KreditKraft: US Entry Sequencing

What to prove, in what order, to make the round obvious

PREPARED FOR

Klára Novotná, Co-Founder & CEO, KreditKraft
July 30, 2026

KreditKraft

KREDITKRAFT

Contents

00	Note on Scope	3
01	The Answer, In One Page	4
02	Where Your Analysis Needs Adjusting	5
03	The Sequence	8
04	Implications for the Financing Decision	10
05	Primary Deal-Breakers	11
06	Next 30 Days	12
A	Annex A: Deck and Narrative Issues	13
B	Annex B: What Would Change This Report	14
	Closing	15

Note on Scope

You sent a market analysis that already covers sizing, competition, route to market and risk, with sources and caveat flags. Repeating that would waste your time. So this report does one thing: it pressure-tests sequencing, which is what you asked for.

It adjusts your analysis in four places where I think the current read will not survive US diligence, then sets out the order of proof I would run. The tone is the straight version you asked for. I apologize in advance if some of it is unwelcome.

One observation sits underneath everything that follows, so it is worth stating plainly up front rather than leaving you to piece it together. Across the deck and your own analysis, the plan seems to be sized for a later round and a larger team than the one you are currently running:

- four business models under a single raise,
- an ask that asks for both product-market fit and scale in the same sentence,
- five markets in motion at once,
- a Q4 outcomes slide that asks for four hard things in five months,
- and a product surface spanning real estate, insurance, estate documents, held-away assets and crypto with no stated order of defense.

None of that means the underlying bet is wrong. It means the materials describe a company built for later, being pitched now.

Everything in this report is the same recommendation applied at different levels of detail: pick one budget line, one channel, one institution, prove it, before adding the next thing.

Several items in the report are marked **[OPEN]**. They are questions only you can answer, and each one materially changes the plan. Please treat this draft as directionally right and precisely wrong until those close.

1. The Answer, In One Page

> **You cannot prove US revenue inside this round, and you should stop trying to.**

A $2M raise, split across the UK and the US, against a credit union sales cycle your own analysis puts at six to eighteen months, does not reach $1M ARR by Q4 2026. Q4 2026 ends five months from now. You have $400k signed and, by your own account, nothing yet collected. That milestone could be missed in public, and missing it is worse than never having promised it.

What makes the round obvious is not revenue. It is removing the three things a US investor cannot underwrite today.

1. **Distribution is not repeatable.** Everything runs through Mastercard. One relationship, one set of champions, no independent path. The first question a serious investor would likely ask is what happens if Mastercard goes quiet.
2. **The value is not countable.** “Detects risks and prevents losses” is a claim. There is no findings-per-household number, no dollars-surfaced number, no conversion-to-institution-revenue number. Without those you cannot price, cannot build an ROI case, and cannot survive a procurement committee.
3. **The category does not exist in the US buying taxonomy.** There is no budget line called risk governance layer. And the whitespace claim in your analysis is, as written, needing adjustment. More on that in section 2.

So the sequence is: category, then countable output, then channel, then revenue. Revenue last, and mostly after this round rather than inside it.

An ROI number from one named US credit union in production is worth more to a US investor than $400k of unconverted European contracts. That is the trade this plan makes.

2. Where Your Analysis Needs Adjusting

Your document is better than most founder-supplied material. These four points are where I would push back, in order of how much damage they do if left alone.

2.1 The whitespace claim is wrong as written

Your analysis says no incumbent owns cross-domain, household-level proactive risk detection delivered into the institutional workflow. The competitor table lists Personetics, Envestnet, Wealth.com, Vanilla, Pontera, Flourish, Quantifind, Vise and Capitalize.

It does not list Carefull.

Carefull is a New York company, founded in 2019, that runs AI-driven proactive detection across a household's linked accounts and delivers it through banks, credit unions and wealth advisors. It has raised roughly $16.5M plus growth financing from CIBC Innovation Banking. More importantly, it is already inside the exact channel you have identified as your wedge:

- Partnership with the California and Nevada Credit Union Leagues since 2024
- Distribution through CUNA Strategic Services
- Named credit union deployments including SECU in Maryland and Merck Employees FCU
- A core-platform partnership with CSI announced March 2026, putting it in front of CSI's bank and credit union client base

Also absent: Eversafe, True Link, Trustworthy, Everplans and FutureVault, all of which occupy adjacent ground.

