

Multi-Cloud and Cloud-Native Strategies for Digital Human Capital Management in Saudi Arabia: Building Scalable and Future-Ready HR Technology Platforms

Faraz Jameel Siddiqui

Published October 2026

Abstract

Digital human capital management increasingly depends on interconnected recruitment, payroll, learning, workforce planning and employee-service applications. Yet distributing these functions across cloud providers does not automatically produce integration, resilience or trustworthy personnel decisions. This integrative review examines how multi-cloud and cloud-native strategies can support scalable human resource technology platforms in Saudi Arabia while preserving information accountability and employee confidence. Thirty peer-reviewed publications issued during 2020–2025 are examined across software architecture, electronic human resource management, cyber security and Saudi digital transformation. The synthesis distinguishes application portability from data portability, technical availability from operational continuity, and analytic capability from legitimate data use. It proposes a layered reference architecture combining domain-oriented services, controlled integration, sovereign data stewardship, interoperable identity and observable operations. A staged implementation pathway is developed for organisations modernising heterogeneous enterprise systems, together with indicators for evaluating capability, fairness, resilience and economic sustainability. The evidence suggests that selective workload placement and disciplined platform engineering are more defensible than indiscriminate replication across providers. The principal contribution is an HR-specific governance framework that links technical architectural choices to human capital outcomes and Saudi institutional requirements. The framework remains conceptual and requires empirical evaluation before performance or compliance benefits can be asserted.

Keywords: digital human capital management; multi-cloud; cloud-native architecture; Saudi Arabia; human resource information systems; data governance; Vision 2030

1 Introduction

Human capital functions now process interdependent events rather than isolated administrative transactions. A recruitment decision changes identity records; onboarding creates payroll and access entitlements; learning histories inform mobility; and departures require timely revocation of privileges. As Saudi organisations pursue greater workforce participation, skills development and digital service delivery, inconsistencies between these events threaten both operational performance and employee trust. Research on digital human resources demonstrates that technology adoption is shaped by organisational resources, user capability and perceived usefulness, not software acquisition alone [21, 22]. Consequently, architecture should be evaluated by the quality of decisions and services it enables.

Cloud-native computing offers independently deployable services, automated infrastructure and programmable interfaces. Multi-cloud arrangements add the possibility of using complementary

providers, separating critical workloads and managing concentration risks. However, these concepts are not synonymous: a containerised application can remain tightly coupled to one provider's databases, identity services and monitoring tools. Alonso and colleagues identify unresolved architectural and data-portability questions in genuinely distributed multi-cloud applications [1]. Migration research similarly indicates that moving monolithic software into smaller services often exposes hidden dependencies and quality shortcomings [2, 3]. Such findings caution against presenting multi-cloud adoption as a guaranteed route to resilience.

Human resources constitutes a particularly demanding application domain because employee records combine sensitive identifiers, compensation, performance evaluations, disability-related adjustments, disciplinary material and sometimes predictive assessments. Data are collected for different purposes and held under different retention obligations. A technically successful integration can therefore become a governance failure if it permits unrestricted aggregation, opaque profiling or uncontrolled transfer. Studies of people analytics describe gaps between declared ethical commitments and operational controls [26, 27]. Similarly, the literature on AI-assisted HRM warns that improved automation does not independently demonstrate equitable or accepted outcomes [24, 25].

The Saudi context adds significant strategic relevance. Studies of local cloud adoption identify trust, provider reputation and verifiable service commitments as determinants of acceptance [14, 15]. Research on Saudi digital government further highlights regulatory alignment, leadership support and integration complexity [16, 17]. Vision 2030 gives human capital transformation a broader policy rationale, but neither an international architecture nor a Saudi digital strategy substitutes for organisation-specific legal and operational assessment. Cloud location, contractual terms, processing purpose and applicable cybersecurity controls must be determined for each deployment rather than presumed from provider branding.

Existing scholarship is commonly divided between infrastructure optimisation and HR technology outcomes. Architectural studies discuss resource scheduling, service separation and security, whereas HR studies assess adoption, analytics and employee experiences. Relatively little work specifies how architectural design decisions should be traced to accountable HR outcomes in Saudi organisations. This review bridges that division through a critical synthesis, proposes a testable reference design and identifies conditions under which multi-cloud complexity may outweigh its benefits. It reports no new survey, experiment or observed deployment results.

