

Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review

Kehinde Hamed ^{1,✉}, Moshood Sorinola ²

1. *Financial Crime and Compliance Management, Utica University, Utica, USA*

2. *Finance and Business Analytics, Georgia State University, Georgia, USA*

Received: July 13, 2026 | Review began: July 29, 2026 | Review ended: September 03, 2026 | Published: October 12, 2026

© Copyright 2026

This is an open access article distributed under the terms of the Creative Commons Attribution License CC-BY 4.0., which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract

Money laundering moves an estimated 2%-5% of global gross domestic product through the financial system each year. Its central tactic, layering funds across multiple institutions, limits detection because each bank observes only part of the transaction chain. Machine learning has improved single-institution transaction monitoring, and graph-based methods exploit the network structure of illicit flows, but models trained within one institution remain constrained when laundering activity spans several institutions. Centralizing standardized and appropriately linked transaction data could improve cross-institutional visibility, but legal, privacy, operational, and competitive constraints make such pooling difficult in the United States. Federated learning offers an alternative by allowing institutions to train a shared model on local data while exchanging model updates rather than raw records. This narrative review synthesizes 45 scholarly sources published between 2013 and 2025, together with US statutes and supervisory guidance retrieved separately from official sources. It examines the evolution of anti-money laundering detection from rules to graph neural networks, compares horizontal, vertical, and graph-federated architectures, evaluates privacy risks and protections against information leakage, considers separate defenses against poisoning and backdoor attacks, and maps these techniques onto the US regulatory environment. The reviewed evidence shows that the core federated-learning, cryptographic, and statistical methods are well established, but AML-specific validation remains limited. Existing studies rely mainly on synthetic data, cryptocurrency networks, challenge datasets, or single-institution records. Standard federated learning can support shared model training but does not reconstruct cross-institutional transaction paths, which additionally require privacy-preserving entity resolution and secure cross-boundary graph protocols. None of the studies included in this review validated privacy-preserving federated anti-money laundering detection using real multi-institutional US banking data.

Categories: Regulatory Compliance and Governance, Security and Privacy, Machine Learning (ML)

Keywords: anti-money laundering, federated learning, differential privacy, graph neural networks, financial crime, privacy-preserving machine learning, transaction monitoring, homomorphic encryption

Introduction And Background

Money laundering is the process of disguising the proceeds of crime so that they appear to originate from legitimate sources [1]. Commonly cited estimates place its annual global scale at between 2% and 5% of gross domestic product, or approximately US\$800 billion to US\$2 trillion [2]. In the United States, financial institutions carry much of the responsibility for detecting suspicious activity under the Bank Secrecy Act and its implementing regulations. Transaction

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. *Cureus J Comput Sci* 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

monitoring has traditionally relied on deterministic rules, including transaction thresholds, velocity checks, and typology-based filters. Although these rules are transparent and auditable, they can generate large numbers of false-positive alerts, requiring compliance teams to devote substantial investigative resources to activity that ultimately proves benign [3,4].

Machine learning has been studied as a means of improving detection and reducing unnecessary alerts. Supervised classifiers, anomaly-detection methods, and sampling strategies designed for highly imbalanced data have often outperformed static-rule baselines in retrospective evaluations. However, the magnitude and reliability of these improvements depend strongly on dataset construction, label quality, class distribution, and evaluation design [3,5,6]. Graph-based methods extend this work by representing transactions as networks. Because laundering often moves funds through chains of intermediary accounts, transaction relationships and network structure may reveal patterns that individual transaction features do not capture [7-10].

In conventional banking systems, however, relevant transaction information is divided across institutional boundaries. When funds move through accounts held at several banks, each institution observes only part of the wider transaction path. A model trained within one institution, regardless of its sophistication, therefore operates with incomplete information about activity that crosses institutional boundaries.

One possible response would be to centralize standardized and appropriately linked transaction records in a shared repository. Such an arrangement could improve visibility across institutions, but it would still require reliable entity resolution, data harmonization, and representation of interinstitutional relationships. It would also raise substantial legal, privacy, operational, cybersecurity, and competitive concerns in the United States. Section 314(b) of the USA PATRIOT Act permits voluntary information sharing for specified anti-money laundering (AML) and counterterrorist-financing purposes, but it does not by itself establish a general authorization for institutions to pool broad transaction datasets for collaborative model development.

Federated learning offers an alternative approach by allowing multiple institutions to contribute to the training of a shared model while retaining their underlying transaction data locally [11,12]. Each participant computes local model updates, which are combined to improve a shared model. This arrangement can support collaborative parameter learning without requiring the direct centralization of raw transaction records.

Standard horizontal federated learning does not, however, reconstruct cross-institutional transaction paths. Combining model parameters from separate institutional datasets cannot by itself identify an account relationship or transaction edge that no single participant fully observes. Cross-institutional graph analysis additionally requires privacy-preserving entity resolution, representation of relationships spanning institutional boundaries, and secure protocols for exchanging information across those relationships. Federated learning must also be combined with appropriate safeguards because model updates may reveal information about local training data [13,14]. Relevant protections include differential privacy, secure aggregation, homomorphic encryption, secure multiparty computation, and separate defenses against poisoning and backdoor attacks [15-17].

This review examines the development of AML detection from deterministic rules to graph-based learning, compares horizontal, vertical, and graph-federated architectures, evaluates privacy and adversarial risks, and considers the regulatory and deployment implications of these approaches in the United States. Evidence derived from real transaction data is distinguished from evidence based on synthetic, simulated, challenge, or public cryptocurrency datasets because these settings provide different levels of support for real-world cross-institutional deployment.

Review

Review methodology

This review is narrative in design and aims to synthesize the state of the field. It does not claim the reproducibility of a systematic review or meta-analysis, although the sources were assembled through a structured and documented process. We searched the OpenAlex scholarly database and the arXiv repository to identify relevant literature. Search terms combined the application domain, including money laundering, AML, transaction monitoring, financial crime, and

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. *Cureus J Comput Sci* 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

financial fraud, with the methods of interest, including federated learning, vertical and horizontal federation, graph neural networks, differential privacy, homomorphic encryption, secure multiparty computation, and secure aggregation. Terms relating to adversarial risks, including gradient leakage, membership inference, model inversion, poisoning, and backdoor attacks, were also included.

