignment, and a rewritten agreement at a valuation tied to proof rather than an AI label.



Strengths

MODERATE The buyer pain is documented rather than asserted: 36 interviews, 31 citing variant creation, a measured 6 to 10 hour setup time per campaign, and 694,000 German vacancies supporting recurring workload.

SIGNIFICANT Alexandra Hartmann (CEO) brings concrete fundraising scope, having helped close EUR 31M across two rounds at FunnelNova GmbH plus diligence work at RetailPilot AG and Northbridge Ventures.

MODERATE The data room is unusually complete and self-critical, disclosing its own margin error, cash sign error, moat gaps and SHA defects instead of hiding them.

MODERATE Use of funds is allocated by workstream, owner, gate and month across six lines totalling EUR 8.75M, with a contingency released only on paid-customer milestones.

MODERATE Market sizing is sourced to named public reports (IAB Europe, OVK/BVDW, Federal Employment Agency) and correctly narrowed from a EUR 118.9B TAM to a EUR 420M German wedge.

Areas for Development

SIGNIFICANT Unit economics are negative at every published price point: EUR 0.15 token cost per impression against EUR 0.012 to 0.019 of revenue, so growth destroys cash.

SIGNIFICANT No technical ownership exists: no CTO, ML, security or DPO hire is identified, and prototype IP assignment remains conditional on paying an external agency invoice.

SIGNIFICANT Commercial proof is absent: zero paying customers, zero paid pilots after 36 interviews, and pipeline LOIs that are non-binding, related-party or compliance-risk.

SIGNIFICANT German recruitment targeting instructions and privacy controls are not deployment-ready, with no impact assessment, no signed processing addendum and shared admin credentials.

Risks

SIGNIFICANT The corrected cash bridge shows the round plus token prepayment does not cover modelled Q3 and Q4 2026 burn, while the investor-facing figure overstates ending cash by EUR 9,050k through a sign error.

SIGNIFICANT Google, Meta and Microsoft can absorb the same recruitment copy and targeting assistance natively into products the buyer already pays for, compressing the standalone workflow layer.

SIGNIFICANT Discriminatory recruitment targeting instructions and incomplete GDPR controls could block pilots or expose the company and its customers to claims under German law.

SIGNIFICANT Under the company's own 16 and 24 month slip cases the round runs out before paid-customer proof, forcing a raise at a lower valuation while the product is still in development.

SIGNIFICANT Core prototype IP is not owned outright at fundraising, and the only engineering capability is an external agency capped at two demos per month.

ExampleForge Marketing AI Ltd sits in deep Elevated Risk, and the position reflects material risks against a return outlook that is currently unattractive rather than a judgement on the underlying idea. It is pulled down by negative contribution margin at every price point, zero paying customers against an EUR 8.75M ask at EUR 35M pre-money, no technical owner, unresolved IP assignment, an investor-hostile shareholders' agreement that conflicts with the deck on round size and shareholder identity, and German recruitment targeting instructions that are not deployment-ready. The central tension is that the company's own documents state these problems clearly, which improves diligence quality

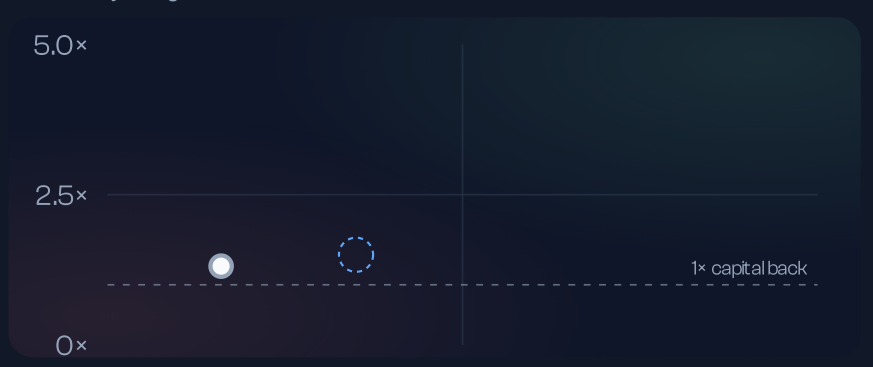
but does not improve the investment case. Revisiting the view would require paid pilots at corrected pricing with positive contribution margin and retention evidence, a signed full-time technical lead, completed IP assignment, a rewritten shareholders' agreement reviewed by external counsel, and a valuation anchored to proof. With those in place, the documented wedge and the CEO's fundraising record would support a materially different conversation.

Funding Stage Benchmark

Stronger than **29** of 100 Seed-stage companies



Risk & return

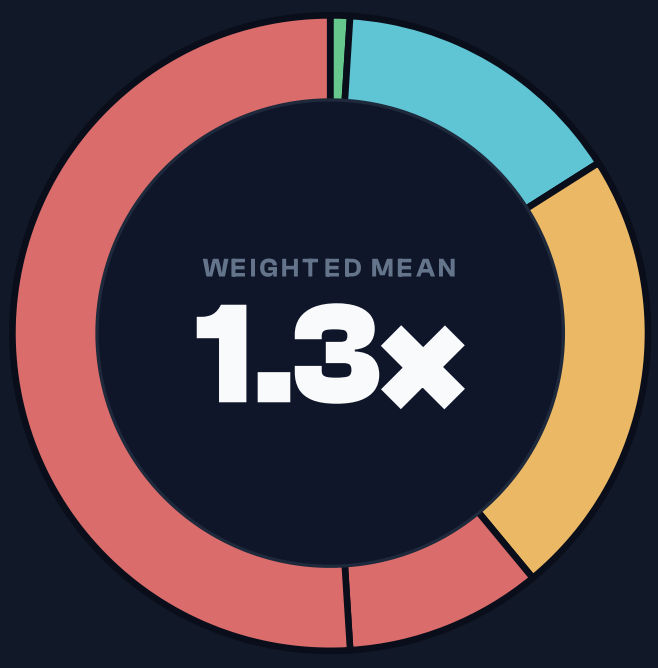


EXAMPLEFORGE MARKETING AI LTD VS SEED AVERAGE

84%	1.3×	70%
LOSS PROBABILITY	WEIGHTED RETURN	RELIABILITY
Average 65%	Average 1.5×	

Research-calibrated model, not an observed average. It shows where a typical Seed company sits over a five-year holding period. This company's exit figures use a ten-year horizon.

Exit Scenarios



MIDDLE 50% RANGE

0–0.7×

Built on thousands of company analyses: the direction holds, but small input errors can move the numbers. A model of possible outcomes, not a promise of future returns.



EXIT SCENARIO	LIKELIHOOD	RETURN RANGE
Breakout recruitment-workflow platform acquisition	1%	11–53×
Modest trade sale after corrected economics	15%	1.1–11×
Survives but the stake stays illiquid and impaired	23%	0.4–0.8×
Distressed sale of prototype and IP below entry	10%	0–0.3×
Operations cease before paid traction	51%	0–0.5×

Read more on exit options

Analysis confidence

Confidence here is driven by unusually detailed and internally cross-checkable data-room documents rather than by external corroboration. Sections built on printed figures (unit economics, cash bridge, pricing, SHA clauses, use of funds) carry high confidence because the numbers are explicit and arithmetically verifiable. Market sizing confidence is moderate: the top-line sources are credible public reports, but the EUR 420M SAM and EUR 36M Year-5 ARR are management assumptions that are not externally audited.

Business Idea

74% analysis confidence



The problem is real and specific. The market supplement records 36 discovery conversations in which 31 respondents cited insufficient ad variants, 27 cited cross-platform setup complexity, and 22 cited legally safe recruitment copy, with a median manual setup time of 6 to 10 hours per campaign. That is a repeated, budgeted workflow in a regulated niche, not a manufactured pain, and 694,000 registered German vacancies in 2024 give the workload a durable base. The wedge is also correctly narrowed from a generic AI-marketing claim to German recruitment campaign operations, which is the strongest strategic move in the material. Uniqueness, however, does not survive comparison. The same job is already served by Google Performance Max, Meta Advantage+ and LinkedIn Campaign Manager, which own the distribution and the performance data, and by ChatGPT, Gemini and Claude for the copy step itself. The company's own defensibility memo concedes a well-funded competitor could clone the current demo scope in six months. Timing is mixed: AI marketing spend is forecast to grow at roughly 25% CAGR, but that same growth funds better-funded rivals and platform-native absorption. Strategy direction is stated more clearly than most early decks (narrow ICP, milestone gates, fail gate at month 9 to 12), yet it rests on hires that are not identified and on pricing that loses money on every impression, so credibility against resources and sequencing is weak.