2 Study aim, objectives and review questions

The study aims to establish an evidence-informed architectural and governance framework for scalable, secure and human-centred digital human capital management in Saudi Arabia. Its contribution is design-oriented: to clarify decisions that organisations can evaluate, rather than to claim that any particular technology stack has already delivered national outcomes.

Four objectives structure the enquiry. First, distinguish multi-cloud deployment models and cloud-native application properties relevant to HR information systems. Second, evaluate integration, identity, data protection, resilience and cost trade-offs across the employee lifecycle. Third, connect architecture choices to digital HR effectiveness, workforce inclusion, institutional trust and Saudi transformation priorities. Fourth, derive an implementation sequence and measurable evaluation criteria for future field research.

Three review questions follow: which architecture characteristics enable reliable HR service composition; what governance mechanisms constrain cross-provider risks without undermining usability; and which evidence and metrics are required to judge adoption success? The questions deliberately concern mechanisms and conditions rather than a single hypothesised positive relationship.

3 Review methodology

An integrative, critical literature review is appropriate because the question crosses empirical HR research, software engineering, information security and national digital policy. The approach permits comparison of heterogeneous study designs while retaining differences in context and evidential strength. It is not a statistical meta-analysis: the selected studies do not share equivalent interventions, populations or outcome metrics, and pooling their numerical results would be misleading.

The review concentrates on 30 identifiable peer-reviewed publications dated 2020–2025, each accompanied by a DOI. Evidence was sought using combinations of the terms multi-cloud, cloud-native, microservices, HR analytics, digital HRM, workforce, employee data, Saudi Arabia, security, governance and migration. Publisher and bibliographic records were consulted to check title, publication year and DOI correspondence. The supplied technical review informed broad thematic organisation but is not used as a substantive reference, preventing the analytical contribution from relying on an unexamined secondary summary.

Eligibility required a direct contribution to one of four domains: distributed application engineering; cloud security and service operations; digital people-management practice; or Saudi organisational and institutional transformation. Studies were excluded where relevance depended solely on generic enthusiasm for technology or where a persistent scholarly identifier could not be established. This is a purposeful, cross-disciplinary synthesis rather than a census of all available work. The manuscript therefore makes no claim about exhaustive retrieval counts or a reproducible screening population.

Each work was read against common analytical dimensions: research setting, evidence type, architectural or organisational mechanism, observed or proposed benefit, limitation, and applicability to sensitive HR workloads. Findings were grouped into workload boundaries, cross-provider integration, platform security, digital employee experience and transformation governance. Contradictory implications were preserved. For instance, replication can reduce dependence on one availability zone yet introduce identity inconsistency and additional personal-data copies. Design propositions are explicitly separated from documented study findings.

Assessment of trustworthiness prioritised peer-reviewed research with intelligible methods and traceable publication details. An article's DOI establishes identification, not methodological quality, and the inclusion of a publication does not certify its relevance to every Saudi HR setting. Policy references are treated as contextual requirements to be checked against current official guidance before implementation. This limitation is especially important where duties vary by organisation, sector, contractual relationship or data classification.

4 Evidence base: from digital HRM to cloud-native platforms

Digital HRM research places organisational alignment before technical sophistication. A meta-analysis links electronic HR adoption to system usefulness, organisational resources, user knowledge

and social influence, with differentiated implications for performance [21]. The broader HRM 4.0 literature identifies digital competencies and human-machine complementarity as essential, rather than peripheral, transformation variables [22]. Bansal and colleagues similarly conceptualise human resource digital transformation as a capability shaped by individual and organisational factors [23]. These contributions imply that infrastructure availability is merely an enabling condition for effective workforce services.

Studies of talent management indicate that stronger digital maturity can support attraction and retention activities, although organisational circumstances mediate those relationships [28]. AI-assisted HRM reviews further distinguish organisational productivity from employee-centred consequences and identify significant variation in outcomes across adoption patterns [24, 25]. Thus, service design must protect human review, explanation and contestability where automated outputs influence careers or remuneration. For architecture teams, these are functional requirements: systems need decision histories, role-sensitive access, override records and explainable data lineage, not only high-throughput databases.

On the engineering side, multi-cloud-native literature differentiates replicated applications from truly distributed component arrangements [1]. Studies of architectural migration show that a service boundary should reflect cohesive business responsibilities and stable interfaces, not arbitrary code fragmentation [2, 3]. A review of microservice principles finds that decoupling offers flexibility but complicates communication and service discovery [4]. Research on customised software services demonstrates practical options for preserving tenant-specific behaviour during migration [5]. For HR platforms, such evidence favours bounded services for recruitment, core personnel records, payroll calculation, learning and case management, integrated through explicit contracts.

