ECOFED-X: Zero-Trust, Compliance-Aware Orchestration for Federated Learning in Critical Applications*

Francesco Avella

DIAG Department
Sapienza University of Rome, Italy
francesco.avella@uniroma1.it

Abstract

Federated Learning (FL) is increasingly adopted in healthcare and other critical domains as a privacy-preserving paradigm for collaborative model training across decentralized and sensitive data silos. Real-world FL deployments require more than algorithmic privacy: they must demonstrate compliance with data protection laws (e.g., GDPR, HIPAA), enforce runtime security guarantees, and maintain auditability across heterogeneous cloud infrastructures. Current FL frameworks typically focus on communication, aggregation, or differential privacy and often ignore infrastructure-level constraints, producing regulatory risk, inefficient resource usage, and weak forensic evidence. We present ECOFED-X, an extension of the ECOFED FinOps orchestration platform (Avella and Silvestri 2025), designed to run FL pipelines in regulated, multicloud environments. ECOFED-X introduces a formal Compliance Metric (Cm) that quantifies a resource's adherence to jurisdictional, contractual, runtime, and auditability requirements. Cm is integrated into a unified scheduler that jointly optimizes compliance, utility, cost, and carbon emissions. We formalize the orchestration problem and present practical workflows including policy-first filtering, attestation gating, and just-in-time provisioning. A healthcare scenario across AWS, Azure, and GCP illustrates expected benefits: reduced compliance violations, improved accountability via signed metadata and attestations, lower idle time and carbon footprint through FinOps integration, and preserved model performance via utility-aware scheduling.

Code — https://gitlab.com/francesco.avella/ecofedx

Introduction

In domains like healthcare and finance, where data privacy is non-negotiable, organizations face a dilemma: how to collaborate without compromising sensitive information. Federated Learning (FL) offers a promising answer by enabling multiple institutions to train models jointly while keeping data; this property is particularly critical in regulated and safety-sensitive domains (Singh et al. 2025; Moran 2025). This collaborative potential is especially valuable in healthcare, where privacy and data locality are critical (Kovačević 2025). While algorithmic techniques (secure aggregation, differential privacy) address data exposure at the model and communication level, real-world FL deployments for critical applications require additional operational guarantees: demonstrable regulatory and contractual compliance (for example, GDPR and HIPAA), verifiable runtime integrity of execution environments, and robust, auditable orchestration across heterogeneous cloud providers (Intel 2025; Yan et al. 2024). Auditability mechanisms such as blockchain-based verification have also been proposed to strengthen trust in FL orchestration (Liu, Zhang, and Rao 2023).

These guarantees are especially crucial when handling sensitive data types such as Personally Identifiable Information (PII) and Protected Health Information (PHI), which are subject to strict regulatory controls.

The original ECOFED framework introduced a FinOps-oriented orchestration approach for FL that reduces idle time, monetary cost and carbon emissions by jointly optimizing placement and scheduling across multiple cloud providers (Avella and Silvestri 2025). ECOFED-X extends ECOFED by elevating compliance and runtime trust to first-class scheduling objectives and by integrating a dedicated Compliance Module and a zero-trust runtime verification layer. In ECOFED-X, the orchestrator is not only cost- and carbon-aware but also compliance-aware: scheduling decisions are driven by explicit, machine-readable policies encoding legal, contractual, and organizational constraints such as data residency, Data Processing Agreements (DPAs, as required under GDPR), Business Associate Agreements (BAAs, as required under HIPAA), KMS ownership semantics, and required certifications. Every selected execution environment must cryptographically prove its state and configuration before and during training using attestation mechanisms such as Trusted Execution Environments (TEEs) (Yan et al. 2024).

