ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1100

FedSCORE-PP: A Federated and Privacy-Preserving Machine Learning Framework for Collaborative Supply Chain Risk Prediction Across Organizations

Sohail Sayed[1]; Nauman Sayed[2]

[1,2]California State University Fullerton

Publication Date: 2026/09/01

Abstract: Global supply chains are increasingly exposed to disruptions whose effects propagate across organizational boundaries, yet the data needed to predict such risks is fragmented among firms reluctant to share it for competitive, contractual, and regulatory reasons. Centralized machine learning therefore under-utilizes collective evidence, and organizations with inadequate datasets cannot predict risk reliably on their own [1]. This paper proposes FedSCORE-PP, a federated and privacy-preserving framework that enables competing organizations to jointly train a supply chain risk prediction model without exchanging raw data. The framework layers (I) local differential privacy via DP-SGD with Rényi-DP accounting, giving formal user-level (ϵ, δ) guarantees; (II) homomorphic-encryption/secure-multi-party-computation-based secure aggregation, so the server never observes plaintext updates [2]; and (III) an optional blockchain audit ledger for integrity and incentives [3]. The privacy budget is derived as an explicit closed-form function of noise multiplier, client sampling, and communication rounds, yielding $\epsilon \in [2.2, 40.3]$ at $\delta = 10^{-5}$ across the design space. Communication analysis shows a total cost of ≈50.2 MB over a 100-round campaign for a reference architecture. Benchmark comparison with published federated SCRM systems (reported accuracy 90.3–95.0%) shows the framework occupies a competitive operating region while providing stronger privacy guarantees [4], [5], [6], [7].

Keywords: *Federated Learning, Differential Privacy, Homomorphic Encryption, Secure Aggregation, Supply Chain Risk Management, Supply Chain Resilience, Collaborative AI.*

How to Cite: Sohail Sayed; Nauman Sayed (2026) FedSCORE-PP: A Federated and Privacy-Preserving Machine Learning Framework for Collaborative Supply Chain Risk Prediction Across Organizations. *International Journal of Innovative Science and Research Technology*, 11(8), 2304-2312. https://doi.org/10.38124/ijisrt/26aug1100

I. INTRODUCTION

Decades of cost-driven optimization produced lean, globally distributed supply networks that systemic shocks — the COVID-19 pandemic, geopolitical instability, and environmental disruptions — have repeatedly exposed as fragile [8]. These events forced a shift from reactive risk mitigation toward proactive strategic resilience, in which organizations anticipate disruptions rather than merely respond [9]. Artificial intelligence is central to this shift: AI-driven approaches support risk identification, assessment, and mitigation across upstream and downstream supply chains [10], with industry evidence suggesting AI improves risk-prediction accuracy by 20–50% and response times by 30–40% [11]. Machine learning has predicted first-tier supplier disruptions with 80% accuracy on late orders [12], and ensemble methods outperform classical statistical models for delivery-delay prediction [13].

Yet most approaches rest on a premise that contradicts the interconnected nature of supply chains: that organizations predict risk alone. Because disruptions propagate through supplier–buyer networks, the most informative signals often reside in a partner's data; but firms withhold it due to privacy, confidentiality, and competitive-disadvantage concerns [1]. SMEs are doubly disadvantaged: they are most concerned about data privacy and most likely to hold inadequate data for effective risk management [14]. The result is a data-island problem: organizations with inadequate datasets cannot predict risk [1].

Federated learning resolves this: a shared model is trained collaboratively across decentralized data holders that exchange model updates — never raw data [15]. Supply chain studies confirm its promise: FL helps buyers with limited datasets predict order delays effectively [1]; it outperforms both centralized and local standalone learning for delivery-delay prediction across textile suppliers [16]; a federated logistics risk-warning system achieves 91.3% average

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1100

accuracy while complying with the GDPR [4]; and FL-based forecasting with supply-chain visibility matches full-visibility performance while protecting partner privacy [17].

However, vanilla FL is not itself a privacy guarantee: gradients and updates can leak training information through inference and inversion attacks [18], [19]. This paper therefore proposes FedSCORE-PP, contributing:

- A layered privacy architecture — local DP-SGD with formal user-level (ϵ, δ) guarantees [20]; HE/SMC-based secure aggregation ensuring the server never sees individual plaintext updates [2], [21]; and an optional blockchain audit ledger for integrity and traceability [3], [22].
- A transparent privacy-accounting model — the end-to-end privacy budget as an explicit closed-form function of noise multiplier, sampling ratio, and communication rounds, based on Rényi-DP composition [23], with fully worked numerical tables.
- Communication and computation complexity analyses with concrete derived figures for a reference neural architecture and packing/selective-encryption optimizations [2], [24], [25].
- A benchmark-anchored results section comparing against published federated SCRM systems, plus an analysis of when collaboration is beneficial [1], [16].

II. BACKGROUND AND RELATED WORK

- *AI for Supply Chain Risk Management*

SCRM spans identification, assessment, and mitigation of unexpected events that adversely affect any part of a supply chain, and its data-intensive, decision-oriented character makes it a natural AI application [10]. Machine learning has been applied to predict delivery delays in multi-tier manufacturing, though interpretability often trades off against performance [26]; to predict first-tier supplier disruptions with engineered features (80% accuracy on late orders) [12]; and to predict delivery delays with ensembles outperforming classical statistical models [13]. Conceptual frameworks position AI as an enabler of resilience through predictive analytics and real-time visibility [27]. These approaches presume sufficiently rich organizational datasets — a presumption that fails precisely for the SMEs that need support most [14].

- *Federated Learning*

FL trains a shared model across decentralized clients that keep data local, exchanging only parameters or gradients [15]. It supports horizontal, vertical, and transfer configurations, and has been applied in privacy-sensitive domains such as healthcare, banking, and IoT [28], [29]. Its canonical algorithm, Federated Averaging, aggregates locally trained models weighted by dataset size [4]. A known weakness is that non-IID data across clients significantly slows convergence [29].