Cloud orchestration literature adds a second lesson: service elasticity must be balanced against response-time objectives, energy demand and state persistence [8, 9]. Connectivity across clusters can enable independent services to cooperate, yet network design and segmentation become part of the application itself [7]. Serverless functions may simplify irregular workflows such as document validation, but security assessment must extend beyond function code to triggers, permissions and data paths [10]. Private-cloud orchestration also remains relevant where policy or contractual constraints preclude deployment of particular datasets on public infrastructure [11].

Cybersecurity evidence warns that infrastructure heterogeneity can fragment protection. Reviews of cloud threats and machine-learning defences cover misconfiguration, compromised privileges and monitoring gaps, but detection performance observed in one environment does not generalise automatically to payroll or employee records [12, 13]. Process-mining research illustrates how event sequences can support identification of abnormal file operations across providers [6]. These insights support a shift from individually secured applications towards consistent enterprise controls, combined with context-specific verification of their coverage and failure modes.

Interdisciplinary reviews show that HR scholarship and computational research often apply different definitions of successful automation [29]. This mismatch matters for platform investment: studies of workplace outcomes recommend assessing consequences across processes and organisational levels, rather than equating adoption frequency with employee value [30].

Table 1. Evidence-to-design synthesis for digital HR platform decisions.

Evidence theme	Sources	Design implication	Critical limitation
Architecture and portability	[1–5]	Bounded service design and executable workloads	Data portability and end-to-end quality remain difficult
Cross-cloud reliability	[6–11]	Event visibility, orchestration and selective placement	Stateful replication and network complexity create new dependencies
Cyber security	[12, 13]	Privilege controls and cross-provider detection	Model efficacy is context-dependent; configuration remains critical
Saudi readiness	[14–20]	Trust, skills, governance and institutional fit	Sector and organisation differences constrain generalisation
Digital HR outcomes	[21–25, 28–30]	User capability, strategy and people outcomes	Productivity does not establish employee fairness
Ethical HR analytics	[26, 27]	Purpose controls, audit and review rights	Formal codes may not ensure operational safeguards

5 Architecture decisions for HR service integration

A future-ready platform should begin with the employee lifecycle as its organising principle. A canonical workforce identifier must be governed centrally even where applications own separate records. The architecture should distinguish authoritative systems from derived analytical stores and establish explicit rules for corrections, retention and deletion. Core employment status and payroll entitlements warrant stronger consistency guarantees than frequently updated training recommendations. Treating all data flows alike creates either unnecessary latency or unacceptable reconciliation risk.

The proposed reference architecture comprises five interacting layers, illustrated in Figure 1. The experience layer provides accessible employee, manager and HR interfaces. A domain-services layer separates accountable operations such as hiring, attendance, benefits and case resolution. An integration layer implements an API gateway, policy-enforced event exchange and versioned data contracts. A data stewardship layer governs master records, lineage, classification, retention and encryption. Finally, a platform-control layer delivers identity federation, workload placement, observability, incident management and recovery testing. The separation is logical: it does not require five different products or providers.

Multi-cloud placement should follow workload criticality and demonstrable provider advantages. A provider may host transactional HR services where approved; another may supply analytical compute for appropriately minimised and authorised data; and a private environment may retain records with particularly restrictive controls. Yet placing a payroll database in two clouds does not automatically deliver continuity. It creates replication lag, reconciliation choices, additional administrative privileges and potentially conflicting jurisdictional exposures. Data movement must therefore be justified by a risk-based availability objective rather than aesthetic architectural symmetry [1, 8].

Integration patterns require similar selectivity. Synchronous APIs suit employee queries requiring immediate validation, whereas asynchronous events are appropriate for communicating non-critical changes to downstream services. Events should contain the minimum information needed by a recipient and reference authoritative records when full details are unnecessary. Idempotent


processing, durable message storage, replay controls and dead-letter handling are necessary because delivery failures are routine distributed-system events. The business consequence of duplicate onboarding messages, for instance, may be incorrect access provisioning or benefits enrolment, not merely a technical exception.

Cloud-native decomposition also raises ownership questions. HR retains accountability for policy decisions and record definitions; information security controls privilege boundaries; data stewards approve permitted reuse; and platform teams operate deployment pipelines and standard tooling. Vendor-managed software does not eliminate these responsibilities. The organisation should map each service to an owner, lawful purpose, recovery objective, dependencies and exit strategy. Lin and colleagues' enterprise service-stack research suggests value in coherent platform abstractions, but internal teams must resist hiding provider-specific constraints behind a misleading universal interface [9].

