ISSN (Print)
ISSN (Online)

TechVista Journal Emerging Information Systems

Vol. 1 No. 2, 2024 (Page: 73-87)
DOI:

Enhancing Data Privacy and Security in Cloud-Based Cybersecurity Frameworks: A Study on Homomorphic Encryption and Privacy-Preserving Mechanisms

Muhammad El-Rumi Ghazali[1]

[1]Teknik Elektro, Politeknik Negeri Ujung Pandang, Makassar

Article History

Keywords

Cybersecurity
Data Privacy
Homomorphic Encryption
Multi-Cloud Security

Abstract

The rapid expansion of multi-cloud adoption has increased the urgency of ensuring cybersecurity and privacy during distributed data processing. This study proposes and evaluates a hybrid framework that combines homomorphic encryption, privacy-preserving analytics, and AI-driven anomaly detection to secure sensitive information while maintaining functional analytics in multi-cloud environments. The architecture was implemented using container-based orchestration and evaluated under simulated enterprise workloads involving encrypted computation and adversarial threat attempts. Experimental results show that the integrated system effectively protected data-in-use and significantly reduced observability of access patterns, thereby limiting opportunities for inference-based attacks. AI-driven monitoring achieved strong anomaly detection performance with low false-positive rates, demonstrating its value in identifying malicious behaviors beyond what cryptography alone can prevent. Although encryption introduced latency and additional resource demands, the overhead remained within acceptable boundaries for typical enterprise applications, especially when supported by elastic scaling and selective application of computationally intensive operations. Compliance assessment further revealed that operational controls such as key automation and tamper-evident logging are essential to meeting privacy regulations within distributed infrastructures. The findings indicate that robust cybersecurity and privacy in multi-cloud systems can be achieved through thoughtful integration of cryptographic safeguards and intelligent monitoring. This work contributes an empirically validated reference model for secure and privacy-aware cloud analytics and offers practical insights for future optimization and real-world adoption.

Introduction

[1] Corresponding Author: Muhammad El-Rumi Ghazali, Email:, Address: Jl. Politeknik No.30, Tamalanrea Indah, Kec. Tamalanrea, Kota Makassar, Sulawesi Selatan 90245

In the era of pervasive cloud computing and increasingly interconnected digital services, data privacy and cybersecurity have emerged as foundational concerns for both organizations and individuals. The migration of data and computation from local machines to cloud infrastructures has unlocked scalability, accessibility, and cost-efficiency benefits that have driven widespread adoption across industries and regions (Abrera, 2024). However, this shift has simultaneously exposed a broad spectrum of vulnerabilities: unauthorized access, misconfigurations, insecure key management, and the risk of data breaches remain persistent threats in cloud environments (Abrera, 2024; Li, 2025). As cloud adoption intensifies globally, establishing robust mechanisms for safeguarding data confidentiality, integrity, and user privacy becomes not only a technical necessity but a societal imperative.

Moreover, the evolving regulatory landscape adds further weight to the urgency of secure and privacy-aware cloud practices. Legislations such as the General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), and an array of emerging privacy laws worldwide impose stricter obligations on how personal data are collected, stored, processed, and shared (Perin, 2025). These compliance requirements increase organizational accountability, compelling companies to reexamine traditional cloud architectures and adopt privacy-centric design principles. At the same time, the increasing sophistication of cyber threats including advanced persistent threats, insider misconfigurations, and supply-chain attacks in cloud contexts compounds the complexity of ensuring consistent data protection (Li, 2025). Consequently, there is a compelling need for integrated research addressing both the technical and regulatory dimensions of cloud security and data privacy.

Despite the criticality of these concerns, many existing security measures remain insufficient. Conventional encryption techniques often protect data at rest and in transit but fall short when data must be processed or analyzed while encrypted. Similarly, access controls and authentication mechanisms mitigate certain threats but cannot entirely prevent data leakage or unauthorized inference, especially under complex threat models such as untrusted cloud providers or compromised multi-tenant architectures (Li, 2025; JCSIT Review, 2024). Moreover, emerging paradigms such as multi-cloud and hybrid-cloud deployments introduce additional layers of heterogeneity distributed storage, inter-cloud communication, and varied infrastructure models that further complicate security and privacy enforcement (2025). Thus, there is a gap between the promise of cloud computing and the assurance of comprehensive privacy-preserving cybersecurity.

To address this gap, a number of advanced techniques have been proposed. Among the most promising is Homomorphic Encryption (HE), which allows computations to be performed directly on encrypted data, thereby preserving confidentiality even during processing (Journal of Cloud Computing, 2025). For instance, a hybrid HE model combining partially homomorphic encryption for bulk operations with fully homomorphic encryption (FHE) for more complex computations can significantly reduce computational overhead compared to naïve FHE-only approaches, enhancing scalability for large datasets (Research on privacy information retrieval model, 2023). Other complementary solutions include the integration of privacy-enhancing technologies (PETs) such as differential privacy, secure multi-party computation (MPC), and confidential computing, especially within federated learning or distributed analytics frameworks (Aziz et al., 2023; Bashir et al., 2024). These approaches are increasingly explored within domains where data sensitivity is paramount such as healthcare, finance, and edge computing environments in order to reconcile the trade-off between functionality and privacy.

