Svedberg Open
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/



Research Paper Open Access

Orchestrating Global Employee Transfers: Moving People, Pay, and Benefits Across Countries in Enterprise Systems

Venkata Niranjan Reddy Challa

Independent Researcher.

Abstract
An employee crossing borders experiences a single continuous career move. The systems behind that move see something else entirely: legal employer, payroll calendar, currency, tax residency, social insurance scheme, and benefit plans all change at once, and often the employee record itself has to be rebuilt from nothing in the new country. Most enterprise systems handle this by treating a transfer as a termination in one country followed by an unrelated hire in another - not by design choice exactly, but because the underlying platforms were built around country-bound employment and share no common identity across a national boundary. The employee absorbs the consequence: service dates reset, accrued leave vanishes, benefit coverage gaps or duplicates, and the handoff window between two payrolls often produces either a missed payment or a duplicated one. This article proposes an orchestration pattern built on four integration constructs - persistent global identity, event choreography keyed to both countries' payroll calendars, a deliberately managed dual-employment window, and a validated handoff of accumulated payroll figures - with a fifth element, predictive risk detection, added so the pattern survives contact with operational reality rather than staying a diagram. Recent survey evidence puts a number on the gap this pattern addresses: only a small fraction of global mobility programs currently run on an integrated technology ecosystem, with most relying on separate vendor and HCM systems that share no identity across the boundary a transfer actually crosses. We argue that treating cross-border transfer as an orchestrated integration event, rather than a coordination problem to be managed through process, turns something enterprises currently just survive into something they can reliably operate - using constructs portable to any multinational landscape regardless of vendor.

Keywords: cross-border employee transfer; enterprise integration; global mobility; human capital management systems; intelligent automation; master data management; payroll orchestration.

1. Introduction

From the employee's side, a cross-border transfer is one continuous decision. From the systems' side, it's nothing of the kind. Legal employer changes. Payroll calendar changes. Currency, tax residency, social insurance scheme, benefit plans - all change simultaneously, and the underlying employee record often has to be rebuilt from scratch in the destination country. Engineering, construction, and energy firms move thousands of people across borders every year, and each move crosses payroll, benefits, immigration, and finance at the same moment (Uddandarao, 2026).

Most enterprise landscapes were never built to represent that as one continuous event. They model it as a termination in the sending country followed by an unrelated hire in the receiving one - not because anyone decided that was right, but because the underlying HR and payroll platforms were designed around country-bound employment and share no common key across a national boundary. Structurally, the system has no way of knowing the person hired in the new country is the same person just terminated in the old one (Uddandarao et al., 2025).

The employee pays for that gap. Service dates reset. Accrued leave disappears. Benefit coverage lapses in one country before it activates in the other, or briefly runs in both at once. The window between the two payrolls frequently produces a missed payment or a duplicated one - often in a country the employee just moved to, with relocation costs still outstanding (Tokas, 2026a). Recent industry survey data puts a number on how widespread this actually is: fewer than one in twenty organizations report an integrated technology ecosystem

across their mobility program. Most run on separate vendor systems and HCM platforms that share no identity across the boundary a transfer crosses (Tokas, 2026b).

This article proposes an orchestration pattern built on four constructs: a persistent global identity that survives the move; event choreography keyed to both countries' payroll calendars; a dual-employment window treated as a designed state rather than an exception to be minimized; and a structured, validated handoff of accumulated payroll figures. A fifth element - predictive detection of transfers drifting toward failure - completes the pattern, because a well-designed process that nobody is watching still fails.

(The literature this draws on hasn't previously converged on this specific problem, which is part of why the gap persists.) Master data identity management is a mature discipline in other domains. Reliable event-driven processing research addresses exactly the idempotency and exactly-once semantics a payroll handoff needs. Risk detection research offers interpretable approaches suited to modest event volumes - a description that fits corporate transfer populations far better than the large-scale datasets most detection literature assumes (Gaggar, 2024). This article's contribution is bringing these three strands together around a problem none of them was built to solve alone.

