Agentic AI for End-to-End Automation of the Data Quality Lifecycle: A DAMA- and NDMO-Aligned Governance Framework for Saudi Enterprises.

Abstract:

Agentic artificial intelligence advances data-quality automation beyond isolated profiling or cleansing by enabling goal-directed systems to plan, invoke tools, evaluate evidence, coordinate specialist agents, and escalate consequential decisions. This review develops a governance-centered framework for end-to-end automation of the enterprise data quality lifecycle within Saudi organizations. A structured integrative review integrates peer-reviewed research on data-quality assessment, monitoring, repair, data-centric artificial intelligence, autonomous agents, and data governance. The synthesis demonstrates that lifecycle automation is technically feasible when autonomous reasoning is separated from authoritative change. The proposed framework organizes specialized agents around quality definition, profiling, rule discovery, continuous supervision, root-cause analysis, repair, validation, publication, and learning. Concurrently, deterministic controls ensure policy enforcement and human accountability. DAMA-oriented responsibilities establish the enterprise operating model for ownership, stewardship, quality management, metadata, architecture, and security, while NDMO-oriented governance introduces Saudi-specific requirements for organizational accountability, classification, traceability, measurable controls, and evidence. The review contends that the practical objective should be maximum trustworthy quality throughput—the proportion of quality issues detected, explained, resolved, validated, and evidenced within accepted risk and service thresholds—rather than maximum autonomous repair. A risk-tiered assurance loop, provenance-rich action records, confidence-aware escalation, segregation of duties, and reversible state changes are identified as essential conditions for safe scaling.

Keywords: agentic AI; data quality lifecycle; data governance; DAMA; NDMO; Saudi Arabia; data-centric AI; autonomous agents; data quality automation

1. Introduction

Data quality has evolved from a back-office cleansing concern to a continuous control challenge. Enterprise decisions, analytics, regulatory reporting, artificial intelligence systems, customer operations, and automated workflows all require data that are sufficiently accurate, complete, consistent, timely, valid, and fit for their intended use. Current environments complicate this requirement, as data are distributed across cloud platforms, operational applications, warehouses, lakehouses, streaming systems, APIs, and machine-learning pipelines, with ownership and quality rules frequently fragmented. Recent reviews consequently conceptualize data quality as a multidimensional, contextual property rather than a single score, emphasizing measurement, monitoring, and organizational governance in addition to correction (Ehrlinger & Wöß, 2022; Reda et al., 2023; Priestley et al., 2023).

Automation has advanced primarily as a set of specialized capabilities. Large-scale verification systems can infer and enforce constraints, machine-learning approaches support error detection and repair, active learning reduces expert-labeling effort, and data-centric AI emphasizes systematic data improvement over model design alone (Schelter et al., 2018; Mahdavi et al., 2019; Abedjan & Mahdavi, 2020; Li et al., 2024; Jakubik et al., 2024). Patel et al. (2023) demonstrate that exploratory analysis and data-quality tasks can be partially automated, while Mecca et al. (2024) highlight the value of learning repair actions through human interaction. However, these capabilities typically address only one stage of the lifecycle. Enterprises continue to depend on manual intervention to connect signals across profiling, issue diagnosis, remediation, validation, stewardship, and control evidence.

Agentic AI significantly changes the architectural approach. Large language model (LLM)-based agents integrate perception, planning, memory, tool use, feedback, and coordination, while multi-agent systems enable role specialization and division of labor (Wang et al., 2024; Guo et al., 2024; Xi et al., 2025). In the context of data quality, agents can translate business expectations into candidate rules, determine which profilers to execute, correlate anomalies with lineage and recent changes, propose repairs, perform low-risk corrections via controlled tools, retest affected assets, compile stewardship evidence, and learn from approved resolutions. However, these same capabilities bring new risks, as a probabilistic planner may misdiagnose defects, apply unsafe repairs, propagate unsupported assumptions, or exercise excessive privileges. Consequently, agentic autonomy requires a more solid governance architecture than conventional workflow automation (Acharya et al., 2025; Barredo Arrieta et al., 2020).

This governance imperative is especially important for Saudi enterprises. Data governance distributes decision rights, accountability, and control mechanisms throughout organizational and technical boundaries (Abraham et al., 2019; Janssen et al., 2020; Walsh et al., 2022). DAMA-oriented practices supply a comprehensive enterprise model encompassing data governance, data quality, metadata, architecture, security, integration, and stewardship. Goel et al. (2024) illustrate how DAMA dimensions can structure governance responsibilities for analytical data. Simultaneously, the Saudi National Data Management Office (NDMO) emphasizes defined institutional mechanisms, control evidence, classification, traceability, and accountable implementation. Empirical research at King Saud University spotlights the strategic, executive, and organizational aspects of NDMO-aligned governance (Aljarallah & Alfadhel, 2025).

The main challenge is not whether AI can automate individual quality tasks, but how to integrate autonomous capabilities into an end-to-end lifecycle without compromising authority, explainability, or compliance. This review addresses this challenge by synthesizing research in data-quality engineering, data-centric AI, autonomous-agent architecture, and governance. It proposes a Saudi-enterprise reference framework in which agents orchestrate evidence and reversible actions, while deterministic policy gates and accountable human roles govern consequential state changes. The contribution is conceptual and design-oriented, identifying where autonomy is appropriate, where approval is still essential, how DAMA and NDMO frameworks might be jointly operationalized, and which metrics are necessary to evaluate trustworthy quality automation.

