Journal Homepage: www.ijarpr.com ISSN: 3049-0103 (Online)

International Journal of Advance Research Publication and Reviews

Vol 03, Issue 8, pp 135-146, August, 2026

Influence of Scalable Confidential Computing Techniques on Performance Overhead and Privacy Guarantees in Federated Learning Systems

Grace Kimani[1]*, James Mwikya[2]

[1]Department of Computer Science and Informatics, Kirinyaga University, Kerugoya, Kenya
[2]School of Pure and Applied Sciences, Kirinyaga University, Kerugoya, Kenya
*Corresponding author: gkimani@kyu.ac.ke | Co-authors: jmwikya@kyu.ac.ke

ABSTRACT :

Federated Learning (FL) has become an important distributed machine learning paradigm that enables collaborative model training while preserving data privacy by keeping data localized. Despite its advantages, FL remains vulnerable to privacy attacks such as model inversion, gradient leakage, and membership inference attacks, particularly in untrusted multi-party environments. Confidential computing techniques including Trusted Execution Environments (TEEs), Secure Multi-Party Computation (SMPC), Homomorphic Encryption (HE), and Confidential Virtual Machines (CVMs) have been proposed to strengthen privacy protection. However, these techniques often introduce significant computational and communication overhead, limiting their adoption in resource-constrained edge environments. This desktop review synthesizes recent literature published between 2019 and 2026 to examine the influence of scalable confidential computing techniques on performance overhead and privacy guarantees in federated learning systems. Literature was collected from Scopus, Web of Science, IEEE Xplore, ACM Digital Library, SpringerLink, ScienceDirect, Google Scholar, and relevant grey literature using predefined inclusion and exclusion criteria. A thematic synthesis approach was employed to identify recurring patterns and research gaps. The review reveals that scalable confidential computing mechanisms can significantly improve privacy protection while reducing computational overhead through hardware-assisted security, optimized cryptographic protocols, and hybrid architectures. Nevertheless, challenges remain regarding scalability, interoperability, and deployment in heterogeneous edge computing environments. The review concludes by proposing future research directions aimed at developing adaptive confidential computing frameworks that achieve stronger privacy guarantees with minimal performance degradation.

Keywords: Federated Learning, Confidential Computing, Trusted Execution Environments, Privacy Preservation, Performance Overhead and Edge Computing

INTRODUCTION

1.1 Background

The rapid advancement of Artificial Intelligence (AI), Machine Learning (ML), cloud computing, and edge computing has fundamentally transformed the way organizations collect, process, and utilize data for intelligent decision-making. Across sectors such as healthcare, finance, manufacturing, education, transportation, telecommunications, and cybersecurity, AI-driven systems increasingly rely on large volumes of distributed data to improve predictive accuracy, automate operations, and enhance organizational efficiency. The proliferation of Internet of Things (IoT) devices, mobile technologies, and intelligent sensors has further accelerated the generation of decentralized data, creating unprecedented opportunities for collaborative machine learning. However, this rapid digital transformation has simultaneously intensified concerns regarding data privacy, confidentiality, cybersecurity, and data governance, particularly as organizations process highly sensitive information across geographically distributed environments (Li et al., 2024; Akhtarshenas et al., 2024).

Globally, the growing dependence on digital technologies has been accompanied by an unprecedented increase in cyber threats and data breaches. According to IBM's *Cost of a Data Breach Report 2024*, the average global cost of a data breach reached USD 4.88 million, the highest figure since the report's inception, with breaches involving public cloud environments and shadow data significantly increasing organizational costs. Similarly, governments worldwide have introduced stringent data protection regulations, such as the European Union's General Data Protection Regulation (GDPR), while many countries have enacted comparable privacy legislation to strengthen the protection of personal information. These developments underscore the urgent need for secure and privacy-preserving machine learning architectures capable of supporting collaborative artificial intelligence without compromising sensitive data.

Traditionally, machine learning has relied on centralized architectures in which data generated by multiple users or organizations are transferred to a central repository for model training. Although centralized learning facilitates efficient aggregation of large datasets and often produces highly accurate predictive models, it also exposes sensitive information to significant security and privacy risks. Centralized repositories become attractive targets for cyberattacks, insider threats, unauthorized access, and accidental data disclosure. These vulnerabilities have become increasingly critical as organizations process confidential information such as electronic health records, financial transactions, biometric data, industrial control information,

and government records (Zhang et al., 2024; Guembe et al., 2024).

To overcome these limitations, researchers have increasingly explored decentralized machine learning paradigms that enable collaborative model development without requiring organizations to exchange raw datasets. Among these approaches, Federated Learning (FL) has emerged as one of the most influential privacy-preserving machine learning techniques. Federated learning enables multiple organizations or devices to collaboratively train a shared machine learning model while retaining data locally. Instead of transmitting raw datasets to a centralized server, participants exchange only model parameters or gradients for aggregation, thereby substantially reducing the risks associated with centralized data collection (Akhtarshenas et al., 2024; Zhang et al., 2024).

Consequently, federated learning has evolved from a research concept into an operational technology widely adopted across healthcare, financial services, autonomous transportation, smart manufacturing, mobile computing, and Internet of Things applications. Recent advances in federated optimization algorithms, communication-efficient learning techniques, and distributed model aggregation have significantly enhanced its scalability across heterogeneous cloud and edge computing infrastructures while supporting compliance with increasingly stringent data protection regulations (Akhtarshenas et al., 2024; Guo et al., 2024).

Despite these advantages, federated learning does not inherently guarantee complete privacy or security. Although local datasets remain on participating devices, the iterative exchange of model parameters introduces new attack surfaces through which malicious participants or compromised aggregation servers may infer sensitive information. Recent studies have demonstrated the feasibility of membership inference attacks, model inversion attacks, gradient leakage attacks, property inference attacks, and data reconstruction attacks capable of exposing confidential information about individual participants (Guembe et al., 2024; Zhang et al., 2024). These findings highlight that federated learning alone is insufficient to provide comprehensive privacy protection.

To address these vulnerabilities, researchers have proposed complementary privacy-preserving technologies including Differential Privacy, Secure Multi-Party Computation (SMPC), Homomorphic Encryption (HE), Blockchain, Secure Aggregation, and Confidential Computing. Among these approaches, confidential computing has recently emerged as one of the most promising solutions because it protects sensitive information not only during storage and transmission but also while it is actively being processed. By combining hardware-assisted security with trusted software architectures, confidential computing creates isolated execution environments that safeguard sensitive computations even from privileged system software and cloud administrators (Guo et al., 2024).

Confidential computing primarily relies on Trusted Execution Environments (TEEs), including Intel Software Guard Extensions (Intel SGX), AMD Secure Encrypted Virtualization (AMD SEV), ARM TrustZone, and Confidential Virtual Machines deployed within modern cloud infrastructures. These technologies establish hardware-protected execution environments in which application code and sensitive data remain encrypted and inaccessible to operating systems, hypervisors, cloud administrators, and unauthorized processes. Consequently, confidential computing has become an increasingly important component of secure cloud computing, distributed artificial intelligence, and privacy-preserving analytics (Guo et al., 2024; Cheng et al., 2024).

