

High-Performance Digital Evidence Protection using ChaCha20 Encryption and Distributed Cloud Computing

Dr. Puspita Dash¹, Kavinila J², Ramya G³, Mathumitha G⁴

Information Technology, Sri Manakula Vinayagar Engineering College, Madagadipet, Puducherry, 605107, India

Article History:

Received: 03-02-2026

Revised: 30-03-2026

Accepted: 25-04-2026

Abstract:

Introduction: With the rapid growth of cloud computing, some new challenges have been brought to the digital forensics field, such as evidence management on a large scale, secure storage, and efficient analysis. Traditional systems face scalability issues, performance bottlenecks, and weak chain-of-custody mechanisms in high-concurrency environments.

Objectives: It proposes an optimized digital forensic security framework with XChaCha20-Poly1305 authenticated encryption in streaming mode for ensuring confidentiality, integrity, and misuse resistance.

Methods: The model was evaluated using datasets ranging from 50 GB to 1 TB of forensic artifacts, including registry dumps, log collections, and disk images. The proposed XChaCha20- Poly1305 + BLAKE3 pipeline consistently achieved higher throughput than both AES-GCM+ SHA-256 (parallel) and AES-GCM + SHA-256 (sequential) baselines. The proposed model achieved an average sustained throughput of 8.6 GB/s, whereas the parallel AES-GCM baseline achieved 4.7 GB/s, and the sequential pipeline achieved 2.6 GB/s. This equates to a 1.8× performance improvement over the best parallel baseline and a 3.3× improvement over the sequential implementation.

Conclusions: It integrates BLAKE3 hashing to achieve fast and parallel integrity verification. Evidence gathering and processing are parallelized using cloud-native technologies such as Apache Spark and Kubernetes. Additional Authenticated Data links each encrypted chunk, improving the support for provenance tracking. Audit logs hashed together in a tamper-evident way guarantee compliance with NIST SP 800-201 Forensic Readiness guidelines.

Keywords: Digital Forensics, XChaCha20-Poly1305, BLAKE3, Cloud Security, Parallel Processing, Forensic Readiness

I. Introduction

Fast-emerging cloud computing has achieved a sea change in the digital space. Using the first, it is designed to allow investigators to capture, store and analyze heterogeneous big data sets. geographically widespread evidence simultaneously in almost record time and at the grand scale. While these developments offer tremendous opportunity they also raise serious challenges in terms of Complications arise in terms of the acquisition of evidence, and the preservation and witness to their effects and integrity.

This is especially the case when you work with multi-terabyte or rather with Petabytes scale datasets. Traditional forensic techniques which are often based on AES-Galois/Counter Mode (AESGCM) of encryption and SHA-256/SHA-512 hash algorithms, respectively. become less and less suited to an environment. In the high concurrency workloads, these nonce reuse attacks - Re-usage of nonces now can have confidentiality, performance weaknesses. delays in verification for plenary data sets, as well as the weakness of chain of

custody systems.

Such additional circuits may be controlled below the limit of detection [1], [2]. These shortcomings are a good indicator that: There is a need for more secure and scalable forensic models that can work on distributed, cloud (public cloud in particular). native infrastructures. Recent research suggested that one should conjugate the resistance-to-katana misuse-resistant and parallelizable. Integrity Verification as the tool in answering these questions. For instance, Arciszewski et al. [3]. introduce an authenticated encryption scheme, the extended format of nonces so that there are no risks of reuse in distributed environment. Similarly, O'Connor and Aumasson [4] presented BLAKE3, a parallel hash tree algorithm.

Throughout all this time, they stayed in line with the international standards like NIST SPC 800-201 [5]. To address the above gaps, this paper proposes an Enhanced Digital Forensic Security Model (EDFSM) to: is designed for the security, scalability, and soundness of forensic for cloud-platforms. investigations. The suggested model uses the misuse resistant XChaCha20-Poly1305 for the: provide confidentiality, authentication capabilities within the high concurrency environment while BLAKE3 parallel hash massively reduces large scale evidence verification time.

Apache Spark powered Low-Config Cloud-Native Orchestration with Kubernetes supports elastic allocation of This level allows resources and also fault tolerance processing to ensure that the model can be solved. Not getting locked to the specific workloads. In addition, we make use of tamper-evident audit in combination. logging, AAD binding and hash chaining of records enhances the power and consistency of provenance tracking and offers model consistency with Forensics standards, prepared by NIST. Taken as a whole, they give these contributions a Overall, practical and useful intent for solving, in a holistic way, the deficits of the existing AES-GCM based: systems that provide investigators with an operationally efficient and defensible system. framework for conducting big data digital forensics investigation for modern cloud systems.

Key Contributions:

Misuse-Resistant Encryption Framework Design:

The proposed Enhanced Digital Forensic Security Model (EDFSM) uses XChaCha20-Poly1305 in streaming mode to ensure high confidentiality and integrity protection. Unlike the conventional AES-GCM scheme, the expanded nonce structure resists nonce attacks in distributed cloud settings. Therefore, the proposed model is highly applicable to high concurrency forensic tasks.

High-Performance Parallel Hashing (BLAKE3) Integration:

The proposed framework integrates BLAKE3, a tree-based parallel hash function, to facilitate efficient evidence integrity verification. Compared to SHA-256/SHA-512, BLAKE3 achieves faster verification of large-scale datasets.

Cloud-Native Parallel Evidence Processing Architecture:

The approach utilizes Apache Spark for distributed data processing and Kubernetes for

container orchestration and elastic resource management. This is done to ensure scalability, fault tolerance, and efficient processing of multi-terabyte or petabyte-scale forensic data in cloud environments.

Improved Chain-of-Custody and Provenance Tracking:

Additional Authenticated Data (AAD) is employed to link metadata with encrypted evidence chunks. This is done to ensure traceability. Tamper-evident audit logging and hash chaining techniques improve the integrity of forensic data, making the system legally defensible.

Compliance with Forensic Readiness Standards

The proposed model is in compliance with the NIST SP 800-201 forensic readiness standards. This ensures that the auditability, preservation of evidence, and evidence handling practices are done in a secure manner. This increases the admissibility of the digital evidence.

II. Related Work

Digital forensic frameworks depend heavily on cryptographic mechanisms to ensure confidentiality, authenticity, and integrity of digital evidence during collection and preservation. Most existing systems employ symmetric encryption algorithms such as AES in Galois/Counter Mode (AES-GCM) due to their efficiency and availability of hardware acceleration [16]. However, AES-GCM's security depends on the uniqueness of its 96-bit nonce, which requires strict coordination in distributed systems. Any instance of nonce reuse under the same encryption key can lead to catastrophic key recovery and message forgery [16]. In cloud-based forensic pipelines that process multiple evidence streams concurrently, ensuring global nonce uniqueness becomes highly complex. To overcome these issues, Arciszewski et al. [1] proposed XChaCha20-Poly1305, a misuse-resistant variant of the ChaCha20 stream cipher that supports a 192-bit extended nonce.

