



Scalable Cloud Models for Banking Transaction Monitoring

Niklas Andersson

Asst Professor

Abstract

Electronic banking offers new opportunities for credit unions and other community financial institutions. By allowing members to bank 24 hours a day from virtually anywhere in the world, the benefits of electronic banking must be weighed against the potential for increased fraud and illicit transactions. Regulatory requirements mandate that as a result of the Money Laundering Control Act of 1986, banks need to have a formal anti-money laundering program that includes procedures to monitor electronic banking activity. Cloud computing technologies provide scalable computing resources at a reduced operational cost, but their use for anti-money laundering systems has not yet been demonstrated. The characteristics of digital banking transactions indicate these systems may benefit from cloud-based resources. Appropriate cloud architectures and data ingestion and transformation methods are proposed to support real-time rule-based compliance system detection and machine-learning-based anomaly detection. Cloud processing scalability creates processing capacity for compliance requirements without negatively impacting the underlying economic advantage.

Cloud computing has emerged as a disruptive technology enabling scalable and affordable data processing in a rapidly evolving environment, and organizations have seized the opportunity to move data storage and processing outside the enterprise firewall. Scalable data processing can be a significant advantage for organizations with irregularly occurring, high-volume access patterns. For compliance monitoring and operational reporting, however, regulatory requirements necessitate monitoring and reporting within a defined period. Such rigid time constraints reduce the potential operational cost benefit of cloud-scale processing resources. Digital transaction characteristics suggest compliance monitoring and anomaly detection can be separated and deployed in a hybrid processing model. The rule-based detection models of the operational system provide anomaly detection ground truth and a source of training data for machine-learning models. Cloud processing enables a scalable architecture to meet current and future capacity demands.

Keywords : Digital banking transaction monitoring; a cloud data management and processing system model; a bank transaction monitoring system that adheres to regulatory requirements; cloud processing of banking data; banking data digital transformation; real-time transaction monitoring of digital banking transactions; near real-time monitoring of banking transactions; rule-based monitoring of banking transactions; machine-learning-based monitoring of banking transactions; full-text data storage for banking transactions.

1. Introduction

Digital banks need scalable technology to support their transaction monitoring requirements, while monitoring systems for digital banking lack data characteristics necessary for compliance. Application and infrastructural profiles guide the design of cloud data processing models that meet technology, cost, and compliance constraints. Knowledge tells architectural patterns for scalable cloud processing, along with cloud computing service and deployment models and security and data locality principles. Examining close-to-cloud data ingestion techniques reveals connectors for data federation, quality-enforcing schema evolution,



and tooling for real-time data feeding. Patterns for connecting machine learning models into the monitoring system, coupled with the rulesets governing deployed models, guide compliance-dedicated model management.

Digital banking is growing rapidly thanks to industry-specific technological and regulatory developments. Major cloud solution providers support a combination of technological fields that are becoming a reference for scalable solutions. Still, transaction monitoring systems that accommodate these technological advancements and scale with business operating volumes are becoming a limiting factor for the digital banking business. Monitoring systems for digital banks require very specific processing capabilities regarding ingested data volume, transformation characteristics in relation to storage volume, and a volume of stored data that can vary quickly. Classic vertical scaling cannot solve these problems, so cloud solutions, especially based on near-real-time data feeding techniques, are being considered. However, the data requiring near-real-time monitoring differ from the usual characteristics: they are not serial, nor do they follow a traditional periodic pattern. This section identifies and analyses the architectural, design, and implementation aspects of the limitations present in existing digital banking transaction monitoring systems.




Fig 1: Modelling Trust and Commitment in Digital Banking Platforms

1.1. Background and Significance

The banking and financial services industry has embraced digital banking, which offers quick responses when customers need bank services, making it a true 24/7 service. To support the strong demand for online banking features, banks have started to launch services targeted at Agile and DevOps development solutions with small-scaled release shipments. Digital transactions include fund withdrawal, transfer within and without banks, and payment to E-KYC service provider. Recently, business models that allow banks to store customer data with acceptable risk profiles have emerged. A growing number of banks offer monitoring systems to mitigate compliance risks related to anti-money laundering, a system description for digital banking monitoring, and enlarge compliant detection services.



Although transaction volumes are lower than those of major credit card providers, the compliance cost-to-monitoring ratio is relatively high due to the company's scale. A cloud subscription-based operation model for monitoring service is proposed, utilizing three digital banking transaction monitoring systems that align and integrate real-time rule-based monitoring and machine-learning-based anomaly detection system prediction for continuous transaction monitoring. Proposed operational models for data ingestion and processing are critical for such a low-cost structure and are capable of running on an external cloud service. For compliance needs, banks can be active at the rule definition phase with automatic connecting model training for an ML-positive validate operation.