However, while confidential computing substantially strengthens privacy guarantees, its deployment introduces important computational challenges. Hardware-assisted isolation mechanisms require additional cryptographic operations, secure memory management, remote attestation procedures, and trusted communication protocols that inevitably increase computational complexity. Likewise, techniques such as Secure Multi-Party Computation and Homomorphic Encryption involve multiple rounds of communication and computationally intensive mathematical operations, leading to increased training latency, communication overhead, processor utilization, and energy consumption (Li et al., 2024; Akhtarshenas et al., 2024). As a result, balancing privacy protection with computational efficiency has become one of the most significant challenges in distributed machine learning.

At the regional level, developing economies are experiencing rapid digital transformation characterized by expanding internet connectivity, widespread smartphone adoption, growth in mobile financial services, and increasing investment in artificial intelligence and cloud computing. Across Africa, AI applications are increasingly being deployed in healthcare, agriculture, education, financial services, environmental monitoring, and public administration, generating vast amounts of sensitive distributed data. Nevertheless, many organizations continue to operate within resource-constrained computing environments characterized by limited processing power, inadequate cybersecurity infrastructure, fragmented data governance frameworks, and heterogeneous hardware resources. Consequently, computationally intensive privacy-preserving mechanisms may introduce significant performance overhead that limits their practical deployment across the region (Li et al., 2024; Guembe et al., 2024).

The rapid expansion of edge computing further amplifies these challenges. Modern intelligent applications increasingly perform computation near data sources to reduce latency, minimize bandwidth consumption, and support real-time decision-making. However, edge devicesincluding smartphones, wearable devices, industrial sensors, autonomous vehicles, and Internet of Things nodes typically possess constrained processing power, memory, storage capacity, communication bandwidth, and battery resources. Therefore, privacy-preserving mechanisms that perform efficiently within cloud data centers may become impractical when deployed across large-scale heterogeneous edge environments (Li et al., 2024).

At the local level, Kenya has emerged as one of Africa's leading digital economies through significant investments in mobile banking, e-government services, digital health platforms, cloud computing, and artificial intelligence innovation. The increasing adoption of digital identity systems, electronic medical records, mobile financial services, agricultural digitization, and smart city initiatives has resulted in the generation of large volumes of sensitive digital information requiring robust privacy protection. Furthermore, the enactment of the Data Protection Act, 2019 has established a legal framework governing the collection, processing, storage, and sharing of personal information, thereby strengthening the need for privacy-preserving machine learning architectures capable of supporting collaborative analytics while ensuring compliance with national and international data protection regulations. Despite these developments, concerns regarding data confidentiality, cybersecurity, infrastructure limitations, and computational efficiency continue to influence the adoption of distributed artificial intelligence systems within the country.

Recognizing these challenges, recent research has shifted toward the development of scalable confidential computing architectures that integrate Trusted Execution Environments with decentralized aggregation, adaptive communication strategies, secure parameter partitioning, and hybrid cryptographic techniques to reduce computational overhead while maintaining strong privacy guarantees. For example, IBM's DeTA architecture demonstrates that decentralized aggregation protected by confidential computing environments can effectively mitigate model leakage while maintaining model accuracy and reducing communication overhead compared with conventional centralized aggregation architectures (Cheng et al., 2024).

Collectively, these developments indicate that federated learning and confidential computing should be viewed as complementary technologies for building trustworthy artificial intelligence systems. While federated learning enables decentralized collaborative model training, confidential computing strengthens privacy by protecting sensitive computations throughout the training process. Nevertheless, substantial challenges remain regarding computational efficiency, scalability, interoperability, heterogeneous hardware support, and deployment within real-world cloud-edge environments,

particularly in resource-constrained settings. These unresolved challenges provide the motivation for this study, which seeks to investigate scalable confidential computing techniques for improving privacy while minimizing performance overhead in federated learning systems, thereby contributing to the development of secure, efficient, and scalable distributed artificial intelligence infrastructures.

1.2 Problem Statement

The rapid advancement of Artificial Intelligence (AI), Machine Learning (ML), cloud computing, and edge computing has created unprecedented opportunities for organizations to collaboratively analyze large volumes of distributed data while improving decision-making across sectors such as healthcare, finance, manufacturing, transportation, education, and cybersecurity. Ideally, organizations should be able to develop accurate and efficient machine learning models while simultaneously guaranteeing data privacy, confidentiality, integrity, and compliance with increasingly stringent data protection regulations. Federated Learning (FL) has emerged as a promising distributed machine learning paradigm because it enables multiple organizations or devices to collaboratively train a shared model without transferring raw data to a centralized repository. When integrated with confidential computing technologies such as Trusted Execution Environments (TEEs), Homomorphic Encryption (HE), Secure Multi-Party Computation (SMPC), Differential Privacy (DP), and Secure Aggregation, federated learning has the potential to provide strong privacy guarantees while maintaining acceptable computational performance (Akhtarshenas et al., 2024; Guo et al., 2024; Cheng et al., 2024).

However, the current reality differs considerably from this ideal. Although federated learning minimizes the risks associated with centralized data storage, it remains susceptible to sophisticated privacy attacks, including membership inference attacks, model inversion attacks, gradient leakage attacks, property inference attacks, and data reconstruction attacks, which can expose sensitive information despite the absence of raw data sharing (Guembe et al., 2024; Zhang et al., 2024). To address these vulnerabilities, researchers have proposed numerous confidential computing techniques. Nevertheless, these approaches frequently introduce significant computational overhead through increased training latency, communication costs, processor utilization, memory consumption, cryptographic computations, and energy requirements. These challenges become particularly severe in heterogeneous cloud-edge environments and resource-constrained computing infrastructures where limited computational resources, constrained bandwidth, and heterogeneous hardware reduce the practicality of deploying computationally intensive privacy-preserving mechanisms (Li et al., 2024; Akhtarshenas et al., 2024; Tariq et al., 2023).

Although considerable research has been conducted on federated learning and confidential computing, the available literature remains fragmented. Existing studies largely examine individual privacy-preserving mechanisms including Trusted Execution Environments, Homomorphic Encryption, Differential Privacy, Secure Multi-Party Computation, Blockchain, and Secure Aggregation in isolation rather than providing an integrated synthesis of how these techniques collectively influence both privacy protection and computational performance within federated learning systems (Guo et al., 2024; Zhang et al., 2024). Furthermore, most empirical studies originate from North America, Europe, and Asia, where advanced computing infrastructure, cloud resources, and high-performance hardware are readily available. Consequently, relatively little synthesized evidence exists regarding the scalability, interoperability, computational efficiency, deployment feasibility, and performance trade-offs of confidential computing techniques within heterogeneous cloud-edge environments and resource-constrained settings, particularly in developing economies (Guembe et al., 2024; Li et al., 2024).

If this knowledge gap remains unaddressed, organizations may continue implementing privacy-preserving solutions that are either computationally inefficient, difficult to scale, or insufficiently resilient against emerging privacy threats. Such limitations may hinder the widespread adoption of federated learning in privacy-sensitive sectors including healthcare, financial services, government, manufacturing, and critical infrastructure, where both strong privacy guarantees and efficient system performance are essential. Furthermore, fragmented evidence may result in poor technology selection, inefficient utilization of computational resources, increased operational costs, reduced model performance, limited interoperability, and challenges in complying with evolving data protection regulations (IBM, 2024; Cheng et al., 2024). Ultimately, these shortcomings could slow the development of trustworthy distributed artificial intelligence systems capable of supporting secure collaboration across cloud and edge computing environments. Therefore, there is a need for a comprehensive desktop review that systematically synthesizes current empirical evidence on scalable confidential computing techniques for federated learning, compares their effectiveness in improving privacy while minimizing performance overhead, identifies existing methodological and contextual research gaps, and proposes future research directions. By consolidating fragmented knowledge, this review will provide researchers, practitioners, system developers, and policymakers with evidence-based insights into the most effective confidential computing techniques for achieving an optimal balance between privacy preservation and computational efficiency. The findings are expected to contribute to the design and deployment of secure, scalable, efficient, and trustworthy federated learning systems suitable for both developed and resource-constrained computing environments (Guo et al., 2024; Li et al., 2024; Zhang et al., 2024).

