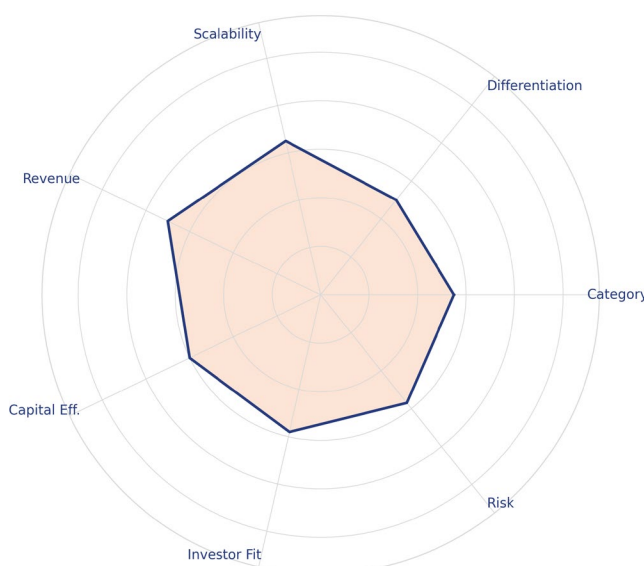


FUNDRLY

Investor Alignment Diagnostic

See where conviction breaks, what to fix first, and how to raise sharper

Company	NovaFlow AI
Sector	AI workflow automation / enterprise software
Stage	Seed to pre-Series A
Purpose	Illustrative sample of a premium Fundrly report
Prepared by	Swire Consulting Group
Date	26 March 2026



All company-specific facts, financials and illustrative materials in this report are constructed for demonstration purposes. Market statistics, investor behaviours and external observations are derived from publicly available sources and SwireIQ, the market intelligence arm of Swire Consulting Group.

What Fundrly surfaces in the first five minutes

Primary conviction break

Narrative-investor mismatch

The current story positions broad platform ambition before proving the narrow wedge investors need to underwrite.

Highest-probability investor fit

AI workflow and enterprise application funds

Most likely to respond positively to exception-handling wedge dominance and expansion logic.

First change to make

Tighten the wedge and rewrite use of funds around proof milestones

Re-anchor the raise around proof thresholds that unlock Series A decision clarity.

First 100-day conviction unlock

3 lighthouse enterprise deployments + expansion density proof

The fastest path to stronger institutional confidence is proof of repeatable wedge expansion.

Likely capital impact if corrected

Higher meeting conversion
Reduced valuation pressure
Stronger Series A narrative readiness

Potential valuation confidence uplift: low-mid 20s → low 30s range

1. How to read this report

Fundrly is designed to diagnose the gap between an underlying business and how that business is likely to be interpreted by capital. This sample report is intentionally detailed: it does not merely say that a narrative is weak, crowded or unclear; it attempts to show precisely why, where that weakness is likely to appear in an investor conversation, and what commercial consequences follow from it.

The report uses a synthetic company, NovaFlow AI, to mimic the sort of materials management teams usually have before a raise: a product deck, a short financial model, headline operating metrics, a target round size and an indicative valuation expectation. On top of those assumed materials, the report overlays external market intelligence so that each major assertion is grounded either in the assumed facts of the company or in observable market conditions.

What this report is trying to answer

If a credible early-stage AI software company enters the market today, what elements of the story will create investor conviction, which elements will trigger hesitation, and what changes are most likely to improve the quality of capital conversations?

2. Company snapshot and assumed pre-raise materials

NovaFlow AI is a London-headquartered software company founded in 2023. The business sells an AI-enabled workflow automation platform to mid-market and enterprise customers that want to automate internal approvals, recurring service requests, process orchestration and exception handling across finance, operations and customer support teams. The company is raising a \$6.0 million round to scale commercial coverage, deepen product capability and extend runway to a point where the business can show clearer evidence of category leadership.

Founded	2023
Headquarters	London, United Kingdom
Employees	68
Current ARR	\$2.4 million
Gross margin	82%
Net revenue retention	118%
Average ACV	\$22k SMB / \$68k mid-market
Monthly net burn	\$280k
Cash runway	~7 months
Target raise	\$6.0 million

Assumed materials reviewed for this sample: (i) a 15-slide investor deck, (ii) a three-statement operating model, (iii) headline SaaS KPI summary, (iv) draft use-of-funds schedule, and (v) management commentary on target customers, product roadmap and round rationale.

