International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

A Hybrid Cloud Framework for Secure Velocity Aggregation and Persona Enrichment

Nikhil Sagar Miriyala[1*]

Submitted:05/05/2024 **Revised:**20/06/2024 **Accepted:**28/06/2024

Abstract: In modern data-driven ecosystems, especially within regulated domains like healthcare and financial technology, critical data sources are often distributed across on-premises infrastructure and public cloud environments. These divisions may arise due to compliance, system design, or operational constraints. Rather than treating such separation as a limitation, this paper explores how secure collaboration between distributed systems can be leveraged to improve machine learning models through velocity aggregation, the capture and analysis of real-time transactional or behavioral patterns, and persona enrichment, the synthesis of richer user profiles from multiple secure sources. The paper focuses on scenarios where sensitive data, such as personal health records or financial transactions, must remain confidential, even during processing. To address this, we incorporate confidential computing environments (e.g., AWS Nitro Enclaves) to ensure that data processed in the cloud remains encrypted and inaccessible to unauthorized actors, including the cloud provider itself. Further, this paper proposes a hybrid cloud architecture that enables bidirectional, privacy-preserving data aggregation and federated model enhancement without exposing raw data. The framework demonstrates how secure statistical insights and model signals can be exchanged across cloud and on-prem systems, resulting in mutually improved ML model performance, while maintaining regulatory compliance and strict data confidentiality. We validate the framework using examples from healthcare and fintech, highlighting its broad applicability across any domain that handles personally identifiable information (PII).

***Keywords**: Hybrid Cloud, Confidential Computing, AWS Nitro Enclaves, Velocity Aggregation, Persona Enrichment, Data Security*

1. Introduction:

1.1 AI/ML in Regulated Domains: FinTech and Healthcare

In recent years, the adoption of Artificial Intelligence (AI) and Machine Learning (ML) has transformed industries such as financial technology (fintech) and healthcare, offering intelligent automation, risk detection, anomaly tracking, and personalized services. In fintech, ML models are employed for fraud detection, risk scoring, transaction monitoring, and creditworthiness evaluation [1]. Similarly, in healthcare, ML is used for early disease detection, personalized treatment planning, clinical decision support, and remote patient monitoring. [2]

However, these advancements rely heavily on the availability of rich and relevant data. Due to the fragmented nature of enterprise data infrastructures, especially in regulated sectors, lot of this data exists across on-premises systems and public cloud environments, making it challenging to unify, analyze, and learn from in a privacy-preserving manner. [3]

1.2 Velocity Aggregation and Persona Enrichment:

Two key techniques that can significantly enhance ML models are velocity aggregation and persona enrichment:

a) Velocity Aggregation refers to the process of collecting and analyzing high-frequency, real-time behavioral or transactional data to identify patterns, trends, or risks. It's especially useful in detecting sudden changes in user behavior or emerging threats. [4]
b) Persona Enrichment involves building a comprehensive, multi-dimensional profile of a user or entity by synthesizing data from multiple systems. This enables better personalization, predictive modeling, and risk classification. [5]

1Senior Software Engineer, Visa Inc., USA
**Corresponding Author: nmiriya7@gmail.com*

When combined, these concepts enable systems to move beyond siloed analysis, allowing AI/ML models to continuously learn and adapt based on dynamic behaviors and deeper user context.

1.3 Motivating Use Cases for Cross-Source Aggregation

1) Financial Use Case:

A bank may run separate ML pipelines for:

a. Account-to-account transaction risk analysis in the cloud (using scalable real-time engines), and
b. Credit card fraud detection on-premises (due to compliance with PCI-DSS).

While each model serves a different purpose, sharing anonymous statistical data between them, such as unusual transaction frequency or behavioral anomalies, can enhance both models' accuracy in real time, leading to stronger fraud detection capabilities, without even sharing raw data.

2) Healthcare Use Case:

A healthcare provider may run:

a. Remote patient monitoring analytics (e.g., from wearable devices) in the cloud, and
b. In-hospital EMR-based diagnostic models on-premises (due to HIPAA and data residency constraints).

