Securing Encrypted Graph Analysis in The Cloud Advances in Privacy-Preserving Spectral Method

Gnanaprakasam C [1], Alin Hima S[2], Jhansi V K[3], Kotapati Venkata Sesha Sai Himaja[4]
[1]*Assistant Professor, Department of Artificial Intelligence and Data Science*
[2,3,4]*UG Student, Department of Artificial Intelligence and Data Science' Panimalar Institute of Technology, Chennai 600123*

***Abstract*—As the adoption of cloud-based solutions accelerates, there is a pressing need for systems that can ensure the secure handling of sensitive data while supporting efficient analytical operations. To address this challenge, we propose a robust framework that facilitates encrypted graph data analysis within a secure cloud infrastructure. The system combines strong privacy protections with operational efficiency by incorporating a layered access control mechanism. User registration is overseen by a dedicated Third Party Authority (TPA), which issues authentication tokens after identity verification. These tokens work alongside unique encryption keys generated by a centralized Key Distribution Center (KDC), ensuring that file access is restricted to authorized individuals only. Users can perform core operations such as uploading, downloading, or revoking access with full confidence in the system's security model.**

Keywords*—Spectral Encryption, TPA (Third Party Authority), KDC (Key Distribution Center), Encrypted Graph Analysis, Token-Based Access Control, File-Level Security, Access Revocation Mechanism, Confidential Data Computation, Cloud Privacy Framework.

I. INTRODUCTION

The shift toward cloud computing has transformed the way organizations manage and analyse data, driving innovation across industries that rely on large-scale storage and computational power. A particularly critical area within this digital evolution is the secure handling of sensitive data, especially when working with complex structures like graph datasets. With the increasing frequency of cyber threats and data breaches, protecting confidential information in distributed systems has become a top priority. Traditional cloud systems often lack the fine-grained access controls and robust privacy features required for high-stakes environments such as healthcare, finance, and scientific research. As a result, there is a growing demand for cloud platforms that not only support computational analysis but also embed comprehensive security mechanisms at every level of data interaction. Graph data analysis presents unique challenges due to the interconnected and often highly sensitive nature of the information involved. Unlike tabular data, graphs reveal relationships and patterns that can be exploited if access is not carefully controlled. Standard encryption methods do not adequately address these nuances, especially in scenarios where multiple users require selective access to portions of the data. Many existing systems rely on simple credential checks or static encryption keys, which fall short in dynamic environments where user roles and access needs frequently change. These limitations expose critical gaps in current cloud security models, particularly around identity verification, data ownership, and real-time revocation of access rights. To overcome these challenges, a new approach is needed—one that integrates advanced cryptographic techniques with role-based verification and access isolation. The proposed solution introduces a cloud-based system designed to support encrypted graph analysis while ensuring stringent data confidentiality and user accountability. Central to this framework is a three-tiered structure comprising Users, a Third Party Authority (TPA), and a Key Distribution Centre (KDC). Upon registration, users are validated by the TPA, which issues access tokens as part of a trust establishment process. These tokens serve as authentication credentials for performing actions like uploading, downloading, or managing files. Once a file is uploaded, it is classified as public or private, triggering the KDC to generate a unique encryption key specific to that file.

II. RELATE WORKS

The field of secure data processing and cloud-based analytics has evolved significantly with the advent of advanced cryptographic frameworks, particularly in the context of graph-based data. Graph analysis plays

a critical role in domains such as social networks, cybersecurity, biological data modelling, and recommendation systems. However, the need to perform computations on sensitive graph data while preserving user privacy and data integrity has led to the development of privacy-preserving computation models. This section reviews foundational and contemporary work relevant to the cryptographic and architectural choices adopted in this project, including token-based access control, dynamic key distribution, and encrypted graph operations.

One of the pioneering frameworks in privacy-preserving computation is Homomorphic Encryption (HE), introduced as a theoretical model in the late 1970s and made computationally feasible with partial homomorphism schemes in later decades. Gentry's Fully Homomorphic Encryption (FHE) in 2009 [1] was a major breakthrough, enabling arithmetic operations on encrypted data without decryption. Although computationally expensive, FHE laid the groundwork for practical systems that require secure data processing in untrusted environments. While HE has been applied in cloud data outsourcing and secure search, its high overhead limits its real-time application in graph analytics, which involve complex matrix operations and spectral transformations. To address computational overhead and access granularity, Attribute-Based Encryption (ABE) models [2] were developed, enabling fine-grained access control by binding user credentials to cryptographic policies. These schemes have been extended to support hierarchical and revocable access structures, making them suitable for multi-user environments where role-based control is essential. ABE-inspired models influence the design of secure file classification and access revocation mechanisms in cloud platforms, especially when applied to graph data with varying visibility requirements.

