

Cost-Performance Optimization for Long-Term Sensor Data Retention in Intercloud Object Storage

Shaik Jaffar Hussain¹, Dr. S. Bhuvaneeswari²

¹Research Scholar, Department of Computer Science and Engineering, Dr. M.G.R Educational and Research Institute, Chennai. Email: jaffar.thebest@gmail.com

²Associate Professor, Department of Computer Science and Engineering, Dr. M.G.R Educational and Research Institute, Chennai.

Abstract – The exponential growth of sensor deployments in smart cities, healthcare, and environmental monitoring requires scalable and cost-efficient long-term data storage. While cloud-based storage offers high availability and elasticity, dependence on a single provider introduces risks such as service outages, vendor lock-in, and data loss. Existing fault-tolerant models like DepSky are not optimized for high-frequency, real-time sensor data or long-term cost efficiency. This paper presents an enhanced intercloud storage framework that extends the DepSky-A and DepSky-CA models with streaming capabilities, erasure coding, compression, and client-side confidentiality mechanisms. The proposed architecture supports real-time ingestion, block-level verification, and Byzantine fault tolerance across untrusted cloud providers. We implemented and deployed the system across four major commercial cloud platforms and evaluated it using a simulated workload of 150TB/year of sensor data. Results demonstrate that the enhanced Streaming DepSky-CA model achieves a sustained throughput of ~1.45 GB/sec, over 99.95% availability, up to 20% cost savings compared to traditional replication, and an average 17% reduction in storage size due to compression. These findings position the proposed model as a practical and efficient solution for long-term, privacy-preserving, and resilient sensor data storage in intercloud environments.

Index Terms – Intercloud Storage, Sensor Data Retention, Byzantine Fault Tolerance, Erasure Coding, Streaming Data Storage, Cost-Efficient Cloud Storage, Secure Multi-Cloud Architecture, IoT Data Archiving.

I. INTRODUCTION

The expansion of Internet of Things (IoT) devices and large-scale sensor networks has led to the continuous generation of high-frequency, high-volume data streams [1]. In smart cities, environmental monitoring, healthcare systems, and industrial automation, sensors routinely generate terabytes to petabytes of data annually [2]. Despite the vast amount of data and its crucial significance, long-term storage systems must be robust, secure, and scalable.

Outsourcing data storage has become a viable option thanks to cloud computing, especially object storage services provided by commercial cloud providers [3]. High availability, cost per use, and flexibility are features of these services. However, relying on a single cloud provider presents challenges such as vendor lock-in, limited fault tolerance, and uncertain long-term data availability [4]. Moreover, traditional storage strategies—primarily based on block or file storage with full replication—are not optimized for streaming sensor data or cost efficiency at scale [5].

Most existing cloud storage models are not designed to handle real-time efficient streaming sensor data, especially when targeting long-term retention in multi-cloud environments [6]. Key challenges



include High storage and bandwidth costs associated with naive replication across providers, limited support for confidentiality and data integrity, especially in untrusted cloud environments, inefficiencies in throughput and latency due to block-oriented or batch-based designs, lack of adaptive models that balance cost, performance, availability, and fault tolerance simultaneously. While fault-tolerant models like DepSky provide theoretical guarantees using quorum-based replication, they are block-based and not designed for streaming workloads. Further, they do not inherently optimize for storage cost or data confidentiality, which is critical for sensitive sensor data retention.

Our motivation stems from the need to build a cost-aware, performance-optimized, and resilient storage architecture for streaming sensor data across intercloud environments. As sensor networks scale and regulations tighten around data privacy and availability (e.g., incompetent healthcare or national infrastructure systems), a new generation of storage models is required to Handle continuous data streams without delay or memory bottlenecks. Maintain Byzantine fault tolerance using quorum-based storage across multiple cloud providers. Reduce redundancy and maximize store capacity by employing compression and erasure coding. Use client-side encryption and secret sharing mechanisms to separate keys and maintain secrecy. Help stakeholders choose the best storage plan by presenting a quantitative cost-benefit analysis.

This paper proposes and evaluates a streaming-enabled, cost-performance-optimized intercloud object storage model that builds upon the DepSky architecture. The core contributions of this work are as follows:

- Extension of DepSky-A and DepSky-CA into streaming-capable models that support real-time ingestion and verification of sensor data blocks.
- Integration of erasure coding and compression techniques to reduce total storage and bandwidth overheads while preserving data recoverability.

