INTELLIGENT DECISION AUTOMATION

Intelligent Decision Automation for Saudi Government and Banking Services under Vision 2030: Integrating Robotic Process Automation, Business Process Management, and Human-in-the-Loop Decision Governance

Mohammad Nauman

Published November 2026

Abstract

Intelligent decision automation promises faster Saudi public and banking services, yet automating a task is not equivalent to governing a decision. This critical integrative review examines how robotic process automation, business process management and human oversight can be combined without weakening procedural fairness, operational resilience or accountability. Thirty peer reviewed studies published between 2020 and 2025 are synthesised alongside the institutional context of Vision 2030, including digital government, regulated finance and electricity related service workflows. The analysis differentiates interface automation from process orchestration and consequential decision authority, and develops an original layered governance architecture with risk dependent human intervention, traceable rules and continuous process assurance. Applications considered include public licences, customer onboarding, banking exceptions, electricity connections, billing disputes and outage compensation. The review finds that automation benefits depend on prior process simplification, dependable data, documented decision rights and meaningful rather than symbolic human review. Machine predictions can inform routing and anomaly detection, but do not by themselves establish entitlement, legal compliance or equitable treatment. A proposed evaluation scorecard integrates throughput, accuracy, appeal outcomes, accessibility, cyber resilience and auditability rather than relying on labour savings alone. The paper identifies research gaps concerning Arabic language service design, cross agency interoperability, automation bias, electricity sector dependencies and longitudinal evidence in Saudi institutions. It proposes staged implementation and independent evaluation rather than claiming demonstrated national performance gains. The contribution is a context sensitive decision governance model that distinguishes technical execution, institutional responsibility and contestable public or financial outcomes.

Keywords: banking; business process management; electricity services; human-in-the-loop; intelligent decision automation; public governance; robotic process automation; Saudi Arabia; Vision 2030

1 Introduction

Saudi Arabia's Vision 2030 places service integration, digital capability and economic diversification at the centre of institutional reform [1]. Government agencies and banks increasingly manage journeys that move across identity records, forms, payment platforms, supervisory controls and customer communication. Those journeys may appear digital at the front end while retaining fragmented checks and manual decisions behind the interface. Intelligent decision automation therefore raises a question more demanding than whether software can imitate an employee: which decisions should be automated, under whose authority, with what evidence, and with what route to correction?

Contemporary research warns that technical efficiency can coexist with fragile public legitimacy when the relevant decision process remains poorly understood [2–4].

Robotic process automation, hereafter RPA, executes repeatable interactions across existing applications. Business process management, hereafter BPM, models and controls end to end workflows, responsibilities and exceptions. Human in the loop governance, hereafter HITL, assigns accountable personnel the capacity to examine, reject, amend or escalate system recommendations. These are complementary rather than interchangeable constructs. RPA can transfer a field without deciding whether it is correct; BPM can route a case without determining whether the rule is just; a human approval click can occur without an effective opportunity to scrutinise the evidence. The distinction matters in regulated finance and public administration, where an erroneous decision may deny access, delay payments or change entitlements [5–8].

Electrical infrastructure gives this debate operational urgency. Electricity connection requests involve permits, network capacity checks, identity verification, contractor scheduling and payments. Billing disputes combine metering evidence, tariff logic, customer communication and possible manual adjustment. Outage notifications and restoration claims require prioritisation without permitting an administrative bot to issue unreviewed operational commands to a live network. Accordingly, the electricity sector is analysed as a service and infrastructure interface, not as an assertion that administrative automation can replace engineering control systems. Work on smart grids and emergency demand response reinforces the need to respect reliability constraints when information is shared across organisational boundaries [9–11].

The paper asks: how can Saudi institutions integrate RPA, BPM and human discretion into an auditable decision architecture that advances speed, fairness and resilience simultaneously? Its aim is to develop a defensible critical synthesis and implementation framework, not to estimate causal effects using uncollected Saudi field data. The objectives are to classify automation layers, evaluate reported tensions in governance, map sector specific decision rights, propose practical controls and define a testable research agenda. The argument is that the most valuable unit of transformation is the governed service journey, rather than the quantity of deployed bots.

2 Conceptual and theoretical foundations

Automation research distinguishes the digital execution of an activity from the redesign of the surrounding organisation [12]. Syed and colleagues [5] identify technical and managerial challenges in RPA, while Hofmann and colleagues [6] frame software robots as components embedded within existing information systems. Mapping work by Wewerka and Reichert [13] shows why RPA descriptions require systematic classification rather than vendor vocabulary. A screen scraping bot may be helpful where a secure application interface does not exist, but it inherits the fragility of changing screens, inconsistent data formats and credentials. The correct performance comparison is therefore not bot execution against employee keystrokes alone, but a redesigned service against the full previous service.