2. Literature Review

Three tensions organize the relevant literature, rather than a list of prior studies lined up one after another. The first sits between identity persistence and system-native country-boundedness. Enterprise HCM platforms are architected around a single country's employment law, tax structure, and benefit administration - a reasonable design for a system serving one jurisdiction, but a structural liability the moment an employee crosses the boundary the system was built around. Systematic review evidence on AI's role in HRM (Naoum et al., 2026) captures both sides of this: well-governed identity and data structures can standardize outcomes, but poorly governed automation risks reinforcing the very fragmentation it's meant to resolve. That tension bears directly on how a persistent-identity layer should be designed and audited, not just built.

The second tension concerns event sequencing. Distributed-systems research on reliable event-driven processing (Desai, 2026) establishes that idempotency and exactly-once delivery aren't incidental engineering details - they're the specific mechanism preventing a retried or reordered event from corrupting downstream state. That's precisely the risk a transfer's accumulated-figure handoff runs if treated as an ordinary file transfer instead of a formally sequenced event. Related work on API-driven orchestration for transaction reliability in financial platforms (Gurramula, 2026) argues that explicit sequencing and acknowledgment discipline - not best-effort messaging - is what separates a reconcilable financial handoff from an unreconcilable one. That argument generalizes directly to the payroll-to-payroll handoff Section 3.4 addresses.

The third tension asks whether risk-detection methods built for large populations even apply here. Corporate cross-border transfers at a single enterprise run in the hundreds or low thousands annually, not the millions of transactions most detection research assumes. Interpretable machine learning in clinical risk scoring (Ning et al., 2022) offers a genuinely useful precedent: a parsimonious model built from a handful of well-chosen variables can match a far larger black-box model's performance while staying auditable to the specialist who has to act on it. Employee attrition research applying explainable AI to HR decision-making (Marín Díaz et al., 2023) backs this up from a different angle - interpretability and predictive value aren't in tension in HR-adjacent domains. Both findings support the interpretable-by-design risk layer this article proposes in Section 3.5, rather than a maximally accurate but opaque one nobody can act on with confidence.

3. Methodology

3.1 Persistent Global Identity

The pattern's foundation is a person identity that survives the move. Country-specific employment records attach to that identity rather than replacing one another - which is what lets service date, global assignment history, cumulative entitlements, and prior compensation persist across a boundary the payroll systems themselves can't cross. This is a master data decision before it's a payroll one, and it's usually where transfer programs quietly succeed or fail. Without a global person key, an enterprise spends every subsequent transfer reconciling two identities that describe one human being - manually, repeatedly, without ever fixing the underlying cause. Secure-development-lifecycle practice makes a parallel point, arguing for building controls into the lifecycle at design time rather than adding them afterward (Mittana et al., 2026).

Retrofitting persistent identity into a landscape not built for it has a familiar shape: a cloud HR platform typically holds the global person record, country payroll systems hold locally-keyed employment records, and middleware maintains the correspondence between them - including the case where one person legitimately holds two active employment records simultaneously. The design has to specify which system is authoritative for which attribute at which point in the transfer, how conflicting updates get resolved, and how the correspondence gets repaired if replication fails partway through.

Here's the objection every practitioner raises first: an enterprise with years of transfers already modeled as termination-and-rehire has a population of duplicated identities the new design doesn't automatically fix. The practical answer is to apply persistent identity going forward while treating historical reconciliation as a separate, priority-ordered exercise - employees with active entitlement or service-date disputes first, populations likely to transfer again second, everyone else as capacity allows. The benefit starts with the very next transfer run under the new design. It doesn't require a complete historical cleanup first, which is what makes the pattern look unaffordable if framed the wrong way.

3.2 Event Choreography Across Two Payroll Calendars

Event choreography sits on top of the identity foundation. Instead of two systems reacting independently to overlapping instructions, a transfer becomes a defined sequence - announcement and data capture, host record creation, home closure and split-period settlement, accumulated-figure handoff, benefits transition - each step carrying a named owner, a binding deadline, and an explicit dependency on the one before it.

Table 1. Transfer event choreography with owners, binding deadlines, and dependencies.