Strengths

MODERATE Pain is evidenced by 36 interviews with concrete frequency data (31 of 36 cite variant creation, 27 cite cross-platform setup) and a measured 6 to 10 hour manual setup time per campaign.

MODERATE The thesis was correctly narrowed from generic AI marketing to German recruitment campaign operations, giving a specific buyer, budget line and regulated workflow.

MODERATE Compliance-aware recruitment copy could become a genuine wedge if the product is redesigned with auditable prompts and documented human review, as the market supplement itself states.

Areas for Development

SIGNIFICANT Replace the generic workflow with one defensible step the incumbents do not own, most plausibly an auditable German AGG and GDPR compliance pre-check with a documented review trail, and prove it on a single paid design partner before raising.

SIGNIFICANT Reprice or re-architect so contribution margin is positive before any growth spend; the current idea only works if token cost per impression falls far below the EUR 0.15 now assumed.

MODERATE State the how-to-win sequence against named platforms explicitly: which buyer, which budget line, which channel, and why that buyer adds a paid tool instead of extending Google, Meta or agency spend.

Risks

SIGNIFICANT Platform absorption risk is existential: Google Performance Max, Meta Advantage+ and LinkedIn Campaign Manager can add recruitment copy generation and targeting assistance natively at no incremental price to the buyer.

SIGNIFICANT The workflow is a generic prompt-to-copy pattern that ChatGPT, Gemini or Claude can reproduce for a fraction of the subscription price, leaving no structural reason to pay EUR 499 to 6,500 per month.

MODERATE Budget competition is unresolved: the deck itself notes many target customers already spend through Google, Meta, LinkedIn or agencies and may not add another paid tool.

Offering

84% analysis confidence



Naming is consistent across sources: the deck, profile and data room all present ExampleForge Marketing AI Ltd and describe the same five-step workflow (brief, generate, pre-check, export, learn), with the domain exampleforge.ai reserved but parked. Maturity is prototype. The demo UI is an investor-only web app in agency cloud with no production hosting, monitoring or access-control review; ad serving is simulated from CSV; export to Google, Meta and LinkedIn is planned, not live; the learn step has no verified performance loop; docs, signup and support do not exist. Commercial viability is the weak point. Four tiers are published from EUR 499 to 6,500 per month, but no buyer has paid, the Enterprise Beta tier allows a 60-day free pilot to be booked as ARR, and the first German recruitment customer is offered a 100% discount for two months with a success-fee option, so any early revenue figure will mix free and discounted usage with list price. Willingness to pay is unvalidated: 14 of 36 interviewees were budget-aware and 9 asked for pilot access, but 0 paid. Integration risk is high because the product depends on ad-platform APIs it does not yet call and on an OpenAI key still held by the agency. The roadmap is milestone-gated and honest, but months 0 to 8 require a CTO who is not hired and an IP handover that is not paid.

Strengths

MODERATE The five-step workflow is easy to understand and maps directly to the pain measured in interviews, so offering clarity is high even though maturity is low.

MODERATE The roadmap is milestone-gated with an explicit fail gate at months 9 to 12 if paid pilots do not convert, which is better discipline than most pre-revenue decks show.

LOW Pricing tiers are published with included volumes and overage rates, so the packaging logic is legible and testable.

Areas for Development

SIGNIFICANT Ship one production integration, most plausibly CSV plus a single ad platform, with monitoring and access control, before adding tiers or features.

SIGNIFICANT Separate free pilots, discounts and success fees from ARR in every reported figure, and stop allowing a 60-day free pilot to be booked as ARR.

MODERATE Publish a live pricing page with terms and a signup path; the domain is parked and pricing exists only inside the deck.

Risks

SIGNIFICANT The export and learn steps, which are the parts that would create switching cost, are not built, so the current demo does not yet deliver the promised outcome.

SIGNIFICANT No customer has run a live campaign, so output quality, compliance flagging accuracy and performance lift are all untested against real platform conditions.

MODERATE The FAQ draft claims ExampleForge makes legacy ad managers unnecessary, a competitive claim with no channel control or evidence behind it that could damage credibility with buyers.

Team

76% analysis confidence



Alexandra Hartmann (CEO) supported Series A and B preparation at FunnelNova GmbH and helped close EUR 31M across two rounds, built board materials and acquisition models at RetailPilot AG, and sourced enterprise SaaS investments at Northbridge Ventures, giving EUR 52M of total fundraising exposure. That is genuine hiring-bar evidence for investor communication and finance discipline, and it is the reason the data room reads as well as it does. The rest of the team does not yet clear the bar for the plan. Oskar Lehto built the demo with no-code tooling and an external agency, interrupted an MSc, and has a three-month product internship; he has never shipped production software or managed engineers, yet the resourcing plan loads him with product rebuild, AI workflow, demo reliability, sales engineering, enterprise demos and pilot onboarding simultaneously. Mia Rosen became marketing director after six months as an assistant with no owned budget and no B2B SaaS demand-generation record. Daniel Weber has business-law coursework but no legal qualification, no DPO experience, and owns the German recruitment targeting design that carries the largest legal exposure. Alexandra Hartmann is simultaneously fundraising, enterprise sales, board materials, hiring and budget control. There is no CTO, ML lead, security owner, DPO or dedicated sales lead, and the board's own Gate 1 makes the round conditional on a CTO who is not yet identified.

Heavily skewed to commercial and administrative capability with almost no technical depth. One strong commercial profile (Alexandra Hartmann) is offset by three junior profiles and a complete absence of engineering, security and data-protection ownership, which is the wrong balance for an AI product in a regulated employment domain. Scored low.

AH

KEYPERSON

CEO

Alexandra H.

DECK-STATED BACKGROUND: Commercial. Public personal web research was not performed because consent was not granted; assessment is limited to uploaded materials and company-provided profile information.

Carries fundraising, enterprise sales, board materials, hiring and budget control at once, so the investor process competes directly with commercial proof. No technical background to challenge architecture or token-cost assumptions. Prior roles are in the uploaded CV pack only and were not independently verified.

OL

FOUNDER

Product lead / co-founder

Oskar L.

DECK-STATED BACKGROUND. Public personal web research was not performed because consent was not granted; assessment is limited to uploaded materials and company-provided profile information.

No shipped production software at scale and no engineering management record, yet he is slated to own the production rebuild. Simultaneously the product and sales bottleneck. The shareholders' agreement also assigns him product delivery, technical supplier management, enterprise sales, fundraising and investor reporting for twelve months. Prior roles are in the uploaded CV pack only and were not independently verified.

MR

KEYPERSON

Marketing lead

Mia R.

DECK-STATED BACKGROUND: Marketing assistant at BrightLoop Agency in 2025 for six months, preparing reports, resizing creatives and joining client calls. Freelance social media support in 2024 for two small local accounts with monthly spend below EUR 1k. BA Communications 2021-2024 with student projects in brand storytelling. Appointed marketing director after the six-month assistant role. Public personal web research was not performed because consent was not granted; assessment is limited to uploaded materials and company-provided profile information.

DW

KEYPERSON

Operations / Compliance

Daniel W.

DECK-STATED BACKGROUND: Operations at ExampleForge 2025-2026, preparing the target-customer list and recruitment segmentation rules. Business development volunteer in 2024, introducing two staffing agencies with no signed pilots. University studies 2020-2024 with a business law minor, described as introductory coursework only with no professional legal qualification. Germany-based, responsible for employer partnerships, vendor paperwork and the recruitment-marketing launch. Public personal web research was not performed because consent was not granted; assessment is limited to uploaded materials and company-provided profile information.

Owns the age, nationality, family-status and health-adjacent targeting design that creates the largest legal exposure, with no legal qualification. Compliance work likely needs external counsel, adding cost and delay not fully reflected in the plan. Prior roles are in the uploaded CV pack only and were not independently verified.

