

Leveraging Confidential Computing for Secure Multi-Party Analytics in the Public Cloud

Jahanzeb Jamil^{1*}

¹Department of Information Technology, University of Karachi

DOI: <https://doi.org/10.36347/sjet.2025.v13i11.003>

| Received: 26.09.2025 | Accepted: 21.11.2025 | Published: 27.11.2025

*Corresponding author: Jahanzeb Jamil

Department of Information Technology, University of Karachi

Abstract

Original Research Article

The extensive implementation of the multi-party analytics cloud computing faces significant security issues especially with sensitive information. The paper presents a confidential computing infrastructure designed to secure multi-party analytics on the Google Cloud Platform, and based on AMD technology SEV-SNP. We have a performance overhead of 22.3 per cent compared to non-homomorphic methods, which is significantly less than the overhead incurred using homomorphic encryption (1,500 per cent) and insecure multi-party computation (800 per cent), and ensures complete data confidentiality throughout computation. The system is able to attain an encryption throughput of 260MB/s, the attestation success rate of 99.98 percent and it can be effectively scaled to support ten subjects. Practical use in experimental assessments of the framework, it is shown to be applicable to privacy preserving analytics through the use of healthcare, financial, and research data and meet regulatory compliance standards.

Keywords: Confidential Computing, Secure Multi-Party Analytics, Google Cloud, AMD SEV-SNP, Privacy-Preserving Analytics, Cloud Security, Data Protection.

Copyright © 2025 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY-NC 4.0) which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

INTRODUCTION

The advent of the cloud computing has transformed the data storage and processing environment and provided effective analytics possibilities in various fields, including healthcare, finances, and scientific studies (Sun, 2019; Mouchet *et al.*, 2023). However, such swift development has caused major fears concerning the security and privacy of data, especially in the context of a situation where sensitive data are exchanged between two or more parties (Geppert *et al.*, 2022; Sun *et al.*, 2019). Confidential computing is an encouraging paradigm that attempts to resolve these issues by providing a way to safely process data in closed systems, so-called Trusted Execution Environments (TEEs) (Zobaed and Salehi *et al.*, 2025). Confidential computing lowers the chance of data breach and unauthorized access by making sure that the data is encrypted at each stage of the life cycle, i.e. during transit, during rest and during computed, which leads to the development of trust in multi-party analytics (Mehrtak *et al.*, 2021).

Multi-party analytics is the situation when two or more parties are in need of access to sensitive data, but this must not interfere with the privacy of a person or the

confidentiality of an organization (Cho *et al.*, 2018). As an example, in the healthcare industry, group research of institutions may require some sensitive data to be analyzed without violating the privacy of patients (Sheller *et al.*, 2020; Cho *et al.*, 2018). Gated technologies, including Secure Multi-Party Computation (MPC) and homomorphic encryption, can be used by entities that would like to carry out common computations on shared datasets without disclosing the personal data to other organisations (Casaletto *et al.*, 2022; Babenko *et al.*, 2022). These privacy-saving methods are necessary because they comply with regulatory policies, such as the HIPAA and GDPR, which require strong protection of personally identifiable information (Asif-ur-Rahman *et al.*, 2019).

Considering the risks and opportunities of the convergence of confidential computing and multi-party analytics, the implementation of these emerging technologies in cloud systems, including Google Cloud, would generate significant improvements in privacy-conserving models and analytics (Gao *et al.*, 2019). Such integration can be effective that the future of data sensitivity and accessibility could be shaped that provides the ability to extract insights out of the otherwise limited datasets within the confines of strict

confidential computing policies (Podschwadt *et al.*, 2022).

Research Objectives

The primary objectives of this research on leveraging confidential computing for secure multi-party analytics in Google Cloud are as follows:

- A strict analysis of confidential computing models, namely Trusted Execution Environments (TEEs) and Secure Multi-Party Computation (MPC), at the Google cloud ecosystem.
- Discuss the efficiency of security technologies, such as homomorphic encryption and differential privacy in securing sensitive data when processing them.
- Use case studies that are specific to the industry to determine the applicability and usefulness of the technologies in the field including the healthcare and the finance industry.
- Make practical suggestions regarding the implementation of secure multi-party analytics into practice.

