PORTFOLIO CASE STUDY

ANONYMISED / CONFIDENTIAL DATA REDACTED

Autonomous Finance Platform

Product Strategy & Vision

Rafael Pardo, Product Manager
Treasury & Cash Automation

Original document: January 2026

Anonymised for portfolio use. All company-identifying information, financial figures, and individual names have been redacted.

Context: This document was authored as the product strategy for a treasury automation initiative at a European HR-tech unicorn (€2.3B valuation, 10 countries, €2B+ annual receivables). It represents the 95% final version, completed just before the author's departure from the company. The strategy was never executed.

The Challenge

The company's finance operations were hitting a scalability ceiling. With €2B+ in annual revenue across 10 countries, treasury and collections were managed through fragmented ERPs, 20+ Google Sheets, and manual processes that created audit risk, limited visibility, and consumed senior finance time on low-value tasks.

DSO lagged industry leaders (Adecco, Randstad, Manpower operate at 52-54 days). Every day of DSO improvement unlocks significant working capital, critical for covenant compliance and growth funding.

Our Solution

The Autonomous Finance Platform combines HighRadius (market-leading treasury automation) with proprietary AI agents trained on the company's unique staffing data. This isn't a generic implementation; it's a competitive advantage that positions the company as an AI-native tech company.

Metric	Target
DSO Reduction	-5 to -10 days, unlocking significant working capital
Cash Forecast Accuracy	85%+ (from estimated baseline)
Process Automation	50%+ reduction in manual touchpoints
Audit & Compliance Risk	Full audit trail, SOX-ready controls

Timeline & Investment

9-month phased rollout starting with NetSuite countries (5 markets) in Phase 1, then legacy ERPs. First measurable ROI expected by Week 12. Multi-year investment with projected ROI exceeding 100%.

Why Now

- Competitive Gap: Adecco, Randstad, and ManpowerGroup already operate at 52-54 days DSO. Every quarter of delay widens the efficiency gap with industry leaders.
- Audit & Compliance Exposure: Manual processes with 20+ Google Sheets create material audit risk. Each audit cycle without proper controls increases exposure.
- Scalability Ceiling: Current processes don't scale with growth. Adding countries or volume requires proportional headcount without automation, eroding unit

economics.

- Covenant Pressure: Securitisation facilities require ongoing DSO performance. Proactive improvement provides buffer; reactive improvement after covenant stress is harder and more expensive.

Product Strategy Overview

Product Ownership & Accountability

The Product Owner is the CEO of this product, accountable for its success or failure. This means owning the outcome, not coordinating activities. If this platform doesn't deliver the committed business results, that failure belongs to the Product Owner.

Success Definition

This initiative succeeds if and only if:

- DSO improves by 5+ days within 6 months of Phase 1 go-live (measurable, time-bound)
- Cash forecast accuracy reaches 85%+ within 4 months (from estimated baseline)
- Finance team manual touchpoints reduced by 30%+ within 3 months of adoption
- User adoption reaches 90%+ of target users actively using the platform weekly

Strategic Context

The company operates across 10 countries with fragmented finance systems, manual processes, and inconsistent data quality. With a negative working capital structure and active securitisation facilities, optimising receivables management is critical to maintaining cash covenants and funding growth.

Approach

Implement HighRadius as the foundation layer for Treasury automation, with AI agents built on top to deliver differentiated automation capabilities. Sequencing prioritises HighRadius go-live (Months 1-4) followed by AI agent development (Months 3-6), with initial AI prototypes running in parallel.