This matters for one reason. An investor will make one call to a credit union CTO and hear the name Carefull. If that name is not already in your deck, everything else in the deck becomes suspect. You lose the room on credibility, not on strategy.

It also closes off the differentiation you have been reaching for. Your positioning against "the answering AI everyone is buying" works against Personetics-style engagement copilots and against the generic GenAI tools the MIT NANDA statistic targets. It does not work against Carefull, because Carefull is not conversational either. It watches accounts continuously and pushes an alert when something fires. No chat surface, no agent the member talks to. If "we don't make you chat with an AI" reaches an investor who knows Carefull, it reads as a founder who has not checked the nearest competitor, not as a differentiator.

The axis that does hold is domain coverage, not interaction model. Carefull watches account and transaction behavior: aging-related fraud, scams, money mistakes, unusual activity. KreditKraft watches assets that generate no transaction at all: a house insured below reconstruction cost, a title with a defect, a forgotten 401(k), an expiring policy, a crypto wallet. None of those show up in a bank feed. That gap requires document ingestion and third-party data Carefull's model does not touch, which is more product surface to build and more procurement friction, but it is real and it survives contact.

The defensible version of your claim is:

> No US player detects across property, insurance, estate, held-away and digital assets at household level in a single layer. Carefull and others own transaction-level behavior; KreditKraft owns asset-level state. Every incumbent owns one domain or one data source.

Name Carefull yourself, in your own deck, before anyone else does, and make the distinction asset-level versus transaction-level rather than proactive versus conversational. Founders who name their sharpest competitor unprompted gain more credibility than the concession costs.

2.2 There is no US budget line called “risk governance layer”

US institutions buy from named budget lines. Four are plausible for KreditKraft, and they are not interchangeable. Each has a different buyer, a different cycle, a different price point and a different competitor.

Budget line	Buyer inside the CU	Competitor you meet	Reality
Elder financial exploitation and vulnerable member protection	Fraud and risk	Carefull	Real budget, regulatory tailwind from the Senior Safe Act and NCUA guidance. Fastest to funded. Most crowded.
Member engagement and digital banking	Digital banking executive	Personetics, via Alkami and Q2	Largest budget. Hardest place to look different.
Wealth transfer and next-generation retention	Wealth or member growth executive	Wealth.com, Vanilla	Thinnest budget at a credit union today. Strongest fit with your $124T narrative.
Financial wellness	Marketing	Everyone	Weakest budget. Avoid.

My recommendation is to lead with wealth transfer and next-generation retention, and treat elder and vulnerable-member protection as a coexistence story with Carefull rather than a head-on fight. Carefull already owns transaction-level behavior inside that budget line at a growing list of credit unions. Walking into an institution that has already bought Carefull and pitching a competing detection layer for the same budget is the hardest available entry. Walking in and pitching the asset-level gap Carefull structurally cannot see, an underinsured house, a forgotten 401(k), a defective title, sitting alongside what they already have, closes faster and does not ask the credit union to rip anything out.

The wealth transfer narrative is also where the asset-level distinction does the most work, because inheritance and next-gen retention hinge on exactly the property, estate and held-away findings Carefull does not touch. The detection engine is the same product either way. Only the framing, the first buyer, and whether you are positioned to displace or to complement change.

What you cannot do is stay abstract. Your first two US customers will have a named budget owner, and a deck that promises four things at once will not reach any of them.

2.3 The 1033 framing is spin, and it is the kind investors catch

Your analysis presents the CFPB 1033 injunction as a hedge that favors KreditKraft, since a read-only architecture does not require it. That reads as advantage-hunting in a document that is otherwise admirably honest.

The current state is that the rule is enjoined and effectively unenforceable, the CFPB is rewriting it, the reopened rulemaking explicitly includes whether banks may charge fees for data access, and individual states are now legislating separately. For a product whose value depends on aggregating household data, that is a cost risk and a fragmentation risk. It is not a tailwind.

