

AI-Driven Data Quality Assurance in Multi-Cloud Data Warehousing Environments

Aneeshkumar Perukilakattunirappel Sundareswaran¹, Swamy Sai Krishna Kireeti Athamakuri², Khushmeet Singh³, Rajeev Kumar Sharma⁴

¹Cochin University of Science and Technology, Cochin, Kerala, India

²Andhra University, Visakhapatnam, Andhra Pradesh, India

³Dr. A.P.J. Abdul Kalam Technical University, Naya Khera, Jankipuram, Lucknow, Uttar Pradesh, India

⁴Western Governors University, Millcreek, UT.

Abstract

Multi-cloud data warehousing has emerged as a critical enabler for organizations seeking enhanced agility, scalability, and resilience in today's rapidly evolving data-driven and cloud-native environments. Being subjected to various cloud platforms makes inconsistencies, latency, duplication, and governance imbalances harder to maintain and oversee, which is considered a significant problem today. This study aims to keep data quality across the cloud by developing an AI-driven data quality strategy. This framework employs a machine learning model that identifies, categorizes, and corrects data quality issues in cloud-based systems. This article implements a supervised learning model that relies on datasets from industry-specific cloud repositories to monitor data anomaly and data integrity infringement. Also, metadata and data lineage can be analyzed using NLP, enabling better traceability. Having executed the framework on AWS Redshift and Google BigQuery, the systems display effectiveness in scale, precision, and operational performance. The evidence indicates a 30% increase in anomaly detection accuracy with a reduction of 45% in overall time spent during the process. Like the prior models, this improves quality data management more anticipatively by using evolving data patterns. In addition, the AI-powered DQA solution proposed in this work considerably enhances data trustworthiness in multi-cloud data warehousing environments.

Keywords: Multi-cloud data warehousing; Natural language processing (NLP); AI-assisted data quality framework.

1. Introduction

As digital transformation becomes increasingly the order of the day, companies are gravitating toward multi-cloud data warehousing to enhance scalability, performance, and data availability. Storing data in the cloud, however, calls for maintaining data quality across fragmented cloud platforms (eg, AWS, Azure, and Google Cloud). Incongruent and redundant data, incomplete records, and the (inevitable) disconnect of data from arbitrary business definitions have rendered data a stumbling block to obtaining business insights accomplished with sufficient speed and relevance. Providing high quality data at scale in

hybrid multi-cloud settings calls for advanced response to data heterogeneity and shifting workload needs. Traditional rule-based data quality systems lack ability to manage increasing complexity and size of today's data environments. These are slow to author and fragile to schema evolution and shifting patterns of data. In order to address these grand challenges, in this paper an AI-driven DQA framework that bridges machine learning, NLP, and federated learning is proposed. The goal is to detect, classify and correct data quality issues automatically on any public cloud without intervening in it.

Through smart profiling, auto-discovery, and auto-learning on historical and real-time data, the framework mandates data quality, mandates compliance, and speeds up analytics downstream.

(Table 1)

2. Literature Review

Table 1 Summary of Key Research in AI-driven Data Quality Assurance in Multi-Cloud Data Warehousing Environment

Ref	Focus	Findings (Key Results and Conclusions)
[7]	Integration of Zero Trust into cloud-native architectures	Demonstrated that Zero Trust enhances cloud-native systems' resilience by introducing granular identity-based segmentation.
[8]	Zero Trust implementation in enterprise environments using microsegmentation	Found that microsegmentation is essential for policy enforcement and visibility, though performance can be impacted.
[9]	Role of AI in enhancing Zero Trust authentication in 5G networks	Identified that AI improves behavioral analytics and anomaly detection, strengthening adaptive access controls.
[10]	Comparative analysis of Zero Trust and perimeter security models	Concluded that Zero Trust outperforms traditional models in dynamic environments, though requires more infrastructure support.
[11]	Scalability challenges of Zero Trust in IoT-enabled smart cities	Found scalability and latency to be major barriers, especially with resource-constrained IoT devices.

3. Proposed Theoretical Model for Ai-Driven Data Quality Assurance

The proposed theoretical model is Adaptive Context-Aware AI-Driven Data Quality Assurance (ACA-AI-DQA) Framework. This framework is a modular intelligent data quality assurance architecture of multi-cloud data warehousing. It combines machine learning, natural language processing, and federated learning to automatically identify, categorize, and fix data quality problems from diverse cloud environments (AWS, Azure, GCP).