Portability must be tested at three distinct levels: executable workloads, persistent data and operational procedures. Container images can simplify redeployment without transferring databases, encryption keys, identity policies or monitoring histories. Consequently, a credible exit exercise should demonstrate a supported export format, an alternate execution environment, restored access rules and validated business transactions. Migration literature demonstrates why post-migration quality testing is indispensable [2, 5]. A successful demonstration should include payroll reconciliation and controlled employee self-service, not simply a container that starts.

A governed, modular HR platform

Service boundaries and controls are logical; cloud placement remains workload-specific



Conceptual framework; deployment boundaries require organisational policy and legal review.

Figure 1. A five-layer conceptual architecture and illustrative placement zones; not a deployed system.

6 Identity, privacy and trustworthy analytics

Unified identity governance is the central security prerequisite for a distributed HR platform. A user authenticated through a corporate directory may still acquire excessive rights when role mappings differ among software services. Federated sign-on must therefore be complemented by least-privilege authorisation, short-lived service credentials, separation of duties and automated offboarding. Role changes are particularly consequential: a former line manager should not retain access to subordin-

ates' pay or medical accommodations after transfer. Continuous access reviews should examine effective privileges across all vendors, not only directory group membership.

Confidentiality also depends on data minimisation and purpose separation. HR analytics frequently combines absences, turnover signals, engagement responses and productivity measures; individually defensible datasets can become intrusive when linked at scale. Research on big data ethics identifies risks of surveillance, profiling and asymmetric employee power [27]. An implementation should record why each attribute is collected, who can process it, the applicable retention schedule and whether an individual can challenge an inference. Pseudonymisation reduces exposure but does not remove obligations where records can be reidentified.

Artificial intelligence introduces governance demands that ordinary cloud monitoring cannot satisfy. Reviews of HRM and AI highlight the need to integrate human judgement, organisational accountability and employee consequences into adoption strategies [24, 25]. Comparative research on HR analytics projects finds that formal ethics boards and codes may coexist with unresolved data ownership and operational assurance gaps [26]. A credible model-governance process therefore requires documented training-data provenance, bias evaluation relevant to the affected workforce, material-change approval, independent review and a mechanism for correcting harmful outcomes.

Technical evidence must support such oversight. Every consequential automated recommendation should carry a versioned record of input categories, model version, applicable business rules, reviewer action and final decision. Logging must avoid reproducing complete sensitive profiles in security tools. Event trails should be tamper-evident, time-synchronised and linked across services using protected correlation identifiers. Research applying process analysis to cross-cloud file events indicates the value of lifecycle visibility, although its specific detection findings should not be treated as proof of suitability for personnel systems [6].

Zero-trust principles are useful only when translated into concrete checks. Mutual service authentication, workload identity, encryption-key segregation, configuration baselines and policy enforcement should operate consistently across providers. The attack-surface analysis in cloud security research supports attention to permissions and exposed interfaces [12, 13]. However, duplicating every security tool in every cloud can itself create blind spots. A unified evidence model should aggregate validated events while retaining local controls that continue functioning during network or central-management failure.

Resilience must also include confidentiality and recoverability. An outage response that reroutes employee applications to an unapproved processing location can trade operational continuity for unlawful disclosure. Recovery exercises must therefore validate permitted locations, administrative entitlements, deletion requests, audit logs and data integrity alongside recovery time. Quantitative work on serverless security reinforces the importance of evaluating systems and threats at service boundaries rather than assuming managed platforms remove risk [10]. A truly dependable HR environment returns to a compliant and reconciled state, not merely an accessible one.

7 Saudi institutional fit and human capital value

Saudi organisations should interpret architectural choice through their own governance landscape. The Personal Data Protection Law, its implementing instruments, applicable data-transfer requirements and relevant cybersecurity policies may affect HR records differently according to controller status, sector and activity. This review does not issue legal conclusions. Rather, it proposes a traceable

control register through which compliance teams map each application, dataset, processor, storage location, transfer path, purpose and retention obligation to an approved policy decision.

Domestic research underscores that provider trust and credible service assurance influence Saudi cloud adoption [14, 15]. Alshahrani's study of government technology adoption links progress to strategic alignment while documenting legacy integration, financial and organisational challenges [16]. Research on cybersecurity practices in Saudi public services similarly draws attention to training, regulatory adherence and technology readiness [17]. These are relevant contextual indicators, although results from public-sector services cannot be transferred uncritically to private payroll or outsourced HR service providers.