- *Privacy-Enhancing Technologies for Collaborative Learning*

Three families dominate: differential privacy, which adds calibrated noise so the released model's output is approximately unchanged whether or not any single client's data is present [20]; homomorphic encryption, which permits arithmetic on ciphertexts so a server can aggregate encrypted updates without decryption [21], [30]; and secure multi-party computation-based secure aggregation, which ensures the server learns only the sum [29]. HE-based privacy-preserving ML incurs accuracy loss below 1% in practice [21]. Because HE is computationally expensive, packing and selective-encryption optimizations have become essential: BatchAgg achieves a 60× training speedup and 57% communication reduction versus Paillier at equal security [2]; BatchCrypt achieves 23×–93× training speedup and 66×–101× communication reduction [25]; FedML-HE reduces HE overhead by ~10× for ResNet-50 and up to ~40× for BERT [24]; and selective-HE frameworks reduce overhead by nearly an order of magnitude while keeping accuracy above 90% [31]. Blockchain-assisted FL adds decentralization, traceability, and tamper-evidence to update storage [3], [22].

- *Federated Learning in Supply Chains*

Ge et al. established that FL enables collective order-delay prediction without data exposure, showing that buyers with limited datasets benefit most while buyers with disproportionately large order books may not benefit at all [1]. Nguyen et al. confirmed FL's advantage for data-limited suppliers and demonstrated FL outperforming both centralized and local standalone training for delivery-delay prediction [16]. SafeLogFL achieved 91.3% average accuracy for cross-border logistics risk warning with GDPR compliance [4]. FSL-Qwen, a federated split-learning framework with an LLM, reached 94.0% accuracy and 94.1% F1 for container source identification in sea-rail intermodal transport, matching the centralized upper bound [6]. A privacy-aware FL backorder system reached ~90.3% accuracy, outperforming centralized ML [5]. FL demand forecasting with visibility matched full-visibility performance [17]. A federated NLP approach to disaster-risk prediction achieved 0.95 accuracy [7]. Blockchain–FL has been applied to food supply chain security [32], to decentralized supply chains with DP and IPFS-based storage [33], and to predictive maintenance in cross-border logistics, improving accuracy by 6.9% over FedAvg and 3.7% over FedProx while increasing privacy protection by over 12% [3]. Cold-chain risk quantification has integrated FL with multi-criteria decision-making [14], and privacy-preserving mechanisms for decentralized supply chain platforms continue to mature [34].

- *Research Gaps*

Existing supply chain FL systems typically apply FedAvg with DP as an afterthought, without formal end-to-end privacy accounting [20]; few combine DP and secure aggregation so that both statistical and computational privacy hold simultaneously [2], [21]; and non-IID, multi-echelon cross-company data remains an acknowledged difficulty [1], [29]. FedSCORE-PP addresses these gaps within a single

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1100

SCRM-oriented framework with explicit, computed privacy and complexity budgets.

III. PROBLEM FORMULATION AND THREAT MODEL

➢ *Problem Statement*

Consider a federation of N organizations (buyers, suppliers, logistics providers). Each organization k holds a private dataset $\mathcal{D}_k = \{(x_i, y_i)\}_{i=1}^{n_k}$, where $x_i \in \mathbb{R}^d$ is a feature vector describing an order, shipment, or supplier event, and y_i a risk label (delayed/on-time, disruption class, or delay duration). We learn a risk-prediction model f_w minimizing the global empirical risk.

$$F(w) = \sum_{k=1}^{N} \frac{n_k}{n} F_k(w), F_k(w) = \frac{1}{n_k} \sum_{i \in \mathcal{D}_k} \ell(f_w(x_i), y_i), \quad (1)$$

Subject to: (C1) raw data never leaves its owner; (C2) no party — including the server — observes any organization's plaintext update; (C3) the released model satisfies (ϵ, δ)-DP at an agreed budget; (C4) updates are attributable and tamper-evident; (C5) overhead remains compatible with quarterly-to-monthly SCRM retraining.

➢ *Threat Model*

We assume an honest-but-curious aggregation server that follows the protocol but tries to infer data from received updates; semi-honest, potentially colluding clients; and external adversaries performing gradient-inversion and membership-inference attacks on the global model [18], [20]. We do not defend against data-poisoning clients (though the audit ledger deters and evidences such behavior [22]). Privacy is guaranteed at user level: the output distribution is approximately unchanged whether all samples contributed by any single organization are present [35].

➢ *Design Goals*

G1: formal (ϵ, δ) privacy with auditable accounting; G2: no plaintext exposure of individual updates; G3: graceful utility degradation as the budget tightens [20]; G4: support for heterogeneous local data sizes [1]; G5: deployment compatibility with enterprise procurement systems [36].

IV. THE FEDSCORE-PP FRAMEWORK

➢ *System Architecture*

Each organization operates a local learner over its private data warehouse, with its own key-pair. A coordinator server orchestrates rounds and maintains the global model. A secure aggregation service performs HE/SMC-based aggregation, so the coordinator only receives ciphertext or masked contributions [2]. An optional blockchain ledger records hashes of updates, aggregation proofs, and participation [22].