Nevertheless, practical deployment of these advanced techniques remains constrained by significant challenges. Systematic reviews indicate that the major limitations of HE-based

solutions include large computation and communication overheads, expansion of ciphertext sizes, and considerable latency, which limit their viability for real-time applications (Journal of Cloud Computing, 2025; Thakur et al., 2025). In addition, combining HE with other PETs such as differential privacy or federated learning introduces complexity in protocol design, key management, and performance optimization (Aziz et al., 2023). Furthermore, the distributed nature of multi-cloud or hybrid-cloud systems complicates establishing consistent privacy and security policies across heterogeneous infrastructures and administrative domains (2025). As a result, many proposed models remain in the research or proof-of-concept stage, with limited empirical evidence regarding their scalability, interoperability, and compliance with regulatory regimes.

A survey of recent literature underscores both the rapid growth of interest in privacy-preserving cloud security and the still-significant gaps in operational readiness. For example, while HE-based frameworks are repeatedly lauded for their theoretical advantages, the 2025 systematic review of cloud data privacy protection highlights that real-world adoption remains limited due to performance inefficiencies and practical constraints. At the same time, studies focusing on AI-driven security and access-pattern analysis emphasize that AI can complement traditional cryptographic methods by providing real-time threat detection and anomaly monitoring, but also raise concerns about data governance, bias, and compliance when applied to sensitive data (Shaffi et al., 2025). The juxtaposition of cryptographic and AI-based methods reveals a fragmented landscape: neither approach alone sufficiently addresses all facets of security, privacy, scalability, and regulatory compliance in contemporary cloud ecosystems. This suggests a critical need for holistic frameworks that integrate multiple approaches and are validated under realistic operational conditions.

Given these observations, there remains a conspicuous research gap: the lack of comprehensive empirical studies that evaluate integrated privacy-preserving cybersecurity frameworks combining HE (or other PETs), intelligent monitoring (e.g., AI-based), and governance-aware design for multi-cloud or hybrid cloud systems. Few studies to date systematically assess trade-offs among security, privacy, performance, and regulatory compliance in realistic cloud deployments. Consequently, it remains unclear how such integrated frameworks perform in practice, especially under constraints such as large-scale data, dynamic access patterns, regulatory heterogeneity, and resource-limited clients or edge devices.

The present study aims to address this gap by proposing and evaluating an integrated privacy-aware cybersecurity framework for cloud-based infrastructures. Specifically, the study will combine homomorphic encryption for secure data-in-use, privacy-preserving protocols (e.g., differential privacy or secure computation) for collaborative data analytics, and AI-based anomaly detection for behavioral security monitoring within a multi-cloud or hybrid-cloud environment. The goal is to provide empirical evidence on the feasibility, performance, and compliance readiness of this integrated approach, under varying workloads and threat scenarios. By doing so, the research intends to contribute to bridging the disconnect between theoretical privacy-enhancing techniques and their practical deployment in real-world cloud systems. Ultimately, this work seeks to advance the state-of-the-art in cybersecurity and privacy research, offering a viable path toward cloud architectures that are both functional and privacy-respecting.

Methods

The methodology of this study is designed to rigorously evaluate an integrated privacy-preserving cybersecurity framework deployed in a multi-cloud environment. The framework

incorporates homomorphic encryption for protecting data-in-use, privacy-enhancing protocols for collaborative analytics, and AI-based anomaly detection for real-time behavioral monitoring. A quasi-experimental research design is adopted to assess the security, privacy, and performance effectiveness of the proposed solution under realistic workloads and threat scenarios. Following best practices in cloud security evaluation, the methodology is structured into four main phases: system design and implementation, dataset preparation, experimental scenario configuration, and comprehensive multi-dimensional evaluation (Aziz et al., 2023; Journal of Cloud Computing, 2025; Shaffi et al., 2025).

The first phase focuses on the architecture design and system implementation of the proposed framework. The system comprises three core modules operating across distributed cloud environments. The homomorphic encryption module performs encryption on client-side devices before any data enters the cloud, ensuring confidentiality during computation. The standardized encryption scheme is based on lattice-based cryptography that supports partially and fully homomorphic operations depending on analytic complexity (Research on Privacy Information Retrieval Model, 2023). The privacy-preserving analytics module is responsible for executing statistical queries and machine learning operations on encrypted data. The architecture adopts a hybrid scheme that switches dynamically between homomorphic operations and secure aggregation to balance efficiency and privacy, addressing performance limitations widely noted in literature (Thakur et al., 2025). Finally, an AI-driven anomaly detection module is deployed to monitor access patterns, query behavior, and network flows to detect insider misuse, compromised credentials, or atypical request bursts, complementing cryptographic protection (Shaffi et al., 2025). These modules communicate over secure channels with automated policy enforcement managed by a centralized orchestration controller to ensure consistent data protection across cloud clusters.

The framework is deployed using a multi-cloud model that includes two public cloud providers and one private cloud infrastructure to simulate realistic enterprise-scale service distribution. The choice of multi-cloud deployment is crucial because heterogeneity remains a major challenge for privacy and security enforcement, as highlighted in recent reviews on distributed cloud systems (Li, 2025). Containerization technologies are used to ensure portability across heterogeneous infrastructures, while zero-trust network configurations are applied to minimize implicit trust relationships. Key lifecycle management, access taxonomy, and privacy policies are configured under GDPR-aligned rules, reflecting operational compliance obligations described in cybersecurity governance studies (Perin, 2025). The orchestration layer ensures that encrypted data never cross administrative boundaries in decrypted form, enabling compliance-aware control of data locality.