BPM offers the missing systems perspective through process discovery, modelling, orchestration, exception handling and feedback. El Gharib and Amyot [14] show that process mining can help locate automation candidates while identifying data preparation limitations. Egger and colleagues [15] extend this logic by linking bot logs and process logs. Beerepoot and colleagues [16] identify persistent BPM challenges involving human involvement, process context and knowledge intensive work. These insights imply that institutions should map the actual sequence of cases, including rework, waiting,

informal interventions and rejected submissions, before deciding whether to automate a task. A polished process model without observed exceptions remains a partial representation.

Decision governance adds a normative dimension. In public services, authority derives from laws, delegated responsibilities and review procedures rather than model confidence. Peeters [7] demonstrates that algorithm design shapes the real discretion available to officials. Roehl and Hansen [8] identify tensions among accountability, efficiency, fairness, resilience and transparency. Enarsson and colleagues [17] likewise show that placing a human in an automated chain cannot be evaluated independently of legal and institutional context. Taken together, these studies challenge the simplistic division between entirely automatic and entirely manual work. Decision authority can vary within a single case as new evidence, uncertainty or consequences emerge.

Adoption also depends on organisational capability. The review by Madan and Ashok [3] links public sector AI diffusion to institutional conditions and value tensions; Straub and colleagues [4] caution that operational performance and normative alignment must be distinguished. Selten and Klievink [18] identify organisational choices between separate innovation teams and integrated capabilities. Saudi research on government technology adoption highlights the importance of implementation conditions rather than treating national ambition as proof of local readiness [19]. The resulting theoretical position combines process redesign, socio technical capability and public value governance, with human judgement retained wherever decision consequences or evidential uncertainty justify it.

3 Critical integrative review methodology

This study employs a structured, critical integrative review rather than a statistical meta analysis. The analytical corpus consists of thirty identifiable, DOI linked peer reviewed publications dated 2020 to 2025. They were selected purposively to cover RPA design and limitations, BPM and process mining, automated administrative decision making, explainability, public sector adoption, Saudi digital services, banking processes and electricity system constraints. Searches and bibliographic cross checks used scholarly publisher records, institutional repositories and research indexing pages with combinations of terms including robotic process automation, business process management, automated decision making, human oversight, government services, banking, Saudi Arabia and smart grid. The reference set supports conceptual comparison; it is not represented as an exhaustive census of database results.

The inclusion logic required an explicit relation to automation mechanisms, human algorithm interaction, institutional governance or a relevant Saudi service context, plus a publication year within the specified range and a resolvable DOI record. Purely promotional reports and untraceable outcome claims were not treated as research evidence. The synthesis follows review conventions, including an abstract, numbered analytical sections, an explicit method, thematic synthesis, figures, tables and practical implications; transparency is maintained without reproducing another study's sample, results or proposed model. This design avoids incorrectly borrowing another article's search sample or claiming independent fieldwork.

For synthesis, each source was assessed along five interpretive dimensions: locus of automation; decision consequence; organisational and information dependencies; human review mechanism; and evidence strength. Studies were then compared across task, journey and institution levels. Strongest weight was given to findings directly supported by empirical designs and to conceptual distinctions corroborated across fields. Single context observations were not generalised to Saudi Arabia without qualification. In particular, electricity optimisation studies establish operational sensitivities, not

measured benefits from government bots. Disagreements are retained where researchers emphasise different values, such as speed versus procedural access.

The review has deliberate limitations. Search coverage was targeted rather than exhaustive, no bibliometric population or pooled effect size was calculated, and no interviews, surveys or Saudi transaction logs were collected. Therefore, illustrative service cases and the proposed scorecard are design propositions requiring validation, not observed deployment results. DOI verification concerns the identification of bibliographic records; it does not certify that each study has equivalent methodological strength. These restrictions make the resulting framework more transparent and prevent spurious precision.

4 Thematic synthesis and proposed architecture

The first synthesis finding is that task execution and decision legitimacy occupy different layers. In a government licence journey, an RPA bot may retrieve a document, but the authorisation rule belongs to an approved policy source and the final decision may require a designated officer. In banking, a bot can populate a customer file, whereas sanctions alerts, contested identity matches and unusual transaction patterns require controls established by risk policy. Human judgement cannot be made redundant merely because a workflow is digital. Grimmelikhuijsen [20] finds that algorithmic transparency can affect perceived trustworthiness, and Haesevoets and colleagues [21] report a preference for AI in an advisory rather than dominant decision role.