2. Aim of the Study

The study aims to develop a critical, implementation-oriented framework for using agentic AI to automate the enterprise data quality lifecycle, from quality definition and discovery through monitoring, diagnosis, remediation, validation, publication, and continuous improvement. The framework is designed to preserve accountable human ownership and deterministic control while aligning operational responsibilities with DAMA-oriented practice and the traceability, classification, evidence, and institutional-control expectations associated with NDMO governance in Saudi Arabia.

3. Research Objectives

Objective 1: Synthesize current evidence on data-quality assessment, monitoring, error detection, repair, data-centric AI, and autonomous agents to identify capabilities that can be composed into an end-to-end quality operating model.

Objective 2: Distinguish lifecycle activities suitable for autonomous execution from activities that require deterministic validation, steward approval, or accountable owner decisions, based on consequence and reversibility.

Objective 3: Map agent roles, data-quality controls, evidence requirements, and escalation routes to DAMA-oriented responsibilities and NDMO-oriented governance expectations relevant to Saudi enterprises.

Objective 4: Define a reference architecture and assurance loop that support provenance, least privilege, segregation of duties, confidence-aware escalation, rollback, and measurable quality outcomes.

Objective 5: Identify empirical research gaps and evaluation measures for testing whether agentic automation improves quality throughput without degrading correctness, explainability, security, or governance accountability.

4. Research Questions

RQ1: Which stages of the enterprise data quality lifecycle can specialized agents automate effectively, and where should authority remain approval-gated?

RQ2: How can agents combine rules, metadata, lineage, profiling results, historical incidents, and commercial context to diagnose and remediate quality defects without converting probabilistic reasoning into ungoverned change?

RQ3: How should DAMA-oriented roles and NDMO-oriented controls shape permissions, evidence, escalation, and accountability for agentic data-quality operations in Saudi enterprises?

RQ4: Which technical and managerial metrics can demonstrate that end-to-end agentic automation improves trustworthy quality throughput, resilience, and control assurance?

5. Review Methodology

A structured integrative review was selected because the research problem crosses database engineering, information quality, machine learning, autonomous agents, and organizational governance. A narrowly bounded systematic review of one discipline would omit necessary mechanisms, while an unstructured narrative approach would make source selection and synthesis difficult to audit. The protocol therefore combined concept-driven searching, backward and forward citation checking, publisher-level bibliographic verification, and thematic synthesis. The two supplied research papers were used only as structural and academic-writing references; their wording, argument sequence, tables, figures, and evidentiary claims were not treated as source material for this review.

Searches targeted peer-reviewed literature available through publisher platforms and major scholarly indexes associated with ACM, IEEE, Elsevier, Springer, Taylor & Francis, Frontiers, and conference proceedings. Search concepts were combined iteratively, including “data quality lifecycle,” “data quality assessment,” “data quality monitoring,” “data profiling,” “data cleaning,” “error detection,” “data repair,” “active learning data quality,” “data-centric AI,” “LLM agents,” “autonomous agents,” “multi-agent systems,” “data governance,” “DAMA,” “NDMO,” “Saudi data governance,” “explainable AI,” and related combinations. We prioritized publications from 2020-2025. Earlier studies were retained only when they provide foundational, still-relevant evidence for automated verification, error detection, repair, or governance.

Inclusion required a peer-reviewed journal or conference contribution with a resolvable DOI and direct relevance to at least one analytical domain: quality definition and assessment; profiling and monitoring; anomaly or error detection; correction and validation; data-centric AI; autonomous or multi-agent systems; enterprise data governance; DAMA-aligned responsibilities; NDMO implementation; or explainability and control. Studies were excluded when they were vendor marketing, unreviewed preprints without a verified peer-reviewed version, opinion-only articles, duplicates, or sources with bibliographic details that could not be independently checked. The final evidentiary corpus was deliberately limited to thirty DOI-verified peer-reviewed works. No database hit, screening, or exclusion counts are reported because a reproducible subscription-database export log was not preserved; inventing a PRISMA-style flow would create false methodological precision.

Quality appraisal considered publisher and venue credibility, clarity of method, evidence-to-claim fit, reproducibility of technical descriptions, relevance to enterprise-scale data operations, and acknowledgment of limitations. We gave greater interpretive weight to systematic reviews, established journal articles, empirical or design-science studies, and well-described conference research. The evidence was coded in two layers. The first mapped each study to lifecycle functions: define, discover/profile, assess, monitor, diagnose, repair, validate, publish, and learn. The second captured enabling or constraining mechanisms, including planning, tool use, active learning, confidence estimation, provenance, human review, least privilege, rollback, and governance decision rights.

Synthesis was comparative rather than vote-counting. Findings were examined for complementarities and boundary conditions, particularly where flexible AI reasoning can increase coverage but reduce determinism. Data-centric AI research was compared with traditional quality-management approaches; agent architecture was examined against evidence from automated quality tools; and governance literature was used to determine where autonomy should be limited. The resulting framework is therefore an inferential synthesis of validated component capabilities, not a claim that longitudinal production settings have already proven fully autonomous enterprise data-quality operations. This distinction is a methodological limitation and a design principle: proposed autonomy levels should be treated as hypotheses requiring controlled enterprise evaluation.