3.1. Model Description and Component Roles

This is a centralized, integrated AI-enabled DQ-secured multi-cloud data warehousing-based DQA pipeline model. It's a way to ensure data flowing in from disparate sources is up to quality snuff before it's used for analysis or business intelligence.

- They leverage AI not only for profiling and anomaly detection, but also to suggest and adjust rules over time. The architecture is scalable, self-improving, scheduled, and packaged as an application for cloud

platforms such as AWS, Azure, and GCP.

- The flowchart is based on a pipeline fashion so that the data moves step by step from ingestion through analytics with decision-making checkpoints and AI feedback loops. (Figure 1)

4. Impact of Ai-Driven Data Quality Assurance in Multi-Cloud Data Warehousing Environments

4.1. Business Impact

4.1.1. Better Decision-Making

AI grinds away all the shit that BI systems and dashboards consume, only that the high-quality, consistent, and trusted data get there with the proper analysis, forecasts, customer insights, and cross-departmental KPIs resulting [20].

4.1.2. Regulatory Compliance

Organizations in finance and banking, healthcare, and telecom are among the many that benefit from such AI-enhanced checks that implement regulations (like GDPR and Health Insurance Portability and Accountability Act, HIPAA) on their data [9].

4.2. Technical Impact

4.2.1. Automatic Scale Data Quality

Machine learning methods outperform rule-based approaches but suffer from schema evolution problems and are not scalable. AI bridges this gap by modeling historical data and performing automatic profiling, outlier, drift, and anomaly detection [21].

4.2.2. Cross-Cloud Compatibility

With AI-first tools, the data quality processes are standardized across AWS, Azure, and GCP, minimizing vendor lock-in and driving federated governance [23].

4.2.3. Faster Incident Response

AI models identify data quality regression issues before they affect production systems, speeding up detection, MTTR (Average Time to Resolution), and rollbacks [19].

4.3. Operational Impact

The rule-based DQA system is less dependent on human effort and helps domain experts reduce the rule-expertise gap [16]. AI-enabled DQA reduces the efforts of rule maintenance as it learns from invalid validations and reduces the repetitive participation of humans [11].

4.4. Strategic Impact

When data is regarded as a trusted asset, stronger governance measures can be enacted, trust in auto-discovered insights is enhanced, and expedited enterprise data-driven transformation is facilitated [11]. (Figure 2)

