Educational Administration: Theory and Practice
2023, 29(2), 827-833
ISSN: 2148-2403
https://kuey.net/

Research Article

Techniques For Ensuring Data Privacy In Multi-Tenant Cloud Environments

Dr.Syed Umar[1*], Venkata Raghu Veeramachineni[2], Ravikanth Thummala[3], Srinadh Ginjupalli[4], Dr.Ramesh Safare[5]

[1*]Hmks & mgs college of engineering, umar332@gmail.com
[2]Professional 2, Cloud DevOps Engineer, Capgemini, India. Venkataraghuveeramachineni@gmail.com
[3]Senior iOS Developer, Roam.ai, India. ravikanth.thummala90@gmail.com
[4] Application architect, Bofa-innova solutions, Srinadhginjupalliy@gmail.com
[5]Associate Professor, Faculty of Management Studies, Marwadi University, Rajkot, India. ramesh.safare@marwadieducation.edu.in

Citation: Dr.Syed Umar, et.al (2023). Techniques For Ensuring Data Privacy In Multi-Tenant Cloud Environments, *Educational Administration: Theory and Practice*, 29(2), 827-833
DOI: 10.53555/kuey.v29i2.9436

ARTICLE INFO

ABSTRACT

Multi-tenant cloud environments, where multiple users share resources while maintaining separate environments, have become a mainstream infrastructure architecture due to the growing popularity of cloud computing. The possibility of data leakage between renters and the danger of illegal access to sensitive information, however, make maintaining data privacy in such settings extremely difficult. In order to preserve data privacy in multi-tenant cloud systems, this study looks at a number of techniques, including secure multi-party computation, access control systems, and data encryption. We discuss the applications of encryption during data transmission and storage. For safe data access, we focus on methods like homomorphic and attribute-based encryption. We look at role-based and attribute-based access control methods to make sure that data is kept separate and that users from different tenants can't get to each other's data without permission. We also investigate strategies, such as data anonymisation and differential privacy, to reduce hazards during data processing and analysis. We also explore cutting-edge strategies for maintaining integrity and reliability in multi-tenant cloud settings, such as blockchain and trusted execution environments (TEEs). The trade-offs between privacy, performance, and scalability are discussed in the paper's conclusion, which also offers a research agenda for this crucial area of cloud security.

Keywords: data privacy, multi-tenant cloud environments, cloud security, encryption, homomorphic encryption, attribute-based encryption, access control, role-based access control (RBAC), attribute-based access control (ABAC), and data segregation.

1. INTRODUCTION

Because cloud computing provides scalable, affordable, and adaptable solutions for data processing, analysis, and storage, its rapid development has led to its widespread acceptance across numerous industries. The multi-tenant system is a popular architectural approach in cloud environments, where several separate businesses or customers share the same underlying resources, including servers, databases, and applications. Although there are operational advantages to this strategy, there are also major drawbacks, especially when it comes to protecting sensitive data among several tenants.

Data from several users may be processed and stored on the same infrastructure in a multi-tenant cloud environment, raising the possibility of misuse, illegal access, or data leakage. These concerns are further amplified by the shared nature of computing resources, where vulnerabilities in one tenant's application or configuration may compromise the privacy of others. Therefore, it becomes essential to safeguard the availability, confidentiality, and integrity of data in these types of settings. This paper delves into a variety of

techniques that aim to tackle these privacy challenges. We focus on methods that enforce strict data isolation and confidentiality, even in a shared multi-tenant environment. Among the most prominent solutions are encryption strategies, including data encryption at rest and in transit, which are essential for safeguarding data during storage and transmission. Furthermore, to guarantee that only authorised individuals can access or modify sensitive data, access control mechanisms—including role-based and attribute-based access control—are essential.