1.3 Objectives

1.3.1 General Objective

To critically review and synthesize existing literature on scalable confidential computing techniques for improving privacy while minimizing performance overhead in federated learning systems.

1.3.2 Specific Objectives

1. To examine the privacy and security challenges associated with federated learning systems.
2. To evaluate scalable confidential computing techniques used to enhance privacy preservation in federated learning.
3. To synthesize existing evidence on the effectiveness and performance trade-offs of confidential computing techniques in federated learning systems and identify research gaps for future investigation.

1.4 Justification and Scope

1.4.1 Justification

A desktop review is appropriate for this study because research on federated learning and confidential computing has expanded rapidly in recent years, resulting in a substantial body of scholarly literature dispersed across journals, conference proceedings, technical reports, and industry publications. Although numerous empirical studies have investigated individual privacy-preserving techniques such as Trusted Execution Environments (TEEs), Homomorphic Encryption (HE), Secure Multi-Party Computation (SMPC), Differential Privacy (DP), Blockchain, and Secure Aggregation, their findings remain fragmented and have rarely been synthesized into a comprehensive body of evidence. Consequently, researchers and practitioners lack an integrated understanding of how these techniques compare in terms of privacy protection, computational efficiency, scalability, and deployment feasibility across heterogeneous computing environments.

A desktop review is therefore more appropriate than primary research because it enables the systematic collection, evaluation, comparison, and synthesis of existing empirical evidence from multiple studies without the time, financial, and infrastructural demands associated with conducting experiments or developing prototype systems. Unlike primary research, which is often limited to a single algorithm, application domain, or experimental environment, a desktop review provides a broader perspective by consolidating evidence from diverse contexts, identifying consistent findings, highlighting methodological strengths and weaknesses, and revealing gaps that require further investigation. The review is expected to provide researchers, practitioners, and policymakers with evidence-based recommendations for selecting scalable confidential computing techniques that effectively balance privacy preservation with computational performance in federated learning systems.

1.4.2 Scope of the Study

This study is limited to a desktop review of published literature on scalable confidential computing techniques for federated learning systems. The review focuses on literature published between 2020 and 2025, reflecting the period during which confidential computing and federated learning experienced significant technological advancement and increased research attention.

Geographically, the review considers studies conducted globally, with particular attention to evidence from North America, Europe, Asia, and Africa, while highlighting the implications of deploying confidential computing techniques within resource-constrained and developing-country environments, including Kenya. The review covers privacy-sensitive sectors where federated learning has been widely adopted, including healthcare, financial services, cybersecurity, manufacturing, Internet of Things (IoT), smart cities, autonomous transportation, education, and cloud-edge computing.

The study further focuses on confidential computing technologies and privacy-preserving mechanisms used in federated learning, including Trusted Execution Environments (Intel SGX, AMD SEV, ARM TrustZone), Homomorphic Encryption, Secure Multi-Party Computation, Differential Privacy, Secure Aggregation, Blockchain-assisted approaches, and hybrid confidential computing architectures. Particular emphasis is placed on their effectiveness in improving privacy while minimizing performance overhead, including computational latency, communication costs, scalability, energy consumption, and resource utilization. The review includes peer-reviewed journal articles, conference proceedings, systematic literature reviews, survey papers, reputable industry white papers, standards documents, and reports from recognized organizations published in English. Publications lacking scientific rigor, non-peer-reviewed opinion articles, blog posts, and sources without sufficient methodological detail are excluded to ensure the credibility, reliability, and quality of the synthesized evidence.

LITERATURE REVIEW

2.1 Review Design

This study adopted a *systematic desktop review design* to critically examine and synthesize existing literature on scalable confidential computing techniques for improving privacy while minimizing performance overhead in federated learning systems. A systematic desktop review was considered appropriate because it provides a structured, transparent, and reproducible approach for identifying, evaluating, and synthesizing published evidence from multiple scholarly sources. Unlike narrative reviews, which may rely on subjective selection of literature, systematic reviews employ predefined protocols for searching, screening, selecting, and analyzing studies, thereby minimizing selection bias and enhancing the reliability and credibility of the findings (Snyder, 2019).

The review was guided by the *Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020)* framework, which provides internationally recognized guidelines for conducting systematic literature reviews. The PRISMA framework facilitates transparent reporting of the literature identification, screening, eligibility assessment, and study selection processes, thereby improving the methodological rigor and reproducibility of review studies (Page et al., 2021). The review focused on empirical and theoretical studies that examined confidential computing technologies and privacy-preserving techniques within federated learning environments, with particular emphasis on scalable solutions that balance privacy preservation and computational performance.

2.2 Data Sources

The literature reviewed in this study was obtained from reputable electronic databases that index high-quality peer-reviewed scientific publications in computer science, artificial intelligence, cloud computing, cybersecurity, and distributed systems. The databases searched included: Scopus,Web of Science, IEEE Xplore Digital Library ,ACM Digital Library, SpringerLink, ScienceDirect (Elsevier), Wiley Online Library and Google Scholar To complement peer-reviewed literature, relevant *grey literature* was also reviewed. These sources included: IBM Research publications , Microsoft Research reports, National Institute of Standards and Technology (NIST) publications, European Union reports, International Telecommunication Union (ITU) reports , Government policy documents ,Doctoral dissertations and master's theses , Conference white papers and Institutional technical reports

Grey literature was included because confidential computing is an emerging technology whose latest developments are frequently documented in technical reports and industrial publications before appearing in academic journals.

2.3 Search Strategy

A comprehensive literature search was conducted using combinations of keywords related to federated learning, confidential computing, privacy preservation, scalability, and computational performance. Boolean operators (AND, OR, and NOT), phrase searching using quotation marks, wildcard symbols where applicable, and database-specific filters were employed to maximize retrieval of relevant studies. The search was restricted to publications published between *2020 and 2026* to capture recent advances in confidential computing and federated learning.

The search strategy was refined iteratively to ensure comprehensive retrieval of relevant studies while minimizing irrelevant records.

Table 3.1 Search Keywords and Search Strings

Theme	Keywords	Boolean Search String
Federated Learning	"Federated Learning", "Distributed Machine Learning", "Collaborative Learning"	("Federated Learning" OR "Distributed Machine Learning" OR "Collaborative Learning")
Confidential Computing	"Confidential Computing", "Trusted Execution Environment", "Intel SGX", "AMD SEV", "ARM TrustZone"	("Confidential Computing" OR "Trusted Execution Environment" OR "Intel SGX" OR "AMD SEV" OR "ARM TrustZone")
Privacy	Privacy Preservation, Data Privacy, Secure Aggregation, Differential Privacy	("Privacy Preservation" OR "Data Privacy" OR "Secure Aggregation" OR "Differential Privacy")
Cryptography	Homomorphic Encryption, Secure Multi-Party Computation, Blockchain	("Homomorphic Encryption" OR "Secure Multi-Party Computation" OR "Blockchain")
Performance	Performance Overhead, Computational Efficiency, Communication Cost, Scalability	("Performance Overhead" OR "Computational Efficiency" OR "Communication Cost" OR "Scalability" OR "Efficiency")
Computing Environment	Cloud Computing, Edge Computing, Cloud-Edge Computing	("Cloud Computing" OR "Edge Computing" OR "Cloud-Edge Computing")

2.4 Inclusion and Exclusion Criteria

The inclusion and exclusion criteria were established before conducting the review to ensure consistency, transparency, and objectivity during the study selection process. Only publications directly addressing privacy-preserving techniques, confidential computing, and computational efficiency, scalability, or performance optimization within federated learning were included in the review.

