IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 6, June 2026

DOI: 10.17148/IJARCCE.2026.156100

A Survey on Network-Layer DDoS Resilience and Cryptographic Data Integrity Verification in Cloud Storage Infrastructures

Rachana V Murthy[1], Mohd Shaif[2], Ronaldo G[3], Satvik R[4]

Assistant Professor, CSE-IoT Cybersecurity with Blockchain, KS Institute of Technology, Bengaluru, India[1]

Student, CSE-IoT Cybersecurity with Blockchain, KS Institute of Technology, Bengaluru, India[2-4]

Abstract: Cloud storage systems increasingly face two distinct but interacting classes of threat: Distributed Denial-of-Service (DDoS) attacks that degrade or remove service availability, and silent data tampering carried out by ransomware or insiders that compromises data integrity without triggering an outage. Existing literature largely treats these as separate research problems — DDoS detection and mitigation studies evaluate only availability metrics, while data-integrity studies assume an always-reachable verification channel. This survey reviews the current state of research across four areas relevant to closing that gap: network- and application-layer DDoS detection, software-defined-network-based automated mitigation, client-side encryption for zero-knowledge cloud storage, and Merkle-tree/blockchain-anchored integrity verification. Eleven peer-reviewed and preprint sources are examined, organised thematically, and compared across detection scope, mitigation capability, and integrity guarantees. Four research gaps are identified, the most central being the absence, across all reviewed systems, of a single pipeline that simultaneously guarantees service availability and cryptographically verified data integrity. The survey concludes by outlining how a unified Defense-in-Depth architecture — combining adaptive DDoS mitigation with AES-256 client-side encryption and Merkle-tree integrity hashing — can close this gap.

Keywords: cloud security, DDoS mitigation, Merkle tree, AES-256 encryption, zero-knowledge storage, blockchain-anchored integrity, ransomware resilience, software-defined networking.

1. INTRODUCTION

Cloud storage has become the default substrate for critical data across healthcare, finance, government, and enterprise software — a shift that has concentrated risk as much as it has concentrated convenience. The same centralisation that makes cloud storage efficient also makes it an attractive target, and the threats it now faces fall into two categories that are frequently studied in isolation: attacks that take a service offline, and attacks that quietly corrupt what a service stores while leaving it apparently online.

Distributed Denial-of-Service attacks remain the most direct route to the first category. Volumetric Layer-4 floods exhaust bandwidth or connection-table capacity, while Layer-7 floods mimic legitimate application traffic closely enough that simple rate limiting fails to distinguish attacker from genuine user. Detection and mitigation research in this space has matured considerably, moving from static thresholding toward machine-learning classification and software-defined-network-based automated response.

Cryptographic integrity-verification research, built around hashing and Merkle-tree constructions, offers a mature toolkit for detecting this kind of tampering, but it is almost universally evaluated under the assumption that the network path to the verifier remains intact.

This survey examines both literatures side by side, with the specific aim of identifying where they fail to connect — the blind spot in which a server can appear "online" while the data it holds has already been compromised. Section II presents a taxonomy of the threat types involved. Section III surveys defense approaches across DDoS detection/mitigation and cryptographic integrity verification. Section IV compares representative systems directly. Section V concludes with the research gaps this comparison exposes and outlines how a unified architecture can close them.

Taxonomy of Threats to Cloud Storage Availability and Integrity

Threats to cloud-stored data can be classified along two axes —the layer they target and whether the resulting compromise is immediately observable. Table I summarises five representative threat types.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 6, June 2026

DOI: 10.17148/IJARCCE.2026.156100

Table I: Taxonomy of Threats to Cloud Storage Availability and Integrity

Threat Type	Mechanism	Target Layer	Example
Layer-4 Volumetric DDoS	High-rate UDP/ICMP/SYN floods saturate bandwidth or connection-table capacity	Network / Transport	SYN flood, UDP reflection
Layer-7 Application DDoS	Floods of seeminglylegitimate HTTP requests exhaust application resources	Application / Web	HTTP GET flood, Slowloris
Ransomware Tampering	Malicious software encrypts or overwrites stored files	Storage / Data	Crypto-ransomware payload
Insider / Unauthorised Modification	Privileged user or compromised credential alters files without detection	Storage / Data	Unauthorised edit or deletion
Silent Corruption (Bit-Rot)	Hardware-level degradation flips bits without triggering alerts	Storage Media	Disk / SSD degradation