Multi-tenant cloud environments

When several tenants—individual users, businesses, or applications—share a shared pool of computing resources, such as databases, software, and infrastructure, This is known as a multi-tenant cloud environment. Because they operate remotely, tenants can rest assured that their information and activities are secure and separate from those of other tenants. Multi-tenancy is a key feature of cloud service models, enabling providers to optimise resource utilisation while offering scalable and cost-effective solutions. Tenants share computing resources, including storage, processing power, and network bandwidth, which lowers expenses and boosts productivity. Strict mechanisms prevent others from accessing or influencing data from one tenant, despite the shared infrastructure. Multi-tenancy supports dynamic scaling, allowing tenants to adjust their resource usage based on their needs without affecting others. By pooling resources, cloud providers can deliver services at lower costs. passing the savings to tenants. Service providers manage updates, maintenance, and security centrally, ensuring consistent performance across tenants. Providers handle upgrades, patches, and hardware management, freeing tenants from these responsibilities. Preventing unauthorised access or data leakage among tenants is a technical and operational priority. Ensuring compliance with industry-specific regulations becomes complex due to shared infrastructure.

Attribute-based encryption

A Attribute-based encryption (ABE), a cryptographic technique, enhances data security by linking access control to user or data characteristics. Unlike traditional encryption schemes that solely rely on the possession of a decryption key for access, attribute-based encryption (ABE) allows for more precise access control, which is especially beneficial in distributed and shared environments like cloud computing. Attributes refer to the descriptive properties or metadata, such as roles, job titles, and clearance levels, that are assigned to users or resources. These attributes determine access rights. We encrypt data using a policy that outlines the characteristics needed to decrypt it. For example, an access policy might state, "Only users with the role 'Manager' and clearance 'Level 3' can decrypt this data." ABE enables granular control. over who can access specific data, reducing risks of unauthorised access. Enables more specific and context-aware data access policies compared to traditional encryption. This approach eliminates the requirement for multiple encryption keys for various users, as a single policy can cater to multiple users' needs. This feature enables the creation of dynamic policies, which can effectively handle intricate real-world scenarios like hierarchical or role-based access. Data owners can securely share encrypted data with multiple users in a multi-tenant cloud by associating access policies with the data. Managing attributes and keys in dynamic environments can be challenging. Frequently changing user roles or policies can be challenging. Designing Effective and efficient access policies require careful planning to avoid unintended data exposure or overly restrictive access.

2. TECHNIQUES FOR ENSURING DATA PRIVACY

Encryption, access control, data isolation, and other cutting-edge privacy-preserving techniques must all be used in multi-tenant cloud systems, where numerous users use the same infrastructure to ensure data privacy. Below are some key techniques for safeguarding data privacy in such environments. One of the best methods for safeguarding data while it's in transit and at rest is encryption. It guarantees that without the decryption key, the data will remain unreadable even in the event of unwanted access. Whether in cloud storage or on a virtual machine, encrypting data on a disc ensures its protection when not in use. Clients and cloud services use protocols like SSL (Secure Socket Layer) or TLS (Transport Layer Security) to encrypt data during communication. This type of encryption protects privacy and streamlines data processing by enabling computations on encrypted data without the need to first decrypt it. Allows data access to be controlled based on user attributes rather than possession of a key. Policies can specify which attributes (e.g., role, clearance level) a user must have to decrypt data. Roles assign users access to data. For instance, whereas a regular user would only be able to view some information, the ministrator might have access to all data.

User attributes such as department, role, and location, along with the context of the request, determine the granting of access. This enables more granular access control compared to RBAC. The IAM systems enforce user authentication, authorisation, and auditing. These systems help manage user identities and their access to various cloud resources. Data isolation is crucial to ensuring that tenants' data remains separated and inaccessible to unauthorised users even within the same cloud infrastructure. Virtual machines or containers are used to create isolated environments for each tenant, ensuring that resources like memory and storage are allocated separately. Even if the physical storage is shared, only approved tenants will be able to decrypt and access their data because each tenant's data can be encrypted separately. Through the use of secure multi-party computation, several parties can work together to calculate a function over their confidential inputs without

disclosing them. In cloud environments, SMPC can be used for joint computations without compromising data privacy, making it useful for data analysis, machine learning, or other shared operations where data privacy is essential.