The proposed architecture in Figure 1 contains six linked control layers. First, a unified service entry records consent status where relevant, requested outcome, accessible contact channel and a case identifier. Second, BPM decomposes the journey into authorised steps with deadlines and accountable owners. Third, APIs or controlled RPA perform bounded exchanges while keeping original record provenance. Fourth, a decision service separates codified business rules from probabilistic predictions and applies versioned thresholds. Fifth, risk based routing sends uncertain, disputed or high consequence cases to empowered reviewers. Sixth, an evidence layer logs inputs, model or rule versions, interventions, explanations and outcome corrections. None of these layers is itself a claim of regulatory compliance.

The central distinction is between eligibility rules and predictive signals. A published tariff or eligibility condition can be represented as a deterministic rule if it is sufficiently complete and maintained by an authorised policy owner. A classifier that estimates fraud, priority or likely documentation errors produces uncertainty, not an entitlement decision. Combining both into a single opaque score would conceal which part of a case is legally required and which part is statistical inference. Fischer Abaigar and colleagues [22] warn of mismatches between technical machine learning assumptions and public sector decision realities. Consequently, decisions should expose rule provenance, confidence and missing information separately.

Operational assurance requires traceability across distributed applications. Each service event should carry a stable case identifier, actor identity, reason code, input version, activity timestamp and approval state. BPM event logs reveal delays and process deviation; bot logs identify failed selectors, duplicate updates and credential issues. Reconciliation checks prevent a bot from marking a case complete when its downstream posting has failed. Incident procedures should specify when to stop automation, restore a known state and notify affected service users. The integrated monitoring proposed by Egger and colleagues [15] makes this feasible as a research direction, although the controls still require practical testing within each institution.

Human involvement must be designed as a capability rather than a ceremonial checkpoint. A reviewer needs intelligible evidence, protected review time, authority to disagree, a route to request more data and a recorded rationale. Haitsma and Brink [23] show why meaningful involvement exceeds minimal intervention, while Hillo and colleagues [24] find legitimacy implications associated with transparency and human discretion. Training should therefore develop critical assessment of automated recommendations, rather than simply instructing employees to accept or reject alerts. Supervisors should monitor both excessive overrides and suspiciously uniform agreement, since either may signal a failing control.

A further design issue is the allocation of responsibility for rules that originate outside the automation team. A tariff amendment, banking directive or administrative eligibility change needs a named policy steward who translates it into a controlled change request, confirms its effective date and signs off the resulting tests. A technically successful bot release is inadequate if it applies superseded rules. Version identifiers should appear in every consequential decision record and support retrospective review when an interpretation changes. This separates business rule custody from information technology deployment and prevents silent changes in substantive public policy.

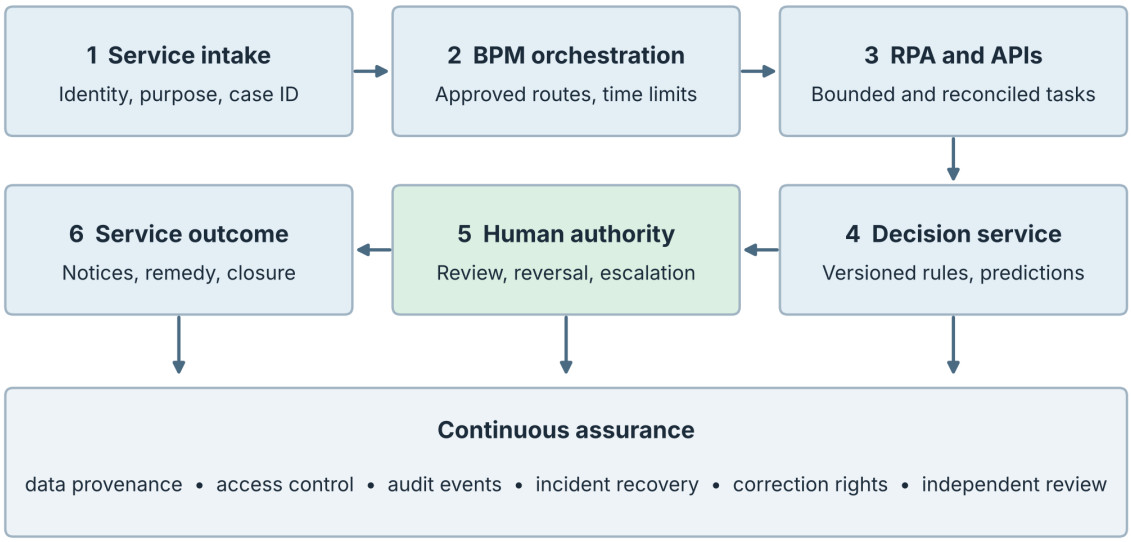


Figure 1. Proposed six-layer architecture for governed intelligent decision automation. Original conceptual model; flows are illustrative, not an implemented system.

