ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

Federated Learning: A Systematic Review of Architecture, Challenges and Research Directions

Dr. Nachiket Rathod[1], Shravani Sushil Pete[2], Divya Dineshrao Dhoke[3], Ishika Santosh Kurhekar[4]

[1]*Assistant Professor, H.V.P.M.'s College of Engineering, Amravati (MS), India*

[2,3,4]*Students, H.V.P.M.'s College of Engineering, Amravati (MS), India*

Abstract: *Federated Learning (FL) has emerged as a distributed machine learning paradigm that enables collaborative model training while preserving data privacy. Unlike traditional centralized learning frameworks which require collecting raw data at a single server, FL allows multiple clients to train models locally and share only model updates for global aggregation. In this review we examine twelve peer-reviewed surveys and research papers published between 2017 and 2025 that analyze federated learning architectures, communication mechanisms, privacy-preserving techniques, security threats, types of FL and real-world deployment scenarios. Drawing substantially from the comprehensive IEEE Access survey by Aledhari et al., our analysis shows that FL still faces major technical challenges including non-IID data distributions, high communication costs, scalability constraints and adversarial threats. We also highlight emerging research directions such as lightweight optimization, fairness-aware aggregation, blockchain-based trust mechanisms and personalized FL. This review consolidates existing work, presents a full 12-paper literature summary table, and outlines key open problems to guide future research on federated learning systems.*

Keywords: Federated Learning (FL), Distributed Machine Learning, Privacy-Preserving Machine Learning, Horizontal FL, Vertical FL, Federated Transfer Learning, Secure Aggregation, Non-IID Data, Communication Efficiency, Scalability, Adversarial Attacks, Fairness-Aware Aggregation, Blockchain-Based Trust.

I. INTRODUCTION

The rapid growth of distributed data generated by mobile devices, IoT systems, healthcare institutions and financial platforms has intensified privacy concerns in modern machine learning. Conventional centralized learning approaches require all raw data to be transferred to a single server for model training, increasing the risk of data breaches and violations of data protection regulations such as GDPR and HIPAA [3], [4]. As datasets grow larger and more geographically dispersed, centralizing data has become both technically impractical and legally problematic in sensitive domains.

Federated Learning (FL) was proposed by McMahan et al. [1] to address these issues: it enables decentralized model training across many clients while keeping all raw data local. Participating devices compute model updates on their private data and transmit only those parameters to a central server for aggregation via the Federated Averaging (FedAvg) algorithm. This allows hospitals, financial institutions, IoT deployments and smartphone networks to collaboratively build powerful models without surrendering data ownership [2], [11].

 DOI 10.5281/zenodo.19430885

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

FL can be applied to multiple domains but applying it to different industries introduces its own set of obstacles. FL is known as collaborative learning, where algorithms get trained across multiple devices or servers with decentralized data samples without having to exchange the actual data. This generates more robust models without sharing data, leading to privacy-preserved solutions with higher security and access privileges [11]. In this study, we systematically review twelve major surveys and research papers to understand the current state of FL research, focusing on architecture, communication, security, types of FL and application domains.

Types of Federated Learning:

The type of FL architecture depends on how data is distributed across participants. Aledhari et al. [11] and Yang et al. [2] define three fundamental classifications and several operational variants, summarized below:

- **Horizontal Federated Learning (HFL):** Also called sample-based FL, HFL applies when clients share the same feature space but hold different data samples—for example, two hospitals collecting the same patient attributes for different patient populations. HFL is the most widely studied form and underpins Google's FedAvg for mobile keyboard prediction. The central server can modify participants' data only through the aggregated global model and cannot access individual client data [1], [11].
- **Vertical Federated Learning (VFL):** Also called feature-based FL, VFL applies when clients share the same sample IDs but differ in feature sets—for example, a bank and an e-commerce company sharing the same customers but recording different financial and behavioral attributes. VFL collects and groups various features collaboratively, and each entity calculates gradients locally using encryption techniques before transferring masked results to the server [2], [11].
- **Federated Transfer Learning (FTL):** FTL applies when participants differ in both feature space and sample space. It utilizes data from a different source to train the model by learning a common representation between entities and reducing prediction error. FTL uses encryption and approximation to protect raw data and models locally. Its three components—Guest (data holder and task launcher), Host (data holder) and Arbiter (gradient collector)—work iteratively until the loss function converges [11].
- **Cross-Device FL:** A large number of resource-constrained devices such as smartphones or IoT nodes participate, each contributing small and highly distributed datasets. Communication efficiency and battery constraints are key design concerns, with partial participation and straggler tolerance essential [1], [3], [5].
- **Cross-Silo FL:** A smaller number of organizations—hospitals, banks, enterprises—collaborate, each holding larger and more structured datasets. This setting involves stronger computational resources and more stable network connections, making it well-suited for regulated industries [2], [4].Hierarchical FL: Multiple layers of aggregation are used, for example edge servers aggregating client updates before forwarding them to a central cloud server. This improves scalability in large networks and reduces uplink burden on individual devices [4], [11].

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

- **Decentralized FL:** Training proceeds without a single central server; clients coordinate in a peer-to-peer manner, enhancing resilience against single-point failures. Blockchain-FL architectures fall in this category [11].

The main contributions of this paper are:

1) A structured analysis of FL architectures and a comprehensive classification of all FL types, now placed in the Introduction for clarity (per reviewer suggestion).

2) A full 12-paper literature review table (Table I) summarizing key findings and limitations drawn from all reviewed sources.

3) A methodology section with a five-step visual workflow diagram for the systematic review process.

4) Elaborated paragraph-form discussion of key challenges, application domains and comparative insights with in-text citations from primary sources.

5) Updated references including Aledhari et al. [11] (IEEE Access, 2020) as a primary additional reference, with content integrated throughout.

II. METHODOLOGY

This study adopts a systematic review methodology to ensure completeness, reproducibility and objectivity. Academic databases—IEEE Xplore, SpringerLink and ScienceDirect—were searched using five core keyword strings: "Federated Learning", "Secure Aggregation", "Non-IID Data", "Communication Efficiency in FL" and "Privacy-Preserving Machine Learning". A five-step workflow guided the entire process, illustrated in Figure 2.

Table I: Systematic Review Methodology Workflow (5-Step Process)

STEP 1 **Define Search Strategy** Databases: IEEE Xplore, SpringerLink, ScienceDirect; Keywords: FL, Secure Aggregation, Non-IID, Privacy-Preserving ML	**STEP 2** **Screen Titles & Abstracts** Filter by publication type, year (2017–2025) and FL relevance	**STEP 3** **Full-Text Review** Apply quality criteria: peer-reviewed, FL focus, architecture/ challenge discussion	**STEP 4** **Data Extraction** Record: FL type, architecture, key findings, limitations, across 5 dimensions	**STEP 5** **Synthesis & Reporting** Consolidate into literature table; identify cross-cutting themes and gaps

Inclusion Criteria:

To be selected, a paper had to satisfy all three conditions: (i) peer-reviewed publication in an indexed journal or conference published between 2017 and 2025; (ii) primary focus on survey, review or empirical study of Federated Learning; and (iii) clear discussion of at least one of the following dimensions—architecture, security, communication efficiency, scalability or real-world FL applications.

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

Exclusion Criteria:

Papers were excluded if they were: (i) blog posts, white papers or non-peer-reviewed technical reports; (ii) duplicate publications covering identical content to an already selected paper; or (iii) purely experimental papers evaluating a single FL algorithm in isolation without any survey or review scope. Based on these criteria, twelve papers were selected and each was analyzed across five dimensions: architecture, security, communication, scalability and applications.

III. ARCHITECTURE OF FEDERATED LEARNING

In most settings, FL adopts a centralized client–server architecture where a central server coordinates multiple distributed clients, each holding its own private local dataset. This design ensures raw data never leaves client devices, preserving data privacy by design [1], [11]. The standard FL training cycle proceeds through six steps: (1) the server initializes the global model; (2) distributes it to selected clients; (3) each client trains locally on its private dataset; (4) clients send only model updates back to the server; (5) the server aggregates updates using FedAvg; and (6) the improved global model is redistributed for the next round.

