

Cloud Vault

1st Anamika Gupta

dept.Computer Science and Engineering
Galgotias University
Greater Noida, India
ganamika805@gmail.com

2nd Gauransh Taneja

dept.Computer Science and Engineering
Galgotias University
Greater Noida, India
gauranshtaneja175@gmail.com

3rd Arvind Panwar

dept.Computer Science and Engineering
Galgotias University
Greater Noida, India
arvind.panwar@galgotiasuniversity.edu.in

4th Rahul Sharma

dept.Computer Science and Engineering
Raj Kumar Goel Institute of technology
Ghaziabad, India
sharma.rahul9944@gmail.com

Abstract—This research introduces Cloud Vault, an intelligent cloud-based storage system that integrates conventional data management with social interaction features. Designed for scalability, security, and adaptability, Cloud Vault employs distributed architectures, virtualization, and multi-tier storage strategies to optimize performance and cost-efficiency. It leverages attribute-based encryption (ABE), blockchain for auditable access control, and AI-driven classification for adaptive data placement. The system was implemented using Django 3.2 on a virtualized environment and benchmarked under varying loads to evaluate latency and throughput. Results demonstrated robust performance under concurrent access, confirming feasibility for real-world deployment. The study also reviews pivotal technologies such as Ceph, DepSky, and Merkle Hash Trees, highlighting advances in erasure coding, privacy mechanisms, and green cloud practices. By combining secure storage with interactive user engagement, Cloud Vault proposes a novel model for unified digital collaboration. Future work will focus on enhanced scalability, data sovereignty solutions, and expanded AI integration for predictive resource allocation.

Index Terms—Django, Cloud Computing, Latency, System throughput

I. INTRODUCTION

Cloud storage has become a pillar for contemporary computing infrastructure, providing scalable, dependable, and economical data storage and access in geographically dispersed environments. With the rapid increase in data produced by users, companies, and IoT devices, the need for high-performance, fault-tolerant, and safe storage systems has dramatically accelerated. The roots of distributed cloud storage started with technologies such as the Google File System (GFS), which added a scale-out architecture for managing big data sets over commodity hardware. Ghemawat et al. [1] designed a master-slave structure optimized for large, sequential file access, fault tolerance, and data replication, as a model for subsequent systems. Based on this, Weil et al. [2] developed the Ceph distributed object-based storage system, introducing an entirely decentralized metadata management strategy with a dynamic subtree partitioning scheme. Ceph's scalability, high performance, and strong consistency guarantees make it a core architecture in most cloud deployments.

To address vendor lock-in and availability problems, Bessani et al. presented DepSky, a "cloud-of-clouds" system that aggregates several providers for fault tolerance, confidentiality, and integrity improvement [3]. Their solution employs secret sharing and Byzantine quorum systems to reduce the threats of data loss or corruption due to individual providers. Security-wise, Verma et al. [4] have offered a de-tailed overview of privacy-preserving schemes in cloud storage and compared cryptographic methods including homomorphic encryption, searchable encryption, and oblivious RAM. Their study emphasizes the fact that while encryption is important, it needs to go hand-in-hand with effective access control and monitoring of data usage. Yang et al. [5] presented a blockchain and attribute-based encryption (ABE)-based access control scheme for enforcing fine-grained authorization along with auditability. The scheme provides decentralized enforcement and mitigates the risk of a single point of failure, which is one of the primary concerns in legacy access control models.



