PERSPECTIVE

Fighting Financial Crime Together: Federated Computing for Cross-Bank AML Compliance

Money laundering is the lifeblood of organized crime, terrorism, and corruption. Globally, an estimated $2 trillion is laundered annually, yet less than 1% of illicit funds are ever seized.[1] For financial institutions (FIs), the cost of compliance has ballooned to tens of billions of dollars, but the criminals continue to adapt.

The core problem isn't a lack of effort—it's a lack of collective intelligence. Sophisticated money laundering schemes, like layering through multiple banks or across borders, exploit the fact that FIs mostly operate in isolation. This is where Federated Computing (FC) offers a pathway to collective defense without compromising the confidentiality of private data.

The State of AML Today: A Battle Fought in Isolation

Anti-Money Laundering programs are under constant pressure. Banks are caught between escalating regulatory demands and increasingly sophisticated criminal networks. The current state is characterized by three main challenges that federated computing is uniquely positioned to solve:

The Data Silo Problem: Collaboration is limited by strict regulation, technical challenges, and competitive dynamics

The most significant challenge is the inability of banks to collaborate effectively. Money laundering schemes rarely use just one bank. They move funds across multiple institutions, payment platforms, and even jurisdictions. To connect the dots and see the full criminal network, institutions need to share data, but in many cases, they are restricted by:

- Privacy Regulations: Strict laws like the General Data Protection Regulation (GDPR) and various national banking secrecy acts prohibit the sharing of raw, Personally Identifiable Information (PII) like transaction details or customer profiles. FIs risk massive fines and reputational damage by pooling data in a central repository.
- Security Concerns: Banks often face reluctance to share sensitive customer data due to fundamental security. Even when trying to collaborate on AML efforts, an organization is naturally concerned about its counterparty's security standards and data governance practices. Sharing data requires banks to entrust their sensitive information to another party, which creates a new attack surface.
- Disparate Data: Disparate data formats create a significant barrier to cross-bank collaboration against financial crime. Even if regulatory and security hurdles are overcome, financial institutions use widely different internal data schemes, formats (e.g., legacy systems), and definitions for transactional data and customer profiles.

This means that data from one bank is often structurally incompatible with data from another. The effort required to move, standardize, and reconcile these data formats centrally is technically complex, time-consuming, and prohibitively expensive, effectively extinguishing the possibility of effective large-scale, real-time cooperative analysis.

- Competitive Concerns: Banks are competitors. Sharing raw customer or transaction data is a non-starter due to the fear of revealing proprietary business insights.

- Cost and Technical Barriers: Sharing and consolidating large volumes of financial data for AML purposes is an expensive technical undertaking. Banks are often unwilling to bear the substantial cost burden associated with moving and storing redundant data sets. Specifically, they face two main hurdles: egress fees and duplicative storage fees. These costs and the technical complexity of re-centralizing data from disparate systems often create a decisive financial barrier to collaborative AML solutions

2. The Problem of False Positives

Current AML systems are often based on rigid, rules-based logic and historical data from a single institution. When dealing with the immense volume of daily transactions, this leads to an overwhelming number of false positive alerts - legitimate transactions incorrectly flagged as suspicious.

- Analyst Fatigue: High false positive rates (often over 90%) burden compliance teams with manual reviews, leading to burnout, increased operational costs, and, critically, a higher risk of analysts missing a real crime buried in the noise.[2]

- Cost and Inefficiency: The manual investigation of false positives is a massive drain on resources, with institutions spending billions annually to clear alerts.

3. The Evolving Threat Landscape

Financial crime tactics are evolving faster than AML systems can adapt. Criminals are exploiting new technologies and financial channels, including:

- Cryptocurrency and DeFi: New, borderless instruments are used for layering and integration.

- Mule Networks: Sophisticated networks of individuals are used to move illicit funds across accounts in multiple FIs.

- Trade-Based Laundering (TBL): Complex invoicing and trade mechanisms mask the movement of funds, often involving international partners.

Traditional AML systems, which only see part of the activity, struggle to keep pace with these multi-institutional and cross-border threats. Combining that with the fact that the financial sector has recently seen an increase in both the volume and severity of AML enforcement actions globally, banks are exposed more than ever to AML risk and regulatory sanctions.