6. Thematic Literature Review and Critical Analysis

6.1 From Fragmented Quality Tools to an Agent-Orchestrated Lifecycle

Traditional data-quality programs separate quality definition, profiling, monitoring, remediation, and stewardship into different tools and teams. That division reflects the field's history but creates latency when defects cross system boundaries. Ehrlinger and Wöß (2022) show that tools vary materially in profiling, measurement, and continuous-monitoring capability, while Reda et al. (2023) confirm that assessment models remain context-dependent. Albrecht et al. (2023) similarly frame quality management as a sequence involving definition, assessment, and improvement rather than a single cleansing step. The implication is that end-to-end automation requires coordination across lifecycle states, not merely a more powerful detector.

An agentic model can treat a quality objective as a governed case. A quality-definition agent translates business expectations, critical-data-element status, prior controls, and usage context into candidate rules and service levels. A discovery agent profiles assets and records observed distributions, null patterns, keys, referential relationships, and freshness. A monitoring agent continuously evaluates rules and change patterns. When an issue appears, a diagnosis agent correlates the failure with lineage, schema changes, upstream incidents, deployments, and historical resolutions. A repair agent proposes or executes bounded corrections; a validation agent reruns independent checks; and a stewardship agent assembles evidence, assigns unresolved cases, and records the authoritative disposition. Table 1 summarizes this functional decomposition.

Table 1. Agentic roles across the end-to-end data quality lifecycle

Lifecycle stage	Primary agent role	Evidence and actions	Recommended authority
Define	Quality-policy agent	Business terms, critical data elements, prior rules; proposes dimensions, thresholds and SLOs	Steward/owner approval for new or changed authoritative rules
Discover & profile	Profiler agent	Schemas, distributions, keys, nulls, referential patterns, freshness	High autonomy for observation; evidence retained
Assess & monitor	Monitoring agent	Executed rules, baselines, drift, anomalies, service levels	High autonomy for detection and alerting
Diagnose	Root-cause agent	Lineage, changes, incidents, ownership, historical resolutions	Autonomous hypothesis generation; confidence-aware escalation
Repair	Remediation agent	Approved playbooks, deterministic transforms, candidate corrections	Automatic only for reversible low-risk actions; approval for consequential changes
Validate	Independent validator	Re-tests, reconciliations, downstream impact checks	Separate from proposing agent; deterministic pass/fail where possible
Publish & close	Stewardship agent	Decision packet, approvals, provenance, before/after state	Human approval for material closures and rule changes
Learn	Calibration agent	Accepted/rejected proposals, recurrence, performance metrics	Update governed playbooks and thresholds; no unconstrained memory

This decomposition should be specialized rather than monolithic. Multi-agent research indicates that role separation can improve task decomposition and coordination, but it also introduces communication and propagation risks (Guo et al., 2024). The appropriate enterprise response is to couple specialization with segregation of duties: the agent that proposes a repair should not be the sole validator of that repair, and the agent that interprets a policy should not directly change access or retention controls. Autonomous intelligence becomes operationally useful when it reduces coordination cost while preserving independent checks.

6.2 Profiling, Rule Discovery and Continuous Monitoring

Profiling, measurement, and monitoring are the strongest candidates for high autonomy because they primarily observe rather than alter authoritative data. Large-scale verification research demonstrates that constraints can be executed automatically and that historical quality metrics can support anomaly detection (Schelter et al., 2018). More recent work on automated exploratory analysis confirms that algorithmic systems can recommend and execute parts of data preparation and quality assessment (Patel et al., 2023). Streaming-quality research adds an important temporal dimension: controls must operate continuously and use contextual information because defects emerge through changing sources, operational states, and data distributions (Mirzaie et al., 2023; Alwan et al., 2022).

Agentic AI adds decision-making around these capabilities. A profiler agent can choose relevant metrics by data type and use case, compare current results with baselines, and request deeper tests when signals conflict. A rule-discovery agent can propose constraints from historical observations and business descriptions, but those constraints should remain candidates until validated against domain expectations. Priestley et al. (2023) show why stage-specific quality requirements matter in machine-learning pipelines; a field acceptable for exploratory analysis may be unacceptable for training, production scoring, or regulated reporting. Quality rules therefore need declared scope, purpose, owner, severity, and evidence, not merely a technical expression.

The first governance principle is evidence hierarchy. Direct observations, executed tests, approved rules, and verified metadata should carry greater authority than agent inference. The second is state separation. A system should distinguish observed defect, suspected defect, confirmed issue, proposed repair, validated repair, and approved closure. This prevents a model-generated hypothesis from becoming a silent operational fact. Continuous monitoring then becomes a controlled event stream from which agents may reason, rather than an unrestricted permission to alter enterprise records.

6.3 Root-Cause Diagnosis, Repair and Human-in-the-Loop Resolution

Diagnosis and repair offer greater value but also greater risk. Error-detection research shows the benefit of learning from limited expert feedback: Raha reduces configuration dependence by combining multiple detectors and actively selecting examples (Mahdavi et al., 2019). Baran extends learning to correction by exploiting contextual information and transfer learning (Abedjan & Mahdavi, 2020). Active-learning research more broadly shows how quality-control systems can focus scarce expert attention on informative cases (Li et al., 2024). BUNNI demonstrates a complementary pattern in which machine learning recommends repair actions while users resolve ambiguous rule violations (Mecca et al., 2024). Together, these studies support semi-autonomous remediation, not unrestricted self-healing.