The scholarly literature search was restricted to works published from January 1, 2013, through December 31, 2025. The database searches were conducted on July 5, 2026. Government statutes, regulatory guidance, interagency supervisory materials, and authoritative documents from international standards bodies were retrieved separately from official sources through July 11, 2026, and were not subject to the 2013-2025 scholarly publication limit.

Terms were combined with OR within each conceptual group. Records were required to match at least one application-domain term and at least one term from either the methods group or the adversarial group. The general Boolean structure was therefore (domain terms) AND (method terms OR adversarial terms). For example: ("money laundering" OR "anti-money laundering" OR "transaction monitoring" OR "financial crime" OR "financial fraud") AND ("federated learning" OR "vertical federated learning" OR "horizontal federated learning" OR "graph neural network" OR "differential privacy" OR "homomorphic encryption" OR "secure multiparty computation" OR "secure aggregation" OR "gradient leakage" OR "membership inference" OR "model inversion" OR "poisoning" OR "backdoor attack"). Searches were run over titles and abstracts.

Inclusion and exclusion criteria were applied during title-and-abstract screening. We included peer-reviewed research articles, field surveys, conference papers, and benchmark-dataset descriptions that addressed a federated or privacy-preserving learning technique relevant to financial-crime detection, a graph-based method for laundering or fraud detection, or an attack or defense within the federated-learning threat model. We excluded non-English records, duplicate records, and studies without a financial-crime or federated-learning or privacy connection. Duplicate records were identified by DOI and, where different versions had separate identifiers, by normalized title. Single-institution fraud studies were retained when they described methods relevant to cross-institutional use.

One reviewer assessed relevance. Seed studies for citation expansion were selected according to their relevance to the review questions and their citation counts in OpenAlex rather than through a formal centrality measure. One level of references and cited works was then examined, and additional records were screened for relevance. This subjective component is one reason the article is presented as a narrative review.

The final review included 45 scholarly sources comprising primary methods studies, field surveys, conference papers, and benchmark-dataset descriptions. Exhaustive retrieval and exclusion counts were not retained, and no formal study-quality or risk-of-bias assessment was performed. The directness of evidence to the target setting was evaluated narratively and is reflected in the tiered structure of Table 7; the tiers do not represent a formal assessment of methodological quality. Primary statutes, FinCEN guidance, interagency supervisory materials, and authoritative background sources from the Financial Action Task Force and the United Nations Office on Drugs and Crime were retrieved separately from official sources.

From rules to graphs

The False-Positive Economics of Rule-Based Monitoring

Any account of machine learning in AML has to start with the operational failure it was introduced to address. Conventional transaction-monitoring systems apply deterministic rules, amount thresholds, structuring detectors, and velocity and peer-group deviation checks that flag transactions for human review. These rules are auditable and map cleanly onto regulatory typologies, which is why they remain the industry backbone, but they are also indiscriminate. Alerts raised are overwhelmingly false positives, and compliance teams spend most of their investigative capacity clearing benign activity [3,4]. Reviews of the field frame this as the defining inefficiency of AML operations and the metric against which learning-based methods are implicitly judged. Raw accuracy matters less here than the ability to maintain recall while reducing the false-positive volume that dominates investigative costs [4,6].

Supervised, Unsupervised, and the Class-Imbalance Problem

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. *Cureus J Comput Sci* 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

Machine-learning approaches to AML divide along the familiar supervised and unsupervised axis, and both must contend with a defining property of the data. Laundering is extremely rare relative to legitimate activity, so labels are scarce and severely imbalanced. Chen et al. [3] cataloged supervised classifiers, decision trees, support vector machines, ensemble methods, and neural networks trained on labeled suspicious-activity data, alongside unsupervised anomaly-detection methods for the much larger unlabeled majority.

Empirical studies subsequently made class imbalance itself an object of investigation. Zhang and Trubey [5] showed that sampling and resampling strategies materially affect detection performance and that naïve training on the original class distribution may understate achievable recall. The systematic review by Ali et al. [6] reaches a consistent conclusion. Learning-based methods can outperform static rules on retrospective benchmarks, but reported performance is highly sensitive to dataset construction, label quality, class distribution, and evaluation protocol, which complicates comparisons across studies. Labanca et al. [4] address the labeling bottleneck through the Amaretto active-learning framework, which selects informative transactions for expert review so that a detection model can improve within a limited annotation budget.

Graph-Based and Graph Neural Network Methods

Graph-based methods model laundering as a transaction-network problem. Layering routes value through chains of intermediary accounts, so the structure of who pays whom can reveal patterns that transaction-level features do not capture. FlowScope, for example, represents laundering as high-volume, multistep movement through intermediary accounts and detects it from the topology of money flows [7].

Guo et al. [8] developed long-term bi-graph layer attention convolutional network, a graph-neural-network architecture that uses both the transaction graph and its reverse graph to capture information about the origins and destinations of blockchain transactions. The model also incorporates long-term layer attention to reduce the loss of useful information across graph-convolution layers. On the evaluated blockchain transaction data, long-term bi-graph layer attention convolutional network reported an accuracy of 0.9776, a precision of 0.9317, a recall of 0.8494, an F1-score of 0.8887, and an area under the receiver operating characteristic curve of 0.9806. Reporting precision, recall, and F1-score alongside accuracy is particularly important in laundering detection because the illicit class is substantially smaller than the legitimate class.

Other studies have extended graph-based AML detection in several directions. Lo et al. [9] use self-supervised node embeddings to reduce dependence on scarce labels, while Alarab and Prakoonwit [10] combine a long short-term memory network with a graph convolutional network to model temporal ordering in Bitcoin transactions. Egressy et al. [18] adapt graph neural networks to directed multigraphs and recover structural information that simpler message-passing models cannot represent. On their synthetic AML benchmark, the adaptations increased the minority-class F1 score of a graph isomorphism network baseline from 28.7% to 57.2%, with reverse message passing and port numbering accounting for much of the improvement. Cheng et al. [19] review the methodological trade-offs in graph-based financial-fraud detection, while mutual-learning and related architectures continue to be evaluated on blockchain transaction data [20].