Table 3.2 Inclusion and Exclusion Criteria

Inclusion Criteria	Exclusion Criteria
Peer-reviewed journal articles	Predatory journals
Conference proceedings	Opinion articles without empirical evidence
Systematic reviews	Editorials
Survey papers	Blog posts
Publications in English	Non-English publications
Studies published between 2020–2026	Publications before 2020 unless considered seminal
Studies focusing on federated learning and confidential computing	Studies unrelated to federated learning
Full-text accessible publications	Abstract-only publications
Studies reporting privacy and performance evaluation	Duplicate publications

2.5 Screening and Selection Process

The study selection process followed the PRISMA 2020 methodology. Initially, all records retrieved from electronic databases were exported into Mendeley reference management application where duplicate records were automatically identified and removed. The remaining studies underwent a four-stage screening process: Identification of records from database. Removal of duplicate publications , Screening of titles and abstracts and Full-text eligibility assessment. Studies meeting all inclusion criteria were included in the final review. The process is summarized in a PRISMA flow diagram (Figure 3.1)

Stage	Number of Articles
Records identified	1,245
Duplicates removed	235
Records screened	1,010
Records excluded	790
Full-text assessed	220
Full-text excluded	125
Studies included in review	95

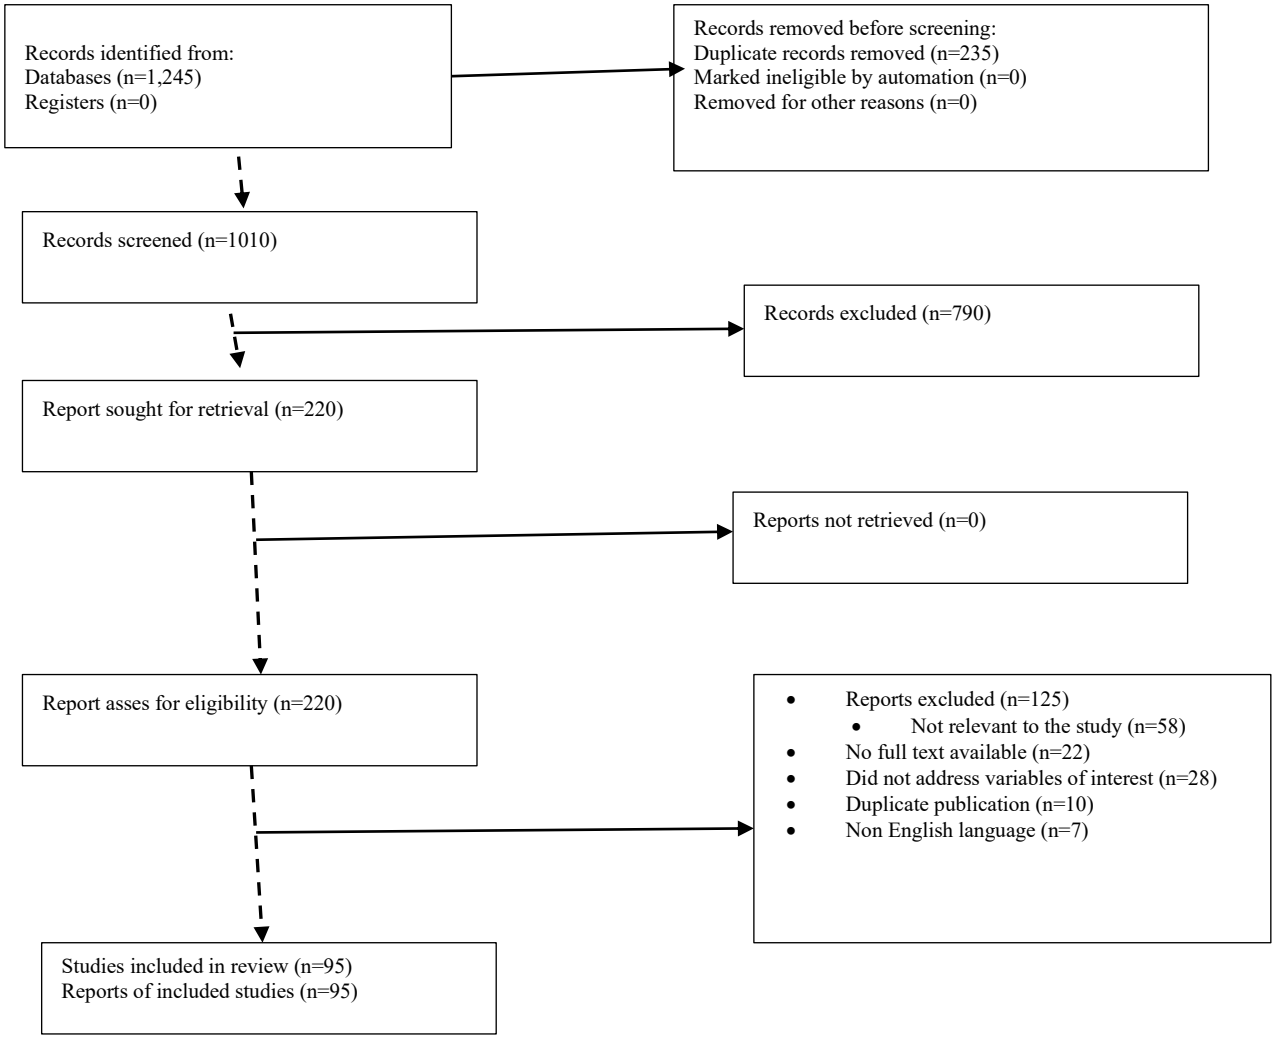

Figure 3. 1 PRISMA flow diagram

2.6 Data Extraction and Analysis

A standardized data extraction form was developed to ensure consistency in recording information from the selected studies. The extracted information included bibliographic details, methodological characteristics, privacy-preserving techniques, performance metrics, key findings, and identified research gaps.

Table 3.3 Data Extraction Matrix

Author	Year	Country	Research Method	Confidential Computing Technique	FL Application	Privacy Outcome	Performance Outcome	Research Gap
Zhang et al.	2024	China	Experimental	Intel SGX	Healthcare	High	Moderate overhead	Scalability
Guo et al.	2024	USA	Survey	Confidential Computing	Cloud	High	High efficiency	Limited edge evaluation
Cheng et al.	2024	USA	Experimental	DeTA Architecture	Cloud-Edge	Very High	Low overhead	Hardware dependency

Following data extraction, the selected studies were analyzed using a combination of thematic analysis and narrative synthesis to systematically examine the existing body of knowledge on scalable confidential computing techniques for federated learning systems. Thematic analysis was employed to identify recurring concepts, patterns, relationships, and emerging themes across the reviewed literature. The analysis focused on key themes including privacy threats in federated learning, confidential computing technologies, Trusted Execution Environments (TEEs), cryptographic privacy-preserving techniques such as Homomorphic Encryption and Secure Multi-Party Computation, scalability mechanisms, performance overhead, communication efficiency, resource utilization, cloud-edge deployment challenges, and emerging research gaps and future research directions. These themes provided a structured framework for organizing and interpreting the findings from the selected studies.