This classification matters because a mitigation effective against one threat type provides no guarantee against another: a Layer-4 scrubbing appliance does nothing to detect a ransomware payload already inside the vault, and a Merkle-tree audit running on a fixed schedule does nothing to keep a service reachable during a volumetric flood. The two halves of the taxonomy — availability threats and integrity threats — are therefore not substitutes for one another, which is the starting point for the gap analysis in Section III.F.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 6, June 2026

DOI: 10.17148/IJARCCE.2026.156100

III. SURVEY OF DEFENSE APPROACHES

a. Network- and Application-Layer DDoS Detection

Much of the DDoS-detection literature treats the problem as traffic classification. Sharif, Beitol and Fazeli [1] built a machine-learning detector aimed specifically at application-layer floods produced by widely available toolkits such as HOIC and LOIC, and reported accuracy above 99% after a feature-selection step that also reduced computational overhead. This result is directly applicable to a Layer-7 detection stage in a unified pipeline, though the work evaluates detection in isolation and stops short of testing any corresponding mitigation or data-protection stage.

A complementary line of work targets the network layer using entropy and clustering rather than supervised classification. Hassan, Abd El Reheem, and Guirguis [3] combined entropy-based statistical detection with k-means clustering, evaluating the approach across three current intrusion-detection datasets and showing it can flag the sudden traffic anomalies characteristic of volumetric floods. The method is lightweight relative to deep classifiers, which makes it attractive for real-time monitoring, but —as with [1] —the evaluation is confined to detection accuracy and does not extend to what happens to stored data once a flood is underway.

Two broader reviews place these point solutions in context. A systematic review of cloud-specific DDoS defenses [4] concluded that effective protection requires layering controls across network, application, traffic-management, and monitoring levels rather than relying on any single technique. A companion survey of SDN-based defenses [5] catalogued detection approaches across cloud, IoT, and conventional network deployments, and noted that attack scale and sophistication continue to outpace static defenses — citing the hyper-volumetric HTTP floods that major cloud providers had to mitigate in late 2023. Neither review, however, connects network-layer resilience to anyguarantee about the integrity ofthe data the network is protecting.

b. SDN-Based and Automated Mitigation Architectures

Detection is only half of an effective defense; the response also needs to be automated if it is to keep pace with a live attack. Wang and Wang [2] proposed SDN-Defend, a hybrid CNN–Extreme Learning Machine architecture that performs anomaly detection and then closes the loop by issuing IP-traceback instructions and controller-level flowrules to filter malicious traffic in real time. This demonstrates that detection-to-mitigation can be fully automated within the SDN control plane, supporting the closed-loop design required at the network-defense stage of a unified architecture. The work is, however, confined to the control plane and offers no mechanism for verifying whether data behind the network layer remains untampered during or after an attack — the same limitation noted for the detection-only studies in Section III.A.

c. Client-Side Encryption and Cryptographic Integrity for Cloud Storage

A separate research stream addresses confidentiality and tamper-detection once data has reached the cloud. Chincholkar et al. [6] combined AES-256 client-side encryption with SHA-256 hashing in a framework benchmarked on AWS S3, generating a per-file cryptographic manifest that records hashes and encryption metadata for later verification. Their evaluation found that the combined AES-256/SHA-256 approach provided reliable tamper detection — caught through either AES-GCM authentication failure or hash mismatch — without materially affecting performance. This is the closest published precedent for a client-side AES-256 encryption stage, though the per-file manifest design does not aggregate hashes into a single, scalably verifiable structure.