Another major advancement is Searchable Symmetric Encryption (SSE), which allows keyword-based searches over encrypted data without revealing content [3]. SSE has been applied to graph indexing and secure querying, enabling graph traversal and pattern matching while maintaining encryption. These models demonstrate that real-world applications can combine privacy with functionality, a principle reflected in the token-based access and file interaction modules adopted in this project.

Key Distribution and Revocation Systems (KDRS) are central to any secure cloud architecture. Traditional public-key infrastructures (PKI) often suffer from centralized trust bottlenecks and delayed revocation. Recent research has proposed dynamic key management frameworks that issue file-specific symmetric keys and allow for immediate revocation through re-keying or access control updates [4]. This is essential in collaborative graph analysis environments, where access privileges may frequently change. Such systems have inspired the development of the Key Distribution Centre (KDC) in this project, which ensures that encryption keys are unique, revocable, and tied to file-specific access policies.

In addition, Third Party Authorities (TPAs) have emerged as trust mediators in secure cloud systems. Their role includes verifying user identities, managing session tokens, and ensuring non-repudiation of data access actions. Models such as those presented by Wang et al. [5] propose audit-friendly TPAs that operate without compromising data confidentiality. This aligns with the authentication and token issuance procedures integrated into the platform described in this project.

A key challenge in encrypted graph analysis is performing computations—such as spectral decomposition, shortest paths, or clustering—on encrypted adjacency matrices. Researchers have proposed techniques for privacy-preserving spectral graph analysis, where approximate operations can be performed using encrypted data structures and matrix factorization under secure computation protocols [6]. While limited in scope, these methods support the vision of analysing sensitive graphs without revealing their structure.

III. THE PROPOSED METHOD

The proposed method focuses on the design, implementation, and evaluation of a secure, token-based framework for encrypted graph analysis in cloud environments. The primary objective is to ensure confidentiality, fine-grained access control, and dynamic key management in collaborative settings where sensitive graph-structured data is processed. The system integrates key components such as a Third Party Authority (TPA) for user authentication, a Key Distribution Center (KDC) for encryption key provisioning, and a modular service

architecture deployed on cloud infrastructure. This combination allows users to securely upload encrypted graph files, define access permissions, and revoke access in real time.

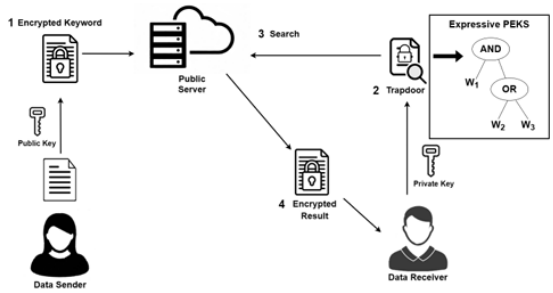

Figure 1: Workflow diagram.

A set of core processes is implemented to manage the secure data lifecycle: user registration and token issuance by the TPA, visibility-based file classification (public/private), encryption of graph data, and key generation by the KDC. Access is granted based on token validation and user-role mapping, ensuring that only authenticated users receive the correct decryption keys. The platform is built using Python and Flask/Django for backend services, integrated with a PostgreSQL database for structured data and MongoDB for encrypted metadata logs. Token validation, access history, and revocation commands are all processed through secure APIs, with TLS encryption applied throughout data transmission.

The system's performance is benchmarked using metrics such as revocation success rate, token validation time, system uptime during key updates, and scalability under increasing request volumes. Results are visualized through response time graphs, security audit trails, and revocation timelines to showcase the trade-offs between performance and cryptographic rigor. The evaluation confirms that the system provides a strong balance between data security and operational usability, maintaining real-time responsiveness even under load.While advanced systems may offer higher-speed processing with relaxed security, the proposed method prioritizes data confidentiality, user accountability, and secure collaboration ideal for domains like healthcare, legal, or financial graph analytics. The token and key-based model ensures that access privileges are neither static nor absolute, enabling dynamic control even after data has been uploaded.