Capability Gaps

No CTO, ML engineer, infrastructure owner, security lead or DPO, despite an AI product handling candidate personal data. No finance function, which is why the shareholders' agreement defers all reporting and the cash model omits taxes and VAT. No qualified privacy counsel engaged; the board notes no signed outside-counsel engagement letter exists.

Strengths

SIGNIFICANT Alexandra Hartmann (CEO) has concrete, dated fundraising scope: EUR 31M closed across two rounds at FunnelNova GmbH plus board and diligence work at RetailPilot AG and Northbridge Ventures.

MODERATE The team is small, four people, and the plan is honestly gated on a technical hire rather than pretending current capability is sufficient.

LOW Daniel Weber is Germany-based with local recruiter access, which matters for a market where language and works-council context gate every pilot.

Areas for Development

SIGNIFICANT Hire a full-time CTO or technical lead with production SaaS and ad-tech API experience before any product-rebuild budget is released, as the board's own Gate 1 requires.

SIGNIFICANT Add a dedicated sales lead so Oskar Lehto stops owning product delivery and enterprise demos in parallel; the resourcing document names this as the core execution risk.

SIGNIFICANT Bring in qualified privacy counsel or a DPO before any pilot, since Daniel Weber has no legal qualification and owns the highest-risk workstream.

Risks

SIGNIFICANT Oskar Lehto is a single point of failure across product rebuild, sales engineering and enterprise demos; the sensitivity plan shows the 12-month plan becoming 16 to 24 months if he remains the bottleneck.

SIGNIFICANT The critical CTO, ML, security, DPO and sales hires are unnamed and uncommitted, yet EUR 1.65M of the round and the entire product gate depend on them.

Market

72% analysis confidence



The headline numbers are sourced rather than invented: IAB Europe puts European digital advertising at EUR 118.9B in 2024 growing 16% year on year, OVK/BVDW put German online display above EUR 6B, Grand View forecasts an USD 82.23B AI-marketing market by 2030 at 25% CAGR, and the Federal Employment Agency recorded 694,000 average registered German vacancies in 2024. The company then correctly telescopes: the serviceable pool is management's estimate of about 7,000 reachable German agencies and employers at roughly EUR 60k of annual campaign-operations budget, or EUR 420M, and the Year-5 target is 1,800 accounts at EUR 20k ARR for EUR 36M. That is a disciplined narrowing and the SAM assumption is labelled as unaudited. The realistic capture picture is far weaker than the target implies. Reaching EUR 36M ARR means acquiring 1,800 accounts in five years from zero, roughly 30 net new paying accounts per month every month, at a paid CAC the deck itself benchmarks at EUR 1,800 to 3,500, which is EUR 3.2M to 6.3M of acquisition spend alone against a EUR 1.8M GTM budget. A base case of 150 to 250 accounts at EUR 12k to 18k ARR (EUR 2M to 4M) is more defensible; an upside of 600 to 900 accounts (EUR 10M to 15M) requires channel partners that are only hypotheses. Demand is cyclical: recruitment ad spend tracks hiring volume, and German vacancy counts fall quickly in a downturn. The market is not winner-take-all; niche compliance tooling can coexist, but only if the compliance layer is actually built.

Market Dynamics

growing but contested by platform-native features

Problem Significance

Painkiller for a narrow job: repeated, time-consuming and compliance-sensitive recruitment campaign setup is a real weekly cost for the target buyer, but the buyer already has working alternatives and has not yet paid to remove it.

Strengths

MODERATE Top-line market figures are attributed to named public sources (IAB Europe AdEx 2024, OVK/BVDW 2024, Grand View Research, Federal Employment Agency) rather than asserted.

MODERATE The company narrowed from a EUR 118.9B TAM to a EUR 420M German wedge and explicitly labelled the SAM as an unaudited management estimate.

MODERATE Pain frequency data from 36 interviews supports a real recurring workload rather than a one-off need.

Areas for Development

SIGNIFICANT Replace the EUR 36M Year-5 ARR target with a bottom-up plan showing monthly account adds, tier mix, channel and CAC, because 1,800 accounts from zero implies about 30 net new paid accounts per month for 60 months.

MODERATE Add a downside demand case tied to German hiring cycles, since recruitment advertising spend contracts when vacancy counts fall and no cyclicality assumption appears in the material.

Risks

MODERATE The EUR 420M SAM rests on an unaudited assumption of 7,000 reachable buyers at EUR 60k each, with no evidence that this budget is discretionary or available to a new tool.

SIGNIFICANT Zero paid pilots after 36 interviews and 9 pilot requests indicates the gap between acknowledged pain and willingness to pay has not closed.

MODERATE Demand is cyclical with German hiring volume, and the 694k vacancy figure is a peak-cycle input that can reverse quickly.

Competitors

64% analysis confidence



On that lens the tightest threats are the platforms the buyer already logs into. Google Performance Max and Meta Advantage+ compete for the same campaign-operations budget with native copy generation, targeting automation and the performance data the subject company does not have. Generic assistants (ChatGPT, Gemini, Claude) undercut the copy step at near-zero marginal cost for the same buyer. Closer-priced martech peers include Copy.ai and AdCreative.ai, which sell SMB AI creative subscriptions in a comparable band, while Jasper, Smartly.io and Persado are real but sell to larger marketing organisations with different procurement motions, so they are budget competitors rather than direct ICP overlap. Agencies and freelancers remain the trusted human substitute. The deck names most of these honestly and does not use a flattering self-serving quadrant, which is unusual and creditable. Web research returned no diligence-grade funding, scale or pricing data for any competitor in this batch, so scale comparisons rest on the deck's own characterisations rather than independent figures.

Service	Owner	Differentiator
LinkedIn Campaign Manager	Microsoft	Same ICP and same budget line, with owned first-party professional data, existing procurement relationships and enterprise support. Ahead of the subject on distribution, data and trust; the subject claims only faster variant generation and a compliance pre-check.
Google Performance Max	Alphabet	Same buyer and same budget, with native AI copy and asset generation plus the conversion data the subject must import. Far ahead on distribution and measurement; the subject has no live integration to Google at all.
Meta Advantage+	Meta Platforms	Same advertiser and same media budget, owning distribution and performance data. The subject's own pipeline account KioskMarkt24 conditions purchase on beating Meta cost-per-acquisition by 40%, which is unproven.
ChatGPT, Gemini and Claude	OpenAI / Google / Anthropic	Cheaper and immediately available for the copy step. The subject's only claimed advantage is workflow structure and compliance pre-checking, neither of which is yet built or audited.
Copy.ai and AdCreative.ai	Copy.ai / AdCreative.ai	Comparable price band and similar self-serve buyer, but horizontal rather than recruitment-specific and without German compliance logic. Direct price-band competition for the Launch tier.
Jasper	Jasper	Better-funded with an established customer base, but aimed at content teams rather than German recruitment campaign operations, so overlap is on budget rather than ICP.
Smartly.io	Smartly.io	Different segment and procurement motion; relevant as a budget and feature benchmark rather than a head-to-head rival for the subject's buyer.
Persado	Persado	Enterprise motion and pricing, so weak direct overlap; relevant only as evidence that language-generation features commoditise quickly at scale.

Service	Owner	Differentiator
Agencies and freelancers	-	More expensive per campaign but trusted and accountable; the subject must beat them on cost per hire, not on feature count, and has no referenceable performance data to do so.

Competitive Threat Map

Competitors closer to the upper-right corner compete for the same customer and the same budget.

Competition geography

Local | National · Germany | Continent | Global ·



Bubble area represents relevant annual revenue · not established. Dashed bubbles: relevant revenue is unknown or cannot be reliably estimated.

Strengths

MODERATE The deck names real substitutes including Google Performance Max, Meta Advantage+, LinkedIn Campaign Manager, generic LLMs, Jasper, Smartly.io, Persado and agencies instead of a flattering self-positioning chart.

LOW The market supplement maps each substitute to the buyer job it already serves, which is a more honest competitive frame than most Seed decks provide.

Areas for Development

MODERATE Add named German or European recruitment-marketing and HR-tech tools that sell to the same staffing agencies, with pricing and traction, because the current set is dominated by global platforms and omits vertical peers.