Related Work:

Confidential computing is a critical advancement in the data analytics paradigm that enables the implementation of multi-party analytics using security measures to ensure that privacy is preserved, particularly on such clouds as Google Cloud. This literature review sums up the available literature on the topic focusing on methodologies, architectural designs, and implications with a range of computational environments.

1. Understanding Confidential Computing

Confidential computing refers to the use of trusted execution environments (TEEs), implemented in hardware, to secure sensitive data, and process such data in a way that prevents its disclosure even to the host environment. This value cannot be done without in the areas, where data privacy is the foremost consideration, such as healthcare, finance, and sensitive governmental applications (Bhalla, 2025; Reddy, 2025). The Execution Environments are considered to be safe because they protect computations and prevent external and internal threats and, in such a way, reduce the risk of the data breach during its implementation to a considerable degree (Shah *et al.*, 2025; Gogineni *et al.*, 2024; Feng *et al.*, 2024). Experience shows that such an architectural solution does not only increase the level of security but also enables the adherence to the strict regulatory frameworks, including the General Data Protection Regulation (GDPR) and the Health Insurance Portability and Accountability Act (HIPAA) (Reddy 2025; Vanga *et al.*, 2025).

2. Multi-Party Computation and Privacy-Preserving Analytics

One of the key concepts that inform multi-party analytics is the so-called secure multi-party computation (MPC), a distributed computing paradigm that operates sensitive data without revealing the information behind it to the involved parties (Ejaz *et al.*, 2024). This approach is especially beneficial when it comes to a group environment in which information exchange is a necessity but confidentiality is essential. Studies have also emphasized the effectiveness of frameworks that use homomorphic encryption, differential privacy, and federated learning, which increases the privacy-protective nature of cloud-based analytics (Sharma *et al.*, 2025; Eboseromen *et al.*, 2022). These methods will guarantee that there is no leakage of confidential dataset attributes even in cases whereby more than one party is involved in the analytical activity.

3. Applications in Cloud Platforms

Various powerful confidentiality functions that make use of confidential computing technologies have been built into cloud providers, especially Google Cloud. Cloud-analytics solutions will need to be designed in such a manner that it tackles the issues of scalability, performance, and privacy. An example of this approach is the confidential virtual machines provided at Google, which allows users to process big data with a high level of confidentiality (Bhalla *et al.*, 2025; Valadares *et al.*, 2021). Studies show that the combination of MPC and TEEs in the cloud environment can significantly improve security without facing exorbitant latency charges (Gogineni *et al.*, 2025; Gao *et al.*, 2023).

4. Use Cases in Various Domains

Confidential computing can be applied in the current healthcare industry to enable various healthcare organizations to work together on predictive healthcare but do so with the highest level of privacy protection of patients (Ejaz *et al.*, 2024; Reddy 2025). Federated learning, a decentralized paradigm, where models are trained on multiple servers that store local data examples and do not exchange them, is one of the most essential applications of this technology and enables the generation of data-driven insights in fragmented healthcare systems (Dendukuri *et al.*, 2025). Similarly, in the financial field, the capacity to perform fraud analytics in a privacy-preserving way will guarantee that sensitive transaction data will be safe yet providing useful information (Malkoochi *et al.*, 2025).

5. Challenges and Future Directions

Even though there have been steps towards the use of confidential computing in the secure analytics process, a number of setbacks still exist. Scalability, computational load, and complexity of deploying those technologies in the existing infrastructure issues should be considered (Eboseremen *et al.*, 2022; Sathar *et al.*, 2025). Scientists suggest adopting a holistic model that combines various security paradigms, such as

differential privacy and encryption, and confidential computing to create strong security models that can endure other cyber threats (Rathi *et al.*, 2025; Li *et al.*, 2023; Ejaz *et al.*, 2024). The optimization of hybrid architectures should become a major priority in future work, and it is necessary to improve the performance and make it easier to deploy in dedicated cloud environments (Gao *et al.*, 2023; Adelusi *et al.*, 2022).