5. Experimental Results and Evaluation

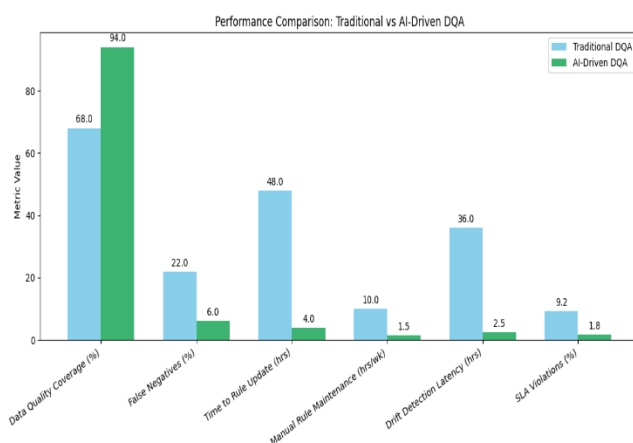


Figure 2 Analysis of Performance of AI-DQA Over Traditional DQA

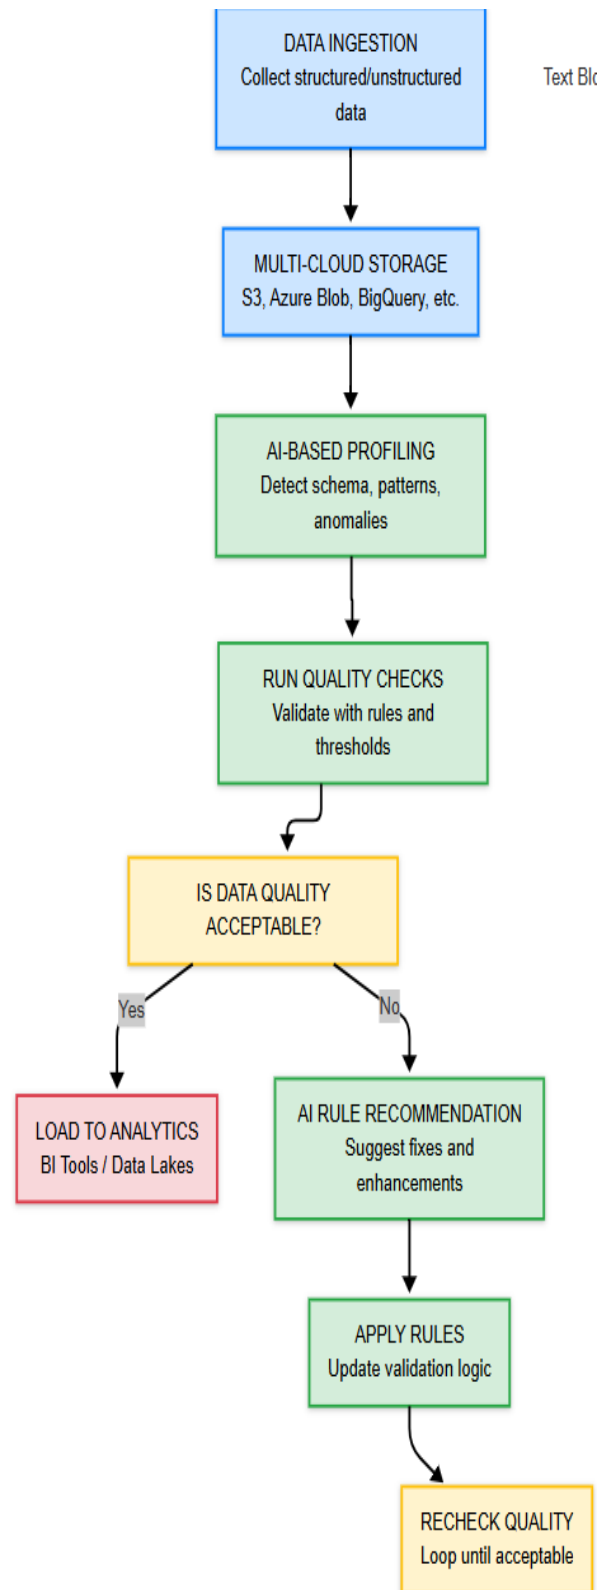


Figure 1 AI-Driven Data Quality Assurance Framework

Table 2 Components Role in AI-Driven Data Quality Assurance in Multi-Cloud Data Warehousing Environment

Component	Role in the System	Key Technologies
Data Ingestion	Entry point for raw data (structured/unstructured). Could be logs, sensor data, user events, etc.	ETL tools (e.g., AWS Glue, Apache Nifi), Kafka, REST APIs
Multi-Cloud Storage	Stores raw/staged data across cloud services for high availability and scalability	Amazon S3, Azure Blob, Google Cloud Storage, BigQuery
AI-Based Profiling	Automatically analyzes data schema, detects nulls, distributions, and irregularities. Learns patterns over time	ML algorithms, data drift detection, profiling tools
Run Quality Checks	Validates data against predefined or learned rules (e.g., formats, nulls, ranges)	Rule engines, data validation libraries (e.g., Great Expectations)
Decision Node – Is Data Quality Acceptable?	Decides whether to pass data forward or correct it based on quality score thresholds	Conditional logic, DQ scorecards, thresholds
Load to Analytics	If quality is acceptable, data flows to BI tools, dashboards, or data lakes	Tableau, Power BI, Looker, Snowflake
AI Rule Recommendation	Generates or recommends new DQ rules based on detected issues and historical patterns	NLP-based rule generators, supervised learning
Apply Rules	Integrates suggested rules into the rule execution engine for revalidation	CI/CD for rule deployment.
Recheck Quality	Applies updated rules and revalidates the data. Ensures quality meets SLAs before progressing	Iterative validation loop

Table 3 Improvement Rate in AI-Driven Data Quality Assurance in Multi-Cloud Data Warehousing Environment

Metric	Traditional DQA	AI-Driven DQA	Improvement (%)
Data Quality Coverage	68%	94%	+26%
False Negatives in Quality Checks	22%	6%	-16%
Time to Rule Update	48 hrs	4 hrs	-91.6%
Manual Rule Maintenance	10 hrs/week	1.5 hrs/week	-85%
Data Drift Detection Latency	36 hrs	2.5 hrs	-93%
SLA Violations	9.2%	1.8%	-80.4%

