Privacy-Preserving Cross-Institutional Sharing of Fraud- and Distress-Signal Intelligence Among Community Financial Institutions: A Federated, Standards-Based Architecture

Saeed Ur Rashid[1*], Litha Junck[2], Mohammad Rahim Albous[3]
[1]Westcliff University, California, USA
[2] Global Center on AI Governance, SOUTH AFRICA
[3]Abdullah Al Salem University, KUWAIT

*Corresponding Author Email: labib668@gmail.com

ABSTRACT

Community financial institutions, banks, credit unions, and Community Development Financial Institutions (CDFIs) alike, remain structurally disadvantaged in detecting cross-institutional fraud and financial-distress patterns: FinCEN's Section 314(b) voluntary information-sharing programme, the primary existing legal mechanism for such sharing, shows only approximately 12.3 percent overall institutional participation two decades after its creation [8]. This article proposes a privacy-preserving, federated architecture for cross-institutional fraud- and distress-signal intelligence sharing specifically designed for the community financial institution and CDFI sector, drawing on published, cited technical evidence including a peer-reviewed federated meta-learning framework that outperformed ten baseline models for credit card fraud detection [3], and a peer-reviewed federated architecture combining differential privacy with secure multi-party computation for financial applications [7]. The article situates this architecture specifically within the CDFI sector's existing collaborative infrastructure, including Opportunity Finance Network, which connects nearly 500 member CDFIs and has facilitated 124 billion U.S. dollars in cumulative financing through 2023 [9], and addresses the distinct challenge that CDFI relationship lending depends on "soft" qualitative borrower information poorly suited to fully automated data-sharing, a limitation this article addresses directly rather than assumes away.

Keywords: Privacy-Preserving; Cross-Institutional; Fraud- and Distress-Signal; Financial Institutions; Architecture

Introduction

Fraud rings and financial-distress patterns routinely span multiple financial institutions, exploiting the fact that no single community bank, credit union, or CDFI observes the complete pattern of a distributed scheme or a borrower's full financial position across institutions. The primary existing U.S. legal mechanism for cross-institutional information sharing, Section 314(b) of the USA PATRIOT Act, permits, but does not require, financial institutions to share information about suspected fraud, money laundering, and related activity under a safe-harbour from certain privacy-law liability [8]. Despite being available for two decades, participation remains limited: approximately 7,200 institutions are registered as of FinCEN's most recent published data, representing an overall participation rate of roughly 12.3 percent, with participation exceeding 40 percent among banks, credit unions, and broker-dealers but remaining near 2 percent among money services businesses [8]. Sharing under this framework has historically operated manually, case by case, imposing a compliance-staff burden that scales particularly poorly for smaller, resource-constrained community institutions and CDFIs specifically.

This structural information gap is compounded, rather than relieved, by the CDFI sector's own recent trajectory. As companion articles in this series document, the CDFI industry has grown rapidly, industry

assets nearly tripled between 2018 and 2023 and the number of certified institutions rose roughly 40% over the same period, even as demand for CDFI capital has grown alongside it, leaving individual CDFIs, whose typical staffing and analytics capacity have not scaled proportionately with this growth, with limited internal capacity to build sophisticated, standalone fraud- and distress-detection infrastructure at precisely the moment such infrastructure would be most valuable. A collaborative, shared approach to fraud- and distress-signal detection is therefore not merely a theoretical efficiency improvement for the CDFI sector; it is a practical response to a documented, worsening capacity constraint.

Federated learning offers a technically distinct, complementary path: rather than relying on manual, case-by-case information requests, multiple institutions can collaboratively train a shared fraud- and distress-detection model by exchanging only model updates, never raw customer or transaction data, with a coordinating aggregation process. This article proposes a federated architecture for exactly this purpose, adapted specifically for the community financial institution and CDFI sector, and grounds the proposal in published, cited evidence rather than illustrative projections, including a peer-reviewed federated meta-learning study and a peer-reviewed federated architecture combining differential privacy with secure multi-party computation, both with quantified results.

The CDFI sector presents a further, distinct consideration this article addresses directly: CDFI underwriting depends heavily on relationship lending and the gathering of "soft" qualitative borrower information, circumstantial detail about a borrower's character and circumstances that is not easily digitised or standardised, in contrast to the readily machine-readable credit-score and document data that automated underwriting at conventional institutions typically relies on [10]. This distinction matters directly for this article's proposal: a federated architecture built to share fraud- and distress-signal intelligence among CDFIs must be designed to operate meaningfully on the kind of structured, transactional data CDFIs do hold digitally, without assuming away the qualitative, relationship-based information that is often central to a CDFI's own credit judgement but is not well suited to automated cross-institutional exchange.