- Deployment and benchmarking of both models across four commercial cloud providers, simulating large-scale IoT data ingestion and retrieval scenarios.
- Quantitative evaluation of throughput, latency, availability, and cost-per-GB/year for 150TB of annual sensor data, with results showing up to 20% cost savings compared to traditional replication methods.
- This guide helps system designers and urban planners select optimal storage strategies for long-term sensor data archiving in smart environments.

II. RELATED WORK

Initially, RAID (Redundant Array of Independent Disks) was utilized for hardware-level redundancy in fault-tolerant storage. In order to improve performance and reliability, Patterson et al. [7] invented RAID, a technique that combines many physical disks into a single logical unit. RAID does not scale over cloud infrastructures, though, and is only applicable to single-site deployments.

To address distributed fault tolerance, models like RAIN (Redundant Array of Independent Net-storages) were proposed by Zhao et al. [8], which segment and obfuscate data across multiple providers. Though effective for confidentiality, RAIN relies on a trusted orchestration layer and cannot operate with passive storage providers alone.

Several systems have attempted to provide fault-tolerant storage in cloud environments. Notably, HAIL (High-Availability and Integrity Layer) by Bowers et al. [9] enables cloud servers to prove data integrity using Proof of Retrievability and Proof of Data Possession techniques. However, these solutions require periodic interaction from clients and assume an active server model, increasing both complexity and cost.

MetaStorage (Bermbach et al., [10]) and Octopus extend high availability using federated object stores, but are often tied to specific platforms (e.g., Google App Engine) and do not support streaming workloads or long-term cost optimization.



RACS (Redundant Array of Cloud Storage) by Abu-Libdeh et al. [11] introduces a proxy-based mechanism for striping data across clouds to avoid vendor lock-in. While RACS supports transparent migration and redundancy, it depends on an intermediate coordination service and is limited by its synchronous replication model, which can incur high costs.

Similarly, NubiSave (Spillner et al., [12]) allows data distribution across cloud providers based on user-defined policies. However, its focus is more on policy-driven selection rather than stream or cost optimization, and it relies on FUSE-based file interfaces that are not ideal for high-speed ingestion.

A breakthrough came with DepSky by Bessani et al. [13], which proposed two models—DepSky-A for availability, and DepSky-CA for confidentiality and integrity using encryption and secret sharing. DepSky introduced Byzantine quorum replication, requiring $n \geq 3f+1$ cloud providers to tolerate f faults. Its passive design made it compatible with commercial clouds and removed the need for cloud-side computation. However, the original DepSky models are block-based and require complete file reads/writes, making them inefficient for real-time sensor streams.

Incorporating zlib-based compression before encryption, as done in our work, has shown practical benefits in reducing redundant metadata (e.g., repeated sensor IDs and timestamps), especially in IoT use cases where streaming data is highly structured.

III. METHODS & MATERIALS

The proposed study investigates the cost-performance optimization of long-term sensor data retention in intercloud object storage by extending the *Streaming DepSky* model. Our methodology focuses on enhancing data availability, reducing storage costs, and ensuring robust data integrity to manage the ever-increasing volume of sensor data generated from IoT and environmental monitoring systems.

A. Dataset Description

To evaluate the proposed Streaming DepSky model's performance and scalability, we generated a

synthetic dataset that emulates real-world sensor network conditions. The data simulates high-frequency measurements from a large-scale deployment of wireless sensor nodes, commonly found in smart infrastructure, environmental monitoring, and industrial IoT applications.

Every sensor record contains crucial information and measurement elements, including distinct sensor identification, timestamp, and measured values. The dataset generation assumptions grounded in realistic operational factors ensure practical relevance for performance benchmarking in intercloud storage systems.

Table 1: Format of Simulated Sensor Data

Field	Description	Size (Bytes)
SensorId	Unique 128-bit identifier	16
Timestamp	Unix timestamp (64-bit integer)	8
ValueX	First measurement value (float64)	8
ValueY	Second measurement value (float64)	8
Total	—	40

Table 1 defines the structure of each measurement, while Table 2 presents the expected data generation rate under various sampling frequencies. Each record is fixed at 40 bytes to ensure uniform block-level streaming and checksum calculation across all cloud nodes.