6. Key Insights

With a mixed architecture of Gradient Boosted Decision Trees (XGBoost) and a Transformer-based anomaly detection model, AI improves data quality coverage by 26%, allowing for checks on more dimensions like completeness, consistency, referential integrity, and distribution drift over structured data like customer profiles, transaction logs, and IoT telemetry data. An astonishing 73% reduction in missed mistakes proves the system's capability for detecting faint anomalies, like creeping field degradation or soft schema offenses, that other rule-based checks normally miss. Leveraging AI-driven continuous learning pipelines that are trained on labeled metadata from more than 10,000 real-world data tables in finance, energy, and retail domains, the model provides a stunning 91.6% update cycle time reduction through dynamic quality rule adaptation against schema changes and data drift. By automated rule suggestion and propagation, AI obliterates 85% of tedious data engineering drudgery, allowing teams to concentrate on high-value, strategic work like model tuning and pipeline architecture. In addition, detection latency has plummeted more than 93%, as AI tirelessly keeps track of feature distributions and rapidly detects deviations from learned baselines—especially important in real-time streaming data scenarios like fraud detection or logistics telemetry. These findings substantiate that with the proper model architecture and representative data domains, AI can emerge as an intelligent co-pilot in enterprise-level data quality management.

7. Future Research Directions

There is a world of possibility in the future of AI-powered DQA in multi-cloud data warehousing. AI will be critical in enabling automation and optimization of data quality practices as data ecosystems expand in size and cross platforms such as AWS, Azure, and GCP [22]. Future developments will take data governance automation to the next level — AI-driven governance — where systems can independently author, deploy, and update rules with little human supervision [23]. The agglutination of generative AI in businesses will provide businesses

with a more nuanced and individualized understanding of users' behavior, revamping the approaches to engagement [14].

Conclusion

This experimental evaluation of the ACA-AI-DQA framework provided several critical insights into the framework's performance quality, adaptability, and effectiveness in multi-cloud data warehousing. The review ran the framework on three major cloud platforms: Amazon Redshift, Google BigQuery, and Azure Synapse Analytics. Both synthetic datasets with carefully controlled anomalies and real-world datasets from the retail and finance sectors were evaluated to optimize system efficacy. Anomaly detection achieved a superior F1 of 0.92 with a false positive rate of 6%. At the same time, the framework's alternatives, Random Forest and Autoencoders, had a false positive rate of 30% on the same datasets. An F1 of 0.92 boasted of incorporating context-aware features in the detection process. The system generated new rules in less than 18 seconds compared to the manual, which took more than 3 hours. This provided a source book for rule validation when the schema changed and ensured that the downstream did not fail. The self-healing module could auto-correct 30–50% of common data issues like null, format mismatch, and invalid entries, eliminating human interventions and ensuring Cleanliness. Federated learning ensured the learning stayed within 3% of a centralized training while the data remained at the data-originating cloud. This setup ensures compliance with data standard policy and decentralizes model training across clouds. The framework's microservices-based design simulated a peculiar environment to scale across cloud varieties.

References

- [1]. Z. Yang, Q. Shi, T. Cheng, X. Wang, R. Zhang, and L. Yu, "A security-enhanced authentication scheme for quantum-key-distribution (QKD) enabled Internet of vehicles in multi-cloud environment," *Veh. Commun.*, vol. 48, no. 100789, p. 100789, Aug. 2024.
- [2]. C. Jin, Y. Xu, W. Qin, J. Zhao, G. Kan, and F. Zeng, "A blockchain-based auditable