Agentic diagnosis should reason across evidence sources that conventional cleaners often treat separately. A sudden completeness failure may originate in an upstream API change; a consistency defect may reflect duplicated master-data identifiers; a freshness failure may result from orchestration delay rather than corrupt records. Lineage, deployment events, ownership metadata, incident history, and quality metrics can therefore reduce false repairs. An agent may form competing root-cause hypotheses, test them with controlled queries, and rank them by evidence. However, the system should preserve uncertainty rather than fabricate a single narrative when evidence is incomplete.

Repair authority should be determined by consequence, reversibility, and confidence. Recomputing a derived field from an approved deterministic formula is lower risk than replacing a customer master value inferred from neighboring records. Reprocessing a failed pipeline is more reversible than deleting records. The framework therefore uses graduated autonomy: low-impact, deterministic, reversible repairs may execute automatically within scoped transactions; probabilistic corrections require confidence thresholds and post-repair validation; changes affecting legal status, financial reporting, sensitive classification, external sharing, or master-record identity require accountable human approval. The objective is trustworthy throughput, not an impressive autonomous-action rate.

6.4 Data-Centric AI and the Architecture of Governed Quality Autonomy

Data-centric AI provides the conceptual bridge between quality engineering and agentic automation. Jarrahi et al. (2023) argue for systematic, iterative data improvement, emphasizing human-centered data work and interaction between domain experts and AI. Jakubik et al. (2024) likewise define data-centric AI as the deliberate engineering of data quality and quantity, while Kumar et al. (2024) identify opportunities and challenges in shifting attention from models to data. Zha et al. (2025) extend this view across training-data development, inference-data development, and data maintenance. These contributions imply that quality automation should operate as a feedback system rather than a one-time cleaning project.

An enterprise agentic loop can convert validated outcomes into governed learning. When a steward accepts a repair, the system may update a reusable playbook, strengthen a rule, or adjust a triage policy; when a proposal is rejected, the reason becomes calibration evidence. Crucially, learning should not mean unconstrained accumulation of model memory. The durable knowledge base should consist of approved rules, quality dimensions, ownership, thresholds, lineage, resolution patterns, and performance statistics under version and retention control. This approach makes organizational learning inspectable and reversible.

Figure 1 presents the proposed architecture. Operational data and metadata feed deterministic profilers, rule engines, lineage services, and observability collectors. Specialist agents plan and coordinate quality work but act through a controlled tool gateway. A shared evidence layer stores test results, hypotheses, actions, confidence, approvals, and before-and-after states. A governance plane then applies policy, identity, risk tier, and approval requirements before changing any authoritative data or quality status. This separation between reasoning and authority is the framework's architectural core.

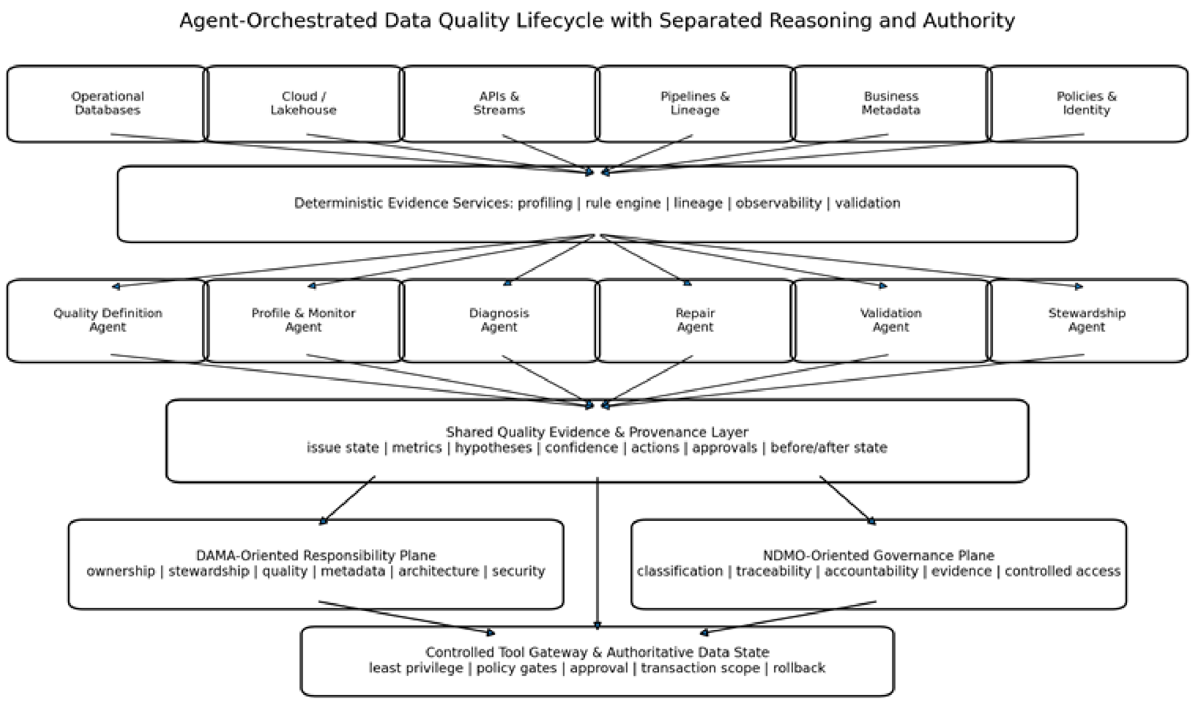

Figure 1. Agent-orchestrated data quality lifecycle architecture with separated reasoning and authority.