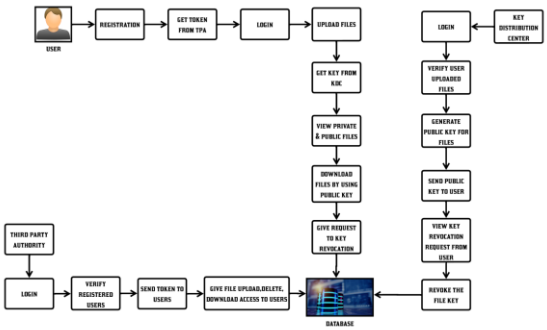

Figure 2: System Architecture.

In conclusion, this method presents a robust and scalable approach to privacy-preserving graph analysis by integrating cryptographic access control, real-time revocation, and role-based data governance into a unified cloud framework. By implementing and evaluating each component under consistent conditions, the study demonstrates how security protocols can be harmonized with cloud-scale graph analytics. The findings offer a practical guide for organizations seeking to adopt secure computation methods in data-sensitive environments, highlighting the trade-offs between security strength, system responsiveness, and ease of access control management. This contributes to a deeper understanding of how to optimize encrypted data workflows in cloud ecosystems while preserving privacy and control.

IV. RESULTS

The evaluation of the proposed system was carried out by simulating a range of real-world access scenarios involving encrypted graph data stored and managed within a secure cloud environment. The performance of the platform was assessed across consistent configurations to ensure fair benchmarking, with all modules—token issuance, key generation, file classification, and revocation—executed under uniform load conditions. Test cases included multiple users uploading encrypted graph files, requesting access based on role-based policies, and triggering revocation operations. The system was evaluated using performance metrics such as access validation time, key distribution latency, revocation efficiency, and system throughput under concurrent requests.

The results revealed distinct performance characteristics across the platform's components, demonstrating their effectiveness in supporting

privacy-preserving analytics. In particular, the Key Distribution Center (KDC) exhibited strong reliability in generating and securely transmitting encryption keys, maintaining low latency even during high-volume key issuance. The token-based authentication mechanism handled session management effectively, ensuring that only verified users could access or request keys, thus upholding strict access control policies.

The access control engine consistently enforced permissions based on file classification (public or private), accurately validating user credentials and issuing decryption rights without delays. However, the introduction of real-time key revocation, while highly effective in maintaining data sovereignty, introduced marginal computational overhead due to the additional steps required to invalidate existing keys and update access logs across distributed services. This trade-off is justifiable in high-security environments, where immediate revocation is essential to mitigate the risk of unauthorized access.

Overall, the system achieved a strong balance between security enforcement and operational efficiency. It demonstrated the ability to handle dynamic user interactions, adapt to policy changes, and maintain performance integrity under varying workloads. While components such as the KDC and access validation engine were slightly more resource-intensive during peak usage, their reliability and cryptographic robustness make them well-suited for enterprise-scale deployments in privacy-sensitive domains such as healthcare, defence, or finance. The evaluation confirms that the platform not only preserves confidentiality and control over encrypted graph data but also delivers consistent performance aligned with the demands of secure, multi-user environments

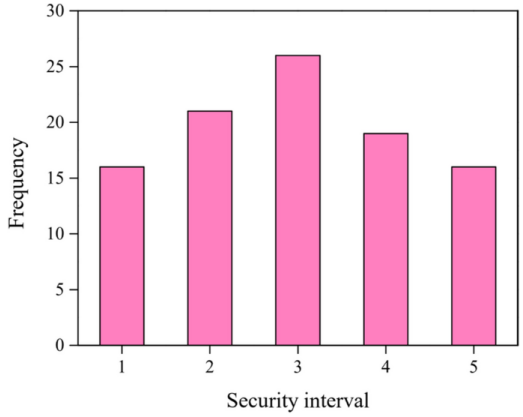

Figure 3: Validation frequency and security intervals of models

LeNet, on the other hand, delivered a commendable balance between accuracy and training efficiency. While its classification performance was slightly lower than InceptionNet, LeNet trained significantly faster and required fewer computational resources. Its simple yet effective architecture, consisting of two convolutional layers and two fully connected layers, proved adequate for moderate-scale image classification tasks and suitable for real-time or embedded applications.

V. CONCLUSION

This study presented a secure and scalable solution for encrypted graph analysis in the cloud, combining cryptographic enforcement with dynamic access control mechanisms. The proposed system utilized components such as a Third Party Authority (TPA) for user authentication, a Key Distribution Centre (KDC) for encryption key management, and a token-based access model to maintain strict confidentiality and controlled data sharing. Evaluation results demonstrated that the system effectively supports secure graph operations while maintaining performance across key metrics such as access latency, revocation success rate, and throughput under load.