PRODUCT VISION

Vision Statement

For	Finance teams across 10 operating countries and HQ Treasury & Collections
Who Struggle with	Fragmented systems, manual reconciliation, inconsistent collection practices, and limited visibility into cash positions and receivables risk
The Platform is	An integrated treasury automation ecosystem
That Provides	A single source of truth for receivables, automated collection workflows, intelligent cash forecasting, and AI-driven credit decisioning, enabling Finance teams to shift from manual execution to strategic oversight
Unlike	The current state of fragmented ERPs (NetSuite in 5 countries, legacy systems in 5), Google Sheets-based workflows, and manual dunning processes
Our Platform	Combines enterprise-grade HighRadius capabilities with custom AI agents tailored to the multi-country staffing business model, securitisation requirements, and growth trajectory

Strategic Pillars

- Unified Data Foundation: Single source of truth for receivables across all ERPs, eliminating reconciliation overhead and enabling real-time visibility
- Intelligent Automation: Replace manual touchpoints with automated workflows, from dunning to cash application, with full audit trail
- Predictive Insights: ML-powered forecasting and payment prediction to enable proactive cash management and risk mitigation
- Scalable Platform: Reusable components and standardised playbooks that scale with geographic expansion and business growth

Competitive Differentiation Architecture

Why can't competitors easily replicate this? The structural advantage comes from the combination of three elements that create compounding value:

Advantage Layer	What We Build	Why It's Defensible
Staffing-Native Data Model	Custom ontology for temp staffing: worker cycles, client payment patterns, seasonal demand, invoice disputes specific to HR-tech	HighRadius out-of-the-box doesn't understand staffing. Competitors would need 12-18 months + domain expertise to replicate
Multi-ERP Orchestration	Single abstraction layer across NetSuite (5 countries) + legacy systems (5 countries) with real-time sync and conflict resolution	Integration debt becomes a moat: solving the hard problem competitors avoid
AI Agents on Proprietary Data	ML models trained on historical payment patterns, dunning effectiveness, and seasonal cash flows across 10 markets	Data exhaust from 10 countries + 7 years of operations creates training data that no competitor can access

The Compounding Effect

Unlike buying HighRadius alone (commodity) or building custom (slow), this architecture creates a flywheel:

- More transactions > Better AI predictions > Higher automation rate > More transactions processed > Richer data for AI
- Each new country added strengthens the model (cross-market patterns) rather than creating isolated silos
- Time advantage: 6-month head start on AI agents while competitors are still implementing vanilla HighRadius

Investment & Expected Returns

Investment figures redacted for confidentiality. The business case demonstrated strong ROI:

Benefit Category	Confidence
Working Capital Release (DSO -5 days)	High
Bad Debt Reduction	Medium
FTE Efficiency Gains (30% automation)	High
Audit & Compliance Risk Reduction	Medium
Interest Savings (Improved Cash Position)	High

Cost of Inaction

Maintaining the status quo carries significant and growing risks:

- Covenant Risk: Securitisation facilities require accurate, timely reporting. Manual processes increase the risk of covenant breaches or reporting errors
- Audit Exposure: Spreadsheet-based processes lack audit trails. As the company scales, auditors increasingly flag control deficiencies
- Scalability Ceiling: Current processes require linear headcount growth. Each new country adds complexity without economies of scale
- Opportunity Cost: Finance teams spend 40%+ time on manual tasks vs. strategic analysis
- Competitive Gap: Industry leaders (Adecco, Randstad) operate at 52-54 day DSO with automated processes

User Personas

Persona 1: Local Finance Lead

Role: Oversees Treasury, Collections, and Credit for a country operation. Reports to local MD and dotted-line to HQ Treasury. Manages Finance Analysts who execute day-to-day operations.

Tools: NetSuite, Google Sheets (dashboards and reporting), email

Pain Points:

- "I don't have real-time visibility into what my team is working on until they escalate issues"
- "When a major client escalates to me, I have to ask my analysts for context; there's no single view of the relationship"
- "HQ asks for reports in formats that don't match how our local ERP exports data"

Persona 2: HQ Treasury Analyst

Role: Consolidates cash positions across 10 countries, manages securitisation reporting, supports investor/bank communications.

Tools: Redshift (via SQL), Google Sheets, PowerPoint, email

Time Allocation: 50% data consolidation/cleaning, 25% forecasting, 15% reporting, 10% analysis

Pain Points:

- "I can't trust the numbers until I've manually validated them against each country's source"
- "Our 13-week forecast is really just a 2-week forecast with guesswork appended"
- "Securitisation covenant reporting takes 2 days every month because data lives in 15 different places"

Persona 3: Local Finance Analyst

Role: Executes daily reconciliations, customer collections calls, payment processing, and data entry. Primary user of finance systems.