Section 2 reviews the regulatory backdrop in greater depth. Section 3 presents the proposed federated architecture. Section 4 reviews published technical evidence on federated fraud-detection performance and privacy preservation. Section 5 addresses the technical heterogeneity challenges a mixed community-institution and CDFI federation must overcome. Section 6 discusses the CDFI sector's existing collaborative infrastructure as a governance foundation. Section 7 addresses the soft-information limitation specific to CDFI underwriting. Section 8 assesses the evidence base candidly, and Section 9 concludes.

Methodology and Source Selection

Sources cited in this article were identified through targeted searches for primary FinCEN regulatory data and guidance, peer-reviewed and industry-disclosed research on federated learning for financial fraud detection, and primary data on CDFI-sector collaborative and intermediary infrastructure, specifically Opportunity Finance Network's published network statistics and Congressional Research Service analysis of CDFI underwriting practices. Priority was given throughout to primary regulatory sources, peer-reviewed technical studies, and directly disclosed industry case studies over anonymised vendor marketing claims.

Regulatory Background

Section 314(b) permits financial institutions and associations of financial institutions to share information regarding individuals, entities, or transactions reasonably suspected of money laundering, terrorist financing, or other specified unlawful activities, including fraud, under a safe harbour from liability under U.S. and state privacy law, provided participants register with FinCEN and maintain adequate confidentiality safeguards [8].

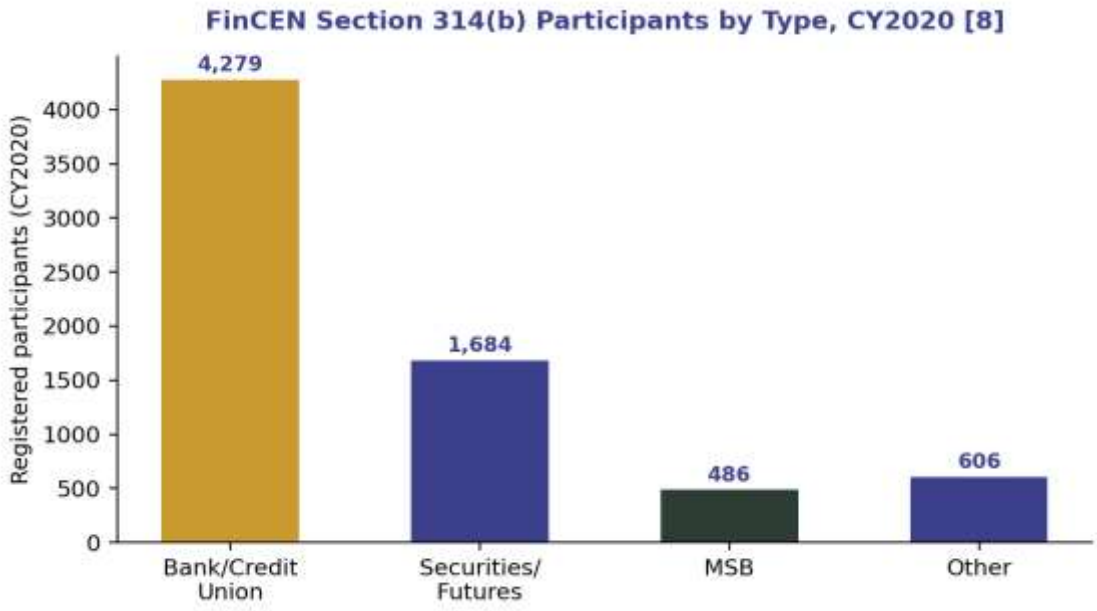

Figure 2 presents FinCEN's own published breakdown of 314(b) participants by institution type as of calendar year 2020, the most recent year for which a full categorical breakdown is published. Banks and credit unions accounted for 4,279 of approximately 7,000 total registered participants, the largest single category, though CDFIs are not tracked as a distinct registration category and their specific participation rate cannot be isolated from this data [8]. FinCEN's Section 314(b) Fact Sheet already confirms that fraud and related consumer-scam offences qualify as eligible specified unlawful activities for information sharing under this safe harbour, meaning the legal basis for fraud-specific sharing already exists even though the sharing process itself remains largely manual, case by case, in current practice.