1.2. Research design

A combination of theoretical research and applied design science was used to achieve an overview of scalable data processing solutions to support digital banking transaction monitoring. The results describe the general characteristics of the data, present the task landscape, and specify the requirements for scalable, accurate, and efficient solutions. Low- and near-real-time implementations capable of continuously and automatically adapting to data changes are discussed. The use of data federation, transformation, enrichment, and quality techniques is analyzed in relation to local-cloud and cloud-cloud patterns. Finally, monitoring suitable for achieving regulatory compliance is considered, including real-time digital and rule-based monitoring as well as struggling online-machine-learning detection approaches that adapt to gradual and recurrent data refreshment.

Data processing model overview aims for digital banking business monitoring based on a regular process through five layers (i.e. source, ingestion, core, sinks, and monitoring). A transaction-monitoring application has two basic characteristics. Each transaction is independent, and the volume of transactions grows (increases) to very large values. When a monitoring task is set up, real-time or near-real-time Digital Business Monitoring is desired. Transaction data can have many variations, and new types of transaction can even appear overtime. последнее.

Equation 1: Real-time stream capacity derivation

Step 1: If one worker can process μ transactions per second, then c workers can process:

$$\text{System capacity} = c\mu \text{ tx/s}$$

Step 2: For stability, the incoming rate must remain below capacity:

$$\lambda < c\mu$$

Step 3: Define utilization ρ as the fraction of capacity being consumed:

$$\rho = \lambda / (c\mu)$$

Step 4: Rearranging for the minimum worker count gives:

$$c \geq \lambda / (\rho_{\text{target}} \mu)$$

Step 5: With $\lambda = 200$ tx/s, $\mu = 25$ tx/s, and a target utilization of 70%:

$$c \geq 200 / (0.70 \times 25) = 200 / 17.5 = 11.43$$

Step 6: Because the worker count must be an integer, round upward:

$$c = 12 \text{ workers}$$

2. Background and Context



Digital banking has developed rapidly, as a result of customer demand and increased competition between banking institutions. In particular, digital banking transactions, which incorporate online bank authentication for e-commerce and online remittance, have grown significantly. Digital banking transaction monitoring refers to the continual monitoring of transactions and the detection of fraud in near real-time. Fraudulent transactions can be identified by comparing transaction characteristics or patterns with past historical data and are classified into two types: rule-based monitoring methods and machine learning monitoring methods. Machine learning methods are used less frequently because of difficulties in explaining the reasons for false alerts.

In addition to the difficulty of interdiction fraud, the financial services industry is heavily regulated. Financial service authorities require periodic file reporting for anti-money laundering compliance, and the data required for these files is often very large in volume and recorded without stable form. The data required for the file is processed periodically, unlike transaction monitoring, which requires nearly real-time processing. These fundamentally different characteristics result in different data processing models for approaching the two types of requirements. Buzzword-based framework and technology choices, such as “big data” and “the cloud,” are considered and compared along several important dimensions that together affect the success of the underlying data processing model.

2.1. Digital Banking Transaction Monitoring Landscape

Banking transaction monitoring systems support the supervision and compliance of banking transactions. Compliance with AML, KYC, and other region-specific regulations requires monitoring transactions in near real-time for suspicious behaviour. These systems detect the vast majority of normal transactions and a small percentage of suspicious transactions that require further manual investigation. Digital banking transaction monitoring systems are understandably gaining increasing attention. Blockchain platforms are newly emerging technologies that act as clouds. By integrating cloud computing with blockchains, banking institutions can execute smart contracts with an automated task manager and depend on the consensus mechanism of blockchains to establish cloud data service providers. The fundamental digital banking transaction monitoring system facial detection model's performance is adequate and requires further improvement. The model's training time is slightly longer due to the larger dataset used in this study.

The data characteristics are diverse and require many transformations during ingestion to support monitoring rule execution and machine learning model prediction in compliance with the LEDGER models. The LEDGER Pattern provides a cloud computing architectural reference for processing the new generation of banking transactions as required by regulations with near real-time integration of batch processes and proposal solutions for scalable, high-performance cloud banking transaction monitoring systems. Public cloud service providers are not suitable for services that require data control. For banking transaction monitoring systems, particular consideration must be given to data locality, security, privacy, and licensing. Public cloud service providers may be used to host third-party machine-learning services for automatic rule proposals, training, and testing, supported by a data connector and federation service that allows data exchange with other cloud providers. The transactions of next-generation Ethereum-based Digital Banking and LEDGER-compliant Digital Banking Transaction Monitoring systems can support the proposal of Darknet Link Prediction Models and Scalable-Rendering Darknet Poster Production Processes.