METHODOLOGY

It was a comprehensive experimental study that developed, deployed, and tested a safe multi-party analytics system involving confidential computing strategies on Google Cloud Platform. The methodological framework was based on a logical empirical methodology which included the configuration of the environment, the deployment of the system, the benchmarking of performance and the validation of the security.

Experimental Environment and Infrastructure Configuration

The experimental hardware was rolled out on the Google Cloud Platform with n2d Standard-8 instances and AMD EPYC 7B12 processors, and 8 virtual CPU cores and 32 GB of RAM. All of them were programmed to make use of the AMD Secure Encrypted Virtualization with Secure Nested Paging (SEV-VSNP) technology, which made the confidential computing achievable. The operating conditions included Ubuntu 20.04 LTS with 500GB persistent disks of equal size such that the storage performance was always constant across the trials. The topology or the network relied on the Premium Tier of Google Cloud, which could support the services at a 10 Gbps throughput and greatly reduced the variability in the performance measurements caused by the network itself.

System Implementation and Architecture

The secure multi-party analytics was written in C++ and it took advantage of low-level access to the hardware and tuned performance. Its design was in the form of a modular architecture that consists of four major modules namely: secure enclave management, cryptographic engine, attestation service, and analytics engine. The secure enclave element provided hardware-enforced isolation through the AMD SEV-SNP and the cryptographic engine was used to encrypt data as it was sent and computed using AES-256-GCM. The attestation service allowed remote check integrity verification based on challenge-response protocols and the analytics engine did privacy-conserving computations on encrypted data sets.

The implementation was a development process based on an iterative development cycle, starting with key cryptographic primitives and adding more advanced multi-party analytics functionalities. Every part was unit tested and integration tested to ensure functional correctness and security standards. The development

process incorporated a persistent security evaluation in form of static code analysis, dynamic vulnerability tests as well as in-depth manual review by the security experts.

Dataset Composition and Preparation

Three different synthetic datasets were created to evaluate the performance of the systems in different application areas. The healthcare data was a simulated medical record dataset that included demographic data, diagnosis, cost of treatment, laboratory results and had a size of 100,000 up to 1,000,000 records (approximately 50MB to 500MB). The financial data was a table that modeled the transaction data, which included transaction amounts, merchant details, geographical locations, and date time, ranging between 100MB and 1GB. The data set of the research consisted of patterns of academic collaboration with profile of researchers, publication history, citation network, and collaboration network with a size of 25MB-250MB. Synthetic generation of all datasets was done to make sure no privacy was violated and still satisfy realistic data distributions and interrelationships.

Performance Evaluation Methodology

Performance evaluation was done in a multi-faceted manner, which used encryption efficiency, computational overhead and scalability characteristics. The encryption speed was measured by conducting a series of classic experiments involving AES-256-GCM encryption and decryption cycles with sizes of 10MB to 500MB of data and each individual run of the experiment was done 100 times to determine statistical significance. The computational overhead analysis was used to compare the execution time of standard analytics functions, including aggregates, join operations, and statistical functions, to confidential computing and standard cloud deployments.

Scalability testing was used to test how the system behaves in response to increasing computational load on two dimensions: horizontal scaling with 2 to 20 parties making use of the system and vertical scaling with 100MB to 5GB of data. Individual scalability tests measured the execution time, memory use, CPU use, and network throughput, and averaged them taken in a 50-iteration range to guarantee the reliability of measurements. Instrumentation was collected by loading performance metrics into the application code, and topped by the Google Cloud Monitoring to provide infrastructure-level metrics.

Security Validation Framework

In security assessment, a broad methodology was used that included the combination of formal verification, penetration testing, and compliance auditing. Formal verification used prover if protocol verifier to test the cryptographic protocols against possible vulnerabilities and model checking using TLA+ used to verify that the system was correct in all possible

execution paths. Penetration testing included API interface fuzz testing, side-channel testing which includes timing attacks and information leakage using cache, and simulated adversary environments, which involved malicious cloud providers and compromised host environments.

The remote attestation protocol was then validated extensively by testing it with 10,000 successive rounds of attestation, the measure of success, rates of various false-positive, and consistency in the performance. Checks against cryptographic implementation involved entropy checks on random number generation, key-strength checking and resistance testing against common cryptographic attack vectors. The assessment of compliance aligned deployed security controls with current industry standard requirements like the NIST Cybersecurity Framework, HIPAA security requirements and the GDPR data protection requirements.