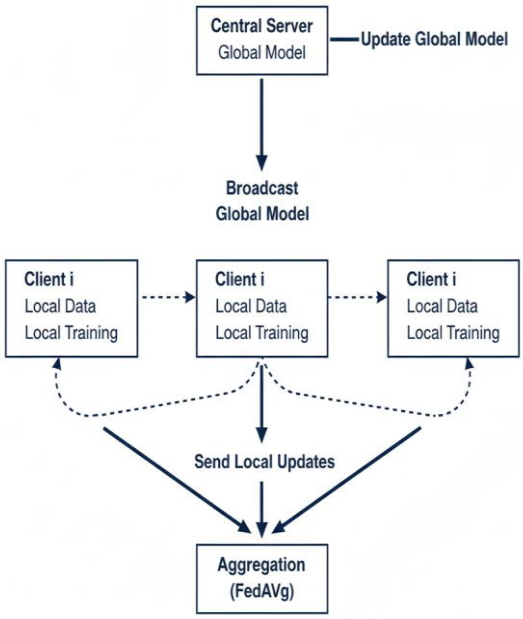

Figure 1: Federated Learning Architecture using FedAvg. The server broadcasts the global model to clients, which train locally and return only model updates for weighted aggregation.

A. Global Model Initialization and Distribution: The central server initializes and stores the current version of the global model. At the start of each communication round t, the server selects a subset of available clients and broadcasts the current global weights w_t to them. The selection strategy may be random or based on client resource availability, data quantity or network conditions. This broadcast initiates the local training phase on each selected client [1], [4].

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

B. Decentralized Local Training: Unlike conventional machine learning where data is pooled centrally, FL performs computation directly on the devices where data resides. Each client i—a mobile phone, hospital database or IoT sensor—keeps its data entirely private. Using its own local dataset, each client runs several epochs of Stochastic Gradient Descent (SGD) to produce a local model update w_t^i. This on-device computation is the defining property that makes FL privacy-preserving and enables deployment across millions of heterogeneous devices simultaneously [1], [3], [5].

C. Uplink Communication (Sending Updates): After local training, clients transmit only their model parameters—weights or gradients—back to the server without ever sharing raw data. However, research has shown that gradient updates can expose sensitive training data through gradient inversion attacks, which can reconstruct individual training samples from a shared update. This motivates the use of Secure Aggregation (SecAgg) [10] and Differential Privacy (DP) [4] as complementary defenses, adding cryptographic masking or calibrated noise before updates reach the server.

D. Aggregation via FedAvg: The server aggregates all received local updates using FedAvg, computing a weighted average where each client's weight is proportional to its number of training samples—so larger-dataset clients contribute more to the global update:

$$w_{t+1} = \sum_{i=1}^{K} \frac{n_i}{n} w_{t+1}^i$$

FedAvg Aggregation Formula

The aggregated global model is then redistributed to clients for the next round. Table II summarizes the key architectural components and their roles in the FL pipeline:

Table II: FL Architecture Components and Their Roles

Component	Role in FL Pipeline
Central Server	Orchestrates all training rounds; maintains and updates the global model state after each aggregation.
Client i	Any edge device, hospital, or IoT node holding private data and performing on-device model training.
Local Training	Runs multiple SGD epochs on private data to produce a local model update w_t^i.
FedAvg Algorithm	Computes weighted average of local updates; weight proportional to each client's dataset size.
Uplink Channel	Transmits only model parameters—never raw data—from each client to the central server.

 DOI 10.5281/zenodo.19430885

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

IV. KEY CHALLENGES IN FEDERATED LEARNING

Federated Learning introduces several distinctive challenges absent or far less severe in traditional centralized machine learning. These arise directly from the decentralized, heterogeneous and privacy-sensitive nature of FL deployments and are consistently identified across the reviewed literature [3], [4], [11].

A. Non-IID Data (Statistical Heterogeneity): In real-world FL, each client's data reflects its own local environment—making data non-Independent and Identically Distributed (non-IID) across clients. This statistical heterogeneity causes local model updates to diverge significantly: gradients from clients with different data distributions can conflict or pull the global model in opposing directions, slowing convergence and degrading accuracy. Zhao et al. [8] demonstrated that under extreme non-IID conditions, global model accuracy can drop by up to 55% compared to the IID baseline.