The "X" in ECOFED-X denotes:

- eXtended capabilities: compliance, trust, and auditability
- eXtreme environments: healthcare, finance, mobility
- eXplainable operations: real-time visibility into cost, carbon and compliance trade-offs

This fusion of policy-driven scheduling and cryptographic runtime verification,which we term compliance-aware

*Accepted at the 1st Workshop on Federated Learning for Critical Applications (FLCA), in conjunction with AAAI26
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

zero-trust, closes two operational gaps. First, declarative policies alone are insufficient if runtime misconfiguration or compromise can subvert stated guarantees. Second, vendor attestation evidence alone is insufficient if it is not evaluated against legal and contractual constraints: attestation can show that a host is technically intact while the host may still be located in a disallowed jurisdiction or lack required contractual assurances. ECOFED-X operationalizes both sides by (i) normalizing heterogeneous provider semantics and attestation formats via the Compliance Module, (ii) encoding legal and contractual requirements as machine-readable, versioned policies, (iii) enforcing attestation gating before provisioning and continuous attested heartbeats during execution, and (iv) centralizing only minimal signed, non-sensitive metadata (hashes, attestation outcomes, timestamps) to support auditable evidence without centralizing sensitive payloads.

This work first reviews related works on federated learning in regulated domains, focusing on compliance-aware aspects, runtime attestation mechanisms, and policy-driven orchestration strategies. It then formalizes the orchestration challenge in critical applications and introduces the Compliance Metric used to evaluate client eligibility. The proposed solution, ECOFED-X, is presented in detail with the selection workflow and integration with cloud-native compliance interfaces. To assess its practical impact, the paper reports experiments using a healthcare dataset, analyzing how compliance thresholds affect model performance and auditability. The paper concludes with a summary of contributions and outlines directions for future development and validation.

Related Work

Federated Learning (FL) is a distributed learning paradigm that enables collaborative model training without centralizing data, making it particularly suitable for regulated domains such as healthcare, finance, public administration, and telecommunications (Abbas, Mehta, and Zhou 2024; Kaimal and Thompson 2025). In these settings, compliance with legal frameworks like GDPR, HIPAA, or PCI-DSS is essential, and client selection must respect jurisdictional, contractual, and runtime constraints. The adoption of FL in critical applications has therefore prompted increasing interest in solutions that embed compliance directly into the training lifecycle (HIPAA Partners Team 2025).

Recent literature on FL in sensitive domains can be grouped into three main areas: (i) architectures and protocols that ensure regulatory compliance and data residency, (ii) technical mechanisms for attesting and verifying runtime capabilities (e.g., TEE, disk encryption, KMS binding, checkpoint signing) (Casella 2025; Gupta 2021), and (iii) governance models and policy-driven selection strategies that translate legal requirements into reproducible and auditable operational rules (Tahir and Ali 2025; Zhang et al. 2022).

Several studies explore how FL can satisfy privacy and compliance constraints by combining privacy-preserving techniques such as differential privacy and secure aggregation with distributed compliance controls. In healthcare, federated frameworks have been proposed for collaborative analysis of sensitive clinical data, integrating access control, auditability, and policy enforcement. In finance, other works address regulatory violation detection and contract-aware training orchestration. For example, FedLC introduces a federated framework for legal compliance detection in financial institutions (Zhang and Chen 2024).

From a systems engineering perspective, recent research describes how to integrate hardware attestation and contractual evidence into FL orchestration pipelines, highlighting trade-offs between communication efficiency and security guarantees. Methodological contributions introduce composite eligibility metrics that aggregate jurisdiction, certifications, and runtime capabilities into weighted scores used to select trustworthy clients.

The original ECOFED paper contributed to this line of work by presenting a practical architecture that includes a Compliance Module for evidence normalization, a canonical client profile, and a signed, versioned selection logic for auditability. Our work builds on this foundation by formalizing the Compliance Metric, specifying normalization constraints, and clarifying the role of governance in defining and versioning the selection weights, ensuring that operational decisions are reproducible, auditable, and aligned with data sensitivity.

Problem Statement

In safety-critical Federated Learning applications, we face the problem of selecting, at each round, a subset of clients distributed across multi-cloud environments so that processing complies with legal, contractual, and technical constraints without centralizing sensitive data. The practical challenge is to translate compliance requirements provided by the data owner (e.g., data residency, contractual clauses, attestable runtime capabilities) into reproducible and verifiable participation decisions: automatically exclude clients that do not guarantee the required legal and technical conditions, apply stricter rules according to data sensitivity, and record all choices for non-intrusive auditability (Tahir and Ali 2025; Dakhia and Merenda 2025; Kumar, Singh, and Patel 2025).