For human capital development, the most defensible value proposition is improved access to timely, accurate and actionable workforce services. Research on Saudi HR efficiency describes the organisational promise of digitisation but does not establish that a particular multi-cloud topology caused an improvement [19]. The broader Saudi human capital literature connects technology integration with institutional support and training [20]. These studies imply that workforce data quality, employee participation, accessible interfaces and supervisory competence require investment alongside infrastructure.

Vision 2030 alignment should be operationalised rather than asserted. Potential indicators include time to onboard new employees, processing errors in payroll, completion of job-relevant learning, resolution time for workforce cases, consistency of workforce reporting and employee perceptions of procedural fairness. Disaggregate measurements where appropriate to test whether systems serve workers across language, ability, location and job category. Findings on future workforce readiness support treating digital literacy and adaptability as outcomes and adoption conditions rather than automatic products of digitisation [18].

A final policy-sensitive decision concerns the boundary between local hosting and global service functionality. A provider's regional presence cannot establish that all backups, telemetry, support access, disaster recovery or subcontracting remain within the same boundary. Procurement must obtain specific architectural evidence and contractually enforce permitted processing. These controls increase implementation effort but make national digital ambitions compatible with accountable employee-data stewardship.

8 Staged implementation and evaluation framework

Figure 2 presents a six-stage implementation pathway designed to prevent premature provider diversification. The initial stage inventories workforce processes, authoritative data, legal obligations, unresolved interfaces and business-critical failure points. Teams then define target outcomes with HR leaders and employee representatives. This sequence prevents technology departments from optimising infrastructure metrics while overlooking inaccurate job records or confusing staff workflows.

The second stage establishes a governed integration foundation: canonical identifiers, service catalogues, reusable authentication patterns, approved encryption, infrastructure-as-code templates and observability conventions. Before any critical migration, organisations should demonstrate automated provisioning and revocation across at least two application boundaries. Architecture principles should be recorded as testable requirements. For example, a recruitment application must not directly modify payroll entitlements, and a learning platform should receive only the employee attributes necessary for enrolment.

Stage three selects a bounded, lower-risk pilot, such as employee case tracking or training catalogue integration. A side-by-side comparison with the existing service should measure end-to-end completion, data reconciliation, accessibility and operational support workload. Migration studies report that transitions towards microservices are long-running organisational journeys, not instantaneous technical substitutions [3]. Therefore, the pilot should include incident simulations and cost accounting before the target architecture is expanded.

Stage four progressively separates high-value domain functions and introduces event-driven integration where its reliability advantages exceed its complexity. Domain extraction must follow business ownership and well-defined service contracts. Stateful components, especially compensation and employment records, should move only after integrity checks, recovery rehearsals and audit access are verified. Organisations can also retain a well-governed monolith where decomposition adds little functional value. Cloud-native architecture is a means of achieving service outcomes, not a mandatory description of every component.

Stage five introduces selective second-provider use. Candidate workloads should be assessed against availability criticality, interoperability, approved data geography, licensing, migration cost, staff expertise and dependency concentration. Multi-cloud deployment research demonstrates that optimisation requires explicit resource and quality constraints [8]. The decision should document why additional providers improve the outcome after integration, network egress, duplicated tooling and incident coordination costs are included. A second provider should not be added simply to satisfy a corporate technology slogan.

The sixth stage institutionalises evidence-based operation. Table 2 connects architectural controls with accountable owners and measurable outcomes. A small dashboard should combine engineering measures, such as service-level attainment, failed identity revocations and restoration tests, with employee-centred measures, such as case resolution accuracy and perceived fairness. Thresholds must be agreed before deployment and calibrated against a baseline. Claims about increased resilience are credible only when recovery tests, not theoretical redundancy, substantiate them.

Evaluation must accommodate causality carefully. A before-and-after improvement may reflect staffing changes, revised policies or seasonal workload rather than architecture. Future pilots should record such confounders and compare similar business units or staged migrations where feasible. Interviews with employees and HR professionals should complement operational logs because technical dashboards cannot reveal whether access is understandable or decisions feel procedurally legitimate. Baseline documentation would also permit transparent reporting of implementation failures and corrective action.

From architecture decisions to measurable value

Six gates prevent premature workload duplication and unsupported benefit claims




Figure 2. Proposed six-stage transformation pathway with evidence-based decision gates.