6.5 DAMA- and NDMO-Aligned Governance for Saudi Enterprises

DAMA and NDMO contribute complementary governance views. Data-governance research consistently shows that decision rights, accountability, and organizational scope determine whether technology creates sustainable value (Abraham et al., 2019; Walsh et al., 2022; Bližnák et al., 2024). Bernardo et al. (2024) further connect data governance and quality management, reinforcing that quality outcomes depend on governance structures rather than isolated tools. A DAMA-oriented operating model therefore defines who owns critical data, who stewards definitions and rules, who manages architecture and integration, who controls security, and who is accountable for quality performance. Agents should operate as delegated processors within these responsibilities, not as new unowned decision makers.

For Saudi enterprises, NDMO alignment raises the standard of demonstrability. The King Saud University study indicates that implementation involves strategic and executive mechanisms as well as organizational challenges (Aljarallah & Alfadhel, 2025). Consequently, an agentic quality system should show which asset was evaluated, which rule or policy was applied, what evidence supported the issue, which agent and model version acted, which tools were invoked, what changed, who approved it, and whether validation succeeded. Governance becomes a machine-observable evidence chain rather than a retrospective narrative.

Goel et al. (2024) explicitly structure governance considerations using DAMA dimensions, demonstrating how a recognized management framework can organize data responsibilities around analytical processes. The same logic applies to quality automation. Table 2 maps common agentic controls to governance purposes: least privilege protects authority boundaries; independent validation supports quality management; lineage and action provenance support traceability; confidence-aware escalation protects stewardship decisions; and immutable logs provide control evidence. Metadata is itself part of trust: Qin and Yu (2024) argue that trustworthy AI depends on metadata describing data, models, pipelines, and artifacts. Agentic quality systems consequently require metadata about the agents and quality actions themselves.

Table 2. Governance mapping for DAMA- and NDMO-aligned agentic quality controls

Control objective	DAMA-oriented responsibility	NDMO-oriented governance emphasis	Agentic implementation
Accountability	Owner/steward decision rights	Institutional responsibility and evidence	Named accountable owner; agent identity; approval trail
Quality management	Dimensions, rules, measurement, issue management	Measurable controls and traceable execution	Versioned rules, SLOs, issue states, independent validation
Metadata & lineage	Definitions, metadata, integration context	Classification and traceability	Evidence graph linking asset, rule, lineage, action and outcome
Security & authority	Access control and segregation of duties	Controlled access and demonstrable compliance	Least privilege, allow-listed tools, scoped transactions, kill switch
Stewardship	Exception handling and semantic adjudication	Documented governance mechanisms	Confidence thresholds, escalation queues, reason-coded decisions
Assurance	Monitoring and governance reporting	Auditability and control evidence	Immutable action logs, before/after state, rollback and evidence completeness metrics

6.6 Assurance, Explainability and Failure Containment

The main failure modes arise when reasoning is mistaken for authority. An agent can hallucinate a business rule, infer the wrong root cause, select an inappropriate repair, or propagate an unsupported claim through other agents. Autonomous-agent surveys emphasize planning, memory, tool use, adaptation, and evaluation as core capabilities, but those same mechanisms enlarge the operational attack and error surface (Wang et al., 2024; Xi et al., 2025; Acharya et al., 2025). Explainability therefore needs to be operational, not rhetorical. Barredo Arrieta et al.

(2020) show the broader importance of interpretable and responsible AI; for data quality, the most useful explanation is compact evidence: failed tests, source records, lineage paths, applied rules, alternatives considered, confidence, and reason for escalation.

Structured enterprise context can strengthen grounding. Knowledge graphs can connect data elements to business concepts, systems, owners, policies, and lineage. Pan et al. (2024) show that knowledge graphs and LLMs are complementary: graphs can ground and constrain model reasoning while models can help construct or query graph structures. In quality operations, such grounding can prevent semantic mistakes, such as applying a "customer identifier" rule to a technically similar but legally distinct identifier. Nevertheless, graph context should guide reasoning rather than automatically legitimize it; stale or incorrect governance metadata can produce confidently wrong actions.

Figure 2 therefore formalizes an assurance loop: sense, plan, simulate or propose, policy-check, execute within a bounded transaction, verify independently, then promote, escalate, or roll back. Each cycle emits evidence and performance feedback. Risk tier determines whether execution is automatic or approval-gated. The loop also supports operational kill switches, rate limits, change windows, and transaction scopes. This design treats agent autonomy as a sequence of auditable state transitions instead of a binary label applied to the whole system.