Collectively, these studies indicate that graph structure can provide information beyond transaction-level features when the relevant network is observable. Their reliance on public blockchain or synthetic transaction networks, however, limits direct transfer to conventional banking systems, where transaction paths are fragmented across institutions.

The Single-Institution Ceiling

The graph methods reviewed above generally assume that the relevant transaction network is observable. Public blockchain datasets approximate this condition because the ledger is visible, which helps explain their frequent use in AML research [8-10]. Conventional banking data do not provide the same visibility. When funds move through accounts held at several institutions, each bank observes only a fragment of the wider transaction network. A model trained within

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. Cureus J Comput Sci 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

a single institution therefore operates on an incomplete network and may fail to recognize patterns that become apparent only when transaction paths across institutions are considered together. Figure 1 organizes the principal techniques proposed to support cross-institutional learning without centralizing raw transaction data.




FIGURE 1: A taxonomy of privacy-preserving techniques for cross-institutional anti-money laundering detection

Source: Original conceptual figure created by the authors from the literature reviewed
DP-SGD, Differentially Private Stochastic Gradient Descent; FL, Federated Learning; GNN, Graph Neural Network

Federated architectures for cross-institutional collaboration

The Federated Learning Primitive

Federated learning enables multiple decentralized data holders to collaborate in training a shared model while retaining their underlying data locally [11,12]. Each participating institution trains the model on its local records and transmits a model update to a coordinating server. The server aggregates the updates and redistributes the revised global model for further local training. Federation therefore changes the information exchanged between institutions: raw transaction records remain local while model parameters or updates are shared. Federated learning also introduces challenges involving communication efficiency, heterogeneous local data, privacy, and coordination across participants [12,21].

Matching Architecture to Data Topology

The appropriate federated architecture depends on how the participating institutions’ data relate to one another. Horizontal federated learning is suitable when institutions hold broadly similar feature sets for different customers or transactions. Each participant trains the shared model on its own observations, and the resulting updates are aggregated [12]. Because this arrangement generally does not require institutions to align overlapping customer records or disclose feature columns, it is operationally simpler than other forms of federation. However, local information may still be inferred from inadequately protected model updates.

Vertical federated learning applies when institutions hold different features relating to overlapping customers or entities. For example, a bank and a payment processor may observe different aspects of the same customer’s financial activity. The parties first identify corresponding records and then train a model using their separate feature sets without directly exchanging those features. Liu et al. [22] describe the entity-alignment and secure-computation mechanisms required for this arrangement. Vertical federation can incorporate information unavailable through horizontal federation, but record alignment introduces additional privacy, technical, and governance challenges.

How to cite this article:

Federated Graph Learning

Federated graph learning addresses settings in which participating institutions hold separate portions of a wider transaction network. A federated graph neural network could, in principle, learn from distributed transaction structures while the underlying institutional subgraphs remain locally controlled. Its feasibility, however, depends on identifying relationships that cross institutional boundaries and enabling information to move across those relationships without exposing the accounts or transactions involved.

Cross-institutional edges must first be established through privacy-preserving entity resolution or record linkage. Secure graph-learning protocols must then support message passing across institutional boundaries. Ordinary parameter aggregation cannot reconstruct an interbank transaction path that no single institution fully observes. Federated graph learning therefore involves capabilities beyond conventional horizontal federated learning.

These systems also combine unresolved challenges from both federation and graph learning, including institutional heterogeneity, cross-edge coordination, topology leakage, malicious updates, and backdoor attacks. Research on attacks against federated graph neural networks shows that robustness remains an open problem [17,23]. Figure 2 therefore presents federated graph learning as offering high potential for cross-institutional signal and detection while retaining comparatively low technical maturity.

Evidence for federated detection

Evidence for federated financial-crime detection remains limited, and much of it concerns credit-card fraud rather than money laundering. Salam et al. [24] evaluated conventional classifiers after resampling and a separate federated convolutional neural network. The reported 99.99% accuracy was achieved by a conventional random-forest classifier, not by the federated model. The federated model achieved accuracies of approximately 90% to 93% across the evaluated frameworks and optimization settings. The dataset contained 284,807 transactions, of which only 492, or approximately 0.17%, were fraudulent. Because the federated results were reported principally in terms of accuracy, they provide limited evidence about performance on the minority fraud class. Tang and Liu [25] also evaluated federated tree-based methods on transaction data. Aljunaid et al. [26] reported 99.95% accuracy, 99.95% sensitivity, 98.23% specificity, a positive predictive value of 100%, and a miss rate of 0.05% for their best-performing model. These results require caution because accuracy can remain high when fraudulent transactions represent only a small proportion of observations. Precision, recall, false-positive volume, and investigator workload are more informative measures of operational performance.

Kanamori et al. [27] provide stronger cross-institutional evidence using real transaction data from five Japanese banks. Under the DeepProtect protocol, each bank encrypts locally computed gradients, allowing the coordinating server to aggregate contributions without directly accessing the individual transaction records. The study demonstrated that privacy-preserving federated training could detect fraudulent transactions that were not detected by models trained on data from a single bank. It also detected some criminal accounts before the participating banks froze them. However, the study addressed financial fraud in Japan rather than cross-bank money-laundering typologies in the United States.

AML-specific federated evidence remains limited [28]. The available evidence therefore supports the technical feasibility of federated financial-crime detection more strongly than its effectiveness for cross-institutional AML. This distinction is reflected in Figure 2, where federated approaches show substantial potential for improving cross-institutional detection but only low-to-moderate technical maturity for the target setting.

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. *Cureus J Comput Sci* 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>




FIGURE 2: Comparative trade-offs among cross-institutional AML approaches across five dimensions
Source: Original conceptual figure created by the authors from the literature reviewed.
AML, Anti-Money Laundering; FL, Federated Learning; GNN, Graph Neural Network; ML, Machine Learning

Privacy guarantees and adversarial robustness

Federation Is Not Privacy

Federated learning reduces the movement of raw data but does not, by itself, provide a complete privacy guarantee. Model updates are derived from local training data and may reveal information about those data. Zhu and Han [13] demonstrated that training examples can be reconstructed from shared gradients without direct access to the original records. Nasr et al. [14] further showed that passive and active adversaries can conduct white-box inference attacks in both centralized and federated learning settings. In a financial consortium, inadequately protected updates could allow a coordinating server or another participant to infer sensitive information associated with customers or transactions. Practical deployments must therefore combine federated learning with additional statistical and cryptographic safeguards.