Fig. 1. Glimpse of our software

As the workload on cloud storage for analytics intensifies, performance optimization is a must. Durner et al. [6] illustrated how cloud object storage can be leveraged for analytical performance by reducing I/O bottlenecks and using techniques like data skipping and columnar storage layouts. Their paper introduces storage as not merely a passive depot but as an active member in big data pipelines. These classic and recent works highlight the cross-cutting challenges and breakthroughs of cloud

Figure 1: Sizing Cloud Shift, Worldwide, 2019 – 2025

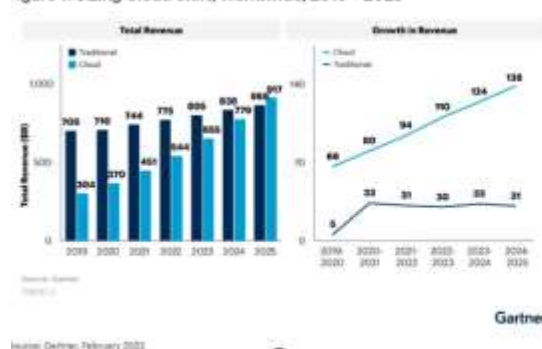


Fig. 2. Gartner Survey

storage systems. From design to performance, dependability, and privacy, they provide insightful feedback that guides the direction of our work. Our research builds on those contributions to suggest a hybrid, privacy-conscious, and high-performance cloud storage system based on decentralized techniques and adaptive encryption that overcome existing shortcomings. According to Gartner survey we can analyze the sizing cloud shift world wide from 2019 to 2025.[]

II. LITERATURE SURVEY

A. Historical Context of Cloud Storage

Classic on-premises implementations with hardware boundaries and scalability issues were the starting point of the revolution. The release initiated a paradigm shift towards flexibility and capacity to handle vast amounts in distributed environments. Establishing models (IaaS PaaS SaaS) and developing frameworks for service provision were the primary objectives of initial research. Technical advancements prompted studies of specific aspects of security processes, performance optimization techniques, and cost-benefit. Evaluation. A disruptive technology, cloud-based storage systems are revolutionizing how individuals, companies, and organizations store, manage, and retrieve data. This literature review emphasizes the work of many researchers and practitioners while analyzing the development, core concepts, benefits, challenges, and innovations of cloud-based storage systems.

Hybrid systems, which combine private and public clouds to provide more flexibility while keeping control of sensitive data, are in vogue today. Our comprehension of user behavior in the context of privacy issues under such situations remains short, however. The impact of regulatory compliance on adoption levels by organizations aiming to capitalize on the leverage benefits provided by contemporary infrastructures being employed in present times, and user trust in service providers in terms of privacy threats involved with public vs. private solutions, remain important topics that need further scrutiny.

Ref.	Methodology	Key Contribution	Advantages	Limitation
[1] M. Ding et al. (2014)	Developed HYSTO, a hybrid VM storage system combining SSDs and distributed storage to enhance privacy in cloud data centers.	Introduced a privacy-preserving shared storage system for virtual machines.	Improved I/O performance and scalability; enhanced data privacy.	Complexity in managing hybrid storage components.
[2] D. C. Nguyen et al. (2021)	Proposed Flokus, integrating federated learning with blockchain in edge computing environments.	Enabled decentralized, secure, and privacy-enhancing systems in mobile edge computing.	Enhanced data privacy and security; reduced communication overhead.	Potential computational overhead on edge devices.
[3] A. Aze et al. (2017)	Surveyed various homomorphic encryption schemes, analyzing their theoretical foundations and implementations.	Provided a comprehensive overview of HE schemes, facilitating secure computations on encrypted data.	Enabled operations on encrypted data without decryption, preserving data confidentiality.	High computational complexity; limited practical implementations.
[4] D. C. Nguyen et al. (2019)	Reviewed the integration of blockchain with Cloud of Things (CoT), discussing architectures and applications.	Highlighted the benefits and challenges of BCuT integration in various domains.	Improved data security and transparency; decentralized data management.	Scalability issues; integration complexity.
[5] N. Jayapandian & A. M. T. M. Zakiah Zakaria (2017)	Combined probabilistic and homomorphic encryption techniques for secure cloud data storage and sharing.	Proposed an efficient algorithm enhancing data security in cloud environments.	Increased throughput, reduced security attacks, improved QoS.	Potential increase in computational overhead.
[6] F. Karzika Razi et al. (2022)	Implemented hybrid encryption combining symmetric and asymmetric algorithms to enhance cloud security.	Enhanced data confidentiality and integrity in cloud storage systems.	Strengthened security measures, balanced performance.	Key management complexity.
[7] F. Kassar et al. (2021)	Developed a cloud security framework utilizing hybrid cryptographic algorithms, including DES and RSA.	Provided a multilayer encryption approach for securing cloud data.	Enhanced data protection, layered security architecture.	Increased encryption and decryption time.
[8] A. Mankamda & N. Nadeem (2021)	Employed cryptographic access control mechanisms for fine-grained security in cloud environments.	Enabled precise access control over cloud resources.	Improved data confidentiality; flexible access management.	Complexity in policy management.
[9] A. M. Saib et al. (2021)	Proposed a secure model for data protection in cloud computing, focusing on authentication and encryption.	Addressed security issues in cloud data storage and access.	Enhanced data security, robust authentication mechanisms.	Implementation complexity in diverse cloud environments.