MODERATE Quantify the switching argument: state what a buyer gives up by staying on LinkedIn Campaign Manager or an agency, in hours saved or cost per hire, rather than asserting differentiation.

Risks

SIGNIFICANT Microsoft and Google own both the recruitment audience and the campaign interface, so they can bundle the same variant-generation and compliance assistance into products the buyer already pays for.

MODERATE Web research produced no independent competitor funding, scale or pricing evidence in this batch, so the true density of better-funded vertical rivals cannot be ruled out.

MODERATE Human agencies and freelancers already deliver compliant campaign operations with trust and accountability, which is exactly the attribute a regulated recruitment buyer values most.

Technology & IP

86% analysis confidence



The architecture is stated plainly and it is a wrapper. The flow is user input to UI to backend orchestration to the OpenAI API to generated copy and recommendations, with prompt templates held in front-end configuration. There is no proprietary model, no fine-tuning evidence, no patent, no proprietary dataset and no independent benchmark. The defensibility memo answers its own moat test: a well-funded competitor could probably clone the current demo scope in six months, and a model provider could add campaign prompt generation and variants natively because those are generic large-language-model workflows. Ad-serving is simulated from CSV files, the data store is a single SQLite file exported weekly with no backup or retention design, tests are a manual QA checklist, and security review is deferred until the production rebuild. Vendor dependency is direct and priced against the company: token cost of EUR 0.15 per impression-generation exceeds the EUR 0.012 to 0.019 charged, so the provider controls both the margin and the rate limits. The claimed closed-loop learning step has no live performance data behind it. Nothing here constitutes a defensible technology position today.

Strengths

MODERATE The company documents its own wrapper risk, moat gaps and vendor dependency honestly in a dedicated defensibility memo rather than claiming AI as IP.

MODERATE A credible moat path is named: exclusive conversion data, embedded workflow, verified lift against Meta and Google, and a hard-to-replace compliance workflow.

Areas for Development

SIGNIFICANT Complete IP assignment and pay out the agency invoice before any production code is reused, and document commercial rights for all stock assets in the demo.

SIGNIFICANT Build one measurable proprietary asset, ideally an audited compliance pre-check with a documented review trail plus a conversion dataset no competitor holds, and publish an independent benchmark.

MODERATE Move prompt management out of front-end configuration into a versioned, auditable store with an access-control and logging design before any customer data is processed.

Risks

SIGNIFICANT Model-provider risk is two-sided: OpenAI can raise token prices or rate-limit the service, and can ship the same campaign-generation feature natively, compressing the business to zero margin.

SIGNIFICANT Core prototype IP is not owned outright at fundraising; assignment depends on paying the final invoice and success bonus to an external agency.

SIGNIFICANT No automated tests, no load tests, no backup or migration design and a shared admin password mean the prototype cannot safely process candidate personal data.

Scalability

78% analysis confidence



The intended model is product-led SaaS with a later channel-partner motion, but the economics scale in the wrong direction. Token cost is charged per impression-generation at EUR 0.15 against EUR 0.012 to 0.019 of revenue, so infrastructure cost grows linearly with usage while revenue per unit is an order of magnitude smaller; every incremental customer increases cash loss. That is the opposite of the sub-linear cost curve a software business needs. People are the second constraint. The resourcing document shows Oskar Lehto at 0.8 FTE product plus 0.5 FTE sales support, Alexandra Hartmann at 0.7 FTE fundraising plus 0.4 FTE sales plus 0.2 FTE hiring, Mia Rosen at 0.6 FTE marketing and Daniel Weber at 0.5 FTE operations plus 0.3 FTE compliance, which totals roughly 4.0 FTE of stated effort across four people already carrying more

than full-time loads. Product and sales cannot progress independently when one person owns both, so schedule slips compound: the plan's own sensitivity moves first pilots from month 12 to month 16 in the moderate case and month 24 in the severe case. Compliance review is a hidden human-in-the-loop cost: recruitment copy in Germany needs documented human oversight, and no qualified reviewer exists on the team. Geographic expansion beyond Germany multiplies legal review without any evidence the workflow transfers. The agency currently provides the only engineering capability, capped at two investor demos per month.

Scale Model

Intended product-led SaaS with a later recruiting-advisor reseller channel; currently a manual prototype supported by an external agency.

Outsourcing Potential

Prompt iteration, QA and infrastructure could stay outsourced in the short term, and compliance review could be provided by external German counsel, but engineering cannot remain agency-dependent once customers process personal data.

Scale Barriers

Negative per-unit contribution margin, single-person product and sales bottleneck, absence of any internal engineering or security capability, and country-specific legal review that must be repeated for each new market.

LIMITED SCALABILITY

Strengths

MODERATE The intended delivery is software rather than services, so once contribution margin is positive the marginal cost of an additional account is genuinely low.

MODERATE A channel-partner motion through recruiting advisors is identified as a phase-three route, which could decouple growth from direct headcount if retention gates are met.

Areas for Development

SIGNIFICANT Re-architect so token cost per served impression falls below the price charged, for example by capping generations, caching outputs or moving personalisation behind a platform fee, before scaling usage.

SIGNIFICANT Separate product and sales ownership with a dedicated sales hire so schedule slippage in one workstream no longer stalls the other.

MODERATE Budget for qualified compliance review as a delivery cost, since German recruitment copy needs documented human oversight that no current team member can provide.

Risks

SIGNIFICANT Cost scales linearly with usage while revenue per unit is below cost, so growth destroys cash rather than creating operating leverage.

SIGNIFICANT One person owns product delivery, sales engineering and enterprise demos, which the company's own sensitivity shows can push first pilots from month 12 to month 24.

SIGNIFICANT The only engineering capability is an external agency limited to two investor demos per month, with no internal team to absorb growth.

Legal & Regulatory

84% analysis confidence



This business sits in one of the most sensitive areas of German advertising law. The targeting brief instructs the system to prioritise users inferred to be 22 to 35 while suppressing late-career profiles, down-rank family-pause and caregiver indicators, prefer German-language profiles while excluding relocation signals, and avoid health-related keyword clusters. Each of these maps to a protected characteristic or a proxy for one under the German General Equal Treatment Act, and the brief pushes works-council communication onto the customer without evidence that the customer will accept it. Privacy readiness is absent: no data protection impact assessment has been started, the data processing addendum is a template with blank subprocessors, prompt transcripts are sent to OpenAI with no standard contractual clauses, candidate CV snippets are stored in prototype logs with no retention schedule, and there is no consent management platform or privacy notice. The EU AI Act is marked

not applicable until scale, yet employment-related candidate handling may be a high-risk use case requiring oversight and monitoring before deployment. The compliance memo's own stance, that legal design should not slow the first pilots, is the wrong sequencing for employment advertising.

Regulatory Risks

Highest exposure is German equal-treatment law applied to recruitment advertising, where the current brief instructs proxy discrimination on age, family status, origin and health.□□□ data protection: candidate CVs and prompt transcripts are processed and retained without a completed impact assessment, signed processing addendum or subprocessor schedule. The EU AI Act may classify employment-related handling as high risk, requiring oversight, monitoring and documentation the product does not have. Security controls over personal data are inadequate. Any one of these could block a pilot or trigger customer and regulator scrutiny.

Regulatory Opportunities

A properly built, auditable compliance layer for German recruitment advertising could become a genuine differentiator and a barrier to entry, because generic platform tools will not carry German-specific legal review, works-council handling or documented human oversight. That advantage only exists after the work is done; today the compliance gap is a liability, not a moat.

Entry Barrier from Regulation

Regulation currently hinders this entrant. GDPR, the General Equal Treatment Act and the EU AI Act all raise the build cost and sequencing burden for a four-person team with no qualified privacy owner, while incumbents with existing legal and compliance functions absorb the same requirements at lower marginal cost.

Strengths

MODERATE The company has written a compliance workplan with named gates (DPIA, DPA, subprocessor map, human review policy, incident response) before customer deployment.

MODERATE The deck and compliance memo identify the German AGG and GDPR exposures explicitly rather than leaving them for an investor to discover.

MODERATE A genuinely compliance-aware recruitment workflow could become a barrier to entry once built and audited, as the market supplement notes.