The synthetic deck positions NovaFlow as 'the AI operating layer for repetitive enterprise workflows'. It describes strong customer satisfaction, good logo momentum and a large market, but it does not define a single dominant category anchor. The financial model shows healthy early growth, but the bridge from current traction to category-scale revenue is not yet explicit. The intended headline round is \$6.0 million at an indicative pre-money valuation of \$32-38 million.

Best-fit capital archetypes

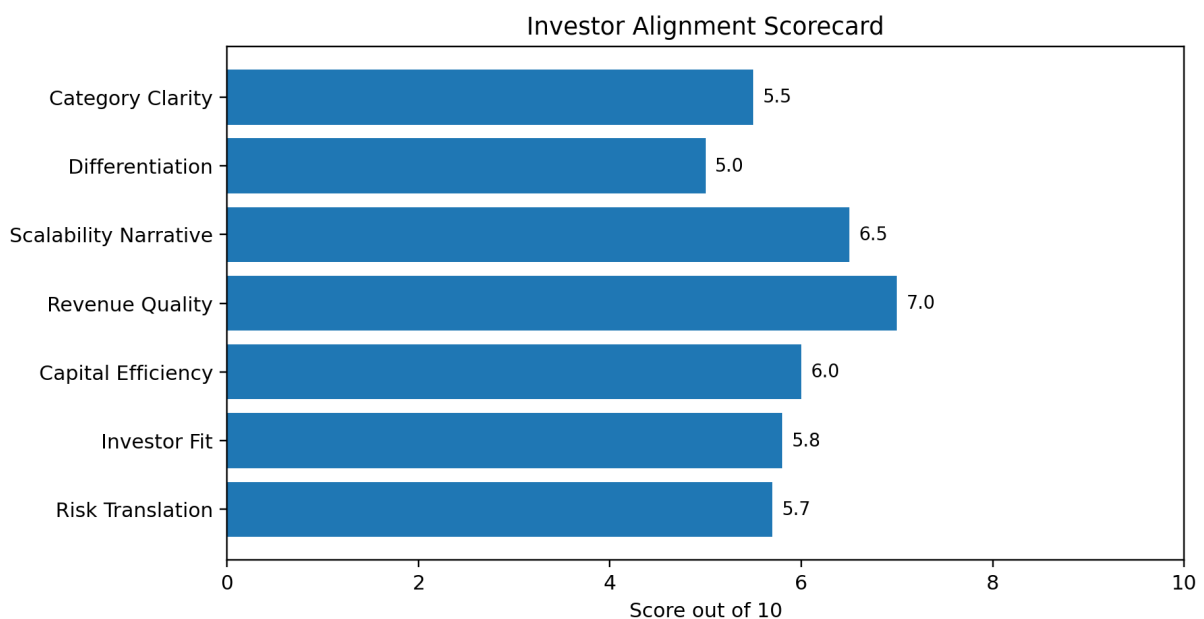
Highest probability: AI / enterprise workflow VCs
Medium probability: Process software strategics
Opportunistic: Family offices with software comfort

Immediate read from a first-pass investor screen

This looks like a credible team operating in an attractive market, but the company could easily be interpreted as another AI-enabled workflow tool unless the narrative is sharpened. In crowded AI application categories, that difference is not cosmetic; it directly affects investor conviction and valuation framing.

3. Executive summary

Overall Investor Alignment Score: 5.9 / 10



NovaFlow is investable in the sense that there is enough market pull, early revenue quality and product relevance to support a credible financing process. However, the company is materially under-optimised from an investor-positioning perspective. The current story does not yet convert business quality into a sufficiently sharp investment case.

At a practical level, this means the company is likely to achieve polite interest and some second meetings, but it will struggle to generate consistent high-conviction momentum unless it improves category clarity, makes differentiation explicit and reframes scale in investor language rather than product language.

- The market is attractive enough to create attention, but also crowded enough that generic AI automation narratives are penalised.
- Product differentiation probably exists in workflow depth, implementation speed and operational outcome, but the current positioning leaves too much of that differentiation for investors to infer.
- Commercial metrics are good enough to support an institutional conversation, but not yet strong enough to overcome narrative ambiguity on their own.
- The round is financable, but the likely quality of capital and valuation outcome will be materially worse than the company could achieve with a tighter story.