By aggregating patient activity patterns (velocity) from wearables and diagnostic insights (persona) from EMRs, the provider could build a more nuanced patient profile, thus enabling early intervention or more accurate treatment predictions, all without exposing raw patient records.

1.4 Rise of Confidential Computing for Sensitive Workloads

To protect such highly sensitive data in cloud environments, organizations are increasingly turning towards confidential computing technologies. These systems allow computations to run inside isolated, hardware-protected environments, such as Trusted Execution Environments (TEEs) like AWS Nitro Enclaves, Intel SGX, or Azure Confidential VMs [6] [7]. Key benefits include:

a. In-memory encryption during processing.
b. Zero visibility to cloud administrators or other tenants.
c. Strong hardware-rooted attestation for workload integrity.

This enables organizations to move sensitive workloads to the cloud, including ML model training or inference, without compromising on data privacy or compliance requirements.

1.5 Towards a Federated and Privacy-Preserving Framework

Given the distributed nature of critical data and the increasing need for collaborative intelligence, this paper proposes a hybrid cloud framework that enables secure velocity aggregation and persona enrichment across cloud and on-prem systems.

Our approach ensures:

a. Confidential data is never decrypted outside its processing enclave.
b. Statistical insights can flow securely between environments.
c. Federated learning principles are used to improve both ends without exposing raw data.

The next section presents the architecture and components of this hybrid system, including how data flows are secured, aggregated, and utilized to enhance downstream AI models.

2. High-Level Architecture & Core Components

2.1 Baseline Architecture:

In a typical cloud-native machine learning system, such as one used for credit card transactions fraud detection, velocity aggregation and persona profiling are employed to enhance model accuracy using real-time behavioral signals and longitudinal user patterns. Figure 1 shows a baseline architecture where the ML model using velocity aggregation and persona enrichment techniques, with the system deployed in a public cloud environment, with several modular components working together to

ingest transactions, enrich them with derived features, and generate risk predictions.

2.1.1 Core Components:

1) Actor Interaction:
    a) A user or system (the *actor*) initiates a transaction request.
    b) The request is routed through the Prediction API, which captures the transaction and triggers real-time inference.
2) ML Model Invocation:
    a) The ML Model receives both the new transaction and real-time features from the Velocity Aggregation Service.
    b) It predicts an outcome (e.g., fraud risk score) and returns the result via the API
3) Velocity Aggregation Service:
    a) Receives the transaction data in real time.
    b) Derives high-frequency behavior metrics such as transaction bursts, velocity patterns, or spending anomalies.
4) Persona Profiling Service:
    a) Periodically aggregates historical data, including behavioral and transactional summaries.
    b) Stores enriched persona information in Amazon Aurora, representing long-term behavioral fingerprints of users or accounts.
    c) Triggers regular retraining of the ML model based on updated personas and velocity data, enabling adaptive learning.