Time Allocation: 40% manual reconciliation, 35% customer calls & dunning, 15% data entry & reporting, 10% escalations

Pain Points:

- "I spend the first 3 hours of every Monday reconciling cash positions across bank accounts"
- "I call the same customers every week because I don't know which ones actually need follow-up"
- "I have to copy-paste data between 5 different spreadsheets just to update one ageing report"

Technical & Organisational Challenges

Technical Challenges

Challenge	Mitigation Strategy
Multi-ERP Integration	Leverage existing Redshift data warehouse + APIs for NetSuite countries; SFTP connectors for legacy ERPs
Data Quality & Consistency	Implement data validation rules in HighRadius; create reconciliation dashboards for exception management
Limited Engineering Bandwidth	Primary reliance on HighRadius Professional Services; Data team for integration support; AI agents built with low-code tools where possible
Legacy ERP Constraints	Prioritise NetSuite countries for Phase 1; develop standardised templates for legacy ERP ingestion

Customer (User) Pain Points

Pain Point	Solution Approach
Manual reconciliation consumes 40%+ of time	Automated cash application and bank reconciliation via HighRadius
Reactive collections process	Proactive dunning workflows with payment prediction scoring
Inconsistent reporting formats	Unified data model with standardised report templates
Limited visibility for HQ	Real-time dashboards with drill-down by country/customer/ageing bucket
No audit trail for decisions	Full workflow history and approval tracking in HighRadius

GTM & Organisational Risks

- Change Management: Local Finance teams may resist centralised tooling. Mitigation: Involve country leads early; demonstrate time savings; maintain local autonomy where appropriate
- Parallel Workload: Teams must run legacy processes during transition. Mitigation: Phased rollout starting with highest-volume countries; extended parallel period
- Vendor Dependency: Heavy reliance on HighRadius PS. Mitigation: Build internal expertise through knowledge transfer; document all configurations

APPROACH

Platform Architecture

The Autonomous Finance Platform consists of three layers:

- Data Foundation Layer: Redshift data warehouse as the single source of truth, fed by ERP integrations (APIs for NetSuite, SFTP for legacy systems). HighRadius syncs with Redshift for unified receivables data.
- Automation Layer (HighRadius): Core modules for Cash Forecasting, Collections Management, and Credit Risk. This layer handles 80% of use cases out-of-the-box with configuration.
- AI Differentiation Layer: Custom AI agents for company-specific use cases: staffing industry payment patterns, multi-country dunning optimisation, securitisation portfolio management. Built in-house leveraging LLM capabilities.