Better framing, and it happens to be true: KreditKraft does not depend on a data-access mandate that may never arrive. That is precisely why the core and digital-banking channel partnership is not only a distribution decision: it is also how you obtain data access inside the institution's own perimeter, without paying aggregation tolls that may be about to get worse.

That version turns the same fact into an argument for the channel strategy, which is the strategy you want funded anyway.

2.4 Your note and your analysis disagree about credit union speed

Your note says credit unions are faster to procure than banks. Your own analysis says credit union purchasing is slow and committee-driven, and cites Filene testers building in three extra months for discovery alone.

Both cannot headline. Best correct it now, internally, because the calendar you plan against depends on it.

Credit unions are the right wedge for mission alignment, referenceability and the fact that one credit union will happily take a reference call from another. They are not the right wedge for speed. Plan on nine to twelve months direct, four to six through a platform channel, and budget the runway accordingly.

One further note on the data moat. Seventeen thousand European consumer users, unmonetized, with retention undisclosed, is not a moat against a competitor already installed inside US credit union channels.

It is genuine product validation and a useful training corpus for detection rules.

Call it that.

Overclaiming here is cheap to fix and expensive to leave in, because it invites an investor to discount everything else you assert.

3. The Sequence

Four phases. Phases 0 and 1 overlap. Phase 2 runs parallel to Phase 1. Nothing here assumes US revenue before 2027.

Phase 0: unblock the structural gates (start this week, complete by September)

These are calendar-bound rather than effort-bound. They cost little, they cannot be accelerated later with money, and every one of them gates something downstream. This is where foreign founders routinely lose two to three months during a term sheet.

SOC 2 Type II. This is the hardest constraint in the entire plan and it is not a sales problem. A Type II requires a continuous observation window, typically three to six months, before a report exists. If the window opens in August, you have a report between November and February and you can close credit union contracts in 2027. If it opens in January, you close nothing in 2027. No amount of sales effort moves this. [OPEN] Whether the window has started is the single most important fact I do not have.

The Delaware question. One of your existing investors is a public regional fund, which matters here. Before a US term sheet, you need written answers to: does any instrument restrict redomiciliation to a Delaware topco, and does any grant or subsidy carry a clawback if IP or economic activity moves out of Czechia. This routinely takes three to six months to resolve and it is discovered, badly, during diligence. Recommend resolving it now, in writing, while it is a housekeeping item rather than a deal item. [OPEN]

Your status. Start Path does not confer a visa. Czechia is an E-2 treaty country, and O-1 is the other common route. Suggest confirming the path and the filing timeline before September. [OPEN]

US entity, banking, and the data path. Delaware C-corp, US bank account, and a decision on where US institutional data lives. A European data path is a procurement objection you do not want to be arguing about in month nine.

Category decision. Section 2.2. Pick the budget line before the first US meeting, not after the fifth.

Phase 1: prove countable output on one referenceable institution (September to December)

Not five pilots. Better to have one, in production, instrumented.

Convert one of Stanford FCU or Wings into a paid pilot governed by a measurement agreement. **The measurement agreement is the deliverable, not the software.** It should fix, in advance:

- Number of households analyzed
- Findings surfaced per household, by domain
- Dollar value of assets, coverage gaps and inefficiencies identified
- Percentage of findings the member or the institution acted on
- The institution’s own consequence: retention, deposit movement, non-interest income, product attach

Then instrument the Generali and Partners Banka engagements retroactively into the same schema, so you arrive at the round with three data points rather than one, and so the European traction finally converts into something a US investor can read.

This is the asset that does not exist today and that everything else depends on. Pricing depends on it. The ROI case in procurement depends on it. The differentiation claim against Carefull depends on it, because "we find things in domains they do not cover" is only persuasive with a count attached.

Pick Stanford or Wings on one criterion: which one already runs on a digital banking platform you intend to partner with. That answer is free and it links Phase 1 to Phase 2.

Phase 2: land one channel agreement (October to Q1 2027, parallel)

Q2 Innovation Studio, Alkami, or Jack Henry FIN. One, not three. Each has its own integration burden and a seven-person team cannot carry three.