Fig 2: Digital Banking Transaction Monitoring Landscape

2.2. Data Characteristics and Compliance Considerations

Digital banking transaction monitoring must handle a variety of data types to track compliance with regulatory requirements. These systems primarily monitor transaction data but also process source data or additional typologies to verify users’ identities, assess risks associated with transaction destinations, detect suspicious behaviour, and evaluate transferred amounts and frequencies. Generally, these operation characteristics enable a specific model to maintain adequate results across data types. Consequently, a monitoring system dedicated to a specific category, such as credit card transactions or foreign trade operations, can be deployed and focused on detected events. However, misbehaviour remains difficult to automatically classify into patterns, pessimistically leading to whole-business-rule-based systems. Compliance rules vary based on local regulations, but some monitoring models remain consistent and serve as process references among institutions.

Beyond transaction processing or customer monitoring, banks must fulfill regulatory reporting requirements. These often include daily reports of significant or high-risk transactions, recurrent transactions above a pre-defined amount, and periodic reporting for money laundering or other illicit finance investigations. Data needed for such reports must undergo a different monitoring process to confirm their accuracy. Automated operations, therefore, require such additional typologies for compliance purposes. Automating daily operations not only addresses near-real-time requirements but also frees operators to focus on critical analysis of complex cases. The need to adhere to complex compliance processes reinforces stringent data governance, validation, and quality requirements. Data source profiling plays a critical role in regulating data quality inputs. Fulfilling such management demands usually require a data quality management project distinct from other initiatives.

Equation 2: End-to-end latency derivation



execution, feature lookup, and alert emission delays.

$$T_{\text{total}} = T_{\text{ingest}} + T_{\text{q}} + T_{\text{rule}} + T_{\text{feature}} + T_{\text{emit}}$$

For a simplified M/M/1-like approximation of the active worker pool, expected waiting plus service time behaves as:

$$W \approx 1 / (c\mu - \lambda)$$

If time is measured in seconds and $c = 12$, $\mu = 25$, $\lambda = 200$:

$$W = 1 / (12 \times 25 - 200) = 1 / 100 = 0.01 \text{ s} = 10 \text{ ms}$$

Assume the remaining components are: ingestion = 12 ms, rule evaluation = 40 ms, feature lookup = 8 ms, alert emission = 15 ms.

$$T_{\text{total}} = 12 + 10 + 40 + 8 + 15 = 85 \text{ ms}$$

3. Cloud Computing Fundamentals for Monitoring Systems

Cloud computing is essential for modern digital banking transaction monitoring due to its scalable infrastructure, optimal resource usage and pay-per-use model. Cloud computing comprises three service models (SaaS, PaaS and IaaS) that define various functionalities of the components and systems. Different deployment models (public, private, hybrid and community clouds) define the locality of the components and systems.

In general, data locality is not an issue for transaction data analytics; however, for clients residing in other regions or countries, data locality, privacy and security are essential. The cloud services infrastructure provider generally includes additional features to handle data privacy and security in a particular region or country by movements of the data processing and storage in defined regions or countries. Some geo-specific regulations such as the General Data Protection Regulation place restrictions related to data processing and movement across borders. To comply with these additional regulations, sensitive data can be stored in a private cloud in a specific region with additional costs, while the other region-specific transactional data can be stored in a public cloud for reduced cost.



Fig 3: Fundamentals Of Cloud Computing

3.1. Cloud Service Models and Deployment Models

Virtualized services provided by cloud suppliers are often categorized into three service models: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). SaaS makes applications accessible over the internet, hence externally for organizations and their users. PaaS provides a platform (hardware, operating system, storage, network configuration, middleware, web services development components, programming languages, databases, and more) that enables customers to build and deploy applications in the cloud without having to worry about purchasing, configuring, and managing these elements. IaaS permits customers to construct their own computing environment, enabling access to virtualized networking, server, and storage infrastructure as well as the automation of service provisioning.

The three service models can be expanded into five categories, including Host as a Service (HaaS) providing any platform but that is accessible by remote access and the customer is mainly responsible for only the Data Management and Application levels. HaaS can be seen as an application-oriented variant of IaaS in the sense that any (cloud-compliant) software that can connect to the platform via the HaaS API can manage it, such as a remote desktop, a terminal emulator, or local machine management software. SaaS hosting services simplify the use of application because data and operating system issues are handled by the service provider. Telephony services that bypass the bulk of control and switch management, implementing a similar abstraction between service provider and user as for computing virtualization, are also beginning to appear.



Equation 3: Daily transaction volume and storage derivation

The paper distinguishes between live monitoring and historical storage. Storage sizing begins with transaction count per day.