Regulators have signaled their intolerance for systemic control failures and outdated technology, demanding institutions step up as financial crime becomes increasingly complex. This message was delivered starkly with the recent $3.09 billion fine against TD Bank, a record penalty imposed because the bank's prolonged AML failures enabled organized crime to transfer hundreds of millions of dollars unnoticed.[3]

What is Federated Computing? Collaboration without Centralization

Federated Computing (FC) is an emerging technique of collaborators across organizations performing computations locally where data resides, and only sharing aggregations of the results back to a centralized location. (Contrast this with centralizing data from multiple collaborators and then running computations.) These computations can range from the simple (e.g. counts of items across different silos) to slightly more complicated (e.g. transform the units of those items from pounds to kilograms) to more complicated still (e.g. train an ML model on those data across multiple silos) to very complex (e.g. deploy one partner's software on another partner's data behind their firewall).

FC addresses critical challenges regarding data privacy, data security, and data sovereignty. In FC, data custodians maintain complete control over their data. This means their organization's existing security controls remain in place. By not sending data to collaborators outside of their organization, custodians reduce any risk of misuse of the data by their partners. Collaborators can also introduce additional privacy-enhancing techniques such as differential privacy or k-anonymization into their FC projects for further protection of data privacy.

Federated Learning, one of the primary subsets of FC, is a decentralized machine learning approach that enables multiple organizations to train a shared prediction model without ever sharing their raw, sensitive data. It's a "bring the code to the data" model, instead of the traditional "bring the data to the code."

How Federated Learning Works:

- Local Training: A shared, initial machine learning model is distributed to all participating banks. Each bank trains this model locally using its own proprietary, customer-specific transaction data. The raw data never leaves the bank's secure environment.

- Sharing Model Updates: Instead of sending the raw data, each bank sends only the model updates (the learned parameters, weights, and insights - *the patterns of crime*) to a central, secure aggregator. These updates are typically encrypted or anonymized.

- Global Aggregation: The central server averages and combines all the model updates from the participating banks to create a single, more robust global model.

- Distribution of the Global Model: The improved global model is sent back to all participating banks, where it is used to refine their local models.

This cycle repeats continuously, allowing the model to learn from the collective experience of all participants—the entire ecosystem—while maintaining strict data privacy. The resulting model is exponentially more powerful than any one bank could build in isolation because it has learned from a vast, diverse, and representative sample of global financial transactions and criminal typologies.

A Collective Defense: How FC Fights Money Laundering

FC directly addresses the core challenges of modern AML, offering a path to "coopetition"— cooperation among competitors.

Connecting the Unconnectable: Detecting Cross-Institutional Crime

Identifying Mule Networks - sharing labeled accounts: A criminal network often spreads suspicious deposits across accounts in several different banks to evade local detection limits. By leveraging FC and confidential computing, financial institutions can securely collaborate on identifying and tracking mule accounts while preserving privacy. FC ensures that no raw data is shared, while still allowing institutions to flag, verify, and cross-check potentially fraudulent accounts.

- Payment Fraud - training Anomaly Detection models on global data: By their nature, sophisticated fraud rings operate across multiple jurisdictions and correspondent banks, making patterns hard to spot to any single institution. Training models on a single bank data often results in limited detection capabilities and high false positive rates. FL directly addresses this challenge: multiple global banks can collaboratively train a shared anomaly detection model to identify these cross-border schemes. Crucially, the sensitive, raw transaction data never leaves the confines of the originating bank's secure environment, ensuring complete confidentiality and adherence to strict data sovereignty laws while harnessing collective intelligence.

Enhancing Accuracy and Cutting Costs: Reducing False Positives

- Richer Datasets: By learning from a broader, more diverse dataset of legitimate and illicit transactions across the industry, the federated model becomes far better at distinguishing normal customer behavior from a genuine financial crime threat.

- Optimizing the Rulebook: The improved model leads to a significant reduction in false positives, freeing up human analysts to focus their time and expertise on the truly high-risk alerts. This not only cuts compliance costs but also improves the job satisfaction and effectiveness of the AML team. Early tests have shown FL can reduce false positives by over 60%.[4]

Assuring Regulatory Compliance and Trust

- Data Sovereignty: FC / FL ensures that customer data remains within the originating bank's physical and jurisdictional boundaries, solving the cross-border data transfer problem that plagues global banks trying to comply with regulations like the GDPR.

- Auditability and Explainability: Modern FCL platforms are designed to be explainable, meaning the model's decision-making process can be audited, which is crucial for regulatory reporting and filing a Suspicious Activity Report (SAR).

Real-World Adoption: Federated Computing in the Financial Sector

FC is moving beyond the theoretical and into active pilot programs and live deployments, spearheaded by innovative FIs and RegTech partners.

Singapore: Regulatory Backing for Collaborative Intelligence

Singapore's central bank, the Monetary Authority of Singapore (MAS), has been a major driver for innovation in financial crime prevention. Instead of enforcing a central data pool, MAS has fostered an environment where privacy-enhancing technologies like FC can thrive, enabling collaboration against sophisticated cross-border crime.