First three actions


1. Rewrite slide 1 category anchor
 2. Replace TAM slide with wedge expansion logic
 3. Reframe use of funds into 2 inflection milestones
-

Bottom-line judgement

NovaFlow should not go to market with the current story unchanged. The business has enough substance to raise, but not enough narrative precision to raise efficiently. That is the gap Fundrly is designed to address.

4. Market intelligence and investor context

The first reason NovaFlow's positioning matters so much is that the company is operating in a market that is both expanding rapidly and becoming increasingly noisy. Mordor Intelligence estimates the workflow automation market at \$23.77 billion in 2025, growing to \$40.77 billion by 2031, while Fortune Business Insights estimates the AI SaaS market at \$22.21 billion in 2025 and \$367.6 billion by 2034.[1][2] Grand View Research similarly sizes intelligent process automation at \$14.55 billion in 2024 and \$44.74 billion by 2030.[3] The directional message is consistent across sources: demand is real and expanding.



That is the good news. The more difficult reality is that attractive markets usually produce crowded narratives. In AI application software, investors are increasingly looking through surface-level claims of efficiency, automation and orchestration. Bessemer's State of AI 2025 argues that top-performing AI companies are now judged against much more aggressive growth expectations, while also warning about low switching costs and commoditised offerings in parts of the market.[4][5] In other words: plenty of investors want exposure to AI software, but they are not paying premium prices for companies that feel substitutable.

This harder filter is reinforced by buyer-side behaviour. Andreessen Horowitz's 2025 survey of 100 enterprise CIOs found that generative AI budgets were growing materially faster than expected, with enterprises expecting around 75% budget growth over the next year.[6] That is bullish for demand. But it also means founders are competing for investor capital in an environment where buyers are real, budgets are increasing, and investors therefore expect category winners - not simply useful products. Higher demand does not reduce the burden of narrative precision; it raises it.

McKinsey's Technology Trends Outlook 2025 points to a dramatic surge in AI-related equity investment, again signalling abundant capital interest in the broad theme.[7] Yet abundant thematic capital creates a paradox: more money chases the category, but the marginal company must work harder to prove why it belongs in the scarce subset of opportunities that investors believe can compound. For NovaFlow, the implication is simple. It is not enough to be in a good market; the company must demonstrate why it should emerge from a crowded set of AI workflow tools as a preferred risk-adjusted bet.

Why this matters for NovaFlow

In a quieter software category, a decent story can survive because the market does some of the work. In AI workflow automation, the market does the opposite: it attracts attention but also forces comparison. That is why category definition, differentiation and capital efficiency become disproportionately important.

5. Fundrly scoring methodology

Fundrly scores investor alignment across seven dimensions. The point is not to produce a pseudo-scientific number for its own sake. The score is a device to force specific judgements on the issues that most often determine whether a business is understood cleanly by capital.

Dimension	Weight	NovaFlow score	What the dimension tests
Category clarity	15%	5.5	Can an investor describe the company in one sentence without qualification?
Differentiation	20%	5.0	Is there a clear reason this company wins relative to similar narratives?
Scalability narrative	20%	6.5	Is the route from current traction to meaningful scale framed convincingly?
Revenue quality	10%	7.0	Do pricing, retention and gross margin support a durable software outcome?
Capital efficiency	10%	6.0	Does the company look like it can turn capital into value inflection efficiently?
Investor fit	10%	5.8	Is the story tailored to the investor type most likely to back it?
Risk translation	15%	5.7	Are actual risks being framed in a way that avoids unnecessary discounting?

A score below 6.0 does not mean a company cannot raise. It usually means the business will be misread, under-valued or subjected to more diligence friction than its fundamentals justify. NovaFlow lands precisely in that zone: likely financable, but likely to leave value on the table.

6. Detailed diagnostic: category, differentiation and scale

6.1 Category clarity

NovaFlow's assumed deck uses three phrases interchangeably: 'AI operating layer', 'workflow automation platform' and 'process orchestration engine'. Each of these phrases is defensible in isolation. The problem is that investors do not reward descriptive breadth; they reward definitional sharpness. A broad set of labels gives management flexibility, but it gives investors too many ways to place the company into the wrong bucket.

