PREVENANCE

CASE STUDY • INTERNATIONAL EXPANSION SUPPORT

$0 → $2.4M US ARR in 14 months

A pre-Series B data governance RegTech entering the United States from Australia without a US office, a US hire, or a Series B raise.

$2.4M

US ARR IN 14 MONTHS FROM A STANDING START

Prévenance

The AI-native operating partner for PE-backed SaaS

ENGAGEMENT SUMMARY

At a glance

Client	Australian-headquartered data governance RegTech scale-up
Product	AI-driven data classification, lineage, and regulatory reporting platform for financial services
Headline metrics at engagement	$11M ARR (AU/NZ) • Series A closed 14 months prior • 27 customers
Strategic context	US entry mandated by the Series A thesis; Series B raise planned at month +18, contingent on US traction
Constraint	No US-based hires; no US office; no incremental funding before Series B
Prévenance engagement	International Expansion Sprint (8 weeks) + 14-month Retainer
Headline outcome at month +14	US ARR $2.4M • 11 US customers (avg ACV $218K) • Series B closed at $24M (target $20M)

SITUATION

An AU RegTech with a US-shaped problem

The client had built a data governance platform that automates classification, lineage tracking, and regulatory reporting for structured and unstructured data within financial services organisations. The product was particularly strong on automated compliance reporting against APRA (Australia), MAS (Singapore), and FCA (UK) regimes. The AU/NZ customer base — by month +14 of their Series A — included three of the four major Australian banks and several large insurers.

The Series A had been raised against a thesis that explicitly required US market entry. Eighteen months in, no US revenue had been booked. The Series B raise was scheduled for month +18 from the engagement start; without US traction, the raise would be a down-round or delayed. The board's diligence committee had identified US entry as the highest-risk, highest-reward variable over the next eighteen months.

Three constraints made the problem structural rather than tactical. First, the founders were not willing to hire into the US until the model was proven — reasonable, but it precluded the standard "hire a US head of sales and let them figure it out" motion. Second, the runway between engagement start and Series B did not support a 12-month sales cycle if entry began from zero. Third, the regulatory shape of the US data governance market is materially different from AU/UK — sector-specific rather than

principle-based, with state-level overlays — and the product positioning that worked at home was unlikely to work without modification.

INTERVENTION

Distributed entry, not transplanted entry

Conventional international expansion advice is variations on "hire a US head of sales, open a New York office, run a US playbook." For a $11M ARR company without the capital or the founder bandwidth for that, the advice is inapplicable. The engagement was architected around a different model: distributed entry executed from Australia with AI-augmented coverage, rather than physical entry executed from the United States.

Phase 1. ICP and regulatory positioning

The first six weeks were spent on positioning. The product's strongest single argument in AU — automated regulatory reporting — had no US analogue. The strongest US analogue — OCC and FDIC reporting for federally regulated banks — required additional product configuration that the company had not built. Three potential US ICP segments were identified and tested against the existing product capability: (1) regional banks with $5B–$50B in assets, where the principal compliance pressure is similar to AU mid-size regulated banks; (2) credit unions in the $1B–$10B AUM range, similar shape; (3) state-regulated insurance carriers, regulatory pattern most similar to the existing UK customer profile.

Within four weeks, the priority segment was confirmed as regional banks ($5B–$50B), with credit unions as a secondary motion. The product roadmap was adjusted to add OCC-aligned reporting templates as a 12-week build, gated to the first signed US customer.

Phase 2. AI-augmented outbound from Sydney

The outbound function for the US market was built and run from Sydney. Two SDRs were assigned to the US motion (one a new hire, one redeployed from the AU team). They operated against a target account universe of approximately 1,100 US regional banks scored against the refined ICP, with AI agents performing account research, regulatory context research, multi-touch sequence drafting, and meeting coordination across time zones. The agents' outbound was timed for US business hours; the SDRs ran live conversations during the overlap window between Sydney and US East/Central time.

The asymmetry of the working day was managed deliberately. Inbound responses to a US prospect coming in at 3 a.m. Sydney time was handled by an inbound qualification agent within minutes, which booked qualifying calls into AE calendars for the next mutual working window. No US prospect ever encountered a 12-hour silence — the standard experience of contacting an AU vendor from the US.

Phase 3. Distributed AEs and the first US customer

Two existing AEs were repositioned to the US motion, with primary accountability. Both were AU-based. Both worked predominantly across the Sydney evening / US Morning window. The first US customer signed in month +5 of the engagement — a $180K ACV regional bank in the Pacific Northwest, sourced entirely through AI-augmented outbound, closed across three video calls and a 90-minute on-site visit by the founder. The on-site was the only US travel the engagement required for that first deal.