The Threat Model

Relevant threats include gradient leakage, membership inference, poisoning, and backdoor attacks. Gradient-leakage and model-inversion attacks seek to reconstruct information from model updates [13]. Membership-inference attacks attempt to determine whether a particular record was included in the training data [29], which may be especially sensitive when the record concerns a customer associated with suspicious-activity monitoring.

Poisoning and backdoor attacks target model integrity rather than confidentiality. A malicious or compromised participant may submit manipulated updates to reduce overall model performance or suppress the detection of selected transaction patterns. Xu et al. [23] demonstrate this risk in federated graph neural networks. Privacy attacks and integrity attacks therefore require different defenses, and measures designed to limit information leakage do not necessarily prevent malicious model updates [17].

Differential Privacy

Differential privacy provides a formal framework for limiting what a model or released update can reveal about an individual training record. Under the framework developed by Dwork and Roth [30], the inclusion or exclusion of one record has a limited effect on the probability distribution of a randomized computation's outputs. Federated implementations commonly clip model updates and add calibrated noise before the updates are released or aggregated.

Wei et al. [15] examine the effect of differential privacy on federated-learning performance, while El Ouadrhiri and Abdelhadi [31] review the broader privacy-utility trade-off in deep and federated learning. Noise may be added centrally by a trusted aggregator or locally by each participant. Local perturbation reduces reliance on the coordinating server but commonly imposes a greater loss of model utility. This trade-off is particularly important in AML because laundering cases are rare, and additional noise may obscure already weak minority-class signals.

Cryptographic Protection

Cryptographic methods protect model updates during transmission and aggregation. Secure aggregation allows a coordinating server to calculate the combined updates of participating institutions without viewing each institution's individual contribution. Bonawitz et al. [16] provide a practical secure-aggregation protocol that also accommodates participant dropout.

Homomorphic encryption permits specified computations to be performed on encrypted values. Fang and Qian [32] combine homomorphic encryption with federated learning, while Ou et al. [33] apply homomorphic encryption in a vertical federated-learning setting. Because fully homomorphic computation may impose substantial processing and communication costs, MaskCrypt selectively encrypts model parameters to reduce overhead [34]. Secure multiparty computation similarly allows multiple parties to calculate a joint function without directly revealing their individual inputs.

Cryptographic protection and differential privacy address different risks. Encryption and secure aggregation can prevent a server from viewing individual model updates during computation, but they do not necessarily limit what may be inferred from the final model. Differential privacy limits the influence of an individual record or participant on released updates or outputs, but it does not conceal the released values during transmission. Truex et al. [35] combine differential privacy with secure multiparty computation, and the broader literature supports layered protections when the threat model includes both untrusted participants and an untrusted coordinator [36].

Federated and Privacy-Preserving Financial-Crime Systems

Several studies apply federated or privacy-preserving methods to financial-crime detection, although not all provide the same level of formal privacy protection. Tang and Liang [37] combine federated learning with graph-based transaction modelling, graph extension, and adaptive model aggregation for credit-card fraud detection. Their approach supports collaborative learning across simulated institutional partitions, but it does not establish a formal differential-privacy or cryptographic guarantee beyond retaining the participating datasets locally.

Abadi et al. [38] develop Starlit for privacy-preserving cross-institutional fraud detection. Fed-RD combines vertical federated learning, differential privacy, and secure multiparty computation on realistically generated synthetic data [39]. Arora et al. [40] evaluate federated financial-anomaly detection with secure multiparty computation and noisy aggregation, while BAML [41] combines federated learning with Hyperledger Fabric to support decentralized compliance

How to cite this article:

checks. Baabdullah et al. [42] examine the combined use of federated learning and blockchain for credit-card fraud detection, and Nguyen and Dang [43] present a privacy-preserving federated model specifically for money-laundering detection.

These studies show that federated training can be combined with statistical, cryptographic, and distributed-ledger mechanisms. Their evaluation settings remain an important limitation. Most rely on synthetic data, public cryptocurrency records, challenge datasets, simulated institutional partitions, or data from a single institution. None of the studies included in this review validates privacy-preserving federated AML detection using real multi-institutional US banking data. Formal privacy guarantees also apply only to the specific privacy budgets, encryption schemes, aggregation protocols, and threat assumptions evaluated in each study. The preservation of minority-class detection performance therefore remains empirical and dependent on the deployment setting. Table 1 summarizes the available evidence and organizes the studies according to how directly they support privacy-preserving cross-institutional AML detection in the United States.

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. *Cureus J Comput Sci* 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

Tier	Study	Setting and data	Approach	Headline result	Principal caveat
1. Real multi-institutional AML	—	No eligible study identified	—	—	No study using real cross-bank, laundering-labeled data was identified in the literature included in this review
2. Real multi-institutional fraud	Kanamori et al. [27]	Real transaction data from five Japanese banks	Federated learning over homomorphically encrypted gradients under the DeepProtect protocol	Demonstrated privacy-preserving multi-bank fraud detection using real transaction data	Non-US setting; fraud rather than cross-bank money-laundering typologies
3. Synthetic or simulated financial crime	Tang and Liang [37]	Public credit-card fraud datasets partitioned to simulate multiple institutions, including IEEE-CIS	Federated graph learning with graph extension and adaptive model aggregation	Achieved the highest recall and AUC across the evaluated datasets; on IEEE-CIS, recall improved by 11.87%-33.9% and AUC by more than 3% over the compared baselines	Fraud rather than AML; simulated institutional partitioning rather than deployment across real banks
3. Synthetic or simulated financial crime	Abadi et al. [38]	Synthetic transaction data supplied by a global financial-industry partner	Starlit, a privacy-preserving federated-learning system	Implemented a cross-institutional system with an emphasis on scalability	Synthetic data; performance was reported mainly in qualitative terms
3. Synthetic or simulated financial crime	Khan et al. [39]	Realistically generated synthetic data with vertically partitioned features	Fed-RD combining vertical federated learning, differential privacy, and secure multiparty computation	Reported limited performance loss as privacy protection increased	Synthetic data; vertical-partition setting only