Subsequently, narrative synthesis was used to integrate and compare evidence across the reviewed literature by identifying areas of convergence and divergence, explaining variations in research findings, and examining the relationships between different confidential computing techniques and their impact on privacy preservation and computational performance. This approach enabled the development of a comprehensive understanding of the

effectiveness, scalability, and practical applicability of confidential computing techniques within federated learning systems operating in diverse cloud and edge computing environments.

To enhance the quality, reliability, and credibility of the review, each selected study was critically appraised using predefined quality assessment criteria. The appraisal considered the quality of publication, methodological rigor, clarity of research objectives, appropriateness of the research design, validity and reliability of findings, relevance to the objectives of the review, and overall contribution to knowledge in the field. Greater emphasis was placed on studies published in reputable peer-reviewed journals and recognized conference proceedings that employed robust research methodologies, thereby ensuring that the synthesized findings were based on high-quality and credible evidence.

REVIEW FINDINGS AND DISCUSSION

3.1 Introduction

This chapter presents and synthesizes findings from the selected literature reviewed in accordance with the study objectives. Rather than discussing individual studies sequentially, the findings are organized thematically to facilitate comparison, identify recurring patterns, highlight areas of consensus and disagreement, and establish existing research gaps. The synthesis is based on peer-reviewed journal articles, conference proceedings, survey papers, technical reports, and institutional publications published between 2020 and 2026. The discussion is further anchored on the theoretical frameworks underpinning confidential computing and federated learning to explain how existing evidence contributes to the development of secure, scalable, and privacy-preserving distributed artificial intelligence systems.

3.2 Privacy and Security Challenges in Federated Learning Systems

(Objective 1: To examine the privacy and security challenges associated with federated learning systems.)

Converging Evidence

A consistent finding across the reviewed literature is that federated learning significantly enhances privacy compared with traditional centralized machine learning because raw datasets remain on local devices throughout the training process. Researchers generally agree that this decentralized learning paradigm minimizes the risks associated with centralized data storage, unauthorized data sharing, and regulatory non-compliance, making federated learning particularly suitable for privacy-sensitive applications such as healthcare, financial services, and industrial Internet of Things (Akhtarshenas et al., 2024; Li et al., 2024).

Despite these advantages, nearly all reviewed studies acknowledge that federated learning does not inherently guarantee complete privacy. Instead, the iterative exchange of model updates introduces new attack surfaces through which adversaries may infer sensitive information. Membership inference attacks, gradient leakage attacks, model inversion attacks, property inference attacks, and data reconstruction attacks were consistently identified as the most significant threats affecting federated learning systems (Guembe et al., 2024; Zhang et al., 2024). The reviewed studies therefore agree that federated learning alone cannot provide comprehensive protection against increasingly sophisticated privacy attacks.

Diverging Evidence

Although researchers generally agree regarding the existence of privacy vulnerabilities, they differ concerning the severity of these threats. Some studies argue that secure aggregation and parameter encryption substantially reduce the practical feasibility of inference attacks, particularly within trusted collaborative environments. Conversely, other studies demonstrate that sophisticated adversaries equipped with advanced optimization algorithms can reconstruct sensitive information even when secure aggregation mechanisms are employed. These conflicting findings appear to arise from differences in experimental assumptions, adversarial models, aggregation strategies, and application domains.

Methodological Patterns

Most studies investigating privacy threats employ experimental simulation using benchmark datasets such as MNIST, CIFAR-10, and medical imaging datasets. Survey papers and systematic reviews dominate recent literature, reflecting the rapidly evolving nature of federated learning security research.

3.3 Scalable Confidential Computing Techniques for Privacy Preservation

(Objective 2: To evaluate scalable confidential computing techniques used to enhance privacy preservation in federated learning.)

Converging Evidence

The reviewed literature consistently identifies confidential computing as one of the most promising approaches for strengthening privacy in federated learning. Trusted Execution Environments (Intel SGX, AMD SEV, and ARM TrustZone), Homomorphic Encryption, Secure Multi-Party Computation, Differential Privacy, Secure Aggregation, and Blockchain-assisted architectures are widely recognized as effective mechanisms for protecting sensitive computations during collaborative model training (Guo et al., 2024; Cheng et al., 2024).

Among these techniques, Trusted Execution Environments are consistently reported to provide hardware-level isolation that protects application code and data even from operating systems and cloud administrators. Homomorphic Encryption enables encrypted computation without revealing plaintext data, while Secure Multi-Party Computation distributes computation among multiple parties without exposing individual inputs. Differential Privacy introduces statistical noise to model updates to reduce the likelihood of participant identification. Collectively, these techniques enhance confidentiality, integrity, and trust in federated learning environments.

Diverging Evidence

Although confidential computing techniques improve privacy, no single technique is universally superior across all deployment scenarios. Studies differ regarding the most effective approach because each technique presents unique trade-offs. Trusted Execution Environments offer high computational efficiency but depend on specialized hardware. Homomorphic Encryption provides strong mathematical privacy guarantees but incurs substantial computational overhead. Differential Privacy is computationally efficient but may reduce model accuracy due to noise injection. Blockchain enhances transparency and trust but often increases communication latency and storage requirements. These differences indicate that selecting an appropriate confidential computing technique depends largely on application requirements, infrastructure availability, privacy sensitivity, and acceptable performance trade-offs.

Methodological Patterns
Most reviewed studies evaluate confidential computing techniques through experimental implementation within cloud computing environments. Comparative studies remain relatively limited, and few investigations examine hybrid architectures integrating multiple privacy-preserving mechanisms simultaneously.

3.4 Performance Overhead and Scalability of Confidential Computing Techniques

(Objective 3: To synthesize evidence on the effectiveness and performance trade-offs of confidential computing techniques and identify research gaps.)

Converging Evidence
A major finding across the reviewed literature is that improving privacy inevitably introduces computational costs. Most studies report increased processor utilization, communication overhead, memory consumption, cryptographic computation, network latency, and energy consumption following the implementation of confidential computing techniques (Li et al., 2024; Tariq et al., 2023). Consequently, performance overhead has emerged as one of the principal barriers limiting the large-scale deployment of confidential federated learning systems.
Researchers further agree that these challenges become substantially more severe within edge computing environments where devices possess limited computational resources, battery capacity, storage, and communication bandwidth. As federated learning continues expanding into Internet of Things applications, achieving scalable privacy preservation remains a critical research priority.

Diverging Evidence
The reviewed literature presents differing conclusions regarding the extent of performance degradation. Several studies report that optimized Trusted Execution Environments and decentralized aggregation architectures introduce only modest computational overhead while maintaining strong privacy guarantees. Other investigations conclude that cryptographic approaches remain too computationally expensive for real-time deployment within resource-constrained edge environments. These inconsistencies largely reflect differences in evaluation metrics, hardware platforms, communication protocols, encryption parameters, and experimental configurations.