$$N_d = \lambda \times 86400$$

Using $\lambda = 200$ tx/s:

$$N_d = 200 \times 86400 = 17,280,000 \text{ transactions/day}$$

If each stored transaction plus metadata consumes $s = 2.5$ KB:

$$S_d = N_d \times s$$

$$S_d = 17,280,000 \times 2.5 \text{ KB} = 43,200,000 \text{ KB/day}$$

Convert to gigabytes:

$$S_d \approx 43,200,000 / 1024 / 1024 = 41.20 \text{ GB/day}$$

For 90-day retention in the hot-plus-warm zone:

$$S_{90} = 90 \times 41.20 = 3708 \text{ GB} \approx 3.62 \text{ TB}$$

If a replication factor $r = 3$ is used for resilience:

$$S_{rep} = r \times S_{90} = 3 \times 3.62 = 10.86 \text{ TB}$$

3.2. Data Locality, Security, and Privacy

Data locality is an aspect of performance that must be considered during CCloud Data Processing Model Selection. Digital activity is highly distributed, as are the Cloud Resources required to process it. Wherever possible, Computing Resources should be located Close to the data in order to minimize transfer Costs. Cloud Service providers do offer user's the ability to select from a number of regional data centers in which they wish to Deploy their Services. However, the Data Sharing mechanisms used to enable Cross-region Communication, for example between the Edge locations of a Content Delivery Network and the Cloud hosting the originating data, may introduce geographical proximity Challenges.

Security and privacy considerations can also be determining Factors in the Design and Deployment of Cloud-Based Solutions. Although Data resided on third-party Infrastructure with its inherent risks, the Financial Services sector is accustomed to Sharing data with third Parties, in the form of money laundering and terrorism financing like GAFI recommendations. Should data be shared with third Parties, the most secure and compliant manner to do so shall be utilized. Keep in mind that data designated as sensitive or personal should not leave its Country of origin without prior Permission of the Data subject or its legal representative(s). Data not stored in the Country of origin must be treated as foreign data and Complex Control assessments must be executed before deploying the Solution in Production.

4. Architectural Patterns for Scalable Processing

The volume of financial transactions processed by digital banks has enabled models for online transaction processing (OLTP) with adequate availability and reliability. However, factors such as international exposure, digital criminals operating with malicious intentions, and demands from regulators for transaction monitoring have generated a need for transactional monitoring systems capable of scaling. Therefore, the corresponding architectural patterns that support monitoring transaction data in a continuous and massive manner are essential.

A real-time transaction-monitoring system needs to consume a high volume of data with very low latency, so a stream-processing architecture that can meet these requirements is a perfect fit. Typical stream-processing frameworks with a defined architecture, such as those based on the Lambda or Kappa architecture patterns, provide the operational requirements. Nevertheless, stream-processing systems usually use simple rules tagged in the transactions for risk detection or acknowledge a high volume of legitimate transactions as normal with only a reduced set of transactions marked as suspicious. The acceptable false-negative rates allow banking institutions to use statistical machine-learning models to monitor transactions in near real time. In this case, the batch-processing architecture is applied with near-real-time integration of the results into the monitoring system.




Fig 4: Scalable Web Architectures Concepts & Design

4.1. Stream Processing Architectures

In a pure stream processing architecture, transaction monitoring systems exploit System as a Service (SaaS) components to detect potentially fraudulent transactions as soon as operations are executed. Transaction streams are analyzed in real time, and alerts are generated as messages when predefined monitoring rules are violated. Monitoring message streams and executing rule actions rely on a set of prediction models hosted as services in the cloud infrastructure. When necessary, alerts are sent to clients through



messaging systems like WhatsApp. Such a stream-based approach is appropriate for high-velocity transaction streams for which any latency introduced by batch computation would hinder the detection of potential fraud.

Although stream processing architectures provide support for monitoring digital banking operation flows in real time without introducing buffering delays, they also exhibit limitations. In particular, real-time rule-based monitoring alone is not sufficient to satisfy the new digital banking transaction monitoring requirements mandated by the Financial Action Task Force. Hence, regulatory compliance necessitates the continual development and deployment of machine learning models for the detection of unusual and suspicious behavior, a task for which a purely stream-oriented architecture is unsuitable. To meet these demands, a monitoring architecture that offers near-real-time fraud services and supports batch-style prediction is required.

4.2. Batch Processing with Near-Real-Time Integration

Builders of monitoring systems that do not depend on constant low-latency processing of transaction data may prefer batch processing. Real-time inspection of incoming transactions proceeds in accordance with defined rules, while an auxiliary module processes batched historical data. Processing of incoming transaction data can follow such architectural pattern. A continuous operator inspects the data stream for signatures of established rules that require fast reaction to protect valuable assets. Detected alerts create records for long-term storage and notify clients who can take immediate preventive action. A second process inspects historical transaction data with machine learning algorithms for suspect trends or that deviate from users' normal behavior.

With Cloud Service Providers offering storage in different locations around the world, the available data align with the production and suppression yearning for faster computation. Components that ingest, prepare, and preprocess the data relinquish the constraints of being on-site in order to obtain fast geography-oriented storage. The data do not feed directly into machine learning algorithms for model training, but are used as recent training data that substitute the older ones. The Cloud Service Provider can automatically train models that build when new data arrive, even in the case that no label for the training is given. Alerts are usually produced through set rules and together with historical data can yield knowledge that can be more relevant for the business.