Recent work has proposed policy-driven client selection strategies that incorporate contextual constraints such as jurisdiction, certification, and runtime trustworthiness. Tahir and Ali formalize selection logic based on legal and operational profiles (Tahir and Ali 2025), while Dakhia and Merenda explore game-theoretic models for resource-constrained environments (Dakhia and Merenda 2025). Kumar et al. introduce mechanisms for auditable and verifiable participation, enabling trust in federated orchestration pipelines (Kumar, Singh, and Patel 2025). These contributions highlight the need for selection frameworks that are not only efficient but also legally sound and operationally transparent.

To solve this problem, we propose a selection algorithm based on compliance that, given a task:

$$T = (\mathrm{id}, R, P, M)$$

Table 1: Glossary of symbols used in ECOFED-X

Symbol	Description
$T = (\mathrm{id}, R, P, M)$	FL task: ID, requirements, policy, sensitivity
id	Task identifier
R	Min HW/SW requirements
$P = (J, S, Q)$	Policy: jurisdictions, certifications, runtime capabilities
J	Allowed jurisdictions
S	Required certifications/contracts
Q	Required runtime capabilities
M	Data sensitivity level (e.g., public, sensitive, PHI, PII)
C_{all}	Set of candidate clients
c	Canonical client profile, composed of:
loc	Client location
certs_valid	Certifications/contracts held
secCaps_verified	Verified runtime capabilities
$V \in [0, 1]$	Auditability readiness score
$L_T(c)$	Jurisdiction compliance
$S_T(c)$	Certification compliance
$Q_T(c)$	Runtime compliance
$V_T(c)$	Auditability component
w_L, w_S, w_Q, w_V	Weights for each component
$\mathrm{Cm}_T(c)$	Compliance Metric score
τ	Selection threshold
k	Clients selected per round

and a candidate set:

$$C_{\mathrm{all}},$$

uses for each client the canonical profile:

$$c = (\mathrm{loc},\ \mathrm{certs_valid},\ \mathrm{secCaps_verified},\ V)$$

The meaning and role of each symbol used in the client profile and selection process are summarized in Table 1. This glossary clarifies the notation used throughout the compliance-aware orchestration algorithm.

The selection process then unfolds as follows:

1. Pre-filter on R to remove clients that do not meet minimal HW/SW requirements.
2. Compute compliance components for each client c:

$$L_T(c) = \begin{cases} 1, & \text{if loc} \in P.J \\ 0, & \text{otherwise} \end{cases}$$

$$S_T(c) = \begin{cases} \dfrac{|\mathrm{certs_valid} \cap P.S|}{|P.S|}, & \text{if } |P.S| > 0 \\ 1, & \text{if } |P.S| = 0 \end{cases}$$

$$Q_T(c) = \begin{cases} \dfrac{|\mathrm{secCaps_verified} \cap P.Q|}{|P.Q|}, & \text{if } |P.Q| > 0 \\ 1, & \text{if } |P.Q| = 0 \end{cases}$$

$$V_T(c) = V \in [0, 1]$$

3. Compute the compliance metric:

$$\begin{aligned} \mathrm{Cm}_T(c) = {} & w_L L_T(c) + w_S S_T(c) \\ & + w_Q Q_T(c) + w_V V_T(c) \end{aligned}$$

$$w_L,\ w_S,\ w_Q,\ w_V \geq 0;$$
$$w_L + w_S + w_Q + w_V = 1$$

4. Threshold filter and selection: keep only clients with

$$\mathrm{Cm}_T(c) \geq \tau$$

where τ depends on sensitivity M. Compliant clients are sorted by $\mathrm{Cm}_T(c)$ (descending) and the top k are selected.