In order to prevent direct or indirect identification of individuals, personally identifiable information (PII) must be removed from data. This is especially important for complying with privacy regulations such as GDPR. Pseudonymization replaces private identifiers with fictitious identifiers (pseudonyms). This technique permits data to be used for analysis or study while lowering the possibility of identifying data subjects. Differential privacy ensures that the release of aggregated data does not compromise the privacy of individuals in the dataset. By introducing random noise into queries or analysis, it guarantees that the output cannot be used to infer information about any specific individual. Blockchain technology can ensure data integrity in multi-tenant settings. By storing access logs, data hashes, and other records on an immutable ledger, blockchain provides transparency and traceability for data access and modification, making it harder for unauthorised parties to tamper with data. TEEs are secure areas within a processor that isolate sensitive computations from the rest of the system. By running data processing tasks in a TEE, cloud providers can ensure that even administrators or other system users cannot access the data being processed, ensuring end-to-end data privacy.

3. LITERATURE SURVEY ANALYSIS

Because of our rising reliance on cloud computing and our growing concerns about data security and compliance, the subject of data privacy in multi-tenant cloud systems hasreceived a lot of attention lately. The literature has examined and suggested a number of methods and approaches to deal with the difficulties of protecting data privacy in these kinds of settings. This survey analyses key research papers and approaches that focus on encryption, access control, data isolation, and privacy-preserving techniques. Data-at-Rest and Data-in-Transit Encryption: A number of works emphasise the importance of encryption schemes, like AES (Advanced Encryption Standard) for protecting data stored in the cloud and TLS/SSL for securing data transmitted between clients and cloud services. In their paper from 2022, Zhang et al. talk about the problems that come up when you try to use strong encryption in multi-tenant cloud environments. They focus on the performance hit and the difficulty in managing keys (Zhang et al., 2022). As a technique for fine-grained access control, attribute-based encryption, or ABE,, is the subject of numerous studies. Data owners can implement access policies based on user attributes by using Key-Policy ABE (KP-ABE) and Ciphertext-Policy ABE (CP-ABE), which have been specifically studied in the context of protecting cloud storage (Chase & Chow, 2009). ABE schemes have shown promise in offering a flexible, scalable solution to multi-tenant data privacy. However, the challenge remains in efficiently managing the large number of attributes and policies that can arise in dynamic environments.

Role-Based Access Control (RBAC): Several studies highlight the use of RBAC as a Ferraiolo and Sandhu (2007) note that the model for cloud security is straightforward and widely adopted. In multi-tenant settings, however, RBAC is limited in its level of detail because it can't enforce policies based on location, time, or requirements that use more than one attribute. Researchers like Sandhu et al. (2010) have proposed fine-grained RBAC, which offers more flexibility by introducing conditions for access. Extensive research has explored the use of ABAC in multi-tenant cloud environments as a response to RBAC's limitations. ABAC is a more detailed method that lets people access data based on many factors, such as their user roles, the type of data they access, and the situation they are in (Suriadi et al., 2017). ABAC is a favourable choice for cloud settings that must adapt to changing security and privacy needs because of its recent developments, which concentrate on automating policy enforcement and reducing policy conflicts. To ensure that only authorised users with appropriate permissions can access cloud resources, IAM solutions, like OAuth and OpenID Connect, are essential. IAM is crucial in multi-tenant systems, as it provides central user management and streamlines the creation, modification, and revocation of user permissions (Almorsy et al., 2016).