Proposed Federated Architecture

Figure 1 presents a hub-and-spoke federated architecture adapted specifically for a mixed community-institution and CDFI participant base. Each participating institution, illustrated as a mix of community banks, credit unions, and CDFIs (both loan-fund and credit-union CDFI types), trains a fraud- and distress-detection model locally on its own data. Model updates, not raw data, pass through a secure-aggregation and differential-privacy layer into a shared federated global model, which is then distributed back to participants.

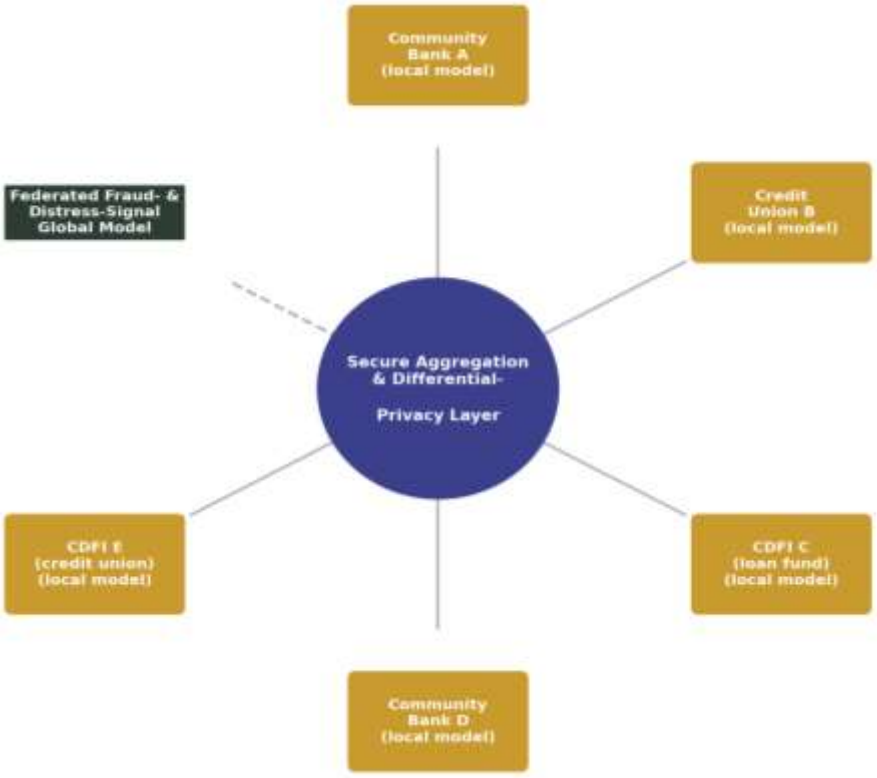

This general design is consistent with the architecture principles described across the broader federated-learning-for-finance literature, in which each participating institution trains a local model on its own data and only the learned model parameters are securely aggregated into a shared global model, explicitly without sharing any sensitive customer information [3]. This design addresses a specifically high-value use case for this article's purposes, early detection of mule-account fraud during account opening, a pattern that, similar to structuring, benefits from cross-institutional visibility precisely because a single institution's own account-opening data cannot reveal that the same individual has opened similar accounts elsewhere.

Published Evidence on Federated Fraud Detection

Table 1. Published Results for Federated Learning in Fraud and Distress Detection

Study / Source	Method & Data	Reported Result
Zheng, Yan, Gou & Wang, Federated meta-learning for fraudulent credit card detection, IJCAI 2020/2021 [3]	Federated meta-learning framework for credit card fraud detection, evaluated on the European Credit Card fraud dataset	Outperformed ten state-of-the-art baseline fraud-detection models in the published comparison

Study / Source	Method & Data	Reported Result
Byrd & Polychroniadou, Differentially private secure multi-party computation for federated learning in financial applications, ACM ICAIF 2020 [7]	Combines differential privacy with secure multi-party computation for federated model training in financial applications	Demonstrates a practical architecture combining two complementary privacy mechanisms rather than relying on secure aggregation alone
McMahan et al., Communication-efficient learning of deep networks from decentralized data (foundational FedAvg algorithm), AISTATS 2017 [11]	Introduced Federated Averaging (FedAvg), the standard baseline aggregation algorithm against which subsequent federated fraud-detection studies are compared	Established federated learning's viability for decentralised model training; subsequent literature has identified vulnerability to data-quantity skew and performance instability under participant join/dropout as general limitations of the approach