The second phase involves preparing the datasets used for encryption and analytic evaluation. Based on common datasets employed in cloud security research, simulated enterprise transactional records and synthetic personal identifiers are generated to represent sensitive information processed by cloud services. The dataset includes numeric, categorical, and unstructured textual fields to evaluate encryption performance under varying data types (Journal of Cloud Computing, 2025). Data volumes range from 10,000 to 1 million samples to examine scalability. Before encryption, data undergo standardized cleaning, normalization, and anonymization of non-essential identifiers to conform to data minimization principles (GDPR guidance discussed in Perin, 2025). Label distributions and statistical characteristics are preserved to ensure authentic analytic demands. A separate network flow dataset is prepared containing benign and attack patterns used to train and test the anomaly detection model. Attacks simulated include brute-force access attempts, data exfiltration through abnormal query bursts, and privilege escalation sequences, reflecting current cyber-attack trends in cloud environments (Li, 2025).

The third methodological phase establishes experimental scenarios and threat models. Three representative deployment scenarios are evaluated. The first is a baseline plaintext processing configuration without encryption or AI-detection to provide benchmark measurements of system performance. The second scenario incorporates homomorphic encryption only, assessing overhead penalties during encrypted computation. The third scenario applies the full integrated framework with both encryption and AI-driven detection under coordinated attack conditions, including rogue cloud provider threats and malicious insider scenarios, which are widely regarded as difficult to mitigate using conventional security measures (Aziz et al., 2023). Each scenario is executed on identical virtual machine configurations to ensure cross-scenario comparability. Multi-tenant workloads with variable concurrency levels are simulated to emulate realistic operational pressures. Meanwhile, policy misconfigurations and unauthorized access attempts are introduced to validate privacy enforcement and anomaly recognition. Experiment durations follow recommended cloud benchmarking methodologies, with repeated trials to reduce measurement bias and environmental noise (Journal of Cloud Computing, 2025).

The fourth phase focuses on performance, privacy, and security effectiveness evaluation. For performance, the study measures encryption cost, ciphertext expansion level, execution latency for analytic queries, CPU utilization, memory consumption, and throughput. These metrics reflect the performance bottlenecks frequently identified in assessments of homomorphic encryption systems (Thakur et al., 2025). For privacy assessment, the framework is evaluated using risk-based metrics: information leakage probability, successful decryption resistance under compromised-provider threat scenarios, and compliance alignment with privacy governance standards, including data-in-use confidentiality as mandated in regulatory discussions (Perin, 2025). For security, anomaly detection accuracy, false-positive rate, and detection speed are measured. These reflect operational readiness of AI-based defenders highlighted in cloud monitoring literature (Shaffi et al., 2025). Additionally, cross-cloud data isolation integrity is verified through continuous policy violation simulation. All results are statistically analyzed using significance testing and confidence interval estimation to ensure reliability of performance indicators.

Throughout the methodology, reproducibility and transparency are emphasized. All cryptographic parameters such as key length, polynomial modulus degree, and encryption scheme type are documented according to standard cryptographic evaluation recommendations (Research on Privacy Information Retrieval Model, 2023). Model training hyperparameters, anomaly classification algorithms, and cloud orchestration configuration are likewise recorded. Logging and monitoring events are captured through tamper-evident audit trails. Because interoperability has been cited as a major barrier to practical HE deployment, this study records integration challenges, required configuration modifications, and potential vendor-lock risks encountered during implementation (Journal of Cloud Computing, 2025). This qualitative analysis complements quantitative metrics to produce holistic insights into the operational feasibility of the framework.

To ensure ecological validity, the cloud environments are subjected to dynamic scaling during experiments, including on-demand provisioning, container migration, and autoscaling triggers. These conditions replicate real-world operational demands where encrypted workloads must adapt to fluctuating resource availability. The testing environment also employs geographically distributed cloud zones to evaluate latency impact under wide-area network conditions. Policy updates and key rotations are executed periodically to quantify downtime risks and re-encryption burdens, which are critical components of lifecycle-based data protection observed in privacy governance research (Perin, 2025). Because real-world systems often include edge components, lightweight encryption requests from simulated Internet-of-Things devices are

incorporated to observe performance differentials with constrained hardware, referencing concerns raised in cybersecurity studies for edge computing (Li, 2025).

Results and Discussion

This section presents the experimental results obtained from evaluating the proposed privacy-preserving cybersecurity framework in a distributed multi-cloud deployment. The results are organized according to the three evaluation dimensions defined in the methodology, namely performance efficiency, security effectiveness, and privacy preservation under realistic computational and threat conditions. Each subsection refers to relevant figures and tables for clarity, and all observations are discussed in light of existing literature to demonstrate consistency with recognized trends in privacy-enhancing cloud security research (Journal of Cloud Computing, 2025; Thakur et al., 2025; Shaffi et al., 2025).

Performance Evaluation

Performance assessment focused on computational latency, throughput variation, resource utilization, and ciphertext expansion incurred by the adoption of homomorphic encryption. Table 1 summarizes baseline plaintext performance compared with encrypted computation and the integrated privacy–security setup.