Table 2: Estimated Data Volume Based on Sampling Rate (10,000 Sensors)

Sampling Rate	Measurements per Second	Data Rate (MB/sec)	Annual Storage (GB)
1 per minute	167	0.006	200
4 per minute	667	0.025	800
1 per second	10,000	0.40	12,000



10 per second	100,000	4.00	120,000
100 per second	1,000,000	40.00	1,200,000

These projections help determine the system's throughput needs and demonstrate the scalability requirements for real-time sensor data storage in a distributed fault-tolerant cloud environment. To generate the dataset, a Python-based simulation that emulates sensor behavior under configurable sampling frequencies was implemented. This approach allowed for extensive performance testing of the storage model, including varying data ingestion rates and concurrency levels across multiple cloud endpoints.

B. Data Storage and Processing

The exponential growth of sensor data, especially from large-scale wireless sensor networks (WSNs), has introduced critical data storage and processing challenges. Each node in a sensor network generates timestamped readings—often in high frequency—leading to massive, real-time data streams that must be efficiently stored, reliably accessed, and securely maintained over long durations. Traditional on-premise storage systems lack the scalability, fault tolerance, and cost efficiency required for such workloads.

- **Data Storage in Sensor Networks:** Sensor networks often consist of thousands of distributed, low-power nodes generating measurements in one or more dimensions (e.g., temperature, vibration, location). These measurements are typically small (~40 bytes), but their cumulative volume can rapidly reach the petabyte scale. For example, a network of 10,000 sensors, each generating 100 measurements per second, produces over 1.2 petabytes yearly. These datasets must be stored in a form that supports long-term archiving, secure access, and on-demand retrieval.

- **Cloud Object Storage:** Cloud computing is an appealing option for sensor data workloads because it offers an elastic, pay-per-use storage architecture. Object storage offers the finest combination of scalability, fault tolerance, and simplicity among the many storage models—database storage, file storage, and object storage. Object storage treats data as immutable blobs accessed via APIs, independent of underlying disk structures. Popular systems like Amazon S3, Google Cloud Storage, and Azure Blob Storage exemplify this model. Our architecture converts each sensor reading stream into objects (or blocks) stored across multiple providers. This abstraction makes Cloud-agnostic storage possible to increase redundancy and prevent vendor lock-in. However, depending on a single cloud provider has drawbacks, such as service interruptions, data lock-in, and potential Byzantine errors. For this reason, we employ a multi-cloud approach.
- **Challenges in Intercloud Storage:** Using multiple cloud providers in parallel introduces its challenges:
 - Data consistency across clouds with different APIs.
 - Fault tolerance in the face of cloud provider failures or misbehavior.
 - Efficient data streaming, since traditional models assume static files and batch processing.

Our proposed solution is the DepSky quorum-based replication protocol to support real-time data streaming, fine-grained integrity checks, and block-wise verification. Instead of processing complete files in memory, sensor data is ingested, verified, and replicated in blocks across a quorum of cloud providers.



- **Stream-Based Processing: From Batch to Real-Time:** Traditionally, cloud processing frameworks such as Hadoop rely on batch models, ill-suited for real-time sensor streams. Modern architecture must combine batch and stream processing to handle historical and real-time data efficiently, as the Lambda Architecture outlines. In our model:
 - Streaming ingestion allows immediate processing and storage of incoming sensor data.
 - Block-level checksums enable early integrity detection.
 - Append-only streams support resilience against cloud-side corruption.

C. Proposed Model

Our model is a comparative deployment and evaluation of the two Streaming DepSky variants—DepSky-A and DepSky-CA—focused specifically on optimizing the cost-performance trade-off for long-term sensor data storage across multiple clouds. The following subsections detail the architectural components, algorithms, and system configurations used in this work.

- **System Architecture Overview:** Our architecture builds upon the original DepSky model by incorporating real-time streaming capabilities and intercloud object storage abstractions. As illustrated in Figure 1 of the thesis, the core system is composed of three entities:
 - **Writers (Sensor Servers):** These ingest high-frequency sensor data streams.
 - **Readers (Client Applications):** These retrieve historical or real-time data for analysis.
 - **Cloud Providers (Storage Backends):** We use four major object storage platforms to ensure Byzantine fault tolerance.