Equation 4: Rule-based alert volume derivation

Rule systems typically approve most traffic and escalate only a small fraction. Let p_s be the suspicious fraction.

$$\text{Alerts/day} = p_s \times N_d$$

With $p_s = 0.5\% = 0.005$ and $N_d = 17,280,000$:

$$\text{Alerts/day} = 0.005 \times 17,280,000 = 86,400 \text{ alerts/day}$$

If only 8% of these require human analyst review after secondary prioritization:

$$\text{Human review/day} = 0.08 \times 86,400 = 6,912 \text{ cases/day}$$

If one analyst handles 220 reviewed cases/day:

$$\text{Analysts needed} = 6912 / 220 = 31.42 \approx 32 \text{ analysts}$$

5. Data Ingestion and Transformation Techniques



When ingesting data from data sources hosted externally, cloud-based transaction monitoring systems can employ connectors optimized for their information storage and delivery format. Doing so may mitigate operational complexities as well as cost, ultimately resulting in total cost of ownership estimates nearing cost-effectiveness. Federators are also available for use with cloud-based database services, enabling cloud transaction monitoring systems to access data residing across multiple repositories hosted both on-premises and within the cloud. Data analytics offerings with such capabilities also tend to provide transformation-based connectors to either ingest or stream data to or through the service. By transcending static federators, these connectors visually abstract the transformation workflow supporting the ingestion process, saving valuable initiatives, paving the way for quicker internal evaluations, and supporting or acting upon a more dynamic perspective of data hosted in repositories marked by cloud services.

A common challenge in the context of data preparation is managing data quality and schema evolution, all while reducing engineering overhead. Modern data processing systems provide enhanced cost-effective capabilities in these areas by intelligently grouping, merging, and optimizing storage as well as compute resources. Automatic data optimization helps guarantee best performance and lowest cost for query processing, while smart management of data life cycles helps keep data fresh and relevant to a given data analytics operation.

Modern-scale data analytics engines also incorporate advanced capabilities for dealing with schema evolution. Whenever a monitoring system generates a new type of structured information, the data processing system capable of speedily discovering its schema and automatically adding it. Machine learning techniques help automatically validate the quality of ingested, stored, and served data, triggering alerts any time data inconsistency or recorded behavior is detected.

Transaction ingestion uses uniform table structures and data formats. Non-uniform schemas have to be transformed and adapted, either once during integration or on-the-fly. Integration also needs automatic schema validation and update detection, with provable data quality.

Equation 5: Confusion-matrix metrics derivation

The article discusses rule-based monitoring, anomaly detection, and false positives. The standard metrics are derived as follows.

$$\text{Precision} = TP / (TP + FP)$$

$$\text{Recall} = TP / (TP + FN)$$

$$F1 = 2 \times \text{Precision} \times \text{Recall} / (\text{Precision} + \text{Recall})$$

Illustrative example for a hybrid system over one evaluation batch: TP = 950, FP = 1485, FN = 50.

$$\text{Precision} = 950 / (950 + 1485) = 950 / 2435 = 0.390$$

$$\text{Recall} = 950 / (950 + 50) = 950 / 1000 = 0.950$$


$$F1 = 2 \times 0.390 \times 0.950 / (0.390 + 0.950) = 0.553$$

5.1. Connectors and Data Federation

Digital banking transaction monitoring systems are connected to various sources (e.g., different banking applications, AML systems), deployed on cloud services, or a mixture of both. Internal connectors enable reliable data acquisition, as they often use

specific APIs that control data flow and quality. External sources do not provide such guarantees, meaning periodic data acquisition is preferable.

Transaction source databases are usually specific to each bank and institution, containing at least the same transaction and account properties. Fraud detection (FD) and know-your-customer (KYC) systems offer their own information. A data lake, at least partly on the same vendor, simplifies integration. A federated query engine provides access to both consolidated and external systems. Data from external systems or vendors, such as on-premise legacy systems, on SFTP servers, or in proprietary formats, are uploaded or integrated on ad-hoc schedules.



5.2. Schema Evolution and Data Quality

Consequently, the ingestion and transformation stage must support the consistent evolution of source schemas as well as the observation of incoming data quality. To provide these capabilities, data ingestion, transformation, and storage are defined and implemented in two complementary layers: an external layer, operating independently of the core data-processing architecture, and an internal layer, natively integrated. The external layer is responsible for ingesting replicated source data using custom-built Microservices. Each Microservice exposes a database’s schema on a dedicated HTTP(S) endpoint. Data source metadata are stored in configuration files managed by a command-line utility. Data ingestion involves frequently invoking the Microservices within a personal computer to download all new data from each database into flat files, which are subsequently transformed




through an external ETL tool into a star schema data warehouse on Microsoft Azure. The internal layer manages ingestion and transformation for replicated data sources. Custom connectors feed the external tables of an Azure Synapse Analytics serverless SQL Pool by establishing federation with the raw data. These routes implement parsing, type mapping, quality assurance, and automated data-level schema evolution for parquet files stored in Azure Blob Storage.

