

Universal Data Engineering Frameworks for Cross-Platform Fraud Detection

Ravi Kiran Alluri

USA

ABSTRACT

Developing, deploying, and operating on various platforms, including mobile banking apps and enterprise payment gateways, fraud detection systems must rise to meet the escalating threats presented by advanced fraud vectors. As a result, building a unified fraud detection system is challenging and necessary due to the heterogeneity of data sources, infrastructure ecosystems, and compliance standards. To address these challenges, this paper presents a generalized data engineering framework that enables scalable, platform-neutral, and real-time dynamic fraud detection through modular design, metadata-driven processing, and cross-platform interoperability. The framework brings together the structure, semi-structured, and unstructured data across a variety of systems through a common ingestion layer, standardized transformation pipelines, and abstracted machine learning integration points.

In this paper, we assess modern technologies and patterns used in data engineering – e.g., schema-on-read ingestion, federated data access, event-driven processing, and cloud-native orchestration – to establish a common base that enables fraud analytics across diverse stacks. It promotes the use of DataOps practices toward automated pipeline deployment, lineage tracing, and incident tracing. Furthermore, this framework not only combines real-time streaming processing (using Apache Kafka, Spark Streaming, or Flink) with historical batch processing (using Delta Lake or Apache Iceberg) for real-time fraud response, but also long-term fraud patterns. It adds interoperability through data abstraction layers, RESTful APIs, and GraphQL endpoints, making the system's backend agnostic and allowing for custom interfaces to the analytical layers.

Furthermore, the proposed framework utilizes adaptive learning methods to customize fraud detection rules and adapt to localized platform activities. This flexibility enables the architecture to be adopted in various domains, including fintech, insurance, healthcare, and e-commerce, where the sharing of knowledge through federated learning is made possible. This paper also discusses cross-jurisdictional data privacy, compliance (GDPR, HIPAA, PCI DSS), and regulatory reporting concerns by incorporating metadata-based policy enforcement and secure data masking techniques into the pipeline.

A prototype for the framework was deployed in three domains—banking, ridesharing, and retail—to provide proof of concept. The findings indicated gains in accuracy, the elimination of latency, and cost reduction resulting from the reuse of a pipeline and orchestration efficiency. Metrics such as fraud recall, false favourable rates, and execution throughput were calculated and compared to those of their pre-existing, siloed systems. It also reflects on some potential bottlenecks, including schema drift management, microservices dependency management, and integration overheads with vendor platforms.

This paper provides a comprehensive view on establishing universal data engineering frameworks applicable for fraud detection tasks, offering insights into modular pipeline design, cloud deployment architectures, and domain-specific customizations. Through the lens of data interoperability, pipeline observability, and real-time reaction responses, the research provides an architecture roadmap for constructing enterprise-grade fraud detection environments that can adapt to modern digital infrastructures.

*Corresponding author

Ravi Kiran Alluri, USA.

Received: June 06, 2025; **Accepted:** June 10, 2025; **Published:** June 22, 2025

Keywords: Cross-Platform Fraud Detection, Data Engineering Frameworks, Metadata-Driven Pipelines, Stream Processing, Federated Learning, Real-Time Analytics, DataOps, Schema Evolution, Platform Interoperability, Compliance Automation, Microservices, Event-Driven Architecture, Data Mesh, Edge Data Processing.

Introduction

The advent of digital ecosystems and the adoption of distributed platforms—from mobile finance applications and IoT payment terminals to cloud-based insurance and healthcare platforms—have amplified the challenge of detection-based fraud prevention to unprecedented levels. Fraud is no longer tied to typical behaviours

or limited datasets. It materializes in fragmented data trails, which savvy threat actors increasingly generate using AI-driven evasion, proxy identities, or cross-platform execution. In this context, the demand for a general, scalable, and generic data engineering library for fraud detection has increased.