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. Cureus J Comput Sci 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

3. Synthetic or simulated financial crime	Arora et al. [40]	Financial-anomaly-detection data from the US-UK Privacy-Enhancing Technologies Challenge	Federated learning with secure multiparty computation and noisy aggregation	Adding cross-institutional information increased AUPRC from approximately 0.6 to 0.7	Challenge setting; restrictive assumptions concerning account-level flags
4. Task-adjacent single-institution fraud or cryptocurrency AML	Egressy et al. [18]	Synthetic AML transaction multigraphs	Graph-neural-network adaptations for directed multigraphs; not federated and not privacy preserving	Increased the minority-class F1-score of a graph isomorphism network baseline from 28.7% to 57.2%	Synthetic data; centralized, single-data-holder setting
4. Task-adjacent single-institution fraud or cryptocurrency AML	Guo et al. [8]	Elliptic Bitcoin transaction graph	Bidirectional transaction graphs with graph convolution and long-term layer attention	Accuracy 0.9776; precision 0.9317; recall 0.8494; F1-score 0.8887; AUC 0.9806	Cryptocurrency with a fully observable ledger; single public graph
4. Task-adjacent single-institution fraud or cryptocurrency AML	Salam et al. [24]	Public credit-card fraud dataset containing 492 fraudulent transactions among 284,807 records	Conventional random forest with hybrid resampling; separate federated convolutional-neural-network experiment	The conventional random forest reported 99.99% accuracy; the federated model reported approximately 90%-93% accuracy across the evaluated settings	The 99.99% result was not federated; class-specific precision, recall, and F1 were not reported for the federated results; extreme class imbalance
4. Task-adjacent single-institution fraud or cryptocurrency AML	Aljunaid et al. [26]	Financial-fraud dataset	Federated learning with SHAP and LIME explainability	Accuracy 99.95%; sensitivity or recall 99.95%; specificity 98.23%; positive predictive value 100%; miss rate 0.05%	Not a real multi-institutional deployment; reported performance depends on the study's class distribution and evaluation design
4. Task-adjacent single-institution fraud or cryptocurrency AML	Khan et al. [41]	Blockchain-recorded transactions	BAML combining Hyperledger Fabric and federated learning	Proposed a decentralized architecture for privacy-preserving financial-compliance checks	No comparison with a trusted-party baseline using real banking data

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. Cureus J Comput Sci 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

5. Foundational methods	Zhang et al. [11]; Bonawitz et al. [16]; Wei et al. [15]; Zhu and Han [13]	General federated-learning survey, secure aggregation, differential privacy, and gradient leakage	Federated-learning, cryptographic, statistical, and attack-analysis foundations	Establish the principal technical building blocks and their limitations	Not specific to financial crime; transfer to AML remains unmeasured
-------------------------	--	---	---	---	---

TABLE 1: Evidence relevant to privacy-preserving cross-institutional financial-crime detection, organized by directness of evidence for US AML

The tiers are ordered according to how directly the evidence supports the target setting: privacy-preserving federated AML detection using real multi-institutional US banking data.

Tier 1 represents the target setting; no eligible study was identified in the literature included in this review.

Tier 2 contains evidence from real multi-institutional data involving a related financial crime.

Tier 3 contains synthetic, simulated, or challenge-based financial-crime studies.

Tier 4 contains task-adjacent studies involving single-institution fraud, simulated collaboration, or cryptocurrency AML.

Tier 5 contains the foundational methods on which the applied systems depend. An em dash indicates that no applicable entry was available.

AML, Anti-Money Laundering; AUC, Area Under the Receiver Operating Characteristic Curve; AUPRC, Area Under the Precision-Recall Curve; LIME, Local Interpretable Model-Agnostic Explanations; SHAP, Shapley Additive Explanations.

US regulatory and data context

Mapping Federation Onto Regulation

US deployment is shaped by three principal regulatory concerns: data confidentiality, permitted information sharing, and model governance. The Gramm-Leach-Bliley Act requires financial institutions to protect customers’ nonpublic personal information and restricts disclosure to nonaffiliated third parties [44]. These obligations make broad pooling of raw transaction histories difficult.

Section 314(b) of the USA PATRIOT Act provides a safe harbor for voluntary information sharing intended to identify and report activity that may involve money laundering or terrorist financing. FinCEN’s 2026 guidance indicates that the safe harbor may cover information relating to suspected fraud, real-time electronic information sharing, and related financial activity, provided participating institutions maintain appropriate security and confidentiality procedures [45]. Suspicious activity reports and information revealing the existence or nonexistence of a suspicious activity report remain subject to separate confidentiality restrictions, except in limited circumstances such as joint filings.

The guidance supports purpose-limited information sharing concerning suspected activity and related financial trails. It does not, however, expressly address whether financial institutions may contribute broad transaction datasets to a general-purpose collaborative training repository. It also does not specifically determine whether exchanged model gradients or parameters fall within the Section 314(b) safe harbor. Any implementation would therefore need to account for other applicable federal and state privacy, cybersecurity, contractual, and supervisory requirements.

The Anti-Money Laundering Act of 2020 encourages coordination, information sharing, technological innovation, and the adoption of new tools in financial-crime compliance [46]. Federated learning may align with these policy objectives by supporting collaborative model development while allowing participating institutions to retain their underlying transaction records locally. This policy alignment does not, by itself, establish legal authorization for a particular federated-learning arrangement.

Model governance creates a separate set of requirements. In April 2026, the Federal Reserve, the Office of the Comptroller of the Currency, and the Federal Deposit Insurance Corporation issued revised model-risk-management guidance that superseded the 2011 framework and the 2021 interagency statement on model risk in Bank Secrecy Act and AML compliance [47]. The revised guidance retains expectations concerning conceptual soundness, outcomes analysis, ongoing monitoring, documentation, and independent challenge while adopting a risk-based and materiality-focused approach. It excludes generative and agentic artificial intelligence from its scope but continues to apply to conventional statistical and machine-learning models.