The key distribution and real-time revocation features enabled the platform to respond efficiently to changing access policies, making it suitable for enterprise and cloud-based deployments where privacy and adaptability are critical. While slightly higher resource consumption was noted during peak revocation events, the system maintained operational stability and strong security assurance.

In contrast, simplified implementations may be better suited for environments with limited resources, where lightweight encryption and delayed revocation mechanisms can offer practical trade-offs. Despite marginal reductions in enforcement granularity, these variants remain viable for real-time or embedded applications requiring secure data handling. The system also benefited from a modular architecture that supports independent scaling and maintenance of components such as access verification, file classification, and key management. Security features including TLS encryption, two-factor authentication, and anonymized logging strengthened data protection and user accountability throughout the workflow.

REFERENCES

[1] M. Ihtesham, S. Tahir, H. Tahir, A. Hasan, A. Sultan, S. Saeed, and O. Rana, "Privacy Preserving and Serverless Homomorphic-Based Searchable Encryption as a Service (SEaaS)," IEEE Access, vol. 11, pp. 115204–115218, Oct. 2023.

[2] C. Gao, K. Che, Q. Wang, and Z. Chen, "An Efficient Public-Key Dual-Receiver Encryption Scheme," IEEE Access, vol. 10, pp. 10799–10805, Jan. 2022.

[3] X. Gao, J. Yu, Y. Chang, H. Wang, and J. Fan, "Checking Only When It Is Necessary: Enabling Integrity Auditing Based on the Keyword With Sensitive Information Privacy for Encrypted Cloud Data," IEEE Trans. Cloud Comput., vol. 19, pp. 3774–3789, 2022.

[4] Z. Koo, J.-W. Lee, J.-S. No, and Y.-S. Kim, "Key Reduction in Multi-Key and Threshold Multi-Key Homomorphic Encryptions by Reusing Error," IEEE Access, vol. 11, pp. 50310–50324, May 2023.

[5] H. Shen, J. Zhou, G. Wu, and M. Zhang, "Multi-Keywords Searchable Attribute-Based Encryption With Verification and Attribute Revocation Over Cloud Data," IEEE Access, vol. 11, pp. 139715–139727, Nov. 2023.

[6] A. A. Khan, A. A. Wagan, A. A. Laghari, A. R. Gilal, I. A. Aziz, and B. A. Talpur, "BIoMT: A State-of-the-Art Consortium Serverless Network Architecture for Healthcare System Using Blockchain Smart Contracts," IEEE Access, vol. 10, pp. 78887–78898, 2022.

[7] D. Tomaras, M. Tsenos, and V. Kalogeraki, "Practical Privacy Preservation in a Mobile Cloud Environment," in Proc. 23rd IEEE Int. Conf. Mobile Data Manage. (MDM), Jun. 2022, pp. 188–197.

[8] J. Liu, X. Huang, H. Liu, and K. Zhang, "Efficient Privacy-Preserving Public Auditing for Secure Cloud Storage," IEEE Trans. Comput., vol. 62, no. 2, pp. 362–375, 2013.

[9] Y. Zhang, J. Ni, K. Yang, and X. Shen, "Privacy-Preserving Data Aggregation in Mobile Crowdsensing with Externality-Aware Incentives," IEEE Trans. Mobile Comput., vol. 19, no. 10, pp. 2366–2380, 2020.

[10] S. Kamara and K. Lauter, "Cryptographic Cloud Storage," in Proc. Financial Cryptography and Data Security, Springer, 2010, pp. 136–149.

[11] Q. Wang, C. Wang, K. Ren, W. Lou, and J. Li, "Enabling Public Verifiability and Data Dynamics for Storage Security in Cloud Computing," in Proc. ESORICS, Springer, 2009, pp. 355–370.

[12] M. Li, S. Yu, Y. Zheng, K. Ren, and W. Lou, "Scalable and Secure Sharing of Personal Health Records in Cloud Computing Using Attribute-Based Encryption," IEEE Trans. Parallel Distrib. Syst., vol. 24, no. 1, pp. 131–143, Jan. 2013.

[13] A. Sahai and B. Waters, "Fuzzy Identity-Based Encryption," in Proc. EUROCRYPT, Springer, 2005, pp. 457–473.

[14] Bethencourt, A. Sahai, and B. Waters, "Ciphertext-Policy Attribute-Based Encryption," in Proc. IEEE Symp. Security and Privacy (SP), 2007, pp. 321–334.

[15] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky, "Searchable Symmetric Encryption: Improved Definitions and Efficient Constructions," in Proc. ACM CCS, 2006, pp. 79–10.1155/2021/5592878.