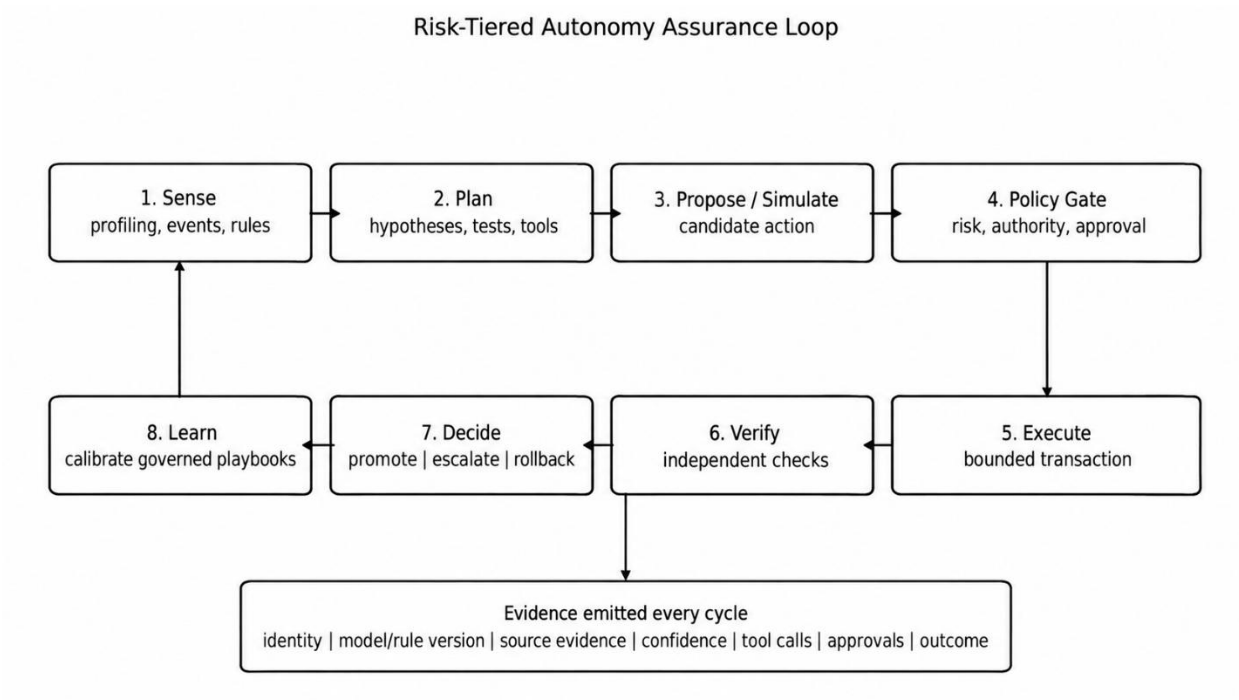

Figure 2. Risk-tiered autonomy assurance loop for data quality state transitions.

7. Discussion

The synthesis indicates that agentic AI is most credible as an orchestration and reasoning layer on top of established quality, metadata, lineage, and governance infrastructure. The literature does not support replacing deterministic quality checks with language-model judgment. Instead, agents add value where quality work is coordination-intensive: selecting tools, interpreting context, linking signals across systems, generating hypotheses, sequencing tests, preparing decision packets, and routing exceptions. Existing quality technologies provide observable evidence and constrained actions; agents provide adaptive planning around them.

A second finding is that lifecycle stages differ structurally in their suitability for autonomy. Profiling and monitoring are observation-dominant and therefore can be highly automated. Rule creation, root-cause diagnosis, and semantic interpretation are inference-dominant and require calibration and review. Repair is action-dominant and must be governed by reversibility and business consequence. Validation should be independent of the proposing agent, while closure and rule changes should preserve accountable stewardship. This stage-specific model is more defensible than a single enterprise-wide autonomy level.

A third finding is that governance quality is part of data quality. A repair that makes values statistically consistent can still be invalid if it violates ownership, approved definitions, classification, retention, or reporting policy. DAMA-oriented responsibilities provide the organizational context for legitimate quality decisions, while NDMO-oriented requirements increase the importance of traceability and evidence in Saudi enterprises. The framework therefore treats every consequential agent action as a governed data event with identity, policy basis, evidence, and outcome.

Finally, the relevant performance target is trustworthy quality throughput. Conventional automation metrics such as number of rules executed or percentage of tickets auto-closed can reward unsafe behavior. A better measurement set combines detection precision and recall, time-to-detect, time-to-diagnose, validated repair rate, recurrence, false-repair rate, rollback rate, steward effort, service-level compliance, evidence completeness, and policy-exception frequency. These measures make it possible to test whether autonomy improves enterprise outcomes rather than simply increasing machine activity.

Implementation evidence should therefore be collected as a longitudinal governance dataset linking each quality event to asset criticality, detected dimension, root cause, proposed action, approval path, execution result, validation result, recurrence, business impact, and control exception. This evidence enables enterprises to compare autonomous and manual handling fairly, calibrate risk thresholds, identify systematic upstream weaknesses, and show whether automation improves reliability rather than shifting effort into hidden review or rework. It also supports governance learning because recurring exceptions can be converted into better rules, better metadata, stronger integration controls, or redesigned source processes instead of repeated downstream correction in Saudi environments.

8. Research Gaps and Future Research

First, peer-reviewed evidence from longitudinal production deployments of fully agentic data-quality operations remains scarce. Research should compare rule-based workflows, single-agent orchestration, and multi-agent architectures on shared enterprise-like benchmarks containing heterogeneous schemas, evolving pipelines, ambiguous business rules, and controlled governance constraints. Evaluations should report both quality outcomes and operational costs, including human review time and rollback frequency.

Second, research is needed on uncertainty-aware quality provenance. Agentic systems should distinguish observed facts, inferred causes, proposed rules, simulated repairs, approved changes, and superseded decisions. Future work should examine how confidence propagates across collaborating agents and how independent validators detect correlated model errors. Third, governance requirements need stronger machine-readable representations so that DAMA responsibilities and NDMO controls can be compiled into permissions, evidence schemas, approval thresholds, and continuously testable policies without reducing governance to superficial checklists.

Fourth, Saudi-enterprise studies should evaluate Arabic-English business terminology, national data-classification contexts, cross-entity data sharing, and regulated reporting scenarios. Fifth, organizational research should examine how agentic automation changes stewardship skills and accountability: whether stewards become more effective exception managers or instead defer excessively to automated recommendations. These questions require mixed-method research combining technical benchmarks, audit evidence, interviews, and operational outcome measures.