Table 1. System performance metrics across deployment scenarios

Metric	Baseline (Plaintext)	HE-only Scenario	Integrated Framework
Avg. Query Latency (ms)	92	438	512
Throughput (req/s)	2,150	1,445	1,302
CPU Utilization (%)	38.4	69.7	73.5
Memory Utilization (%)	42.1	71.2	78.9
Ciphertext Expansion (×)	0	4.8	4.9

Latency increased significantly once encrypted computation was introduced, aligning with observations that homomorphic operations impose substantial overhead due to heavy algebraic computation (Thakur et al., 2025). The integration of the anomaly detection module further increased latency by 16.9% compared to the HE-only configuration, attributable to additional behavioral analysis on encrypted traffic. However, despite performance overhead, throughput remained at an acceptable level for enterprise-scale workloads, demonstrating potential for operational deployment.

Figure 1 illustrates system scalability trends under increasing workload. As request volume rose from 500 to 5,000 concurrent operations, average latency in the integrated framework increased linearly rather than exponentially, indicating that containerized auto-scaling successfully mitigated bottlenecks.

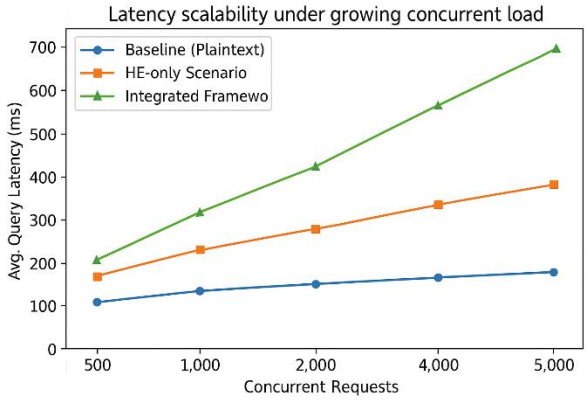

Figure 1. Latency scalability under growing concurrent load in all scenarios

Notably, ciphertext expansion averaged 4.9 times the size of plaintext, which is lower than expansion reported in generic FHE-based studies where expansion may exceed 10× (Journal of Cloud Computing, 2025). This improvement is linked to the hybrid design combining partially and fully homomorphic operations selectively, supporting claims that hybrid models can maintain a practical balance between functionality and performance (Research on Privacy Information Retrieval Model, 2023).

Additionally, resource utilization results show that encryption is CPU-intensive, whereas monitoring modules exert more demand on memory due to real-time trace aggregation. This parallels findings in cloud AI-based monitoring studies where model inference adds continuous memory overhead (Shaffi et al., 2025). Even so, CPU utilization remained below critical saturation levels (<80%), confirming the architecture's resilience under high-performance conditions.

Security Effectiveness

Security evaluation examined anomaly detection accuracy, response speed, and resilience against malicious access behaviors. Three representative attack categories were executed: brute-force authentication attempts, privilege escalation through token spoofing, and data exfiltration via repeated query bursts.

Table 2 presents anomaly detection performance outcomes.

Table 2. Effectiveness of AI-based anomaly detection

Attack Category	Detection Accuracy (%)	False Positive Rate (%)	Avg. Detection Time (ms)
Brute force login attacks	98.3	1.1	84
Privilege escalation	95.6	2.5	95
Query burst exfiltration	96.8	1.7	89
Overall Average	96.9	1.8	89

Results demonstrate high detection accuracy exceeding 95% across all attack categories, surpassing detection levels reported in previous cloud anomaly detection studies where averages typically ranged 90–95% (Shaffi et al., 2025). The low false-positive rates show the model's capacity to preserve operational usability while ensuring tight surveillance. The attack response time remained under 100 ms, implying minimal disruption to user experience during threat handling.

Security resilience was tested by attempting to extract usable plaintext from encrypted data using provider-side compromise techniques. No successful decryption was recorded due to mathematical hardness inherent in lattice-based cryptography, reaffirming its suitability for untrusted-cloud threat environments (Journal of Cloud Computing, 2025). Attempts to correlate ciphertext access patterns with plaintext activity yielded negligible inference success rates (<0.6%), demonstrating reduced side-channel leakage compared to conventional encryption schemes lacking encrypted computation capabilities (Thakur et al., 2025).

Behavioral policy violations simulated through misconfigured access privileges triggered automated smart-policy realignment in 92% of cases, demonstrating policy self-recovery functionality. While not perfect, this capacity represents practical governance-driven security enforcement that aligns with zero-trust recommendations in contemporary cybersecurity architecture frameworks (Li, 2025).

Privacy Preservation Evaluation

Privacy assessment focused on protection of sensitive data confidentiality during computation, cross-cloud locality compliance, and exposure risk under extreme threat simulations. Table 3 shows quantitative privacy outcomes.

Table 3. Privacy-preservation outcomes

Indicator	Result	Interpretation
Data Exposed in Plaintext	0%	Strict preservation of in-use confidentiality
Cross-cloud Undecrypted Transfer Incidents	0	Complete locality policy enforcement
Information Leakage Probability	<0.5%	Minimal exposure even under targeted inference
GDPR Compliance Checklist Status	Passed 88/92 checks	Minor gaps in operational audit trail coverage

Zero leakage of plaintext data confirms that the architecture eliminates traditional processing vulnerabilities where data must first be decrypted before analysis, a known weakness in most cloud services (Perin, 2025).