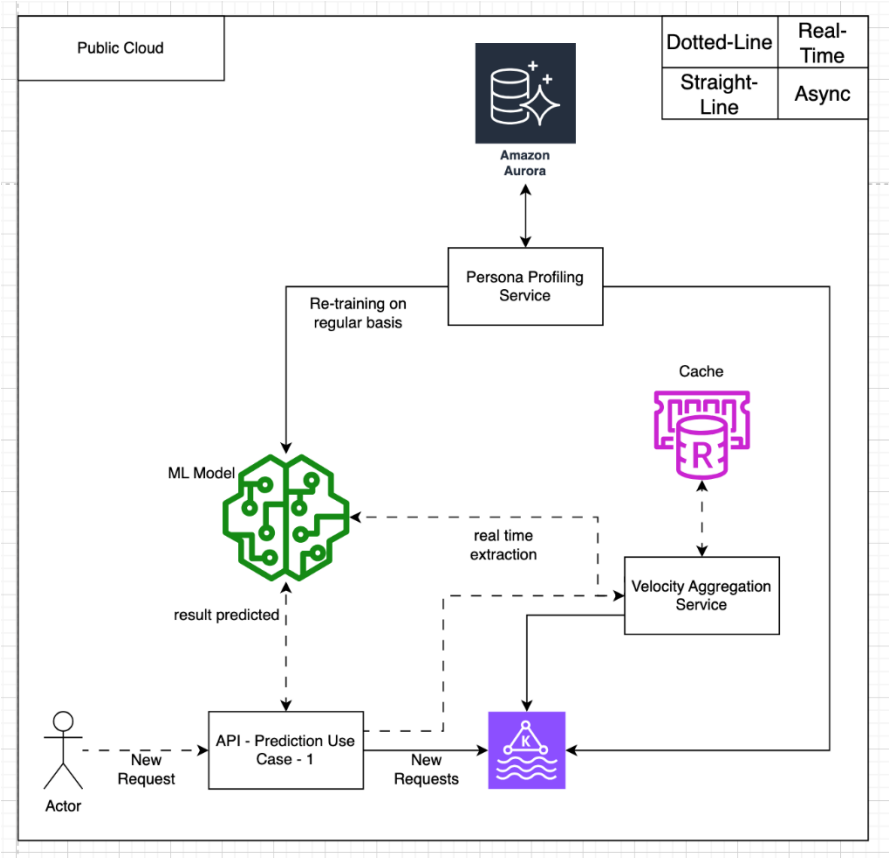

Figure 1: Baseline Architecture of a Single-System

2.1.2 System Characteristics:

1) Works entirely within a single cloud domain.
2) No sharing of data across systems or regulatory boundaries.
3) Velocity and persona data are only derived from internal transactions.
4) Provides foundational support for fraud detection, behavior modeling, and transaction scoring.

2.2 Federated Learning Architecture for Cross-Use Case Enrichment:

To enable collaborative intelligence across systems handling sensitive data, we extend the baseline design into a federated architecture that supports cross-use-case learning between public cloud-based and on-premises ML systems. This is particularly applicable in scenarios where different product lines (e.g., account-based and card-based systems) operate under regulatory constraints that prevent direct raw data sharing.

As shown in Figure 2, this architecture allows both environments to share encrypted, aggregated insights and persona enhancements to improve local ML models, while preserving strict data locality, compliance, and confidentiality guarantees. Let us understand the system in more detail in the subsections below.