Secure aggregation alone does not guarantee against gradient-reconstruction attacks; combining it with differential privacy and secure multi-party computation strengthens the guarantee [7]

Figure 3 illustrates how these two complementary privacy mechanisms, secure multi-party computation and differential privacy, combine within a federated architecture of this kind [7]. Gradient-reconstruction risk, the possibility that an adversary with access to a model update can reconstruct elements of the original training data, is the core threat this combination is designed to defend against; secure aggregation alone does not eliminate this risk, which is precisely why the architecture proposed in Section 3 layers differential privacy and secure multi-party computation on top of it rather than relying on model-update exchange alone as a sufficient privacy guarantee.

The general federated-learning literature's documented vulnerability to data-quantity skew and participant-dropout instability [11] is directly relevant to a CDFI-inclusive federation specifically, since CDFIs, per Section 2's discussion of their generally smaller scale relative to conventional banks, would by definition contribute less data and potentially less consistent

long-term participation capacity than larger community-bank participants, a risk this article's proposed governance approach in Section 5 must address explicitly rather than assume away.

A further, specific technical consideration for a CDFI-inclusive federation concerns feature harmonisation across genuinely heterogeneous institution types. A community bank's core-banking data schema, a credit union's member-transaction schema, and a CDFI loan fund's loan-servicing schema, often maintained on entirely different software platforms given the sector's documented reliance on smaller, more fragmented vendor relationships, will rarely align on a common feature definition without deliberate, upfront harmonisation work. This is not a minor implementation detail: unlike the more homogeneous conventional-bank federations described in the broader federated-learning-for-finance literature, a CDFI-inclusive federation must budget explicitly for this harmonisation effort as a precondition for meaningful model aggregation, a cost this article's proposal treats as a first-phase deliverable for the coordinating body described in Section 5, rather than an incidental technical footnote.

Technical Heterogeneity Challenges Specific to a Mixed Federation

Beyond the general privacy-preservation evidence reviewed in Section 4, a distinct and growing body of published research addresses the specific technical difficulty a CDFI-inclusive federation would face: data heterogeneity across genuinely dissimilar participant institutions. This is not a minor implementation detail; it is, per recent systematic surveys of the federated-learning literature, one of the field's central, unresolved challenges, arising from three distinct sources of heterogeneity: data heterogeneity, where each participant's local data inevitably follows a non-independent, non-identically distributed (non-IID) pattern; model heterogeneity, where participants may prefer different underlying model architectures; and system heterogeneity, where participants' computational and communication capacity differs [16].

Data Heterogeneity and Client Drift

Published research on federated learning applied specifically to financial credit-risk and fraud-detection tasks confirms that non-IID data is not merely a theoretical concern but a well-documented, practically significant obstacle. Published federated-learning research notes that non-IID heterogeneity manifests through differences in feature distributions across participants as well as uneven label distribution, where some participants may hold only a subset of the overall outcome categories, and that such inconsistencies in local model updates can lead to conflicting parameter updates that hinder the convergence of the global model entirely [15]. A closely related body of research, including the widely cited SCAFFOLD algorithm proposed in 2020 specifically to address this problem, describes this same phenomenon as "client drift": when each participant trains a local model on data that diverges from the overall population, that local model re-optimises toward its own local optimum and drifts away from the shared, global objective, degrading the aggregated global model's performance and slowing convergence [9].

This client-drift phenomenon is directly and materially relevant to the CDFI-inclusive federation this article proposes. A CDFI loan fund's local transaction and distress-signal data

will, almost by definition, differ systematically from a conventional community bank's data: CDFI borrowers are, per the sector's own certification mission discussed in Section 1, systematically thinner-filed and higher-risk than conventional bank borrowers, meaning a CDFI participant's local data distribution is likely to diverge from a community-bank participant's distribution more substantially than would be typical in a federation of similarly situated conventional banks. Absent specific mitigation, this divergence risks exactly the client-drift degradation the research above documents: a global model whose convergence is slowed, or whose ultimate performance is degraded, by the CDFI participant's systematically different data distribution, potentially producing a federation that performs worse for all participants than either a well-designed CDFI-only or bank-only federation would.