9. Practical, Managerial and Policy Implications

Organizations should begin with bounded, evidence-rich use cases: automated profiling, quality-rule execution, freshness monitoring, anomaly triage, duplicate detection, diagnostic evidence gathering, and stewardship routing. Repair automation should initially be limited to deterministic and reversible actions. Before scaling, enterprises should establish agent identities, scoped tool permissions, transaction boundaries, independent validation, rollback procedures, and a provenance store that records the complete quality case.

Managers should define an agent RACI alongside the human data-governance RACI. Each agent requires an accountable owner, permitted domains, authorized tools, quality objectives, escalation thresholds, and shutdown procedures. Quality dashboards should distinguish machine-proposed, machine-executed, human-approved, rolled-back, and unresolved outcomes. Procurement should require exportable evidence, open interfaces, model and rule versioning, and separation between agent inference and authoritative data state.

For Saudi organizations, policy and control alignment should be designed into the workflow. NDMO-relevant classifications, responsible roles, quality rules, approvals, lineage, and remediation evidence to affected data assets, and retain them for assurance. DAMA-oriented domains can organize operational ownership, while NDMO-oriented controls can define proof requirements. This combination allows enterprises to pursue automation without treating governance as a post-implementation documentation exercise.

10. Limitations

This review is integrative rather than exhaustive. The thirty-source corpus was deliberately curated and DOI-verified, and direct evidence on end-to-end autonomous quality systems remains limited. Several recommendations therefore combine validated component capabilities into a reference architecture that still requires production testing. Peer-reviewed NDMO-specific research is comparatively sparse, so Saudi alignment is interpreted through available implementation evidence and governance principles rather than claimed as a complete legal or regulatory analysis. Rapid evolution in agent platforms may also make specific technical patterns obsolete faster than the underlying governance principles.

11. Conclusion

Agentic AI can extend data-quality automation from isolated tools into a coordinated lifecycle, but only when autonomy is treated as governed delegation. Evidence from profiling, monitoring, error detection, repair, active learning, and data-centric AI supports automation of substantial operational work. Agent research adds planning, tool use, memory, and multi-agent coordination. Together, these capabilities can reduce detection and resolution latency, but they also create new failure paths when probabilistic reasoning can mutate authoritative data without independent control.

The proposed framework separates reasoning from authority. Specialist agents define candidate rules, profile and monitor data, diagnose defects, propose or execute bounded repairs, validate outcomes, and learn from approved resolutions. Deterministic gateways enforce least privilege, policy constraints, transaction scope, approval requirements, and rollback. Every quality case preserves evidence, confidence, identity, tool calls, before-and-after state, validation, and ownership. Risk-tiered autonomy then permits fast automation for reversible low-impact actions while retaining human approval for consequential decisions.

DAMA and NDMO provide complementary foundations for Saudi enterprises. DAMA-oriented responsibilities clarify who owns quality, metadata, architecture, integration, security, and stewardship. NDMO-oriented governance strengthens the requirement that those responsibilities and controls be demonstrable through classification, traceability, institutional accountability, and durable evidence. The resulting objective is not maximum machine independence. It is maximum trustworthy quality throughput: faster and broader detection, diagnosis, remediation, and validation while maintaining correctness, explainability, security, and accountable control.

References

1. Abraham, R., Schneider, J., & vom Brocke, J. (2019). Data governance: A conceptual framework, structured review, and research agenda. International Journal of Information Management, 49, 424-438. https://doi.org/10.1016/j.ijinfomgt.2019.07.008
2. Abedjan, Z., & Mahdavi, M. (2020). Baran: Effective error correction via a unified context representation and transfer learning. Proceedings of the VLDB Endowment, 13(12), 1948-1961. https://doi.org/10.14778/3407790.3407801
3. Acharya, D. B., Kuppan, K., & Bhaskaracharya, D. (2025). Agentic AI: Autonomous intelligence for complex goals-A comprehensive survey. IEEE Access, 13, 18912-18936. https://doi.org/10.1109/ACCESS.2025.3532853
4. Albrecht, R., Overbeek, S., & van de Weerd, I. (2023). Designing a data quality management framework for CRM platform delivery and consultancy. SN Computer Science, 4, Article 742. https://doi.org/10.1007/s42979-023-02196-z
5. Aljarallah, Y., & Alfadhel, M. (2025). Data governance evaluation of the Data Management Office at King Saud University based on the National Data Management Office standards. Journal of Information Systems Engineering and Management, 10(10s), 794-805. https://doi.org/10.52783/jisem.v10i10s.1530
6. Alwan, A. A., Ciupala, M. A., Brimicombe, A. J., Ghorashi, S. A., Baravalle, A., & Falcarin, P. (2022). Data quality challenges in large-scale cyber-physical systems: A systematic review. Information Systems, 105, 101951. https://doi.org/10.1016/j.is.2021.101951
7. Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020).

Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82-115. https://doi.org/10.1016/j.inffus.2019.12.012

8. Bernardo, B. M. V., São Mamede, H., Barroso, J. M. P., & Santos, V. M. P. D. (2024). Data governance & quality management-Innovation and breakthroughs across different fields. Journal of Innovation & Knowledge, 9(4), 100598. https://doi.org/10.1016/j.jik.2024.100598
9. Bližnák, K., Munk, M., & Pilková, A. (2024). A systematic review of recent literature on data governance (2017-2023). IEEE Access, 12, 149875-149888. https://doi.org/10.1109/ACCESS.2024.3476373
10. Ehrlinger, L., & Wöß, W. (2022). A survey of data quality measurement and monitoring tools. Frontiers in Big Data, 5, 850611. https://doi.org/10.3389/fdata.2022.850611
11. Goel, K., Martin, N., & ter Hofstede, A. H. M. (2024). Demystifying data governance for process mining: Insights from a Delphi study. Information & Management, 61(5), 103973. https://doi.org/10.1016/j.im.2024.103973
12. Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N. V., Wiest, O., & Zhang, X. (2024). Large language model based multi-agents: A survey of progress and challenges. Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, 8048-8057. https://doi.org/10.24963/ijcai.2024/890
13. Jakubik, J., Vössing, M., Kühl, N., Walk, J., & Satzger, G. (2024). Data-centric artificial intelligence. Business & Information Systems Engineering, 66, 507-515. https://doi.org/10.1007/s12599-024-00857-8
14. Janssen, M., Brous, P., Estevez, E., Barbosa, L. S., & Janowski, T. (2020). Data governance: Organizing data for trustworthy artificial intelligence. Government Information Quarterly, 37(3), 101493. https://doi.org/10.1016/j.giq.2020.101493
15. Jarrahi, M. H., Memariani, A., & Guha, S. (2023). The principles of data-centric AI. Communications of the ACM, 66(8), 84-92. https://doi.org/10.1145/3571724
16. Kumar, S., Datta, S., Singh, V., Singh, S. K., & Sharma, R. (2024). Opportunities and challenges in data-centric AI. IEEE Access, 12, 33173-33189. https://doi.org/10.1109/ACCESS.2024.3369417

17. Li, N., Qi, Y., Li, C., & Zhao, Z. (2024). Active learning for data quality control: A survey. Journal of Data and Information Quality, 16(2), Article 11, 1-45. https://doi.org/10.1145/3663369
18. Mahdavi, M., Madden, S., Abedjan, Z., Ouzzani, M., Tang, N., Castro Fernandez, R., & Stonebraker, M. (2019). Raha: A configuration-free error detection system. Proceedings of the 2019 International Conference on Management of Data, 865-882. https://doi.org/10.1145/3299869.3324956
19. Mecca, G., Papotti, P., Santoro, D., & Veltri, E. (2024). BUNNI: Learning repair actions in rule-driven data cleaning. Journal of Data and Information Quality, 16(2), Article 12, 1-31. https://doi.org/10.1145/3665930
20. Mirzaie, M., Behkamal, B., Allahbakhsh, M., Paydar, S., & Bertino, E. (2023). State of the art on quality control for data streams: A systematic literature review. Computer Science Review, 48, 100554. https://doi.org/10.1016/j.cosrev.2023.100554
21. Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J., & Wu, X. (2024). Unifying large language models and knowledge graphs: A roadmap. IEEE Transactions on Knowledge and Data Engineering, 36(7), 3580-3599. https://doi.org/10.1109/TKDE.2024.3352100
22. Patel, H., Guttula, S., Gupta, N., Hans, S., Mittal, R. S., & N., L. (2023). A data-centric AI framework for automating exploratory data analysis and data quality tasks. Journal of Data and Information Quality, 15(4), Article 44, 1-26. https://doi.org/10.1145/3603709
23. Priestley, M., O'Donnell, F., & Simperl, E. (2023). A survey of data quality requirements that matter in ML development pipelines. Journal of Data and Information Quality, 15(2), Article 11, 1-39. https://doi.org/10.1145/3592616
24. Qin, J., & Yu, B. (2024). Metadata in trustworthy AI: From data quality to ML modeling. Proceedings of the International Conference on Dublin Core and Metadata Applications, 2024. https://doi.org/10.23106/dcmi.953354037
25. Reda, O., Chaouni Benabdellah, N., & Zellou, A. (2023). A systematic literature review on data quality assessment. Bulletin of Electrical Engineering and Informatics, 12(6), 3736-3757. https://doi.org/10.11591/eei.v12i6.5667
26. Schelter, S., Lange, D., Schmidt, P., Celikel, M., Bießmann, F., & Grafberger, A. (2018). Automating large-scale data quality verification. Proceedings of the VLDB Endowment, 11(12), 1781-1794. https://doi.org/10.14778/3229863.3229867

27. Walsh, M. J., McAvoy, J., & Sammon, D. (2022). Grounding data governance motivations: A review of the literature. Journal of Decision Systems, 31(sup1), 282-298. https://doi.org/10.1080/12460125.2022.2073637
28. Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J. (2024). A survey on large language model based autonomous agents. Frontiers of Computer Science, 18(6), 186345. https://doi.org/10.1007/s11704-024-40231-1
29. Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., Yin, Z., Dou, S., Weng, R., Qin, W., Zheng, Y., Qiu, X., Huang, X., Zhang, Q., & Gui, T. (2025). The rise and potential of large language model based agents: A survey. Science China Information Sciences, 68, 121101. https://doi.org/10.1007/s11432-024-4222-0
30. Zha, D., Bhat, Z. P., Lai, K.-H., Yang, F., Jiang, Z., Zhong, S., & Hu, X. (2025). Data-centric artificial intelligence: A survey. ACM Computing Surveys, 57(5), 1-42. https://doi.org/10.1145/3711118