Comparative Analysis Methodology

The comparative analysis contrasted the suggested solution with three other privacy-saving methods that include the usage of traditional cloud computing using basic encryption, homomorphic encryption, and secure multi-party computation. All possible options were tested based on the standard best-practice recommendations and evaluated on the same datasets and performance indicators. The comparison focused on four dimensions of the level of security assurances provided, level of performance overhead incurred, complexity of the operation and the ability to be applied in the actual situations.

Reliability and Fault Tolerance Assessment

System reliability was tested by constantly monitoring the operations of the system over a period of thirty days, with respect to availability, mean time between failure and recovery performance. Fault-injection testing modeled the various types of failures such as network partitions, storage failures, memory corruption and cryptographic errors. Mechanisms of recovery were authenticated with automated failover testing and manual drills of disaster-recovery, recovery rates and recovery time was being tracked in each scenario.

Validation and Reproducibility Measures

All tests were done under controlled environmental conditions to ensure that there was experimental validity and reproducibility of the tests that did not realize undue interference. Statistical analysis was used in performance measurements through the calculation of confidence intervals and outliers to ensure

that the results are reliable. The code validation used a complete test suite with 94.3 percentage code and 91.8 per cent branch coverage, and all the experimental results were confirmed by various team members. The empirical findings have been exhaustively described in the manuscript with respect to the experimental methodology, data generation and measurement protocols in order to enable independent replication of the findings. Such a strict approach of the methodology guarantees that the findings are scientifically valid, statistically significant and can be applied to realistic operation settings. The multi-faceted assessment system provides a characterization of both the performance and security characteristics of the suggested confidential computing solution in a detailed way and, therefore, facilitates an extensive evaluation of its applicability to secure multi-party analytics.

Experimental Setup

Test Environment Configuration

- **Cloud Provider:** Google Cloud Platform
- **Instance Type:** n2d-standard-8 (AMD EPYC 7B12, 8 vCPUs, 32GB RAM)
- **Confidential Computing:** AMD SEV-SNP enabled
- **Operating System:** Ubuntu 20.04 LTS
- **Storage:** Balanced Persistent Disk with 500 GB capacity
- **Network Configuration:** Premium Tier, 10 Gbps throughput

Dataset Characteristics

We utilized three distinct datasets to evaluate system performance across different scenarios:

1. Healthcare Dataset: Synthetic patient records (100,000 - 1,000,000 records)

- **Fields:** Patient ID, Age, Diagnosis Codes, Treatment Costs, Lab Results
- **Size:** 50MB - 500MB
- **Sensitivity:** High (simulated PHI)

2. Financial Dataset: Synthetic transaction records

- **Fields:** Transaction ID, Amount, Merchant, Location, Timestamp
- **Size:** 100MB - 1GB
- **Sensitivity:** Medium-High

3. Research Dataset: Academic collaboration data

- **Fields:** Researcher ID, Publications, Citations, Collaborations
- **Size:** 25MB - 250MB
- **Sensitivity:** Medium

Performance Evaluation

1. Encryption and Data Transfer Performance

Table 1: Data Encryption Performance (AES-256-GCM)

Data Size	Encryption Time (ms)	Decryption Time (ms)	Throughput (MB/s)
10 MB	45.2 ± 2.1	42.8 ± 1.9	221.2 ± 10.3
50 MB	198.7 ± 8.5	192.3 ± 7.2	257.4 ± 11.1
100 MB	385.4 ± 12.3	378.9 ± 11.6	259.3 ± 8.3
500 MB	1892.6 ± 45.7	1876.3 ± 42.1	264.2 ± 6.4

Observation: In table 1 the encryption system maintained a steady throughput rate of about 260MB/s with different data sizes, which demonstrates the efficacy of cryptographic operations and low performance loss.