5 Applications across Saudi government, finance and electricity services

Public administration is a plausible starting point for bounded automation in document intensive services. A licensing workflow may classify incoming documents, check whether mandatory fields are present, retrieve authenticated records and route missing evidence to the applicant. BPM can preserve statutory or administrative sequencing while RPA bridges older registers during system modernisation. Yet the final refusal of a licence must remain explainable, contestable and authorised. Saudi evidence on the barriers to electronic government adoption [25] and multi case technology implementation [19] suggests that citizen acceptance and institutional coordination are material conditions, not optional communication activities. Arabic and English guidance, accessibility and alternative support channels belong inside the service design.

Banking offers high volume but unevenly risky tasks. Customer onboarding may involve routine document comparison, account opening checks, adverse screening and escalation. Payment exceptions can include duplicated references, inconsistent beneficiary details and suspicious transaction patterns. Mamede and colleagues [26] show the value of process improvement before RPA deployment in a banking setting. The Saudi Central Bank's open banking direction provides a relevant interoperability context; nevertheless, authorised data sharing should not be confused with unrestricted automated decision rights. A controlled implementation should distinguish customer initiated access, institution managed risk decisions and independent compliance review, recording each handover and exception.

Risk teams should resist the assumption that perfect consistency implies fairness. Uniformly applied erroneous rules can distribute harm at scale, while a highly accurate anomaly model may create unacceptable false positives for particular customer groups. Aoki and colleagues [27] demonstrate that explanation formats influence perceptions of fairness and trust, and Rizk and Lindgren [28] describe how automation changes the decision space between officials and citizens. For bank credit, account restriction or government benefit cases, proportionate safeguards include reason codes, accessible reconsideration, feedback on corrected decisions and sampling for disparate error patterns. A reviewer's authority must encompass reversal, not just annotation.

Electrical infrastructure illustrates interaction between service automation and critical operations. A new connection application may link land permits, engineering eligibility, inspection scheduling, payment and customer notification. Intelligent routing could flag missing drawings or conflicting addresses, allowing engineers to focus on exceptional or technically complex cases. Billing disputes may require meter readings, tariff records, historical adjustments and customer explanation. These workflows are suitable for BPM coordination and constrained RPA tasks, but network switching, emergency protection and real time grid control require separate engineering safety systems. Related work on smart grid data [9] and emergency demand response [10] underscores why service decisions must preserve the integrity of operational information.

An outage compensation case creates a useful boundary test. Bots may assemble incident identifiers, customer eligibility data and claim documentation; a policy engine may calculate an indicative amount under an authorised rule; a staff member may review exceptions, disputed causation or vulnerability considerations. Customers must be able to contest inaccurate outage periods or account matches. Data used for customer notification should be segregated from systems that dispatch operational commands. Research on smart grid scheduling [11] concerns optimisation under technical constraints, not evidence that administrative automation improves grid reliability. This separation avoids transferring claims between distinct technical domains.

Across these applications, the most credible opportunity is service coordination. BPM can make responsibilities and waiting time visible; RPA can relieve staff of stable repetitive exchanges; decision assistance can prioritise attention; governance can preserve appeal and accountability. A national digital strategy may support shared platforms, but agencies and banks still face differing permissions, risks and service obligations. Table 1 maps the resulting differences into explicit decision boundaries. The lesson is not to standardise every decision, but to standardise the evidence and governance required whenever discretion is exercised.

Consider the transfer of a disputed customer record between a bank and a public service provider. It may be technically possible to match names, numbers and timestamps, but matching does not establish the lawful purpose for sharing or the correctness of the underlying record. An accountable

workflow should show which institution supplied each field, which authority permitted its use and who must correct it. The end user should not be required to pursue several disconnected departments when a single automation chain produces a combined error. Shared administration is valuable only if shared responsibility is defined.

Table 1. Cross-sector decision rights and audit requirements proposed for Saudi service workflows.

Service setting	Automatable activities	Reserved human judgement	Assurance evidence
Government permits	Document completeness; status updates; workflow routing	Contested eligibility; refusal; exceptional circumstances	Rule version; reviewer reason; appeal outcome
Banking operations	File assembly; reconciliations; routine notifications	Unusual risk alerts; restrictions; disputed identity	Consent provenance; exceptions; audit log
Electricity connections	Application matching; scheduling; payment reconciliation	Engineering exceptions; contested service eligibility	Inspection record; sign-off; safe handover
Outages and billing	Claim assembly; meter record comparison; notices	Causation dispute; compensation exception; vulnerability	Incident identifier; correction; customer notice

6 Risk, ethics and human decision governance

Risk classification should begin with the consequences of a mistaken outcome, not with whether the software appears sophisticated. Low consequence examples include formatting acknowledgements and routing non sensitive enquiries. Moderate consequence examples include service corrections with straightforward appeal. High consequence examples include refusal of essential services, financial restrictions, or decisions affecting vulnerable users. Critical infrastructure operations introduce safety and continuity dependencies even when the administrative interface appears routine. Table 2 therefore sets out a proposed control matrix using consequence, reversibility and evidential uncertainty as primary criteria. Its thresholds are illustrative and would require legal and risk owner approval in each participating organisation.