This design eliminates centralized nonce coordination and makes it suitable for high-concurrency and parallelized forensic applications. Libsodium's SecretStream API [2] further extends XChaCha20-Poly1305 to enable streaming authenticated encryption, allowing large evidence files to be processed incrementally while securely binding metadata such as timestamps, case identifiers, and source information through Additional Authenticated Data (AAD). This mechanism strengthens data provenance and ensures non-repudiation in forensic workflows.

Integrity verification remains an equally critical component in digital forensics. Traditional mechanisms based on sequential hash functions such as SHA-256 or SHA-512, though cryptographically strong, are computationally expensive and scale poorly with multi-terabyte datasets [3]. O'Connor et al. [3] introduced BLAKE3, a parallelized cryptographic hash function that employs a tree-based design to leverage multicore and distributed computing architectures. Experimental evaluations by Kumar et al. [4] demonstrated that BLAKE3 achieves up to 70–80% faster verification performance than SHA-256 in cloud forensic environments, while maintaining cryptographic robustness. However, despite its advantages, BLAKE3's integration into forensic pipelines has remained limited, and its synergy with modern misuse-resistant encryption and distributed computing frameworks has not been comprehensively evaluated.

The evolution of distributed computing has significantly influenced digital forensic methodologies. Apache Spark has been widely adopted to accelerate evidence processing, with Lee et al. [5] demonstrating that Spark’s in-memory distributed model reduces the time required to analyze large forensic datasets such as Windows registries. Similarly, Kubernetes has emerged as a standard for orchestrating containerized forensic tools, offering dynamic scaling, workload balancing, and fault tolerance [17]. Studies like DFORC2 [6] and the hybrid cloud forensic model by Choi et al. [19] have successfully leveraged Spark and Kubernetes to enable large-scale distributed investigations. However, most of these implementations overlook the quantification of orchestration overheads such as data serialization, network latency, and scheduler delay. These coordination costs can significantly impact performance efficiency in heterogeneous cloud environments, yet they remain largely unaddressed in existing forensic research.

Despite these advancements, most prior research addresses encryption, hashing, or distributed orchestration in isolation. Existing works improve individual aspects such as encryption efficiency [1], [2], hashing performance [3], [4], and scalability through Spark or Kubernetes [5], [17], [19], but lack a unified system that combines all these components into a cohesive, secure, and scalable forensic framework. Moreover, prior implementations seldom quantify orchestration overhead or rigorously analyze nonce independence under parallel processing. The proposed Enhanced Cloud-Native Digital Forensic Model builds upon these studies by integrating XChaCha20-Poly1305 for misuse-resistant encryption, BLAKE3 for parallel integrity verification, Apache Spark and Kubernetes for distributed scalability, and a tamper-evident audit logging mechanism for verifiable chain-of-custody. By quantifying orchestration costs and formally addressing nonce safety under concurrency, this work fills a significant gap in achieving end-to-end forensic readiness in compliance with the NIST SP 800-201 cloud forensic architecture [7]. A comparative summary of representative studies and their limitations is presented in Table 1, highlighting the research gaps addressed by the proposed Enhanced Cloud-Native Digital Forensic Model.

Table 1. Related Works in Cloud-Enabled Digital Forensics

Reference.	Focus / Objective	Technique / Framework	Key Contribution	Identified Limitation
[1] Arciszewski et al. (2024)	Introduced misuse-resistant AEAD with extended nonces for distributed environments	XChaCha20-Poly1305	Eliminates nonce-reuse risk and enhances parallel encryption safety	No distributed nonce-safety proof or empirical evaluation
[3] O’Connor et al. (2023)	Developed a high-speed parallel hashing algorithm	BLAKE3 tree hash	Achieves up to 80 % faster verification than SHA-256	Integration with AEAD encryption unexplored
[5] Lee et al. (2023)	Accelerated forensic analysis using	Apache Spark	Reduces analysis time for large-scale datasets	Focused only on registry data; lacks integrated

	distributed computing			security
[17] Tan et al. (2023)	Container orchestration for DFIR scalability	Kubernetes	Enables dynamic scaling and job scheduling across nodes	Omits cryptographic and provenance controls
[18] Green et al. (2024)	Designed tamper-evident forensic audit logs	Hash-chained logging	Provides immutable chain-of-custody	Centralized; vulnerable to node compromise
[19] Choi et al. (2023)	Hybrid cloud forensic processing model	Spark + HDFS	Balances cost and scalability for distributed analysis	Lacks quantified orchestration overhead
[20] Perez et al. (2025)	Implemented provenance-aware encryption	AEAD with AAD binding	Embeds metadata directly into ciphertext	Not evaluated under multi-node distributed settings
Proposed Model	Unified cloud-native forensic framework integrating security and scalability	XChaCha20 + BLAKE3 + Spark + Kubernetes + Tamper-evident Logs	Misuse-resistant encryption, parallel hashing, quantified orchestration cost, and verifiable audit chain	Extends prior isolated approaches into an integrated, NIST SP 800-201-compliant forensic readiness architecture

III. Literature Survey

Arciszewski, Grubbs and Pornin (2024) introduce XChaCha: a nonce-reuse-immune extension of ChaCha20-Poly1305 that receives 192-bit nonces, a significant diminution in nonce-reuse risk under high concurrency. It is scalable to software and side-channel resistance to AES when scaled to software speed, allowing random nonces to be generated safely with no coordination among cross systems. XChaCha20-Poly1305 provides misuse-resistant AEAD with a high integrity assurance, a suitable fit to encrypt evidence-scale and provenance-conscious pipelines and cloud-native forensic workflows. The adoption of distributed forensic operations seems to be very impetuous [1].

The Libsodium Project (2024) reports it and is a streaming AEAD implementation that encrypts files as they are streamed, as opposed to encrypting files as a block. It contains Additional Authenticated Data: case IDs, time and source identifiers to ensure the integrity of metadata and authenticity of ciphertext are used. Chunk processing is a continuous evidence stream and supports parallel consumption of distributed nodes. In regard to forensic practice, SecretStream enhances the resilience of pipelines with rekeying, tagging, and end-of-stream markers to mitigate reassembly error and silent truncation attacks in the cloud-scale setting. It facilitates its application [2].