2. Enclave Operations Overhead

Table 2: Enclave Operation Latency Comparison

Operation	Standard VM (ms)	Confidential VM (ms)	Overhead (%)
Memory Allocation	0.45 ± 0.02	0.52 ± 0.03	15.6%
Data Sealing	12.3 ± 0.5	14.8 ± 0.6	20.3%
Attestation	N/A	156.7 ± 8.9	N/A
Secure Analysis	45.6 ± 2.1	53.2 ± 2.4	16.7%

Observation: In table 2 confidential computing environment had an average of 17.5% overhead when performing operations common to all, and attestation (which is the most resource-intensive process) took 156.7ms.

3. Multi-Party Analytics Performance

Table 3: Secure Analytics Execution Time

Analytics Type	Parties	Data Size	Execution Time (s)	Traditional Approach (s)
SUM_ANALYSIS	2	100MB	0.89 ± 0.04	0.72 ± 0.03
SUM_ANALYSIS	4	500MB	3.45 ± 0.12	2.89 ± 0.10
JOIN_ANALYSIS	3	300MB	2.78 ± 0.11	2.31 ± 0.09
STATS_ANALYSIS	5	1GB	7.89 ± 0.25	6.45 ± 0.21

Observation: In table 3 secure analytics operations had an average overhead of 22.3 2 performance overhead in comparison with the traditional methods, which is considered as acceptable given the enhanced security assurances.

Security Evaluation

1. Attestation and Integrity Verification

Table 4: Attestation Performance and Security

Metric	Value	Description
Attestation Success Rate	99.98%	Successful remote verifications
False Positive Rate	0.001%	Incorrect attestation approvals
Attestation Time	156.7 ± 8.9 ms	End-to-end verification
Key Establishment	89.3 ± 4.2 ms	Secure channel setup

Observation: In table 4 the remote attestation protocol checked the integrity of the enclaves with a success rate of 99.98, and recorded insignificant false positives. This attestation process took less than 160ms, thus making it feasible in practice.

Scalability Analysis

1. Horizontal Scaling Performance

Table 5: Multi-Party Scaling Characteristics

Number of Parties	Memory Usage (MB)	CPU Utilization (%)	Processing Time (s)
2	128 ± 8	45.2 ± 3.1	0.89 ± 0.04
5	215 ± 12	68.7 ± 4.2	1.78 ± 0.08
10	389 ± 18	82.3 ± 5.1	3.45 ± 0.15
20	723 ± 32	91.5 ± 6.3	6.89 ± 0.28

Observation: In table 5 the system displayed linear scalability to a point of 10 parties, after which it slowly degraded which can be attributed to increased overhead in coordination.

Comparative Analysis

1. Performance vs Security Tradeoffs

Table 6: Comparison with Alternative Approaches

Approach	Security Level	Performance Overhead	Data Privacy
Traditional Cloud	Low	0%	Limited
Homomorphic Encryption	High	1500%	Complete
Secure MPC	High	800%	Complete

Observation:

In table 6 the confidential computing solution provides full data privacy and a significantly lower performance penalty compared to strictly cryptographic solutions, including homomorphic encryption.

real practice situations of multi-party cooperation, which often include two to eight parties. In addition, its memory usage doubles with the number of additional parties, which is solvable in modern cloud instance architectures (35MB).

DISCUSSION

Performance Analysis

Experimental findings show that the cost of the proposed methodology of confidential computing is a performance overhead that is manageable at 22.3 percent in comparison to traditional cloud analytics. This overhead is significantly reduced compared to the overhead of other privacy-preserving methods, like homomorphic encryption (1500% overhead) or secure multi-party computation (800% overhead). The primary performance costs arise from:

- **Memory Encryption:** AMD SEV-SNP memory encryption introduces approximately 15-20% performance penalty for memory-intensive operations.
- **Attestation Overhead:** The remote attestation process adds 156.7ms to initial setup, but this is a one-time cost per session.
- **Cryptographic Operations:** AES-GCM encryption/decryption accounts for the remaining performance difference.

Security Assessment

Our implementation successfully addresses the core security requirements for multi-party analytics:

- **Data Confidentiality:** All data remains encrypted during transit, at rest, and during computation using hardware-enforced memory encryption.
- **Integrity Verification:** Remote attestation provides cryptographic proof of enclave integrity before data processing.
- **Access Control:** Fine-grained authorization ensures only authorized parties can access specific analytical results.
- **Audit Trail:** Comprehensive logging of all operations enable’s compliance verification.