A companion paper from an overlapping author group [7] extended the same cryptographic primitives toward ransomware resilience specifically, adding machine-learning anomaly detection based on file-entropy and modification-velocity heuristics, alongside versioned, retention-locked

backups. The system reported 94% ransomware-detection accuracy and 100% restoration success when reverting to pre-attack versions, with encryption overhead held under five percent of baseline latency. This work targets the “silent tampering” half of the threat taxonomy directly and shows that behavioural anomaly detection can complement cryptographic verification by flagging ransomware before corrupted data propagates into backups — though detection here is still bounded by the anomaly-scoring interval rather than continuous cryptographic re-verification.

d. Merkle-Tree and Blockchain-Anchored Integrity Verification

Merkle trees have a well-established lineage as a scalable, tamper-evident structure for cloud-storage auditing. Rao, Zhang, and Tu [8] introduced a Batch-Leaves-Authenticated Merkle Hash Tree to support outsourced auditing in settings where the third-party verifier cannot be fully trusted, improving on the efficiency bottlenecks of earlier proof-of-storage schemes. He, Shi, Huang, and Hu [9] proposed a “T-Merkle tree” variant anchored to a Hyperledger Fabric blockchain that supports binary search within tree branches; their prototype demonstrated improved storage utilisation alongside both security and efficiency gains over conventional binary structures.

Liu, Ren, Feng, Wang, and Wei [10] advanced this further with a Quad Merkle Tree audit scheme that removes the need for a trusted third-party auditor entirely by anchoring the tree root to a blockchain and verifying integrity through smart contracts, reporting significantly improved computing and storage efficiency relative to binary-tree schemes. This is the

DOI: 10.17148/IJARCCE.2026.156100

strongest direct precedent for anchoring a Merkle root to an append-only log, since it solves the same core problem — eliminating reliance on a single trusted verifier — through a structurally similar mechanism, while also showing that branching factor materiallyaffects audit performance.

An arXiv-hosted framework for cloud file versioning [11] extends the basic Merkle-proof scheme to support version control alongside tamper detection, allowing a system to distinguish legitimate updates from unauthorised modification rather than treating every change as a potential attack — a capability complementary to, but distinct from, pure tamper-evidence.

e. Summary of Reviewed Studies

Table II summarises the eleven sources reviewed above by focus area, methodology, and principal reported result.

Table II: Summary of Reviewed Studies

Ref	Focus Area	Methodology	Key Result
[1]	Application-layer (L7) DDoS detection	ML with feature selection on toolkit-generated floods	99.9% accuracy, 96% precision, 98% recall
[2]	SDN-based DDoS detection & mitigation	CNN-ELM anomalydetection + IP-traceback flowrules	Automated, real-time detection and mitigation
[3]	SDN DDoS detection	Entropy + k-means clustering across 3 IDS datasets	Effective against sudden/rapid traffic anomalies
[4]	Cloud DDoS defense (systematic review)	SLR of 2020–2024 ML/DL techniques	Confirms need for multi-layered defense
[5]	SDN DDoS survey	Taxonomy of entropy :ML/SDN- assisted detection	Documents record-scale 2023 attacks, defense gaps
[6]	Cloud confidentiality + integrity	AES-256 client-side encryption + SHA-256 manifest, AWS S3	Reliable tamper detection, minimal performance impact

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 6, June 2026

DOI: 10.17148/IJARCCE.2026.156100

[7]	Ransomware-resilient backup	AES-GCM + SHA-256 + ML anomaly detection + versioning	94% ransomware detection, 100% recovery, <5% overhead
[8]	Outsourced cloud audit	Batch-leaves-authenticated Merkle hash tree	Efficient auditing without a fully trusted auditor
[9]	Blockchain Merkle audit	T-Merkle tree on Hyperledger Fabric	Improved storage utilisation, binary search support
[10]	Blockchain-anchored integrity audit	Quad Merkle tree + smart contracts, removes trusted TPA	Higher compute/storage efficiency than binary-tree audits
[11]	Cloud file integrity + versioning	Merkle proofs with version control	Combines tamper detection with version tracking

f. Research Gaps

Synthesising across Sections III.A–D exposes four gaps that no individual study closes.