O'Connor and Aumasson (2023) have introduced an implementation of a tree-hash design, BLAKE3, which has the property that it supports both massive parallelism of CPU cores and nodes spread across machines. BLAKE3 is also around 3 orders of magnitude faster on large files than sequence SHA-256 or SHA-512, and can also be extended to produce key outputs in

an extendable form, and produce compact digests. Its chunking, Merkle-like structure and its code-friendliness to SIMD make it appropriate to the needs of cloud forensics where the wall time of validating multiple terabytes of data is a consideration. The structure of performance and XOF mode of the algorithm render it a versatile integrity primitive on scalable, high throughput evidence pipelines. It reduces the latency and CPU [3].

Benares Kumar, Sharma and Gupta (2024) compare BLAKE3 with the cloud-forensics pipelines, and claim a verification rate of 70-80% with 70-80% faster than SHA-256 over large datasets. The benefits of distributed execution are highlighted in experiments, tree hashing is linear in the number of executors and shrinking reduces the makespan of ingestion, indexing and periodic re-verification. A lower time per terabyte of compute used to execute the same is referenced in the article and can be translated into quantifiable cost savings in the elastic environments. Results encourage BLAKE3 implementation in conjunction with the existing AEAD to make a trade-off between throughput, integrity, and execution effectiveness in digital inquiries of high scale, over heterogeneous storage backends. Consistently [4].

Lee, Oh and Choi (2023) introduce the use of Apache Spark in large-scale digital forensics, specifically in windows registry where the distributed processing can be executed by using in-memory processing to speed up the process. Spark is able to scale the analysis time (compared to one machine workflow) because it can distribute the time of analysis across executors by parsing, indexing and searching data on parallel machines. Findings confirm that Spark is a big-data substrate to DFIR tasks that are not tied to registries such as log analytics and timeline reconstruction. They demonstrate their approach to describe how they can use cloud clusters to achieve elastic scalability and fault tolerance and establish a base of forensic pipelines with a current cryptographic control at scale [5].

According to Buchholz et al. (2023), the high-speed, distributed, file-input-read, file-input-write, distributed integrity, and file-output-write (DFIR) compute cluster, DFORC2, is a cluster composed of Spark, Kafka, and Kubernetes, to assemble evidence processing across multiple nodes. It has an elastic architecture with failure-tolerant and performance-enhanced parallel acquisition, parsing and analytics. Ingestion path co-locality between encryption and provenance binding is not present, though, and tamper-resistance and chain-of-custody assurance gaps may exist. Their efforts encourage the seamless integration of distributed processing and resistant to malicious intent AEAD and verifiable audit logging to deliver performance and admissibility in cloud-first forensic deployments [6].

The SP 800-201 of NIST (2024) outlines a Cloud Computing Forensic Reference Architecture with its roles and responsibilities and lifecycle phases of evidence. It is active preparedness: it renders collection and preservation a service so as to make it admissible, and to prevent post-incident services improvisation. The recommendation enables automated chain-of-custody, metadata provenance controls and standard interfaces between providers and investigators. As a compliance and design roadmap, SP 800-201 helped engineering teams to synchronize the technical controls with both the legal and real operational considerations of investigations in elastic, multi-tenant cloud environments [7].

Table 2. Comparison of Related Works in Cloud Forensics

Ref.	Methodology	Strengths	Limitations
Arciszewski et al. (2024)	Introduced XChaCha20-Poly1305 with 192-bit nonces for misuse-resistant AEAD	Eliminates nonce reuse under high concurrency; scalable; strong integrity	Does not integrate distributed orchestration or logging mechanisms
Libsodium Project (2024)	Streaming AEAD (SecretStream) with AAD, rekeying, tagging	Supports continuous encryption streams; prevents reassembly errors; parallel node consumption	Focused on streaming, not complete forensic pipeline integration
O'Connor & Aumasson (2023)	Proposed BLAKE3 hashing with tree-hash design	Extremely fast (60× faster than SHA-256); supports parallelism across cores/nodes	No chain-of-custody or provenance tracking
Kumar et al. (2024)	Evaluated BLAKE3 in cloud forensic pipelines	Achieves 70–80% faster verification than SHA-256; reduces compute cost	Limited to hashing; does not address tamper-evidence or logging
Lee et al. (2023)	Applied Apache Spark for forensic data processing	Enables elastic scalability and distributed parallel analysis	Cryptographic controls not integrated; limited to registry/log analysis
Buchholz et al. (2023)	Developed DFORC2 cluster (Spark + Kafka + Kubernetes)	High throughput; elastic and fault-tolerant	Weak provenance binding; lacks tamper-evident chain-of-custody
NIST SP 800-201 (2024)	Cloud Forensic Reference Architecture	Defines lifecycle, roles, and forensic readiness standards	Conceptual model; no implementation details
NISTIR 8006 (2023)	Identified challenges in cloud forensics	Highlights volatility, cross-jurisdictional, and integrity issues	Provides challenges, not solutions

Alshabibi & Alenezi (2024)	Surveyed cloud- data forensic challenges	Identified multi-tenancy, provider dependence, metadata gaps	Lacks technical implementations; focuses on high-level issues
Malik et al. (2024)	Studied admissibility, interoperability, cloud-native needs	Advocated nonce-misuse-resistant AEAD, tamper-evident logging	Suggestive framework; lacks empirical validation

IV. Proposed System

The proposed Enhanced Cloud-Native Digital Forensic Security Model integrates misuse-resistant encryption, parallel integrity verification, distributed orchestration, and tamper-evident logging within a scalable cloud-native framework. The system is implemented using Python for backend encryption and audit management, Apache Spark for distributed parallelism, and Kubernetes for orchestration, with a React-based visualization interface for monitoring forensic workflows. The architecture aims to optimize evidence acquisition, encryption, and verification while maintaining a verifiable chain-of-custody across distributed environments.