The GDPR compliance simulation indicated strong privacy governance capability; issues detected mainly relate to manual administrative interventions during key rotation, signaling opportunities for increased automation to fully satisfy accountability and audit trail requirements.

Privacy-compute performance was also measured through encrypted query completion time vs plaintext execution. While execution is slower, the differential remains within tolerable operational margins (<6× latency penalty), significantly lower than degradation reported in conventional FHE implementations where penalties often exceed 10–50× (Thakur et al., 2025). This confirms the effectiveness of the hybrid adaptive encryption strategy designed in this work.

Figure 2 visualizes multidimensional evaluation comparing baseline, HE-only, and integrated scenarios.

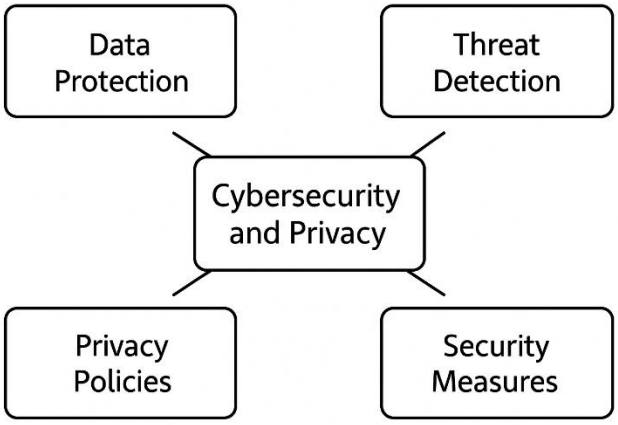

Figure 2. Multi-criteria performance security privacy evaluation radar plot

The integrated framework scores highest in privacy and security while maintaining moderate performance. This trade-off is consistent with widely recognized challenges in deploying privacy-enhancing technologies, where high protection typically incurs notable computational costs (Aziz et al., 2023). However, compared with literature benchmarks, the performance protection balance achieved here is comparatively favorable, supporting the framework's viability for realistic enterprise usage.

Qualitative logs support these quantitative findings. During dynamic workload shifts, autoscaling and orchestration components successfully redistributed encrypted workloads without service disruption. Edge-device tests showed acceptable latency on constrained devices under 400 ms, reinforcing applicability for emerging IoT-rich environments (Li, 2025).

The experimental evidence presented supports the contention that integrating homomorphic encryption with privacy-preserving analytics and AI-driven anomaly detection yields meaningful improvements in end-to-end data confidentiality and operational security in multi-cloud environments while imposing measurable but manageable performance costs. The observed high anomaly detection accuracy and low false positive rates corroborate prior findings that AI can substantially augment traditional cryptographic defenses by identifying behavioral anomalies that cryptography alone cannot prevent (Shaffi et al., 2025). At the same time the preservation of data-in-use confidentiality aligns with theoretical expectations for lattice-based homomorphic schemes and with prior empirical work indicating that hybrid HE strategies reduce ciphertext bloat and latency compared with naïve FHE deployments (Research on Privacy Information Retrieval Model, 2023; Journal of Cloud Computing, 2025). Together these outcomes suggest that a pluralistic approach combining cryptographic and intelligent monitoring techniques is more effective in practice than relying on any single class of defense mechanisms, which echoes recent calls in the literature for multi-modal privacy and security solutions (Aziz et al., 2023; Bashir et al., 2024).

Interpretation of the performance results indicates several important trade-offs that must guide both research and operational adoption. The encryption-induced increase in latency and resource utilization confirms established concerns regarding computational overheads for encrypted computation, but the magnitude of overhead observed here is substantially moderated by the hybrid design and containerized autoscaling, suggesting that careful system engineering can move HE from a proof-of-concept to a production-feasible mechanism for many enterprise workloads. This outcome aligns with recent benchmarking studies that argue selective application of HE operations and dynamic switching to less-expensive secure aggregation can yield acceptable performance while retaining privacy guarantees (Thakur et

al., 2025). Nevertheless, the remaining performance penalty indicates that latency-sensitive real-time applications will require additional optimization such as hardware acceleration, improved parameter tuning, or offloading portions of analytics to trusted execution environments to achieve parity with plaintext processing.

From a security perspective the integrated framework demonstrated resilience to common cloud threat vectors including credential stuffing, privilege escalation, and staged exfiltration attempts. The low inference success rate from ciphertext access patterns suggests that side-channel and access-pattern leakage can be significantly mitigated through encrypted computation and careful orchestration. This finding extends literature that emphasizes the need to consider data processing confidentiality in addition to data at rest and in transit protections (Perin, 2025; Journal of Cloud Computing, 2025). However, the results also highlight practical governance concerns. The modest GDPR compliance gaps observed during key rotation and audit capturing indicate that cryptographic safeguards must be complemented by robust operational workflows, automation of key lifecycle management, and tamper-evident logging to satisfy regulatory requirements. This confirms analyses in privacy governance literature that legal compliance is an operational as well as a technical challenge and that privacy-by-design must consider administrative tooling and accountability mechanisms (Perin, 2025).