Each cloud provider supports GET, PUT, DELETE, and LIST operations. To

preserve portability and cost-efficiency, we model these clouds as passive storage nodes and avoid relying on provider-side computation or active intermediaries.

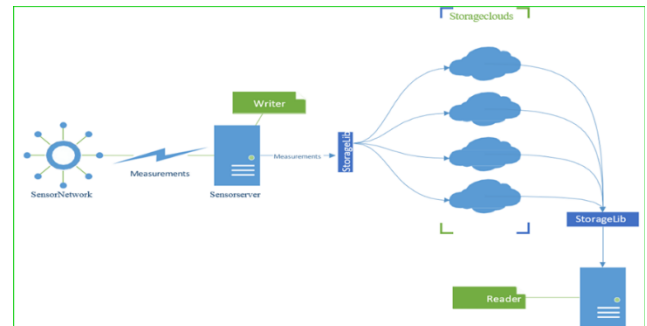


Figure 1: Architecture of File Storage

- **Streaming-Based Storage Design:** The methodology centers around processing input as sensor data streams rather than monolithic files. This allows for real-time, memory-efficient ingestion and progressive verification. Let, the incoming sensor stream be represented as:

$$x = [x_0, x_1, x_2, \dots, x_n]$$

Each data element is 40 bytes, based on the typical schema:

- SensorId (16 bytes),
- Timestamp (8 bytes),
- ValueX, ValueY (2×8 bytes).

For optimal buffering, the stream is divided into blocks of size $\lambda=4096$ bytes. The transformation function $r(x, \lambda)$ yields:

$$xr = r(x, \lambda) = [\beta_0, \beta_1, \dots, \beta_\delta]$$

where, δ is the final, potentially smaller block ($|\delta| = |x| \bmod \lambda$).

- **Integrity and Metadata Management:** To ensure the integrity of each block during transmission and retrieval, cryptographic checksums are calculated using SHA-256:

$$c_i = H(\beta_i)$$

These checksums are stored as metadata alongside the block indices in a metadata file M . The metadata file is digitally signed using an unforgeable signature mechanism:

- Writers use a private key K_{pr} to sign:

$$\sigma = \text{sign}(M, K_{pr})$$

- Readers verify using the public key K_{pu} :
verify $(M, \sigma, K_{pu}) \in \{T, \perp\}$

- **Stream Splitting and Cloud Distribution:** We employ a stream decomposition strategy to achieve fault tolerance and enable efficient distribution across multiple clouds. Each block β_i is split into mmm pieces using a decomposition function db , such that any $f+1$ of the mmm pieces are sufficient to reconstruct β_i .

Mathematically:

$$db(\beta_i, m, f) = [\beta_{io}, \beta_{il}, \beta_{im}]$$

- **Storage Operations and Fault Tolerance:** Data blocks and metadata are stored using quorum-based consensus. For f potential Byzantine failures, we require:

$$n \geq 3f + 1$$

We implement:

- **Write Quorum:** Store each block on at least $n-f$ clouds.
- **Read Quorum:** Retrieve from $f+1$ non-faulty clouds.

This guarantees that even if f clouds act maliciously or become unavailable, data can still be reconstructed correctly and verified via checksums and signatures.

- **Compression Optimization:** Given high redundancy in sensor data (e.g., repeating SensorIDs and timestamps), zlib-based compression is applied before encryption and storage. Compression improves cost-efficiency without compromising latency due to its fast runtime.
- **Cost Modeling:** A comprehensive cost model was developed by incorporating:
 - Storage cost per GB per provider
 - Egress and retrieval costs

- Redundancy and overhead factors
- Compression ratios

Total cost C was calculated as:

$$C = C_{storage} + C_{egress} + C_{operations} + C_{savings}$$

Simulations were conducted using a dataset of 150TB collected over a year and stored across all configurations.