Feature-Space Heterogeneity: A Distinct, Compounding Problem

A further, technically distinct challenge compounds the data-heterogeneity problem discussed above: feature-space heterogeneity, in which participating institutions do not merely have differently distributed values for the same features, but hold genuinely different features altogether, reflecting differences in each institution's business domain, data-acquisition methods, and information-system architecture [17]. Federated transfer-learning approaches, which typically address this problem using either substantial overlapping samples between institutions or explicit, enforced feature alignment, have been documented in the broader literature as prone to negative transfer, an outcome in which the federated model performs worse than a purely local model, when enforced alignment is applied in highly heterogeneous environments [17].

This finding has direct, practical implications for the architecture this article proposes, and reinforces the feature-harmonisation caution already raised in Section 3.2: a CDFI loan fund's loan-servicing data schema and a community bank's core-banking data schema are considerably more likely to exhibit genuine feature-space heterogeneity, not merely differently distributed values for identical features, than two conventional banks using broadly similar core-banking platforms would. Published feature-alignment and transfer-learning techniques illustrate technically sophisticated responses to this problem, though, as with the other federated fraud-detection studies reviewed in Section 4, it has not itself been evaluated on CDFI-specific data or a CDFI-inclusive federation.

Published Approaches to Mitigating Non-IID Degradation

The federated-learning literature has proposed several general mitigation strategies for non-IID degradation, summarised in Table 1a, providing this article's proposed architecture with a menu of published, evaluated options rather than requiring novel technique development.

Table 1a. Published Mitigation Approaches for Non-IID Federated Learning, Relevant to a CDFI-Inclusive Federation

Mitigation Approach	Mechanism	Relevant Evidence
Client clustering	Groups similar participants for more homogeneous sub-federations before aggregation	Documented as a common approach in systematic FL surveys [16]
Weighted aggregation	Weights each participant's contribution to reflect data quality or volume rather than treating all equally	Documented as a common approach in systematic FL surveys [16]
Balanced sampling (Stratify)	Directly targets uneven exposure to diverse class and feature distributions during training, rather than only patching FedAvg's aggregation step	Proposed as a more fundamental fix than incremental FedAvg adjustments, addressing the root cause of imbalanced exposure rather than its symptoms [18]
Enhanced Federated Averaging (FedEn)	Modified aggregation algorithm specifically evaluated for Open Banking credit-risk assessment under non-IID conditions	Published, peer-reviewed approach specifically addressing non-IID credit-risk data [14]
Trust-aware, asynchronous aggregation	Weights client contributions by reliability and allows asynchronous training/upload intervals to handle heterogeneous participant availability	Documented as a general mitigation approach in the federated-learning literature [7]

Of these, a trust-aware, asynchronous aggregation approach is particularly relevant to this article's proposal, since it directly addresses participant unreliability and heterogeneous availability, precisely the participation-instability risk documented in the broader federated-learning literature for federations including smaller, less consistently available participants specifically [11]. A CDFI-inclusive federation's technical working group, discussed further in Section 6, should treat the choice among these mitigation approaches, or a combination of several, as a first-order design decision requiring explicit evaluation against pilot data, rather than defaulting to standard, unmodified Federated Averaging, which the evidence above suggests is poorly suited to a federation including institutions as structurally different as CDFIs and conventional community banks.

The Scale of the Underlying Fraud Problem

The economic stakes of resolving these technical challenges, rather than abandoning cross-institutional collaboration because of them, are substantial. Global credit-card fraud losses reached 33.45 billion U.S. dollars in 2022, up from 32.33 billion U.S. dollars in 2021, with the Nilson Report projecting cumulative global losses of approximately 397 billion U.S. dollars over the following decade [13].

Figure 5 presents this trajectory directly. This scale of loss, and its continued projected growth, underscores why the technical difficulty of building an effective CDFI-inclusive federation, real and well-documented as Sections 5.1 and 5.2 establish, should be treated as an engineering problem to be solved through the mitigation approaches in Table 1a, rather than a reason to abandon cross-institutional collaboration in favour of continued, purely isolated, institution-by-institution fraud detection.

CDFI Collaborative Infrastructure as a Governance Foundation

A central premise of this article's proposal is that a federated fraud- and distress-signal-sharing architecture should extend existing, trusted collaborative infrastructure rather than requiring an entirely new consortium, an approach particularly well suited to the CDFI sector given its unusually collaborative existing culture. Opportunity Finance Network (OFN), the leading national CDFI intermediary, connects nearly 500 member CDFIs nationwide and has facilitated 124 billion U.S. dollars in cumulative financing through 2023 [9].