Traditionally, these fraud detection systems were developed as monolithic platform-based applications, relying on batch ETL-based processes and siloed rule engines. These classical systems are now considered obsolete in dynamic, heterogeneous environments that require high-throughput, low-latency, and context-aware detection strategies. In addition, the growing number of data formats (JSON, Parquet, Avro, etc), data sources (APIs, event

streams, user sessions), and compliance regulations (e.g., GDPR, PCI DSS, HIPAA) forces a re-examination of how we design data to fuel fraud analytics. The traditional “one-size-fits-all” architecture, often associated with classic approaches, cannot effectively accommodate the distribution of fraud signals across mobile, web, cloud, and hybrid infrastructures.

To address these challenges, we present a generalizable data engineering framework that separates data ingestion, transformation, storage, and analytics through a modular pipeline design. The ultimate aim is to create a cross-platform fraud detection system based on easily pluggable interfaces (Kafka topics, REST APIs, metadata catalogues). The interfaces will help connect diverse data sources, preserving traceability, data reliability, and observability. Our proven foundation in modern data engineering practices, including schema-on-read processing, streaming-first architecture, metadata-driven orchestration, and cloud-native deployment, enables this.

This follows the guiding premise of this work, which is that the pursuit of universality in data engineering does not lie in a monolithic platform, but in defining data pipeline functionality as reusable and exchangeable components. These modules can be called into use and orchestrated to address specific fraud use cases—such as card-not-present fraud, synthetic identity fraud, transaction laundering fraud, and social engineering fraud—without the need to rearchitect the entire system by domain or platform. Additionally, by combining federated learning and edge analytics, organizations can collaboratively train models without migrating sensitive data across regulatory lines, ensuring real-time fraud detection even in compliance-heavy industries.

Another aspect that is investigated is the extent to which such a framework can enable hybrid execution, which involves combining stream processing and batch processing into a single pipeline, allowing for both real-time alertness and historical trend analysis. The framework is based on Apache Kafka, Apache Flink, Spark Structured Streaming, and Delta Lake, providing the capability of scalable execution pipelines with backpressure handling, event replay, schema evolution, and consistent state management. We also acknowledge the difficulties in building general systems – schema drift, data format heterogeneity, lineage tracking, and semantic drift from data producers to consumers are nontrivial technical challenges. This paper acknowledges the difficulty of building general systems (eg, schema drift, data format heterogeneity, lineage tracking, and semantic drift). The framework addresses these problems by introducing a metadata management layer that centralizes control over data contracts, transformation logic, and versioned schema mappings.

The rest of this paper is structured as follows: literature review reviews recent state-of-the-art in fraud detection architectures and cross-platform data engineering; methodology details the technical frameworks and operational tactics of the novel approach; results shows a multidomain use case demonstrating deploying the approach; discussion critically discusses performance trade-offs and limitations; and conclusion states future utilities of the universal data engineering in a fraud detection management ecosystem.

Literature Review

The field of fraud detection has evolved significantly over the past decade, moving from rigid rule-based systems to flexible, data-driven, and machine learning-powered architectures. However, a

critical gap remains in enabling fraud detection systems to operate seamlessly across disparate platforms. Literature addressing this problem highlights the importance of decoupling data architecture from platform-specific dependencies to ensure interoperability, scalability, and responsiveness in fraud analytics.

A foundational study argues that real-time fraud detection requires a stream-first approach, where event data is ingested and processed with minimal latency using technologies such as Apache Kafka and Apache Flink. This approach has been extended in, where the authors introduced a hybrid pipeline model that combines stream and batch workloads to handle both transactional fraud detection and historical behavior analysis. These models, while effective in isolated deployments, lack a universal structure that can span multiple domains or integrate with federated systems across enterprises [1,2].

Data engineering frameworks supporting fraud detection have also undergone a transition from monolithic ETL processes to modular orchestration tools such as Apache Airflow, Dagster, and Prefect. As discussed in, orchestration tools are now essential in supporting metadata-driven workflows, automated retries, DAG-based job execution, and lineage tracking—critical features for fraud investigations that require traceable audit trails and contextual analysis. However, these tools, in their native form, do not inherently support cross-platform deployments unless coupled with abstraction layers [3].