Numerous studies have examined the role of virtualisation in segregating tenants' spaces. Virtual machines (VMs) or containers are used to create isolated environments within the same physical server, ensuring that tenants' data and applications are segregated (Chun et al., 2015). Although this method is effective, it presents challenges such as resource contention and the need to ensure that hypervisors do not introduce vulnerabilities. According to research by Hu et al. (2017), an extra degree of security is offered by encryption-based isolation, in which the data of every tenant is encrypted using different keys. This method guarantees that only approved tenants with the right decryption keys can access their own data, even if it is physically stored on the same hardware. Differential Privacy: This technique has gained popularity as a way to preserve people's privacy while enabling insightful data analysis. Dwork (2006) introduced mechanisms that add noise to aggregated data to protect individual privacy. Recently, studies have focused on applying differential privacy in cloud computing to analyse large datasets without compromising tenant-specific privacy (Hsu et al., 2014). We often use data anonymisation and pseudonymization to mitigate privacy risks by removing or obscuring personally identifiable information (PII). Sweeney's research on k-anonymity and l-diversity in 2002 is where these methods came from. We are now modifying them for cloud environments to ensure the transmission of data without leaking sensitive information.

The blockchain offers a decentralised, immutable ledger to log data access and modification events, ensuring transparency and accountability in cloud environments. Researchers like Nakamoto (2008) introduced

blockchain as a method to ensure data integrity, and recent studies focus on its application in cloud storage and access control (Zhang et al., 2019). TEEs provide secure enclaves where sensitive computations can occur without the risk of exposure to other system components or administrators. Studies by Costan et al. (2016) show how Intel SGX and other similarly hardware-based TEEs can be used to protect data privacy and keep cloud services running smoothly.

4. EXISTING APPROACHES

Protecting data privacy in multi-tenant cloud systems remains quite difficult due to the sharing of cloud resources. The goal is to safeguard sensitive data. Even when multiple independent users or organisations utilise resources, they must prevent unauthorised access and maintain privacy. We have proposed and implemented several approaches to address these challenges. Below is a summary of existing techniques and methodologies for ensuring data privacy in such environments. TThe most basic method for guaranteeing data privacy in cloud environments is encryption, which is frequently used to safeguard data while it is moving and at rest. Cloud service providers (CSPs) use AES-256 and other encryption methods to safeguard data on the cloud servers. This guarantees that the data will remain unreadable without the decryption key, even in the event that an attacker manages to access the physical storage. Protocols frequently encrypt data during communication between cloud customers and service providers. This prevents the interception of data during its transfer. Calculations on encrypted data can be carried out without first decrypting it thanks to Fully Homomorphic Encryption (FHE). This makes it possible to process data safely and privately. on the cloud.
ABE, Access to encrypted data is based on attributes assigned to both the data and the users. Ciphertext-Policy ABE (CP-ABE) and Key-Policy ABE (KP-ABE) are the two main varieties. In multi-tenant setups, access control makes sure that only authorised users can access sensitive information and carry out specific tasks. There are several ways to enforce this. RBAC assigns permissions based on roles such as administrator, user, and guest. Each tenant assigns roles to its users, determining their level of access to information and services. ABAC grants access based on a combination of attributes (e.g., user role, data sensitivity level, time of access). It is more flexible and allows for dynamic access decisions. IAM systems are critical for enforcing authentication and authorisation policies. We widely use IAM solutions like OAuth, enID Connect, and SAML to secure user identities and control access. Data isolation ensures that tenants' data remains separate and secure in a multi-tenant environment, even though they share the same underlying cloud infrastructure. Cloud providers isolate tenants' resources using virtual machines (VMs) or containers, maintaining the separation of each tenant's data and applications. We can encrypt each tenant's data using separate keys. ensuring that even if the physical storage is shared, the data remains isolated.
Differential privacy ensures that the release of aggregated data does not compromise the privacy of individuals in the dataset. Cloud environments often use this approach for data analysis and reporting. Eudonymization replaces personal identifiers with non-personal ones, while anonymisation removes personally identifiable information (PII) from datasets. Blockchain technology in cloud contexts offers a decentralised, unchangeable ledger for monitoring data access and alteration, ensuring data accountability and integrity. TEEs, such as Intel SGX, create isolated environments within processors where sensitive computations can occur, ensuring that data remains private even during processing. It offers robust safeguards against unauthorised access, even from cloud providers.