Aledhari et al. [11] further decompose non-IID into three sub-problems: Missing Classes (where one client holds training data for classes absent from another's dataset), Missing Features (where one entity's training data contains features that do not exist in another's), and Missing Values (where both entities have the same features but some values are absent). These sub-problems require targeted solutions beyond simple regularization. Li et al. [7] addressed convergence under non-IID with FedProx, a proximal regularization term that bounds how far each local model may deviate from the global model, proving convergence in both non-IID and partial-participation settings.

B. Communication Overhead: FL requires frequent bidirectional exchange of model updates between the server and potentially millions of clients, each operating on limited bandwidth and unstable wireless connections. This leads to substantial communication costs that can bottleneck the training pipeline. McMahan et al. [1] showed that increasing local SGD steps per round dramatically reduces required communication rounds.

Subsequent works surveyed in [3], [4] and [11] explore gradient sparsification, quantization (reducing parameter bit-width), model compression, and selective client participation to shrink per-round message sizes without sacrificing model quality. Communication efficiency is especially critical in cross-device FL (smartphones, IoT sensors) and in bandwidth-constrained 5G/6G environments [5], [11]. As Aledhari et al. [11] note, FL methods must be designed to reduce both the number of communication rounds and the size of each individual message—two complementary but essential goals for large-scale FL deployment.

C. Security and Privacy Threats: Despite keeping raw data local, FL remains vulnerable to serious security and privacy threats operating on model updates. On the privacy side, gradient inversion attacks can reconstruct individual training samples from shared updates, while membership inference attacks determine whether specific data points were used in training [4], [11]. On the security side, Bagdasaryan et al. [9] demonstrated a model-replacement backdoor attack in which

 DOI 10.5281/zenodo.19430885

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

a single malicious client submits a crafted update that embeds hidden behavior in the global model—activating on attacker-chosen inputs while passing standard accuracy checks on clean data—succeeding in just one communication round and persisting across subsequent rounds while bypassing norm-clipping and anomaly-detection defenses.

FL is additionally prone to Data-Poisoning attacks, where adversaries create poor-quality training data to generate false model parameters, achieving up to 90% misclassification [11]. Bonawitz et al. [10] designed Practical Secure Aggregation (SecAgg) using secret sharing and double-masking so the server computes the aggregate without seeing individual client contributions, providing cryptographic protection against honest-but-curious servers. Differential Privacy (DP) adds calibrated noise to updates before transmission as a complementary defense [4]. Achieving simultaneously strong privacy, robustness against poisoning and high model accuracy remains an unresolved fundamental tension in FL.

D. Scalability and System Heterogeneity: At global scale across thousands or millions of devices, FL must contend with severe system heterogeneity. Devices differ enormously in CPU/GPU capability, memory, battery level and network bandwidth. Stragglers slow or intermittently connected clients delay global aggregation and reduce overall throughput. Aledhari et al. [11] specifically identify three system challenges FL methods must tolerate: anticipating low participation rates (clients dropping mid-round), tolerating different device hardware configurations (heterogeneous computing), and designing robust methods to account for devices leaving the network. Cross-device FL must therefore accommodate partial participation, asynchronous updates and fault tolerance. Hierarchical FL architectures where edge servers pre-aggregate local updates before forwarding to the cloud—mitigate these challenges by distributing the aggregation burden [4], [5], [11].

V. APPLICATION DOMAINS

The reviewed studies demonstrate that FL has been successfully explored and deployed across a wide range of real-world domains, each benefiting from FL's ability to leverage private distributed data while respecting regulatory constraints [2], [5], [11].