Weights (w_L, w_S, w_Q, w_V) are defined by governance (data owner, legal team, security/infrastructure leads, operational stakeholders). The Compliance Module applies versioned weights in computing $\mathrm{Cm}_T(c)$, provides support tools (simulations, recommendations), and logs every change (author, version, timestamp) in the Policy Logs. An example involving three clients is presented below to clarify the proposed algorithm (see Table 2).

Policy sets:

$$\begin{aligned} P.J &= \{\mathrm{EU}\} \\ P.S &= \{\mathrm{GDPR},\ \mathrm{BAA}\} \\ P.Q &= \{\mathrm{TEE},\ \mathrm{KMS_customer}\} \end{aligned}$$

Selection parameters:

$$\begin{aligned} & k = 3, \quad \tau = 0.95 && \text{(threshold)} \\ & w_L = 0.35, \quad w_S = 0.20 && \text{(legal, certs)} \\ & w_Q = 0.35, \quad w_V = 0.10 && \text{(security, auditability)} \end{aligned}$$

Client A profile:

$$\begin{aligned} \mathrm{loc} &= \mathrm{EU} \\ \mathrm{certs_valid} &= \{\mathrm{GDPR},\ \mathrm{BAA}\} \\ \mathrm{secCaps_verified} &= \{\mathrm{TEE},\ \mathrm{KMS_customer}\} \\ V &= 1.0 \end{aligned}$$

$$L = 1, \quad S = 1.0, \quad Q = 1.0, \quad V = 1.0$$
$$\Rightarrow \mathrm{Cm} = 1.00 \quad \text{(admitted)}$$

Client B profile:

$$\begin{aligned} \mathrm{loc} &= \mathrm{EU} \\ \mathrm{certs_valid} &= \{\mathrm{GDPR}\} \\ \mathrm{secCaps_verified} &= \{\mathrm{diskEnc}\} \\ V &= 0.6 \end{aligned}$$

$$L = 1, \quad S = 0.5, \quad Q = 0.0, \quad V = 0.6$$
$$\Rightarrow \mathrm{Cm} \approx 0.50 \quad \text{(excluded)}$$

Client C profile:

$$\begin{aligned} \mathrm{loc} &= \mathrm{EU} \\ \mathrm{certs_valid} &= \{\mathrm{GDPR},\ \mathrm{BAA}\} \\ \mathrm{secCaps_verified} &= \{\mathrm{TEE}\} \\ V &= 0.95 \end{aligned}$$

$$L = 1, \quad S = 1.0, \quad Q = 0.5, \quad V = 0.95$$
$$\Rightarrow \mathrm{Cm} \approx 0.82 \quad \text{(excluded)}$$

Table 2: Compliance evaluation for three clients under threshold $\tau = 0.95$. Abbreviations: G = GDPR, B = BAA, T = TEE, K = KMS_customer, D = diskEnc. Only Client A meets the required score and is selected.

ID	Loc.	Certs	Runtime	Audit (V)	Cm
A	EU	G,B	T,K	1.00	1.00
B	EU	G	D	0.60	0.50
C	EU	G,B	T	0.95	0.82

Table 3: Signed metadata fields logged for each selection decision

Field	Description
task.id	Unique identifier of the federated learning task
client.id	Identifier of the selected or evaluated client
Cm	Compliance Metric score computed for the client
timestamp	UTC time of the selection decision
reason	Justification for inclusion or exclusion (e.g., below threshold, missing certifications)
policy.version	Version of the policy set used for evaluation (includes weights and constraints)

All selections and the policy version used are stored as signed metadata:

$$\{\,\text{task.id},\ \text{client.id},\ \text{Cm},\ \text{timestamp},\ \text{reason},\ \text{policy.version}\,\}$$

to ensure traceability and auditability without exposing application data. Table 3 lists the metadata fields that ECOFED-X logs for each selection decision.

Selection based on the Compliance Metric provides a reproducible and verifiable rule to include only clients that meet legal, contractual, and technical requirements. Governance defines the weights, while the Compliance Module ensures versioning and traceability, aligning operational decisions with data sensitivity and auditability.