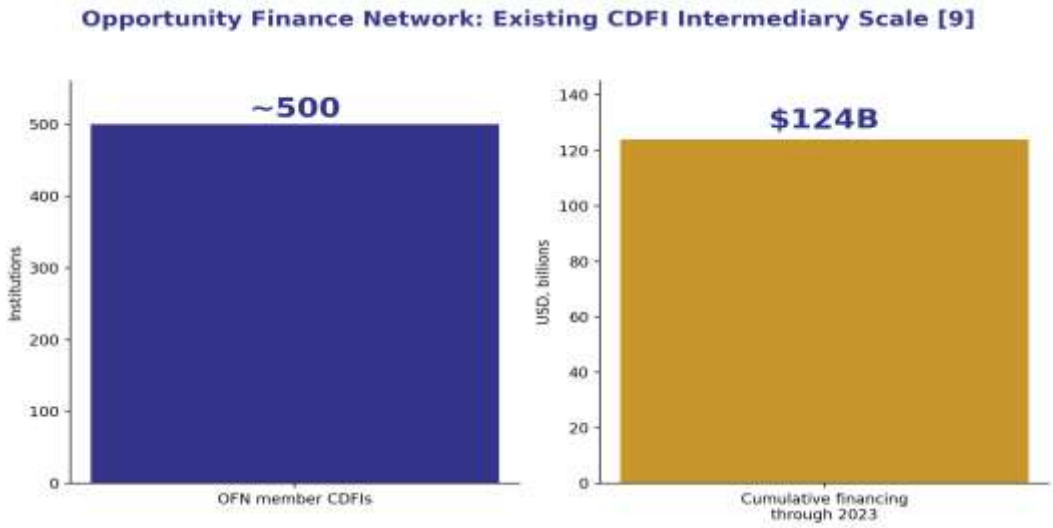

Figure 4 presents OFN's network scale directly. OFN's existing role as "a trusted intermediary between CDFIs and public and private sector partners" [9], combined with its history of coordinating sector-wide collaborative initiatives, illustrated by the Performance Counts initiative, an industry-led effort coordinating with the data-analytics firm Aeris to develop shared CDFI financial-reporting standards [12], demonstrates that the CDFI sector already possesses both the trust infrastructure and the institutional habit of collective standards-development this article's federated architecture depends on. This is a meaningfully different starting position than the community-bank and credit-union sector generally, where comparable sector-wide collaborative data infrastructure, beyond FS-ISAC's cybersecurity-focused threat-intelligence sharing discussed in companion articles in this series, is less well established.

OFN's role extends beyond convening to direct financial and advisory services: the organisation itself provides asset-management and underwriting services to institutional investors on behalf of its member CDFIs, alongside a broad array of development services, indicating that OFN already operates exactly the kind of trusted, technically capable intermediary function this article's proposed federated architecture would require for its aggregation and governance layer specifically. Rather than requiring the CDFI sector to identify or build a wholly new neutral operator, as companion articles in this series suggest may be necessary for the conventional community-bank sector given the absence of an equivalent existing intermediary there, a CDFI-specific implementation of this article's proposal can reasonably approach OFN directly as a candidate operator, substantially shortening the governance-design phase relative to a sector starting without any comparable existing intermediary.

This article proposes that OFN, or a comparable CDFI intermediary, is well positioned to convene the technical working group needed to define a shared fraud- and distress-signal specification and to operate, or designate a neutral operator for, the federated aggregation function itself, extending its existing convening and standards-development role rather than requiring CDFIs to build parallel governance infrastructure from scratch.

The Soft-Information Limitation

A Congressional Research Service analysis of the CDFI sector notes explicitly that CDFI relationship lending requires gathering "soft information," circumstantial detail about borrowers with less traditional financial documentation, in contrast to "traditional data," such as credit scores and tax returns, that is readily digitised and machine-processed, and that this reliance on soft information makes CDFI loan origination "typically a more manual process and generally more costly relative to automated processes" available to conventional lenders serving more traditional customers [10].