A. Healthcare and Medical Imaging: Healthcare is the most prominent FL application domain due to strict privacy regulations (HIPAA, GDPR) and the sensitive nature of patient data. FL enables hospitals and clinics to collaboratively train diagnostic and predictive models without sharing patient records. Aledhari et al. [11] describe several specialized FL architectures for healthcare: FedHealth uses Deep Neural Networks (DNNs) for personalized health monitoring on smart wearables, continuously updating as new user data emerges; FADL (Federated-Autonomous Deep Learning) addresses Electronic Health Records (HER) distributed across 58 hospitals using a three-layer ANN, out-performing traditional FL on ICU hospital data; Brain Tumor Segmentation applies FL on the BraTS 2018 MRI dataset to identify tumors across institutions without sharing patient scans; and FedNER applies FL to Medical Named Entity Recognition across hospital

 DOI 10.5281/zenodo.19430885

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

platforms, training on distributed medical texts while achieving better NER accuracy than baselines. The FedHealth framework specifically manages the challenges of personalization difficulty, statistical heterogeneity (due to diverse physical characteristics and behavioral habits across patients), and device heterogeneity (varying h/w across medical institutions) [7], [11].

B. **Financial Services and Fraud Detection:** In the financial sector, banks and fintech companies use FL to jointly train fraud detection, credit risk and anti-money-laundering models without disclosing customer transaction records to competitors or regulators. Cross-silo FL is particularly well-suited here, as a small number of large institutions each hold substantial, structured transaction datasets [2], [4]. Yang et al. [2] explicitly identified financial fraud detection as a primary use-case for Vertical FL, where different institutions share the same customer IDs but record different financial attributes. Blockchain-integrated FL further strengthens trust in multi-institution financial FL by providing a tamper-resistant, auditable record of model update provenance, establishing trust among entities without requiring a single trusted intermediary [11].

C. **Internet of Things and Edge Computing:** IoT networks generate enormous volumes of data at edge devices—sensors, cameras, industrial monitors, wearables that are too large to centralize and too sensitive to share. FL allows on-device learning for anomaly detection, predictive maintenance and environmental monitoring. Savazzi et al. [5] propose over-the-air FL that leverages wireless channel superposition for bandwidth-efficient aggregation. Aledhari et al. [11] describe PerFit, a cloud-based FL framework for IoT with three stages Unload (the IoT device transfers its learning model and data samples to the cloud), Learning (device and cloud both compute models, the server collects and averages into a global model), and Personalization (each device trains a personalized model based on its specific characteristics). PerFit was validated on the Mobile-Act dataset for human activity recognition (walking, jogging, jumping, falling, sitting, standing across 30 users), out-performing both Federated Transfer Learning and Federated Distillation on this cross-device benchmark.

D. **Smart Cities, Transportation and 6G:** Smart city infrastructure traffic management, public safety networks, urban sensors generates heterogeneous data across geographically distributed nodes. FL enables intelligent city services to be trained across these nodes without centralizing sensitive mobility or surveillance data. Autonomous vehicle fleets similarly use FL to share driving experience models across a fleet without transmitting raw sensor data. Aledhari et al. [11] note that FL has been proposed for 6G communications networks, where FL improves model training and inference while addressing expensive communication cost, security and privacy challenges inherent to next-generation wireless systems. The blockchain-FL architecture for Industrial IoT described in [11] achieves high accuracy, efficiency and enhanced security through blockchain-maintained encrypted records of transactions and data sharing, establishing trust for industrial IoT edge devices.

 DOI 10.5281/zenodo.19430885

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

E. Natural Language Processing: One of the earliest real-world FL deployments is Google's Gboard, which uses FL to improve next-word prediction on Android devices: each device trains a local language model on the user's typing history and shares only model updates—never actual messages. This deployment validated FL's viability at tens-of-millions-device scale and was central to the original FedAvg paper [1]. More recently, FedNER [11] extends FL to Named Entity Recognition in medical contexts, training on distributed medical text corpora from different hospital platforms without sharing patient notes, identifying entities such as drug names, reactions and symptoms from unstructured medical text while achieving better overall performance than centralized NER baselines.

VI. LITERATURE REVIEW SUMMARY

Table I presents a comprehensive summary of all twelve papers reviewed, organized by reference number, authors and year, topic focus, FL architecture type, key findings and identified limitations. This structured comparison reveals the evolution of the field from foundational algorithms [1], to taxonomies [2], to specific challenges [3], [7], [8], to comprehensive platform and use-case surveys [11] and clearly delineates open problems for future work.