Proposed Solution

To address the challenges outlined in the previous sections, we present ECOFED-X, an extended framework for Federated Learning in multi-cloud environments designed to address both operational efficiency and regulatory compliance. ECOFED-X builds on the original ECOFED architecture, which was conceived to minimize cloud execution costs through a FinOps-oriented approach, aiming to reduce redundant resource usage and lower CO2 emissions associated with large-scale distributed training (Avella and Silvestri 2025).

The core objective of ECOFED-X is to enable secure, compliant, and cost-aware Federated Learning across heterogeneous cloud infrastructures, where clients may reside in different jurisdictions and operate under diverse contractual and runtime constraints. In such settings, selecting which clients may participate in each training round is non-trivial: it must respect legal boundaries (e.g., data residency), contractual obligations (e.g., DPA, BAA), and technical guarantees (e.g., attested runtime environments) (Dakhia and Merenda 2025; Kumar, Singh, and Patel 2025; Zhou et al. 2023).

To address this, we introduce a Compliance Selection Module within ECOFED-X that performs client selection based on a formalized Compliance Metric. For each training task, the module evaluates candidate clients against a set of requirements, including jurisdiction, certifications, runtime capabilities, and auditability readiness, and computes a weighted score that determines eligibility. Only clients that meet or exceed a sensitivity-dependent threshold are admitted to the round.

The compliance attributes used in the selection process are retrieved dynamically from cloud providers via standardized APIs or metadata interfaces. These include:

- **Jurisdiction and residency**: inferred from region tags or data location policies.
- **Contractual status**: validated through linked agreements or service-level metadata.
- **Runtime capabilities**: attested via cloud-native APIs (e.g., TEE availability, disk encryption status, KMS bindings).
- **Auditability signals**: derived from freshness and integrity of compliance evidence (Services 2025; Azure 2025; Platform 2025).

ECOFED-X is designed to operate across major cloud platforms, including Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP), each of which exposes compliance-related metadata through APIs such as AWS Config, Azure Policy Insights, or GCP Resource Manager. These interfaces allow ECOFED-X to query and verify client attributes without accessing application data (Katasani 2025; Albshaier, Almarri, and Albuali 2025).

Compliance is especially critical when handling sensitive data such as: PII (Personally Identifiable Information): data that can identify an individual, including names, addresses, ID numbers, or biometric records; PHI (Protected Health Information): medical data linked to individuals, regulated under laws like HIPAA in the US and GDPR in the EU.

In these contexts, Federated Learning must not only preserve privacy but also enforce strict participation rules to ensure that only clients operating under compliant conditions are involved. ECOFED-X enforces these rules through its Compliance Selection Module, and all decisions are versioned, signed, and logged for auditability ensuring that participation is not only compliant but also reproducible and traceable.

Experiments

To evaluate the impact of compliance-aware selection on model performance, we conducted a series of experiments using the NIH ChestX-ray14 dataset (Wang et al. 2017), a large-scale medical imaging dataset released by the National Institutes of Health. It contains over 100,000 frontal-view

Table 4: Impact of client selection on model accuracy. Lower τ admits more clients but reduces compliance strength.

τ	**Clients**	**Avg** Cm	**Accuracy (%)**
0.95	2	0.98	89.5
0.90	3	0.93	89.2
0.85	4	0.89	88.9
0.80	5	0.85	88.6
0.75	6	0.81	88.2

chest radiographs from more than 30,000 patients, annotated with 14 common thoracic pathologies such as pneumonia, cardiomegaly, and infiltration. Due to its clinical relevance and privacy sensitivity, ChestX-ray14 is widely adopted in federated learning research for healthcare applications. Other studies have used federated learning to analyze medical images across global healthcare networks while preserving privacy (Yadav, Veeramanickam, and Kumar 2025). To simulate a realistic non-IID scenario, we partitioned the dataset across 6 synthetic clients, distributed across various cloud platforms, using a Dirichlet distribution with concentration parameter $\alpha = 0.5$, inducing heterogeneous label distributions and reflecting real-world data silos in hospitals or clinics. Each client was assigned a randomized compliance profile including jurisdiction, certifications, runtime capabilities, and auditability score.