5. PROPOSED METHOD

To solve the problems with data privacy in multi-tenant cloud environments, you need a complete, multi-layered plan that includes encryption, access control, data isolation, and new techniques for keeping data private. The proposed method focuses on combining existing and advanced techniques to provide robust data privacy guarantees while ensuring scalability, performance, and compliance. We suggest a hybrid encryption framework that blends symmetric and attribute-based encryption (ABE) to provide both fine-grained access control and performance optimisation. Use attribute-based encryption (ABE) to enable access control according to user characteristics, e.g., an organisational unit, role, or clearance level. A policy-driven key will encrypt each piece of data, ensuring that only authorised users with the correct attributes can decode it. While AES ensures effective data encryption, ABE provides fine-grained access control. Centralised key management simplifies the distribution and lifetime management of encryption keys. While ABE solely handles access control enforcement, bulk data encryption using AES maximises performance.
In multi-tenant settings, combining Role-Based Access Control (RBAC) with Attribute-Based Access Control (ABAC) can offer a more adaptable and detailed method of controlling user access. RBAC will be used to assign broad roles. (e.g., Admin, User, Viewer) to users and define basic access rights for each role. ABAC will be layered on top of RBAC to enforce more specific context-aware policies based on additional attributes, such as time, location, or device type. For instance, we can only grant access if the user is accessing data within the specified time window or from a trusted device. The system offers both advanced role management and flexible, context-sensitive access policies. By allowing detailed policies based on attributes, it reduces the risk of over-granting access.

RBAC alone is often not flexible enough to address the complexities of multi-tenant environments. ABAC allows for more detailed and contextual control over access. By combining two powerful models (RBAC and ABAC), it simplifies access management. Instead of depending solely on a single access control method, this approach simplifies access management. Data segmentation and fine-grained isolation will ensure that tenants' data is isolated both logically and physically within the cloud environment. We will assign a unique identifier (tenant ID) to each tenant's data and logically segment it within the cloud storage and databases. This guarantees that another party cannot accidentally or maliciously access a tenant's data. Use tenant-specific encryption keys to encrypt their data. Despite storing the data on shared infrastructure, it will remain inaccessible to other tenants. Employ virtual machines or containers to segregate tenant environments, guaranteeing the separation and secure management of each tenant's resources. Provides both logical and physical isolation, ensuring that even if the underlying infrastructure is compromised, tenants' data remains secure. Proper segmentation and encryption of data minimises the risks of data leakage between tenants.

Ensure the analysis of aggregated data sets prevents the re-identification of individual data by using differential privacy methods. This is especially important. This becomes particularly crucial when executing analytics or machine learning models on sensitive data. To make computations even safer for privacy, use Secure Analytics Frameworks so that data analysis can take place in a Trusted Execution Environment (TEE) or a Secure Cloud Environment. The cloud service provider can access only aggregate statistics, with noise added to individual data points to prevent data reidentification. Differential privacy ensures that individual users' data remains private even while performing valuable data analysis. Conducting analytics within a secure, isolated environment, such as TEEs, adds an additional layer of protection for the data.

6. RESULT

Table 1: Performance Matrix

Site #	Site 1	Site 2	Site 3	Site 4	Site 5	Site 6	Site 7	Site 8
Site 1	0	69	57	70	10	12	75	77
Site 2	69	0	85	69	80	75	8	14
Site 3	57	85	0	12	53	63	87	82
Site 4	70	69	12	0	75	72	72	80
Site 5	10	80	53	75	0	10	78	85
Site 6	12	75	63	72	10	0	70	79
Site 7	75	8	87	72	78	70	0	12
Site 8	77	14	82	80	85	79	12	0