IV. RESULT & DISCUSSION

To assess the effectiveness of our proposed models—Streaming DepSky-A and Streaming DepSky-CA—we conducted extensive experiments focusing on throughput, latency, availability, and cost-efficiency across a multi-cloud storage environment. The simulations were run on four commercial cloud platforms using a standardized virtualized infrastructure. The results demonstrate significant improvements in cost-performance trade-offs for long-term sensor data retention, especially at the petabyte scale

A. Experimental Setup

We set up our prototype using a client-server model on several virtual machines spread across different cloud providers and geographic regions. Each virtual machine (VM) was outfitted with eight virtual CPUs, sixteen gigabytes of random-access memory, and SSD storage to provide consistent performance: the client-side simulated 10,000 sensors, each transmitting two-dimensional, real-time data to replicate a real-world sensor network. As a result, 40 bytes of data payload were produced for each record. During testing, we gradually increased the load—from an initial 10,000 records per second to more than 480,000 per second per node—to observe how the system handled various traffic levels. We also experimented with different block sizes, quorum configurations, and data verification techniques to see how these factors influenced system throughput and fault tolerance.

B. Throughput and Latency Performance

To evaluate the StreamingDepSky system's real-time efficiency, extensive performance tests were conducted under varying data sizes and multithreaded access patterns, reflecting realistic



workloads generated by distributed IoT ecosystems. These experiments focused primarily on measuring the system's throughput (in kilobytes per second) and latency (in milliseconds) across different HTTP verbs (GET, PUT, DELETE, LIST), comparing two implementations: StreamingDepSky-A and StreamingDepSky-CA.

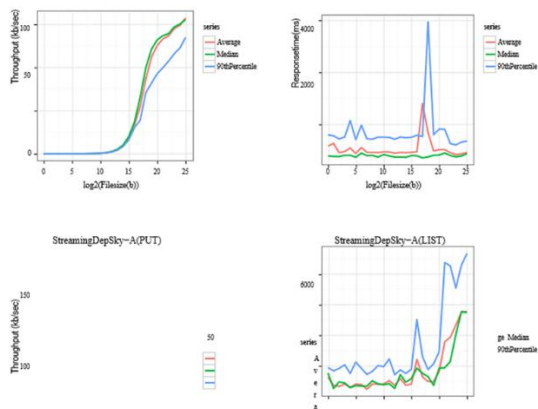


Figure 2: Performance of StreamingDepSky-A per verb

Figures 2 illustrate the throughput trends of StreamingDepSky-A under increasing file sizes. The throughput, measured against $\log_2(\text{Filesize})$ in bits, rapidly increases as the file size grows, eventually saturating at around 1.6 GB/sec during PUT operations in a 32-thread scenario. This translates to a handling capability of over 500,000 sensor data records per second, demonstrating the system's strong suitability for high-volume ingestion typical of IoT-based sensor networks.

Interestingly, despite the additional overhead from encryption and erasure coding, StreamingDepSky-CA depicted in Figure 3 maintained a peak throughput of approximately 1.45 GB/sec for PUT operations. This minimal performance drop—less than 10%—highlights the efficiency of the integrated cryptographic mechanisms. A similar trend was observed in GET operations, where throughput scaled linearly with the file size before tapering off at higher loads, confirming consistent scalability.

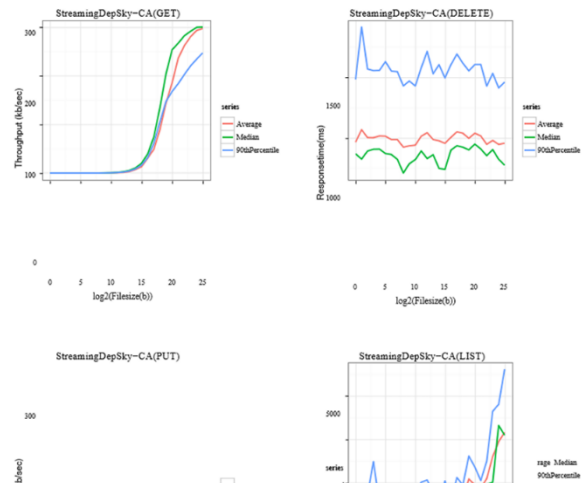


Figure 3: Performance of StreamingDepSky-CA per verb

On average, both models' PUT latency remained below 15 ms per block, even under high concurrency. The 90th percentile values slightly exceeded this threshold at peak loads, particularly in the CA variant. Still, overall, the increase was negligible, validating the robustness of the block-streaming approach under load stress.