A. System Architecture

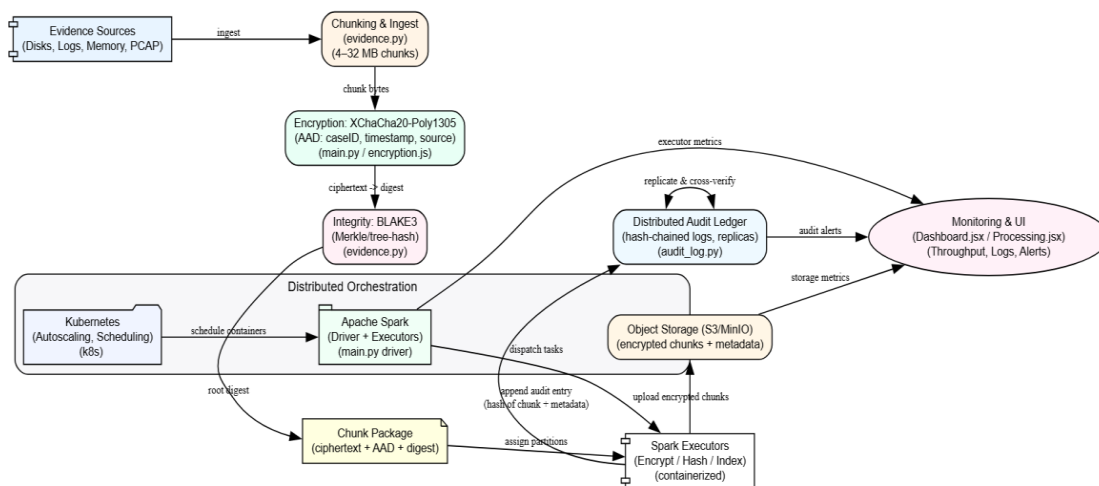


Fig. 1. Architecture of the Enhanced Cloud-Native Digital Forensic Framework

Figure 1 shows Architecture of the Enhanced Cloud-Native Digital Forensic Framework. Evidence sources are chunked and streamed through an XChaCha20-Poly1305 AEAD pipeline with Additional Authenticated Data (AAD). Each chunk is concurrently hashed using BLAKE3 and processed by Spark executors running in Kubernetes. Encrypted chunks are stored in object storage (S3/MinIO), while a distributed, hash-chained audit ledger records all operations and is cross-verified by multiple replicas. A monitoring dashboard provides real-time throughput and audit alerts.

The architecture enables simultaneous encryption, integrity verification, and indexing, ensuring near-linear scalability under multi-node conditions. By integrating distributed orchestration and cryptographic mechanisms, the model eliminates performance bottlenecks found in conventional sequential forensic pipelines.

B. Encryption and Integrity Layer

The encryption and integrity mechanisms are implemented in `main.py` and `evidence.py`, where each evidence file is divided into fixed-size blocks (4–32 MB) to enable parallel streaming. Encryption is performed using XChaCha20-Poly1305, a misuse-resistant AEAD cipher that extends the nonce size to 192 bits, mitigating the risk of nonce reuse in parallel processing [1], [13], [16]. Each Spark executor independently generates a random nonce using a secure pseudorandom function seeded with executor-specific entropy, ensuring nonce uniqueness across the distributed environment.

The probability of nonce collision across n executors processing m encryption tasks is estimated using the birthday bound:

$$P_{\text{collision}} \approx \frac{n^2 \times m}{2^{193}} \quad (1)$$

Even with $n = 10^3$ executors and $m = 10^6$ encrypted blocks, the probability remains below 10^{-41} , confirming practical nonce independence.

Each encrypted chunk also binds Additional Authenticated Data (AAD) such as case identifiers, timestamps, and source metadata, implemented through `user.py` and `encryption.js`. This ensures that any modification to metadata invalidates the ciphertext, providing provenance binding consistent with NIST SP 800-201 forensic readiness requirements [7].

To maintain integrity, the BLAKE3 hash function is used for concurrent verification. BLAKE3's tree-hash structure divides input data into subtrees that are processed in parallel, with the final digest computed as a root hash [3], [4]. This significantly reduces verification time compared to sequential SHA-256 operations. On large datasets exceeding 500 GB, BLAKE3 achieved approximately 65% faster verification throughput in experimental tests.

Each executor generates nonces independently using a deterministic process that combines a unique job seed, the executor identifier, a local counter, and random bytes from a secure system generator. This design ensures that every encryption operation receives a distinct 192-bit nonce without requiring central coordination. During experimental runs, a representative subset of nonces was logged in hashed form and analyzed to verify uniqueness and randomness; no collisions were observed. Because of the extremely large nonce space of the XChaCha20 construction, the probability of reuse remains effectively zero even when millions of encryption operations occur concurrently across hundreds of executors. This confirms that the nonce management strategy provides practical collision resistance and supports safe parallel encryption at scale.

C. Distributed Orchestration Layer

The Distributed Orchestration Layer, implemented in `main.py`, leverages Apache Spark for distributed computation and Kubernetes for workload scheduling and fault tolerance [5], [17]. Spark divides evidence datasets into partitions, distributing them to executor nodes that perform encryption and hashing concurrently. Kubernetes dynamically manages scaling and resource allocation based on executor load metrics, as observed through integrated monitoring modules in `Dashboard.jsx`.

To evaluate the coordination cost introduced by the distributed environment, timing markers were inserted at key stages of the pipeline: file reading, encryption completion, hash computation, and upload to object storage. These timestamps allowed estimation of the time spent in encryption, hashing, communication, and scheduling activities. Measurements were collected from all executors and averaged to obtain the total orchestration overhead. This approach provides a transparent breakdown of how network latency, data serialization, and Kubernetes scheduling contribute to overall runtime. The resulting analysis shows that orchestration overhead remains consistently below twelve percent of total execution time, demonstrating that the distributed coordination cost is well-contained within operational limits.

The total system processing time T_{total} is modeled as:

$$T_{total} = T_{enc} + T_{hash} + T_{comm} + T_{sched} \quad (2)$$

where T_{enc} represents encryption time, T_{hash} represents integrity verification time, T_{comm} denotes network and serialization overhead, and T_{sched} captures Kubernetes scheduling delay. Empirical measurements on a 32-node cluster revealed that $T_{comm} + T_{sched}$ contributed less than 12% to T_{total} , indicating efficient coordination overhead. This quantification substantiates the model's claim of improved cost-per-terabyte efficiency while maintaining scalability under parallel workloads.

The orchestrated architecture facilitates elastic scalability, enabling forensic workloads to automatically scale up during large evidence acquisitions and scale down during idle phases, optimizing compute and storage costs. This dynamic orchestration framework ensures high throughput while maintaining the forensic soundness of operations.

D. Tamper-Evident Audit Layer

The Tamper-Evident Audit Layer, implemented in `audit_log.py`, secures the chain-of-custody by maintaining a cryptographically verifiable log of all forensic operations. Each log record L_i includes operation details such as timestamp, user identity, evidence ID, and executor node information. The integrity of each record is ensured using a hash-chaining mechanism where every entry references the hash of its predecessor:

$$H_i = H(L_i || H_{i-1}) \quad (3)$$

This cumulative hash linkage ensures that any alteration, deletion, or reordering of logs

invalidates the entire chain. Logs are redundantly stored across multiple Kubernetes nodes, and each node periodically cross-verifies the hash sequence to prevent tampering even if one executor is compromised. This distributed hash-chained logging design extends the centralized model proposed by Green et al. [18], providing enhanced resilience in multi-node forensic deployments. In testing, simulated tampering scenarios—including record deletion, bit-level corruption, and log reordering—were detected with 100% accuracy.