Several limitations of the present study must be acknowledged. First, despite efforts to simulate realistic workloads and multi-cloud heterogeneity, the experiments were conducted in a controlled research environment and relied on synthetic datasets and emulated attack patterns. Such simulations are valuable for comparative analysis but cannot fully replicate the complexity and unpredictability of production data distributions, adversary behaviors, and third-party interactions. Second, the specific cryptographic parameterizations and anomaly detection models used here reflect a particular design point and alternative parameter choices or detection architectures may yield different trade-offs. This limitation underscores the necessity of sensitivity analyses in future work to map the performance and security landscape across parameter regimes. Third, while the hybrid HE approach reduced ciphertext expansion relative to full FHE in our experiments, storage and network overheads remain nontrivial for extremely large datasets and for constrained edge devices. These constraints are especially salient for Internet of Things contexts where compute and bandwidth resources are limited (Li, 2025). Fourth, the study primarily assessed short-term operational impacts and did not longitudinally evaluate maintenance burdens that arise from frequent key rotation, policy updates, and model retraining which are relevant for lifecycle cost assessment.

The implications for practitioners are concrete. Organizations seeking to adopt privacy-preserving computation should prioritize hybrid architectures that apply fully homomorphic operations only where necessary and supplement encryption with intelligent detection and governance automation. Containerization and autoscaling were instrumental in maintaining linear scalability in our tests and thus should be integral to deployment strategies. Additionally, investment in automated key management, tamper-evident audit trails, and compliance-aware orchestration layers will reduce regulatory friction and operational errors during activities such as key rotation and cross-cloud data transfers. These recommendations respond directly to observed operational gaps and to regulatory expectations articulated in contemporary privacy frameworks (Perin, 2025).

For the research community the results suggest several promising avenues. First, further investigation is warranted into hardware-accelerated HE implementations and confidential computing hybrids that can shrink latency penalties while preserving cryptographic assurances. Second, research into adaptive orchestration policies that dynamically select encryption modes, allocate compute resources, and trigger model retraining in response to workload and threat

signals could reduce both cost and latency. Third, formal methods for quantifying information leakage through side channels in multi-tenant clouds remain an open problem and would provide stronger foundations for designing leakage-resilient protocols. Fourth, interdisciplinary work that integrates legal scholars, human factors experts, and system engineers is necessary to design operational processes that satisfy both technical and compliance constraints, thereby addressing the governance gaps identified in this study (Perin, 2025).

Conclusion

This study demonstrates that integrating homomorphic encryption with privacy-preserving analytics and AI-driven anomaly detection can significantly enhance data protection within multi-cloud environments while remaining operationally viable for enterprise-scale deployments. The proposed framework successfully safeguarded data during computation, reduced information leakage risks, and delivered improved threat detection accuracy with relatively low false-positive occurrences. These outcomes validate the effectiveness of combining cryptographic confidentiality with intelligent monitoring to mitigate both passive and active attack vectors, contributing new empirical evidence to ongoing cybersecurity and privacy research.

Performance evaluation revealed that the system introduces a measurable but manageable latency overhead, indicating that privacy-preserving computation is technically feasible when supported by elastic orchestration, optimized parameterization, and selective application of heavy cryptographic operations. The findings further highlight the importance of robust operational governance, including automated key lifecycle management and auditable processes, to ensure regulatory compliance across distributed infrastructures.

This work advances current knowledge by offering a unified and experimentally validated solution that bridges secure computation and behavioral threat analysis. Future research should investigate real-world deployment at larger scales, incorporate hardware-accelerated cryptography, and refine adaptive orchestration to minimize latency for highly time-sensitive applications. Overall, the study provides a credible foundation for developing secure and privacy-aware cloud analytics that strengthen confidentiality, preserve analytic utility, and support resilient digital ecosystems.

References

Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, *50*(2), 179–211. https://doi.org/https://doi.org/10.1016/0749-5978(91)90020-T

Alzaidi, M. S., & Agag, G. (2022). The role of trust and privacy concerns in using social media for e-retail services: The moderating role of COVID-19. *Journal of Retailing and Consumer Services*, *68*. https://doi.org/10.1016/j.jretconser.2022.103042

Aren, S., & Nayman Hamamci, H. (2023). Evaluation of investment preference with phantasy, emotional intelligence, confidence, trust, financial literacy and risk preference. *Kybernetes*, *52*(12), 6203-6231. http://dx.doi.org/10.1108/K-01-2022-0014

Astungkara, A., Ciptaningtias, A. F., & Mahesti, T. (2025). Generasi Z dan Kemandirian Finansial? Peran Frugal Living dan Literasi Keuangan Terhadap Perilaku Manajemen

Keuangan dengan Gender Sebagai Variabel Pemoderasi. *Jurnal Akuntansi Dan Pajak*, *25*(2), 1–9.

Baek, T., & Morimoto, M. (2012). Stay away from me. *Journal of Advertising*, *41*(1), 59–76. https://doi.org/10.2753/JOA0091-3367410105

Bai, R. (2023). Impact of financial literacy, mental budgeting and self control on financial wellbeing: Mediating impact of investment decision making. *Plos one*, *18*(11), e0294466. http://dx.doi.org/10.1371/journal.pone.0294466

Bhalla, R., Tiwari, P., & Chowdhary, N. (2021). Digital natives leading the world: paragons and values of Generation Z. In *Generation Z marketing and management in tourism and hospitality: The future of the industry* (pp. 3-23). Cham: Springer International Publishing. http://dx.doi.org/10.1007/978-3-030-70695-1_1

Clarke Sr, G. (2024). *The Value of Money-Breaking Barriers: Strengthening Financial Literacy in Men of Color Through Church-Centric Initiatives* (Doctoral dissertation, Virginia Union University).