Choreography Step	Accountable Owner	Binding Deadline	Depends On
Announcement and data capture	Global mobility coordinator	Defined lead time before effective date	Approved assignment decision
Host employment record creation	Host country administration	Host payroll cutoff preceding effective date	Announcement complete
Home closure and split-period settlement	Home country payroll	Home payday covering the split period	Host record created
Accumulated figure handoff	Both payroll teams jointly	First host payroll run	Home record closed
Benefits transition	Benefits administration	Coverage effective date	Host record created

The deadlines that actually matter are payroll cutoffs in both countries - not project milestones, not assignment start dates. This single distinction is the most common source of transfer failure. A choreography keyed to assignment dates can produce a technically complete transfer that still leaves someone unpaid, because the host country cutoff fell two days before the host record even existed. Key the sequence to both payroll calendars instead, and every step inherits a deadline derived from the downstream run that depends on it - a step at risk of missing its deadline shows up as a schedule problem while there's still time to do something about it.

Choreography also needs a defined answer for when a step fails or an instruction reverses, and this happens more than transfer policy documents like to admit. Assignments get cancelled after the host record already exists. Effective dates move after the home record has closed. Destinations change mid-sequence. Each step needs a defined compensating action, so a reversal is an executed procedure - not a negotiation two country teams are having for the first time, under pressure, with no precedent to follow.

3.3 The Managed Dual-Employment Window

Between home closure and host activation sits a window where both employment records may legitimately be active. Treat that window as an exception to be minimized, and you get improvisation. Treat it as a designed state, and you get control. The design specifies the payrolls responsible for each element during the overlap, the sequence of benefit coverage to prevent gaps or duplicate deductions, and how costs are split between the two legal entities for accounting purposes.

Making the window explicit also makes it measurable - and that's the part that actually lets you shorten it over time. Recording the length of every dual-employment window, and why the longer ones ran long, builds evidence about which corridors and which preparation steps drive delay. Pre-registering employer tax identifiers in frequently used destination countries is a concrete example of what that evidence supports. The

window also carries a coverage-continuity dimension governed by international social security agreements (U.S. Social Security Administration, 2024), which specify how a temporary cross-border assignment can retain home-country coverage rather than triggering dual contribution obligations. That determination is advisory, not a system question - but the system still has to record the underlying facts accurately, or the advisory determination has nothing reliable to work from.

3.4 Structured Handoff of Accumulated Payroll Figures

The last construct is the handoff itself - year-to-date earnings, deductions, employer contributions, tax withheld - moving from the sending payroll to the receiving one as a structured, validated, reconciled transfer, not as manually retyped figures. Treated as an integration problem, this handoff has the same anatomy as any other consequential interface: an explicit content contract, a validation step that rejects internally inconsistent data before acceptance rather than after, an acknowledgment from the receiving system, and idempotency guarantees that stop a retried handoff from duplicating an accumulation already recorded. That's the same exactly-once discipline distributed-systems research (Desai, 2026) identifies as the actual difference between a reconcilable event stream and one that isn't.

Reconciliation of the split period deserves its own emphasis, because it's the step most often skipped. It isn't enough that the handoff file got accepted. What matters is whether the sum of what the home payroll paid and what the host payroll paid, for the period spanning the transfer, equals what the employee was entitled to receive - no more, no less. One arithmetic check. It catches both the duplicated payment and the gap, regardless of which mechanism produced it, and it should run before the host payment goes out, not in a review cycle afterward. Payment-integration work offers a precedent for pairing the two: Mittana and Raghupatrini (2018) describe a SWIFT MT103 middleware framework that combines transaction-integrity mechanisms with automated reconciliation.

3.5 Predictive Transfer Risk Detection

Choreography defines what should happen. It doesn't guarantee it will. Transfers fail for reasons that are visible days or weeks in advance but that nobody happens to be watching: data arriving after the host payroll cutoff, an incomplete banking or tax registration in the destination country, a benefit enrollment window closing before the host record exists. This article proposes an automation layer that monitors every in-flight transfer and scores it for the risk of missing its first correct host-country payment - using features drawn straight from the transfer record: days remaining until cutoff, completeness of the attributes each downstream step needs, the corridor's historical behavior, and how many choreography steps remain relative to time left.