In, a metadata-centered architecture for data pipeline construction was introduced, focusing on the role of data contracts and centralized metadata repositories. This model facilitates schema enforcement, versioning, and runtime validation across services. In fraud detection scenarios, this improves consistency when ingesting variable data formats from mobile apps, e-commerce platforms, or POS systems. Similarly, the work in emphasizes the role of schema registries in managing evolving data formats, an essential requirement for systems aiming to achieve universality [4,5].

Several studies explore the integration of federated learning for fraud detection in distributed environments [6,7]. Federated models enable multiple institutions—such as banks, insurers, or tax authorities—to collaborate on fraud pattern recognition without sharing raw data. The potential of this approach was demonstrated in, where a federated fraud detection prototype reduced false positives by 17% and met cross-border privacy regulations. However, the absence of a standardized data engineering infrastructure hinders large-scale adoption [8].

Event-driven architectures (EDAs), as discussed in, have proven particularly relevant to fraud analytics due to their real-time responsiveness and support for asynchronous event propagation. EDAs using tools like Kafka, Pulsar, and Event Bridge enable systems to decouple producers and consumers, making it easier to introduce fraud detection logic as independent services. Coupled with microservices-based deployments, these architectures allow for easy scaling, observability, and versioned experimentation of fraud detection algorithms [9].

From a compliance and regulatory standpoint, and highlight the growing need for data engineering systems to incorporate privacy-preserving features, such as anonymization, encryption at rest, and fine-grained access control. These capabilities are increasingly vital in fraud detection pipelines that operate across multiple

geographies, where differing data privacy laws can significantly impact the system design [10,11].

The emerging concept of a "data mesh"—introduced in—also aligns with the goals of universal data frameworks. A data mesh emphasizes decentralized data ownership, domain-driven pipeline design, and interoperability through well-defined APIs. In fraud detection contexts, this enables faster response to domain-specific threats while allowing federated governance [12].

While these studies lay a strong foundation, none provide an integrated, end-to-end data engineering framework explicitly tailored for universal, cross-platform fraud detection. This paper aims to bridge that gap by synthesizing stream-first ingestion, modular metadata-driven pipelines, event-based microservices, and federated learning—all within a compliant and scalable framework suitable for deployment in 2025's diverse fraud ecosystems.

Methodology

The development of a universal data engineering framework for cross-platform fraud detection involves a multi-layered, modular design approach that abstracts platform dependencies, enabling scalable deployment across heterogeneous systems. This methodology prioritizes three core principles: interoperability, real-time responsiveness, and compliance-aware design. The framework is structured around a layered architecture that separates concerns across ingestion, transformation, analytics, governance, and interface layers. These layers work in unison to collect, normalize, analyze, and respond to fraud-related signals across various environments, including mobile apps, banking systems, ride-hailing platforms, and e-commerce ecosystems.

At the foundation lies the data ingestion layer, which supports the acquisition of both streaming and batch data from a diverse range of systems. This includes Kafka topics for real-time events, RESTful and GraphQL APIs for application-level interaction, change data capture (CDC) systems from transactional databases, and file-based ingestion from data lakes or object storage. Schema-on-read techniques are employed to dynamically interpret the data structure at query time, enabling support for evolving or semi-structured data formats, such as JSON, Avro, and Parquet, without requiring rigid predefinition. The ingestion layer is decoupled from upstream producers using an event-driven model that ensures message durability and scalability.

Data passes into a transformation and enrichment layer that performs essential functions such as data cleaning, normalization, entity resolution, and feature engineering. This layer leverages Apache Spark for distributed batch processing and Apache Flink for real-time enrichment and outlier detection. A metadata management component governs all transformations by maintaining reusable transformation templates, schema registries, and lineage metadata. This approach ensures that transformations remain reproducible and auditable, which is critical for compliance and regulatory reporting. Advanced transformations include the application of hashing and tokenization techniques for privacy preservation, dimensional flattening for analytical workloads, and window-based aggregations for behavioural analysis.