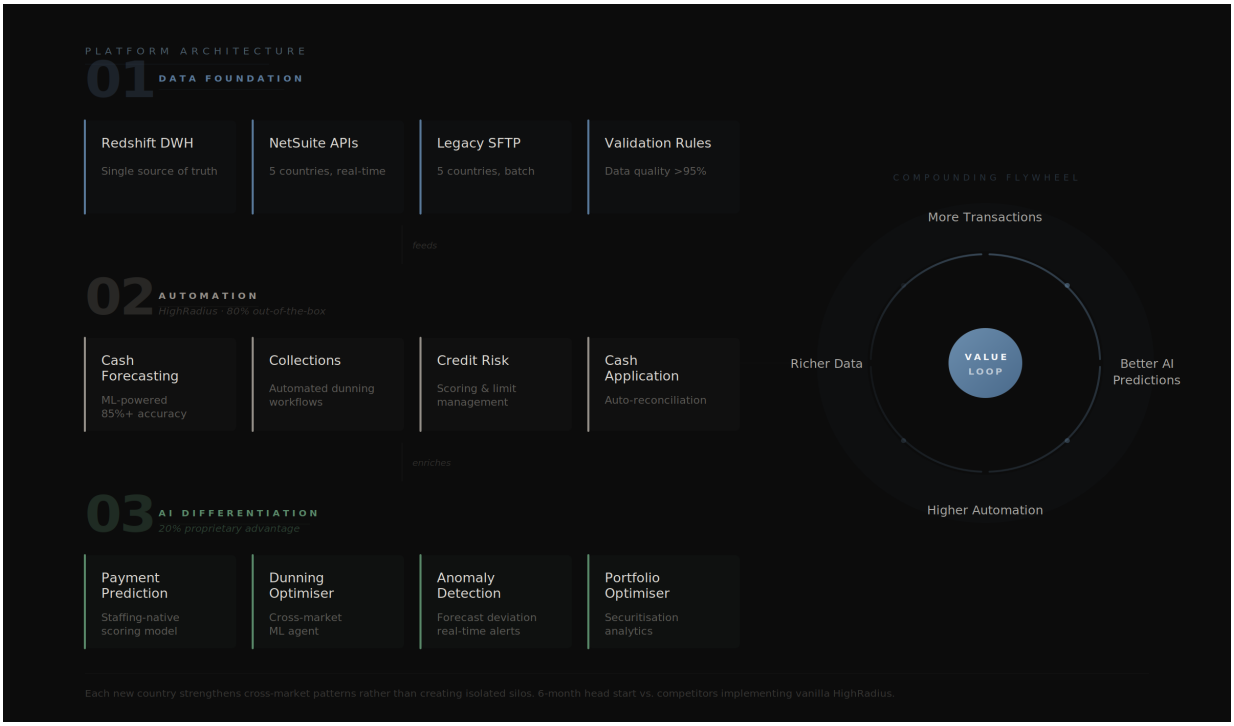

Phased Rollout Strategy

Phase	Scope	Countries	Timeline
Phase 1	HighRadius Core (Forecasting + Collections)	NetSuite countries (5)	Months 1-4
Phase 2	Legacy ERP Integration + AI Agent MVP	Legacy ERP countries (3)	Months 4-6
Phase 3	Full Platform Rollout + Advanced AI	Remaining countries (2)	Months 6-9

Do's & Don'ts

Do	Don't
Start with highest-volume NetSuite countries	Attempt "big bang" rollout across all 10 countries
Prioritise HighRadius go-live before custom AI development	Build AI agents without a stable data foundation
Leverage HighRadius PS for heavy lifting; build internal expertise	Create dependency on external resources without knowledge transfer
Validate each phase with user feedback before expanding	Force adoption without addressing local team concerns
Build reusable components and templates from Day 1	Create country-specific one-off solutions

Milestones & Time-to-Value

Month	Milestone	Deliverable	Success Criteria
M1	Project Kickoff & Discovery	Requirements docs, data mapping	Sign-off from Finance leads
M2	Data Integration Phase 1	NetSuite > HighRadius connectors live	Real-time data sync validated
M3	HighRadius UAT	Forecasting & Collections in UAT	User acceptance sign-off
M3	AI Agent MVP	First AI agent prototype	Demo to stakeholders
M4	Go-Live Phase 1	HighRadius live for 3 markets	Parallel run successful
M5	Optimisation & Legacy	Legacy ERP connectors; workflow tuning	Data quality >95%
M6	Phase 2 Go-Live + AI	3 more countries live; AI in production	Metrics improvements validated
M7-9	Full Rollout	All 10 countries; advanced AI	Platform targets achieved

Time-to-Value Milestones

Enterprise buyers expect positive ROI within 3 months (a16z research: 57% of buyers). Value delivery designed to prove impact early:

Week	Value Delivered	Measurable Outcome
W2	First dashboard live with real-time AR ageing	Finance leads see consolidated view (vs. 4hr manual process)
W4	Automated daily cash position report	Treasury saves 1hr/day on manual consolidation
W6	First AI-prioritised collection list	Collectors work highest-probability accounts first
W8	Cash forecast accuracy baseline	First measurable comparison: AI vs. manual forecast
W12	Phase 1 Go-Live: Full automation for 5 countries	DSO improvement measurable; ROI calculation possible

North Star Metric

> Total Euro Value Influenced by AI-Assisted Actions
>
> This single metric captures the platform's business impact. It compounds as each component improves.