Privacy and cybersecurity are structural requirements, particularly when RPA credentials connect multiple systems. Institutions should enforce least privilege, credential vaulting, purpose limitation, access logging and retention schedules. Case level records must avoid unnecessary replication of sensitive documents into logs. Saudi personal data governance and banking security requirements should be interpreted by authorised legal and compliance teams rather than assumed satisfied by a software configuration. Controls also need to cover external model services, outsourced automation providers, incident reporting and privileged bot accounts. Without explicit ownership, interconnected automation may multiply the consequences of a single compromised credential.

Fairness assurance requires attention to both model behaviour and administrative process. A perfectly calibrated probability estimate may still be unsuitable as the decisive basis for a case affecting fundamental services. Conversely, mandatory human review may increase delays for people least able to navigate the system. Roehl and Hansen [8] show why governance involves trade offs rather than universal optimisation; convenient electronic interfaces can still undermine public values when they displace accessibility, explanation or effective reconsideration. In the proposed architecture, fairness therefore means consistent lawful criteria, effective access to explanation and correction, and monitoring of group level adverse outcomes where data governance allows.

Automation bias is a particular weakness in superficial HITL designs. Reviewers facing heavy queues may accept a recommendation because the interface displays a green indicator, because previous colleagues agreed, or because reversing a system outcome imposes administrative effort. The safeguard is not simply more human clicks. Organisations should include blind quality checks, disagreement rights, decision rationale sampling and periodic exercises using deliberately ambiguous but safe test cases. Enarsson and colleagues [17] and Haitsma and Brink [23] support evaluating the actual conditions of human agency. These tests must protect service users and never insert uncontrolled artificial outcomes into production decisions.

Resilience is equally important. Temporary failures should not silently convert deferred cases into refusals, duplicate payments or incorrect network service notifications. Fallback processes need defined service level objectives, queues capable of replay without duplication, and clear manual operating procedures. Auditors require the ability to reconstruct who or what acted, which information was visible and whether a correction was applied downstream. For continuity sensitive electricity workflows, escalation to network operations must remain distinguishable from routine customer service automation. These controls make the proposed system governable under failure rather than merely efficient when all components work.

Table 2. Consequence-sensitive human oversight matrix (design proposal, not a regulatory classification).

Tier / consequence	Proposed automation boundary	Required human control	Monitoring priority
A — routine / reversible	Execute bounded administrative steps	Sampled review and accessible support	Completion accuracy; duplication
B — material / correctable	Decision assistance and rule validation	Review material exceptions; documented override	Appeals; false alerts; delayed corrections
C — high impact / contested	Prepare evidence; do not issue final refusal automatically	Named decision maker; explanation and reconsideration	Disparate harms; reviewer agreement bias
D — safety / critical service	Administrative coordination only; no operational switching	Engineering authority; independent safety gate	Continuity; fail safe isolation; incident response

7 Evaluation strategy and implementation roadmap

An implementation programme should start with process discovery rather than software procurement. Candidate services can be assessed by frequency, standardisation, data quality, exception rate, regulatory sensitivity and change volatility. A high volume process with frequent policy revisions may be a poorer candidate than a smaller but stable service with obvious customer benefit. For each candidate, process owners should document current handling time, rework, adverse decisions, appeals, staff effort and accessibility barriers before any intervention. Recorded baseline definitions prevent favourable comparisons produced by changing the mix of cases or ignoring work transferred to another department.

A controlled pilot should compare the redesigned workflow with a documented baseline, ideally using comparable cohorts, risk adjustment and observation over enough time to capture exceptions. The scorecard needs to report median and tail processing times, first time completion, manual correction frequency, false alerts, complaint resolution, human override quality and cyber incidents. Labour hours saved should be assessed alongside implementation costs, maintenance burden and

user effort. The retained evidence on AI diffusion [3] and Saudi institutional implementation [19] supports evaluation of managerial capability and the practical conditions of adoption. No numerical improvement is forecast here because the evidence does not justify it.

Figure 2 presents an iterative assurance cycle. An accountable service owner first defines the public or financial value at stake, identifies prohibited automation and approves measurable outcomes. A multidisciplinary team maps the existing process and revises unnecessary steps. Technical staff configure restricted automation and evidence capture. Reviewers then test ambiguous cases and appeal paths before live release. Independent monitoring examines errors, process drift, complaints and unequal effects, after which governance owners decide whether to scale, revise or suspend the service. The cycle runs again whenever a rule, data source, interface or risk profile changes.