The likely investor shortcut is to collapse NovaFlow into a much simpler category: AI-enabled workflow software. Once that happens, the company will be compared against a wide set of automation, RPA, low-code and AI assistant narratives. That is not just a semantic issue. It alters who the peer set is, what growth rates feel sufficient, and how much defensibility investors assume needs to be proven.

The company's intended strategic ambition appears to be larger than that simplified category. It wants to be seen as a core operating layer that sits across multiple enterprise processes. However, that ambition is not yet framed in a way that helps the investor understand why the company earns a larger category. Today the narrative reads like an aspiration, not an earned market definition.

Investor lens: Broad category language increases the risk of being benchmarked against larger, less differentiated AI automation peer sets, which can weaken conviction before diligence begins.

6.2 Differentiation

Differentiation is where the current story underperforms most sharply. Management likely knows that the product is stronger than 'another workflow tool'. Perhaps implementation is materially faster than incumbents. Perhaps the AI layer handles exceptions rather than just happy-path automation. Perhaps customers adopt it because it compresses process turnaround and removes manual escalation work. Those are all plausible differentiators. The problem is that the deck, as assumed, leaves them implicit.

That is a costly mistake. In crowded software categories, investors do not work hard to infer differentiation. They look for one or two points that can be repeated easily in internal debate. The partner who likes the deal needs phrases that survive outside the room. If differentiation only exists in a product demo or in founder explanation, it is too fragile to do its job in a financing process.

Bessemer's warning about low switching costs and commoditised offerings in parts of AI is relevant here.[4] NovaFlow does not need to prove that no competitor can ever emerge. It needs to prove that its route to becoming the system of record for specific workflows is more compelling than the alternatives and hard enough to replicate that investors can believe in pricing power and durability.

Why this matters to capital: If differentiation relies on founder explanation rather than surviving repetition in partner discussions, conviction often stalls between first meeting and internal review.

6.3 Scalability narrative

The current scaling story is believable at a surface level but underdeveloped in the areas that matter to investors. The model suggests a move from mid-market workflow deployments into broader enterprise process ownership. That could be compelling. But the bridge between those two states is not yet sequenced clearly enough. Investors want to know: what is the initial wedge, what adjacent workflows are unlocked next, what changes in ACV as the company expands inside an account, and at what point does the product become strategically difficult to replace?

Without that sequence, 'large TAM' becomes noise rather than support. Many companies in AI workflow automation can cite a large market. Fewer can explain a credible expansion logic from today's product to tomorrow's category control. The absence of that logic does not make the company weak; it makes the company look earlier and more uncertain than it really is.

Decision implication: Without a visible wedge-to-platform sequence, investors are more likely to underwrite NovaFlow as a strong point solution than a future category-defining workflow layer.

Implication for the raise

If NovaFlow goes to market now, the likely investor reaction is: 'This is a good product in a good market, but I am not yet convinced it is the winner.' The financing process then becomes harder, slower and more valuation-sensitive.

7. Detailed diagnostic: commercial model and unit economics

7.1 Assumed operating profile

Metric	Current	Comment
ARR	\$2.4m	Good enough to support an institutional conversation, but not enough to overcome narrative ambiguity alone
YoY growth	165%	Strong, but below the rare 'shooting star' AI outliers investors now reference
Gross margin	82%	Healthy software economics
Net revenue retention	118%	Positive sign; good but not category-best
Average ACV	\$22k SMB / \$68k MM	Implies room to grow if enterprise expansion is proven
CAC payback	18 months	Acceptable, but not exceptional
Magic number	0.7x	Suggests reasonable GTM productivity, not yet elite
Monthly burn	\$280k	Manageable if growth stays robust

These synthetic metrics describe a business that is plausibly attractive, but not one whose metrics are so overwhelming that investors will forgive a weak story. That distinction matters. If NovaFlow were growing at extreme outlier levels, some ambiguity could be tolerated. But with an otherwise healthy but not extraordinary profile, narrative discipline becomes a major determinant of financing quality.

7.2 Revenue quality

Revenue quality is one of the stronger parts of the case. Gross margin at 82% is healthy, and 118% net revenue retention suggests that customers who adopt the product can find additional value over time. That is exactly the sort of evidence a workflow platform needs if it wants to claim it can move from a point solution into a broader process layer.