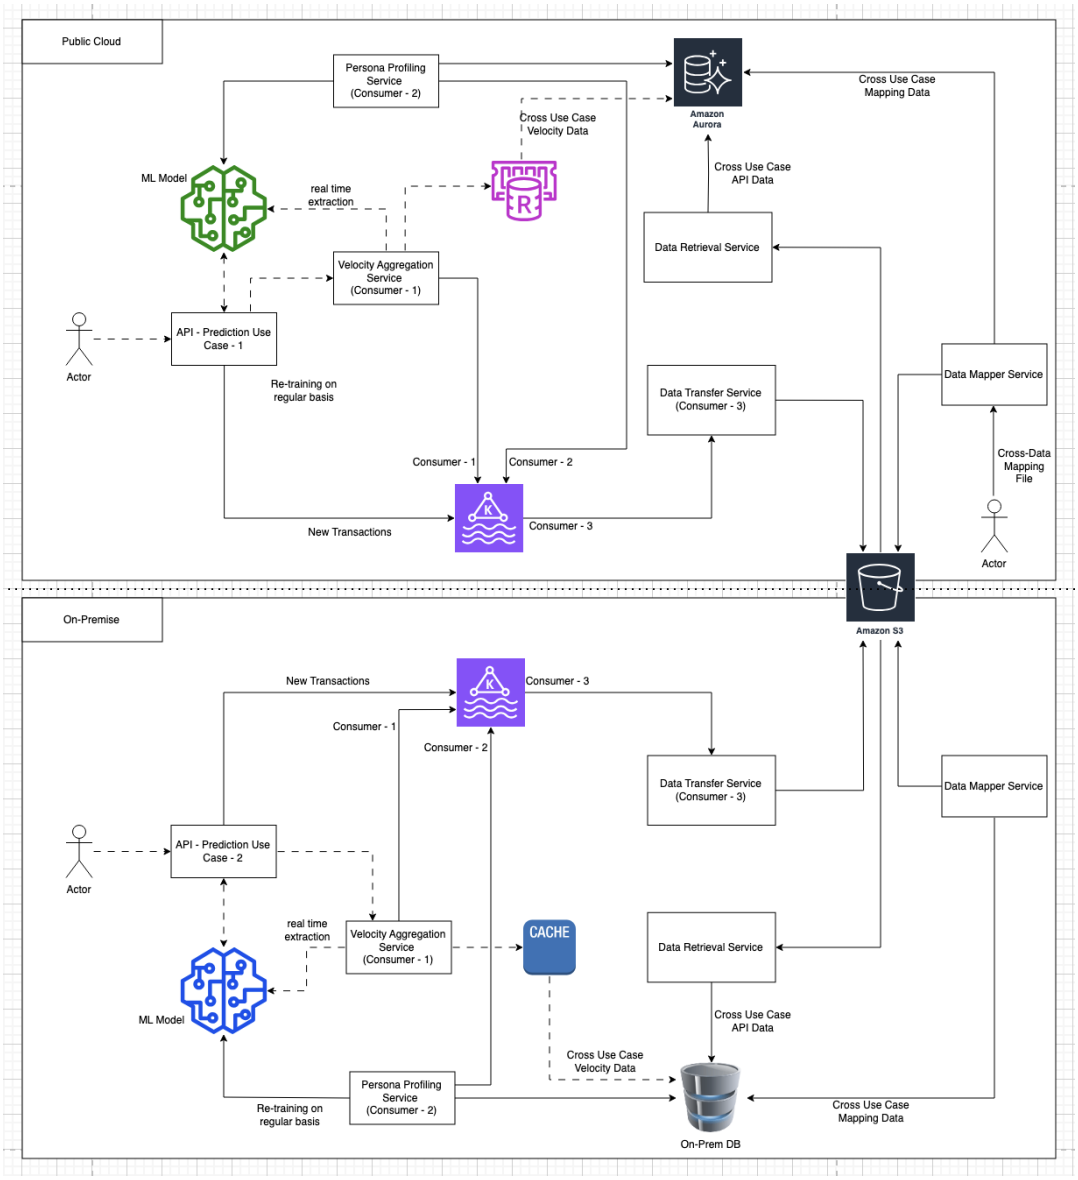

Figure 2: Federated Learning Architecture

2.2.1 Dual Environment Design

1) Public Cloud System:
    a) Hosts the ML pipeline for one use-case (ex: Account-to-Account) and operates inside AWS Nitro Enclaves to ensure that data remains encrypted during processing.

b) Relies on AWS KMS for data-at-rest encryption and PGP for controlled, secure cross-environment transmission.

2) On-Premise System:
   a) Hosts the ML system for the second use-case (ex: Credit Card-based), which must operate in a tightly controlled infrastructure, often behind firewalls and subject to PCI-DSS or similar regulatory standards.
   b) Relies on AES-256 or similar in-house encryption technologies for data-at-rest encryption and PGP for controlled, secure cross-environment transmission.

2.2.2 Secure Data Mapping & Identity Bridging

1) Clients supply a Cross-Data Mapping File, for example, mappings between account numbers and owned credit card IDs.
2) The Data Mapper Service on the cloud ingests this mapping and stores it in Amazon Aurora, encrypted using AWS KMS.
3) A secondary, PGP-encrypted version of the mapping file is pushed to Amazon S3, where it's securely pulled into the on-prem environment.
4) On-prem systems decrypt the file using their private PGP key and populate a local mapping database.