Once transformed, the data is routed to a hybrid storage layer that balances the need for low-latency lookups and long-term analytical storage. Real-time state is maintained using NoSQL stores, such as Apache Cassandra or Redis. At the same time,

historical data is warehoused in columnar formats within systems like Apache Iceberg, Delta Lake, or BigQuery. The system supports incremental updates and time travel queries to trace fraud evolution and maintain historical model performance logs. Data partitioning, clustering, and compaction strategies are applied based on domain-specific characteristics such as customer ID, transaction date, or fraud score thresholds.

The analytical layer interfaces with the stored data using both rule-based engines and machine learning models. Rules are defined in a domain-specific language (DSL) that supports conditional logic, temporal windows, and pattern recognition. Models are deployed using MLFlow and are trained on features generated through a feature store layer that abstracts the engineering logic from the model code. The feature store supports offline training and online inference synchronization, ensuring that model behavior remains consistent across environments. Federated learning components are integrated to allow cross-institutional model training while preserving data sovereignty.

The orchestration and governance layer ensures observability, version control, and access management across the entire framework. All pipeline steps are orchestrated using tools like Apache Airflow or Prefect, which offer DAG visualization, failure recovery, and task retries. Role-based access control (RBAC) policies are enforced at every stage using cloud-native identity and access management (IAM) solutions. Metadata is exposed through APIs for downstream services such as dashboarding, alerting systems, and compliance reporting engines. Observability is enhanced through the use of distributed tracing and metrics collection via Prometheus and Open Telemetry.

Finally, the framework interfaces with external applications via APIs, real-time alerts, dashboards, and audit trails. Fraud alerts can be emitted to CRM systems, payment gateways, or user notification services through Kafka topics or WebSocket channels. Dashboards built using platforms like Superset or Power BI present visual summaries of fraud trends, anomaly clusters, and pipeline health. Compliance audit trails, including data provenance and transformation logs, are maintained in immutable storage and made available to authorized regulatory agents upon demand.

The overall framework is designed to be cloud-native, containerized, and deployable on Kubernetes clusters. Helm charts and CI/CD pipelines enable automated deployment, scaling, and testing. This methodology ensures that the framework can be replicated across departments, institutions, or even geographic regions with minimal code modification, enabling truly universal fraud detection capabilities.

Results

To evaluate the effectiveness of the proposed universal data engineering framework for cross-platform fraud detection, extensive tests were conducted in three real-world environments (i.e., digital banking, ride-hailing platforms, and retail e-commerce) using controlled deployments. Each of them had its data formats, fraud vectors, latency requirements, and regulatory constraints. The goal was to evaluate the framework's capacity to detect malicious activities accurately, minimize response time, support schema evolution, and comply with local data privacy regulations.

In the digital banking example, the system was embedded in a real-time transaction processing engine, which consumed and processed customer transaction streams from mobile and web

applications. This configuration was designed to identify card-not-present fraud, account takeover attacks, and anomalous user purchasing behavior. The fraud detection logic consisted of supervised machine learning models (trained on historical features such as transaction speed and geo-location inconsistency) and real-time rule triggers (e.g., blocked IP addresses or outlier device IDs). The deployment reached 94% precision and 88% recall, with a fraud alert latency of less than 1.8 seconds. Compared to the legacy system, which runs on a batch and delays alerts by 20-30 minutes, the process has dramatically decreased customer impact and operational response time.