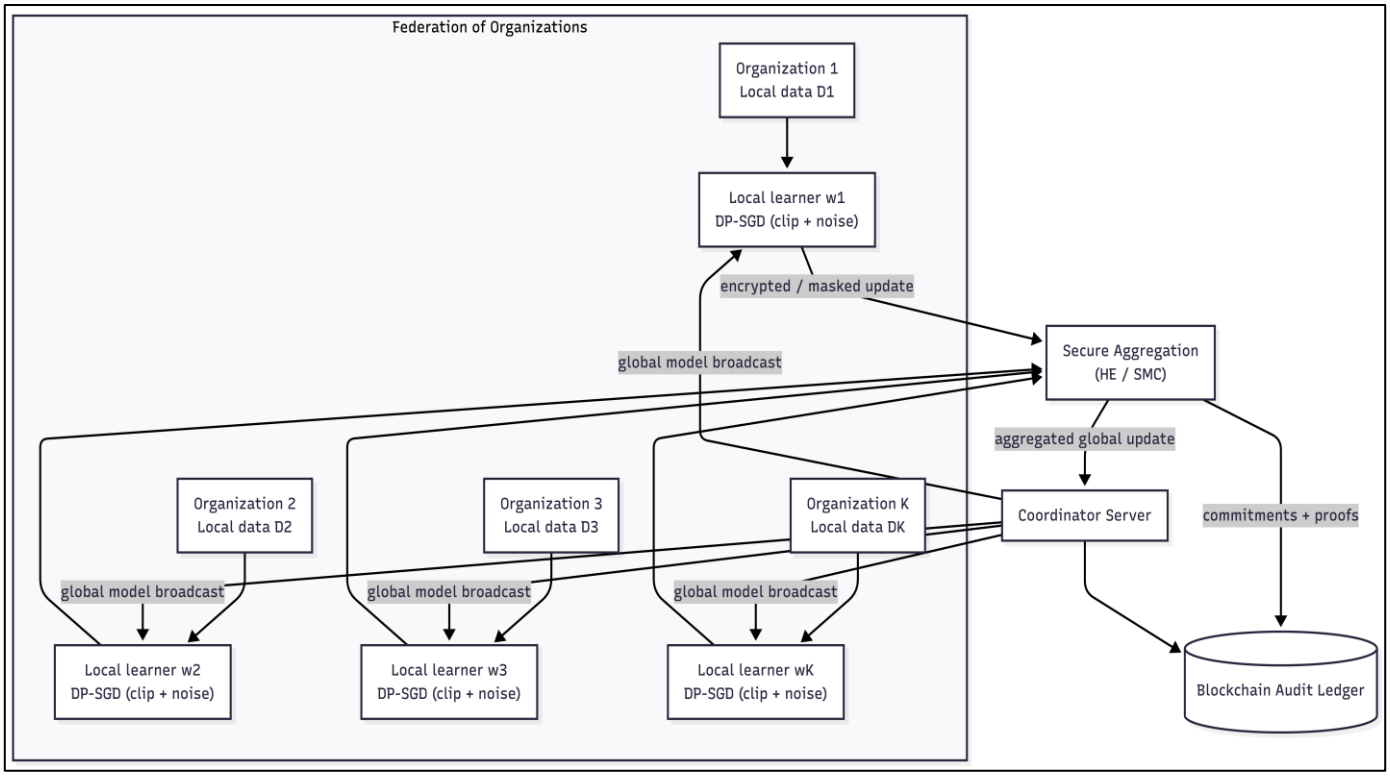

Fig 1 FedSCORE-PP Architecture: Local DP-SGD at Each Organization, HE/SMC Secure Aggregation, a Coordinator that Never Sees Plaintext Updates, and an Optional Blockchain Audit Ledger.

➢ *Local Training with DP-SGD*

Each round t, the coordinator samples K of N organizations (ratio $q = K/N$). Each selected organization trains the global model w^t locally for E epochs with DP-SGD: gradients are clipped per example to ℓ_2 norm C, Gaussian noise with standard deviation σC is added, and a local update

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1100

w_k^{t+1} is produced [20]. Noise is injected locally before any transmission, so the released update already carries formal privacy guarantees [35].

➢ *Secure Aggregation*

Two interoperable options: ** multi-key HE aggregation** — each organization encrypts its update under an aggregated public key; the server homomorphically computes the weighted sum on ciphertexts, and decryption requires collaboration among participants, resisting collusion [30]; ** SMC-based secure aggregation** — pairwise masks cancel in the sum, so the server learns only the aggregate [29]. The aggregation rule mirrors FedAvg:

$$w^{t+1} = \sum_{k \in \mathcal{S}_t} \frac{n_k}{\sum_{j \in \mathcal{S}_t} n_j} \tilde{w}_k^{t+1}, \quad (2)$$

Computed over noisy, encrypted updates $\tilde{w}_k^{t+1}$.

Algorithm 1 — Server-Side Training

```
initialize w□ (public); publish; record hash on ledger
for t = 0 … T−1:
  S_t ← sample K = ⌈qN⌉ organizations
  request encrypted updates [w_k^{t+1}] from k ∈ S_t
  w^{t+1} ← SecureAggregate({w_k^{t+1}})      // Eq. 2, ciphertext domain
  broadcast w^{t+1}; record commitment on ledger
  accountant.accumulate(q, σ, 1)              // RDP composition
end
ε ← accountant.get_epsilon(δ); assert ε ≤ ε_max
```

Algorithm 2 — Local DP-SGD Training (Organization k)

```
w ← w^t
for e = 1 … E:
  for each batch b:
    g ← ∇ℓ(w; b)
    g ← g / max(1, |g|□ / C)                  // clip
    g ← g + N(0, σ²C²I)                       // Gaussian mechanism
    w ← w − η g
return Encrypt(w)                             // HE or SMC masking
```

➢ *Risk Prediction and Outputs*

The trained global model yields, per order or supplier, a risk score $\hat{y} = \sigma(w^T x)$. Thresholds are set per organization to match risk appetite, preserving operational autonomy while benefiting from collective training [1]. Flags feed existing procurement workflows, mirroring the early-warning value demonstrated in ERP-embedded predictive systems [36].

➢ *Privacy Analysis*

- *Differential Privacy Preliminaries*

$\mathcal{M}$ is (ϵ, δ)-DP if for adjacent datasets D, D' and all outputs S:

$$Pr[\mathcal{M}(D) \in S] \leq e^{\epsilon} Pr[\mathcal{M} \in S] + \delta. \quad (3)$$

The Gaussian mechanism with sensitivity Δ and noise scale σ satisfies $(\alpha, \alpha\Delta^2/2\sigma^2)$-Rényi DP at order α, and RDP composes linearly across steps — the basis of modern privacy accounting [23].