Areas for Development

SIGNIFICANT Remove age, family-status, nationality-proxy and health-adjacent targeting instructions from the brief entirely and replace them with job-relevant criteria reviewed by qualified German employment counsel before any pilot.

SIGNIFICANT Complete a data protection impact assessment, a signed data processing addendum with named subprocessors including OpenAI, and a retention and deletion workflow before processing candidate data.

SIGNIFICANT Replace shared admin credentials with individual access control and stand up an incident response process before any pilot, not after funding.

MODERATE Classify the product under the EU AI Act now rather than marking it not applicable until scale, since employment-related handling may be high risk.

Risks

SIGNIFICANT Discriminatory recruitment targeting could expose the company and its customers to claims under German equal-treatment law, and one pipeline prospect has already asked for age and nationality targeting.

SIGNIFICANT Processing candidate personal data through a third-party model without a data processing addendum, standard contractual clauses or a completed impact assessment creates direct GDPR exposure.

SIGNIFICANT Shared admin credentials over a system holding personal data, with security review deferred, is a live control failure rather than a planning gap.

Exit

70% analysis confidence



Base case: the company hires a CTO, ships a production beta, converts a handful of the nine pilot requests into paid design partners at corrected pricing, and reaches 150 to 250 accounts at EUR 12k to 18k ARR, or EUR 2M to 4M, within three to four years. That requires a second or third raise because the corrected bridge shows this round does not cover the 16-month slip case, so ownership dilutes materially from the modeled 20% entry. Upside case: the compliance pre-check becomes an auditable standard that German staffing agencies adopt, channel partners open a reseller motion, and the company reaches 600 to 900 accounts (EUR 10M to 15M ARR) and is acquired by an HR-tech or marketing-platform buyer wanting a German recruitment workflow. Downside is the most likely branch: pilots do not convert, token economics stay negative, and the round is consumed before proof. Exit paths are trade sale only; there is no realistic IPO route for a single-vertical German workflow tool at this scale, and acquihire is the fallback if the team and IP are clean. Unicorn probability is very low because the wedge is a EUR 420M pool with heavy platform overlap. Timing is unfavourable: the plan needs 12 months to first paid beta and the corrected cash model does not reach that point under slippage, while Google, Meta and Microsoft are shipping comparable features now rather than later.

Unicorn potential

Unicorn probability is very low, under 1%. A hundred-million valuation outcome is also unlikely, in the low single digits, because the serviceable pool is EUR 420M with heavy platform overlap and no proprietary data asset. The most probable outcomes are a small trade sale of the workflow and IP or a wind-down before proof.

Exit Scenario

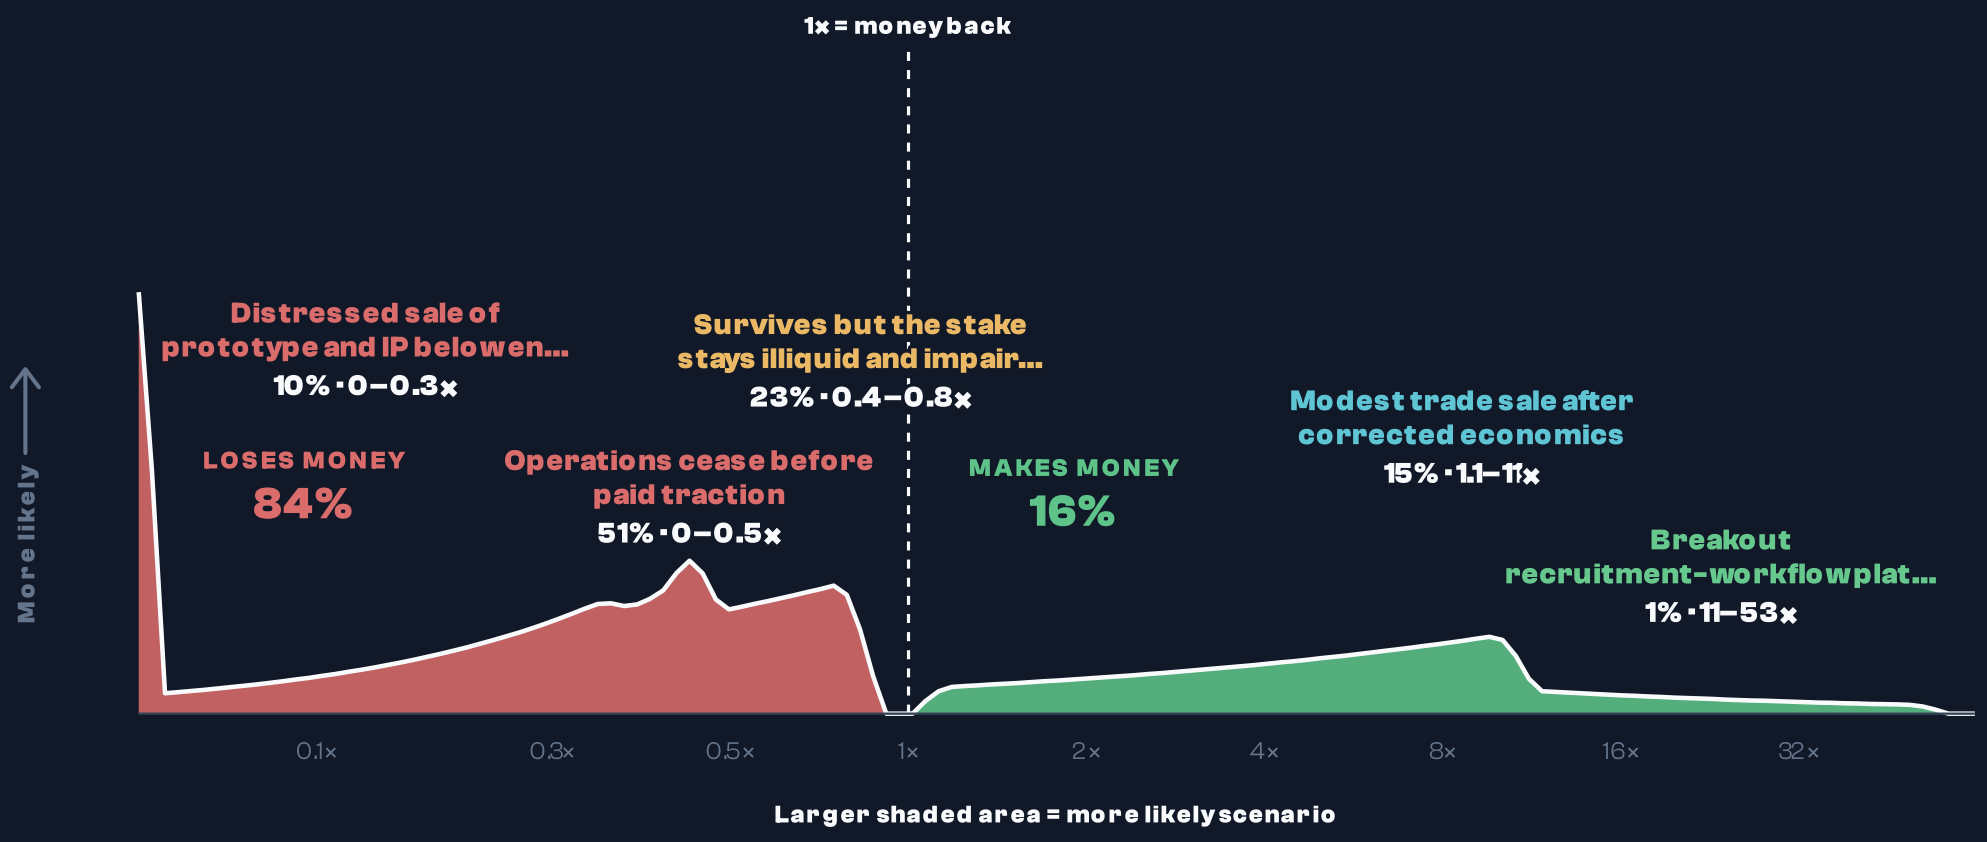
The most realistic exit is a trade sale to an HR-tech or marketing-platform acquirer that wants a German recruitment campaign workflow and an auditable compliance layer, within a modeled holding horizon of 10 years from the 2026 round on a gross investor MOIC basis. Reaching even a modest exit requires the company to first correct contribution margin, complete IP assignment, hire a technical lead and convert paid pilots with retention evidence; without those, the realistic outcome is an asset sale of the prototype and prompt library below entry.