The federated training was orchestrated using the Flower framework (Beutel et al. 2020), which provides modular support for client-server communication, custom selection logic, and integration with cloud-native APIs. ECOFED-X was used to select clients at each round based on the Compliance Metric Cm, with varying thresholds τ to reflect different levels of data sensitivity.

Table 4 shows how increasing the number of selected clients (by lowering τ) leads to higher data diversity but also introduces weaker compliance guarantees. As more clients are admitted, the global model accuracy tends to decrease slightly due to the inclusion of less trustworthy environments, illustrating the trade-off between compliance and performance.

All selection decisions were logged with signed metadata including task ID, client ID, Cm, timestamp, and policy version. This ensures full auditability and reproducibility of the orchestration process. These results confirm that ECOFED-X enables flexible governance over compliance constraints while maintaining acceptable model performance in sensitive domains such as healthcare.

Conclusion and Future Work

In this work, we introduced ECOFED-X, an extension of the ECOFED framework for Federated Learning in multi-cloud environments, designed to combine operational efficiency with regulatory compliance. The system integrates a selection module based on a weighted Compliance Metric that admits only clients meeting legal, contractual, and technical requirements especially in scenarios involving sensitive data such as *PII* (Personally Identifiable Information) and *PHI* (Protected Health Information). The selection process is governed by a policy-defined weight vector, applied by the Compliance Module, which retrieves compliance evidence from cloud providers via standardized APIs and logs all decisions to ensure auditability and reproducibility. ECOFED-X offers a practical solution for enabling compliant Federated Learning in distributed and regulated settings, while also contributing to environmental sustainability through its FinOps-inspired resource optimization strategy inherited from the original ECOFED framework. Future work will focus on extending the proposed framework in several directions:

- Integrating dynamic selection mechanisms based on model quality feedback and operational risk signals;
- Extending the Compliance Module to support edge providers and hybrid infrastructures (cloud and on-premise);
- Exploring multi-objective optimization strategies that balance compliance, energy efficiency, and model performance;
- Adopting interoperable standards for representing compliance policies and client profiles;
- Validating the framework in real-world use cases, particularly in healthcare and finance, with benchmarks on cost, emissions, and regulatory adherence.

These developments position ECOFED-X as a comprehensive platform for responsible, sustainable, and regulation-aware Federated Learning. By integrating policy-driven orchestration with verifiable runtime trust, ECOFED-X enables reproducible and legally compliant training across heterogeneous environments.

References

Abbas, R.; Mehta, K.; and Zhou, L. 2024. SmartHealth-FL: Federated Learning for Privacy-Preserving Clinical Analytics. In *Proceedings of the International Conference on Health Informatics*, 112–120. Springer.

Albshaier, L.; Almarri, S.; and Albuali, A. 2025. Federated Learning for Cloud and Edge Security: A Systematic Review of Challenges and AI Opportunities. *Electronics*, 14(5): 1019.

Avella, F.; and Silvestri, F. 2025. Eco-Friendly Federated Learning: Designing a FinOps Framework for Multicloud Resource Coordination. In *2025 3rd International Conference on Federated Learning Technologies and Applications (FLTA)*.

Azure, M. 2025. Azure Policy Regulatory Compliance controls for Azure API Management. https://learn.microsoft.com/en-us/azure/api-management/security-controls-policy. Accessed: 2025-10-22.

Beutel, D. J.; Topal, T.; Mathur, A.; Qiu, X.; Fernandez-Marques, J.; Gao, Y.; Sani, L.; Li, K. H.; Parcollet, T.; Gusmão, P. P. B. d.; and Lane, N. D. 2020. Flower: A Friendly Federated Learning Research Framework. *arXiv preprint arXiv:2007.14390*.

Casella, B. 2025. A Performance Analysis of VM-Based Trusted Execution Environments for Confidential Federated Learning. *arXiv preprint*.

Dakhia, Z.; and Merenda, M. 2025. Client Selection in Federated Learning on Resource-Constrained Devices: A Game Theory Approach. *Applied Sciences*, 15(13): 7556.