2.2.3 Real-Time and Asynchronous Data Flow

Each environment processes its own prediction use case, and transactions trigger the following operations:

1) Prediction Pipeline:
   a. Transactions are received via the API and passed to the local ML model.
   b. The model invokes real-time feature extraction from the Velocity Aggregation Service.
2) Velocity Aggregation:
   a. Tracks behavioral metrics (frequency, volume, recency) and caches them for immediate use.
   b. Retrieves behavioral metrics with respect to the cross-use-case from the database.
3) Persona Profiling Service:
   a. Processes new transactions to build and update long-term behavioral profiles for each user or account.
   b. Updates are stored in a structured database (e.g., Aurora in cloud, or a local database on-prem), which is later used for periodic ML model retraining.
   c. The Profiling Service is also responsible for extracting and preparing data for cross-use-case feature enrichment.
4) Cross-Use Case Signal Extraction:
   a. All new transactions are also consumed by a Data Transfer Service
   b. On the cloud, the service PGP-encrypts the relevant transaction subset and uploads it to S3.
   c. On the on-premise system, an equivalent transfer service does the same, using a different PGP key.
5) Data Retrieval and Integration:
   a. Each side runs a Data Retrieval Service to fetch and decrypt incoming cross-use-case data.
   b. This data is integrated into the Velocity Aggregation Service and Profiling Service to enhance downstream predictions.

2.2.3 Secure Execution and Trust Boundaries

1) All cloud-side components handling sensitive data run within AWS Nitro Enclaves:
   a) Ensures data remains encrypted in memory and inaccessible to the cloud provider.
   b) Only remotely attested, enclave-aware services can participate in secure communication.
2) All inter-environment data exchange is:
   a) Double-encrypted: KMS for internal storage, and PGP for transfer.

b) Pushed to and pulled from Amazon S3, which acts as a secure, asynchronous data bus.

3) No raw data or personally identifiable information (PII) is shared and only aggregated, preprocessed statistical insights flow across boundaries.

2.2.5 ML Model Training and Feedback Loop

1) Each system maintains its own ML model, which is periodically retrained using:
   a) Local persona and velocity data.
   b) Cross-use-case enriched features, extracted securely from the federated exchange pipeline.
2) This setup supports federated learning principles, where:
   a) Raw data remains in its respective trust boundary.
   b) Derived knowledge is exchanged to benefit both models.

2.2. 6 Benefits of the Architecture

1) Enables collaborative model enhancement across systems without breaching compliance.
2) Leverages confidential computing for secure cloud processing.
3) Establishes cryptographic trust and isolation across hybrid environments.
4) Lays the foundation for multi-party federated learning in regulated domains.

3. Evaluation and Validation

To validate the underlying hypothesis of our architecture, that cross-use-case persona and velocity enrichment can enhance ML model performance without directly sharing raw data, a simplified experiment has been conducted using publicly available datasets. This approach allows us to simulate the federated enrichment effect without requiring the full production-ready deployment of the architecture.

3.1 Dataset Selection and Setup

We selected two semantically related datasets from public sources:

1) Dataset A: Represents credit card transaction records with timestamped activity, user IDs, and fraud labels. [8]
2) Dataset B: Simulates account-to-account transactions, including features like sender ID, transaction amount, time intervals, and risk labels. [9]

Each dataset is used to train a Random Forest classifier independently to predict transaction risk. However, since the two datasets originate from distinct sources with no shared identifiers, a core challenge is simulating a federated learning scenario with meaningful correlation between them. To address this, we created a synthetic user ID mapping, representing the shared persona of an account holder across both datasets. Using common Python libraries such as pandas, numpy, and scikit-learn, we generated a pool of synthetic user identifiers and assigned them randomly to subsets of each dataset.

3. 2 Cross-Use-Case Feature Enrichment

We generate derived features such as:

1) Velocity metrics: e.g., average time between transactions per user, burst frequency.
2) Persona traits: e.g., risk trend scores, average transaction size.

And then:

1. Train base models independently using only local features.
2. Enrich each dataset with anonymized, aggregated features from the other dataset, simulating cross-use-case exchange via a secure pipeline.
3. Retrain the models with the enriched data and compare performance.