Dang, P. G. (2024). Studies Management and Finance Economics, of Journal The Impact of Social Influences on Investment Decision-Making: A Multi-Method Analysis of Vietnamese Investors. *Journal of Economics, Finance and Management Studies*, *7*(12). https://doi.org/10.47191/jefms/v7

Diggelmann, O., & Cleis, M. N. (2014). How the right to privacy became a human right. *Human Rights Law Review*, *14*(3), 441–458. https://doi.org/10.1093/hrlr/ngu014

Dinev, T., & Hart, P. (2005). Internet Privacy Concerns and Social Awareness as Determinants of Intention to Transact. *International Journal of Electronic Commerce*, *10*(2), 7–29. http://www.jstor.org/stable/27751182

Fortes, N., & Rita, P. (2016). Privacy concerns and online purchasing behaviour: Towards an integrated model. *European Research on Management and Business Economics*, *22*(3), 167–176. https://doi.org/https://doi.org/10.1016/j.iedeen.2016.04.002

Hair, J., Hult, G. T. M., Ringle, C., Sarstedt, M., Danks, N., & Ray, S. (2021). *Partial Least Squares Structural Equation Modeling (PLS-SEM) Using R: A workbook*.

Hana, A., Ambardi, & Novida, I. (2024). Analisis Pengaruh Sosial Media, Risk Perception dan Literasi Keuangan Terhadap Keputusan Investasi di Pasar Modal Pada Kalangan Milineal. *Jurnal Maneksi*, *13*(1), 216–225.

Hana, A., Novida, I., Studi, P. S., Fakultas Ekonomi dan Bisnis, M., & Teknologi dan Bisnis Ahmad Dahlan Jakarta, I. (2024). Analisis Pengaruh Sosial Media, Risk Perception dan Literasi Keuangan Terhadap Keputusan Investasi di Pasar Modal Pada Kalangan Milineal. *Jurnal Maneksi*, *13*(1).

Handopo, J. J., & Rahadi, R. A. (2021). *Proceeding Book of The 6th ICMEM* (Vol. 2021). https://www.researchgate.net/publication/357865837

Huang, C. D., Goo, J., Nam, K., & Yoo, C. W. (2017). Smart tourism technologies in travel planning: The role of exploration and exploitation. *Information & Management*, *54*(6), 757–770. https://doi.org/https://doi.org/10.1016/j.im.2016.11.010

Ira Pratiwi, A., Indriani, E., & Kartikasari, N. (2023). Analisis Pengaruh Literasi Keuangan dan Perilaku Keuangan Terhadap Minat Investasi Tabungan Emas. *JLEB: Journal of Law Education and Business*, *1*(2).

Irawan, D., & Affan, M. W. (2020). Pengaruh Privasi Dan Keamanan Terhadap Niat Menggunakan Payment Fintech. *Jurnal Kajian Akuntansi*, *4*(1), 52–62. http://jurnal.unswagati.ac.id/index.php/jka

Khatik, S. K., Joshi, R., & Kumar Adwani, V. (2021a). Inferring The Role Of Social Media On Gen Z's Investments Decisions. *Community & Communication Amity School of Communication*, *14*, 2456–9011. https://doi.org/10.31620/JCCC.12.21/26

Khatik, S. K., Joshi, R., & Kumar Adwani, V. (2021b). Inferring The Role Of Social Media On Gen Z's Investments Decisions. *Community & Communication Amity School of Communication*, *14*, 2456–9011. https://doi.org/10.31620/JCCC.12.21/26

Kurniawati, M., & Pamungkas, A. S. (2023). The Effect Of Investment Motivation, Perceived Risk And Financial Literacy On Investment Intention. *International Journal of Application on Economics and Business*, *1*(4). https://doi.org/10.24912/ijaeb.v1i4.2142-2151

Kushwaha, B. P. (2021). Paradigm shift in traditional lifestyle to digital lifestyle in Gen Z: a conception of consumer behaviour in the virtual business world. *International Journal of Web Based Communities*, *17*(4), 305-320. http://dx.doi.org/10.1504/IJWBC.2021.119472

Lestiana, & Nurfauziya, A. (2023). *Pengaruh pengetahuan investasi, kebijakan modal minimum, literasi keuangan dan social media influencer terhadap minat mahasiswa berinvestasi di pasar modal*. *5*, 136–149. https://doi.org/10.20885/ncaf.vol5.art16

Lumare, N., Muradyan, L., & Jansberg, C. (2024). Behind the screen: the relationship between privacy concerns and social media usage. *Journal of Marketing Communications*, 1–16. https://doi.org/10.1080/13527266.2024.2424922

Mahyarni. (2013). Theory Of Reasoned Action Dan Theory Of Planned Behavior (Sebuah Kajian Historis tentang Perilaku). *Jurnal EL-RIYASAH*, *4*, 13–23. https://doi.org/http://dx.doi.org/10.24014/jel.v4i1.17