Scalability Implications

The system is known to be horizontally scaled, and can handle ten or more involved entities before linear degradation in apparent system performance is noticed. As a result, the method is applicable to the majority of

Practical Application

The performance and security characteristics make this approach particularly suitable for:

- **Healthcare Research:** Multi-institutional medical studies while protecting patient data under HIPAA regulations.
- **Financial Crime Detection:** Collaborative fraud analysis between banks without sharing sensitive customer information.
- **Cross-Organizational Business Intelligence:** Secure analytics across supply chain partners while protecting proprietary information.
- **Government Data Collaboration:** Inter-agency data analysis while maintaining strict access controls and audit trails.

Limitations and Future Work

While the results are promising, several areas warrant further investigation:

- **Very Large-Scale Deployments:** Performance with more than 20 parties requires optimization.
- **Complex Analytics:** Support for more sophisticated machine learning algorithms within enclaves.
- **Cross-Cloud Deployments:** Extending the approach to multi-cloud environments.
- **Quantum Resistance:** Integration of post-quantum cryptographic algorithms.

CONCLUSION

This study shows that it is viable and effective to rely on confidential computing technology to apply secure multi-party analytics to public cloud environments with special attention to the Google Cloud Platform. With a controlled introduction and thorough assessment of an extensive architecture on the foundation of AMD SEV-SNP technology, it has therefore been discovered that confidential computing offers a viable answer to the inherent dilemma of preserving privacy of data and simultaneously allowing collaborative analytics in multiple organizations.

As experimental findings suggest, the suggested strategy has achieved the best balance between security and performance, adding only 22.3% performance overhead to traditional cloud analytics and ensuring the full confidentiality of the data used in the computations at the same time. This breakthrough is an important enhancement to the existing privacy preserving techniques like homomorphic encryption, which has a computational overhead of 1500 0 and secure multi-party computation, which has an overhead of 800 0, making confidential computing a more practical alternative. The consistency in the throughput of the uniform encryption of approximately 260 MB/s at varying data sizes is a witness of the efficiency of the cryptographic implementation. The process of attestation that introduces 156.7 ms in the initialization provides the required cryptographic evidence on the integrity of the enclaves and forms the basis of trust in the multi-party settings.

The security validation proves that the framework meets the basic needs of sensitive data processing, such as memory encryption enforced by the hardware, remote attestation success rate of 99.98, fine-grained access control, and audit. These qualities are especially relevant to the controlled industries, including healthcare and finance, where adhering to standards such as HIPAA and GDPR is obligatory. Confidentiality of the system during storage, transmission and computation is a paradigm shift in cloud security and effectively prevents insider threats, infrastructure compromises and careless practices of cloud providers.

Analysis of scalability proves the efficiency of the framework in that it can handle a range of up to ten parties and that the performance declines in a linear manner, which correlates with the majority of collaborative real-life scenarios. The relatively small memory footprint of about 35MB per extra party is not an issue when it comes to using in current cloud environments, thus making implementation across a wide range of organizational sizes and data volumes an easy task.

The paper provides some essential knowledge that adds to the current body of literature on the topic of cloud security, analytics, and privacy preservation and provides a strong basis in further research and real-world application after the research. First, it shows that trusted execution environments based on hardware can be used to successfully close the tradeoff between high security assurances and reasonable performance cost. Second, it offers an approved architectural design of the implementation of secure multi-party analytics that can be applied in various domains and scenarios. Thirdly, it makes the field of confidential computing a viable alternative to entirely cryptographic methods that have carried prohibitive performance implications historically.

However, there are a few limitations that should be considered in the future. The fact that a drop in performance was seen past ten parties involved indicates that optimization is needed in very large-scale implementations. In addition, the current framework can be extended to include additional machine learning algorithms and cross-cloud environments, which, despite its support of the basic analytics operations, would make it more applicable. Lastly, the incorporation of post-quantum encryption algorithms is significant future research that should be pursued, due to the dynamism of the threat environment.