Federated models may complicate validation because each institution may have limited visibility into the data, model updates, and training influence contributed by other participants. A participating bank would still need to explain model outputs, test performance, monitor changes, document limitations, and support independent challenge over a process that it does not fully control. Explainability methods used in adjacent federated fraud-detection studies offer a partial response [26], but most federated financial-crime studies do not directly address the validation and auditability requirements likely to arise in US supervisory practice.

The Benchmark-Data Problem

The limited availability of representative data remains a major obstacle. The confidentiality concerns that motivate federated learning also make real cross-institutional US banking data difficult for researchers to obtain. The field therefore relies heavily on proxy datasets.

The Elliptic Bitcoin transaction graph is widely used in graph-based AML research [8-10]. Its publicly observable transaction structure supports the evaluation of graph-learning methods, but a cryptocurrency ledger differs substantially from fragmented and privately held fiat-banking networks and raises distinct regulatory considerations [48].

Synthetic datasets provide another option. AMLworld uses an agent-based generator calibrated against real transaction patterns to produce networks containing known laundering typologies and complete labels [49]. SynthAML contains 20,000 alerts and more than 16 million transactions derived from data held by a systemically important Danish bank [50]. Jensen et al. [50] report evidence that models trained on the synthetic data can transfer to real data within the originating institution.

These resources support controlled method development and comparison, but they do not establish performance across heterogeneous banking institutions in the United States. Deprez et al. [51] similarly identify the field's continuing dependence on synthetic and blockchain datasets and the resulting limitations on real-world deployment claims.

Open Gaps

The evidence becomes limited at the point where privacy-preserving federated methods must operate on real cross-institutional AML data. Kanamori et al. [27] demonstrate privacy-preserving federated learning using genuine transaction data from five Japanese banks, but the study concerns financial fraud rather than cross-bank money-laundering typologies.

Most other end-to-end systems included in this review rely on synthetic data, public cryptocurrency records, challenge datasets, simulated institutional partitions, or data from a single institution [37-43]. None of the studies included in this review validates privacy-preserving federated AML detection using real multi-institutional US banking data.

The lack of representative cross-institutional data also prevents reliable estimation of how much AML detection performance is preserved after privacy protections are applied. General studies describe the trade-off between differential privacy and model utility [15,31], but they do not resolve that trade-off in AML settings where the positive class is rare and operational costs depend heavily on false-positive volume.

Two further problems remain. Poisoning and backdoor attacks have been demonstrated against federated graph models relevant to cross-institutional detection [23], but no defense identified in this review has been validated within a consortium of mutually distrustful financial institutions. In addition, the technical literature gives limited attention to

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. Cureus J Comput Sci 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

explainability, model validation, documentation, and independent challenge. Explainable federated fraud-detection studies provide only adjacent evidence [26]. None of the studies included in this review demonstrates a privacy-preserving federated AML system that has been deployed and validated under current US supervisory expectations.

Review Limitations

This narrative review searched OpenAlex and arXiv, and screening and evidence classification were conducted by one reviewer. Exhaustive retrieval and exclusion counts were not retained, and no formal risk-of-bias or study-quality assessment was performed. The evidence tiers in Table 7 indicate the directness of evidence to the target setting rather than methodological quality. Citation expansion and tier assignment involved informed but subjective judgment. The findings should therefore be interpreted as a structured narrative synthesis rather than as the results of a systematic review or meta-analysis.

Conclusions

Privacy-preserving federated artificial intelligence aligns a technical capability with a structural problem. Standard federated learning can support shared model training across institutions without centralizing raw transaction records. However, analysis of cross-bank laundering paths requires additional mechanisms for privacy-preserving entity alignment, representation of interinstitutional edges, and secure graph message passing.

The core components have been demonstrated in experimental, synthetic, and challenge-based settings, but the evidence does not establish effective detection of real cross-institutional money laundering in the United States. Future research should prioritize validation using real multi-institutional data, measurement of privacy-utility trade-offs under AML class imbalance, defenses against poisoning and backdoor attacks, and governance methods that support explanation, validation, monitoring, and independent challenge. Evidence in these areas is needed before broad operational deployment in US AML programs can be supported.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Kehinde Hamed, Moshood Sorinola

Acquisition, analysis, or interpretation of data: Kehinde Hamed, Moshood Sorinola

Drafting of the manuscript: Kehinde Hamed

Critical review of the manuscript for important intellectual content: Kehinde Hamed, Moshood Sorinola

Supervision: Kehinde Hamed

Disclosures

Conflicts of interest: In compliance with [the ICMJE uniform disclosure form](#), all authors declare the following:

- **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work.
- **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work.
- **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. *Cureus J Comput Sci* 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

Data Availability Statements

Data sharing is not applicable. This work is theoretical; no datasets were generated or analyzed.