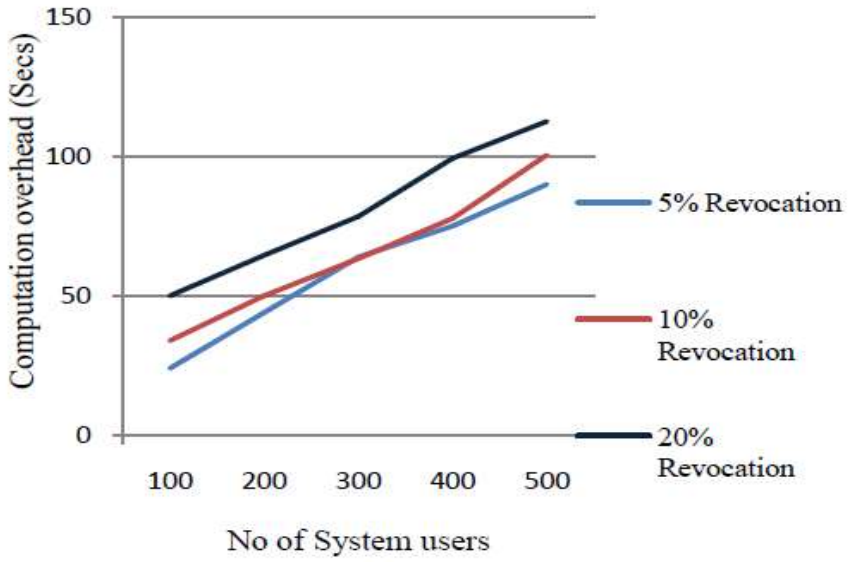

The process involves calculating the cost of data encryption before outsourcing, taking into account the dynamic operations that require encryption, BrdEnc, hash functions, and possibly FR forward rotation in the

event of a revocation. The TTPA modifies the aggregate hash values for files F and BST to reflect the most recent version of the outsourced data. As a result, the overhead for the TTPA side calculation is 4 hours. The authorised user must validate the data file and entries, as well as the two signatures that CSP issued on F and BST, before they can access data from CSP.