REFERENCES

1. Adelusi, B., Ojika, F., & Uzoka, A. (2022). A conceptual model for cost-efficient data warehouse management in aws, gcp, and azure environments. *International Journal of Multidisciplinary Research and Growth Evaluation*, 3(2), 831-846. <https://doi.org/10.54660/ijmrge.2022.3.2.831-846>
2. Bhalla, J. (2025). Safeguarding sensitive data with confidential computing. *World Journal of Advanced Engineering Technology and Sciences*, 15(2), 421-427. <https://doi.org/10.30574/wjaets.2025.15.2.0505>
3. Dendukuri, S. (2025). Federated learning in healthcare: protecting patient privacy while advancing analytics. *Journal of Computer Science and Technology Studies*, 7(7), 840-845. <https://doi.org/10.32996/jcsts.2025.7.7.90>
4. Eboseremen, B. (2022). Secure data integration in multi-tenant cloud environments: architecture for financial services providers. *JFMR*, 3(1), 579-592. <https://doi.org/10.54660/jfmr.2022.3.1.579-592>
5. Ejaz, U., Islam, S., Sarkar, A., & Imashev, A. (2024). Federated learning for secure and privacy-preserving medical collaboration across multi-cloud healthcare systems. *Iosr Journal of Mechanical and Civil Engineering*, 21(5), 36-44. <https://doi.org/10.9790/1684-2105023644>
6. Feng, D., Qin, Y., Feng, W., Li, W., Shang, K., & Ma, H. (2024). Survey of research on confidential computing. *Iet Communications*, 18(9), 535-556. <https://doi.org/10.1049/cmu2.12759>
7. Gao, H., Yue, C., Dinh, T., Huang, Z., & Ooi, B. (2023). Enabling secure and efficient data analytics pipeline evolution with trusted execution environment. *Proceedings of the VLDB Endowment*, 16(10), 2485-2498. <https://doi.org/10.14778/3603581.3603589>
8. Gogineni, A. (2025). Confidential computing architectures for enhanced data security in cloud environments. *IJSAT*, 16(2). <https://doi.org/10.71097/ijst.v16.i2.3172>
9. Li, X., Li, F., & Gao, M. (2023). Flare: a fast, secure, and memory-efficient distributed analytics framework. *Proceedings of the VLDB Endowment*, 16(6), 1439-1452. <https://doi.org/10.14778/3583140.3583158>