References

1. Financial Action Task Force: What is money laundering?. (2024). Accessed: July 11, 2026: <https://www.fatf-gafi.org/en/pages/frequently-asked-questions.html#tabs-36503a8663-item-6ff811783c-tab>.
2. United Nations Office on Drugs and Crime: Money laundering. (2024). Accessed: July 11, 2026: <https://www.unodc.org/unodc/en/money-laundering/overview.html>.
3. Chen Z, Khoa LDV, Teoh EN, Nazir A, Karuppiiah EK, Lam KS: Machine learning techniques for anti-money laundering (AML) solutions in suspicious transaction detection: a review. Knowledge and Information Systems. 2018, 57:245-285. [10.1007/s10115-017-1144-z](https://doi.org/10.1007/s10115-017-1144-z)
4. Labanca D, Primerano L, Markland-Montgomery M, Polino M, Carminati M, Zanero S: Amaretto: An Active Learning Framework for Money Laundering Detection. IEEE Access. 2022, 10:41720-41739. [10.1109/access.2022.3167699](https://doi.org/10.1109/access.2022.3167699)
5. Zhang Y, Trubey P: Machine learning and sampling scheme: an empirical study of money laundering detection. Computational Economics. 2019, 54:1043-1063. [10.1007/s10614-018-9864-z](https://doi.org/10.1007/s10614-018-9864-z)
6. Ali A, Abd Razak S, Othman SH, et al.: Financial fraud detection based on machine learning: a systematic literature review. Applied Sciences. 2022, 12:9637. [10.3390/app12199637](https://doi.org/10.3390/app12199637)
7. Li X, Liu S, Li Z, et al.: FlowScope: spotting money laundering based on graphs. Proceedings of the AAAI Conference on Artificial Intelligence. 2020, 34:4731-4738. [10.1609/aaai.v34i04.5906](https://doi.org/10.1609/aaai.v34i04.5906)
8. Guo C, Zhang S, Zhang P, Alkubati M, Song J: LB-GLAT: long-term bi-graph layer attention convolutional network for anti-money laundering in transactional blockchain. Mathematics. 2023, 11:3927. [10.3390/math11183927](https://doi.org/10.3390/math11183927)
9. Lo WW, Kulatilake GK, Sarhan M, Layeghy S, Portmann M: Inspection-L: self-supervised GNN node embeddings for money laundering detection in bitcoin. Applied Intelligence. 2023, 53:19406-19417. [10.1007/s10489-023-04504-9](https://doi.org/10.1007/s10489-023-04504-9)
10. Alarab I, Prakoonwit S: Graph-based LSTM for anti-money laundering: experimenting temporal graph convolutional network with bitcoin data. Neural Processing Letters. 2023, 55:689-707. [10.1007/s11063-022-10904-8](https://doi.org/10.1007/s11063-022-10904-8)
11. Zhang C, Xie Y, Bai H, Yu B, Li W, Gao Y: A survey on federated learning. Knowledge-Based Systems. 2021, 216:106775. [10.1016/j.knosys.2021.106775](https://doi.org/10.1016/j.knosys.2021.106775)
12. Kairouz P, McMahan HB: Advances and open problems in federated learning. Foundations and Trends in Machine Learning. 2021, 14:1-210. [10.1561/22000000083](https://doi.org/10.1561/22000000083)
13. Zhu L, Han S: Deep leakage from gradients. Lecture Notes in Computer Science. Yang Q, Fan L, Yu H (ed): Springer, Cham; 2020. 12500:17-31. [10.1007/978-3-030-63076-8_2](https://doi.org/10.1007/978-3-030-63076-8_2)
14. Nasr M, Shokri R, Houmansadr A: Comprehensive privacy analysis of deep learning: passive and active white-box inference attacks against centralized and federated learning. 2019 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA. 2019, 739-753. [10.1109/SP.2019.00065](https://doi.org/10.1109/SP.2019.00065)
15. Wei K, Li J, Ding M, et al.: Federated learning with differential privacy: algorithms and performance analysis. IEEE Transactions on Information Forensics and Security. 2020, 15:3454-3469. [10.1109/tifs.2020.2988575](https://doi.org/10.1109/tifs.2020.2988575)
16. Bonawitz K, Ivanov V, Kreuter B, et al.: Practical secure aggregation for privacy-preserving machine learning. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. 2017, 1175-1191. [10.1145/3133956.3133982](https://doi.org/10.1145/3133956.3133982)
17. Liu P, Xu X, Wang W: Threats, attacks and defenses to federated learning: issues, taxonomy and perspectives. Cybersecurity. 2022, 5:4. [10.1186/s42400-021-00105-6](https://doi.org/10.1186/s42400-021-00105-6)
18. Egressy B, von Niederhäusern L, Blanuša J, Altman E, Wattenhofer R, Atasu K: Provably powerful graph neural networks for directed multigraphs. Proceedings of the AAAI Conference on Artificial Intelligence. 2024, 38:11838-11846. [10.1609/aaai.v38i10.29069](https://doi.org/10.1609/aaai.v38i10.29069)
19. Cheng D, Zou Y, Xiang S, Jiang C: Graph neural networks for financial fraud detection: a review. Frontiers of Computer Science. 2025, 19:199609. [10.1007/s11704-024-40474-y](https://doi.org/10.1007/s11704-024-40474-y)

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. Cureus J Comput Sci 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