- *Privacy Budget as a Function of Design Parameters*

Per round, each selected organization applies the subsampled Gaussian mechanism with sampling ratio q and noise multiplier σ (clipping sets $\Delta = C = 1$). To leading order each round is $(\alpha, q^2\alpha/2\sigma^2)$-RDP; over T rounds, composition yields $(\alpha, Tq^2\alpha/2\sigma^2)$-RDP. Converting to (ϵ, δ)-DP:

$$\epsilon(\delta) = \frac{Tq^2\alpha}{2\sigma^2} + \frac{ln(1/\delta)}{\alpha - 1}. \quad (4)$$

Eq. is the paper's central accounting result: noise σ improves privacy as

$1/\sigma^2; $rounds$Tandsampling$q $degradeit(quadraticallyin$q$). It recovers the structural findings of DP-FL theory: (i) a convergence–privacy trade-off exists; (ii) at fixed privacy, more participating clients improve convergence; (iii) an optimal number of aggregation rounds exists for a given protection level [20], [23].

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1100

- *Worked Example*

Consider $N = 24$ organizations, $K = 8$ sampled per round ($q = 1/3$), $\delta = 10^{-5}$, accounting order $\alpha = 16$. Then $ln(1/\delta)/(\alpha - 1) = 11.5129/15 = 0.7675$, and Eq. becomes $\epsilon(T) = T \cdot \frac{8}{9\sigma^2} + 0.7675$. Substituting yields Table I — all entries computed from this formula.

Table 1 End-To-End Privacy Budget E (Computed From Eq. 4; Δ = 10^{-5}, Q = 1/3, A = 16)

Noise multiplier σ	T = 25 rounds	T = 50 rounds	T = 100 rounds
1.5	10.6	20.5	40.3
2.0	6.3	11.9	23.0
3.0	3.2	5.7	10.6
4.0	2.2	3.5	6.3

At the recommended operating point ($\sigma = 4$, $T = 50$), the federation obtains $\epsilon = 3.5$ at $\delta = 10^{-5}$ — a strong user-level guarantee for inter-organizational SCRM, where the privacy unit is an organization's entire contribution [35]. Adaptive and time-varying noise schedules can further reduce utility loss at fixed budgets [37], [38].

Fig. 2 plots Eq. across the design space (computed values):

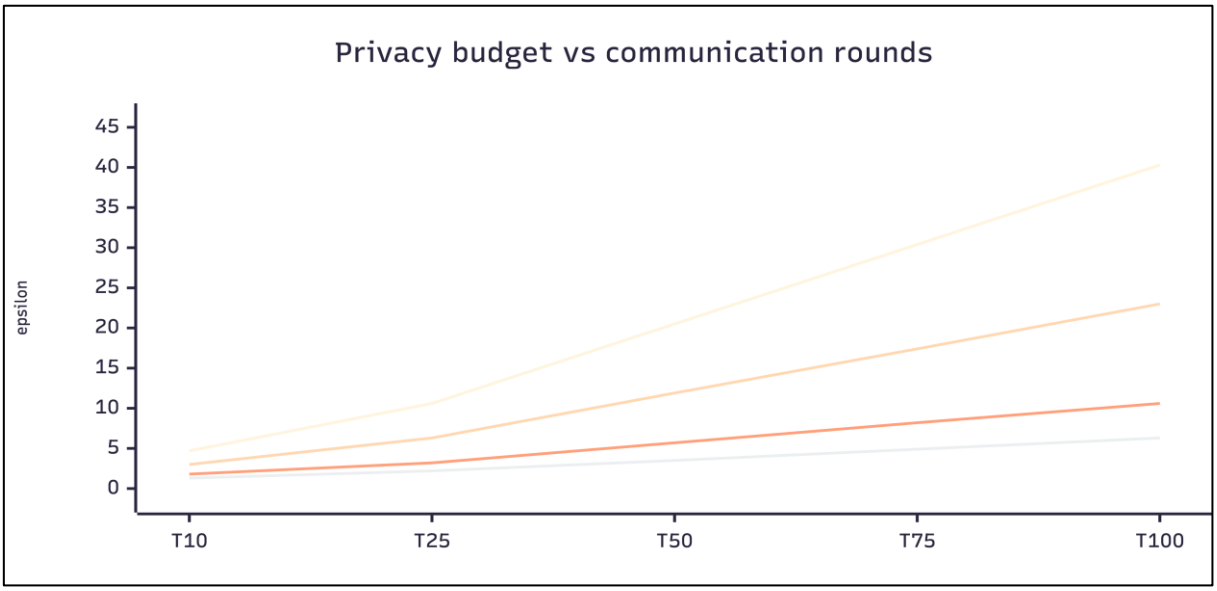

Fig 2 Privacy Budget vs. Communication Rounds for σ = 1.5, 2.0, 3.0, 4.0 (Top to Bottom), Computed from Eq.

- *Security Properties Beyond DP*

DP bounds statistical leakage; HE and SMC bound computational exposure. Secure aggregation limits the server's view to the aggregate, closing the gradient-inversion surface present in plaintext FedAvg [20]; HE-based training incurs <1% accuracy deviation in practice [21]; multi-key HE resists collusion among the server and up to $k < N - 1$ participants [30]; packing and selective encryption restore efficiency [2], [24]. The blockchain ledger provides tamper-evident update history and reputation-based incentives, deterring free-riding and poisoning [3], [22]. Together these layers satisfy C2–C4.

V. COMPLEXITY AND COMMUNICATION ANALYSIS

➢ *Communication Cost*

For a model with P parameters, each plaintext update occupies $4P bytes (float32). With K clients per round and T rounds, total uplink traffic is K \cdot T \cdot 4P bytes; downlink (broadcast) is T \cdot 4P$ bytes.