Table 2. Control ownership and evaluation measures.

Control domain	Accountable function	Example measure
Identity lifecycle	HR operations + IAM	Revocation delay; privileged access exceptions
Workforce records	Data steward + HRIS owner	Reconciliation errors; unauthorised data duplication
Service continuity	Platform engineering	Validated restore time; transaction integrity
Employee service	HR service owner	Resolution accuracy; completion and accessibility
Analytics oversight	HR governance + privacy	Reviewed decisions; documented appeals and model changes
Cloud economics	Architecture + finance	Full ownership cost; provider exit rehearsal
Regulatory alignment	Legal + security + procurement	Closed compliance actions; verified data pathways

9 Critical discussion

The synthesis supports a conditional proposition: cloud-native modularity can help organisations modernise workforce services, whereas multi-cloud deployment is justified only when provider diversity contributes a measurable advantage after governance costs are considered. The two strategies address different problems. Microservices primarily change application structure and release coordination; multi-cloud changes dependency distribution and operating boundaries. Neither inherently improves record quality, strategic HR capabilities or the legitimacy of automated decisions. The literature on electronic HRM and digital transformation therefore acts as an essential corrective to infrastructure-led narratives [21, 23].

One unresolved tension is between interoperability and bespoke optimisation. Standard interfaces make replacement more feasible, but proprietary services may offer superior functionality for particular analytical or operational tasks. A blanket demand for provider neutrality could prevent

beneficial innovation; total dependency on specialised products, however, may undermine continuity and bargaining power. The appropriate governance mechanism is an explicit exit-cost estimate, tested export capability and documented acceptance of unavoidable dependencies. This is more precise than an unsupported claim of complete portability.

A second tension concerns elasticity and personnel-data consistency. Stateless recruitment search can often scale horizontally without changing business meaning. Pay calculations and employment histories depend on stronger guarantees, and simultaneous writes across geographic locations can create conflicts with serious consequences. Selecting storage and replication patterns according to domain invariants is thus preferable to uniform cloud-native design. Performance optimisation should be evaluated against the cost of correcting a wrong salary, access entitlement or termination record, not against latency alone.

Third, compliance cannot be delegated wholesale to technology vendors. Cloud providers can offer certifications, access controls and audit evidence, but organisations determine processing purposes, permissible integrations and decisions affecting employees. Research on ethical HR analytics suggests that formal governance language is insufficient when teams lack clear responsibility for actual data flows [26, 27]. Traceability from policy to infrastructure configuration and operational verification therefore becomes a central component of platform maturity.

Fourth, the strategic value of the proposed architecture depends on institutional capability. HR professionals require data literacy, platform staff require distributed-system engineering expertise, and executives need to interpret risk and service metrics together. Evidence from Saudi digital adoption and workforce research reinforces the role of skills, governance and leadership [16, 18, 20]. Investments in capability development are not separate from the technical programme; they determine whether interfaces, security mechanisms and analytic outputs are used appropriately.

The conceptual framework ultimately challenges the assumption that maximum technological distribution creates maximum resilience. Concentrating a sensitive data store in a properly governed location while distributing stateless services may be more robust and less costly than indiscriminate duplication. Conversely, a single provider may remain appropriate where failure tolerance, service design and contractual protections are strong. Architecture should be selected through a demonstrable workload-level decision, and periodically revised as regulatory conditions and operational evidence change.

10 Limitations and future research

This review integrates publications from related but distinct disciplines and does not estimate an effect size for multi-cloud adoption in Saudi HR systems. Its purposeful literature selection is not an exhaustive bibliography. Several architecture studies investigate generic enterprise, telecommunications or government applications rather than employment records, limiting direct external validity. Likewise, Saudi digital transformation findings frequently address public institutions, which may differ in procurement, exposure and governance from commercial HR providers. These boundaries prevent causal claims about national workforce performance.

Future work should test the architecture in longitudinal organisational deployments, with separate evidence for employee outcomes, security effectiveness and financial sustainability. Comparative designs could examine a single-cloud modernisation against selective multi-cloud placement while holding HR process redesign constant. Researchers should report recovery outcomes, data quality,

access violations, user experience, implementation effort and total ownership cost rather than headline service uptime alone. Independent assessment of employee rights and cross-border processing would add a necessary legal and ethical dimension.

Further enquiry should also investigate bilingual self-service, outsourced workforce arrangements, changing employment relationships and the governance of algorithmic recommendations. Reporting negative findings would be particularly valuable because unsuccessful integrations, contested employee analytics and unexpected data-transfer constraints are often underrepresented in platform success accounts.