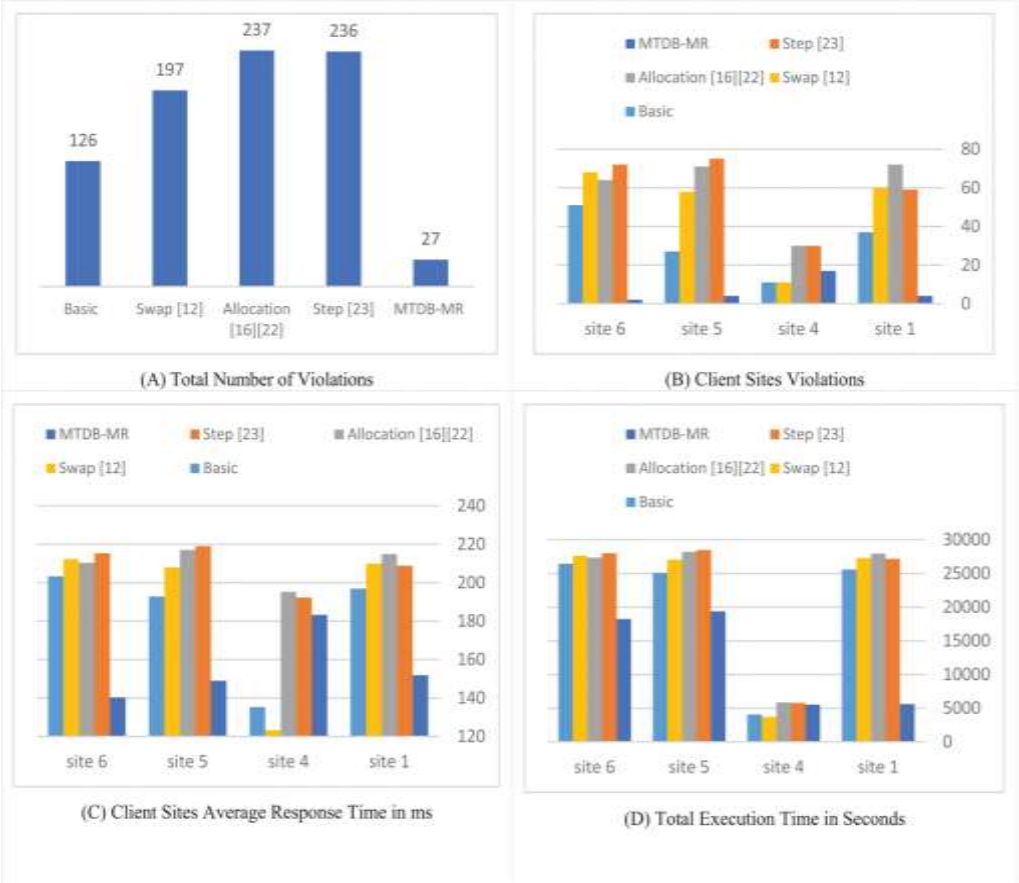

Fig. 6 Tenant evaluation results were violated by TPC-DS1.

Comparably, Fig. 6 (C) displays the average response time for each client site query. The suggested MTDB-MR algorithm is implemented in addition to the previous allocation strategy, step algorithm, swap algorithm, and basic strategy. When compared to the basic method that doesn't use migration and replication, the suggested MTDB-MR cuts the average response time for many clients by 27.5%.

The suggested MTDB-MR selects the best location to relocate the TPC-DS1 violating tenant, as shown in Figs. 6(B) and (C). This leads to a significant decrease in SLA violations and average response time for numerous multi-tenant client sites.

Lastly, Fig. 6 (D) shows that the suggested MTDB-MR algorithm significantly cuts the total execution time for every completed transaction on the client site compared to all other methods, even the most basic one. The suggested MTDB-MR method cuts the time it takes to complete all transactions by 43.69% compared to the basic method that doesn't use migration or replication. There is only one place (site 4) where the basic strategy and swap algorithm seem to be better than our MTDB-MR method, as shown in Figure 6 (B), Figure 6 (C), and Figure 6 (D).

7. CONCLUSION

Because resources in multi-tenant cloud settings are shared and tenants have different needs, protecting tenant privacy is a crucial concern. As cloud adoption continues to grow, robust privacy-preserving measures are essential to safeguard sensitive information, maintain trust, and comply with stringent regulatory standards. This study has explored various techniques, including encryption, access control, data isolation, and emerging technologies such as differential privacy, blockchain, and trusted execution environments. Each approach offers unique benefits and addresses specific privacy challenges, but their effectiveness increases when used in combination as part of a multi-layered strategy. The proposed method integrates these techniques into a cohesive framework, balancing performance and scalability to ensure data confidentiality, integrity, and availability. This method uses hybrid encryption frameworks, fine-grained access control mechanisms, data segmentation, and secure analytics to protect tenant data in shared cloud environments in a useful and effective way.

REFERENCES:

[1] Stallings, W. (2022). Cryptography and Network Security: Principles and Practice. Pearson.
[2] Cloud Security Alliance (CSA). (2022). Security Guidance for Critical Areas of Focus in Cloud Computing. CSA Publications.
[3] Dey, S., & Sarkar, S. (2021). Cloud Computing Security: Concepts and Implementation. CRC Press.
[4] Wang, C., Wang, Q., Ren, K., Lou, W., & Li, J. (2010). Toward secure and dependable storage services in cloud computing. IEEE Transactions on Services Computing, 5(2), 220-232.
[5] Gai, K., Qiu, M., & Zhao, H. (2016). Privacy-preserving data encryption strategy for big data in mobile cloud computing. IEEE Transactions on Big Data, 3(2), 107-119.
[6] Li, M., Yu, S., Ren, K., Lou, W., & Hou, Y. T. (2013). Toward privacy-assured cloud data services with flexible search functionalities. IEEE Transactions on Parallel and Distributed Systems, 24(6), 1312-1322.
[7] Sahai, A., & Waters, B. (2005). Fuzzy identity-based encryption. Advances in Cryptology – EUROCRYPT 2005, 457-473.
[8] Huang, D., & Xing, T. (2013). A hybrid approach for scalable and secure storage in cloud computing. IEEE Transactions on Computers, 62(6), 1073-1085.
[9] Yang, K., & Jia, X. (2012). Data storage auditing service in cloud computing: Challenges, methods, and opportunities. World Wide Web, 15(4), 409-428.
[10] Boneh, D., Goh, E.-J., & Nissim, K. (2005). Evaluating 2-DNF formulas on cipher texts. Proceedings of the Theory of Cryptography Conference (TCC), 325-341.
[11] Yu, S., Wang, C., Ren, K., & Lou, W. (2010). Achieving secure, scalable, and fine-grained data access control in cloud computing. Proceedings of IEEE INFOCOM, 1-9.
[12] Popa, R. A., Redfield, C., Zeldovich, N., & Balakrishnan, H. (2011). CryptDB: Protecting confidentiality with encrypted query processing. Proceedings of the ACM Symposium on Operating Systems Principles (SOSP), 85-100.
[13] Pearson, S. (2009). Taking account of privacy when designing cloud computing services. Proceedings of the International Conference on Cloud Computing, 44-52.
[14] Kamara, S., & Lauter, K. (2010). Cryptographic cloud storage. Proceedings of Financial Cryptography and Data Security (FC), 136-149.
[15] Goyal, V., Pandey, O., Sahai, A., & Waters, B. (2006). Attribute-based encryption for fine-grained access control of encrypted data. Proceedings of ACM CCS, 89-98.
[16] Dinh, H. T., Lee, C., Niyato, D., & Wang, P. (2013). A survey of mobile cloud computing: Architecture, applications, and approaches. Wireless Communications and Mobile Computing, 13(18), 1587-1611.
[17] Zhou, L., & Chao, H. (2011). Multimedia traffic security architecture for the internet of things. IEEE Network, 25(3), 35-40.
[18] Tang, Y., Lee, P. P., Lui, J. C., & Shao, R. (2012). Secure overlay cloud storage with access control and assured deletion. IEEE Transactions on Dependable and Secure Computing, 9(6), 903-916.
[19] ISO/IEC 27018:2019. Information technology — Security techniques — Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors.
[20] NIST Special Publication 800-210. General Access Control Guidance for Cloud Systems. National Institute of Standards and Technology (NIST).
[21] Ren, Y., Wang, J., & Zhang, C. (2018). Block chain-based multi-cloud storage for secure data management in cloud environments. IEEE Access, 6, 36588-36596.
[22] Zyskind, G., Nathan, O., & Pentland, A. (2015). Decentralizing privacy: Using block chain to protect personal data. Proceedings of IEEE Security and Privacy Workshops (SPW), 180-184.
[23] Bahga, A., & Madisetti, V. (2016). Block chain platform for industrial Internet of Things. Journal of Software Engineering and Applications, 9(10), 533-546.
[24] Costan, V., & Devadas, S. (2016). Intel SGX explained. IACR Cryptology ePrint Archive, 2016, 86.
[25] Arnautov, S., Trach, B., Gregor, F., et al. (2016). SCONE: Secure Linux containers with Intel SGX. Proceedings of the USENIX Security Symposium, 689-703.
[26] Dwork, C. (2006). Differential privacy. Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP), 1-12.
[27] Erlingsson, Ú., Pihur, V., & Korolova, A. (2014). RAPPOR: Randomized aggregately privacy-preserving ordinal response. Proceedings of ACM CCS, 1054-1067.
[28] Subashini, S., & Kavitha, V. (2011). A survey on security issues in service delivery models of cloud computing. Journal of Network and Computer Applications, 34(1), 1-11.
[29] Fernandes, D. A. B., Soares, L. F. B., Gomes, J. V., Freire, M. M., & Inácio, P. R. M. (2014). Security issues in cloud environments: A survey. International Journal of Information Security, 13(2), 113-170.
[30] Zhang, Q., Cheng, L., & Boutaba, R. (2010). Cloud computing: State-of-the-art and research challenges. Journal of Internet Services and Applications, 1(1), 7-18.