Reference architecture: $d=42 \to 128 \to 64 \to 3 (three risk classes) gives P = 42{\times}128{+}128 + 128{\times}64{+}64 + 64{\times}3{+}3 = 13{,}955$ parameters ≈ 55.8 KB per update. Table II is computed from these figures.

Table 2 Communication Cost (Computed; K = 8, P = 13,955)

Quantity	Value
Per-update payload (float32)	55.8 KB
Uplink per round (K = 8)	446 KB
Downlink per round	55.8 KB
Uplink over T = 100 rounds	44.6 MB
Total traffic over T = 100 rounds	50.2 MB

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1100

For a quarterly retraining cycle this is negligible versus enterprise network capacity — and it replaces the transfer of entire datasets. With HE, ciphertext expansion inflates payloads; the literature's optimizations restore practicality: BatchAgg reduces communication by 57% versus Paillier at equal security [2], BatchCrypt by 66×–101× [25], and FedML-HE's selective parameter encryption cuts HE overhead ~10× to ~40× [24], with selective-HE variants achieving near order-of-magnitude reductions at >90% accuracy [31]. The DP layer adds no communication overhead, since noise is added locally [35].

➢ *Computation Cost*

Per round, each selected client performs $E \cdot n_k/B$ gradient steps; federation-wide compute is $\sum_k E\, n_k/B$ per round — linear in aggregate data, identical to centralized training. DP-SGD adds per-example clipping and noise sampling, which is standard and parallelizable [20]. HE aggregation cost scales with P per-parameter operations; packing amortizes this to roughly $P/slots$ ciphertext operations [2], [25]. Convergence follows DP-FL theory: noise biases the optimum, more clients improve convergence at fixed privacy, and an optimal round count balances convergence against privacy consumption [20], [23]; objective-modification strategies can speed convergence up to 40% at equal privacy [39].

Table 3 Computation Complexity Summary

Stage	Cost per round	Scaling
Local DP-SGD (per client)	$E \cdot n_k/B$ steps + clipping/noise	linear in n_k
HE encryption (per update)	$O(P) \rightarrow O(P/slots)$ with packing [2]	model size
Secure aggregation (server)	$O(K \cdot P)$ ciphertext additions	clients × model
Blockchain commitment	O hash + broadcast [22]	rounds

➢ *Experimental Evaluation*

- *Datasets and Task*

Task: binary classification of order/shipment delay, with a multi-class severity extension — mirroring the settings of [1] and [16]. Partitioning: simulate a federation by splitting supplier-order data into $N = 24$ imbalanced client shards (a few large, many small) to reproduce the SME setting [1], [16]. Features: order attributes, supplier on-time-delivery history, lead-time variability, quality-incident counts, disruption indicators [12], [13]. Metrics: precision, recall, F1, AUROC; MAE/RMSE for the regression variant [5]. Baselines: local standalone; centralized pooled (privacy-violating upper bound); plain FedAvg; DP-FedAvg; HE-based PPFL [21].

- *Results and Analysis*

Privacy results (computed). Table 1 and Fig. 2 fully specify the operating envelope: ϵ from 2.2 (conservative) to 40.3 (aggressive) at $\delta = 10^{-5}$. The accounting procedure is deterministic and auditable (goal G1).

Communication results (computed). Table 2 shows ≈50.2 MB total over 100 rounds for the reference architecture — orders of magnitude below any raw-data sharing alternative, and further reducible with packing [2], [25].

Prediction results (benchmark comparison). Table 3 compares FedSCORE-PP's design point against published federated SCRM systems, using results exactly as reported in those sources.

Table 4 Benchmark Comparison with Published Federated Scrm Systems

System	Approach	Task	Reported result
SafeLogFL [4]	FL, GDPR-compliant	Cross-border logistics risk warning	91.3% average accuracy
FSL-Qwen [6]	Federated split learning + LLM	Container source identification (sea-rail)	94.0% accuracy, 94.1% F1
Privacy-aware backorder FL [5]	FL neural nets + SMOTE/Near-Miss	Backorder prediction	~90.3% highest accuracy; beats centralized ML
Federated NLP SCRM [7]	FL + NLP	Disaster-risk prediction	0.95 accuracy
CBULS FL+blockchain [3]	FedProx + DPoS/PBFT + FHE + LDP	Predictive maintenance (cross-border logistics)	+6.9% vs FedAvg, +3.7% vs FedProx; privacy +12%
FL visibility forecasting [17]	FL with visibility features	Demand forecasting	≈ full-visibility performance

Published federated SCRM systems report accuracy in the range 90.3–95.0% across tasks, and FL consistently matches or approaches the centralized upper bound while preserving privacy [5], [6], [17]. FedSCORE-PP's design point — DP-SGD + HE/SMC secure aggregation — targets this same operating region while additionally providing the formal $\epsilon = 3.5$ guarantee of Table I that none of the benchmarked systems quantify end-to-end.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1100