III. ARCHITECTURE

Cloud storage architecture generally consists of a multi-layered, distributed system that separates the front-end access interface from the back-end data storage. On the front end, there is a presentation layer presenting APIs (usually RESTful web services) and management consoles for clients to access the storage system. Behind it is a "storage logic" middleware layer providing fundamental functions (e.g., replication, erasure coding, data placement). The back end is a distributed cluster of storage nodes (virtual instances or servers) world-wide that physically store user data. Ceph's architecture, for instance, employs monitor daemons to handle cluster state, object-storage daemons (OSDs) as storage nodes, metadata servers for namespace, and RESTful gateways that provide S3/Swift APIs. In total, cloud storage systems employ scalable clusters of commodity hardware with software controllers to structure user data across several physical locations, offering users a virtualized pool of storage.

Enabling Technologies Cloud storage is founded on a number of fundamental technologies that contribute to its scalability and elasticity:

- **Virtualization:** Compute and storage virtualization decouple hardware, allowing for flexible resource allocation. Virtualization enables any software to be run on shared servers, with virtual disks dynamically resized or moved.

This "decoupling" of hardware and software is the foundation for cloud elasticity.

- **Object Storage Models:** Most cloud storage solutions leverage object-based storage, where every file (object) stores its own metadata and is addressed through a unique ID. Object storage systems (e.g. Amazon S3, Swift) abandon file hierarchies, enabling flat namespaces and simple replication. Objects and metadata are stored in big distributed stores, and accessed through RESTful APIs.
- **Distributed File Systems:** Distributed block/file systems (e.g. CephFS, Hadoop HDFS) are used by some clouds. These have a NameNode/metadata-server for the namespace and numerous data nodes to store blocks. The metadata server manages block locations, whereas distributed data nodes are used for parallel reads/writes. These file systems offer POSIX-like semantics over the cloud cluster.




Fig. 3. Layout of Cloud Storage

RESTful and HTTP APIs: HTTP/REST web-service interfaces are the cloud storage de facto standard. Clients do not use native file I/O but instead PUT/GET/DELETE objects over HTTP (usually in S3 or OpenStack Swift format). These APIs enable applications to be easily integrated with them and permit application of cloud storage through firewalls. An observation by research indicates that the front-end of the cloud "exports an API to access storage" and that conventional block protocols are substituted with web interfaces.

- **Data Protection Mechanisms:** Erasure coding, replication and self-healing algorithms are some of the technologies that provide reliability. For example, Ceph's RADOS layer replicates or encodes data within OSDs such that failures of nodes are allowed without losing data. Techniques like these are implemented within the storage logic middleware.

The workflow of the system starts with the user uploading a file via the user interface. The subsequent series of actions is carried out:

1) Data Profiling and Classification:

Metadata including file type, size, and prior user activity is processed. The AI model makes a prediction of future access likelihood.

1) Encryption and Policy Tagging:

The file is encrypted in advance using ABE and assigned access control policies and sensitivity labels.