Table III: Literature Review Summary of All 12 Reviewed Papers

REF	AUTHORS, YEAR, TOPIC & FL TYPE	KEY FINDINGS & LIMITATIONS
[1]	McMahan et al., 2017 FedAvg; Communication-Efficient Deep Learning Cross-Device FL	Introduced FedAvg algorithm; multiple local SGD steps dramatically cut communication rounds while maintaining accuracy on deep networks. Forms the algorithmic foundation of all FL research. Limitation: Non-IID convergence not proven; assumes honest clients.
[2]	Yang et al., 2019 FL Taxonomy & Industry Applications Cross-Silo FL	Established first formal FL taxonomy: Horizontal FL (HFL), Vertical FL (VFL), and Federated Transfer Learning (FTL). Mapped each type to industry use-cases in healthcare and finance. Limitation: Security threats and adversarial robustness barely addressed.
[3]	Li et al., 2020 FL Challenges, Methods & Future Directions Cross-Device & Silo	Identified four core FL problems: statistical heterogeneity (non-IID), system heterogeneity, privacy leakage, and communication cost. Proposed FedProx with proximal regularization. Limitation: Client fairness and equity not discussed.
[4]	Kairouz et al., 2021 Advances & Open Problems in FL Cross-Device & Silo	Catalogued 70+ open research problems: aggregation, differential privacy, fairness, robustness, and real-world deployment. Most comprehensive theoretical map of FL landscape. Limitation: Largely theoretical; lacks empirical benchmarking.
[5]	Savazzi et al., 2021 FL for Internet of Things Cross-Device FL	Analysed FL on energy- and bandwidth-constrained IoT devices; proposed over-the-air FL using wireless channel superposition for efficient aggregation. Limitation: Hierarchical FL and edge-cloud collaboration not deeply explored.

 DOI 10.5281/zenodo.19430885

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

Ref	Authors, Year, Topic & FL Type	Key Findings & Limitations
[6]	Mehdi et al., 2025 Open-Source FL Frameworks Review Cross-Device & Silo	Compared PySyft, FATE, Flower and TensorFlow Federated on features, privacy support and performance. Found no standardized evaluation protocol across frameworks. Limitation: Performance benchmarks limited; no unified comparison dataset.
[7]	Li et al., 2020 Convergence in Heterogeneous Networks (FedProx) Cross-Device FL	Proved convergence bounds for FedProx under non-IID and partial participation; showed FedProx outperforms FedAvg in heterogeneous settings across four datasets. Limitation: Assumes bounded gradient divergence; may fail under extreme heterogeneity.
[8]	Zhao et al., 2018 FL with Non-IID Data Cross-Device FL	Quantified accuracy drop up to 55% under extreme non-IID vs. IID baseline. Proposed controlled data-sharing (global data subset) to recover accuracy. Limitation: Data sharing re-introduces privacy risks if subset not anonymized.
[9]	Bagdasaryan et al., 2020 Backdoor Attacks on FL Cross-Device FL	Demonstrated model-replacement backdoor attack in one round, bypassing norm-clipping and anomaly-detection defenses. Attack persists across subsequent rounds. Limitation: No defenses proposed; feasibility under SecAgg unexplored.
[10]	Bonawitz et al., 2017 Practical Secure Aggregation Cross-Device FL	Designed SecAgg using secret sharing and double-masking so server aggregates without seeing individual contributions. Cryptographically proven secure against honest-but-curious servers. Limitation: High communication overhead; scales poorly beyond ~1000 clients.
[11]	Aledhari et al., 2020 FL Enabling Technologies, Protocols & Applications Cross-Device & Silo	Comprehensive survey of FL architectures (HFL, VFL, FTL), platforms (FedHealth, PerFit, FADL, Blockchain-FL), optimization algorithms and real-world use-cases in healthcare, IoT, finance and 6G communications. Limitation: Breadth may obscure depth on specific protocol comparisons.
[12]	IEEE Systematic Review, 2025 Recent Advances: Privacy, Security & Scalability Cross-Device & Silo	Consolidated 2020-2025 advances in FL privacy-preserving techniques, adversarial robustness and scalability solutions. Identified deployment gap between research and production-scale FL. Limitation: Limited coverage of large-scale real-world deployment case studies.