Customers two through six followed in months +7 through +11, with ACV ranging from $140K to $ 290K. By month +11, the US motion was operating at roughly 60% of the AU motion's velocity per AE, which was considered acceptable given the cross-timezone constraint and the absence of physical presence.

Phase 4. Selective US presence and Series B preparation

At month +12, with five signed US customers and a US pipeline of approximately $1.6M, the model was sufficiently proven to support selective US presence. The founder relocated to Boston for an extended sixteen-week customer-anchored stay. A US-based commercial advisor (fractional, part-time) was contracted to provide market-relationship coverage. The Series B raise opened at month +15.

OUTCOMES

Fourteen months later

Metric	Baseline	Month +6	Month +14
US ARR	$0	$180K (1 customer)	$2.4M (11 customers)
US customers signed	0	1	11
Average ACV (US)	—	$180K	$218K
US headcount (US-based)	0	0	1 (founder, temporary)
US pipeline	$0	$520K	$4.8M
US customer acquisition cost	—	$84K	$71K (blended)
Sales cycle (US)	—	4.5 months	3.8 months
Total ARR (global)	$11.0M	$14.7M	$21.6M

Series B outcome

The Series B closed at month +16, with a $24M raise at a pre-money valuation that materially exceeded the original target. The lead investor's diligence specifically identified the US traction trajectory and the demonstrated capital efficiency of the entry model as material to the valuation outcome. The company's ratio of US ARR generated to US-deployed capital was approximately 4.8x more efficient than the comparable cohort of AU/EU RegTech firms that had entered the US through conventional means.

> *The Series B lead didn't price an Australian company at $11M ARR. They priced an $11M ARR Australian company with $2.4M in US ARR, a proven distributed entry model, and 4.8x the capital efficiency of comparable RegTech US entries.*

REFLECTION

What worked, what almost didn't, and what this means

Three things worked better than expected. First, the regional bank ICP. The decision was contrarian — most AU RegTechs entering the US target the top-50 banks because the deal sizes are larger. The regional bank segment was less crowded, more receptive to a non-US vendor, and operationally faster to navigate. Average ACV at $218K was lower than the top-50 strategy would have produced per deal, but velocity was materially higher, and capital efficiency was the metric that mattered for the Series B narrative.

Second, the AI-augmented outbound model. The decision to run outbound from Sydney rather than hire US-based SDRs saved approximately $1.4M in personnel costs across the 14 months and produced a faster ramp than the hire-and-train alternative would have. The risk was that US prospects would discount the AU origin; in practice, the contextual depth of the agent-drafted outbound meant most prospects did not register the geographic origin until late in the cycle, by which point product fit had been established.

Third, the founder's deliberate relocation to the US at month +12, after the model was proven. Conventional advice would have been to relocate earlier. Relocating later — with five signed customers, a $1.6M pipeline, and a working motion — meant the founder's time in the US was devoted to closing and expansion rather than to early-stage learning. The capital efficiency of that decision is hard to overstate.

Two things almost didn't work. The OCC-aligned reporting build was 9 weeks late. It nearly cost the second customer; it was completed in time only because the customer's own procurement process slipped. Earlier scoping of the product gap would have made the schedule less precarious. And the time-zone overlap window for AE conversations was structurally constraining; in retrospect, one of the two

AEs assigned to the motion should have been UK-based rather than AU-based to widen the live-conversation window into European mornings and US East afternoons.

This engagement reframes the question of what international entry actually requires. For a $11M ARR company, conventional advice prescribes “double the team, double the capital, then go.” In a 2026 AI-native commercial environment, the actual minimum viable US entry can be conducted with two redeployed SDRs, two repositioned AEs, no US-based headcount, and a founder who travels once a quarter. The capital efficiency of this approach, compared with the conventional alternative, is roughly fivefold. The strategic implication is that international entry should now be evaluated as a candidate move much earlier in a company’s lifecycle than the historical playbook suggests.

About this case study

This case study is illustrative and is based on a real Prévenance engagement. The client is not named in accordance with the confidentiality terms of the engagement. Identifying details — including sector specifics, identifying personnel, and certain financial markers — have been altered to protect client confidentiality. The engagement structure, methodology, timeline, and reported order-of-magnitude outcomes reflect the actual engagement. Named references can be arranged under NDA for qualified prospects during the engagement scoping process.

If this is relevant to a PortCo in your portfolio

Prévenance supports international expansion engagements typically structured as 8-week Sprints (positioning, ICP refinement, market architecture) followed by 12–18-month Retainers (live operation, scaling, in-market presence build).

To discuss: contact@prevenance.ai

To learn more: prevenance.ai

PREVENANCE