Gap 1 —DDoS defense and data-integrity research remain disjoint literatures. DDoS papers

[1]–[5] evaluate only availability metrics and never test whether stored data survives an attack uncorrupted; Merkle-tree and encryption papers [6] –[11] assume a reachable, unobstructed channel to the cloud and do not model verification under network-layer attack.

Gap 2 — Per-file hashing does not scale to fleet-wide tamper localisation. The manifest-based scheme in [6] verifies files individually rather than producing a single, logarithmically-verifiable fingerprint of an entire vault, while the Merkle-tree audit schemes [8]–

[10] are evaluated independently of any client-side encryption pipeline.

Gap 3 — Ransomware detection in [7] is reactive rather than continuous. Detection latency there is bounded by the anomaly-scoring interval rather than by continuous cryptographic re-verification; no reviewed system combines real-time behavioural detection with sub-five-second Merkle-root re-computation.

Gap 4 —Blockchain-anchoring overhead is unexamined for lightweight, single-organisation deployments. The schemes in [9] and [10] target multi-tenant scenarios with full consensus overhead, which may exceed what a single-organisation vault needs if it is only seeking tamper-evidence via an append-only log.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 6, June 2026

DOI: 10.17148/IJARCCE.2026.156100

Table III: Research the Reviewed Literature Gaps Identified in

Gap	Description	Unaddressed By
G1	No unified pipeline linking DDoS resilience to data-integrity verification	[1] –[5] assume reachable network; [6] –[11] ignore network-layer attack
G2	Per-file hashing/manifests don't scale tofleet-wide tamper localization.	[6] uses per-file manifests, not an aggregated Merkle root
G3	Anomaly-based ransomware detection lacks continuous cryptographic confirmation	[7] relies on periodic, not continuous, verification
G4	Blockchain consensus overhead unexamined for single-organisation deployments	[9], [10] target multi-tenant TPA-elimination scenarios

IV. CONCLUSION

This survey examined the current state of research relevant to building a cloud storage system that is simultaneously DDoS-resilient and cryptographically tamper-evident. Reviewing eleven sources across DDoS detection and mitigation, SDN-based automated response, client-side encryption, and Merkle-tree/blockchain-anchored integrity verification shows that each individual sub-problem is reasonably well studied, but that none of the reviewed systems unify availability and integrity guarantees into a single pipeline.

The four gaps identified — disjoint literatures, non-scalable per-file hashing, reactive-only ransomware detection, and unexamined blockchain-anchoring overhead for lightweight deployments — define the design requirements for a unified Defense-in-Depth architecture. A system combining real-time Layer-4/Layer-7 traffic monitoring and automated mitigation with AES-256 client-side encryption and continuously re-verified Merkle-tree hashing, anchored to a lightweight append-only log rather than a full multi-tenant blockchain, directly addresses each of these gaps and represents the most direct path toward closing them collectively rather than in isolation.

V. COMPARATIVE ANALYSIS OF EXISTING APPROACHES

Table IV compares the representative systems discussed above across four practical criteria: DDoS defense capability, data-integrity guarantees, whether the two are unified in a single pipeline, and whether integrity verification is blockchain-anchored.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 6, June 2026

DOI: 10.17148/IJARCCE.2026.156100

Table IV: Comparison of Existing Defense Approaches

System	DDoS Defense	Data Integrity	Unified Pipeline	Blockchain-Anchored
Sharif et al. [1]	Yes (L7 ML)	No	No	No
SDN-Defend [2]	Yes (L4/L7 CNN-ELM)	No	No	No
Chincholkar et al. [6]	No	Yes (AES+SHA manifest)	No	No
Chincholkar et al. [7]	No	Yes (AES+SHA+ML anomaly)	Partial (backup-focused)	No
He et al. [9]	No	Yes (T-Merkle+Hyperledger)	No	Yes
Liu et al. [10]	No	Yes (Quad Merkle+blockchain)	No	Yes
Proposed CRCI	Yes (L4 + L7)	Yes (AES-256 + Merkle)	Yes	Yes (lightweight, append-only)