Probabilities are calculated based on a probabilistic model.

Scenario	Probability	Return range	Description
Operations cease before paid traction	51%	0–0.5×	Modeled horizon is 10 years from the 2026 round, gross investor MOIC on a hypothetical entry of EUR 8.75M at EUR 35M pre-money, implying 20% valuation-implied ownership that is not a cap-table fact. The data room's own corrected bridge shows the round plus a EUR 1.9 M token prepayment does not cover modelled Q3 and Q4 2026 burn, no customer has paid, and contribution is negative at every published price. If pilots fail or token cost stays at EUR 0.15 per impression against EUR 0.012 to 0.019 pricing, the company stops before proof and the stake is written off. Entry terms are modeled, not confirmed company terms.
Distressed sale of prototype and IP below entry	10%	0–0.3×	Same 10-year horizon and hypothetical entry. The agency-built prototype, prompt library and targeting templates are sold in a forced or opportunistic sale once the round is exhausted, with IP assignment still conditional on paying the final invoice and success bonus. Residual value is limited to the code, the domain and any remaining team, and the buyer pays well below the EUR 4.375M post-money because no customers, retention data or clean compliance posture transfer.
Survives but the stake stays illiquid and impaired	23%	0.4–0.8×	Same horizon and entry. The company reaches a small paying base at corrected pricing but never scales: growth stays slow, platform features absorb the generic steps, and no acquirer emerges. Shares are illiquid with no buyer at the entry valuation, so the investor sees partial recovery at best. This branch carries the surviving-but-not-growing weight typical of a below-median seed profile with negative unit economics.
Modest trade sale after corrected economics	15%	1.1–11×	Same horizon and entry. Contribution margin is fixed through repricing or a token-cost reduction, 150 to 300 accounts are reached, and a mid-sized HR-tech or marketing acquirer buys the German recruitment workflow. The band is wide because outcomes between a small trade sale and a strong strategic sale are hard to separate at this stage, and because the company will realistically need at least one further round, diluting the modeled 20% entry ownership before any sale.
Breakout recruitment-workflow platform acquisition	1%	11–53×	Same horizon and entry. The compliance pre-check becomes an auditable standard, channel partners open a reseller motion, and the company reaches several hundred accounts with positive contribution, prompting acquisition by a large platform or HR-tech buyer. This is the rare tail: it requires the technical lead, clean IP, corrected economics and paid retention all to land, and multiple future rounds would dilute the modeled entry ownership substantially, so the multiple is stated net of that expected dilution.

Outcome curve

The shaded area shows the likelihood of each return range. 1.0× means getting your investment back.



The larger the shaded area within a return range, the more likely that scenario is. Built on thousands of company analyses: the direction holds, but small input errors can move the numbers. A model of possible outcomes, not a promise of future returns.

Strengths

MODERATE A trade-sale path exists in principle: HR-tech and marketing-platform acquirers buy German recruitment workflow assets, and the compliance layer would be the asset they lack.

MODERATE The market is growing and the buyer problem is documented, so a corrected product entering now would meet demand rather than create it.

Areas for Development

SIGNIFICANT Prove one paid cohort with retention before modelling any exit, because no exit multiple is credible while contribution margin is negative.

MODERATE Plan the financing sequence explicitly: state how many rounds sit between this ask and the EUR 36M ARR target and what ownership an investor retains after them.

Risks

SIGNIFICANT Platform vendors are shipping comparable recruitment and creative assistance now, so the window for a standalone workflow layer may close during the 12 to 24 month build.

SIGNIFICANT The corrected cash model runs out before paid-customer proof under the company's own slip cases, so the most likely path is a distressed or down-round financing.

SIGNIFICANT Negative contribution margin means growth destroys value, so even a successful commercial outcome does not create equity value until pricing or token cost changes.

Presentation

88% analysis confidence



This is an unusually complete and self-critical investor deck. It states the negative facts on the cover (0 paid customers, EUR 312k consulting sales), shows management's margin calculation beside the corrected version, quantifies schedule slippage at 12, 16 and 24 months, names real substitutes instead of a flattering quadrant, and maps its own shareholders' agreement defects to investor risk. For an investor it answers nearly every first-pass question in one document, which is why presentation scores far above the underlying categories. The weaknesses are accuracy and consistency rather than omission: the cash summary shown to investors contains a sign error worth EUR 9,050k, the round size conflicts with the shareholders' agreement, and the differentiation claim is asserted without data. A deck this honest about bad numbers still needs those numbers corrected before circulation.

Structure and flow

The flow follows a sensible investor sequence from problem through market, product, pricing, GTM, competition, team, funding and risk, with a decision frame at the end. Two structural choices weaken it: the material completeness checklist and decision frame are meta-commentary about the deck rather than investment content, and the corrected unit economics appear late when they should sit next to pricing. Overall the structure is close to investor-ready and needs correction more than reordering.

Investor deck completeness

STATUS	IDEAL SECTION	WHAT INVESTORS EXPECT	MATERIAL EVIDENCE	ACTION
Present	**Problem / customer pain**	Specific pain, buyer, urgency and why now	Slide 04 with 36 interviews, 6 to 10 hour setup time, pain frequency table in the market supplement	Keep; add one named buyer quote with company attribution.
Present	**Solution / product**	What is built, how it works, why it solves the pain	Slide 08 five-step workflow plus slide 09 milestone-gated roadmap; export and learn steps not live	Label clearly which steps are live versus planned on the workflow slide.
Present	**Market / ICP**	Target customer, market size, segment focus, buying context	Slides 05 and 06 with sourced TAM, EUR 42.0M SAM, primary and secondary ICP	Keep; mark the SAM as a management estimate on the slide itself.
Weak	**Business model / pricing**	Revenue model, pricing logic, margin logic, repeatability	Slide 10 tiers EUR 499 to 6,500 with the EUR 0.15 versus EUR 0.012 to 0.019 gap stated	Show the corrected contribution margin directly beside the tier table, not one slide later.
Present	**Traction / validation**	Customers, pilots, revenue, pipeline, usage, retention	Slide 07 states 0 paid pilots, 0 ARR, non-binding LOIs; EUR 312 k consulting on cover	Keep the honesty; separate the EUR 312 k one-off consulting line from any product revenue label.
Weak	**Team**	Founder roles, execution proof, role loading, hiring gaps	-	-
Weak	**Competition / differentiation**	Alternatives, substitutes, uniqueness, moat, why customers choose this	Slide 13 names platforms, generic AI, martech and agencies honestly	Replace the unproven differentiation claim with one measurable comparison against a named alternative.
Present	**Go-to-market**	Sales motion, channels, acquisition costs, partnerships, repeatable distribution	Slide 12 three phases, EUR 2 60 versus EUR 1,800 to 3,500 CAC, sales owner risk	Keep; state which channel is tested versus assumed.
Weak	**Financials / unit economics**	Forecast discipline, costs, margin, runway, key assumptions	Slide 11 corrected math, slide 17 runway sensitivity, appendix cash bridge with a sign error	Correct the EUR 7,950k ending-cash figure before circulation; it currently overstates cash by EUR 9,050k.
Present	**Funding ask / use of funds**	Amount, round context, budget split, milestone allocation, runway, delay sensitivity	Slide 16 six workstreams with owners and gates, slide 17 three runway cases	Reconcile the EUR 8.75M ask with the EUR 4M stated in the shareholders' agreement.
Present	**Risks / regulatory / legal**	Material risks, dependencies, compliance, IP, permits	Slides 18 to 20 covering AGG, GDPR, AI Act, SHA defects and a full risk register	Keep; move the compliance redesign ahead of any pilot commitment in the roadmap.