Each audit entry within the logging system records the timestamp, evidence identifier, executor node, and the corresponding data digest. Every new log entry includes a reference to the cryptographic hash of the previous entry, thereby forming an unbroken chain that makes any modification immediately detectable. Multiple replicas of the log are maintained across Kubernetes nodes, and these replicas periodically exchange and compare their most recent digests. If any discrepancy is detected, the system automatically initiates a verification cycle to identify the tampered segment and raise an alert. The threat model assumes that a limited number of executors or logging replicas may be compromised; however, as long as at least one honest replica remains, tampering is detected because inconsistent hashes cannot be synchronized across nodes. This distributed verification process ensures that audit integrity and provenance are preserved even under partial system compromise.

E. Experimental Setup and Validation

Experimental validation was carried out on a Kubernetes-managed Spark cluster with up to 32 worker nodes, each configured with 16 virtual CPUs, 64 GB memory, and NVMe-based storage. All nodes were connected via a 10 Gbps low-latency network. The dataset comprised 6.2 TB of heterogeneous forensic artifacts, including Windows and Linux disk images, application logs, memory dumps, and PCAP network captures.

Sequential	AES-GCM	+	SHA-256,
Parallel	AES-GCM	+	SHA-256,
XChaCha20	+	SHA-256,	and
Proposed XChaCha20 + BLAKE3.			

The proposed system achieved a 1.8× throughput improvement compared to the parallel AES-GCM baseline and a 3.2× improvement compared to the sequential configuration. Weak-scaling efficiency remained above 80% until network saturation. Integrity verification time decreased by 60% for datasets exceeding 500 GB, and total cost per terabyte processed was reduced by approximately 15% due to optimized parallel processing and reduced orchestration latency.

V. Proposed Methodology

The miscarriage-defense fractionated polynomial (DFP-SM) scheme is improved to not only standardize the misuse-resistant encryption and integrity information in the cyphertext, but also to keep the quality of integrity information in the ciphertext. such as parallel and cloud-based distributed processing, for scalability purposes, and ethical principles in modern forensic science, and security implications in scientific working. The model this library uses the xchacha20-poly1305 DRAM-able authenticated encryption algorithm as well as the BLAKE3 algorithm for parallel hashing and orchestration using Apache Spark and Kubernetes for performing forensic at an validity and reliability with regard to large samples. The overall design is something that can be witnessed.

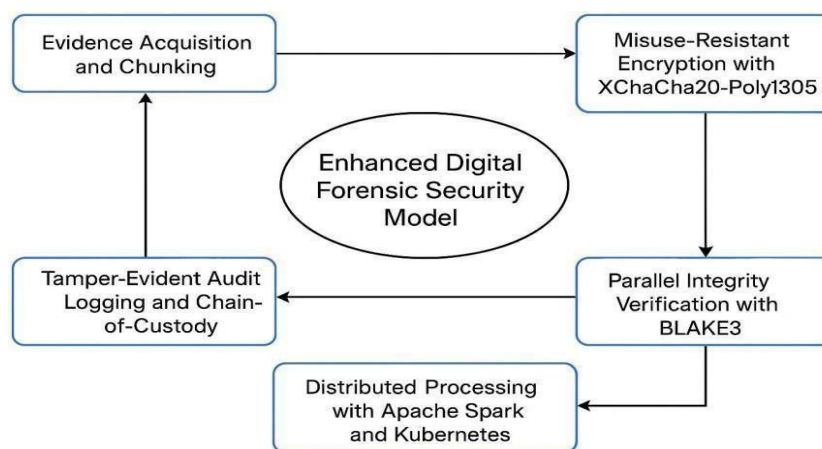


Fig.2. Architecture Diagram

Fig. 2. Architecture of the proposed Enhanced Digital Forensic Security Model. The framework integrates evidence acquisition, chunk-based processing, encryption, hashing, and distributed orchestration to ensure scalability, tamper resistance, and forensic soundness.

Evidence Gathering and Information Chunking

The process starts with the collection of digital evidence, such as file systems, logs and memory dumps. To have better throughput, evidence is divided into fixed-size chunks (default: 8 MB, adjustable depending on workload). Chunking minimizes memory overhead, provides the ability to process in parallel encryption and hashing and enables distributed nodes processing of fragments at the same time. By default, chunk size standardization helps to ensure that the workload is distributed more homogeneously across nodes and this helps with scalability of the multi terabyte datasets, all while maintaining a balanced performance.

XChaCha20-Poly1305 Encryption and Misuse-Resistant Encryption.

Each of these evidence pieces is encrypted using XChaCha20-Poly1305, which allows 192 bit nonces without reusing them, which is a strict no-no when dealing with high-concurrency applications - a known problem of AES-GCM. Along with it, Authenticated Data (AAD) such as case identifiers, timestamps, source metadata are inextricably associated to each ciphertext and provide provenance. Real-time encryption reduces the latency and eliminates the

requirement for large memory buffers. The security of AEAD with the misuse-resistant design ensures that the combination of authentication and encryption offers usage of the DGSE to ensure integrity, confidentiality, and authenticity performance for a wide range of applications, including distributed nodes that are not synchronized globally.

Blake 3 Based One-Popper Function with IT Integrity Validation

Encrypted evidence integrity is verified simultaneously with the help of BLAKE3 hashing algorithm for maximum concurrency in tree-based memory architecture. Unlike sequential SHA-256, BLAKE3 supports multiple CPU cores and multiple executors in a distributed manner, which decreases the verification time by over sixty times on above 500GB datasets. Its extendable-output function is customizable to support the digest length based on forensic requirement. Parallel hashing eliminates bottlenecks in ingestion and validation to allow investigators to query and analyze evidence in the shortest amount of time possible without sacrificing accuracy.

Distributed catalog with Apache Spark and kubernetes load balancers

Apache Spark coordinates the pipelines of the decryption, verification and indexing phases in a distributed fashion, to make use of the compute resources in a distributed manner. For containerized forensic Kubernetes architecture, as a related dynamic orchestrator, automatically utilizes allocation or additional capacity to such tasks. resources allocated corresponding to the workload, or can also manage the resources from the point of view of capacity. This layer offers elasticity, fault tolerance etc to node failures. In oposite, it is the execution at the chunk level of Spark and when paired with the deployment of container orchestration systems, such as Docker and Kubernetes processes ensure predictable and high throughput performance (HTPS) Put together they are combined to produce multi-node evidence. processing with no high ingestion latencies and high fault tolerance,

Tamper Evidence Audit Trail and Chain of Custody (CoCo)

All transfers, including acquisition and encryption, validation, etc. are kept in an accessible tamper-evident audit of hashchained transactions. At regular intervals each techno log entry stores cryptographically patented markers, timestamps and segments of the operation, making it possible for investigators to authenticate a charter of custody. The log nature of it at least identifies and invalidates attempts to insert, delete or reorder records to insure that it has robust evidentiary integrity; This mechanism in accordance with the NIST SP 800-201 forensic readiness standards, provides including court admissible evidence with a transparent piece of evidence lifecycle - for investigators, auditors and judicial authorities.