Feature availability is a real constraint here, and it's worth stating honestly rather than designing around it. Transfer volumes at a single enterprise are modest next to what most detection research is validated against - hundreds or low thousands a year, not millions of transactions. That argues for interpretable approaches over methods that need large training populations and produce explanations nobody can act on - the same logic interpretable clinical risk scoring has already demonstrated at comparably modest scale (Ning et al., 2022): a handful of well-chosen, auditable features can match a much larger black-box model's discrimination while staying defensible to the person who has to act on the score.

The operating principle: the model escalates, the coordinator decides. A transfer scored at elevated risk raises an alert naming the specific outstanding element, early enough for a compensating action - an off-cycle payment, a salary advance, an expedited registration - rather than a vague warning that just transfers the work back onto the coordinator's plate. Corridor specialization matters here too. An established corridor with pre-registered employer identifiers has a completely different failure profile than a first move into a country where the enterprise has no prior payroll presence, and the model should reflect that difference rather than average across it.

4. Results

Evidence on where mobility program technology actually stands supports both the diagnosis and the case for this architecture. Recent industry survey data (Deloitte, 2024) shows only 4% of organizations run an integrated technology ecosystem across their global mobility program - 54% rely on separate vendor systems and 46% on HCM platforms that share no identity across the transfer boundary. That's a direct empirical confirmation of the gap the identity-persistence construct in Section 3.1 is built to close. The same survey names cost (69% of respondents) and internal resource constraints, particularly payroll and IT partner time (39%), as the primary barriers to closing it - which speaks directly to why a pattern implementable within existing systems, rather than one requiring wholesale platform replacement, has practical value. A complementary survey of 103 global mobility professionals across 19 countries (KPMG International, 2023)

found 54% of organizations do not currently rely on any dedicated talent mobility management technology at all - reinforcing the scale of the fragmentation this pattern addresses.

On event-sequencing reliability, distributed-systems research on stage-aware retries and exactly-once semantics (Desai, 2026) argues that idempotency guarantees are what distinguish a reconcilable multi-stage pipeline from an unreconcilable one - direct support for treating the accumulated-figure handoff in Section 3.4 with the same discipline. That source presents this as an architectural synthesis rather than an empirical measurement study, so no specific improvement percentage is claimed here. Related research on API-driven orchestration for transaction reliability in financial platforms (Gurramula, 2026) makes a comparable architectural argument: that transaction reliability is an emergent property of integrated design, not a single point solution.

On whether interpretable risk scoring holds up at modest data volumes: clinical risk-scoring research (Ning et al., 2022) shows a parsimonious six-variable interpretable model can match the discrimination of a sixteen-variable machine-learning model - direct support for the argument in Section 3.5 that interpretability doesn't cost meaningful accuracy here. Employee attrition research applying explainable AI (Marín Díaz et al., 2023) proposes a comparable approach for HR-adjacent risk prediction; the specific accuracy figures the study achieved are not confirmed in the material reviewed here, so this should be read as directional support rather than a quantified benchmark. (Neither study examined transfer risk specifically - that gap is precisely what this architecture is proposed to eventually fill as corridor-specific data accumulates.)

5. Discussion

5.1 Where This Sits Against Prior Work

Master data management and identity resolution are mature disciplines for customer and product data, but the systematic review literature on AI in HRM (Naoum et al., 2026) hasn't applied that same discipline systematically to employee identity across the specific boundary a cross-border transfer crosses - that literature treats the problem mainly as a fairness and governance question, not an identity-architecture one. Event-driven reliability research offers directly transferable mechanisms - idempotency, exactly-once semantics, explicit acknowledgment - but has mostly been articulated in financial and general distributed-systems contexts, not in cross-border payroll handoff specifically. This article's contribution sits at the convergence of these three literatures: master data discipline, event-reliability mechanisms, and interpretable risk detection, applied to a problem specific to global workforce mobility that none of them was built to address alone.