Scale decisions should be conditional on passing explicit gates. Technical readiness requires stable integration and recoverable transactions. Governance readiness requires named owners, approved decision rights, usable explanations, access controls and demonstrable corrective processes. Social readiness requires accessible channels, trained staff and service user communication. Operational readiness for electricity linked cases adds confirmation that administrative tools cannot bypass engineering safety controls. Agostinelli and colleagues [29] illustrate evolving RPA tooling that generates robots from interface logs, but automated construction of a script does not supply legal authority or service accountability. Human responsibility cannot be generated from recorded clicks.

Procurement should therefore specify vendor independent event formats, exportable evidence, testable interfaces and clear exit provisions. Where legacy user interfaces force temporary bots, the transition plan should include eventual API replacement and a register of fragile selectors. Organisations should publish internal service level commitments and report failures without suppressing inconvenient exceptions. This approach supports Vision 2030 objectives through measurable institutional capability, not unverified claims that deployment volume is equivalent to transformation. A robust programme can be slower at the pilot stage but less costly when rules, staff or infrastructure change.

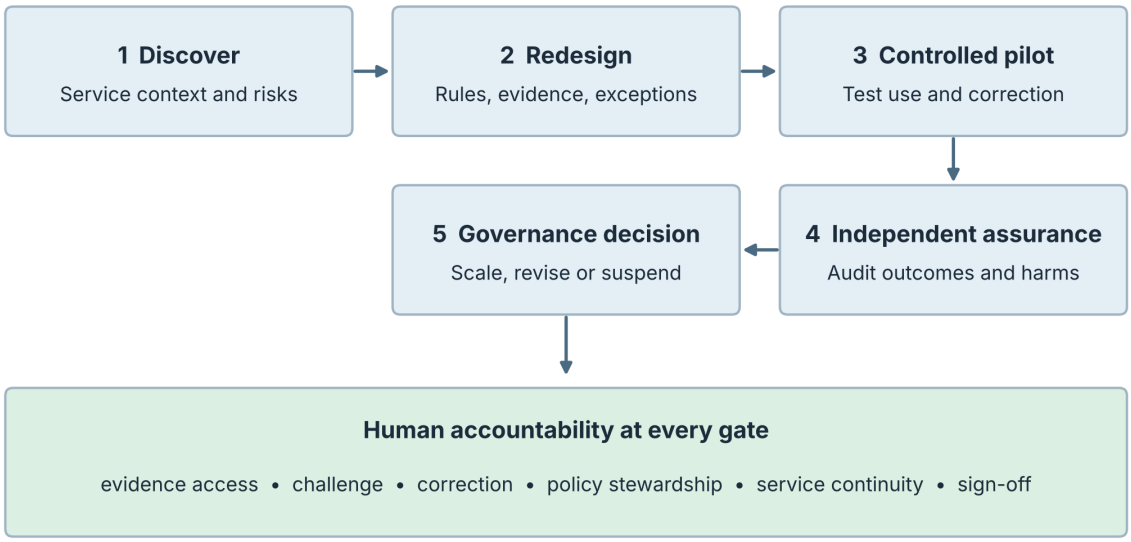


Figure 2. Iterative assurance cycle with continuous human accountability and decision gates. Original conceptual model.

8 Research gaps and future directions

The literature provides strong conceptual warnings but limited longitudinal evidence about integrated RPA, BPM and meaningful HITL governance in Saudi institutions. National reports and case studies describe digital progress, yet they do not establish causal benefits attributable to a particular decision architecture. Future research should compare services before and after redesign, track maintenance and appeals over several cycles, and include independently audited outcome measures. Mixed methods are appropriate because transaction logs show process behaviour while interviews and user studies reveal whether reviewers possess real discretion and whether affected people understand the results [2, 8, 24].

Arabic language and multilingual service contexts warrant targeted assessment. Automated document extraction, transliteration and name matching may have unequal error profiles across scripts, dialects and older records. Researchers should measure how those errors influence routing, identity resolution, delay and appeal, rather than reporting character recognition accuracy in isolation. Accessibility studies should include mobile users, older people, people with disabilities and those who need assisted channels. This is a service justice question as much as an engineering challenge. Evidence collection must follow applicable privacy and research ethics safeguards.

A second gap concerns interfaces between public authorities, banks and electricity providers. Research should model what data is necessary for a specific service, when sharing is lawful, which actor may make a final determination and how a correction propagates across systems. Federated auditing and shared case identifiers deserve evaluation, but common identifiers also introduce correlation and privacy risks. For grid related journeys, future studies should separate customer administration from the control and protection systems that maintain physical reliability. This distinction is essential if claims about digital efficiency are to remain technically credible [9, 10, 30].