Automated batch scheduling ensures that new arrivals in the source databases propagate through to the data warehouse. Statistical data profiling tools produce quality reports and automatically alert the operation team when indicators drop below predefined thresholds. Quality-related procedures for incoming data, such as type mappers and checks, are reusable components deployed through the data-processing architecture in the form of Apache Spark transformation functions. This means that wrappers for third-party quality monitoring tools can also be defined in a relatively straightforward manner.

6. Model Deployment and Monitoring for Compliance

Two parallel approaches are employed to monitor compliance in near real time. The first **Rule-Based Model** represents a straightforward integration into cloud data streaming services, where logical rules can be executed continuously on each transaction processed. Such an implementation covers business rules for transactions' wielder, such as regular receiver, amount, or frequency patterns. The second **Machine Learning (ML) Model** provides an alternative approach to compliance monitoring, where anomalies in the data are detected automatically. The model is trained offline, learning on historical transactions previously categorized as legitimate or fraudulent, and making it able to distinguish between benign and suspicious behaviors. In contrast to rules, which require explicit knowledge of the users' behavior and police definitions of fraud, ML learns normal behavior by itself, exposing those transactions that deviate most from the others.

The results provided by the ML model are typically not used in a standalone way, due to the high volume of false positives even in advanced models. It requires investigation by specialists to determine if the detected transaction is indeed suspicious. To facilitate the investigation task, the exposed transactions are prioritized, so that the human resources may focus on the most suspicious ones. The challenge is to determine the interest ranking with the lower number of features possible, such that the monitoring is human-friendly and the fraud probability remains high.



6.1. Real-Time Rule-Based Monitoring

Real-time transactional monitoring based on business rules and risk indicators helps meet regulatory compliance requirements. Banking institutions monitor customers, payment channels, and geographical regions for abnormal activity patterns. Multiple institutions can co-operate by sharing few granular data points to create similar datasets for collaborative insights, thus serving both narrow and broad perspectives.

Rule definitions are easy to express in DSLs, enabling non-technical regulatory officers to put in place rule definitions. Fintech developments often see ad-hoc rule definitions but with a higher investment and maintenance costs. Real-time predictable monitoring is easy to put in place but operational burden usually leads to had-hoc solutioning patterns. Scalability, update/maintenance management and technical expertise management are challenging aspects of simple rule definition based compliance monitoring. It provides a near-realtime view of alerts but false positive detection and investigation operational burden is high. Insights-driven industries typically invest on these monitoring areas frequently based on board mandates. Despite these burdens, this pattern is widely adapted and few signals shared across institutions can help mitigate compliance cost significantly.

6.2. Machine Learning for Anomaly Detection

Advanced machine learning models based on transactional data of clients are widely deployed for detecting fraud events.



Normally such models require periodic updates, to ensure that they are based on the most recent behavior of clients. Trained on historical labelled events, they rely on historical characteristics of fraudulent actions. Underlying data preparation for training involves label generation, preprocessing, feature engineering and feature selection. For training, data is usually sampled: non-fraud events on the entire period and fraud events on the time window during which they occurred. Different approaches exist for testing and anomaly detection: (i) the models are engaged to assign an anomaly score to each transaction, being necessary a threshold to define a fraud or (ii) each transaction is classified as fraud or non-fraud event. In first case, fraud score is compared with a threshold, derived from an initial evaluation of precision and recall on testing data; in second, the model is engaged for real-time classification. In both cases, clients present portfolio of predicted fraud events, being relevant to analyse these results periodically. Anomaly monitoring ensures clients respect the behavioral normality detected by the models, regardless of whether the model is online deployed or not. A set of unique identifiers is generated and monitor behavior without the model being continuously active.

For anomaly-based supervised learning, experimental studies exploit Dynamic Time Warping (DTW) for feature extraction and Probabilistic Neural Networks for classification. A pre-processing phase applies time-distributed features calculation for all accounts in a time frame, enabling classification using time series as input. The limitation and main hindrance is the need for a labelled dataset including all account activity combined with the corresponding features. A pre-processing stage produces the DTW feature dataset required for any classification model. In temporal segmentation, sequential data is split into ranges using start and end tokens, generating a set of sequences for each category. Segments are sampled but, differently from historical samples, all segments respect their original women/men ratio.

7. Conclusion

In this study, scalable data processing models for digital banking transaction monitoring systems were synthesized. One of the practical implementations is a real-time rule-based monitoring system deployed in a public cloud. In a recent evaluation, no false negatives were found in one month of production data processing, and the whole system handling both data ingestion and monitoring fits on a single machine. Nevertheless, transaction monitoring is a challenging compliance task often involving the detection of rare events and requiring the analysis of an entire data set to identify relevant patterns. Consequently, several scalable enterprise solutions are being considered for compliance use, permitting near-real-time processing with the workload executed in a distributed cloud model. The proposed cloud data processing models enable bank transaction monitoring infrastructure to be deployed in a highly efficient way but not necessarily in a cost-efficient manner.