- The Regulatory Framework: MAS launched the COSMIC (COllaborative Sharing of ML/TF Information & Cases) platform to allow banks to securely share information on customers exhibiting high financial crime risk. While COSMIC uses structured data sharing, the entire regulatory philosophy is highly supportive of the underlying principle of FL: collective defense through shared insights.

- Industry Pioneers: Technology providers have enabled this collaborative model with their platforms, which incorporate FL engines. These platforms are used by various FIs operating in the region, including local institutions and global players in Southeast Asia.

- Key Outcome: Institutions using these FC-enabled platforms in Singapore have reported significant results, including a 60-70% reduction in false positives and three times faster average investigation closure times, demonstrating the model's direct impact on operational efficiency and effectiveness.

Hong Kong: Interbank Collaboration

The Hong Kong Monetary Authority (HKMA) has actively encouraged the use of AI, including federated learning, to bolster AML efforts. Banks like Airstar Bank and livi bank have initiated pilot programs to use FL for interbank cooperation. The specific objective is to enhance the monitoring of suspicious financial activities and reduce the high operational costs associated with traditional AML methods, all while maintaining the stringent data privacy requirements of the region.[5]

European Payments: Banking Circle

Banking Circle, a fully licensed payments bank focused on cross-border transactions, has implemented FC to enhance its transaction monitoring. Faced with the challenge of needing a globally powerful model while complying with the strict data residency requirements of different European nations, Banking Circle leveraged FC frameworks. Initial tests showed a marked improvement in detection precision and accuracy, demonstrating FC's immediate, tangible value in high-volume, cross-jurisdictional financial environments.[6]

Australia: Community Banking and Collective Intelligence

In Australia, the need for a stronger, real-time defense against fast payments has driven smaller institutions to adopt collaborative compliance models. Community-owned banks like Regional Australia Bank and Beyond Bank are utilizing platforms that incorporate a federated learning engine. This allows them to benefit from a "collective intelligence" that strengthens their detection capabilities and provides access to advanced typologies, proving that the technology is accessible and beneficial for institutions of all sizes, not just Tier-1 global players.[7]

The Rhino Federated Computing Platform (RhinoFCP)

Rhino Federated Computing has built the world's leading data collaboration platform, RhinoFCP. While maintaining a focus on security & privacy, we have turned RhinoFCP into a totally secure, totally extensible collaboration sandbox - allowing collaborators to run code on one another's data while ensuring that security, privacy, and legal teams can be comfortable. RhinoFCP offers flexible architecture (multi-cloud and on-prem hardware), end-to-end data management workflows (multimodal data, schema definition, harmonization, and visualization), a privacy screen (custom differential privacy budget, custom k-anonymization values), and allows for the secure deployment of custom code & 3rd party applications via persistent data pipelines.

RhinoFCP can also seamlessly integrate into banks' tech stacks, serving as middleware to any number of federated workloads run on distributed data.

Conclusion: The Future is Shared, Not Centralized

Federated Computing is more than just a technological upgrade; it represents a paradigm shift in the governance of Anti-Money Laundering. By enabling financial institutions to share insights instead of data, FC offers the first scalable, privacy-compliant solution to the data silo problem.

For an industry constantly under threat from sophisticated criminal forces, FC is the key to moving from an isolated, reactive defense to a unified, predictive, and proactive collective intelligence. This collaborative approach is essential to finally outpace the criminal underworld and build a financial system that is not only compliant but truly resilient.

Sources

1. https://amlwatcher.com/blog/what-is-money-laundering-with example/#:~:text=UNODC%20has%20its%20take%20on,seized%2C%20and%20frozen%20by%20authorities.
2. https://www.tookitaki.com/compliance-hub/federated-learning-in-aml-australia
3. https://www.justice.gov/archives/opa/pr/td-bank-pleads-guilty-bank-secrecy-act-and-money-laundering-conspiracy-violations-18b#:~:text=Office%20of%20Public%20Affairs%20%7C%20TD,United%20States%20Department%20of%20Justice
4. https://www.tookitaki.com/compliance-hub/federated-learning-in-aml-australia
5. https://hongkongbusiness.hk/financial-services/news/fcc-analytics-taps-airstar-and-livi-bank-federated-learning-aml
6. https://flower.ai/industries/finance/
7. https://www.tookitaki.com/compliance-hub/federated-learning-in-aml-australia

WANT TO LEARN MORE?

Contact us to explore the possibilities.
contact@rhinofcp.com

Rhino
Federated Computing