10. Malkoochi, R. (2025). Confidential computing for privacy-preserving fraud analytics. *European Journal of Computer Science and Information Technology*, 13(24), 115-228. <https://doi.org/10.37745/ejcsit.2013/vol13n24115228>
11. Rathi, N. (2025). Beyond encryption: a holistic approach to privacy-preserving query processing in modern database systems. *International Journal of Scientific Research in Engineering and Management*, 09(06), 1-9. <https://doi.org/10.55041/ijsem.49974>
12. Reddy, A. (2025). Privacy-preserving ai models for secure healthcare data in the cloud. *FTSHSL*, 3(1), 42-52. <https://doi.org/10.69888/ftshsl.2025.000363>
13. Sathar, G. (2025). Cloud computing for big data analytics: scalable solutions for data-intensive applications. *Journal of Information Systems Engineering & Management*, 10(4), 977-990. <https://doi.org/10.52783/jisem.v10i4.10181>
14. Shah, S. and Choksi, N. (2025). Confidential computing for serverless workloads: secure and scalable data processing in untrusted environments. *World Journal of Advanced Engineering Technology and Sciences*, 14(3), 086-104. <https://doi.org/10.30574/wjaets.2025.14.3.0067>
15. Sharma, A., Khilji, N., Purohit, S., & Saxena, R. (2025). Federated learning in the cloud: a study of privacy-preserving ai architectures and recent implementations across the three major cloud service providers. *Lex Localis - Journal of Local Self-Government*, 23(S4), 3467-3480. <https://doi.org/10.52152/801099>
16. Valadares, D., Will, N., Caminha, J., Perkusich, M., Perkusich, Â., & Gorgônio, K. (2021). Systematic literature review on the use of trusted execution environments to protect cloud/fog-based internet of things applications. *Ieee Access*, 9, 80953-80969. <https://doi.org/10.1109/access.2021.3085524>
17. Vanga, P. (2025). Securing patient data in healthcare cloud systems: a technical overview. *IJSAT*, 16(1). <https://doi.org/10.71097/ijst.v16.i1.2754>
18. Arivarasi, A. and Ramesh, P. (2022). Oelb - ih algorithm for secure data routing to improve the network location advisory privacy performance in wsn. *Tehnicky Vjesnik - Technical Gazette*, 29(4). <https://doi.org/10.17559/tv-20210729085331>
19. Asif-Ur-Rahman, M., Afsana, F., Mahmud, M., Kaiser, M., Ahmed, M., Kaiwartya, O., ... & James, A. (2019). Toward a heterogeneous mist, fog, and cloud-based framework for the internet of healthcare things. *Ieee Internet of Things Journal*, 6(3), 4049-4062. <https://doi.org/10.1109/jiot.2018.2876088>
20. Babenko, M., Tchernykh, A., Pulido-Gaytán, B., Avetisyan, A., Nesmachnow, S., Wang, X., ... & Granelli, F. (2022). Towards the sign function best approximation for secure outsourced computations and control. *Mathematics*, 10(12), 2006. <https://doi.org/10.3390/math10122006>
21. Casaletto, J., Cline, M., & Shirts, B. (2022). Modeling the impact of data sharing on variant classification. *Journal of the American Medical Informatics Association*, 30(3), 466-474. <https://doi.org/10.1093/jamia/ocac232>
22. Cho, H., Wu, D., & Berger, B. (2018). Secure genome-wide association analysis using multiparty computation. *Nature Biotechnology*, 36(6), 547-551. <https://doi.org/10.1038/nbt.4108>
23. Farhan, M., Ayoub, M., & AL-JADER, A. (2023). Gps-based fall detection system for old and specially-abled people. *Indonesian Journal of Electrical Engineering and Computer Science*, 31(3), 1545. <https://doi.org/10.11591/ijeecs.v31.i3.pp1545-1550>
24. Gao, H., Ma, Z., Luo, S., & Wang, Z. (2019). Bfmpc: a blockchain-based fair and robust multi-party computation scheme. *Ieee Access*, 7, 110439-110450. <https://doi.org/10.1109/access.2019.2934147>
25. Geppert, T., Deml, S., Sturzenegger, D., & Ebert, N. (2022). Trusted execution environments: applications and organizational challenges. *Frontiers in Computer Science*, 4. <https://doi.org/10.3389/fcomp.2022.930741>
26. Mehrtak, M., SeyedAlinaghi, S., MohsseniPour, M., Noori, T., Karimi, A., Shamsabadi, A., ... & Dadras, O. (2021). Security challenges and solutions using healthcare cloud computing. *Journal of Medicine and Life*, 14(4), 448-461. <https://doi.org/10.25122/jml-2021-0100>
27. Mouchet, C., Bertrand, E., & Hubaux, J. (2023). An efficient threshold access-structure for rlwe-based multiparty homomorphic encryption. *Journal of Cryptology*, 36(2). <https://doi.org/10.1007/s00145-023-09452-8>
28. Podschwadt, R., Takabi, D., Hu, P., Rafiei, M., & Cai, Z. (2022). A survey of deep learning architectures for privacy-preserving machine learning with fully homomorphic encryption. *Ieee Access*, 10, 117477-117500. <https://doi.org/10.1109/access.2022.3219049>
29. Sheller, M., Edwards, B., Reina, G., Martin, J., Pati, S., Kotrotsou, A., ... & Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10(1). <https://doi.org/10.1038/s41598-020-69250-1>
30. Sun, P. (2019). Privacy protection and data security in cloud computing: a survey, challenges, and solutions. *Ieee Access*, 7, 147420-147452. <https://doi.org/10.1109/access.2019.2946185>
31. Zobaed, S. and Salehi, M. (2025). Confidential computing across edge-to-cloud for machine learning: a survey study. *Software Practice and Experience*, 55(5), 896-924. <https://doi.org/10.1002/spe.3398>