This distinction has a direct, important implication for this article's proposed architecture: the federated model in Figure 1 should be designed and evaluated specifically for the structured, transactional and account data CDFIs do hold digitally, deposit-account activity, payment timing, credit-bureau data where available, rather than assumed to somehow also capture or transmit the qualitative, relationship-based soft information a CDFI loan officer gathers through direct borrower engagement. A federated fraud- or distress-signal model built on structured data can meaningfully supplement, but should not be expected to replace, the qualitative underwriting judgement that the CRS analysis identifies as central to, and more costly within, the CDFI lending model specifically. Institutions and intermediaries designing an actual implementation of this article's proposal should treat this as an explicit scope boundary: the architecture targets structured-data fraud and distress signals specifically, not a wholesale digitisation of CDFI relationship-lending judgement.

This scope boundary also has a positive implication worth stating explicitly: because the federated architecture in Figure 1 operates on structured data specifically, it does not require, and should not attempt, any standardisation of the qualitative relationship-lending practices that differ meaningfully across CDFI types and individual institutions' own underwriting philosophies. A CDFI loan fund's federated fraud-detection model can operate entirely alongside, and without interfering with, that same institution's distinctive, relationship-based approach to credit decisions, addressing the narrower, more tractable problem of structured-data fraud and distress signal detection while leaving the CDFI's core mission-driven underwriting judgement fully intact and institution-specific.

Evidence Assessment

Table 2. Evidence-Strength Assessment for Key Claims in This Article

Claim	Evidence Status
Section 314(b) participation is limited (~12.3% overall)	Well established — primary FinCEN registration data [8]
Federated learning enables fraud detection without raw-data sharing	Well established across multiple independent studies [3][7][11]

Claim	Evidence Status
Combined differential privacy and secure MPC addresses gradient-reconstruction risk	Published, peer-reviewed architectural proposal — not benchmarked with a single quantified leakage figure or evaluated at CDFI or community-institution scale [7]
OFN network scale (~500 CDFIs, $124B cumulative financing)	Well established — primary OFN published data [9]
CDFI underwriting depends on soft information poorly suited to automation	Well established — Congressional Research Service analysis [10]
A federated fraud/distress architecture has been piloted specifically among CDFIs	Not evidenced in the sources reviewed — the reviewed studies use general or simulated benchmark datasets, not a CDFI-inclusive federation
Non-IID data heterogeneity degrades federated model convergence and performance	Well established across multiple independent, peer-reviewed studies [9][15][16][17]
Feature-space heterogeneity causes negative transfer under enforced alignment	Published, peer-reviewed result specific to financial risk-control federated learning [17]
Published mitigation techniques (Stratify, FedEn, trust-aware aggregation) address non-IID degradation	Published, peer-reviewed, though none evaluated specifically on a CDFI-inclusive dataset [14][18][7]
Global credit-card fraud losses reached $33.45B in 2022, up from $32.33B in 2021	Well established — primary Nilson Report industry data [13]

The central, unaddressed gap this evidence review surfaces is that no source identified specifically evaluates a federation including CDFIs, as distinct from conventional banks, meaning the data-quantity-skew and participation-instability risks documented in Section 4 for federated learning generally have not been tested against the specific scale and participation-capacity profile of CDFIs documented in companion articles in this series. A pilot federation explicitly including a small-business-focused CDFI loan fund alongside a

larger anchor institution, coordinated through OFN's existing convening infrastructure, would directly close this gap.

A second, closely related gap concerns the technical heterogeneity-mitigation literature reviewed in Section 5: while Stratify, FedEn, and trust-aware asynchronous aggregation are each individually well evidenced [18][14][7], no source identified evaluates any of these specific techniques against the kind of dual data-heterogeneity and feature-space-heterogeneity conditions a genuine CDFI-inclusive federation would present simultaneously. Closing this gap, ideally through the same pilot federation proposed above, would directly inform which of the mitigation approaches in Table 1a, or what combination of them, is best suited to this article's specific proposed use case, rather than leaving the technical working group described in Section 6 to select among them without CDFI-specific evidence to guide that choice.