Formula Component	Owner	Lever
# Active Users on Platform	Change Management	Training, UX, adoption campaigns
x Tasks per User (weekly)	Product Team	Workflow coverage, feature depth
x AI Resolution Rate (%)	Product Team	Model accuracy, automation scope
x Average Euro Value per Task	Customer Success	Enterprise expansion, high-value use cases

Example: 30 users x 50 tasks/week x 40% AI resolution x 500 avg value = 300K/week influenced by platform

Key Performance Indicators

Metric	Target	Measurement Method
Days Sales Outstanding (DSO)	-5 to -10 days	HighRadius DSO dashboard
Cash Forecast Accuracy	85%+	Forecast vs. actual variance
Manual Touchpoint Reduction	<50%	Process audit/time tracking
Collection Rate (30-day)	+5pp	HighRadius Collections module
User Adoption Rate	>90%	HighRadius login/activity metrics
AI Automation Rate	50%+ for target processes	Agent execution logs

Adoption Strategy

Core Message: "AI handles the repetitive work so you can focus on judgment calls and client relationships."

Concern	How We Address It
"Will AI replace my job?"	AI automates data consolidation and routine follow-ups. It doesn't replace relationship management, negotiation, and judgment. Your role evolves from data entry to strategic decision-making.
"I don't understand AI"	You don't need to. The platform surfaces recommendations in plain language. Training focuses on using the tool, not understanding algorithms.
"My process works fine"	We're not changing your process. We're removing the manual steps that slow it down. You still make the decisions; you just have better information faster.
"What if AI makes mistakes?"	All AI recommendations require human approval for high-stakes actions. The system learns from corrections and improves over time.

Resistance Early Warning System

Adoption signals monitored weekly to catch resistance before it becomes entrenched:

Signal	Threshold	Response
Login frequency	<3 logins/week per user	1:1 with user to understand blockers
Training completion	<80% by go-live	Manager involvement; additional sessions
Feature usage	Key features unused after 2 weeks	Targeted micro-training; workflow review
Feedback sentiment	Negative trend in surveys	Country Champion escalation; PM intervention
Shadow processes	Teams maintaining parallel spreadsheets	Root cause analysis; process redesign if needed

Quick Wins to Build Momentum

- Week 2: Real-time AR dashboard eliminates 4-hour manual consolidation, immediate time savings visible to all
- Week 4: Automated daily cash position report, Treasury team testimonial shared in Finance newsletter
- Week 8: First AI prediction outperforms manual forecast, concrete proof point for skeptics

- Month 4: Phase 1 country presents results at Finance All-Hands, peer validation more powerful than top-down mandate

AI Governance & Trust Framework

Enterprise finance requires the highest standards of AI governance. Trust in AI systems means demonstrating the work is done correctly and showing how.

Human-in-the-Loop Requirements

Decision Type	AI Role	Human Approval Required
Payment prediction scoring	Recommends priority order	No, informational only
Dunning email content	Generates draft	Yes, before sending (Phase 1)
Cash forecast adjustment	Suggests corrections	Yes, Treasury sign-off
Credit limit changes	Recommends increase/decrease	Yes, Credit Manager approval
Escalation routing	Assigns priority tier	No, can be overridden

Audit Trail & Explainability

- Complete Logging: Every AI decision logged with timestamp, input data, model version, output, and confidence score
- Explainable Outputs: Payment predictions include top 3 factors driving the score (e.g., 'payment history', 'ageing bucket', 'seasonal pattern')
- Version Control: All model versions tracked; ability to roll back to previous version within 24 hours
- Retention Policy: Decision logs retained for 7 years to meet audit and regulatory requirements

AI Risk Mitigation

Risk	Mitigation	Monitoring
Hallucination in financial data	AI never creates data; only analyses existing records from source systems	Reconciliation checks against ERP totals
Model drift over time	Monthly accuracy reviews; automated alerts if performance drops >5%	Eval suite runs weekly against holdout data
Bias in collection prioritisation	Regular fairness audits across customer segments	Demographic disparity analysis quarterly

Risk	Mitigation	Monitoring
Over-reliance on automation	Gradual automation increase; mandatory human review for high-value decisions	Track override rates and outcomes

