

GLOBAL REIMBURSEMENTS

COMPLIANCE & RISK REVIEW

LAUREN ORIE

TABLE OF CONTENTS

EXECUTIVE SUMMARY	01
RISK 1:	02
Outsourced Operations & Retained Accountability	
RISK 2:	03
Long Dispute Windows Against US-Calibrated Fraud Controls	
RISK 3:	04
Waitlist to Build Controls & Not Delaying the Problem	
RISK 4:	05
Limited Auditability and Defensible Decision-Making	
CONCLUSION	06

EXECUTIVE SUMMARY

OVERVIEW

The Company's proposed Global Reimbursements feature solves a real and well-documented problem: customers are currently running two reimbursement workflows, and their international employees wait longer to get paid. Working with the Partner to handle non-USD processing is a smart way to reach the market quickly without building international payment infrastructure from scratch.

The proposal is strong on how the feature works under normal conditions. The risks I'd prioritize all sit in what happens when conditions aren't normal, and in what the Company remains accountable for even though the Partner does the processing. Four stand out: oversight of an outsourced operation the Company still answers for, long international dispute windows that the company's US-calibrated fraud controls weren't built for, a waitlist that should be used to build controls rather than just to delay the problem, and documentation gaps that could make decisions hard to defend later.

My overall view is that the feature can and should launch, but on a gated basis, with a clear line drawn between what's acceptable to handle manually at low volume and what has to be built before any funds move. Handling rare exceptions by hand during the waitlist is proportionate and genuinely useful for learning. Launching without settled loss-allocation, basic oversight, and an auditable record is not, because those are the gaps that get expensive quietly. The goal isn't to slow the launch. It's to make sure the controls scale as fast as the product does.

RISK 1

Section Remarks

The Company still answers for what the Partner does, so oversight and the legal characterization need to be settled before any funds move.



RISK 1

OUTSOURCED OPERATIONS & RETAINED ACCOUNTABILITY

THE RISK

Building on the Partner instead of standing up international payments in-house is the right call, and the proposal is correct that it gets the feature to market faster. The risk isn't the outsourcing. It's that the plan leans on the Partner as the place where compliance risk gets managed, when the accountability for that risk never actually leaves the Company.

The proposal's own framing is worth sitting with here. It states that the Company never holds, controls, or transmits customer funds, and that the company doesn't hold international money transmitter licenses, which is exactly why the Partner is doing the processing. That structure is deliberate and sound. But two facts sit in tension with the hands-off posture: The Company owns the entire customer-facing experience, and the Company takes a 15 to 25 percent share of the Partner's FX margin. A company that controls the interface, sets the payment instructions, and earns a cut of the transaction is not as far from the activity as "we never touch the funds" suggests.

WHY IT MATTERS

There are two distinct exposures here, and they're worth separating.

The first is operational visibility. Customers won't distinguish between the Company and the Partner. If onboarding stalls or a payment fails, they'll see a Company problem, because the experience lives on the Company's platform. If the only signal that something broke is a customer escalating it, the Company is operating blind on activity it's answerable for.

The second is the regulatory characterization itself, and this is the one I'd flag for counsel rather than assert myself. Whether the Company's role stays cleanly on the "we only collect instructions" side of the line is a legal determination, not a settled fact, and the revenue share and interface ownership are exactly the features a regulator would look at. The proposal already lists the vendor contract and the Terms of Service as open items requiring outside counsel, which tells me the team senses this. The point is that the answer to "are we comfortably not a transmitter here" should be confirmed before launch, not assumed.

RISK 1

LAUREN ORIE

OUTSOURCED OPERATIONS & RETAINED ACCOUNTABILITY

RECOMMENDATION

Two things, matched to the two exposures.

For visibility, the Company needs a defined oversight arrangement with the Partner before launch: who owns what, how issues get reported and escalated, what status information the Company is entitled to and how fast. The goal isn't to rebuild the Partner's controls internally. It's to make sure the Company isn't dependent on a customer complaint to learn that something failed. In practice that can be regular status reporting, defined escalation channels, and automated status signals on the payment flow. The specific mechanism matters less than the principle that the Company can see what it's accountable for.

For the regulatory question, the existing open items are the right vehicle, they just need to be closed before launch rather than after. Counsel should confirm how the Company's role is characterized given the interface control and revenue share, and the Terms of Service should reflect that characterization clearly. This isn't net-new work. It's making sure the items the team already flagged are resolved on the right side of the launch date.

TIMING

Pre-launch. The Partner sits at the center of the whole model, so both the oversight arrangement and the legal characterization need to be settled before any customer funds move. These aren't things you want to be reasoning about for the first time while responding to a live incident.

RISK 2

Section Remarks

Long international dispute windows can reverse a payment long after it's settled, and the US fraud model has no answer for that.



RISK 2

LAUREN ORIE

LONG DISPUTE WINDOWS AGAINST US-CALIBRATED FRAUD CONTROLS

THE RISK

The proposal is refreshingly direct about this one, and I want to build on what it already admits rather than invent a new concern. In the Risk Mitigation section, the plan states that the Company's existing fraud controls depend on two things: a short two-day dispute window on US direct debits, and third-party verification services that confirm account ownership. It then says plainly that these tools aren't available internationally.