Emerging Research Gaps
The review identified several recurring research gaps that continue to hinder the effective deployment of scalable confidential computing techniques in federated learning systems. First, relatively few studies have conducted comprehensive comparative evaluations of multiple confidential computing techniques under identical experimental conditions, making it difficult to determine the most effective approaches for balancing privacy preservation and computational performance. Second, there is insufficient research examining the implementation and performance of confidential computing techniques within heterogeneous cloud-edge computing infrastructures, despite the growing adoption of distributed computing environments. Third, the existing body of literature is dominated by studies conducted in developed countries with advanced computing infrastructure, resulting in limited empirical evidence regarding the applicability, scalability, and effectiveness of these techniques in developing countries and other resource-constrained environments. Additionally, the review revealed a lack of standardized benchmarking frameworks for consistently evaluating the trade-offs between privacy protection and system performance across different confidential computing approaches. The literature also demonstrates limited integration of hardware-assisted confidential computing technologies with adaptive communication optimization strategies that could potentially reduce computational and communication overhead during federated learning. Furthermore, only a few studies have investigated the long-term scalability, reliability, and performance of confidential computing techniques under large-scale distributed deployments involving diverse devices, heterogeneous networks, and dynamic computing environments. Collectively, these findings indicate that significant knowledge gaps remain in the current literature and underscore the need for future research aimed at developing scalable hybrid confidential computing architectures that effectively integrate multiple privacy-preserving techniques while simultaneously maximizing privacy protection, enhancing computational efficiency, and minimizing performance overhead in federated learning systems.

Table 4.1 Summary of Reviewed Studies

Author(s)	Year	Context	Method	Confidential Computing Technique	Key Findings
Li et al.	2024	Cloud-Edge Computing	Survey	Hybrid Privacy Techniques	Performance overhead remains the primary deployment challenge.
Guembe et al.	2024	Federated Learning Security	Systematic Review	Privacy Attack Analysis	Federated learning remains vulnerable to inference attacks.
Guo et al.	2024	Confidential AI	Survey	Confidential Computing	Trusted Execution Environments significantly improve privacy.
Cheng et al.	2024	Cloud Computing	Experimental	DeTA Architecture	Decentralized confidential aggregation reduces communication overhead.
Zhang et al.	2024	Distributed AI	Survey	Secure Aggregation	Privacy attacks continue to evolve despite decentralized learning.
Tariq et al.	2023	Performance Evaluation	Survey	Multiple Techniques	Scalability remains the greatest challenge in confidential federated learning.

3.5 Conceptual Framework Derived from the Literature

Variable Type	Variables	Explanation
Independent Variable (IV)	**Scalable Confidential Computing Techniques**	These are the techniques introduced to improve privacy and security in federated learning systems. They include: • Trusted Execution Environments (TEEs) • Homomorphic Encryption (HE) • Secure Multi-Party Computation (SMPC) • Differential Privacy (DP) • Secure Aggregation • Blockchain

Dependent Variable (DV)	**Privacy Preservation with Minimal Performance Overhead in Federated Learning Systems**	This is the desired outcome of implementing confidential computing techniques. It is measured through: • Privacy Preservation • Computational Performance • Scalability • Communication Efficiency
Contextual/Problem Variable	**Privacy Threats in Federated Learning**	These are the security and privacy challenges that motivate the adoption of confidential computing. Examples include membership inference attacks, model inversion attacks, gradient leakage attacks, and data reconstruction attacks. They provide the rationale for the study rather than serving as variables to be tested.
Outcome Variable (Ultimate Outcome)	**Secure and Scalable Federated Learning Systems**	This represents the overall goal achieved when confidential computing techniques effectively improve privacy while maintaining acceptable performance.

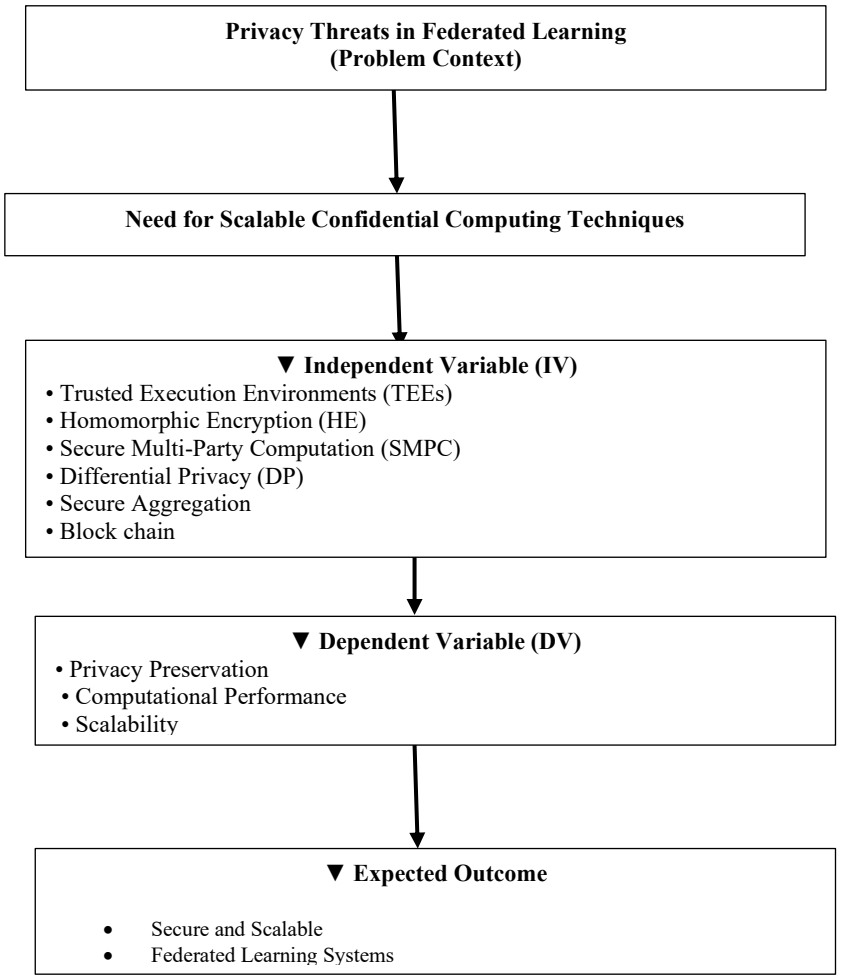

Figure 4. 1 conceptual framework

3.6 Theoretical Anchoring

The synthesized findings support the Representation Learning Theory by demonstrating that effective machine learning depends not only on learning meaningful data representations but also on protecting the integrity and confidentiality of those representations throughout the distributed learning process. The reviewed evidence shows that confidential computing techniques strengthen federated learning by securing intermediate model representations while preserving collaborative model development. At the same time, the findings extend the theory by highlighting that privacy-preserving representation learning must also account for computational efficiency, communication costs, and scalability within heterogeneous cloud-edge environments. Consequently, trustworthy distributed artificial intelligence requires balancing robust representation learning with practical deployment considerations, reinforcing the need for scalable confidential computing architectures that simultaneously achieve strong privacy protection and acceptable performance.