No reviewed system occupies the bottom row of Table IV. Systems with strong DDoS defense ([1], [2]) provide no integrity guarantee at all, while systems with strong integrity guarantees ([6] , [7] , [9], [10]) provide no DDoS defense. The proposed Cyber-Resilient Cloud Infrastructure (CRCI) is positioned specifically to fill that row, combining an adaptive DDoS-mitigation shield with AES-256 client-side encryption and Merkle-tree integrity hashing anchored to a local, append-only log.

REFERENCES

[1]. D. M. Sharif, H. Beitollahi, and M. Fazeli, “Detection of Application-Layer DDoS Attacks Produced by Various Freely Accessible Toolkits Using Machine Learning,” IEEE Access, vol. 11, pp. 51810–51819, 2023, doi: 10.1109/ACCESS.2023.3280122.

[2]. J. Wang and L. Wang, “SDN-Defend: A Lightweight Online Attack Detection and Mitigation System for DDoS Attacks in SDN,” Sensors, vol. 22, no. 21, art. 8287, Oct. 2022, doi: 10.3390/s22218287.

DOI: 10.17148/IJARCCE.2026.156100

I. Hassan, E. Abd El Reheem, and S. K. Guirguis, "An entropy and machine learning based approach for DDoS attacks detection in software defined networks," Scientific Reports, vol. 14, Aug. 2024, doi: 10.1038/s41598-024-67984-w.

[3]. "Advancements in detecting, preventing, and mitigating DDoS attacks in cloud environments: A comprehensive systematic review of state-of-the-art approaches," Journal of King Saud University – Computer and Information Sciences, 2024. [Online] . Available: https://www.sciencedirect.com/science/article/pii/S111086652400080X
"A comprehensive survey on DDoS detection, mitigation, and defense strategies in software-defined networks," Cluster Computing (Springer Nature), 2024. [Online]. Available: https://link.springer.com/article/10.1007/s10586-024-04596-z

[4]. Chincholkar, A. Londhe, M. Joshi, P. Gole, and S. Mirgale, "Ensuring Confidentiality and Integrity in Cloud Storage using AES Encryption and SHA-256 Hashing," International Journal of Engineering Research & Technology (IJERT), vol. 14, no. 10, Oct. 2025, doi: 10.5281/zenodo.18080875.

[5]. Chincholkar, A. Chauhan, Y. Gher, P. Mogarkar, and A. Nirmal, "An Encryption-Based Automated Cloud Backup and Recovery Framework with Ransomware Resistance," International Journal of Engineering Research & Technology (IJERT), vol. 14, no. 10, Oct. 2025, doi: 10.5281/zenodo.18080873.

[6]. L. Rao, H. Zhang, and T. Tu, "Dynamic Outsourced Auditing Services for Cloud Storage Based on Batch-Leaves-Authenticated Merkle Hash Tree," IEEE Transactions on Services Computing, vol. 13, no. 3, pp. 451–463, May 2020, doi: 10.1109/TSC.2017.2708116.

[7]. K. He, J. Shi, C. Huang, and X. Hu, "Blockchain Based Data Integrity Verification for Cloud Storage with T-Merkle Tree," in Algorithms and Architectures for Parallel Processing (ICA3PP 2020),

[8]. Lecture Notes in Computer Science, vol. 12454, Springer, Cham, 2020, doi: 10.1007/978-3-030-60248-2_5.

[9]. Z. Liu, L. Ren, Y. Feng, S. Wang, and J. Wei, "Data Integrity Audit Scheme Based on Quad Merkle Tree and Blockchain," IEEE Access, vol. 11, pp. 59263–59273, 2023, doi: 10.1109/ACCESS.2023.3240066.

[10]. "Towards Immutability: A Secure and Efficient Auditing Framework for Cloud Supporting Data Integrity and File Version Control," arXiv preprint arXiv:2308.04453, 2023