A third gap is the empirical meaning of human control. Studies should compare different reviewer interfaces, explanation types, workload levels and reversal permissions using ethically controlled scenarios. Outcomes could include true error detection, unjustified agreement, time to correction and perceived procedural fairness. The perception evidence of Haesevoets and colleagues [21] and Aoki and colleagues [27] suggests that explanatory design matters; experimental work is still required to determine which formats improve actual oversight rather than confidence alone. Evaluation should explicitly record cases where humans appropriately disagree with model suggestions.

Finally, organisational research should examine maintenance and changing policy. Bots can become brittle as interfaces evolve; decision policies can drift away from written rules; experienced reviewers may leave; and new channels can create unexpected exclusions. Longitudinal governance registers, regular policy reconciliation and post incident learning would allow researchers to assess whether early efficiency survives these changes. Cross sector comparisons could identify which controls generalise and which must remain specific to public administration, banking or infrastructure. Such studies would convert the framework proposed here from a defensible synthesis into a validated implementation theory. The proposed distinction between procedural efficiency and legitimate authority is essential when administrative records shape real access to services.

9 Conclusion

Intelligent decision automation can support Vision 2030 when institutions govern complete service journeys rather than merely multiplying software robots. RPA is strongest at bounded execution, BPM

at coordinating and diagnosing processes, and human oversight at interpreting exceptions, protecting rights and accepting responsibility for consequential outcomes. Their integration is therefore a design problem involving rules, people, information architecture and institutional authority. The critical evidence reviewed here cautions against treating speed as a substitute for fairness or a human approval screen as proof of meaningful discretion.

For Saudi government and banking services, the layered architecture provides a basis for staged design, traceable evidence and risk proportionate review. Electricity service workflows demonstrate its value while clarifying that administrative automation must not intrude into safety critical grid control. The two analytical tables specify candidate decisions and governance safeguards; the two figures distinguish system architecture from iterative assurance. Future longitudinal and comparative research should evaluate actual user outcomes, resilience and correction rights. Until then, this framework should be treated as a testable proposition, not an empirically proven national model; institutional readiness requires documented responsibility, independent scrutiny, effective appeals and safety conscious deployment.

References

[1] Alkorbi S, Alrwais O. An empirical case study of digital government transformation in Saudi Arabia. *Information*. 2025;16(12):1110. https://doi.org/10.3390/info16121110

[2] Zuiderwijk A, Chen YC, Salem F. Implications of the use of artificial intelligence in public governance: a systematic literature review and a research agenda. *Government Information Quarterly*. 2021;38(3):101577. https://doi.org/10.1016/j.giq.2021.101577

[3] Madan R, Ashok M. AI adoption and diffusion in public administration: a systematic literature review and future research agenda. *Government Information Quarterly*. 2023;40(1):101774. https://doi.org/10.1016/j.giq.2022.101774

[4] Straub VJ, Morgan D, Bright J, Margetts H. Artificial intelligence in government: concepts, standards, and a unified framework. *Government Information Quarterly*. 2023;40(4):101881. https://doi.org/10.1016/j.giq.2023.101881

[5] Syed R, Suriadi S, Adams M, et al. Robotic process automation: contemporary themes and challenges. *Computers in Industry*. 2020;115:103162. https://doi.org/10.1016/j.compind.2019.103162

[6] Hofmann P, Samp C, Urbach N. Robotic process automation. *Electronic Markets*. 2020;30(1):99–106. https://doi.org/10.1007/s12525-019-00365-8

[7] Peeters R. The agency of algorithms: understanding human-algorithm interaction in administrative decision-making. *Information Polity*. 2020;25(4). https://doi.org/10.3233/IP-200253

[8] Roehl UBU, Hansen MB. Automated, administrative decision-making and good governance: synergies, trade-offs, and limits. *Public Administration Review*. 2024;84(6):1184–1199. https://doi.org/10.1111/puar.13799

[9] Mohanty A, Ramasamy AK, Verayiah R, et al. Smart grid and application of big data: opportunities and challenges. *Sustainable Energy Technologies and Assessments*. 2024;71:104011. https://doi.org/10.1016/j.seta.2024.104011

[10] Galeela M, Altamimi A, Kopsidas K, Hassan H. Probabilistic emergency demand response planning considering overhead lines time varying thermal ratings for reliability enhancements. *Sustainable Energy, Grids and Networks*. 2024;38:101305. https://doi.org/10.1016/j.segan.2024.101305