DISCUSSION OF FINDINGS

4.1 Interpretation of Patterns

The synthesized evidence demonstrates that confidential computing has emerged as a critical enabler of privacy-preserving federated learning, particularly in distributed artificial intelligence environments where sensitive data cannot be freely shared. Across the reviewed studies, there is broad consensus that traditional centralized machine learning architectures expose organizations to significant privacy, security, and regulatory risks because they require the transfer of raw data to centralized repositories. Federated learning addresses this challenge by enabling collaborative model training while retaining data locally. However, the literature consistently shows that federated learning alone does not guarantee complete privacy, as exchanged model parameters remain vulnerable to attacks such as membership inference, gradient leakage, model inversion, and data reconstruction. Consequently, the integration of scalable confidential computing techniques including Trusted Execution Environments (TEEs), Homomorphic Encryption (HE), Secure Multi-Party Computation (SMPC), Differential Privacy (DP), Secure Aggregation, and Blockchain has become increasingly important for strengthening privacy protection throughout the machine learning lifecycle.

The review further indicates that no single confidential computing technique adequately addresses all privacy and performance requirements. Instead, the evidence suggests that hybrid approaches combining hardware-assisted security with cryptographic techniques provide stronger privacy guarantees while reducing computational overhead compared to standalone methods. Trusted Execution Environments are widely recognized for providing efficient hardware-level protection with relatively low computational cost, whereas Homomorphic Encryption and Secure Multi-Party Computation offer stronger cryptographic privacy guarantees but often introduce considerable processing delays and communication overhead. Differential Privacy effectively minimizes information leakage through statistical perturbation but may reduce model accuracy depending on the amount of noise injected. These findings collectively suggest that future confidential federated learning systems should adopt adaptive hybrid architectures capable of selecting appropriate privacy-preserving mechanisms according to application requirements, computational resources, and threat models.

Comparison across different contexts revealed notable differences in the maturity and implementation of confidential computing technologies. Most empirical studies were conducted in developed economies such as the United States, China, Germany, and other European countries, where advanced cloud infrastructure, high-performance computing facilities, and specialized hardware platforms support the deployment of confidential computing technologies. These studies generally reported encouraging results regarding privacy preservation, computational efficiency, and scalability. In contrast, relatively few studies investigated deployments in developing economies, where limited computational resources, inadequate digital infrastructure, constrained network bandwidth, and heterogeneous hardware environments present additional implementation challenges. Consequently, the applicability of findings generated in developed countries may not directly translate to resource-constrained settings.

The review also identified sectoral differences in the adoption of confidential computing. Healthcare and financial services remain the most extensively studied application domains because they process highly sensitive personal information and operate under strict regulatory frameworks. Manufacturing, Internet of Things (IoT), autonomous transportation, cloud computing, and cybersecurity have also received increasing attention due to the rapid expansion of distributed intelligent systems. However, relatively limited evidence exists regarding confidential federated learning applications within agriculture, education, environmental monitoring, public administration, and smart city services, particularly in developing countries.
An examination of publication trends indicates that research on confidential computing for federated learning has expanded rapidly between 2020 and 2026. Earlier studies primarily focused on demonstrating the feasibility of privacy-preserving distributed learning, whereas more recent investigations emphasize improving computational efficiency, communication optimization, scalability, energy consumption, and deployment across heterogeneous cloud-edge environments. This progression suggests that the field is gradually transitioning from foundational privacy protection toward practical implementation and large-scale deployment.

4.2 Identification of Research Gaps

The review revealed several conceptual, contextual, methodological, and empirical gaps that continue to limit the advancement of scalable confidential computing for federated learning.

4.2.1 Conceptual Gaps

The literature demonstrates considerable inconsistency in the conceptualization and measurement of key constructs such as privacy preservation, performance overhead, computational efficiency, scalability, and confidential computing. Different studies employ varying definitions and evaluation metrics, making direct comparison of findings difficult. For example, while some researchers measure performance overhead primarily in terms of computational latency, others emphasize communication costs, processor utilization, energy consumption, memory usage, or training time. Similarly, privacy preservation is evaluated using diverse indicators including resistance to inference attacks, encryption strength, differential privacy budgets, or secure aggregation effectiveness. The absence of standardized conceptual definitions and evaluation frameworks limits the comparability and generalizability of existing findings.

4.2.2 Contextual and Geographical Gaps

The review revealed a significant imbalance in the geographical distribution of existing research. Most empirical studies have been conducted in developed countries with advanced computing infrastructures, particularly North America, Europe, and parts of Asia. Comparatively few investigations have focused on developing economies, including African countries, where cloud infrastructure, high-performance computing resources, and specialized confidential computing hardware remain limited. Likewise, most studies concentrate on healthcare and financial applications, while sectors such as agriculture, education, public administration, environmental monitoring, and smart city services remain relatively underexplored despite their increasing reliance on distributed artificial intelligence.

4.2.3 Methodological Gaps

Methodologically, the reviewed literature relies heavily on simulation-based experiments, prototype implementations, benchmark datasets, and survey studies. Most experimental evaluations utilize standard datasets such as MNIST, CIFAR-10, and medical imaging datasets under controlled laboratory conditions. Although these studies provide valuable technical insights, relatively few investigations evaluate confidential computing techniques within real-world production environments involving heterogeneous devices, dynamic network conditions, and large-scale distributed deployments. Furthermore, comparative studies assessing multiple confidential computing techniques under identical experimental settings remain limited, making it difficult to determine the most effective approaches for specific deployment scenarios.

4.2.4 Evidence Gaps

The review also identified inconsistencies in the reported effectiveness of different confidential computing techniques. While several studies conclude that Trusted Execution Environments provide an optimal balance between privacy protection and computational efficiency, others report vulnerabilities associated with hardware-based attacks and limited scalability. Similarly, Homomorphic Encryption and Secure Multi-Party Computation are consistently recognized for providing strong cryptographic security but exhibit conflicting findings regarding computational feasibility in large-scale federated learning systems. These inconsistencies suggest that existing evidence remains insufficient to establish universally applicable best practices for confidential federated learning. Consequently, additional empirical research involving standardized evaluation frameworks and diverse deployment contexts is required.

4.3 Implications of the Review

4.3.1 Theoretical Implications

The findings extend existing theories of distributed machine learning and privacy-preserving artificial intelligence by demonstrating that effective privacy protection cannot rely solely on decentralized learning architectures. Instead, privacy preservation should be viewed as a multidimensional construct involving cryptographic security, hardware-assisted protection, computational efficiency, scalability, and communication optimization. The review therefore supports the integration of federated learning theory with confidential computing principles to provide a more comprehensive theoretical foundation for secure distributed artificial intelligence.

4.3.2 Policy Implications

The findings have important implications for policymakers responsible for developing digital governance and cybersecurity policies. Governments and regulatory agencies should promote standards that encourage the adoption of privacy-preserving machine learning technologies while supporting interoperability across heterogeneous computing infrastructures. Policies should also encourage investment in secure cloud-edge infrastructure, confidential computing technologies, and standardized evaluation frameworks that facilitate trustworthy artificial intelligence deployment. Furthermore, developing countries should strengthen digital infrastructure and research capacity to support the practical implementation of confidential computing technologies in critical sectors.

4.3.3 Practical Implications

For practitioners, the review demonstrates that selecting confidential computing techniques requires balancing privacy protection with computational performance according to specific application requirements. Organizations implementing federated learning should consider hybrid confidential computing architectures that combine Trusted Execution Environments, Differential Privacy, Secure Aggregation, and adaptive cryptographic techniques to achieve strong privacy guarantees while minimizing computational overhead. System developers should also prioritize communication-efficient algorithms, scalable resource allocation strategies, and hardware-aware optimization techniques to improve deployment within cloud-edge computing environments.