Scope Definition

In Scope

- HighRadius implementation: Cash Forecasting, Collections Management, Credit Risk modules
- Integration with NetSuite (5 countries) and legacy ERPs (5 countries)
- AI agents for payment prediction, dunning optimisation, and cash forecast anomaly detection
- Dashboards and reporting for HQ Treasury and local Finance teams
- Training and change management for end users

Out of Scope (Explicit Exclusions)

- ERP Replacement: integrates with existing ERPs, does not replace them
- Accounts Payable: focus is on receivables and collections
- Payroll Integration: focus is on client receivables
- Bank Account Management: TMS remains separate
- Securitisation Deal Structuring: project supports covenant reporting only
- New Country Expansion: scope covers existing 10 countries
- Custom ERP Development: no modifications to source ERP systems

Assumptions & Risk Register

Risk / Assumption	Probability	Impact	Mitigation
HighRadius can integrate with legacy ERPs via SFTP	Medium	High	PoC validation before full commitment
Data Team can support integration work in Months 2-3	Medium	High	Pre-align with Data leadership; define scope early
Local Finance teams will adopt new workflows	Medium	Medium	Change management plan; involve users early
AI agents can deliver value without ML engineering hire	Medium	Medium	Start rule-based; use LLM APIs; partner externally if needed
Budget approved as proposed	Medium	High	Phased investment; demonstrate ROI at each stage

Risk / Assumption	Probability	Impact	Mitigation
No major org restructuring during implementation	Low	High	Maintain exec sponsorship; modular architecture

APPENDIX

AI Agent Use Cases (Prioritised)

Use Case	Priority	Complexity	Timeline
Payment Prediction Scoring	P0	Medium	M3-4
Automated Dunning Optimisation	P0	Medium	M4-5
Cash Forecast Anomaly Detection	P1	Low	M5-6
Credit Limit Recommendation	P1	High	M6-7
Securitisation Portfolio Optimiser	P2	High	M7-9

Governance & RACI

Activity	Product Manager	Treasury Director	CFOO	Country Leads
Roadmap & Prioritisation	A/R	C	I	C
Budget Approval	R	A	A	I
Vendor Selection	R	A	I	C
Technical Architecture	A/R	I	I	C
Go-Live Decisions	R	A	I	C
Change Management	R	C	I	A
Escalation Resolution	R	A	A	C

R = Responsible, A = Accountable, C = Consulted, I = Informed

AUTHOR'S REFLECTION

Why This Document Matters

This strategy was designed for execution. Organisational changes prevented implementation, but the thinking behind it represents something I believe is rare and increasingly valuable: the perspective of someone who has operated the systems that AI is about to replace.

I spent seven years inside treasury operations at scale. I reconciled cash positions across 10 countries. I built the unit economics models that went to the board. I watched senior finance professionals spend 40% of their time on tasks that a well-designed agent could handle in seconds. That experience taught me something that most AI product managers don't have: an intuitive understanding of where automation breaks down, where human judgment remains essential, and where the real value sits.

The answer, consistently, is in the last 30%.

The first 70% of finance automation is straightforward: rules-based workflows, standard reconciliations, templated reporting. Any competent implementation of HighRadius or a similar platform handles this. The remaining 30%, the exceptions, the edge cases, the judgment calls that require context across markets and counterparties, is where autonomous systems either create genuine value or become expensive mistakes.

This document reflects that belief. The AI Governance framework isn't theoretical. It comes from watching what happens when financial systems produce outputs that humans trust without verification. The change management strategy isn't generic. It comes from seeing adoption fail when tools are imposed rather than co-designed with the people who use them. The competitive differentiation architecture isn't an exercise. It comes from understanding that proprietary data and domain-specific ontologies are the only sustainable moats in a world where foundation models commoditise general intelligence.

I believe treasury will be one of the first enterprise functions to become fully autonomous. Not because the technology is easy, but because the data is structured, the rules are well-defined, and the cost of manual operations is measurable and growing. The transition from manual to autonomous won't be driven by AI researchers. It will be driven by operators who understand both the domain and the technology deeply enough to design systems that finance teams actually trust.

That is the role I intend to play.

Rafael Pardo
rafael@rpardo.ai / rpardo.ai