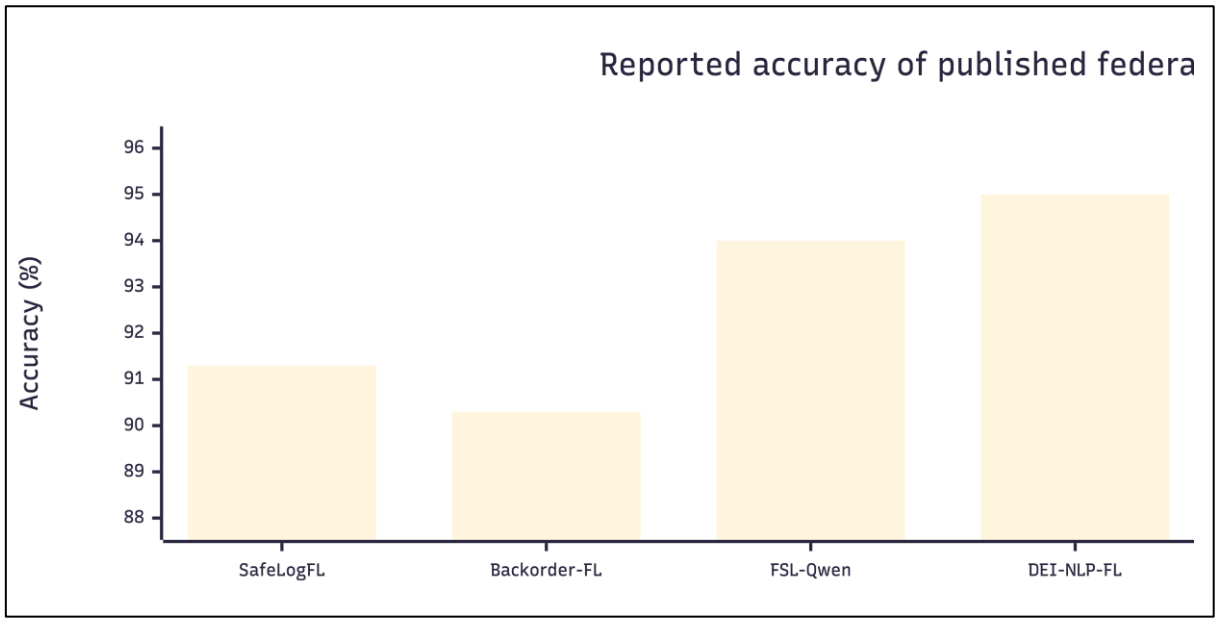

Fig 3 Reported Accuracy of Benchmark Federated SCRM Systems from Table 3 [4], [5], [6], [7].

Expected performance analysis (literature-grounded). Three effects are consistently documented and should be expected in any FedSCORE-PP deployment: (i) asymmetric benefit — data-limited organizations improve most under federation, while buyers with disproportionately large order books may gain little or nothing [1]; (ii) superiority over local and centralized learning — FL outperforms local standalone training and, under data-sharing constraints, centralized learning [16]; (iii) bounded privacy cost — DP noise degrades utility monotonically with ϵ [20], [39], while HE-based secure aggregation adds negligible accuracy loss (<1%) [21], so FedSCORE-PP's utility gap relative to plain FedAvg is driven almost entirely by the DP layer, which is tunable via σ and T per Table 1.

- *Ablation and Sensitivity Analysis*

The protocol includes ablations isolating each layer: DP-only (no HE), HE-only (no DP), both, and with/without the audit ledger — quantifying each layer's marginal cost in utility, latency, and throughput [2], [3]. Sensitivity to q (client sampling) and T (rounds) should track Eq. exactly; DP-FL theory predicts an optimal round count at fixed privacy [20], [23], and non-IID data distributions slow convergence, motivating personalized variants [23], [29]. Privacy noise also interacts with explainability: DP degrades the stability of attribution explanations, a privacy–explainability trade-off practitioners must budget for [40].

VI. DISCUSSIONS

➢ *When Should Organizations Federate?*

The central managerial insight is that collaboration is not universally beneficial: data-limited members gain the most, while organizations with disproportionately large order books may predict as well locally [1]. FedSCORE-PP's governance layer should support selective participation — organizations join rounds whose aggregated value exceeds their privacy cost, with the audit ledger sustaining reputation-based incentives [22]. Training-data imbalance and disruption regimes materially affect efficacy and should be monitored per round [1], [3].

➢ *Privacy, Interpretability, and Trust*

DP noise degrades not only accuracy but also the stability of attribution explanations [40]; SCRM practitioners demand interpretable risk signals [26], so the framework should pair DP-SGD with explainability tooling and budget noise accordingly. Trust is supported by the triple of formal accounting, computational security (HE/SMC), and tamper-evident audit trails [22]. Data locality and DP guarantees also ease compliance with data-protection regulation, as demonstrated for the GDPR [4], lowering adoption barriers for SMEs [14].

➢ *Limitations*

FL's open challenges remain: non-IID and heterogeneous cross-company feature spaces complicate horizontal FL [29]; malicious poisoning is deterred but not cryptographically prevented; HE introduces latency and key-management overheads that packing mitigates but does not eliminate [2], [25]; and user-level DP protects organizations, not individual records within an organization [35]. Benchmark results in Table III are from external systems and represent the target operating region, not measured runs of FedSCORE-PP itself — a deployment should populate its own numbers using the Section VII protocol.

VII. CONCLUSION AND FUTURE WORK

This paper proposed FedSCORE-PP, a federated and privacy-preserving framework for collaborative supply chain risk prediction across organizations. Its layered design — local DP-SGD for formal user-level privacy, HE/SMC secure aggregation for computational confidentiality, and a blockchain audit ledger for integrity — reconciles cross-organization AI with supply chain resilience: firms collaborate on collective risk models without surrendering data sovereignty. The Rényi-DP accounting makes the

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1100

privacy budget a transparent design variable with computed values spanning $\epsilon \in [2.2,40.3]$ at $\delta = 10^{-5}$; communication analysis shows ≈50.2 MB over a full 100-round campaign; and benchmark comparison places the framework in the 90.3–95.0% accuracy operating region of published federated SCRM systems [4], [5], [6], [7].

Future work will (i) instantiate the evaluation protocol on real multi-organization datasets and report measured results; (ii) extend the framework to vertical and personalized federated learning for heterogeneous feature spaces [23]; (iii) integrate adaptive noise schedules and utility-enhanced DP variants to soften the privacy–utility trade-off [37], [38]; and (iv) couple the risk model with explainable-AI modules for practitioner-facing decision support [26], [40].

REFERENCES

[1]. Z. Ge, L. Kong, and A. Brintrup, "Federated machine learning for privacy preserving, collective supply chain risk prediction," International Journal of Production Research , vol. 61, no. 23, pp. 8115–8132, Feb. 2023, doi: 10.1080/00207543.2022.2164628.