In the ride-hailing system, the system was implemented to identify fake driver accounts, GPS spoofing, and promotion abuse fraud. The data was in a semi-structured form and was received from mobile device telemetry, booking APIs, and payment gateways. Augmented by Flink-driven stream processing and dynamic rule chaining, the newly proposed framework detected abnormal rides with better sensitivity compared to previous attempts. Detection recall increased by 22% compared to the legacy microservice pipeline, especially for synthetic identities and geospatial inconsistencies. The modular design of the pipeline enabled us to iterate and deploy new features to detect fraud quickly, without requiring a system shutdown—a crucial requirement when operating at the scale of ride-sharing.

The e-commerce use case included the ingestion of order, payment, customer profile, and delivery data to detect frauds, such as triangulation scams, bulk purchase automation, and return abuse. Fraud behaviours were dynamic and distinct across geographies and SKUs. A novelty was a federated model that learned fraud patterns across multiple vendors while minimizing vendor-specific data privacy. Models trained in MLFlow and deployed through the framework were refreshed nightly through retraining jobs scheduled with Airflow DAGs, and the feature store was updated in sync with real-time features (e.g., cart abandonment sequences, payment token reuse, and quick switching of addresses). The introduction of the deployment resulted in an improved fraud catch rate of 91% with a reduced false positive rate of 5.2%, compared to 11.8% from an existing rules-only system.

Operation metrics in all three aspects showed that the proposed pipeline consistently made the data pipeline reliable, maintained high data freshness, and enabled the model to become observable. The integration of new data sources was straightforward due to the schema-on-read implementation and metadata abstraction. With reusable DAGs, pipeline deployment time was reduced by approximately 37%, and schema drift was mitigated through automated schema registry updates. Furthermore, governance reviews (both internal and third-party) confirmed that the compliance features – including data masking, audit logging, and access control – met the requirements, demonstrating that the framework was regulatory-ready.

The framework was stress-tested for throughput and latency, processing more than 25,000 events per second in the banking domain with an end-to-end latency of less than 2.5 seconds. CPU and memory usage were optimal thanks to Kubernetes auto-scaling container services. The addition of observability tools (such as Grafana and Prometheus) delivered live metrics and provided operators with real-time alerts of data flow deviations and processing hiccups. From the experimental findings, it is evident that the architecture improves the efficiency of fraud detection and also fosters operational flexibility, system scalability, and ecosystem trust.

Discussion

We found it to be very compelling that the results emerging from the use of the generic data engineering framework across digital banking, ride-hailing and e-commerce domains support the idea that a cross platform, modular, data pipeline architecture can indeed be used to enhance our ability to detect financial fraud without compromising the business, regulatory or operational context. The obtained performance metrics – high precision and recall, low latency, and throughput scalability – validate the central hypothesis that a universal and reusable framework, whose purpose is defined by metadata and is powered by event streams, is not only possible but also beneficial in various fraud ecosystems.

Perhaps the most important result is that the fraud alert latencies of all the domains decrease. Traditional ETL-based fraud detection systems are inherently batch-based and have a time lag in response, which does not effectively cover fast-paced fraud patterns. In contrast, the tested framework was able to issue real-time alerts stably with an average latency of under 2.5 seconds. This enhancement is particularly crucial for applications such as banking and ride-sharing, where spotting fraud early can save money and reduce reputational damage. These latency improvements are due to the stream-first ingestion architecture of the framework and the distributed processing engines such as Apache Flink and Spark Streaming, on which it is based.

Moreover, the technique of schema-on-read and template-based “transformation” worked like a charm to address the schema evolution problem, specifically in the context of ride-hailing and e-commerce at a massive scale, where data looks very different all the time. Dynamic schema registration and metadata-driven transformation support in the framework enabled integration of arbitrary sources of new data with minimal engineering effort. This competence aligns with the observations of Choi et al [4]. and Ruiz et al., which emphasize the growing importance of flexible schema evolution in online contexts. By encoding transformation logic in metadata configurations rather than in code, the framework both maintains and reduces technical debt, a problem that has plagued legacy fraud detection systems [5].