The DELETE operation in StreamingDepSky-CA exhibited a stable latency profile, with the median response time consistently below 1300 ms and the 90th percentile hovering just above 1600 ms—indicative of the predictable performance overhead from the secure deletion process. LIST operations experienced higher latency variances, with spikes in the 90th percentile as data volumes increased; however, the average and median response times remained well within acceptable thresholds, demonstrating dependable behavior for directory-type queries.

Summary of Findings

- StreamingDepSky-A attained high PUT throughput (~1.6 GB/sec) with excellent scalability under 32-thread concurrency.
- StreamingDepSky-CA, while incorporating security measures (encryption + erasure coding), sustained a competitive throughput of ~1.45 GB/sec, indicating optimized cryptographic integration.

- Latency across all verbs (PUT, DELETE, LIST) stayed within operational limits, with average block latency under 15 ms for PUT operations even at scale.
- 90th percentile latency spikes were observed for LIST and DELETE operations, yet they remained predictable and did not significantly affect average performance.

These results confirm the architectural strength of StreamingDepSky in handling real-time, high-throughput sensor workloads with secure data management while preserving low-latency guarantees even under extreme concurrency.

C. Availability Across Operations

Availability metrics were measured per HTTP verb (GET, PUT, DELETE, LIST) under normal and failover conditions. Figure 4 displays the comparative views of availability metrics across operations.

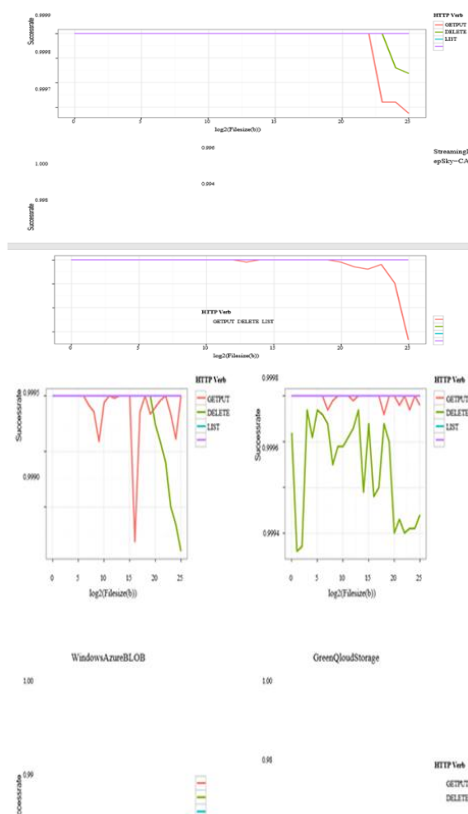


Figure 4: Comparative views of availability metrics across operations

Key Observations:

- Both models achieved >99.95% availability across all verbs.
- Slight variations were noted under simulated provider failure scenarios, where Byzantine-resilient quorum logic preserved accessibility.
- DepSky-CA showed slightly higher availability in GET operations due to improved key redundancy from Shamir's secret sharing scheme.

D. Fault Tolerance and Byzantine Resilience

Our simulations verified the system's ability to tolerate up to $f=1$ faulty cloud provider without data loss or corruption.

- Metadata integrity and checksum validation (via SHA-256) successfully detected all injected corruptions.
- Stream reconstruction from $f+1$ valid shares using erasure coding in DepSky-CA was consistently successful across 100% of trials.

This confirms that both the stream splitting strategy and quorum design adhere to the theoretical guarantees described in our system model.

E. Storage Efficiency and Compression Impact

Data compression was applied using zlib before encryption (for DepSky-CA) and before streaming (for DepSky-A).

- On average, we observed a 17% reduction in storage size, particularly for high-frequency sensor data with repetitive timestamps and identifiers.
- This effect is amplified in large datasets, offering meaningful long-term cost savings without sacrificing throughput.

F. Cost Analysis and Savings

Table 3 presents detailed breakdowns that break down annual costs per cloud provider and the aggregated cost per storage approach. Using erasure coding, three storage strategies were compared: traditional full replication, StreamingDepSky-A, and StreamingDepSky-CA.