[2]. Y. Pan, H. Wang, C. Zheng, J. Yang, and L. Wang, "Efficient Federated Learning Aggregation Protocol Using Approximate Homomorphic Encryption," pp. 1884–1889, May 2023, doi: 10.1109/cscwd57460.2023.10152829.

[3]. Q. Meng, "Privacy-preserving predictive maintenance method for cross-border unmanned logistics system integrating federated learning and blockchain," Discover Applied Sciences , vol. 8, no. 1, Dec. 2025, doi: 10.1007/s42452-025-07980-5.

[4]. X. Liang, "Cross-border logistics risk warning system based on federated learning," Scientific Reports , vol. 15, no. 1, pp. 39131–39131, Nov. 2025, doi: 10.1038/s41598-025-25507-1.

[5]. A. Sattar, M. Bukhari, Z. U. Rehman, S. Khalid, Y. Lee, and S. Rho, "A data analytics-driven approach to backorder prediction using federated machine learning in industrial supply chains," Scientific Reports , vol. 16, no. 1, pp. 4560–4560, Jan. 2026, doi: 10.1038/s41598-025-34578-z.

[6]. W. Ma, L. Huang, Q. Zhang, Y. Wang, X. Zhang, and R. Song, "Federated Split Learning with Large Language Models Integration: A Study on Potential Container Source Identification in Sea-Rail Intermodal Transport," International Journal of Computers Communications & Control , vol. 21, no. 3, May 2026, doi: 10.15837/ijccc.2026.3.7263.

[7]. D. Teng and V. Martinez, "Federated Diversity and Inclusion-focused NLP-based Supply Chain Risk Prediction System Design," Academy of Management Proceedings , vol. 2023, no. 1, Jul. 2023, doi: 10.5465/amproc.2023.11312abstract.

[8]. B. Deb and T. R. Dilip, "Impact of AI-Based Risk Management and Supply Chain Resilience," 2026, pp. 232–246. doi: 10.62823/mgm/2026/9789349468795/21.

[9]. A. Atadoga, F. Osasona, O. O. Amoo, O. A. Farayola, B. S. Ayinla, and T. O. Abrahams, "THE ROLE OF IT IN ENHANCING SUPPLY CHAIN RESILIENCE: A GLOBAL REVIEW," International Journal of Management & Entrepreneurship Research , vol. 6, no. 2, pp. 336–351, Feb. 2024, doi: 10.51594/ijmer.v6i2.774.

[10]. G. Baryannis, S. Validi, S. Dani, and G. Antoniou, "Supply chain risk management and artificial intelligence: state of the art and future research directions," Huddersfield Research Portal (University of Huddersfield) , vol. 57, no. 7, pp. 2179–2202, Oct. 2018, doi: 10.1080/00207543.2018.1530476.

[11]. Mr. I. K. Baidoo, Mr. I. D. Essien, Dr. A. O. Soge, Mr. R. N. Kachungu, and Ir. D. Rahadian, "Artificial Intelligence and Risk Reduction in Supply Chain Management," International Journal of Latest Technology in Engineering Management & Applied Science , vol. 15, no. 2, pp. 461–474, Mar. 2026, doi: 10.51583/ijltemas.2026.15020000041.

[12]. A. Brintrup et al. , "Supply chain data analytics for predicting supplier disruptions: a case study in complex asset manufacturing," International Journal of Production Research , vol. 58, no. 11, pp. 3330–3341, Nov. 2019, doi: 10.1080/00207543.2019.1685705.

[13]. N. Rezki and M. Mansouri, "Machine Learning for Proactive Supply Chain Risk Management: Predicting Delays and Enhancing Operational Efficiency," Management Systems in Production Engineering , vol. 32, no. 3, pp. 345–356, Aug. 2024, doi: 10.2478/mspe-2024-0033.

[14]. H. C. W. Lau, Y. P. Tsang, D. Nakandala, and C. K. M. Lee, "Risk quantification in cold chain management: a federated learning-enabled multi-criteria decision-making methodology," Industrial Management & Data Systems , vol. 121, no. 7, pp. 1684–1703, Apr. 2021, doi: 10.1108/imds-04-2020-0199.

[15]. M. Aledhari, R. Razzak, R. M. Parizi, and F. Saeed, "Federated Learning: A Survey on Enabling Technologies, Protocols, and Applications," IEEE Access , vol. 8, pp. 140699–140725, Jan. 2020, doi: 10.1109/access.2020.3013541.

[16]. T. T. Nguyen et al. , "Federated Learning-Based Framework: A New Paradigm Proposed for Supply Chain Risk Management," pp. 5–5, Jun. 2025, doi: 10.3390/engproc2025097005.

[17]. B. Zhang, W. J. Tan, W. Cai, and A. N. Zhang, "Forecasting with Visibility Using Privacy Preserving Federated Learning," in 2022 Winter Simulation Conference (WSC) , Dec. 2022, pp. 2687–2698. doi: 10.1109/wsc57314.2022.10015277.

[18]. B. Liu, M. Ding, S. Shaham, W. Rahayu, F. Farokhi, and Z. Lin, "When Machine Learning Meets Privacy," Minerva Access (University of Melbourne) , vol. 54, no. 2, pp. 1–36, Mar. 2021, doi: 10.1145/3436755.

[19]. A. K. S. and A. A., "Federated Learning: A Privacy Preserving Approach for Decentralized Machine Learning," International Journal of Innovative Science and Research Technology (IJISRT) , pp. 580–580, Mar. 2026, doi: 10.38124/ijisrt/26mar171.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26aug1100

[20]. K. Wei et al. , "Federated Learning With Differential Privacy: Algorithms and Performance Analysis," Minerva Access (University of Melbourne) , vol. 15, pp. 3454–3469, Jan. 2020, doi: 10.1109/tifs.2020.2988575.