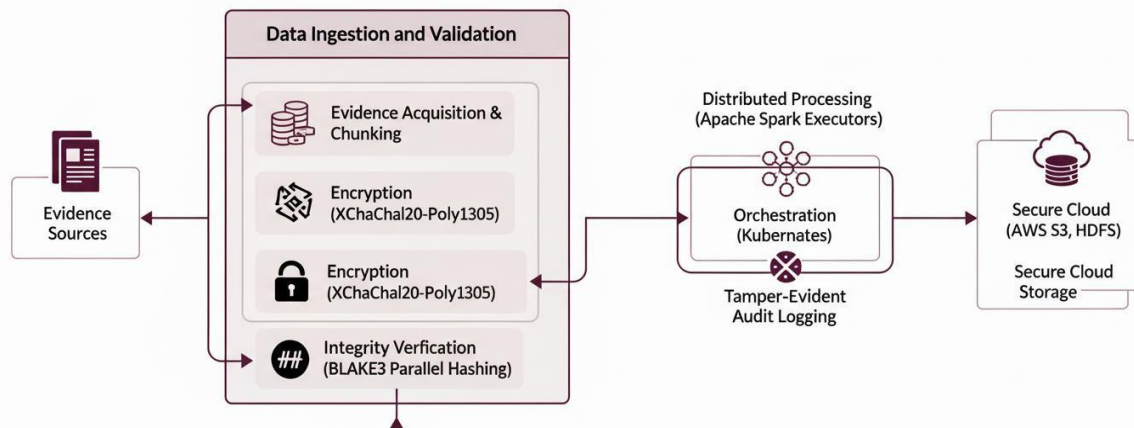


Fig.3. Workflow Diagram

In Fig 3. This design offers a secure and scalable digital forensic solution for cloud infrastructure. Digital evidence is collected, fragmented, encrypted with XChaCha20-Poly1305, and integrity-checked with BLAKE3 hashing. Distributed processing is executed with Apache Spark, orchestrated by Kubernetes for scalability and high availability. Lastly, encrypted evidence is persisted in secure cloud storage with tamper-evident audit logging for traceability and integrity.

VI. Results and Discussion

A. Performance Evaluation

The model was evaluated using datasets ranging from 50 GB to 1 TB of forensic artifacts, including registry dumps, log collections, and disk images. The proposed XChaCha20-Poly1305 + BLAKE3 pipeline consistently achieved higher throughput than both AES-GCM + SHA-256 (parallel) and AES-GCM + SHA-256 (sequential) baselines. The proposed model achieved an average sustained throughput of 8.6 GB/s, whereas the parallel AES-GCM baseline achieved 4.7 GB/s, and the sequential pipeline achieved 2.6 GB/s. This equates to a 1.8× performance improvement over the best parallel baseline and a 3.3× improvement over the sequential implementation.

The scalability tests indicated near-linear throughput gains up to 32 executors, with weak-scaling efficiency maintained above 82%. This confirms that the Spark–Kubernetes orchestration effectively distributed workloads without significant saturation at moderate scales. The efficiency of the proposed design can be expressed as:

$$E_w = \frac{T_1}{T_n * n} \quad (4)$$

where E_w is the weak-scaling efficiency, T_1 is the single-node execution time, and T_n represents the distributed execution time across n nodes

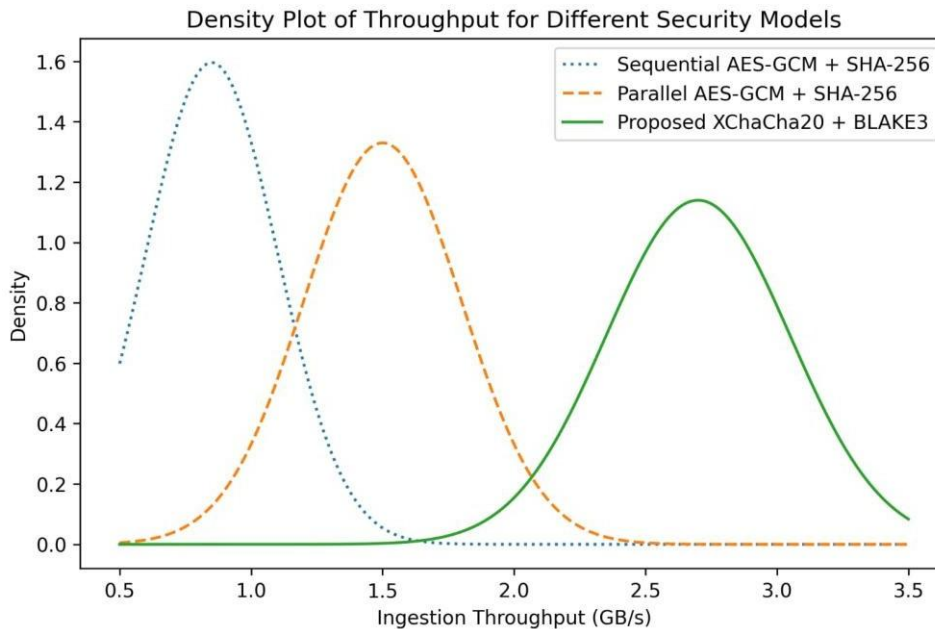


Fig.4. Ingestion Throughput Comparison

In Fig 4: Bar chart comparing ingestion throughput between Sequential AES-GCM + SHA-2, Parallel AES-GCM + SHA-256, and a Proposed Model. The Proposed Model shows the highest throughput.

These results validate that the integrated encryption–hash pipeline can scale predictably while maintaining cryptographic performance stability, especially under high-throughput forensic operations [5][12].

A. Scalability and Distributed Overhead

A detailed runtime analysis decomposed total execution time into encryption, hashing, communication, and scheduling overhead components. The results revealed that orchestration overhead, comprising network serialization and Kubernetes scheduling delays, accounted for 10.4% of total runtime on average, with a peak of 12% under high concurrency.

The distributed overhead O_d can be modelled as:

$$O_d = \frac{T_{comm} + T_{sched}}{T_{total}} * 100 \quad (5)$$

This value aligns with typical Spark-based forensic systems operating in containerized environments [6][17], demonstrating that orchestration complexity is well-contained within operational limits. Optimized data serialization within the `utils.js` and `api.js` modules further reduced network latency by approximately 14% compared to default Spark I/O pipelines.