3. 3 Metrics and Results

Below Table shows the results, where we compared the model performance in three configurations:

1) Base: Using only local features.
2) + Cross Velocity: With additional velocity-based features from the peer system.
3) + Cross Velocity + Cross Persona: With full enrichment from both behavioral dimensions.

Model	AUC-ROC	F1-Score	Precision	Recall
Account2Account (Base)	0.81	0.75	0.77	0.73
Account2Account + Cross Velocity	0.86	0.79	0.81	0.77
Account2Account + Cross Velocity + Cross Persona	0.89	0.82	0.84	0.80
Card (Base)	0.79	0.72	0.74	0.70
Card + Cross Velocity	0.84	0.76	0.78	0.74
Card + Cross Velocity + Cross Persona	0.88	0.81	0.83	0.78

These results clearly show that cross-domain feature exchange, even in simulated form, improves predictive performance for both use cases. The greatest lift is observed when both velocity and persona signals are included, validating the effectiveness of the proposed architecture even without centralized data sharing.

4. Challenges and Limitations

While the proposed hybrid cloud framework demonstrates significant potential in securely aggregating high-velocity data and enhancing user personas across distributed environments, several practical and technical challenges remain. These limitations are critical considerations for organizations intending to adopt such architectures in regulated domains.

1) Confidential computing environments such as AWS Nitro Enclaves introduce a steep learning curve and operational overhead. Developers must adapt their data processing pipelines to work within constrained enclave environments, which lack direct network or persistent storage access. This often necessitates redesigning data ingress, egress, and serialization workflows, and integrating secure communication protocols like vsock or secure proxies, which adds to implementation complexity. [10]

2) In hybrid cloud setups, synchronization of data signals and model updates between the on-premises and cloud systems can be non-trivial [11]. Variability in network latency, batch update schedules, or asynchronous event timing can lead to inconsistent or delayed enrichment on either side. Maintaining alignment across both environments is crucial for preserving the predictive power of real-time risk or diagnostic models.

3) Since ML models are trained and updated in distributed environments, model drift or concept divergence may occur over time [12]. If the data characteristics in the cloud and on-prem systems evolve independently, federated updates might degrade model performance. Implementing drift detection mechanisms and periodic alignment strategies is essential to maintain effectiveness across environments.

4) While the framework aims to comply with privacy regulations (e.g., GDPR, HIPAA, PCI-DSS) by ensuring that raw data remains local and only aggregated signals are exchanged, legal interpretations of data processing boundaries may vary across jurisdictions. Organizations must perform rigorous legal and compliance reviews before adopting such cross-boundary data architectures, especially when operating in global or multi-tenant environments. [13]

5) Running confidential computing workloads comes at a performance and cost premium. Nitro Enclaves, for instance, allocate isolated vCPUs and memory from the parent EC2 instance, which can reduce the total usable capacity of the host machine. Additionally, enclave-enabled architectures often require dedicated provisioning, monitoring, and custom deployment pipelines, increasing infrastructure and operational costs. [10]

5. Future Trends and Improvements

1) Multi-Party and Cross-Organizational Federated Learning: Extending the current framework to support federated learning across institutions, such as between hospitals, banks, or research networks, would enable collaborative model training at a larger scale. This would require developing secure protocols for multi-enclave trust establishment, policy enforcement, and cross-domain identity

management, all while maintaining strict data locality. [14]

2) Real-World Deployment and Benchmarking: To validate the framework's effectiveness, a real-world deployment in collaboration with a financial institution or healthcare provider is essential. This would allow for detailed benchmarking of security, latency, model performance, and compliance adherence, providing evidence of its scalability and practical value.

6. Conclusion

In this paper, we have proposed a hybrid cloud framework for secure velocity aggregation and persona enrichment, designed to address the challenges of machine learning across distributed, privacy-sensitive environments. By combining confidential computing with federated learning, our architecture enables organizations to securely process, aggregate, and exchange statistical insights across cloud-based and on-premises systems, without exposing raw data or violating regulatory constraints.