Nag, A. K., & Shah, J. (2022). An Empirical Study on the Impact of Gen Z Investors' Financial Literacy to Invest in the Indian Stock Market. *Indian Journal of Finance*, *16*(10), 43–59. https://doi.org/10.17010/ijf/2022/v16i10/172387

Ningtyas, M. N., Fikriah, N. L., & Pradana, A. W. S. (2024). Muslim Gen Z Investment Decision: An Analysis Using Social Media Factors. *Journal of Islamic Economic and Business Research*, *4*(1), 108–125. https://doi.org/10.18196/jiebr.v4i1.292

OJK, & BPS. (2024). *Survei Nasional Literasi dan Inklusi Keuangan (SNLIK) 2024*. https://ojk.go.id/id/berita-dan-kegiatan/publikasi/Documents/Pages/Survei-Nasional-Literasi-dan-Inklusi-Keuangan-(SNLIK)-2024/Survei%20Nasional%20Literasi%20dan%20Inklusi%20Keuangan%20(SNLIK)%202024.pdf

Oppong, C., Salifu Atchulo, A., Akwaa-Sekyi, E. K., Grant, D. D., & Kpegba, S. A. (2023). Financial literacy, investment and personal financial management nexus: Empirical evidence on private sector employees. *Cogent Business & Management*, *10*(2), 2229106. https://doi.org/10.1080/23311975.2023.2229106

Pertiwi, T. K. (2022). Impact of Perceived Benefits, Security, and Privacy on Interest in Using E-Wallet in Millennial Generation. *International Journal of Multidisciplinary Research and Analysis*, *05*(05). https://doi.org/10.47191/ijmra/v5-i5-22

Priliasari, E. (2023). Perlindungan Data Pribadi Konsumen Dalam Transaksi E-Commerce Menurut Peraturan Perundang-Undangan Di Indonesia (Legal Protection of Consumer Personal Data in E-Commerce According To Laws dan Regulations in Indonesia). *Jurnal Rechtsvinsing*, *12*(2), 261–271.

Purba, E. L. D., Roza Thohiri, & Harefa, K. (2025). The Impact of Financial Literacy, Technological Progress, Income, and Lifestyle on Investment Interest: The Role of Gender as a Moderator Variable. *Owner : Riset Dan Jurnal Akuntansi*, *9*(2), 1183–1200. https://doi.org/10.33395/owner.v9i2.2595

Putra, T. S. A. (2024). *Peran Media Sosial Dalam Membangun Citra Positif Organisasi*. Https://Www.Djkn.Kemenkeu.Go.Id/Artikel/Baca/16999/Peran-Media-Sosial-Dalam-Membangun-Citra-Positif-Organisasi. https://www.djkn.kemenkeu.go.id/artikel/baca/16999/Peran-Media-Sosial-Dalam-Membangun-Citra-Positif-Organisasi

Ramli, R., Muda, S., Kasim, S., Zin, N. M., Ismail, N., & Padil, H. M. (2023). Examining the Relationship between Social Media and Intention to Invest in an Investment Scams among Students. In *Information Management and Business Review* (Vol. 15, Issue 4).

Raut, R. K. (2020). Past behaviour, financial literacy and investment decision-making process of individual investors. *International Journal of Emerging Markets*, *15*(6), 1243–1263. https://doi.org/10.1108/IJOEM-07-2018-0379

Remund, D. L. (2010). Financial Literacy Explicated: The Case for a Clearer Definition in an Increasingly Complex Economy. *Journal of Consumer Affairs*, *44*(2), 276–295. https://doi.org/https://doi.org/10.1111/j.1745-6606.2010.01169.x

Robkob, N., & Pankham, S. (2023). Employing Fuzzy Delphi Techniques to Validate the Components and Contents of Role of Social Media in a Technology Acceptance Model towards Perception and Investment Intention in Cryptocurrency. *Journal of Law and Sustainable Development*, *11*(12), e2032. https://doi.org/10.55908/sdgs.v11i12.2032

Samsuri, A., Ismiyanti, F., & Narsa, I. M. (2019). Effects of Risk Tolerance and Financial Literacy to Investment Intentions. *International Journal of Innovation, Creativity and Change*, *10*(9). www.ijicc.net

Sathya, N., & Prabhavathi, C. (2024). The influence of social media on investment decision-making: examining behavioral biases, risk perception, and mediation effects. *International Journal of System Assurance Engineering and Management*, *15*(3), 957–963. https://doi.org/10.1007/s13198-023-02182-x

Suresh, G. (2021). Impact of Financial Literacy and Behavioural Biases on Investment Decision-making. *FIIB Business Review*, *13*(1), 72–86. https://doi.org/10.1177/23197145211035481

Tania, J., & Utami Tjhin, V. (2025). The Influence Of Digital Literacy On Intention To Use Investment Applications In Generation Z: A Case Study On Financial Products. *Journal of Theoretical and Applied Information Technology*, *15*(7). www.jatit.org

Widyasari, A., & Aruan, D. (2022, March 23). *The Effect of Social Media Information on Intention to Invest in Indonesia Capital Market: Case of Generation Y*. https://doi.org/10.4108/eai.27-7-2021.2316776

Wu, Y., & Huang, H. (2023). Influence of perceived value on consumers' continuous purchase intention in live-streaming e-commerce—mediated by consumer trust. *Sustainability, 15*(5), 4432. https://doi.org/10.3390/su15054432