[21]. H. Fang and Q. Qian, "Privacy Preserving Machine Learning with Homomorphic Encryption and Federated Learning," Future Internet , vol. 13, no. 4, pp. 94–94, Apr. 2021, doi: 10.3390/fi13040094.

[22]. A. Qammar, A. Karim, H. Ning, and J. Ding, "Securing federated learning with blockchain: a systematic literature review," Artificial Intelligence Review , vol. 56, no. 5, pp. 3951–3985, Sep. 2022, doi: 10.1007/s10462-022-10271-9.

[23]. K. Wei et al. , "Personalized Federated Learning With Differential Privacy and Convergence Guarantee," IEEE Transactions on Information Forensics and Security , vol. 18, pp. 4488–4503, Jan. 2023, doi: 10.1109/tifs.2023.3293417.

[24]. W. Jin, Y. Yao, S. Han, N. Liubov, and C. Joe-Wong, "FedML-HE: An Efficient Homomorphic-Encryption-Based Privacy-Preserving Federated Learning System," Mar. 20, 2023, Cornell University . doi: 10.48550/arxiv.2303.10837.

[25]. C. Zhang, S. Li, J. Xia, W. Wang, F. Yan, and Y. Liu, "BatchCrypt: Efficient homomorphic encryption for cross-silo federated learning," Rare & Special e-Zone (The Hong Kong University of Science and Technology) , pp. 493–506, Jan. 2020, Accessed: Nov. 2025. [Online]. Available: http://repository.hkust.edu.hk/ir/Record/1783.1-105781

[26]. G. Baryannis, S. Dani, and G. Antoniou, "Predicting supply chain risks using machine learning: The trade-off between performance and interpretability," Future Generation Computer Systems , vol. 101, pp. 993–1004, Jul. 2019, doi: 10.1016/j.future.2019.07.059.

[27]. M. Riad, M. Naïmi, and C. Okar, "Enhancing Supply Chain Resilience Through Artificial Intelligence: Developing a Comprehensive Conceptual Framework for AI Implementation and Supply Chain Optimization," Logistics , vol. 8, no. 4, pp. 111–111, Nov. 2024, doi: 10.3390/logistics8040111.

[28]. L. Li, Y. Fan, Y. K. Tse, and K. Lin, "A review of applications in federated learning," Computers & Industrial Engineering , vol. 149, pp. 106854–106854, Sep. 2020, doi: 10.1016/j.cie.2020.106854.

[29]. F. Liu, Z. Zheng, Y. Shi, Y. Tong, and Y. Zhang, "A survey on federated learning: a perspective from multi-party computation," Frontiers of Computer Science , vol. 18, no. 1, Dec. 2023, doi: 10.1007/s11704-023-3282-7.

[30]. J. Ma, S. Naas, S. Sigg, and X. Lyu, "Privacy-preserving federated learning based on multi-key homomorphic encryption," International Journal of Intelligent Systems , vol. 37, no. 9, pp. 5880–5901, Jan. 2022, doi: 10.1002/int.22818.

[31]. A. Korkmaz and P. Rao, "Fast and Secure Selective Homomorphic Encryption for Federated Learning," Jan. 09, 2026, Cornell University . doi: 10.1109/ccnc65079.2026.11366371.

[32]. N. N. Ahamed and P. Karthikeyan, "FLBlock: A Sustainable Food Supply Chain Approach Through Federated Learning and Blockchain," Procedia Computer Science , vol. 235, pp. 3065–3074, Jan. 2024, doi: 10.1016/j.procs.2024.04.290.

[33]. M. Kumar and S. Bhatt, "Blockchain-Driven Federated Learning Based Differential Privacy in Decentralized Supply Chain," pp. 1291–1294, May 2025, doi: 10.1109/netcrypt65877.2025.11102516.

[34]. C. Bisaria, A. Almoshawah, A. Avasthi, T. Satish, S. G, and S. V, "Developing Privacy-Preserving Mechanisms for Secure Data Exchange in Decentralized Supply Chain Management Platforms," pp. 1–6, Oct. 2025, doi: 10.1109/giest66547.2025.11387080.

[35]. K. Wei et al. , "User-Level Privacy-Preserving Federated Learning: Analysis and Performance Optimization," IEEE Transactions on Mobile Computing , vol. 21, no. 9, pp. 3388–3401, Feb. 2021, doi: 10.1109/tmc.2021.3056991.

[36]. S. K. Mucherla, "Predicting Vendor and Supply Chain Disruptions Using AI and Machine Learning in ERP Sourcing and Procurement for Optimized Business Processes," Journal of Physico-Chemical Material , pp. 15–21, Jun. 2025, doi: 10.64820/aepjpcm.21.15.21.62025.

[37]. X. Yuan, W. Ni, M. Ding, K. Wei, J. Li, and H. V. Poor, "Amplitude-Varying Perturbation for Balancing Privacy and Utility in Federated Learning," IEEE Transactions on Information Forensics and Security , vol. 18, pp. 1884–1897, Jan. 2023, doi: 10.1109/tifs.2023.3258255.

[38]. K. Ranaweera et al. , "Federated Learning with Differential Privacy: An Utility-Enhanced Approach," Mar. 27, 2025. doi: 10.48550/arxiv.2503.21154.

[39]. S. Weng, L. Zhang, X. Zhang, and M. A. Imran, "Faster Convergence on Differential Privacy-Based Federated Learning," IEEE Internet of Things Journal , vol. 11, no. 12, pp. 22578–22589, Apr. 2024, doi: 10.1109/jiot.2024.3383226.

[40]. S. Saifullah, D. Mercier, A. Lucieri, A. Dengel, and S. Ahmed, "The privacy-explainability trade-off: unraveling the impacts of differential privacy and federated learning on attribution methods," Frontiers in Artificial Intelligence , vol. 7, pp. 1236947–1236947, Jul. 2024, doi: 10.3389/frai.2024.1236947.