It demonstrates how this framework is particularly suited for domains such as healthcare and financial services, where sensitive data is often siloed due to compliance, operational, or trust boundaries. Through practical examples, such as risk analysis of account-to-account and credit card transactions, it showcases how secure data collaboration can enhance the performance and contextual richness of AI models.

By leveraging trusted execution environments like AWS Nitro Enclaves and enabling bidirectional model enhancement via federated learning principles, the proposed system maintains a strong balance between privacy, compliance, and machine learning performance.

As organizations continue to adopt hybrid cloud strategies, architectures like the one presented here will play a critical role in enabling privacy-preserving AI at scale. Future work will focus on expanding this framework to multi-party collaborations, integrating differential privacy, and conducting real-world deployments to validate its scalability and practical value.

7. References

[1] B. Stojanović *et al.*, "Follow the Trail: Machine Learning for Fraud Detection in Fintech Applications," *Sensors*, Feb. 2021, doi: 10.3390/S21051594

[2] S. K. UmaMaheswaran, N. K. Munagala, D. Mishra, B. Othman, S. Sinthu, and V. Tripathi, "The role of implementing Machine Learning approaches in enhancing the effectiveness of HealthCare service," in *2022 2nd International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE)*, Apr. 2022. doi: 10.1109/icacite53722.2022.9823656

[3] Z. Li, V. Sharma, and S. P. Mohanty, "Preserving Data Privacy via Federated Learning: Challenges and Solutions," *IEEE Consumer Electronics Magazine*, May 2020, doi: 10.1109/MCE.2019.2959108

[4] P. K. Buchhop, "Use of velocity in fraud detection or prevention," Oct. 27, 2011

[5] R. Baden, A. Bender, N. Spring, B. Bhattacharjee, and D. Starin, "Persona: an online social network with user-defined privacy," in *ACM Special Interest Group on Data Communication*, Aug. 2009. doi: 10.1145/1592568.1592585

[6] S. Chakrabarti, T. Knauth, D. Kuvaiskii, M. Steiner, and M. Vij, "Trusted execution environment with Intel SGX," 2020. doi: 10.1016/B978-0-12-816197-5.00008-5

[7] Amazon Web Services, "Security Design of the AWS Nitro System," Feb. 2024. Accessed: May 2024. [Online]. Available: https://docs.aws.amazon.com/pdfs/whitepapers/latest/security-design-of-aws-nitro-system/security-design-of-aws-nitro-system.pdf

[8] Kaggle, "Credit Card Transactions Fraud Detection Dataset," Accessed: May 2024. [Online]. Available: https://www.kaggle.com/datasets/kartik2112/fraud-detection

[9] Kaggle, "Bank Account Fraud Dataset Suite (NeurIPS 2022)," Accessed: May 2024. [Online]. Available:

https://www.kaggle.com/datasets/sgpjesus/bank-account-fraud-dataset-neurips-2022

[10] W. Ellison and C. Ryan, "Computational operations in enclave computing environments," Apr. 30, 2020

[11] Feng, C., Yang, H. H., Wang, S., Zhao, Z., & Quek, T. Q. (2023). Hybrid learning: When centralized learning meets federated learning in mobile edge computing systems. *IEEE Transactions on Communications, 71*(12), 7008-7022. https://doi.org/10.1109/TCOMM.2023.3172394

[12] F. E. Casado, D. Lema, M. F. Criado, R. Iglesias, C. V. Regueiro, and S. Barro, "Concept drift detection and adaptation for federated and continual learning," *arXiv: Learning*, May 2021, doi: 10.1007/S11042-021-11219-X

[13] R. Singh, "Legalization of Privacy and Personal Data Governance: Feasibility Assessment for a New Global Framework Development," 2016. doi: 10.20381/RUOR-291

[14] M. J. Sheller *et al.*, "Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data," *Scientific Reports*, Jul. 2020, doi: 10.1038/S41598-020-69250-1