The use of cloud services for enterprise data processing is expected to grow consistently over the coming years. Public cloud services, such as Amazon AWS or Microsoft Azure, dynamically provide on-demand resources that increase operational flexibility. Local private cloud installations, or hybrid-deployed cloud services that combine local infrastructure with the cloud as an overflow, can also achieve highly scalable data processing but without the extra benefit of paying only for actual resource consumption. For cloud-based data management services, the ability to postpone the selection of the actual data repository until the moment of access (data federation), the possibility of automatic connection points for external data sources, and assistance in data query building also avoid latency in data provisioning. Nevertheless, the fact that cloud service providers handle customer data means that priority must be given to data locality when defining the locations of cloud resources, and that appropriate data security measures must be integrated within and around the cloud infrastructure during its actual deployment.



Design quantity	Formula	Worked result	Why it matters
Minimum workers	$c \geq \lambda / (\rho_{\text{target}} \mu)$	12 workers	Keeps the stream engine stable
End-to-end latency	$T_{\text{total}} = \Sigma \text{ stage delays}$	85 ms	Meets near-real-time monitoring
Transactions per day	$N_d = \lambda \times 86400$	17.28 million/day	Sizes downstream systems
Daily storage	$S_d = N_d \times s$	41.20 GB/day	Drives data lake design
90-day replicated storage	$3 \times 90 \times S_d$	10.86 TB	Explains cloud storage need
Daily alerts	$p_s \times N_d$	86,400/day	Quantifies operational load
Analyst team size	review cases / analyst capacity	32 analysts	Shows need for prioritization
Hybrid precision	$TP / (TP + FP)$	39.0%	Human review still required
Hybrid recall	$TP / (TP + FN)$	95.0%	Few suspicious events missed
Posterior fraud probability	$(TPR\pi) / (TPR\pi + FPR(1-\pi))$	8.68%	Class imbalance challenge
30-day training rows	$\lambda \times 86400 \times 30$	518.4 million	Necessitates scalable batch compute
Illustrative monthly cost	sum of cost components	\$1,838.20	Links architecture to economics

Table : Consolidated implementation table

7.1. Future Directions

Experimental results of a pattern matching method for detecting online banking fraud using extrusion, as leveraged in the Shiny app for the ICCD case study, the visual application prototype for predictive machine learning for credit card fraud detection, and a model for real-time anomaly detection for Internet of Things (IoT) devices, have been explored within previous research. Future research can conduct any of the SCoM processing approaches, examine unsupervised and semisupervised models in any of the configurations, or explore federated machine learning for transaction monitoring in financial services.

The conceptual models can facilitate experimentation using different techniques and data storage and with combinations of supervised and unsupervised methodologies. For large banking institutions, the patterns can be implemented in a cloud or hybrid cloud for near-real-time detection of card fraud, credit fraud, or internet fraud. For smaller banks, a pattern-agnostic deployment



in a private data lake can enable periodic batch monitoring with no upstream machine learning or complex event processing. In both scenarios, the result can improve decision support in regulatory audits and investigations.

8. References

1. Jullum, M., Løland, A., Huseby, R. B., Ånonsen, G., & Meegahapola, L. (2020). Detecting money laundering transactions with machine learning. *Journal of Money Laundering Control*, 23(1), 173–186.
2. Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317–331.
3. Benchaji, I., Douzi, S., El Ouahidi, B., & Jaafari, J. (2021). Enhanced credit card fraud detection based on attention mechanism and LSTM deep model. *Journal of Big Data*, 8, 151.
4. Inala, R. (2020). Building Foundational Data Products for Financial Services: A MDM-Based Approach to Customer, and Product Data Integration. *Universal Journal of Finance and Economics*, 1(1), 1-18.
5. Rtayli, N., & Enneya, N. (2020). Enhanced credit card fraud detection based on SVM-recursive feature elimination and hyper-parameters optimization. *Journal of Information Security and Applications*, 55, 102596.
6. Carcillo, F., Le Borgne, Y.-A., Caelen, O., & Bontempi, G. (2021). SCARFF: A scalable framework for streaming credit card fraud detection with Spark. *Information Fusion*, 41, 182–194.
7. Mangalampalli, B. M. (2021). Scalable Data Warehouse Architecture for Population Health Management and Predictive Analytics. *World Journal of Clinical Medicine Research*, 1(1), 1-18.
8. Nagy, E., Lovas, R., Pintye, I., Hajnal, Á., & Kacsuk, P. (2021). Cloud-agnostic architectures for machine learning based on Apache Spark. *Advances in Engineering Software*, 159, 103029.