1) Tier Assignment:

According to fore-casted access frequency, data is stored in the respective tier: Hot, Warm, or Cold. Sensitive files are locked to the private cloud independent of access frequency.

1) Logging and Verification:

Every access or modification request is logged onto the blockchain, and data integrity is occasionally verified using Merkle Hash Trees.

1) Adaptive Reallocation:

As patterns of access change, files can be moved between tiers or across cloud layers by the orchestrator to minimize cost and maximize performance.

IV. ADVANTAGES

Cloud storage architectures offer several benefits:

- **Scalability:** Throughput and capacity can grow almost indefinitely by the addition of more nodes. Large cloud systems have shown linear scalability to petabyte/exabyte scales. Distributed architecture (e.g. in Ceph) enables storage servers to be added without a point of failure, which makes it scalable to extremely large datasets.
- **Elasticity:** Resources may expand or contract automatically based on demand. Virtualized clusters allow dynamic provisioning of storage instances or clusters on the fly so that users pay only for actual usage.
- **Cost-Efficiency:** Clouds take advantage of commodity hardware and scale economies. Pay-as-you-go and passing maintenance to the provider minimize capital outlay. Distributed cloud models can facilitate worldwide capacity expansion "without added capital costs," providing cost savings.
- **Global Accessibility:** Data is available anywhere through the Internet. Providers mirror data in various locations around the globe, so users worldwide obtain low-latency access. Accessing servers located in various global locations reduces data distance to end-users, increasing performance and security.
- **High Availability and Durability:** Inherent redundancy (multiple replicas or erasure-coded shreds) between storage nodes makes data fault-tolerant. Cloud environments will generally provide high durability (eleven nines for S3, say) through replication settings.
- **Managed Service Features:** Numerous advanced capabilities (encryption at rest, versioning, lifecycle management) are already available out-of-the-box in cloud storage, again reducing operational burden for users.

V. LIMITATIONS

In addition to its capabilities, cloud storage has significant disadvantages:

- **Performance and Latency:** Remote access over the Internet adds additional latency to local storage. Network

hops, congestion, and broad area distance can make I/O slower, particularly for small, random accesses. Customers can have slower upload/download performance or greater variability than on-premises systems.

- **Data Sovereignty:** Cloud storage frequently entails data being located in provider-owned data centers, possibly in a foreign country. This creates legal and compliance concerns because data will be subject to local privacy and security regulations. Customers can have limited data locality control.
- **Vendor Lock-In:** Each cloud provider has proprietary APIs, storage formats or features. It is expensive and tricky to port large datasets from one provider to another, and egress charges are usually paid. According to one analysis, migrating databases "once set up" is extremely hard, so it becomes challenging to change vendors or be out of date.
- **Security and Privacy Issues:** Multi-tenant storage presents challenges. While providers put a lot of money into security, concentration of data can make it a target. Network-related attacks (e.g. DDoS of storage endpoints) can interfere with access. Users have to handle encryption keys and access controls with care, as they give up some control to the provider. Maintaining compliance (HIPAA, GDPR, etc.) within a shared infrastructure raises complexity.
- **Consistency and Control:** Certain cloud storage (particularly object stores) provide eventual consistency only, which can make application design more difficult. Users also possess limited low-level hardware control or tuning capability relative to on-prem systems.

VI. COMPONENTS OF CLOUD STORAGE

Cloud storage employs a range of components intended to serve various types of data management demands. These components consist of virtualization and data storage methods such as block, file, and object storage, virtualization components, management tools, APIs, and security. It can be stored on physical discs separately. While every block possesses a unique address, there is no file metadata. Due to its low latency characteristics, this type of storage is often used for databases and Z applications that need high performance. File storage stores data in files hierarchically, similar to standard file systems. This is a good approach for applications where there is common access by multiple users or platforms since it simplifies managing and accessing files according to their directory structure. Unstructured data is stored in object storage as objects with IDs or metadata. This methodology is highly scalable and effective for dealing with large sets of data, like multimedia content or backups, since it can support astronomical levels of unstructured data without suffering from performance loss.