Observed structure (from materials)

#	SLIDE OR SECTION	INVESTOR THEME
1	Cover: Raise EUR 8.75M	Other / distracting
2	Executive Summary	Combined

#	SLIDE OR SECTION	INVESTOR THEME
3	Material Completeness	Other / distracting
4	Customer Pain	Problem
5	Market Context	Market
6	ICP and Buying Context	Market
7	Customer Discovery and Traction	Traction
8	Product Workflow	Solution
9	Product Status and Roadmap	Solution
10	Business Model and Pricing	Finance
11	Unit Economics	Finance
12	Go-to-Market	GTM
13	Competition and Differentiation	Competition
14	Technology and IP	Competition
15	Team and Resourcing	Team
16	Use of Funds	Finance
17	Runway and Sensitivity	Finance
18	Legal, Regulatory and Data	Risks
19	Governance and Shareholder Agreement	Risks
20	Risk Register	Risks
21	Decision Frame	Other / distracting

Weak or missing investor themes

- Independent technical benchmark
- Signed counsel engagement letter
- Customer reference call

Title and clarity suggestions

- Slide 11: retitle to 'Unit Economics: Management View versus Corrected View' so the correction is visible before a reader reaches the numbers

Identified gaps

Cash figure accuracy HIGH RISK

The investor summary reports EUR 7,950k ending cash while the corrected bridge gives EUR -1,100k, a EUR 9,050k overstatement caused by a sign error.

Round consistency HIGH RISK

The deck and board minutes state an EUR 8,750,000 raise while the shareholders' agreement states EUR 4,000,000 at the same pre-money valuation.

Shareholder identity MEDIUM RISK

Differentiation proof MEDIUM RISK

The differentiation claim of a recruitment-specific workflow and compliance pre-check is asserted with no benchmark, customer evidence or measured comparison.

Team proof depth MODERATE

Independent validation MODERATE

No signed counsel engagement letter, no independent technical benchmark and no customer reference call are attached.

Financials

85% analysis confidence



The financial case fails on its own printed numbers. Management reports 82% gross margin by classifying OpenAI tokens as R&D learning cost below the EBITDA line; the appendix shows that at 1,000,000 monthly impression-generations, EUR 22,000 of revenue less EUR 4,000 hosting leaves EUR 18,000 reported gross profit but EUR 150,000 of token cost, giving correct contribution of EUR -132,000 before any CAC or support. Pricing confirms the structural problem: tiers sell impressions at EUR 0.012 to 0.019 while a key variable cost is EUR 0.15, so scale worsens burn rather than improving margin. The 2026 cash bridge only balances after correction: EUR 8,750k opening, less EUR 3,100k Q3 burn, less EUR 4,850k Q4 burn, less EUR 1,900k token prepayment gives EUR -1,100k, while the investor summary reports EUR 7,950k because a formula adds the Q4 burn back as a positive. That is a EUR 9,050k overstatement of ending cash in the figure shown to investors. Assumptions are unsupported: EUR 260 launch CAC with no paid acquisition evidence against the deck's own EUR 1,800 to 3,500 paid benchmark, 0.5% monthly churn with no live customers, 7-day enterprise payment with no contracts, and taxes and VAT excluded entirely. The Year-5 EUR 36M ARR target needs 1,800 accounts at EUR 20k, roughly 30 net new paid accounts per month for five years, from a standing start. The plan figure of EUR 610k monthly burn plus token prepayments means any slip past 12 months exhausts the round before proof.

Amounts in EUR	2026 P	2027 P	2028 P	2029 P	2030 P
Revenue	312k	-	-	-	-
Expenses	9.9M	-	-	-	-

Cash flow and runway

FORECAST CREDIBILITY

Hockey stick

growth rests on the plan alone

What the forecast rests on: The forecast rests only on the plan: no signed customer contracts, no LOI converts to revenue, and the pipeline document states no customer has completed a paid live campaign. The investor summary figure of EUR 7,950k ending cash is produced by a formula that adds the Q4 burn back as a positive number, so the version shown to investors is arithmetically wrong by EUR 9,050k. The plan figure of EUR 610k monthly burn plus token prepayments is an assumption, not a measured burn.

No cash-flow calculation possible. The material does not give: No current cash balance as of 2026-10-05; No measured historical monthly net cash burn; No cash receipts or cash payments breakdown; No historical cash balances before the planned round; No taxes, VAT or working-capital lines in the model; No forecast beyond 2026 in the cash bridge.

Industry Benchmark

Claimed 82% gross margin versus a corrected negative contribution at all tiers, against a healthy SaaS gross margin of 70 to 80%. Assumed 0.5% monthly churn (about 5.8% annual) versus 10 to 25% annual churn typical for SMB SaaS, and there is no live customer base to support either figure. Assumed EUR 260 CAC versus the deck's own EUR 1,800 to 3,500 paid benchmark and versus typical SMB SaaS CAC of EUR 500 to 2,000. LTV cannot be computed because ARPU, retention and contract length are all unvalidated.

Unit Economics

At the printed example, 1,000,000 monthly impression-generations produce EUR 22,000 revenue, EUR 4,000 hosting and EUR 150,000 token cost, giving EUR -132,000 contribution before CAC and support, or roughly EUR -0.13 per impression served. Because overage is priced at EUR 0.012 to 0.019 while the personalisation path costs EUR 0.15, every incremental unit destroys about EUR 0.13, so contribution margin is negative rather than merely thin. CAC cannot be computed: no customers have been acquired, the EUR 260 figure has no paid acquisition evidence, and the deck's own paid benchmark of EUR 1,800 to 3,500 against a EUR 499 to 1,250 entry tier implies 1.5 to 7 months of gross revenue before any token cost, which is already uneconomic. Payback, LTV and lifetime are all unverifiable until paid pilots produce retention data. The Enterprise Beta tier's 60-day free pilot may be booked as ARR, which would inflate any revenue figure reported later.

Strengths

MODERATE The data room discloses its own margin error, cash sign error and unsupported assumptions rather than hiding them, which makes the numbers verifiable.

MODERATE Use of funds is allocated by workstream, owner and gate across six lines totalling EUR 8.75M, which is more specific than most Seed asks.

MODERATE Runway sensitivity is quantified at 12, 16 and 24 months with euro amounts, so capital sufficiency risk is visible rather than assumed away.

Areas for Development

SIGNIFICANT Rebuild pricing so contribution margin is positive at realistic usage: either cut token cost per impression well below EUR 0.012, cap included generations, or move to a platform fee plus metered pass-through with a documented margin.

SIGNIFICANT Correct the investor cash summary and reissue it with an independently checked bridge, since the version shown overstates ending cash by EUR 9,050k.

SIGNIFICANT Replace the EUR 260 CAC and 0.5% churn assumptions with measured data from paid pilots, and include German VAT, taxes and working capital in the model.

MODERATE Restate the Year-5 target bottom-up by tier and channel with the acquisition spend required, instead of asserting 1,800 accounts at EUR 20k ARR.

Risks

SIGNIFICANT Unit economics are negative at every published price point, so each additional customer increases cash loss until pricing or token cost changes.

SIGNIFICANT The investor-facing cash figure contains a sign error; the corrected bridge shows the round plus token prepayment does not cover Q3 and Q4 2026 burn.

SIGNIFICANT No historical financial years, no measured burn and no current cash balance exist, so runway cannot be computed from evidence and rests entirely on plan assumptions.

MODERATE Taxes and VAT are excluded from the cash model while the company is setting up in Germany, omitting likely working-capital and compliance costs.

Fundability

84% analysis confidence



Revenue status is explicitly pre-revenue: the deck cover states 0 paid customers and EUR 312k of consulting sales, the board minutes record no paying customers, no ARR, no live website and no production deployment, and the pipeline document confirms no customer has completed a paid live campaign. The EUR 312k is one-off consulting, not recurring product revenue, and should be read as roughly half the weight of ARR for round purposes. The ask is EUR 8.75M at EUR 35M pre-money, implying EUR 43.75M post-money and 20% valuation-implied ownership, which is roughly 112 times the one-off consulting sales and unquantifiable against product revenue because none exists. The shareholders' agreement describes a different round: EUR 4,000,000 seed-plus at the same EUR 35M pre-money, a material inconsistency that must be resolved before any term sheet. Use of funds is genuinely detailed across six workstreams with owners and gates, which is a real strength, but the plan depends on a CTO who is not identified and on Gate 3 requiring 10 paid pilots before scale spend. Capital sufficiency fails under the company's own sensitivity: the 16-month slip case needs EUR 9.76M operating burn plus a EUR 1.9M token reserve, and the 24-month case needs EUR 14.64M plus EUR 2.8M, both beyond this round while the product is still in development.