9. Vorobyev, I., & Krivitskaya, A. (2022). Reducing false positives in bank anti-fraud systems based on rule induction in distributed tree-based models. *Computers & Security*, 120, 102786.
10. Mattaparthi, R. (2021). Unified Data Lineage and Quality Governance Framework for Multi-Source Sensor Streams in Heavy-Duty Powertrain Manufacturing. *Online Journal of Mechanical Engineering*, 1(1), 1-15.
11. Jensen, R. I. T., & Iosifidis, A. (2022). Qualifying and raising anti-money laundering alarms with deep learning. *Expert Systems with Applications*, 205, 119037.
12. van der Aalst, W. M. P., Adriansyah, A., & van Dongen, B. F. (2020). Process mining and fraud detection: A data-driven approach to financial transaction monitoring. *Business Process Management Journal*, 26(4), 905–923.
13. Canhoto, A. I. (2021). Leveraging machine learning in the global fight against money laundering and terrorism financing: An affordances perspective. *Journal of Business Research*, 131, 441–452.
14. Loganathan, R. (2021). Integrated Risk and Compliance Frameworks for Global Data Center Operations: A Governance-Centric Approach. *Universal Journal of Computer Sciences and Communications*, 1(1), 1-26.
15. Gupta, A., Dwivedi, D. N., Shah, J., & Jain, A. (2021). Data quality issues leading to sub optimal machine learning for money laundering models. *Journal of Money Laundering Control*, 25(3), 551–555.
16. Ruiz, E. P., & Angelis, J. (2021). Combating money laundering with machine learning: Applicability of supervised-learning algorithms at cryptocurrency exchanges. *Journal of Money Laundering Control*, 25(4), 766–778.
17. Nagy, E., Lovas, R., & Kacsuk, P. (2021). Scalable cloud-based machine learning architectures for big data analytics. *Journal of Cloud Computing*, 10, 1–15.
18. Jullum, M., Løland, A., Huseby, R. B., Ånonsen, G., & Meegahapola, L. (2020). Machine learning for transaction monitoring and suspicious activity detection in financial institutions. *Journal of Money Laundering Control*, 23(2), 247–261.
19. Yandamuri, U. S. (2021). A Comparative Study of Traditional Reporting Systems versus Real-Time Analytics Dashboards in Enterprise Operations. *Universal Journal of Business and Management*, 1(1), 1-13.



20. Błaszczyszki, J., de Almeida Filho, A. T., Matuszyk, A., Szeląg, M., & Słowiński, R. (2021). Auto loan fraud detection using dominance-based rough set approach versus machine learning methods. *Expert Systems with Applications*, 163, 113740.
21. Gai, K., Qiu, M., & Sun, X. (2020). A survey on financial fraud detection using machine learning and big data analytics. *IEEE Access*, 8, 188127–188141.
22. Abdallah, A., Maarof, M. A., & Zainal, A. (2020). Fraud detection system: A survey of machine learning techniques. *International Journal of Computer Applications*, 182(1), 1–10.
23. Mukesh, A., & Aitha, A. R. (2021). Insurance Risk Assessment Using Predictive Modeling Techniques. *International Journal of Emerging Research in Engineering and Technology*, 2(4), 68-79.
24. Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2020). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences*, 479, 448–455.
25. Roy, A., Sun, J., Mahoney, R., Alonzi, L., Adams, S., & Zou, X. (2020). Deep learning detecting fraud in credit card transactions. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2020, 1–5.
26. Awoyemi, J. O., Adetunmbi, A. O., & Oluwadare, S. A. (2021). Credit card fraud detection using machine learning techniques: A comparative analysis. *International Journal of Computer Science and Information Security*, 19(2), 45–52.
27. Reddy, V. A. R. (2021). Challenges in Standardizing Member Eligibility Data Across Multi-Payer Healthcare Ecosystems. *International Journal of Medical Toxicology and Legal Medicine*, 24(3), 1-19.
28. Patil, S., Vanarse, S., & Saini, S. (2021). Detection of fraudulent transactions using SAS Viya machine learning algorithms. *Procedia Computer Science*, 190, 204–209.
29. Kaur, H., & Singh, S. (2021). Identification of non-typical international transactions on bank cards of individuals using machine learning methods. *Procedia Computer Science*, 190, 178–183.
30. Chang, V., Doan, L. M. T., Di Stefano, A., & Sun, Z. (2022). Digital payment fraud detection methods in digital ages and Industry 4.0. *Computers & Electrical Engineering*, 100, 107734.



31. Cardoso, M., Saleiro, P., & Bizarro, P. (2022). LaundroGraph: Self-supervised graph representation learning for anti-money laundering. Proceedings of the 3rd ACM International Conference on AI in Finance.
32. Inala, R. Designing Scalable Technology Architectures for Customer Data in Group Insurance and Investment Platforms.
33. Vanini, P., Rossi, S., Zvizdic, E., & Domenig, C. (2022). Time-aware and interpretable predictive monitoring system for anti-money laundering. Machine Learning with Applications, 8, 100306.