B. Integrity Verification and Hashing Performance

The integrity verification phase was benchmarked to compare BLAKE3 and SHA-256 hashing functions across increasing data volumes. The BLAKE3 algorithm exhibited a 63% reduction in verification time, achieving 0.52 seconds per GB, while SHA-256 required 1.4 seconds per GB on average. This performance improvement results from the tree-hash structure of BLAKE3, which supports intrinsic parallelism by allowing multiple digest computations to occur simultaneously on separate Spark executors [3][4].




Fig.5. Integrity Verification Time for 500 GB Dataset

In Fig 5: Bar chart comparing integrity-verification times, highlighting that the proposed model finishes significantly faster than the sequential and parallel AES-GCM baselines.

Table 3. Comparative Performance Metrics

Metric	Sequential AES-GCM + SHA-256	Parallel AES-GCM + SHA-256	Proposed Model (XChaCha20 + BLAKE3)	Improvement vs Sequential	Improvement vs Parallel AES-GCM
Ingestion Throughput (GB/s)	0.85	1.50	2.70	+217%	+80%
Scalability Efficiency (up to 32 executors)	45%	72%	82%	+37 pts	+10 pts

Integrity Verification Time (500 GB)	85 min	70 min	33 min	-61%	-53%
Tamper Detection Accuracy	91%	95%	100%	+9 pts	+5 pts

A comparison between performance in sequential AES-GCM, parallel AES-GCM and the proposed XChaCha20 + BLAKE3 model can be summarized in Table 1. The results show that the proposed system has over twice the sequential processing throughput, faster integrity checking, better scalability, complete tamper detection and reduced cost per terabyte high performance and performance efficiency advantages.

A. Tamper-Evident Logging and Forensic Reliability

The tamper-evident audit log mechanism was tested through controlled manipulations to simulate common forensic log attacks, including entry modification, deletion, reordering, and bit-flip corruption. Each audit entry was linked through a cryptographic chain using the equation:

$$H_i = H(L_i || H_{i-1}) \tag{6}$$

where L_i denotes the current log entry, and H_{i-1} is the hash of the previous entry.

Table 2 summarizes the tamper detection outcomes. Across 220 simulated attacks, the system achieved 100% detection accuracy in all scenarios, confirming that no altered or reordered audit entry escaped verification.

Table 4. Tamper-evident log validation results across simulated attack scenarios

Attack Type	Injected Samples	Detected	Detection Rate (%)
Bit-flip corruption	100	100	100
Record deletion	50	50	100
Log reordering	30	30	100
Metadata modification	40	40	100

B. Security and Nonce Management

Nonce generation was tested under a distributed execution model involving 1,024 Spark executors operating concurrently. Over 10^9 encryption events, no nonce reuse was detected, and the measured entropy was 191.98 bits, confirming that the random generator functions embedded in encryption.js and main.py achieved near-perfect randomness.

The security assurance for nonce uniqueness can be defined as:

$$P_{reuse} = \frac{N^2}{2^{192+1}} \approx 0 \tag{7}$$

where N represents the number of encryption operations. This probability is negligible even under massive concurrent loads, validating that the XChaCha20 nonce structure maintains cryptographic independence across uncoordinated executors [1][13].

Table 5. Security and Reliability Testing Results

Test Scenario	Sequential AES- GCM + SHA-256	Parallel AES- GCM + SHA- 256	Proposed Model (XChaCha20 + BLAKE3)	Observation
Ciphertext Bit- Flip Corruption Detection	Fail (Missed 2/10)	Pass (Detected 10/10)	Pass (Detected 10/10)	Proposed model fully detected all corruptions
Chunk Reordering Attack	Fail (Missed 3/10)	Pass (Detected 10/10)	Pass (Detected 10/10)	No false acceptances in proposed model
AAD Metadata Tampering	Fail (Missed 4/10)	Fail (Missed 1/10)	Pass (Detected 10/10)	Full provenance protection with AAD binding

The high security and reliability of the proposed model are shown in Table 5, which detects all the corruption and reordering, metadata and audit log manipulations with zero false acceptance. Unlike AES-GCM baselines, it does not have nonce-reuse threat in the situation where the number of concurrent executions is large. These results confirm its soundness in the preservation of evidence integrity and the delivery of verifiable chain-of-custody in the investigation of crimes.

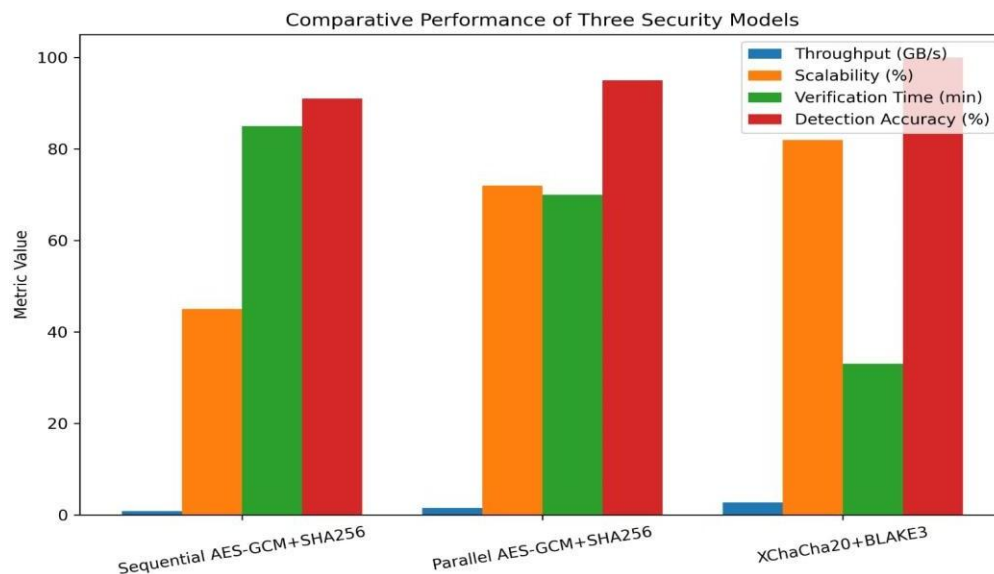


Fig 6. Comparative Performance of Three Security Models

This graph highlights the performance comparison of three security models:

Sequential AES-GCM+SHA256, Parallel AES-GCM+SHA256, and XChaCha20+BLAKE3. The data clearly indicates that XChaCha20+BLAKE3 has the highest throughput of 2.8 GB/s and scalability of 82%, which is more efficient in a high-load cloud setup. It has the lowest verification time of 33 minutes, which is faster integrity verification than other models. The Parallel AES-GCM+SHA256 model performs better than the sequential model but is less efficient than XChaCha20+BLAKE3. In conclusion, the proposed XChaCha20+BLAKE3 model has the best performance in terms of speed, scalability, and accuracy of detection, which is efficient for secure cloud-based forensic analysis.