Investor Accessibility

Easy to read, hard to underwrite. The data room answers standard investor questions in unusual detail and labels its own weaknesses, which speeds first-pass review. Underwriting then stalls on three unresolved items: a round size that differs between documents, an unverified shareholder list, and unit economics that are negative at every published price.

Cap Table

Not clean. The shareholders' agreement shows Oskar Lehto 42%, Milena Krause 23%, Joonas Rantala 15%, Rita Hoffman 10% and a 10% unallocated growth pool that the board may change before or after completion without adjusting the pre-money valuation. This is a potential discrepancy requiring verification, not a finding of fact.

Shareholder Agreement

Legal coherence defects: company name differs between cover and clauses; English law with Berlin courts plus Delaware arbitration for the investor; the agreement may be treated as binding while articles, register, option plan and IP assignments are updated later; compliance obligations deferred until after product-market fit.

Dead Shareholders

No dead equity is disclosed, but the 10% unallocated growth pool can be increased or reduced by the board without changing the pre-money valuation, which effectively transfers value between existing and new holders.

Bad/Good Leaver Provisions

Bad leaver is defined expansively and decided by the board, including disagreement with the product roadmap or conduct the board considers harmful to fundraising, with transfer at the lower of nominal value, subscription price or last price paid. No good-leaver definition exists for death, disability, illness, redundancy, termination without cause, family leave or relocation.

Red Flags

Round size conflict: shareholders' agreement states EUR 4,000,000 while the deck and board minutes state EUR 8,750,000

Full-ratchet anti-dilution with no time limit and no employee-option carve-out

5% vetoes over budget, pricing, hiring and fundraising with no stated basis and no deadlock response rule

Investor cash summary overstates ending cash by EUR 9,050k through a sign error

Valuation of EUR 35M pre-money against zero product revenue and EUR 312k of one-off consulting sales

No identified or committed CTO despite Gate 1 making it a condition of the round

Strengths

MODERATE Use of funds is allocated by workstream, owner, gate and month across six lines totalling EUR 8.75M, with a contingency released only on paid-customer milestones.

MODERATE The board set five explicit investor gates, including a full-time CTO before product spend and a clean shareholders' agreement before a large investor signs.

MODERATE The deck states the negative facts (no customers, no ARR, no live site) on its cover rather than burying them, which supports honest negotiation.

Areas for Development

SIGNIFICANT Reconcile the round size: the shareholders' agreement states EUR 4,000,000 while the deck and board minutes state EUR 8,750,000, and one document must be corrected before diligence continues.

SIGNIFICANT Reset the valuation to one anchored on proof, or stage the round against gates, because EUR 35M pre-money on zero product revenue and EUR 312k of one-off consulting sales is not supported by comparable evidence.

SIGNIFICANT Provide a resource-loaded budget showing what the round buys at 12, 16 and 24 months, including the token reserve, so investors can see where proof arrives relative to cash.

Risks

SIGNIFICANT The valuation is roughly 112 times one-off consulting sales with no product revenue, so a down round or structured terms are likely if any proof milestone slips.

SIGNIFICANT The shareholders' agreement is investor-hostile and internally inconsistent, naming a different round size, a different shareholder set and three conflicting governing-law regimes.

SIGNIFICANT Under the company's own 16 and 24 month slip cases the round runs out before paid-customer proof, forcing a raise at a lower valuation while the product is still in development.

SIGNIFICANT The corrected cash bridge shows the round plus token prepayment does not cover the modelled Q3 and Q4 2026 burn, so the stated runway is not achievable as written.

Glossary

CAC
Customer acquisition cost: total sales and marketing spend needed to win one new customer.

LTV
Lifetime value: gross profit expected from a customer over the whole relationship.

ARR / MRR
Annual or monthly recurring revenue: normalised value of contracted subscription revenue, excluding one-off services.

Churn
Rate at which customers cancel or fail to renew, usually monthly or annual.

Contribution margin
Revenue minus all costs caused by serving it, including hosting and model tokens, before sales and overhead.

Gross margin
Revenue minus cost of goods sold as a percentage; healthy software businesses sit around 70 to 80 percent.

CAC payback
Months of gross profit needed to recover the cost of acquiring a customer.

TAM / SAM / SOM
Total market, serviceable market, and the realistic share the company expects to win.

ICP
Ideal customer profile: the firmographic description of the buyer most likely to buy and succeed.

LOI
Letter of intent: a non-binding statement of interest with no payment or procurement commitment.

DPIA
Data protection impact assessment: required GDPR analysis before high-risk personal data processing starts.

DPA
Data processing addendum: contract governing how a vendor may process personal data for a customer.

EU AI Act
EU regulation classifying AI systems by risk; employment-related uses may be high risk with extra duties.

AGG
German General Equal Treatment Act prohibiting discrimination, including in recruitment advertising.

Full-ratchet anti-dilution
Protection repricing an investor's shares to the lowest later price, including option grants, heavily diluting founders.

Bad leaver
Clause forcing a departing founder to give up shares, often at nominal value, on board-set terms.

MOIC
Multiple on invested capital: exit proceeds divided by amount invested.

Token cost
Fee paid to an AI model provider per generation, scaling directly with product usage.

DDScore.ai interpretation

CURRENT SCORE: **34%** PROFILE: **UNEVEN READ**

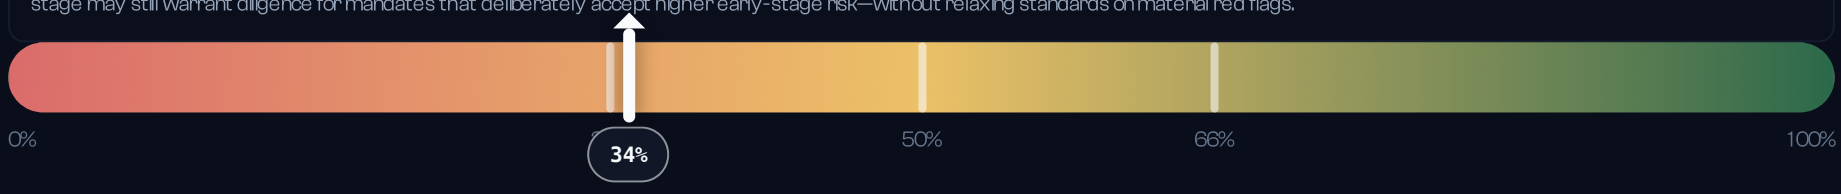
Evidence pulls in both directions. The score is not signalling a clean bull or bear case; it is signalling that the materials do not yet support a confident view.

THE LOWER BAND — HIGH RISK
Scores below about one-third of the scale signal statistically elevated risk. The materials typically show overlapping problems that dominate the picture in aggregate.

THE MIDDLE BAND — UNEVEN READ
Evidence pulls in both directions. The score is not signalling a clean bull or bear case, but that the materials do not yet support a confident view.

THE UPPER BAND — FAVOURABLE READ
The evidence cluster is more constructive across multiple dimensions simultaneously. This is not a guarantee, but it describes a different posture from the left side of the scale.

DDScore reflects investment risk and evidence density, not idea quality alone. Earlier-stage companies (for example strong pre-seed or seed profiles) therefore often read lower than later-stage peers with revenue, even when the story is solid; that is a statistical tendency, and exceptional early-stage cases can still reach strong reads. From about the mid-forties up, a modest headline driven mainly by stage may still warrant diligence for mandates that deliberately accept higher early-stage risk—without relaxing standards on material red flags.



Scores are statistically indicative inputs, not investment advice. Individual variance is high.

Based on the information available for this analysis, the opportunity carries significant risks relative to the return outlook implied by the materials.



Important disclaimer

DDScore does not provide investment advice and does not tell users what decision to make. DDScore provides analytical tooling and quantitative scoring based on submitted materials, available information, benchmarks and the DDScore scoring model. It supports judgement and due diligence workflows. It does not replace investor judgement or a full due diligence process. Investing in private companies involves significant risk, including the possible loss of all invested capital.