The channel does two things that matter more than sales. It answers the investor question about whether distribution survives without Mastercard, and it absorbs a meaningful share of the vendor due-diligence burden that a small foreign vendor would otherwise carry alone.

One recommended correction to your analysis: a marketplace listing compresses the sales cycle. It does not generate demand. Nobody browses these and buys. You still have to create the deal and then run it through the channel to shorten it. Advise to plan the pipeline accordingly.

Phase 3: raise (Q1 to Q2 2027)

At that point the story is: one named US institution in production with published detection economics, one platform channel with pre-built integration, SOC 2 Type II in hand, Mastercard as one of two channels rather than the only one, and a competitive claim that has already survived contact with the incumbent.

That is a fundable US story, and it prices materially better than the same company pitching today on signed European contracts and a whitespace claim that does not hold.

4. Implications for the Financing Decision

You said the sharpest thing would be pressure-testing sequencing because you are at a decision point on financing and timing. Here it is.

The current plan asks investors to fund hope. A $2M raise, deployed across two geographies, against a Q4 2026 milestone of $1M ARR that the sales cycle mathematically cannot deliver. Investors do not fund that. They fund de-risking, and they want to see that you know which risks matter.

There are two coherent options:

Option A: bridge now, real round on US proof. Raise $750k to $1M from existing investors plus the Start Path network, on a SAFE, explicitly to fund Phase 0 and Phase 1. Then raise the full round in Q1 to Q2 2027 against a named US reference and countable detection economics. More total dilution only if the proof fails. Less, and at a much better price, if it lands. This is the option I would take if the runway supports it. **[OPEN: this depends entirely on current runway in months, which I do not have.]**

Option B: raise the $2M now, re-cut the milestones. Same money, but the milestone slide changes from ARR to proof: SOC 2 Type II delivered, one named US institution live with published detection economics, one channel agreement signed, Mastercard concentration reduced below a stated threshold. Investors who understand FI sales will recognize that as a competent plan. Investors who only respond to an ARR number were never going to be good partners for a six-to-eighteen-month sales cycle anyway.

What I would not do: raise $2M against a $1M ARR promise and split it across the UK and the US. That funds two half-motions and proves neither, and it puts you back in the market in twelve months having missed a public number.

The UK question you have not made explicit

The UK has Consumer Duty as a genuine regulatory forcing function, existing Mastercard distribution, no visa problem, no entity problem, and no SOC 2 gate. The US has the larger prize and a much higher cost of entry.

There is a defensible sequencing answer that perhaps nobody has put in front of you: **the UK is your 2026 revenue market and the US is your 2027 revenue market.**

UK revenue funds the US proof. The US work in 2026 is Phase 0 and Phase 1 only, which is cheap, and it is founder time rather than team time.

That is a real answer to “timing,” and it may be a better one than treating the two markets as a single simultaneous push.

5. Primary Deal-Breakers

- **SOC 2 not started.** Nothing closes, the whole plan slips twelve months, and the round slips with it.
- **Mastercard cools during the transition.** You have no independent channel yet, so this is fatal until Phase 2 lands. It is also why Phase 2 should not wait for Phase 1 to finish.
- **A cap table or redomiciliation blocker surfaces during a term sheet.** Deal dies or reprices. Entirely preventable this month.
- **Carefull or Personetics ships the cross-domain version.** Neither is durably blocked from doing it. Your defense is depth across domains, workflow integration, and speed, not category novelty.
- **You become the bottleneck.** Founder alone in New York with delivery in Prague means you personally attend every technical diligence call across a nine-hour split day. This is the failure mode nobody plans for and everybody hits. **[OPEN: just to confirm who covers US technical delivery? It seems Radek is not still with the company, I think we spoke about this on our call. The deck lists Tomáš Bureš as CPTO and a technical lead, and Radek is not on it.]**