20. Yu L, Zhang F, Ma J, Li Y, Yang Y, Jia W: [Who are the money launderers? Money laundering detection on blockchain via mutual learning-based graph neural network](#). 2023 International Joint Conference on Neural Networks (IJCNN), Gold Coast, Australia. 2023, 1-8. [10.1109/IJCNN54540.2023.10191217](#)
21. Wen J, Zhang Z, Lan Y, Cui Z, Cai J, Zhang W: [A survey on federated learning: challenges and applications](#). International Journal of Machine Learning and Cybernetics. 2023, 14:513-535. [10.1007/s13042-022-01647-y](#)
22. Liu Y, Kang Y, Zou T, et al.: [Vertical federated learning: concepts, advances, and challenges](#). IEEE Transactions on Knowledge and Data Engineering. 2024, 36:3615-3634. [10.1109/tkde.2024.3352628](#)
23. Xu J, Wang R, Koffas S, Liang K, Picek S: [More is better \(mostly\): on the backdoor attacks in federated graph neural networks](#). Proceedings of the 38th Annual Computer Security Applications Conference. 2022, 684-698. [10.1145/3564625.3567999](#)
24. Salam MA, Fouad KM, Elbably DL, Elsayed SM: [Federated learning model for credit card fraud detection with data balancing techniques](#). Neural Computing and Applications. 2024, 36:6231-6256. [10.1007/s00521-023-09410-2](#)
25. Tang Y, Liu Z: [A credit card fraud detection algorithm based on SDT and federated learning](#). IEEE Access. 2024, 12:182547-182560. [10.1109/access.2024.3491175](#)
26. Aljunaid SK, Almheiri SJ, Dawood H, Khan MA: [Secure and transparent banking: explainable AI-driven federated learning model for financial fraud detection](#). Journal of Risk and Financial Management. 2025, 18:179. [10.3390/jrfm18040179](#)
27. Kanamori S, Abe T, Ito T, et al.: [Privacy-preserving federated learning for detecting fraudulent financial transactions in Japanese banks](#). Journal of Information Processing. 2022, 30:789-795. [10.2197/ipsjip.30.789](#)
28. Guembe B, Azeta A, Osamor VC, Ekpo R: [A federated machine learning approaches for anti-money laundering detection](#). Proceedings of the International Conference on Information Systems and Emerging Technologies (ICISSET). SSRN. 2023, 1-13. [10.2139/ssrn.4669561](#)
29. Hu H, Salcic Z, Sun L, Dobbie G, Yu PS, Zhang X: [Membership inference attacks on machine learning: a survey](#). ACM Computing Surveys. 2022, 54:1-37. [10.1145/3523273](#)
30. Dwork C, Roth A: [The algorithmic foundations of differential privacy](#). Foundations and Trends in Theoretical Computer Science. 2014, 9:211-487. [10.1561/04000000042](#)
31. El Ouadrhiri A, Abdelhadi A: [Differential privacy for deep and federated learning: a survey](#). IEEE Access. 2022, 10:22359-22380. [10.1109/access.2022.3151670](#)
32. Fang H, Qian Q: [Privacy preserving machine learning with homomorphic encryption and federated learning](#). Future Internet. 2021, 13:94. [10.3390/fi13040094](#)
33. Ou W, Zeng J, Guo Z, Yan W, Liu D, Fuentes S: [A homomorphic-encryption-based vertical federated learning scheme for risk management](#). Computer Science and Information Systems. 2020, 17:819-834.
34. Hu C, Li B: [MaskCrypt: federated learning with selective homomorphic encryption](#). IEEE Transactions on Dependable and Secure Computing. 2025, 22:221-233. [10.1109/tdsc.2024.3392424](#)
35. Truex S, Baracaldo N, Anwar A, Steinke T, Ludwig H, Zhang R, Zhou Y: [A hybrid approach to privacy-preserving federated learning](#). Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security. 2019, 1-11. [10.1145/3338501.3357370](#)
36. Yin X, Zhu Y, Hu J: [A comprehensive survey of privacy-preserving federated learning: a taxonomy, review, and future directions](#). ACM Computing Surveys. 2021, 54:1-36. [10.1145/3460427](#)
37. Tang Y, Liang Y: [Credit card fraud detection based on federated graph learning](#). Expert Systems with Applications. 2024, 256:124979. [10.1016/j.eswa.2024.124979](#)
38. Abadi A, Naseri M, Doyle B, et al.: [Starlit: privacy-preserving federated learning to enhance financial fraud detection](#). 2025 3rd International Conference on Federated Learning Technologies and Applications (FLTA), Dubrovnik, Croatia. 2025, 126-133. [10.1109/FLTA67013.2025.11336591](#)
39. Khan MSI, Gupta A, Seneviratne O, Patterson S: [Fed-RD: privacy-preserving federated learning for financial crime detection](#). 2024 IEEE Symposium on Computational Intelligence for Financial Engineering and Economics (CIFER), Hoboken, NJ, USA. 2024, 1-9. [10.1109/CIFER62890.2024.10772978](#)
40. Arora S, Beams A, Chatzigiannis P, et al.: [Privacy-preserving financial anomaly detection via federated learning & multi-party computation](#). 2024 Annual Computer Security Applications Conference Workshops (ACSAC Workshops), Honolulu, HI, USA. 2024, 270-279. [10.1109/ACSACW65225.2024.00038](#)

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. Cureus J Comput Sci 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>

41. Khan AA, Alsufyani A, Alsufyani N, Mohamed MA: [BAML: a decentralized approach to secure, privacy-preserving financial compliance for enhancing anti-money laundering with blockchain hyperledger and federated learning](#). Peer-to-Peer Networking and Applications. 2025, 18:270. [10.1007/s12083-025-02086-6](#)
42. Baabdullah T, Alzahrani A, Rawat DB, Liu C: [Efficiency of federated learning and blockchain in preserving privacy and enhancing the performance of credit card fraud detection \(CCFD\) systems](#). Future Internet. 2024, 16:196. [10.3390/fi16060196](#)
43. Nguyen HHX, Dang TK: [A privacy-preserving federated learning model for money laundering detection](#). SN Computer Science. 2025, 6:586. [10.1007/s42979-025-04112-z](#)
44. Public Law 106 - 102 - Gramm-Leach-Bliley Act. (1999). Accessed: July 10, 2026: <https://www.govinfo.gov/app/details/PLAW-106publ102>.
45. Financial Crimes Enforcement Network: Section 314(b) Fact Sheet. (2026). Accessed: July 10, 2026: <https://www.fincen.gov/system/files/shared/314bfactsheet.pdf>.
46. Public Law 116 - 283 - William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021. (2021). Accessed: July 10, 2026: <https://www.govinfo.gov/app/details/PLAW-116publ283>.
47. Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency, Federal Deposit Insurance Corporation: SR 26-2: Revised Guidance on Model Risk Management. (2026). Accessed: July 10, 2026: <https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm>.
48. Fletcher E, Larkin C, Corbet S: [Countering money laundering and terrorist financing: A case for bitcoin regulation](#). Research in International Business and Finance. 2021, 56:101387. [10.1016/j.ribaf.2021.101387](#)
49. Altman E, Blanus J, von Niederhausern L, Egressy B, Anghel A, Atasu K: [Realistic synthetic financial transactions for anti-money laundering models](#). Advances in Neural Information Processing Systems 36 (NeurIPS 2023). 2023, 29851-29874. [10.52202/075280-1300](#)
50. Jensen RIT, Ferwerda J, Jorgensen KS, et al.: [A synthetic data set to benchmark anti-money laundering methods](#). Scientific Data. 2023, 10:661. [10.1038/s41597-023-02569-2](#)
51. Deprez B, Vanderschueren T, Baesens B, Verdonck T, Verbeke W: [Network analytics for anti-money laundering—a systematic literature review and experimental evaluation](#). INFORMS Journal on Data Science. 2025, 5:119-154. [10.1287/ijds.2024.0042](#)

How to cite this article:

Hamed K, Sorinola M (October 12, 2026) Privacy-Preserving Federated Artificial Intelligence for Cross-Institutional Anti-Money Laundering Detection in the United States: A Review. Cureus J Comput Sci 3 : es44389-026-00320-7. DOI <https://doi.org/10.7759/s44389-026-00320-7>