11 Conclusion

Digital human capital management requires a dependable relationship between people, processes, decisions and data. This review shows why a cloud-first technical programme is insufficient unless its integration and governance mechanisms are designed around that relationship. Multi-cloud infrastructure can diversify selected dependencies, and cloud-native methods can improve independent evolution of services, but both create coordination obligations that organisations must be capable of fulfilling.

The proposed contribution is a five-layer HR platform architecture linked to an incremental implementation pathway and measurable accountability controls. It combines domain-based design, selective placement, workforce identity governance, data stewardship and verifiable recovery with attention to employee experience. For Saudi Arabia, the framework provides an approach to aligning modern workforce services with digital transformation priorities without assuming that regional hosting or provider certification resolves every regulatory requirement.

A future-ready platform should therefore be judged by accurate services, defensible decisions, resilient operations, manageable costs and justified data use. The proposed framework offers an analytical basis for those judgements, while acknowledging that empirical validation remains necessary before organisational benefits can be claimed.

References

- [1] Alonso J, Orue-Echevarria L, Casola V, et al. Understanding the challenges and novel architectural models of multi-cloud native applications: a systematic literature review. *Journal of Cloud Computing*. 2023;12:6. <https://doi.org/10.1186/s13677-022-00367-6>
- [2] Hasan MH, Osman MH, Admodisastro NI, Muhammad MS. Legacy systems to cloud migration: a review from the architectural perspective. *Journal of Systems and Software*. 2023;202:111702. <https://doi.org/10.1016/j.jss.2023.111702>
- [3] Michael Ayas H, Leitner P, Hebig R. An empirical study of the systemic and technical migration towards microservices. *Empirical Software Engineering*. 2023;28:85. <https://doi.org/10.1007/s10664-023-10308-9>
- [4] Velepucha V, Flores P. A survey on microservices architecture: principles, patterns and migration challenges. *IEEE Access*. 2023;11:88339–88358. <https://doi.org/10.1109/ACCESS.2023.3305687>
- [5] Nordli ET, Haugeland SG, Nguyen PH, Song H, Chauvel F. Migrating monoliths to cloud-native microservices for customizable SaaS. *Information and Software Technology*. 2023;160:107230. <https://doi.org/10.1016/j.infsof.2023.107230>
- [6] Zhang X, Cui L, Shen W, et al. File processing security detection in multi-cloud environments: a process mining approach. *Journal of Cloud Computing*. 2023;12:100. <https://doi.org/10.1186/s13677-023-00474-y>
- [7] Martin R, Vidal I, Valera F. A software-defined connectivity service for multi-cluster cloud native applications. *Computer Networks*. 2024;248:110479. <https://doi.org/10.1016/j.comnet.2024.110479>