However, the story currently under-leverages these strengths. The metrics appear in the deck as evidence of traction, but not as evidence of thesis. In other words, they are shown, but not interpreted. A strong raise narrative would make a more explicit argument: retention is not just positive hygiene; it is proof that the platform can expand inside accounts and that the initial workflow is an effective entry wedge.

Investor interpretation: Retention is not only evidence of customer satisfaction, it is proof that the first workflow can become the economic bridge to broader account ownership.

7.3 Capital efficiency

Capital efficiency is acceptable but again under-explained. At a \$280k monthly burn and roughly seven months of runway, the round is needed on a practical timetable. That does not automatically create a problem. Many good companies raise with similar runway. The issue is whether investors perceive the company as raising from strength or from necessity.

The distinction depends on the use-of-funds story. If capital is framed as enabling specific value inflections - for example, enterprise-grade workflow depth in two target functions, a repeatable expansion motion and a narrower but stronger go-to-market focus - then the same burn profile can look efficient. If the use of funds reads more generically as 'more product and more sales', the company will look less disciplined than it is.

Capital read-through: The same burn profile can signal either urgency or strength, depending on whether the use of funds is tied to explicit proof milestones investors can underwrite.

Hidden weakness in many early-stage decks

Companies often think metrics speak for themselves. They do not. Metrics only change investor behaviour when management interprets them in a way that strengthens the business thesis.

8. Likely investor objections and investment committee simulation

The following objections are not predictions of rejection; they are the questions most likely to arise if a credible early-stage software investor were debating the opportunity internally.

Likely objection	Why it appears	What would neutralise it
This feels crowded	Category labels resemble many AI automation stories	Define a narrower wedge and explicit reason why this workflow layer wins
Where is the moat?	Differentiation is feature-led, not structurally framed	Show why workflow depth, data capture and account expansion compound into defensibility
How big can this get?	Large TAM is asserted, not sequenced	Map a clear wedge-to-platform expansion pathway
Is burn too high for this stage?	Round need and runway may look reactive	Tie use of funds to precise value inflection milestones
Could incumbents copy this?	No explicit argument on why replacing the product becomes hard	Demonstrate implementation speed plus switching cost and embedded process ownership

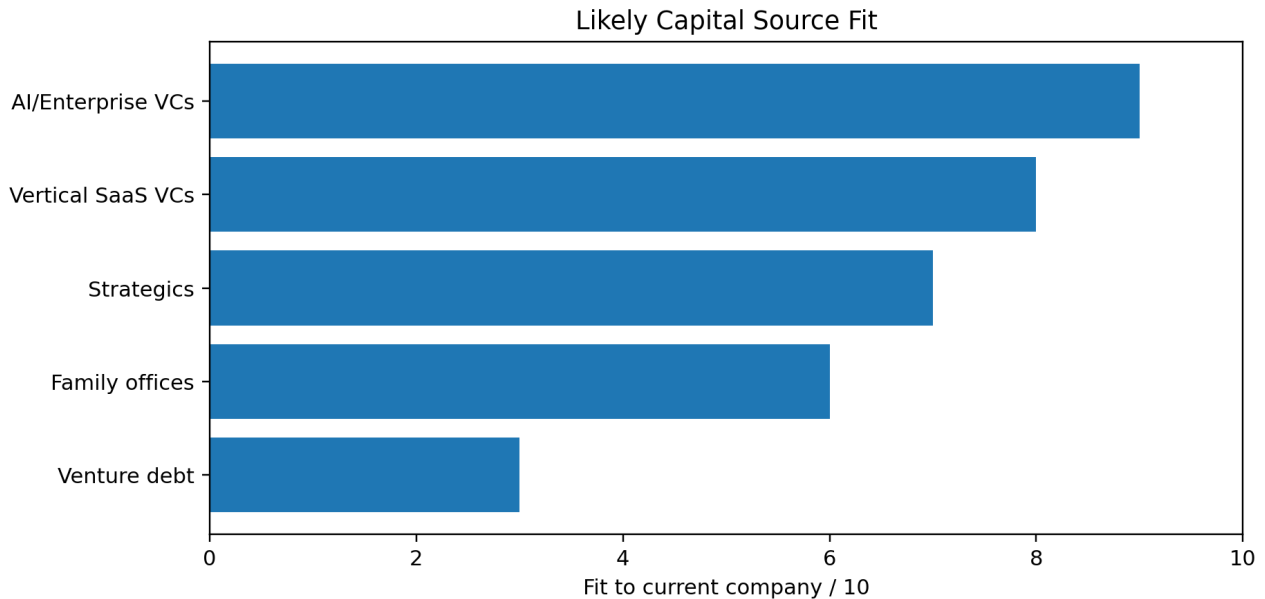
Illustrative investment committee simulation:

- Partner A: 'The market is undeniably large and budgets are moving from experimentation into core IT, which is attractive.[6] But I still need to understand why this is the winner, not just a participant.'
- Partner B: 'The retention metrics are good. That suggests the product has teeth. I would be more interested if the company framed one narrow wedge more aggressively rather than describing itself across three categories.'
- Partner C: 'I worry the company is raising on a narrative that is slightly too broad for its current proof points. Show me how this becomes a repeatable enterprise expansion machine rather than a useful automation product.'

None of those objections is fatal. That is exactly the point. The business is not being rejected for a deep flaw; it is at risk of being discounted for a narrative one. Those are often the most expensive problems, because management does not feel them as operational pain until a raise underperforms.

Fundrly readout: These are not rejection risks, they are momentum risks, the points where capital conversations lose speed even when the business remains fundamentally attractive.

9. Capital strategy: what capital is realistically available



Not all capital is equally available to NovaFlow. The company sits in a part of the market that is attractive enough to draw broad interest, but too crowded to benefit from undifferentiated outreach. The correct capital strategy is therefore selective rather than expansive.

9.1 AI / enterprise software venture funds

This is the most obvious source of capital. Generalist seed-to-Series A funds with a strong enterprise software bias, and AI-focused investors that understand workflow automation, are likely the best primary audience. They can underwrite the software economics and the category opportunity. However, they will also have the highest comparative standard because they see many adjacent companies.

NovaFlow is most likely to resonate with investors who are explicitly comfortable with application-layer AI rather than only infrastructure stories. The narrative therefore needs to answer a simple but demanding question: why is this company not just wrapping existing models, but building a durable operational layer in a workflow category that matters?

9.2 Vertical SaaS and process software strategics

Strategic capital could be relevant, particularly from software businesses adjacent to process management, service operations or enterprise systems of record. Strategics may care less about category purity and more about whether NovaFlow could become a meaningful capability extension. Their appetite depends on whether the company can demonstrate practical deployment value and a route to account-level embeddedness.

9.3 Family offices

Family offices are possible but not obviously first-choice lead capital for this profile. They can participate where the software proposition is clear, the downside appears manageable and there is confidence in governance and execution. They are more likely to respond well once the narrative is simplified and the route to value creation is less 'venture vocabulary' and more concrete commercial progression.

9.4 Venture debt

Venture debt is not a sensible primary solution at this stage. While the company has recurring revenue and decent margins, the better use of capital strategy is to optimise the narrative and

equity process rather than introduce debt complexity before the business has stronger scale proof points.