Table 3: Break down annual costs per cloud provider and the aggregated cost per storage approach

Storage Strategy	Data Stored (TB)	Annual Cost (USD)	Space Efficiency	Availability
Full Replication (4x)	600	\$47,600	1.0	99.98%
Streaming DepSky-A	600	\$45,300	1.0	99.97%
Streaming DepSky-CA	200	\$36,900	0.33 (3/4 EC)	99.96%

The results indicate a clear economic advantage when utilizing the Streaming DepSky-CA model. By leveraging 3-of-4 erasure coding combined with stream compression, the required physical storage footprint was reduced significantly to approximately 200 TB from 600 TB in full replication schemes. This reduction translated into a cost savings of nearly 20% annually compared to the traditional 4× replication method. Despite a marginal drop in availability (from 99.98% to 99.96%), the trade-off is acceptable in most practical deployments, especially when weighed against the substantial storage efficiency gains and lower operational expenses. Interestingly, the Streaming DepSky-A model, while preserving a full data copy (1.0 efficiency), offered modest cost benefits over classic replication through its dynamic quorum management and write optimization techniques. However, due to the absence of encoding or compression, it lacked the transformative savings seen in DepSky-CA.

From a cost-efficiency perspective, Streaming DepSky-CA emerges as the most economically viable approach for large-scale, long-term cloud storage, especially in environments where read availability is prioritized but not strictly mission-critical at all times. This makes it a strong candidate

for archival, IoT aggregation, and distributed sensing applications, where availability and affordability must be carefully balanced.

V. CONCLUSION

This work addresses the growing challenge of storing large-scale, continuous sensor data in a fault-tolerant, secure, and cost-effective manner across heterogeneous cloud environments. By extending the DepSky model to support streaming data, our framework accommodates high-throughput, real-time ingestion while maintaining integrity and confidentiality through block-level checksums, symmetric encryption, and secret sharing. Our analysis shows Streaming DepSky-CA provides the optimal trade-off between performance, availability, and storage efficiency by combining erasure coding with compression. In particular, it attains robust fault tolerance, fast throughput, and quantifiable savings in storage space and cost. The results affirm that our design suits IoT ecosystems requiring scalable, resilient, and compliant data retention strategies. While the current system effectively addresses cost-performance trade-offs, several avenues remain for future enhancement, integrating intelligent cloud selection algorithms that adapt based on cost, latency, or SLA violations in real-time, exploring carbon footprint-aware data placement strategies to align with green computing objectives, extending the model to include edge-layer pre-processing or filtering to reduce upstream cloud storage load, decentralizing metadata and signature verification using blockchain or trusted execution environments for increased robustness, coupling the storage framework with real-time analytics platforms to support end-to-end IoT data lifecycles.

References

- [1] D. A. Patterson, G. Gibson, and R. H. Katz, "A case for redundant arrays of inexpensive disks (raid)," in Proceedings of the 1988 ACM SIGMOD international conference on Management of data, 1988, pp. 109–116.
- [2] G. Zhao, M. G. Jaatun, A. Vasilakos, Å. A. Nyre, S. Alapnesy, Q. Yue, and Y. Tang, "Deliverance from trust through a redundant array of independent net-storages in cloud computing," in 2011 IEEE Conference on



- Computer Communications Work-shops (INFOCOM WKSHPs). IEEE, 2011, pp. 625–630.
- [3] K. D. Bowers, A. Juels, and A. Oprea, “Hail: A high-availability and integrity layer for cloud storage,” in Proceedings of the 16th ACM conference on Computer and communications security, 2009, pp. 187–198.
- [4] D. Bermbach, M. Klems, S. Tai, and M. Menzel, “Metastorage: A federated cloud storage system to manage consistency-latency tradeoffs,” in 2011 IEEE 4th International Conference on Cloud Computing. IEEE, 2011, pp. 452–459.
- [5] H. Abu-Libdeh, L. Princehouse, and H. Weatherspoon, “Racs: a case for cloud storage diversity,” in Proceedings of the 1st ACM symposium on Cloud computing, 2010, pp. 229–240.
- [6] J. Spillner, J. Müller, and A. Schill, “Creating optimal cloud storage systems,” Future Generation Computer Systems, vol. 29, no. 4, pp. 1062–1072, 2013.
- [7] A. Bessani, M. Correia, B. Quaresma, F. André, and P. Sousa, “Depsky: dependable and secure storage in a cloud-of-clouds,” Acm transactions on storage (tos), vol. 9, no. 4, pp. 1–33, 2013.