-
- [8] Singh B, Kaur R, Woodside M, Chinneck JW. Low-power multi-cloud deployment of large distributed service applications with response-time constraints. *Journal of Cloud Computing*. 2023;12:1. <https://doi.org/10.1186/s13677-022-00363-w>
- [9] Lin J, Xie D, Huang J, Liao Z, Ye L. A multi-dimensional extensible cloud-native service stack for enterprises. *Journal of Cloud Computing*. 2022;11:83. <https://doi.org/10.1186/s13677-022-00366-7>
- [10] Ni K, Mondal SK, Kabir HMD, Tan T, Dai HN. Toward security quantification of serverless computing. *Journal of Cloud Computing*. 2024;13:140. <https://doi.org/10.1186/s13677-024-00703-y>
- [11] Park J, Jeong S, Yeom K. Private cloud bespoke orchestrator: techniques for constructing and operating bespoke-private cloud virtual machine environments for cloud users. *Journal of Cloud Computing*. 2025;14:34. <https://doi.org/10.1186/s13677-025-00760-x>
- [12] Alouffi B, Hasnain M, Alharbi A, Alosaimi W, Alyami H, Ayaz M. A systematic literature review on cloud computing security: threats and mitigation strategies. *IEEE Access*. 2021;9:57792–57807. <https://doi.org/10.1109/ACCESS.2021.3073203>
- [13] Nassif AB, Talib MA, Nasir Q, Albadani H, Dakalbab FM. Machine learning for cloud security: a systematic review. *IEEE Access*. 2021;9:20717–20735. <https://doi.org/10.1109/ACCESS.2021.3054129>
- [14] Alghaith W. The role of trust in cloud service providers on the adoption of cloud computing in Saudi Arabia: an empirical investigation. *Journal of Computer Science*. 2023;19(9):1107–1123. <https://doi.org/10.3844/jcssp.2023.1107.1123>
- [15] Alamri N, Alzahrani S. Navigating the clouds: an exploratory research of cloud computing adoption in Saudi Arabia's small and medium enterprises. *Engineering, Technology & Applied Science Research*. 2024;14(6):17859–17869. <https://doi.org/10.48084/etaser.8435>
- [16] Alshahrani A. Adopting emerging technologies in digital government: a multi-case analysis of drivers, enablers, and challenges in Saudi Arabia. *Digital Government: Research and Practice*. 2025;6(3):36. <https://doi.org/10.1145/3719297>
- [17] Al-Hawamleh AM. Investigating the multifaceted dynamics of cybersecurity practices and their impact on the quality of e-government services: evidence from the KSA. *Digital Policy, Regulation and Governance*. 2024;26(3):317–336. <https://doi.org/10.1108/DPRG-11-2023-0168>
- [18] Alrashedi AK. Future workforce skills for Saudi Arabia: a PLS-SEM analysis of Vision 2030 readiness. *Information Development*. 2025. <https://doi.org/10.1177/02666669251357311>
- [19] Hassan AB, Al-Subhi RA, Bajaber N. The impacts of digital transformation on HR efficiency in KSA. *International Journal of Financial, Administrative and Economic Sciences*. 2024;3(10):278–319. <https://doi.org/10.59992/IJFAES.2024.v3n10p8>
- [20] Abubakar AA, Al-Mamary YH, Singh HP, Singh A, Alam F, Agrawal V. Exploring factors influencing sustainable human capital development: insights from Saudi Arabia. *Heliyon*. 2024;10(16):e35676. <https://doi.org/10.1016/j.heliyon.2024.e35676>
- [21] Zhou Y, Cheng Y, Zou Y, Liu G. e-HRM: a meta-analysis of the antecedents, consequences, and cross-national moderators. *Human Resource Management Review*. 2022;32(4):100862. <https://doi.org/10.1016/j.hrmr.2021.100862>
- [22] da Silva LBP, Soltovski R, Pontes J, et al. Human resources management 4.0: literature review and trends. *Computers & Industrial Engineering*. 2022;168:108111. <https://doi.org/10.1016/j.cie.2022.108111>
- [23] Bansal A, Panchal T, Jabeen F, Mangla SK, Singh G. A study of human resource digital transformation (HRDT): a phenomenon of innovation capability led by digital and individual factors. *Journal of Business Research*. 2023;157:113611. <https://doi.org/10.1016/j.jbusres.2022.113611>
- [24] Malik A, Budhwar P, Kazmi BA. Artificial intelligence (AI)-assisted HRM: towards an extended strategic framework. *Human Resource Management Review*. 2023;33(1):100940. <https://doi.org/10.1016/j.hrmr.2022.100940>
- [25] Basu S, Majumdar B, Mukherjee K, Munjal S, Palaksha C. Artificial intelligence–HRM interactions and outcomes: a systematic review and causal configurational explanation. *Human Resource Management Review*. 2023;33(1):100893. <https://doi.org/10.1016/j.hrmr.2022.100893>
- [26] Bar-Gil O, Ron T, Czerniak O. AI for the people? Embedding AI ethics in HR and people analytics projects. *Technology in Society*. 2024;77:102527. <https://doi.org/10.1016/j.techsoc.2024.102527>
-

-
- [27] Manroop L, Malik A, Milner M. The ethical implications of big data in human resource management. *Human Resource Management Review*. 2024;34(2):101012. <https://doi.org/10.1016/j.hrmr.2024.101012>
- [28] Montero Guerra JM, Danvila-del-Valle I, Méndez-Suárez M. The impact of digital transformation on talent management. *Technological Forecasting and Social Change*. 2023;188:122291. <https://doi.org/10.1016/j.techfore.2022.122291>
- [29] Pan Y, Froese FJ. An interdisciplinary review of AI and HRM: challenges and future directions. *Human Resource Management Review*. 2023;33(1):100924. <https://doi.org/10.1016/j.hrmr.2022.100924>
- [30] Pereira V, Hadjielias E, Christofi M, Vrontis D. A systematic literature review on the impact of artificial intelligence on workplace outcomes: a multi-process perspective. *Human Resource Management Review*. 2023;33(1):100857. <https://doi.org/10.1016/j.hrmr.2021.100857>