[11] Hua LG, Alghamdi H, Hafeez G, et al. Optimal energy management via day-ahead scheduling considering renewable energy and demand response in smart grids. *ISA Transactions*. 2024;154:268–284. https://doi.org/10.1016/j.isatra.2024.08.032

[12] Cabello Ruiz R, Jiménez Ramírez A, Escalona Cuaresma MJ, González Enríquez J. Hybridizing humans and robots: an RPA horizon envisaged from the trenches. *Computers in Industry*. 2022;138:103615. https://doi.org/10.1016/j.compind.2022.103615

[13] Wewerka J, Reichert M. Robotic process automation – a systematic mapping study and classification framework. *Enterprise Information Systems*. 2023;17(2). https://doi.org/10.1080/17517575.2021.1986862

[14] El-Gharib NM, Amyot D. Robotic process automation using process mining – a systematic literature review. *Data & Knowledge Engineering*. 2023:102229. https://doi.org/10.1016/j.datak.2023.102229

[15] Egger A, ter Hofstede AHM, Kratsch W, et al. Bot log mining: an approach to the integrated analysis of robotic process automation and process mining. *Information Systems*. 2024;126:102431. https://doi.org/10.1016/j.is.2024.102431

[16] Beerepoot I, Di Ciccio C, Reijers HA, et al. The biggest business process management problems to solve before we die. *Computers in Industry*. 2023;146:103837. https://doi.org/10.1016/j.compind.2022.103837

[17] Enarsson T, Enqvist L, Naarttijärvi M. Approaching the human in the loop – legal perspectives on hybrid human/algorithmic decision-making in three contexts. *Information & Communications Technology Law*. 2022;31(1):123–153. https://doi.org/10.1080/13600834.2021.1958860

[18] Selten F, Klievink B. Organizing public sector AI adoption: navigating between separation and integration. *Government Information Quarterly*. 2024;41(1):101885. https://doi.org/10.1016/j.giq.2023.101885

[19] 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

[20] Grimmelikhuijsen S. Explaining why the computer says no: algorithmic transparency affects the perceived trustworthiness of automated decision-making. *Public Administration Review*. 2022;82. https://doi.org/10.1111/puar.13483

[21] Haesevoets T, Verschuere B, Van Severen R, Roets A. How do citizens perceive the use of artificial intelligence in public sector decisions? *Government Information Quarterly*. 2024;41(1):101906. https://doi.org/10.1016/j.giq.2023.101906

[22] Fischer-Abaigar U, Kern C, Barda N, Kreuter F. Bridging the gap: towards an expanded toolkit for AI-driven decision-making in the public sector. *Government Information Quarterly*. 2024;41(4):101976. https://doi.org/10.1016/j.giq.2024.101976

[23] Haitsma L, Brink B. From human intervention to human involvement: a critical examination of the role of humans in (semi-)automated administrative decision-making. *Digital Government: Research and Practice*. 2025;6(3):33. https://doi.org/10.1145/3716173

[24] Hillo J, Vento I, Erkkilä T. Algorithmic governance: experimental evidence on citizens' and public administrators' legitimacy perceptions of automated decision-making. *Public Administration*. 2025;104(3):536–551. https://doi.org/10.1111/padm.70028

[25] Hasan A, Alenazy AA, Habib S, Husain S. Examining the drivers and barriers to adoption of e-government services in Saudi Arabia. *Journal of Innovative Digital Transformation*. 2024;1(2):139–157. https://doi.org/10.1108/JIDT-09-2023-0019

[26] Mamede HS, Martins CMG, da Silva MM. A lean approach to robotic process automation in banking. *Heliyon*. 2023;9(7):e18041. https://doi.org/10.1016/j.heliyon.2023.e18041

[27] Aoki N, Tatsumi T, Naruse G, Maeda K. Explainable AI for government: does the type of explanation matter to the accuracy, fairness, and trustworthiness of an algorithmic decision as perceived by those who are affected? *Government Information Quarterly*. 2024;41(4):101965. https://doi.org/10.1016/j.giq.2024.101965

[28] Rizk A, Lindgren I. Automated decision-making in public administration: changing the decision space between public officials and citizens. *Government Information Quarterly*. 2025;42(3):102061. https://doi.org/10.1016/j.giq.2025.102061

[29] Agostinelli S, Hohenadl T, Marrella A, Martínez-Rojas A. SmartRPA: generating software robots from user interface logs. *SoftwareX*. 2025;29:101995. https://doi.org/10.1016/j.softx.2024.101995

[30] Hegazy I. Leveraging advanced digital technologies for smart city development in Saudi Arabia: opportunities, challenges, and strategic pathways. *International Journal of Low-Carbon Technologies*. 2025;20:834–847. https://doi.org/10.1093/ijlct/ctaf044