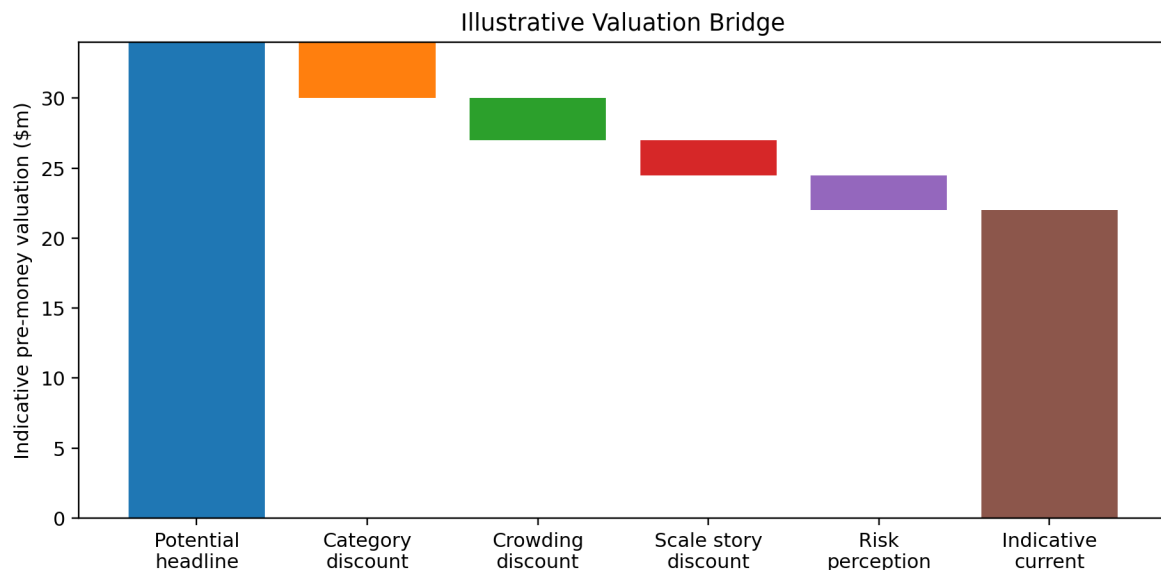
Recommended capital sequencing

Lead with a narrow list of AI- and enterprise-software-focused VCs that can understand both the software economics and the workflow category. Use strategics selectively as validation and optionality. Treat family office capital as supplementary rather than central unless the story becomes much more concrete and downside-readable.

Fit logic: Capital fit strengthens when the same operating evidence is translated differently for enterprise software VCs, strategics, and supplementary family office capital.

10. Round design and valuation framing

The stated intention in the synthetic materials is to raise \$6.0 million at a pre-money valuation in the \$32-38 million range. In the abstract, that is not unreasonable for an AI software company with healthy early revenue and strong growth. The issue is not whether that number is mathematically impossible; it is whether the current narrative earns it.



On current facts, a credible investor could see a headline valuation path in the low-to-mid \$30 million range if the company were framed as a high-conviction workflow platform with a compelling expansion logic. But the present narrative likely introduces enough discounting factors - category ambiguity, crowded comparables, under-explained differentiation and avoidable perceived risk - that the practical pre-money outcome may compress into something like \$22-28 million.

That delta is commercially significant. It changes dilution, negotiating leverage and the quality of investors willing to engage. More importantly, it changes how much 'story debt' the company carries into the next financing round. A suboptimal round can be survivable; it is just expensive.

10.1 Recommended round framing

- **Raise size:** keep the round around \$6.0 million if milestone use is tight and clearly linked to value inflection.
- **Valuation ambition:** anchor to the stronger end of the range only after the category and expansion story are tightened; otherwise preserve flexibility.
- **Milestone framing:** define the raise as the capital required to prove a repeatable wedge and expansion motion, not simply to extend runway.
- **Use of funds:** show the investor exactly what changes in the business because of this money - product depth in selected workflows, enterprise-grade deployment, and a narrower but stronger go-to-market engine.

The highest-return move is not to argue harder for the current valuation. It is to improve the story enough that the valuation feels self-evident to the right investor.

Valuation implication: Narrative precision here does not just influence pricing, it changes who is willing to lead and how much dilution the company carries into the next round.

11. Detailed recommendations and suggested narrative rewrite

11.1 Recommendation 1 - choose a single category anchor

NovaFlow should stop trying to be three things at once. The primary category anchor should be chosen based on where the company's proof points are strongest today and where the next 18 months of value creation are most plausibly earned. For this synthetic company, the strongest anchor is not 'AI operating layer' but something closer to 'enterprise workflow intelligence platform for high-friction approvals and exception handling'.

That is narrower, but strategically stronger. It tells the investor where the company starts, what kind of problem it owns and why the workflow complexity matters. Category breadth should show up later in the scale story, not at the opening line.

11.2 Recommendation 2 - make differentiation explicit and repeatable

Management should identify two differentiators that can survive repetition. For example: (i) speed-to-value because the product can automate messy exception handling rather than just templated flows, and (ii) account expansion because once embedded in one critical workflow, the company captures adjacent process ownership. Those are not feature statements; they are investment statements. They explain why this company can win commercially and why that win can compound.

11.3 Recommendation 3 - rewrite scale in sequence, not in slogans

The current scale story likely jumps too quickly from 'large TAM' to 'big company'. Instead, the narrative should sequence growth: initial wedge in finance and operations workflows; expansion to adjacent cross-functional processes inside the same account; increase in ACV through workflow density and governance features; eventual strategic position as a workflow intelligence layer in the account.