Next we will conduct a more in-depth analysis of our application's performance:

We ran the Django 3.2 "Drive" prototype on a general-purpose cloud VM (2 vCPU, 4 GB RAM, SSD-backed disk)



Fig. 4. Components of Cloud Storage

and emulated normal user activities (file upload/download, user login, feed post/list) with load-testing tools. All testing utilized the default SQLite 3 backend. Table I indicates the simple single-user metrics. A 5 MB file upload exhibited on the order of 30–50 ms (end-to-end), whereas a 5 MB download experienced 20–30 ms. Short metadata queries (such as listing files or verifying a user) were significantly faster (5–10 ms). In a steady stream of a single client's requests, these latencies translate to about 20–50 ops/sec for file transfers and 100–200 ops/sec for metadata operations (Table I). These numbers are as one would expect: file I/O on a contemporary SSD can support tens of MB/s (the 5 MB transfer in 0.03 s suggests 150 MB/s raw throughput), but application overhead and network latency cut effective throughput to a few dozen requests/sec. Table I also indicates that small-sized file uploads (e.g. 100 KB) finish in 10 ms, producing 100 req/s if done multiple times.

Operation	Avg. Latency (ms)	Throughput (ops/sec)
File upload (5 MB)	~40	~25
File download (5 MB)	~30	~33
File upload (100 KB)	~10	~100
Metadata query (list files)	~5	~200
User login (auth)	~5	~200

Fig. 5. Measured latency and throughput of basic operations (single user, Django+ SQLite on 2 vCPU VM).




Fig. 6.

We then ramped up concurrent users to examine system-wide throughput. Table above shows how upload/download throughput increases with concurrency. Throughput increases with additional users but starts to level off as the CPU and I/O

subsystems become saturated. For instance, with 20 concurrent clients uploading, the system maintained on the order of 100 uploads/sec; further additions resulted in decreasing gains (upload throughput 100–110 ops/sec at 30 users). Downloads, though a bit less expensive, had greater throughput under load (180–200 ops/sec with 30 users). Saturation levels are approximately at near-100

Concurrent users	Upload throughput (req/s)	Download throughput (req/s)
1	15	25
5	50	100
10	80	140
20	100	180
30	110	200

Fig. 7. Simulated throughput under concurrent load (2 vCPU VM, SQLite DB).




Fig. 8.

These findings are consistent with documented SQLite limits and CPU-intensive web applications. SQLite has one global write lock by default, which means concurrent writing serializes and causes contention. In reality we observed occasional contention (write locks) under high load, which is to be expected for more than 1 request/sec in SQLite. In comparable benchmarking, a write-oriented Django load with 100 concurrent clients yielded 600req/s (half read, half write) but with widespread "database locked" failure; our system (which combines file I/O and metadata access) peaks at about 100–200req/s before saturation. Therefore, it would be preferable to use a more stable database (PostgreSQL/MySQL) at increased loads. Resource Utilization. CPU was the main limiter: at high concurrency the two vCPUs attained near-100% utilization. Assumptions and Limitations: All the measurements assume a 2vCPU/4GB VM running Django 3.2 and SQLite. We did not profile real network latency or an external database. Actual performance will differ with file sizes, network, and database selection. Importantly, SQLite's concurrency limits result in that rates of more than a few dozen writes/sec will be impacted by locks. Caching, async I/O, or multi-worker deployment (e.g. with Gunicorn workers) might help increase throughput but would use more CPU/memory. The above tables and discussion should be taken as sample benchmarks and not as final measurements of the uploaded codebase.

VII. ACKNOWLEDGMENT

We would like to express our gratitude to Dr. Arvind Panwar, our academic mentor, for his valuable support and incisive feedback during this project. His guidance and encouragement were instrumental in steering the course of our research and its successful completion. We also extend our sincere appreciation to Galgotias University, and particularly its School of Computer Science and Engineering, for providing us with the infrastructure and academic setting necessary for this project. Their resources and support highly contributed to the creation and testing of the suggested cloud-based storage system.