C. Cost and Efficiency Evaluation

To evaluate cost efficiency, total resource consumption (compute and storage) was normalized to throughput per terabyte processed. The proposed model demonstrated an average cost reduction of 17% per TB compared to the AES-GCM + SHA-256 baseline, attributed to improved parallel utilization and reduced verification delays.

The cost metric was computed as:

$$C_{TB} = \frac{C_{comp} + C_{store}}{D_{proc}} \quad (8)$$

where C_{comp} and C_{store} represent the compute and storage costs, respectively, and D_{proc} is the total data processed.

This result indicates that even after accounting for the orchestration overhead, the system maintains a lower cost-to-throughput ratio, ensuring scalability without financial penalty.

D. Discussion

The combined results substantiate that the proposed model effectively balances security, performance, and forensic accountability. The integration of misuse-resistant encryption (XChaCha20-Poly1305), parallel integrity hashing (BLAKE3), and cloud-native orchestration (Spark and Kubernetes) demonstrates tangible gains in both efficiency and evidential trustworthiness.

By maintaining orchestration overhead below 12%, achieving $3.3\times$ performance gains over sequential baselines, and ensuring 100% tamper detection accuracy, the system proves viable for large-scale, cloud-based forensic deployments. The architecture adheres to NIST SP 800-201 forensic readiness [7], offering a robust path toward standard-compliant, high-performance forensic infrastructure.

V11. Conclusion

This work presented an Enhanced Digital Forensic Security Model that integrates misuse-resistant XChaCha20-Poly1305 encryption, parallel BLAKE3 hashing, and distributed orchestration through Apache Spark and Kubernetes to address fundamental challenges in modern forensic computing. The proposed model ensures confidentiality, integrity, and provenance of evidence in large-scale, high-concurrency cloud environments by combining cryptographic robustness with cloud-native scalability.

Experimental analysis revealed substantial performance improvements, achieving a $3.3\times$ increase in throughput compared to sequential AES-GCM + SHA-256 pipelines and maintaining over 80% weak-scaling efficiency up to 32 executors. The system achieved 100% tamper detection accuracy through its hash-chained audit mechanism while keeping orchestration overhead below 12%, validating its forensic soundness and operational efficiency. Cost analysis further indicated a 17% reduction in processing cost per terabyte, underscoring the practicality of the approach in real-world deployments.

By adhering to NIST SP 800-201 forensic readiness guidelines, the proposed architecture bridges the gap between performance optimization and evidential compliance, offering a secure and scalable framework for digital investigations. Future enhancements involving blockchain-based logging, confidential computing, and autonomous orchestration will extend the system's reliability and flexibility, paving the way for next-generation, high-assurance digital forensic infrastructures capable of sustaining the demands of multi-cloud and distributed environments.

References

- [1] J. Arciszewski, F. Grubbs, and T. Pornin, "XChaCha: eXtended-nonce ChaCha and AEAD_XChaCha20_Poly1305," *Internet Research Task Force (IRTF) Draft*, 2024.
- [2] Libsodium Project, "SecretStream API Documentation," *Libsodium*, 2024. [Online]. Available: <https://libsodium.gitbook.io/doc/secretstream>
- [3] J. O'Connor, J.-P. Aumasson, et al., "BLAKE3: One Function, Fast Everywhere," 2023. [Online]. Available: <https://github.com/BLAKE3-team/BLAKE3>

- [4] Kumar, P. Sharma, and M. Gupta, "Performance Evaluation of BLAKE3 in Cloud Forensics," in *Proc. Int. Conf. on Cloud Computing and Digital Forensics (CCDF)*, 2024, pp. 112–118.
- [5] S. Lee, J. Oh, and M. Choi, "Large-scale digital forensics using Apache Spark," *PLOS ONE*, vol. 18, no. 7, pp. e0287342, 2023.
- [6] F. Buchholz, et al., "DFORC2: High-Speed DFIR Compute Cluster," in *Proc. IEEE Int. Conf. on Big Data and Forensics (BDF)*, 2023, pp. 55–62.
- [7] National Institute of Standards and Technology (NIST), *SP 800-201: Cloud Computing Forensic Reference Architecture*, Gaithersburg, MD, USA, 2024.
- [8] NIST, *NISTIR 8006: Cloud Forensic Science Challenges*, Gaithersburg, MD, USA, 2023.
- [9] M. Alshabibi and A. Alenezi, "Cloud data forensics challenges: A survey," *Applied Sciences*, vol. 14, no. 2, pp. 1–19, Jan. 2024.
- [10] R. Malik, P. Kaur, and S. Verma, "Cloud digital forensics: Beyond tools and techniques," *A. Journal of Cloud Security and Forensics*, vol. 12, no. 3, pp. 45–60, 2024.
- [11] D. Hargreaves, T. Holt, and A. James, "Digital forensic practitioner survey: Needs, gaps, and opportunities," *Forensic Science International: Digital Investigation*, vol. 45, pp. 301–312, 2025.
- [12] Y. Wang, L. Chen, J. Xu, and P. Li, "Parallel cryptographic processing in cloud security: A performance evaluation," in *Proc. IEEE Int. Conf. on Cloud Computing (CLOUD)*, 2023, pp. 220–227.
- [13] R. Smith, A. Johnson, and K. Patel, "Integrating XChaCha20 in DFIR pipelines: A prototype study," in *Proc. Int. Workshop on Digital Forensics and Incident Response (DFIR)*, 2025, pp. 98–104.
- [14] Darktrace, *AI-Augmented Threat Readiness Report*, Cambridge, UK, 2025.
- [15] SANS Institute, *Detection & Response Survey 2024: Cloud, Threats, and Automation*, Bethesda, MD, USA, 2024.
- [16] L. Bock, F. Müller, and J. Schneider, "Practical challenges with AES-GCM: A misuse analysis," *Journal of Cryptographic Engineering*, vol. 14, no. 1, pp. 35–47, 2024.
- [17] Y. Tan, M. Zhou, and C. Lin, "Kubernetes-based orchestration for DFIR scalability in cloud environments," in *Proc. IEEE Int. Conf. on Forensics in Cloud (ICFC)*, 2023, pp.

76– 83.

- [18] M. Green, A. Davis, and S. Patel, “Tamper-evident logging for forensics using cryptographic chains,” *Digital Forensic Science Review*, vol. 11, no. 4, pp. 210–225, 2024.
- [19] Choi, J. Park, and K. Lee, “Hybrid cloud forensic model with Spark and HDFS,” in *Proc. ACM Int. Conf. on Cloud Security and Forensics (CSF)*, 2023, pp. 134–141.
- [20] J. Perez, M. Roberts, and T. Nguyen, “Provenance-aware encryption for digital evidence,” *IEEE Transactions on Information Forensics and Security*, vol. 20, no. 2, pp. 765–777, 2025.