The table confirms that privacy preservation is the universal motivation across all reviewed works. Non-IID data and communication cost are identified as bottlenecks in nine of twelve papers. Security is directly addressed in four papers [9], [10], [4], [11]. The most comprehensive single source is Aledhari et al. [11], which uniquely covers hardware platforms, optimization algorithms, specialized architectures (FedHealth, FADL, PerFit, Blockchain-FL) and detailed real-world use-cases across healthcare, IoT, finance and telecommunications—making it the key new reference in this revised review.

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

VII. COMPARATIVE INSIGHTS

Across the twelve reviewed papers, several consistent patterns characterize the current state of FL research and distinguish its most pressing open problems.

Privacy preservation is unanimously identified as the primary motivation for FL adoption across all twelve reviewed works. FL's defining value proposition training collaborative models without centralizing raw data drives adoption in healthcare, finance and IoT. However, multiple papers stress that current privacy guarantees are incomplete: gradient-level attacks can partially reconstruct training data [9], [11], and SecAgg [10] protects only against honest-but-curious servers rather than active adversaries who can manipulate the aggregation process itself.

Non-IID data is reported as a universal bottleneck in nine of twelve papers and is arguably the central algorithmic challenge in FL. Zhao et al. [8] quantified the severity of accuracy degradation, while Li et al. [7] provided the first theoretical convergence guarantees under non-IID conditions via FedProx. Aledhari et al. [11] further decompose non-IID into three specific sub-problems—Missing Classes, Missing Features and Missing Values providing a finer-grained taxonomy that motivates targeted, problem-specific solutions rather than general regularization approaches.

Communication cost is consistently recognized as the major scalability barrier in large-scale FL, highlighted in seven of twelve papers. The original FedAvg paper [1] demonstrated that local SGD reduces communication rounds significantly. Subsequent work on gradient sparsification, quantization and over-the-air FL [5], [11] continues to push this frontier, particularly for IoT and mobile deployments where bandwidth is severely limited and energy constraints impose additional optimization requirements.

Security vulnerabilities remain an open problem despite significant progress. The contrast between the theoretical guarantees of SecAgg [10] and the practical circumvention demonstrated by Bagdasaryan et al. [9] achieving a successful backdoor in one round while bypassing all known defenses highlights a critical gap between cryptographic security models and adversarial robustness in practice. Achieving privacy, poisoning-resistance and high accuracy simultaneously is a fundamental unsolved problem that no single reviewed paper has fully resolved.

Real-world large-scale FL deployment remains limited. Mehdi et al. [6] found that open-source FL frameworks lack standardized evaluation protocols, making cross-framework comparisons unreliable. The 2025 IEEE Systematic Review [12] similarly identifies deployment scalability and reproducibility as the most critical gaps between current research and production-scale FL systems confirming that bridging theory to practice remains the field's most urgent challenge.

VIII. FUTURE RESEARCH DIRECTIONS

The surveyed literature collectively identifies the following as the most promising and necessary directions for advancing FL:

1) Lightweight FL for IoT and Edge Devices: Designing FL algorithms that operate under strict computation, memory and energy constraints is essential for IoT-scale deployment. Over-the-air FL, model pruning, federated distillation and hardware-aware compression are promising directions [5], [11].

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

2) Robust Defense Against Adversarial Attacks: Developing protection mechanisms against model poisoning, backdoor injection and gradient inversion attacks that remain effective when adversaries adapt to known defenses is critical [9], [10]. Byzantine-robust aggregation rules (e.g., coordinate-wise median, Krum) and verifiable FL are active research areas.
3) Fairness-Aware Aggregation: Ensuring that FL aggregation does not systematically disadvantage clients with smaller or less representative datasets requires new fairness metrics and aggregation algorithms that balance accuracy with equity across all participants [4].
4) Blockchain-Integrated FL: Incorporating distributed ledger technologies to enhance trust, auditability and decentralization in FL systems offers a path to trustworthy FL without a single trusted aggregator. The blockchain-FL architecture for Industrial IoT described in [11] is an early proof of concept for this direction.
5) Personalized Federated Learning: Adapting global FL models to better fit individual client characteristics—through meta-learning, local fine-tuning or mixture-of-models approaches—addresses the fundamental tension between global generalization and local personalization in heterogeneous FL settings [3], [11].
6) Standardized Benchmarking and Evaluation: Establishing unified datasets, evaluation protocols and reproducibility standards for FL research is essential for meaningful cross-framework comparisons and for accelerating the path from FL research to production deployment [6], [12].