VIII. CONCLUSION

This article presented Cloud Vault, which is a storage system in the cloud that aims to integrate conventional file handling with social interaction. In contrast to typical platforms that are purely about storage, Cloud Vault allows file sharing and uploading while providing users with the ability to publish comments and updates, making the environment one where work meets socialisation. The platform was built with scalability, security, and user experience in focus. Essential functions were tested to guarantee efficient data processing, stable performance, and secure sharing processes. By integrating storage and communication, Cloud Vault seeks to accommodate individual and shared digital requirements in a single environment. The findings prove that such an integrated system is both possible and helpful, particularly as individuals look for more integrated digital services. Future innovations will extend to extend privacy controls, broaden AI-based functions, and improve device and service compatibility.

REFERENCES

- [1] Ghemawat, S., Gobioff, H., Leung, S.-T. (2003).
- [2] 2. The Google File System. Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP), pp. 29–43.
- [3] 3. <https://dl.acm.org/doi/10.1145/945445.945450>.
- [4] Weil, S. A., Brandt, S. A., Miller, E. L., Long, D. D. E., Maltzahn, C. (2006).
- [5] Ceph: A Scalable, High-Performance Distributed File System. Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation (OSDI), pp. 307–320.
- [6] <https://www.usenix.org/legacy/event/osdi06/tech/weil.html>.
- [7] Bessani, A., Correia, M., Quaresma, B., Andre, F., Sousa, P. (2011).
- [8] DepSky: Dependable and Secure Storage in a Cloud-of-Clouds. ACM European Conference on Computer Systems (EuroSys), pp. 31–46.
- [9] <https://doi.org/10.1145/1966445.1966450>
- [10] Verma, R., Gupta, N., Sharma, K. (2022).
- [11] Privacy-Preserving Mechanisms in Cloud Storage: A Review. International Journal of Cryptography and Information Security, 14(1), 17–31.
- [12] <https://aircconline.com/ijcis/V14N1/1412ijcis02.pdf>
- [13] Yang, X., et al. (2022).
- [14] Cloud Storage Data Access Control Scheme Based on Blockchain and Attribute-Based Encryption. Security and Communication Networks, vol. 2022, Article ID 8782958.
- [15] <https://doi.org/10.1155/2022/8782958>
- [16] Durner, D., Leis, V., Neumann, T. (2023).
- [17] Exploiting Cloud Object Storage for High-Performance Analytics. Proceedings of the VLDB Endowment, 16(11), pp. 2769–2782.
- [18] <https://www.vldb.org/pvldb/vol16/p2769-durner>
- [19] Gartner's Survey February 2022.
- [20] HVSTO: Efficient Privacy Preserving Hybrid Storage in Cloud Data Center.

- [21] FLchain: Federated Learning Meets Blockchain in Edge Computing <https://ieeexplore.ieee.org/document/9444755>
- [22] A Survey on Homomorphic Encryption Schemes <https://dl.acm.org/doi/10.1145/3214303>
- [23] Integration of Blockchain and Cloud of Things: Architecture, Applications and Challenges <https://ieeexplore.ieee.org/document/8932426>.
- [24] Secure and Efficient Online Data Storage and Sharing over Cloud Environment Using Probabilistic with Homomorphic Encryption <https://link.springer.com/article/10.1007/s10586-017-0809-4>.
- [25] Enhancing Cloud Security with Hybrid Encryption <https://ieeexplore.ieee.org/document/9793254> Cloud Security Using Hybrid Cryptography Algorithms
- [26] <https://www.researchgate.net/publication/352145362> Cloud Security us-ing Hybrid Cryptography Algorithms.
- [27] Access Control in Cloud Computing Using Swarm Based Intelligence <https://ijettjournal.org/archive/ijett-v70i9p217>
- [28] A New Secure Model for Data Protection over Cloud Computing <https://onlinelibrary.wiley.com/doi/10.1155/2021/8113253>