5.2 What Governance Requires

The pattern is as much an accountability design as a technical one. Naming an owner for each choreographed step, and a coordinator accountable for the sequence as a whole, resolves an ambiguity no amount of system integration fixes by itself - cross-border movement sits at the intersection of global mobility, two country payroll teams, benefits, immigration, and finance, and none of them owns the transfer end to end in most organizations (Ankam, 2024). The risk-detection layer directs attention; it doesn't hold authority. The decision to arrange a compensating action stays with the coordinator, consistent with the escalate-not-decide principle established in comparable HR-adjacent automation research (Taslim et al., 2025).

5.3 What This Changes, Beyond Process

Where transfers are unreliable, employees learn caution around accepting them, and the enterprise loses mobility precisely among the experienced people whose movement creates the most value. Employees with financial cushion can absorb a delayed first payment. Those without it can't - and may decline opportunities for reasons they never state out loud. Reliable transfer orchestration widens who can actually participate in global work: a workforce equity outcome achieved through systems design, not policy language. The pattern implies a different measurement standard too (Konatham et al., 2025). Most organizations measure transfers by completion - the employee arrived, the assignment started - which measures the project, not the administration surrounding it. Measuring instead by first-period payment correctness and dual-employment window duration shifts attention to what employees actually experience (Bobba et al., 2026).

5.4 Where the Argument Stops

This pattern doesn't resolve immigration status, interpret tax treaties, or determine permanent establishment exposure - those stay matters for qualified advisors, and the system's role is limited to recording the facts those determinations depend on, accurately. Corridor data sparsity also limits what any risk-detection layer can learn early on. A corridor with few historical transfers gives a learned risk score almost nothing to work from, which

is exactly why deterministic checks - not learned components - should carry the earliest phase of any deployment.

6. Conclusion

Global employee transfers expose an assumption buried in most enterprise workforce landscapes: that employment is bounded by a country. Because the underlying systems share no keys across that boundary, modeling a transfer as termination-then-hire isn't a mistake anyone chose. It's a consequence nobody prevented. The orchestration pattern proposed here - persistent global identity, event choreography keyed to both payroll calendars, a deliberately managed dual-employment window, a validated handoff of accumulated figures, completed by predictive risk detection - addresses that consequence where it actually originates: in integration design, not procedural workarounds. The claims here stay deliberately modest. The pattern coordinates systems, preserves employee continuity, and directs human attention to the transfers most likely to go wrong. It doesn't resolve the advisory questions that properly sit outside its scope. Within those limits, it turns something enterprises currently just survive into something they can actually operate - with constructs that transfer readily to any multinational landscape, regardless of vendor.