So the feature extends payouts into markets where the Company's proven fraud defenses don't reach, and where the dispute windows are dramatically longer than the two days the current model is built around.

The reason this is the risk I'd rank second is the interaction between those long windows and how the money moves. The Company pulls funds from the company's account and pays the employee, and for cross-border payments the proposal has the Partner converting at the spot rate and sending funds, with the reimbursement marked confirmed after four business days of silence. But a UK or Australian dispute can land long after those four days, after the employee has been paid and the FX has settled. If the funding debit gets reversed weeks, months, or in Australia's case years later, the money is already gone, and any currency movement in the meantime widens the gap. A two-day-window fraud model has no answer for a seven-year clawback.

There's a quieter version of the same problem worth naming: marking a payment confirmed after four business days of no error notification assumes that silence means success. In a cross-border flow across multiple banking systems, silence can also mean a delayed or dropped notification. Confirming on the absence of bad news is a thin signal when the dispute tail is this long.

WHY IT MATTERS

This is where a customer-experience issue quietly becomes a financial-exposure issue. As long as the Company is the party in the chain that pulled the funds and instructed the payout, a reversal that arrives after settlement is the Company's problem to absorb or chase, even though the Partner moved the money. The longer the window, the longer that exposure stays open on the books, and the proposal currently has no defined process for it, dispute resolution is explicitly listed as an open item to be built reactively. Building a clawback process in the middle of your first major clawback is the worst possible time to build it.

RISK 2

LAUREN ORIE

LONG DISPUTE WINDOWS AGAINST US-CALIBRATED FRAUD CONTROLS

RECOMMENDATION

The honest framing is that the launch gate is doing real work here. The waitlist limits this to trusted, long-standing customers, which genuinely lowers the near-term fraud risk. That buys time, but it doesn't close the gap, so two things should be in place before broad rollout.

First, the Company should clarify with the Partner, in the contract, who bears the loss when a payment is reversed or disputed after settlement, including the currency difference on a reversal. That's a contract and disclosure question, and it's far cheaper to settle before launch than to argue after a loss. It ties directly to the open Terms of Service and vendor-contract items.

Second, the long-window currencies deserve closer handling than a flat rollout. That can mean watching the GBP and AUD corridors more conservatively during the waitlist phase, being deliberate about how much volume rides on them early, and treating the four-day auto-confirm as provisional rather than final on those rails. The broader point is that the fraud controls for this feature can't be ported from the US model. They have to be rebuilt for an environment where verification tools are absent and the dispute tail is long, and that rebuild is what the waitlist period is for.

TIMING

The feature can launch to the gated waitlist before all of this is fully built. That's a proportionate use of the trusted-customer cohort, and it's the right call to retain the accounts that have been asking.

The loss-allocation question, though, should be answered pre-launch. It's contractual, it's cheap to settle now, and it gets expensive once a real reversal has already landed. That belongs with the open Terms of Service and vendor-contract items, closed on the right side of the launch date.

The corridor-specific controls can mature during the waitlist phase, as long as that phase is explicitly being used to build them rather than just to delay the problem. Watch the GBP and AUD rails more conservatively, be deliberate about how much early volume rides on them, and treat the four-day auto-confirm as provisional on those corridors until the dispute behavior is better understood.

RISK 3

Section Remarks

The waitlist is only worth the wait if it's used to build controls and set objective access criteria, not just to cap volume.



USING THE WAITLIST TO BUILD CONTROLS & NOT JUST DELAY THE PROBLEM

THE RISK

The waitlist is the proposal's main risk control at launch, and it's a sound instinct: gate the feature to trusted, long-standing customers while more robust controls get developed. The risk is in how the waitlist gets treated. The plan describes it as a way to buy time, but it doesn't say what that time is for. If the waitlist is just a volume limit, the Company reaches the end of it with the same control gaps it started with, only now under pressure to open up because the trusted customers are happy and the feature works.

There's a second piece the proposal leaves thin: who decides who gets in. Account Managers review requests and selectively enable the feature. AMs are the right people to identify which customers want this, but they're commercially oriented by role, and asking them to also be the risk gate creates an inconsistent standard. One AM's "trusted, long-standing" is not another's, and there's no objective criteria written down.

WHY IT MATTERS

International payments don't hold still. The proposal supports five currencies across markets with different rails, different dispute windows, and different fraud exposure, and a lot of what actually breaks only becomes visible once real money is moving. The waitlist period is the one window where the Company can watch that happen at low volume, with low-risk customers, before the stakes are high. That's a genuinely valuable asset, but only if the company is deliberately capturing what it learns: which corridors are smooth, where settlement lags, what the real exception patterns are, how the Partner actually performs. If that learning isn't structured, the waitlist ends and the expansion decision gets made on instinct rather than evidence.