6. Next 30 Days

1. Engage a SOC 2 auditor and open the Type II observation window. Nothing else on this list matters as much.
2. Get written confirmation from your existing investors on redomiciliation and any clawback exposure.
3. Confirm your visa route and file.
4. Choose the budget line from section 2.2 and rewrite the one-line positioning to match.
5. Add Carefull to your competitive slide, with the narrowed whitespace claim from section 2.1.
6. Ask Stanford FCU and Wings which digital banking platform they run. Use the answer to pick the channel.
7. Draft the measurement agreement for the first pilot before selling it.
8. Retro-instrument Generali and Partners Banka into the same metric schema.
9. Decide between Option A and Option B on financing, against a runway number.
10. Start scoping US-based technical delivery capacity now, ahead of need. If Radek and the Prague team are not going to carry US technical diligence calls and integration work across a six-to-nine-hour offset depending on the prospect's coast, you want that gap named and a plan for it before a pilot exposes it live, not during one. A short, defined scope, such as the Phase 2 platform integration, is a low-risk way to test a US technical partner before committing further.

Annex A: Deck and Narrative Issues

Separate from the strategy, because these are presentation problems with strategic consequences. US seed investors read sloppiness as an operating signal.

The market slide does not reconcile. SOM is given as $8.4B to $48B against a SAM of $300B. That is 16 percent penetration at the top end, which no investor will accept. It sits next to a €2.5B to €4.2B figure for the same markets in a different currency. Pick one currency, one methodology, and a SOM that is defensibly small. A credible SOM built bottom-up from 450 credit unions above $1B in assets, times a realistic per-institution contract value, will be a far smaller number and a far more persuasive one.

"Expected Outcomes by Q4 2026" is five months away. It promises $1M+ ARR, a fully operational AI agent, validated POCs with banks and credit unions, and a signed credit union association covering 500,000 members. Re-cut against section 4.

"Signed ARR." $400k signed with nothing collected is not ARR and a US investor will say so. Either move it to contracted value with a stated conversion timeline, or wait until it converts. **[OPEN: I still need the contract terms, counterparties, and whether any portion is Mastercard-funded rather than institution-funded. If it is Mastercard-funded, it is Mastercard revenue with bank logos on it, and it needs to be presented that way before someone else characterizes it.]**

Minor typos on the two most-read slides. "Affluent American looses annually" on slide one (should be losses). "Baked by VC and industry leaders" on the team slide, which should read Backed. "100.000k ARR" on the traction chart.

The one-liner. "KreditKraft is risk governance layer to get full visibility into financial portfolios and uncover hidden inefficiencies" is missing an article and, more importantly, describes a capability rather than a buyer's problem. Once you pick the budget line, rewrite it to name the buyer and the consequence, for example, tuned to the elder and vulnerable-member protection line recommended in section 2.2, and doing quiet double duty as the asset-level distinction from section 2.1:

> "KreditKraft shows credit unions the property, insurance, and estate risks sitting in a member's life that transaction monitoring can't see, before they become a loss, a claim, or a complaint."

If the wealth-transfer narrative ends up leading instead, the buyer-facing version of that story could be:

> "KreditKraft helps credit unions catch the asset risks that never show up in a transaction, before a household's wealth quietly moves to the next generation, or the wrong institution."

Annex B: What Would Change This Report

The following would change the recommendation, not just the detail.

Runway in months. Drives Option A versus Option B entirely.

SOC 2 status. If a Type II observation window is already open, Phase 1 and Phase 2 can compress by a quarter and a 2026 close becomes possible.

The $400k contracts. Terms, counterparties, funding source, termination rights, and conversion timing.

Stanford FCU and Wings. Stage, whether the contact is a decision maker or an innovation scout, whether there is a written next step, and how each was sourced.

Generali and Partners Banka outcomes. If countable results already exist, Phase 1 shortens dramatically and the round can move forward.

Product reality. Your analysis describes a read-only, on-premise architecture. The deck shows Plaid as a partner. These imply different US builds, different costs and different procurement conversations. On-premise across many institutions is a services burden a seven-person team should think hard about carrying.

Round terms already in motion. Instrument, valuation, and how much of the $2M is soft-circled.

Closing

One thing worth repeating is that everything above is the same argument at different resolutions: the plan is built for a later round and a larger team, and the fix is to prove one thing, not four, before the next raise. While the open questions may change the detail, they do not change that conclusion.