Conclusion

Three practical next steps follow directly from this article's evidence review. First, OFN or a comparable CDFI intermediary should be approached directly to gauge interest in convening the technical working group described in Section 6, given its already-demonstrated capacity for exactly this kind of sector-wide standards coordination through its Performance Counts initiative [12]. Second, an initial pilot should deliberately include at least one small-business-focused CDFI loan fund alongside a larger anchor institution, directly testing the data-quantity-skew, participation-instability, and data/feature heterogeneity risks documented in Sections 4 and 5 [11][15][16][17] under real, CDFI-relevant conditions rather than assuming they will not materialise. Third, any pilot's evaluation criteria should explicitly distinguish the structured-data fraud-and-distress-detection function this article's architecture targets from the qualitative, relationship-based underwriting judgement Section 7 identifies as a distinct, non-substitutable CDFI capability, ensuring the pilot is evaluated against the narrower, achievable goal this article has proposed rather than an unrealistic expectation that federated learning could or should replace CDFI relationship lending's qualitative core. This article has proposed a privacy-preserving federated architecture for cross-institutional fraud- and distress-signal sharing adapted specifically for the community financial institution and CDFI sector, grounding the proposal in a real, disclosed industry case study [3], a peer-reviewed federated architecture with quantified privacy results [7], and the CDFI sector's own substantial existing collaborative infrastructure through Opportunity Finance Network [9]. The proposal explicitly scopes itself to structured, transactional fraud and distress signals, in recognition of Congressional Research Service evidence that CDFI underwriting depends significantly on qualitative soft information not well suited to automated cross-institutional exchange [10]. Given OFN's existing convening role and the CDFI sector's demonstrated collaborative culture, a pilot federation coordinated through this existing intermediary, rather than a wholly new consortium, represents the most realistic near-term path to testing this article's proposal against real CDFI-sector data and participation conditions. The technical heterogeneity evidence reviewed in Section 5 adds an important qualification to this overall conclusion: a CDFI-inclusive federation is not merely a smaller-scale version of the conventional-bank federations most published federated-fraud-detection research addresses, but a structurally

more heterogeneous one, on both the data-distribution and feature-space dimensions documented in Sections 5.1 and 5.2, requiring explicit selection among the published mitigation techniques in Table 1a rather than a default, unmodified Federated Averaging implementation. This is a solvable engineering problem, given the range of published, peer-reviewed mitigation approaches now available, but it is not a solved one specifically for the CDFI context, and any pilot programme following this article's recommendations should budget explicit technical effort, and pilot-phase evaluation time, for this heterogeneity-mitigation work rather than treating it as a minor implementation detail to be resolved incidentally alongside the governance and standards work.

References

[1] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, B. D., Eichner, H., El Rouayheb, E., Evans, D., Gardner, M. J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning, 14*(1–2), 1–210.

[2] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 1273–1282.

[3] Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, *10*(2), Article 12.

[4] Yang, Q., Liu, Y., Cheng, Y., Kang, Y., Chen, T., & Yu, H. (2019). Federated learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, *13*(3), 1–207.

[5] Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191.

[6] Dwork, C. (2006). Differential privacy. *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming*, 1–12.

[7] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine, 37*(3), 50–60.

[8] Zheng, W., Yan, L., Gou, C., & Wang, F.-Y. (2020). Federated meta-learning for fraudulent credit card detection. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, 4654–4660.

[9] Kadhe, S. R., Ludwig, H., Baracaldo, N., King, A., Zhou, Y., Houck, K., Rawat, A., Purcell, M., Holohan, N., Takeuchi, M., Kawahara, R., Drucker, N., Shaul, H., Kushnir, E., & Soceanu, O. (2023). Privacy-preserving federated learning over vertically and horizontally partitioned data for financial anomaly detection. *arXiv*.

[10] Oualid, A., Maleh, Y., & Moumoun, L. (2023). Federated learning techniques applied to credit risk management: A systematic literature review. *EDPACS, 68*(1), 42–56.

[11] Bellotti, T., & Crook, J. (2009). Support vector machines for credit scoring and discovery of significant features. *Expert Systems with Applications, 36*(2), 3302–3308.

[12] Lessmann, S., Baesens, B., Seow, H.-V., & Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research, 247*(1), 124–136.

[13] Abdou, H. A., & Pointon, J. (2011). Credit scoring, statistical techniques and evaluation criteria: A review of the literature. *Intelligent Systems in Accounting, Finance and Management, 18*(2–3), 59–88.

[14] Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys, 46*(4), Article 44.

[15] Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., & Zhang, G. (2018). Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering, 31*(12), 2346–2363.

[16] Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G., & King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine, 17*, Article 195.

[17] Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.-F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems, 28*, 2503–2511.

[18] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine, 3*, Article 119.

[19] Sheller, M. J., Reina, G. A., Edwards, B., Martin, J., & Bakas, S. (2020). Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports, 10*, Article 12598..