The applications of federated learning models also demonstrate that the framework can be efficient in settings where privacy compliance and cross-jurisdictional regulations are a concern. In the e-commerce scenario, data was contributed by various vendors across different jurisdictions, enabling federated learning to support collective intelligence without revealing raw data. This strategy helped decrease the number of false positive calls and was also consistent with the proposal by Kwon et al, regarding the scalability of federated architectures in financial ecosystems [6].

Although these data are promising, some limitations were experienced in the actual use of the process. One of the most significant issues was navigating complex inter-service dependencies at the orchestration level. While there are orchestration tools like Apache Airflow, they offer very high-level orchestration features. Managing such a large number of microservices and keeping them operational across multiple clusters adds significant complexity. These issues were especially pronounced when different feature dependencies, versioned ML models, were being rolled out - marshalling this feature set and scheduling/run-time validation was complicated. 4.3 BURST INPUT MODEL The S network has the problem of behaviour under a bursty scenario. Although auto-scaling via Kubernetes also significantly increased the system's robustness, under extreme

pressure, there were some short-lived performance degradations at specific stream processing loads, especially when consuming high-velocity topics from Kafka, where partitions had fluctuated in their distribution.

Furthermore, the framework's reliance on modularity and abstraction comes at a cost. Modularity has the disadvantage that it is beneficial for reusability in most aspects, except for performance. For example, metadata rule-driven transformation logic may not achieve the raw execution speed of handcrafted, tightly connected transformation scripts directed by human developers. However, the flexibility gained by such a trade-off typically far exceeds the marginal costs of performance for most enterprise environments, where ease of use and auditability are equally important.

Another important observation is that observability tooling in the framework played a critical role in identifying abnormalities in the data pipeline ahead of time. The monitors gathered through Prometheus and visualized with Grafana also had an impact on fraud analytics, pinpointing pipeline inefficiencies, back-pressure events, and memory peaks. This observation supports the assertions of Sharma and Kalita that observability is an essential foundation in any real-time analytics framework [9].

Compliance-wise, being able to obtain on-demand lineage reports, transformation logs, and user access audits helped lower the regulatory overhead. This is consistent with Bartolini et al., who cited that compliance automation is necessary to future-proof fraud detection infrastructure. Insist on Metadata and policy for All Data Activities. Through metadata integration and policy enforcement everywhere data is accessed—both as it is ingested and throughout its lifecycle in flight or at rest—the framework itself is auditable and transparent by default [10].

The paper confirms that the framework is both cross-domain applicable and scalable, and that it is ready for compliance from the outset. Technical and architectural issues, such as orchestration complexity and optimizing bursty workloads, among others, persist; however, the conceptual solution represents a significant step forward in providing an overarching approach to fraud detection across digital platforms. The conclusions of this assessment serve as a foundation for the development of additional capabilities and as a baseline for companies willing to renew their fraud analytics platform by leveraging a future-proof, compatible, and compliant architecture.

Conclusion

The rapid digitalization of banking, e-commerce, mobility, and healthcare ecosystems has increased the attack surface for fraud, while also revealing the limitations of legacy, siloed systems for fraud detection. This paper fills the gap with a proposal for a compatible, scalable, and compliance-oriented universal data engineering framework for cross-platform fraud detection. The research reported in this article demonstrates that a metadata-driven, modular, cloud-native architectural approach is capable of properly reconciling data ingestion, transformation, modeling, and governance in heterogeneous environments.

The industry-specific readiness and performance of the framework in three industry domains (digital banking, ride-hailing, and e-commerce) revealed substantial gains in the accuracy of

fraud detection, the platform's real-time ability to respond, operational efficiency, and compliance with regulations. In particular, the framework consistently provided high precision and recall, substantially lowered fraud alert latency, and kept high observability through monitoring and logging integration. These results provide evidence to support the claim that fraud detection pipelines can be generalized and shared across contexts without compromising the performance and clarity of the models.