On the AM point, the risk is consistency and defensibility. If access decisions are subjective and undocumented, the Company can't easily show why one customer was enabled, and another wasn't, and it can't spot if a high-risk account got waved through on a judgment call. That's a small problem at waitlist scale, and a real one as volume grows.

RISK 3

LAUREN ORIE

USING THE WAITLIST TO BUILD CONTROLS & NOT JUST DELAY THE PROBLEM

RECOMMENDATION

Treat the waitlist as a pilot with a defined purpose, not just a gate. Before broad rollout, the Company should be tracking the things that will inform the real controls: corridor-level performance, dispute and exception patterns, settlement timing, the Partner reliability. The expansion decision should be tied to what that data shows, not just to whether the waitlist customers are satisfied.

For access, give AMs simple, objective criteria rather than open discretion, things like account tenure, payment history, and absence of prior reversals. AMs can still own the commercial side of identifying demand, but the risk threshold should be a written standard applied consistently, with anything outside it escalated rather than waved through. That keeps the gate both faster and more defensible.

TIMING

Pre-launch and ongoing. The pilot structure and the access criteria should be set before the waitlist opens, since that's the whole point of the waitlist. The learning itself continues well past launch, because the risks and costs here keep shifting as the feature expands into more corridors and higher volume.

RISK 4

Section Remarks

DECISIONS MADE ACROSS TWO COMPANIES NEED A DOCUMENTED TRAIL, BECAUSE LONG DISPUTE WINDOWS GUARANTEE SOMEONE WILL ASK ABOUT THEM LATER.



RISK 4

LAUREN ORIE

LIMITED AUDITABILITY & DEFENSIBLE DECISION-MAKING

THE RISK

The proposal says very little about how decisions get documented and retained: payment exceptions, disputes, the AM access calls, escalations with the Partner, anything that falls outside the standard flow. The risk isn't that the Company will make bad calls in the moment. It's that months later, the company may not be able to show clearly why a reasonable call was made, who made it, and what it was based on.

This is harder here than in a normal workflow because the program runs across two companies. The Partner handles parts of onboarding, KYC, payment, and FX, while the Company owns the customer relationship and the records. When something needs to be reconstructed later, that split is exactly where the trail goes cold if the documentation standard wasn't set up front.

WHY IT MATTERS

In a payments program, decisions get revisited long after they're made, and the long dispute windows from Risk 2 guarantee it. A disputed payment in the UK or Australia could surface well after the fact and require the Company to explain what happened and why. If the records are thin or scattered across Slack threads and email, the company can struggle to demonstrate it exercised proper oversight, even when the original decision was perfectly reasonable.

Good documentation does a few things at once: it makes clear which party did what, it keeps similar situations handled consistently as volume grows, it surfaces patterns worth fixing, and most importantly, it lets the Company defend a decision if a customer, auditor, or regulator ever asks. Given that audit readiness is part of the broader compliance picture for a company at this stage, this isn't optional infrastructure; it's the thing those reviews actually examine.

RISK 4

LIMITED AUDITABILITY & DEFENSIBLE DECISION-MAKING

RECOMMENDATION

Before launch, set documentation standards for Global Reimbursements: what gets captured, where it lives, who maintains it, and how long it's kept. At minimum, that covers payment approvals, delays, returns, and cancellations; disputes and outcomes; escalations with the Partner; exceptions to standard processing; the AM access decisions from Risk 3; and higher-risk transactions.

Where it's practical, the strongest version builds this into the workflow itself, so records are generated as a natural byproduct of the process rather than something added manually after the fact. That makes them both more reliable and easier to pull when someone needs them. The standard doesn't need to be elaborate at launch; it needs to exist before the first dispute, so the record is being kept from the first transaction rather than reconstructed after a problem.

TIMING

Pre-launch and ongoing. Auditability works best built in from the start, since you can't retroactively document decisions that weren't captured. The standard can be refined over time as dispute patterns and regulatory expectations develop, but the basic record has to be in place before launch.

CONCLUSION

SYNOPSIS & CLOSING REMARKS

The thread running through all four risks is the same. Global Reimbursements is worth launching, the Partner relationship is a smart way to get there, and the waitlist is a reasonable way to start. The work is making sure the controls scale with the product instead of getting bolted on after a problem forces the issue.

That's where proportionality matters, because not every risk needs the same response on day one. During the waitlist phase, handling rare routing errors, isolated delays, and one-off exceptions by hand is completely reasonable. Volume is low, the customers are trusted, and that hands-on period is exactly where the team learns what the program actually needs. Doing that manually at the start is a feature, not a gap.

What can't be deferred is the foundation: a clear oversight arrangement with the Partner, a settled answer on who bears the loss when a payment reverses, objective access criteria, and an auditable record of decisions. Those are the gaps that compound out of sight and get expensive once real volume is moving. The distinction I'd hold onto is simple: manual effort on edge cases is fine to start with and automate later, but missing structural foundations is not. Set the foundation before the waitlist opens, use the pilot to learn where the real friction is, and build the rest toward what the data actually shows. That's how this launches responsibly and becomes something the Company can stand behind as it scales.