Gupta, A. M. Y. M. R. H. R. D. 2021. Federated Learning Across Trusted Execution Environments. In *IEEE Poster*.

HIPAA Partners Team. 2025. HIPAA Compliance in Federated Learning: Protecting Patient Privacy. https://hipaapartners.com/blog/hipaa-compliance-in-federated-learning-protecting-patient-privacy. Accessed: 2025-10-22.

Intel. 2025. Confidential Federated Learning with OpenFL. https://community.intel.com/t5/Blogs/Tech-Innovation/Artificial-Intelligence-AI/Confidential-Federated-Learning-with-OpenFL/post/1658148. Accessed: 2025-10-22.

Kaimal, P.; and Thompson, E. 2025. HIPAA-Compliant Federated Learning: A Framework for Secure Medical Collaboration. *Journal of Medical AI Systems*, 17(2): 85–97.

Katasani, D. P. 2025. Federated Learning with AI Agents Across Multi-Cloud Data Platforms. *European Modern Studies Journal*, 9(5).

Kovačević, A. 2025. Federated Learning in Healthcare: Improving Collaboration and Privacy. In *International Conference on Marketing and Smart Technologies (ICMarkTech)*, 463–475. Springer.

Kumar, A.; Singh, R.; and Patel, M. 2025. Fortifying Federated Learning Towards Trustworthiness via Auditable Data Valuation and Verifiable Participation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Liu, Y.; Zhang, W.; and Rao, A. 2023. Auditable and Verifiable Federated Learning Based on Blockchain-Enabled Architecture. *IEEE Transactions on Services Computing*, 16(4): 1021–1034.

Moran, O. 2025. Federated Learning & Global Data Sovereignty Compliance. https://dualitytech.com/blog/federated-learning-in-meeting-global-data-sovereignty-regulations. Accessed: 2025-10-22.

Platform, G. C. 2025. GCP Resource Manager Compliance Metadata Overview. https://cloud.google.com/resource-manager/docs. Accessed: 2025-10-22.

Services, A. W. 2025. Viewing Compliance Information and Evaluation Results for your AWS Resources with AWS Config. https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config_view-compliance.html. Accessed: 2025-10-22.

Singh, S.; Verma, S. B.; Gupta, B. K.; and Dubey, S. 2025. Federated Learning for GDPR Compliance. In *Hybrid Intelligent Systems*, 383–392. Springer.

Tahir, M.; and Ali, M. I. 2025. Client Selection in Federated Learning: Challenges, Strategies, and Contextual Considerations. In *Federated Learning Systems*, 43–62. Springer.

Wang, X.; Peng, Y.; Lu, L.; Lu, Z.; Bagheri, M.; and Summers, R. M. 2017. ChestX-ray8: Hospital-scale Chest X-ray Database and Benchmarks on Weakly-Supervised Classification and Localization of Common Thorax Diseases. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3462–3471.

Yadav, S.; Veeramanickam, M.; and Kumar, R. 2025. Federated Learning for Privacy-Preserving Medical Image Analysis Across Global Healthcare Networks. In *International Conference on Data Analytics and Management*, 388–398. Springer.

Yan, J.; Li, Y.; Yin, S.; Kang, X.; Wang, J.; Zhang, H.; and Hu, B. 2024. An Efficient Greedy Hierarchical Federated Learning Training Method Based on Trusted Execution Environments. *Electronics*, 13(17): 3548.

Zhang, M.; and Chen, R. 2024. FedLC: Federated Learning-Based Legal Compliance Detection in Finance. *IEEE Transactions on Knowledge and Data Engineering*, 36(2): 456–468.

Zhang, Y.; Wang, Y.; Liu, J.; and Xu, H. 2022. FedPPO: Reinforcement Learning-Based Client Selection for Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems*.

Zhou, Y.; Wang, J.; Li, H.; and Chen, M. 2023. Client Selection in Federated Learning: Principles, Challenges, and Future Directions. *IEEE Transactions on Neural Networks and Learning Systems*.