At the heart of the framework is its layered nature, which helps to abstract the functional aspects and provides organizations with the possibility to independently scale their ingestion, transformation, analytics, and orchestration layers. The use of schema-on-read, dynamic metadata catalogues, federated model training, and a hybrid storage strategy allows for high flexibility in ingesting structured and semi-structured data sources. These architectural decisions enable seamless integration with existing enterprise systems, allowing for the real-time or near real-time identification of changing fraud patterns while maintaining privacy and governance invariants.

In the current era of regulation, the inclusion of federated learning and privacy-preserving data transformations is particularly essential. Moreover, these mechanisms enable the safe sharing of fraud intelligence among institutions and comply with international compliance, including GDPR, HIPAA, and PCI DSS. The ability to centralize analytic intelligence while localizing data governance provides a basis for collaborative fraud prevention between jurisdictions.

Despite these strengths, challenges remain. Challenges surrounding the orchestration of pipelines, the coordination of microservices, and dynamic resource scaling must be continually addressed to ensure the system remains resilient in the face of workload fluctuations. Although the abstraction of metadata enables the reuse of and auditable pipelines, it may introduce performance overhead in some high-throughput cases. Therefore, the next generation of improvements should focus on smart workload equilibrium, predictive pipeline tuning, and automated metadata optimization.

This paper also creates opportunities for further investigation into fraud analytics. In the future, graph-based fraud detection techniques can be used to identify relational anomalies among accounts, devices, and transaction paths. The promise of edge-based inference, in which models are deployed at the edge (on the device or at the branch), can also be leveraged to detect fraud sooner and reduce data transport costs. Secondly, the adoption of zero-trust data architectures and confidential computing can enhance the security level of fraud detection pipelines in sensitive domains, such as national tax systems and health insurance networks.

The portable data engineering framework presented in this study provides a generic master plan for building fraud detection systems that can combat newly emerging fraud threats. By applying the tenets of modularity, observability, and compliance by design, organizations can rapidly, accurately, and robustly address their exposure to fraud. As fraud vectors adapt to new tactics in digital ecosystems, such broad-reaching frameworks will be essential for protecting transactions, identity, and trust between platforms.

References:

1. Ahmed T, Raza M (2024) Real-Time Stream Analytics for Financial Fraud Detection, IEEE Trans. Ind. Informat 20: 1121-1130.
2. Iyengar S, Hu L (2023) A Dual-Mode Fraud Detection Pipeline: Combining Real-Time Alerts with Historical Analysis, IEEE Access 11: 94321-94332.
3. Martin J, Zhang H (2023) Metadata-Driven ETL Orchestration in Modern Data Warehousing, Proc. IEEE BigData 1223-1232.
4. Choi Y (2023) Design Patterns for Metadata-Centric Data Pipelines, ACM Trans. Data Syst 42: 211-234.
5. Ruiz L, Nogueira F (2025) Schema Management for Evolving Data Streams, IEEE Trans. Knowl. Data Eng 35: 22-34.
6. Kwon R (2025) Federated Fraud Detection Across Distributed Enterprises, IEEE Internet Comput 29: 44-52.
7. Sen P, Ramesh A (2024) Collaborative Machine Learning for Real-Time Fraud Detection, IEEE Conf. Cloud Comput 233-240.
8. Joshi D, Patel V (2023) Evaluating Federated Learning in Compliance-Constrained Financial Environments, Proc. IEEE ICSC 455-463.
9. Sharma N, Kalita M (2023) Event-Driven Microservices for Financial Services, IEEE Software 40: 87-94.
10. Bartolini G (2024) Privacy-Aware Architectures in Cross-Border Fraud Detection, IEEE Trans. Secur. Priv 19: 14-23.
11. Yoon E (2023) Regulatory-Aware Data Engineering Pipelines in the Age of GDPR, IEEE Cloud Comput 10: 41-51.
12. Shafiq Z (2024) Domain-Driven Data Mesh for Large-Scale Analytics, Proc. IEEE ICDE 1224-1235.

Copyright: ©2025 Ravi Kiran Alluri. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.