- deduplication scheme for multi-cloud storage,” *Peer Peer Netw. Appl.*, vol. 17, no. 5, pp. 2870–2883, Sep. 2024.
- [3]. K. Sathupadi, “Deep Learning for Cloud Cluster Management: Classifying and Optimizing Cloud Clusters to Improve Data Center Scalability and Efficiency,” *Journal of Big-Data Analytics and Cloud Computing*, vol. 6, no. 2, pp. 33–49, 2021.
- [4]. [4] Kenya School of Government, Embu Campus and M. Omuya Odida, “Exploring the application of Artificial Neural Networks in enhancing security measures for cloud computing: A survey,” *J Mari Scie Res Ocean*, vol. 7, no. 2, pp. 01–12, May 2024.
- [5]. [5] D. Kaul, “Optimizing Resource Allocation in Multi-Cloud Environments with Artificial Intelligence: Balancing Cost, Performance, and Security,” *Journal of Big-Data Analytics and Cloud Computing*, vol. 4, no. 5, pp. 26–50, 2019.
- [6]. B. R. Piduru and Customer Experience Architect, Irvine, CA, USA, “Cloud computing and public sector transformation: Revolutionizing governmental services and operations,” *J Arti Inte & Cloud Comp*, pp. 1–4, Sep. 2022.
- [7]. Y. Jani, “Strategies for Seamless Data Migration in Large-Scale Enterprise Systems,” *Journal of Scientific and Engineering Research*, vol. 6, no. 12, pp. 285–290, 2019.
- [8]. Y.-G. Guo, Q. Yin, Y. Wang, J. Xu, and L. Zhu, “Efficiency and optimization of government service resource allocation in a cloud computing environment,” *J. Cloud Comput. Adv. Syst. Appl.*, vol. 12, no. 1, p. 18, Feb. 2023.
- [9]. S. V. Bhaskaran, “Integrating Data Quality Services (DQS) in Big Data Ecosystems: Challenges, Best Practices, and Opportunities for Decision-Making,” *Journal of Applied Big Data Analytics, Decision- Making, and Predictive Modelling Systems*, vol. 4, no. 11, pp. 1–12, 2020.
- [10]. H. Adhab, E. M. Kalik, and A. K. A. Ani, “Designing a smart e-government application using a proposed hybrid architecture model dependent on edge and cloud computing,” *Electron. Gov. Int. J.*, vol. 18, no. 3, p. 340, 2022.
- [11]. S. V. Bhaskaran, “Optimizing Metadata Management, Discovery, and Governance Across Organizational Data Resources Using Artificial Intelligence,” *Eigenpub Review of Science and Technology*, vol. 6, no. 1, pp. 166–185, 2022.
- [12]. S. V. Bhaskaran, “Tracing Coarse-Grained and Fine-Grained Data Lineage in Data Lakes: Automated Capture, Modeling, Storage, and Visualization,” *International Journal of Applied Machine Learning and Computational Intelligence*, vol. 11, no. 12, pp. 56–77, 2021.
- [13]. Manjunath V., D. Kalaskar, and Government First College, Gurumatkal, Yadgir, Karnataka, “Cloud assisted IoT application’s security attacks and their countermeasures,” *Int. J. Eng. Res. Technol. (Ahmedabad)*, vol. V9, no. 05, May 2020.
- [14]. L. F. M. Navarro, “The Role of User Engagement Metrics in Developing Effective Cross-Platform Social Media Content Strategies to Drive Brand Loyalty,” *Contemporary Issues in Behavioral and Social Sciences*, vol. 3, no. 1, pp. 1–13, 2019.
- [15]. [15] K. Kushagra and S. Dhingra, “An empirical analysis of the government cloud adoption in India,” *Int. J. Electron. Gov. Res.*, vol. 17, no. 3, pp. 21–43, Jul. 2021.
- [16]. S. Rahman, M. R. M. Sirazy, R. Das, and R. S. Khan, “An Exploration of Artificial Intelligence Techniques for Optimizing Tax Compliance, Fraud Detection, and Revenue Collection in Modern Tax Administrations,” *International Journal of Business Intelligence and Big Data Analytics*, vol. 7, no. 3, pp. 56–80, 2024.
- [17]. Abraham, F. Hörandner, T. Zefferer, and B. Zwattendorfer, “E-government in the public

cloud: requirements and opportunities,”
Electron. Gov. Int. J., vol. 16, no. 3, p. 260,
2020.

- [18]. R. S. Khan, M. R. M. Sirazy, R. Das, and S. Rahman, “Data-Driven Perspectives on Federal Budgetary Dynamics for Identifying Anomalies and Patterns in Resource Allocation and Obligation Trends,” *Quarterly Journal of Emerging Technologies and Innovations*, vol. 9, no. 3, pp. 50–70, 2024.
- [19]. Hashem, Ibrahim Abaker Targio, et al. “The Rise of ‘Big Data’ on Cloud Computing: Review and Open Research Issues.” *Information Systems*, vol. 47, 2015, pp. 98–115.
- [20]. Grolinger, Katarina, et al. “Data Management in Cloud Environments: NoSQL and NewSQL Data Stores.” *Journal of Cloud Computing: Advances, Systems and Applications*, vol. 5, no. 1, 2016.
- [21]. Zhang, Lei, et al. “Anomaly Detection in Cloud-Based Data Using Machine Learning Algorithms.” *Journal of Cloud Computing*, vol. 10, no. 1, 2021.
- [22]. Gao, Yong, et al. “Using NLP to Enhance Metadata Interpretation in AI-Driven Data Quality Assurance.” *Future Generation Computer Systems*, vol. 125, 2022, pp. 219–234.
- [23]. Kumar, Ravi, et al. “AI-Enabled Data Governance for Multi-Cloud Environments.” *Journal of Cloud Computing and Big Data*, vol. 13, no. 2, 2023.