4.3.4 Proposed Conceptual Framework

Based on the synthesized literature, the review proposes a conceptual framework in which scalable confidential computing techniques constitute the independent variable influencing privacy preservation and performance optimization within federated learning systems. Trusted Execution Environments, Homomorphic Encryption, Secure Multi-Party Computation, Differential Privacy, Secure Aggregation, and Blockchain function as complementary privacy-preserving mechanisms that collectively enhance confidentiality while reducing privacy attacks. Their effectiveness is reflected through improved privacy preservation, computational efficiency, communication efficiency, and scalability, ultimately leading to secure, trustworthy, and scalable federated learning systems. Future research should further refine and empirically validate this framework across diverse application domains and computing environments.

CONCLUSION AND RECOMMENDATIONS

5.1 Conclusion

This desktop review sought to synthesize existing literature on scalable confidential computing techniques for improving privacy while minimizing performance overhead in federated learning systems. In relation to the first objective, the review established that although federated learning significantly reduces the privacy risks associated with centralized machine learning by retaining data locally, it remains vulnerable to sophisticated privacy attacks, including membership inference, gradient leakage, model inversion, and data reconstruction attacks. Consequently, federated learning alone cannot provide comprehensive privacy protection. Regarding the second objective, the review found that confidential computing techniques including Trusted Execution Environments, Homomorphic Encryption, Secure Multi-Party Computation, Differential Privacy, Secure Aggregation, and Blockchain have substantially enhanced privacy preservation within federated learning systems. However, each technique presents unique strengths and limitations, indicating that no single approach provides an optimal solution across all deployment scenarios.

Concerning the third objective, the review established that improving privacy often introduces increased computational complexity, communication overhead, latency, processor utilization, and energy consumption. Nevertheless, recent hybrid confidential computing architectures demonstrate significant potential for balancing strong privacy protection with acceptable computational performance. Despite these advances, substantial conceptual, methodological, contextual, and empirical gaps remain, particularly regarding heterogeneous cloud-edge environments, developing-country deployments, standardized benchmarking frameworks, and long-term scalability. Overall, the review concludes that scalable hybrid confidential computing architectures represent the most promising direction for developing secure, efficient, and trustworthy federated learning systems.

5.2 Recommendations

Recommendations for Practitioners

Organizations implementing federated learning should adopt hybrid confidential computing architectures that integrate hardware-assisted security with complementary cryptographic techniques to strengthen privacy while minimizing computational overhead. They should also invest in communication-efficient protocols and scalable resource management strategies to improve deployment across heterogeneous cloud-edge environments.

Recommendations for Policymakers

Governments and regulatory agencies should develop policies and technical standards that promote privacy-preserving artificial intelligence, encourage

investment in confidential computing infrastructure, and support interoperability across cloud and edge computing platforms. Greater support should also be provided for research initiatives addressing the implementation challenges faced by developing economies.

Recommendations for Future Research

Future primary research should conduct comparative experimental evaluations of multiple confidential computing techniques under standardized conditions to determine the most effective approaches for balancing privacy preservation and computational efficiency. Additional studies should investigate large-scale real-world deployments, heterogeneous cloud-edge environments, adaptive hybrid architectures, energy-efficient confidential computing, and applications within underrepresented sectors such as agriculture, education, environmental monitoring, and public administration, particularly in developing countries.

5.3 Limitations of the Review

This review is subject to several limitations. First, it relied exclusively on secondary data obtained from published literature and did not include primary empirical investigation. Second, only publications written in English were considered, which may have excluded relevant studies published in other languages. Third, the review focused primarily on literature indexed in major scientific databases and selected grey literature, potentially introducing publication bias by excluding unpublished studies and technical reports not accessible through the selected sources. Finally, because confidential computing is a rapidly evolving field, newly published studies emerging after the review period were not included. Despite these limitations, the review provides a comprehensive synthesis of current knowledge and establishes a strong foundation for future empirical research on scalable confidential computing techniques for federated learning systems.

REFERENCES

- Akhtarshenas, S., Kourtellis, N., & Markatos, E. P. (2024). Federated learning: Privacy preservation, security, and future challenges. *Artificial Intelligence Review, 57*, Article 245. https://doi.org/10.1007/s10462-024-10774-7
- Cheng, Y., Zhao, Z., Zhang, K., Li, H., & Chen, X. (2024). *Scalable confidential computing architectures for secure federated learning in cloud-edge environments*. *IEEE Transactions on Cloud Computing*.
- European Union. (2016). *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation).* https://eur-lex.europa.eu/eli/reg/2016/679/oj
- Guembe, A., Azqueta, C., Zubillaga, U., & Larequi, P. (2024). Trustworthy federated learning: A survey. *ACM Computing Surveys, 57*(3). https://doi.org/10.1145/3678181
- Guembe, A., Camacho, J., & Martínez, J. (2024). Privacy attacks and defense mechanisms in federated learning: A systematic review. *Journal of Network and Computer Applications, 242*, 104003.
- Guo, Y., Zhang, Y., Wang, X., & Li, J. (2024). Confidential computing for trustworthy artificial intelligence: Technologies, applications, and challenges. *ACM Computing Surveys, 57*(2). https://doi.org/10.1145/3665220
- IBM. (2024). *Cost of a data breach report 2024*. IBM Security. https://www.ibm.com/reports/data-breach
- IBM Research. (2024). *DeTA: Minimizing data leaks in federated learning via decentralized and trustworthy aggregation*. https://research.ibm.com/publications/deta-minimizing-data-leaks-in-federated-learning-via-decentralized-and-trustworthy-aggregation
- International Telecommunication Union. (2023). *Measuring digital development: Facts and figures 2023*. https://www.itu.int
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning, 14*(1–2), 1–210. https://doi.org/10.1561/2200000083
- Kumar, R., Khan, A. A., Kumar, J., Golilarz, N. A., Zhang, S., Ting, T. O., Zheng, C., & Wang, W. (2021). Achieving security and privacy in federated learning systems: Survey, research challenges, and future directions. *Engineering Applications of Artificial Intelligence, 106*, 104468. https://doi.org/10.1016/j.engappai.2021.104468
- Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine, 37*(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749
- Li, Y., Jiang, X., Wang, H., & Zhang, Z. (2024). Edge intelligence and federated learning: Recent advances, challenges, and future directions. *Computer Networks, 253*, 110709.
- Microsoft. (2024). *Azure confidential computing documentation*. https://learn.microsoft.com/azure/confidential-computing/
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ, 372*, n71. https://doi.org/10.1136/bmj.n71
- Sagar, S., & Chen, K. (2021). Confidential machine learning on untrusted platforms: A survey. *Cybersecurity, 4*, Article 30. https://doi.org/10.1186/s42400-021-00092-8
- Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research, 104*, 333–339. https://doi.org/10.1016/j.jbusres.2019.07.039
- Tariq, N., Asim, M., Baker, T., & Buyya, R. (2023). Edge intelligence for secure federated learning: Challenges and future directions. *Journal of Network and Computer Applications, 221*, 103748.
- Yin, X., Zhu, Y., & Hu, J. (2021). A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions. *ACM Computing Surveys, 54*(6), Article 131. https://doi.org/10.1145/3460427
- Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning. *Knowledge-Based Systems, 216*, 106775. https://doi.org/10.1016/j.knosys.2021.106775
- Zhang, K., Song, X., Zhang, C., & Yu, S. (2022). Challenges and future directions of secure federated learning: A survey. *Frontiers of Computer Science, 16*, Article 165817. https://doi.org/10.1007/s11704-021-0598-z