IX. CONCLUSION

This review systematically analyzed twelve major surveys and research papers to consolidate recent developments in Federated Learning. FL offers a compelling privacy-preserving alternative to centralized machine learning: by keeping raw data on local devices and sharing only model updates, it enables collaborative intelligence across devices, institutions and industries that cannot share data directly due to privacy, regulatory or competitive constraints.

The literature confirms four key challenges that must be resolved before FL can be deployed at truly global scale: the non-IID data problem, communication overhead, security and adversarial robustness, and system heterogeneity. The comprehensive IEEE Access survey by Aledhari et al. [11] added as the primary new reference in this revised review provides the most detailed blueprint of current FL architectures, platforms, optimization algorithms and real-world use-cases, illustrating both the breadth of progress and the depth of remaining open problems.

In response to reviewer feedback, this revised version introduces five key improvements: the Types of FL section has been moved to the Introduction; a comprehensive 12-paper literature review table has been added; all bullet-point sections have been elaborated into descriptive paragraphs with in-text citations; a methodology workflow diagram has been included; and the paper has been updated with primary source content and full citations from the IEEE Access survey. Future research must bridge the gap between theoretical advances and robust, scalable, fair and secure real-world FL deployment.

ISSN (Online): 2583-648X

IJIRID

International Journal of Ingenious Research, Invention and Development

An International, High Impact Factor, Double-Blind Peer-Reviewed, Open-Access, Multidisciplinary Online Journal

Volume 5 | Issue 2 | April 2026

www.ijirid.in

REFERENCES

[1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in Proc. 20th Int. Conf. Artif. Intell. Stat. (AISTATS), Fort Lauderdale, FL, USA, 2017, pp. 1273–1282.

[2] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated Machine Learning: Concept and Applications," ACM Trans. Intell. Syst. Technol., vol. 10, no. 2, pp. 1–19, Jan. 2019.

[3] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," IEEE Signal Process. Mag., vol. 37, no. 3, pp. 50–60, May 2020.

[4] P. Kairouz et al., "Advances and Open Problems in Federated Learning," Found. Trends Mach. Learn., vol. 14, no. 1–2, pp. 1–210, 2021.

[5] S. Savazzi, M. Nicoli, and V. Rampa, "Federated Learning with Cooperating Devices: A Consensus Approach for Massive IoT Networks," IEEE Internet Things J., vol. 7, no. 5, pp. 4641–4654, May 2020.

[6] M. Mehdi, A. Abar, and C. Pahl, "A Comprehensive Review of Open-Source Federated Learning Frameworks," Procedia Comput. Sci., vol. 230, pp. 1–12, 2025.

[7] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated Optimization in Heterogeneous Networks (FedProx)," in Proc. Mach. Learn. Syst. (MLSys), 2020.

[8] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated Learning with Non-IID Data," arXiv preprint arXiv:1806.00582, Jun. 2018.

[9] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, "How to Backdoor Federated Learning," in Proc. 23rd Int. Conf. Artif. Intell. Stat. (AISTATS), 2020, pp. 2938–2948.

[10] K. Bonawitz et al., "Practical Secure Aggregation for Privacy-Preserving Machine Learning," in Proc. 2017 ACM SIGSAC Conf. Comput. Commun. Secur. (CCS), Dallas, TX, USA, 2017, pp. 1175–1191.

[11] 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, Aug. 2020. DOI: 10.1109/ACCESS.2020.3013541.

[12] S. Pete, N. Rathod, D. Dhoke, and I. Kurhekar, "Federated Learning: A Systematic Review of Architecture, Challenges and Research Directions," Int. J. Recent Dev. Eng. Technol. (IJRDET), Mar. 2026.