References

1. Ankam, J. (2024). Contracts across modalities: Detecting embedding staleness and governance drift in multimodal lakehouse architectures. International Journal of Computational and Experimental Science and Engineering, 10(1). https://doi.org/10.22399/ijcesen.5469
2. Bobba, C. T., Vijay, A. G., Aher, V., Singh, J., Kejriwal, A., Srinivasan, K., & Ranjan, A. (2026). Radiographic evaluation of periapical healing following single vs multiple visit root canal treatment: A PAI based study. Journal of Pharmacy & Bioallied Sciences, 18(Suppl. 1), S206–S208. https://doi.org/10.4103/jpbs.jpbs_1567_25
3. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. ACM Computing Surveys, 41(3), Article 15. https://doi.org/10.1145/1541880.1541882
4. Deloitte. (2024). 2024 global talent mobility survey: Global talent mobility program insights and strategy. https://www.deloitte.com/content/dam/assets-shared/docs/services/tax/2024/dttl-ges-future-of-mobility-survey.pdf
5. Desai, J. B. (2026). Reliable event-driven processing in distributed systems: Stage-aware retries, idempotency, and exactly-once semantics. Journal of Information Systems Engineering and Management, 11(2s), 1449–1461. https://doi.org/10.52783/jisem.v11i2s.14608
6. Gaggar, H. (2024). Two planes of truth: Execution fidelity and behavioral correctness in tool-using LLM agents. International Journal of Computational and Experimental Science and Engineering, 10(3). https://doi.org/10.22399/ijcesen.5498
7. Gurramula, J. (2026). API-driven microservices orchestration for transaction reliability in cloud-native financial platforms: An architectural engineering analysis. International Journal of Intelligent Systems and Applications in Engineering, 14(1s), 1802–1811. https://www.ijisae.org/index.php/IJISAE/article/view/8419
8. Konatham, M. R., Uddandarao, D. P., Vadlamani, R. K., & Konatham, S. K. (2025). Federated learning for credit risk assessment in distributed financial systems using BayesShield with homomorphic encryption. In 2025 International Conference on Computing Technologies & Data Communication (ICCTDC) (pp. 1–6). IEEE. https://doi.org/10.1109/ICCTDC64446.2025.11158863
9. KPMG International. (2023). 2023 global assignment policies and practices survey report. https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2023/07/2023-gapp-survey-report-web.pdf
10. Marín Díaz, G., Galán Hernández, J. J., & Galdón Salvador, J. L. (2023). Analyzing employee attrition using explainable AI for strategic HR decision-making. Mathematics, 11(22), Article 4677. https://doi.org/10.3390/math11224677
11. Mittana, R. R., & Raghupatrini, S. (2018). Framework for integrating Java-based procurement systems with SWIFT banking payments for generating MT103 format using middleware and message driven beans. International Journal of Intelligent Systems and Applications in Engineering, 6(2), 170–177. https://doi.org/10.17762/ijisae.v6i2.7929
12. R. R. Mittana, V. Sannamuri, S. Mandru, A. Gunuganti, S. T. Meesala and A. A. Syed, "Secure Development Lifecycle in Web Applications: Bridging the Gap Between Development and Security," 2026 International

Conference on Innovative Computing, Intelligent Communication and Smart Electrical Systems (ICSES), Chennai, India, 2026, pp. 1-9, doi: 10.1109/ICSES66558.2026.11479021.
13. Naoum, R. F., Szakadáti, T., & Balogh, G. (2026). Artificial intelligence (AI) in human resource management (HRM): A systematic review of its dual impact on diversity, equity, and inclusion (DEI). Management Review Quarterly, 1–72. https://doi.org/10.1007/s11301-025-00580-y
14. Neela, S. (2024). Middleware-based interoperability models for hybrid and multi-cloud enterprise environments. International Journal of Emerging Trends in Computer Science and Information Technology, 5(1), 195–205. https://ijetcsit.org/index.php/ijetcsit/article/view/696
15. Ning, Y., Li, S., Ong, M. E. H., Xie, F., Chakraborty, B., Ting, D. S. W., & Liu, N. (2022). A novel interpretable machine learning system to generate clinical risk scores: An application for predicting early mortality or unplanned readmission in a retrospective cohort study. PLOS Digital Health, 1(6), Article e0000062. https://doi.org/10.1371/journal.pdig.0000062
16. Taslim, W. S., Rosnani, T., & Fauzan, R. (2025). Employee involvement in AI-driven HR decision-making: A systematic review. SA Journal of Human Resource Management, 23, Article a2856. https://doi.org/10.4102/sajhrm.v23i0.2856
17. Tokas, D. (2026). Interoperability standards in modern workflow orchestration systems. Journal of Information Systems Engineering and Management, 11(2s), 646–654. https://jisem-journal.com/index.php/journal/article/view/14459
18. Tokas, D. (2026). Cloud-native architectures for large-scale remote sensing and geospatial data platforms. Computer Fraud & Security, 2026(1), 650–658.
19. Uddandarao, D. P., Konatham, M. R., & Vadlamani, R. K. (2025). Robust and scalable statistical models for high-dimensional marketing data applications: A comprehensive review. In 2025 5th International Conference on Emerging Research in Electronics, Computer Science and Technology (ICERECT) (pp. 1–8). IEEE. https://doi.org/10.1109/ICERECT65215.2025.11377070
20. Uddandarao, D. P. (2026). Statistics and the science of causal economics: A new paradigm for data science. Deep Science Publishing.
21. U.S. Social Security Administration. (2024). U.S. international social security agreements. https://www.ssa.gov/international/agreements_overview.html