That sequence converts a broad aspiration into an underwritable plan.

11.4 Recommendation 4 - lower perceived risk proactively

Investors will assume higher risk in three areas unless management lowers it explicitly: adoption, competition and switching. The deck should therefore state more clearly why workflows chosen for initial penetration are painful enough to justify adoption, why existing alternatives do not handle the same level of exception complexity, and why account expansion increases embeddedness over time.

11.5 Recommendation 5 - align to investor archetype

The lead narrative should be built for enterprise software and AI application investors first. Strategic and family office versions can then be adapted from that core. One common mistake is trying to tell one broad story that works for everyone. It usually works fully for no one.

Execution principle: The strongest raise narratives are built around the decision logic of the highest-fit lead investor, then adapted outward rather than diluted for universal appeal.

11.6 Sample rewrite: before and after

Current-style language	Suggested investor-ready framing
NovaFlow is the AI operating layer for repetitive enterprise workflows.	NovaFlow helps mid-market and enterprise teams automate high-friction approvals and exception handling - the workflows where generic automation typically breaks.
We operate in a large workflow market.	We start in expensive, failure-prone workflows where speed-to-resolution matters, then expand account value by taking ownership of adjacent processes.
Our AI differentiates us.	Our advantage is not generic AI; it is workflow intelligence that handles exceptions, shortens deployment time and increases account-level process density.

Core narrative shift

Move from describing technology to describing the commercial mechanism by which the company becomes more valuable, more embedded and more difficult to replace.

12. 100-day pre-raise workplan

Phase	Objective	Outputs	Owner
0-30 days	Tighten narrative	Category anchor, 2 core differentiators, wedge-to-platform scale story, revised opening 5 deck slides	CEO / strategy lead
0-30 days	Evidence package	Customer quotes, implementation speed evidence, expansion examples, revised KPI framing	Customer success / finance
31-60 days	Deck and model refresh	Updated deck, milestone-based use of funds, investor-specific versions, clearer valuation rationale	CEO / finance
31-60 days	Target list design	20-30 high-fit VCs, 5-8 relevant strategics, 5-10 selective family offices	CEO / advisor
61-100 days	Market testing	Warm investor conversations, objection capture, final narrative refinements	CEO / advisor

The purpose of the first 100 days is not to produce more material. It is to reduce ambiguity. A financing process works best when every major element of the story answers a known investor question and no core part of the narrative relies on management being in the room to rescue it.

Appendix A. Synthetic operating assumptions

The following assumptions were used to make the sample diagnostic specific enough to feel real:

- Current ARR of \$2.4m, with growth of 165% year-on-year.
- 82% gross margin and 118% net revenue retention.
- Mix of SMB and mid-market customers, with ACVs of ~\$22k and ~\$68k respectively.
- Monthly burn of \$280k and cash runway of roughly seven months.
- Target raise of \$6.0m to fund product depth, go-to-market scale and runway to the next proof point.
- Indicative headline valuation ambition of \$32-38m pre-money.

None of these figures represents a real company. They are assembled to mimic the kind of 'promising but under-optimised' profile that often appears before a seed extension or pre-Series A financing.

Appendix B. Sources

[1] Mordor Intelligence, Workflow Automation Market, published 15 Jan 2026. Used for workflow automation market sizing and 2026-2031 CAGR.

[2] Fortune Business Insights, AI SaaS Market Size, Industry Share, Forecast, 2034. Used for AI SaaS market sizing and growth outlook.

[3] Grand View Research, Intelligent Process Automation Market Size Report, 2030. Used for process automation market sizing and CAGR context.

[4] Bessemer Venture Partners, The State of AI 2025, published 13 Aug 2025. Used for commentary on AI benchmark expectations and market crowding.

[5] Business Insider, 'Q2T3' is the 'freakish' new growth benchmark for AI startups, published Aug 2025. Used to illustrate new growth expectations for AI-native startups.

[6] Andreessen Horowitz, How 100 Enterprise CIOs Are Building and Buying Gen AI in 2025, published 10 Jun 2025. Used for enterprise budget growth and demand commentary.

[7] McKinsey, Technology Trends Outlook 2025, published Jul 2025. Used for commentary on the surge in AI equity investment.