Journal Homepage: -www.journalijar.com

INTERNATIONAL JOURNAL OF ADVANCED RESEARCH (IJAR)

Article DOI:10.21474/IJAR01/23457
DOI URL: http://dx.doi.org/10.21474/IJAR01/23457

RESEARCH ARTICLE

DIGITAL EVIDENCE ACQUISITION FOLLOWING MINING AND NETWORK ATTACKS IN A PRIVATE BLOCKCHAIN

Ali Hamadou[1], René Ndoundam[2], Abdoul Aziz Issaka Hassane[1] and Mbopgue Timo Patrick Noel[2]

1. Department of Mathematics, Dan Dicko Dankoulodo University of Maradi, Maradi, Niger.
2. Department of Computer Science, University of Yaoundé I, Yaoundé, Cameroon.

Manuscript Info

Manuscript History
Received: 10 March 2026
Final Accepted: 12 April 2026
Published: May 2026

Key words:-
Blockchain; Private blockchain; Mining attack; Network attack; Digital forensics; Hyperledger Fabric; Evidence acquisition; Chain of custody.

Abstract

The rapid adoption of blockchain technology has introduced new cybersecurity challenges, particularly in private blockchain environments where mining and network attacks pose significant threats. This study investigates the digital forensics process within private blockchain environments, with a specific focus on methodologies for acquiring digital evidence following mining and network attacks. In this context, we propose two complementary evidence acquisition models applied to Hyperledger Fabric: a post-mortem model using forensic disk imaging and hash verification (dd, sha256sum), and a live acquisition model using Hyperledger Explorer's REST API. A Blockchain Denial-of-Service (BDoS) flooding attack simulation was conducted to validate the proposed models, with quantitative acquisition metrics reported (imaging time, storage overhead, API latency, CPU impact). The proposed models successfully enabled the identification and collection of user generated and machine generated evidence, including cryptographic transactions,node event logs, smart contract interactions, and network traffic metadata, while preserving chain of custody integrity. This work contributes a structured,practical methodology for blockchain forensics in private environments, addressing the unique challenges posed by the decentralized and immutable nature of blockchain systems.

Introduction:-

Blockchain technology has emerged as a transformative paradigm for secure, decentralized data exchange, with adoption spanning finance, healthcare, logistics, supply chain management, and digital identity [1]. The global blockchain market, estimated at USD 17.57 billion in 2023, is projected to reach USD 469.49 billion by 2030 at a compound annual growth rate of 59.9% [13]. This rapid proliferation—particularly of private, permissioned deployments such as Hyperledger Fabric—brings with it an expanding cybersecurity threat landscape. Private blockchains, which operate with a restricted set of known actors and controlled nodes, present internal governance vulnerabilities that differ substantially from those of public chains: their smaller network size amplifies the economic and operational impact of attacks, while their permissioned architecture offers investigators a more tractable surface for forensic intervention.

Corresponding Author:-Ali Hamadou
Address:-Department of Mathematics, Dan Dicko Dankoulodo University of Maradi, Maradi, Niger.

Attacks on blockchain systems fall into two principal categories: mining attacks, which compromise the computational integrity of the chain (e.g., 51% attacks, selfish mining, cryptojacking, timejacking), and network attacks, which disrupt inter-node communication (e.g., eclipse attacks, Sybil attacks, DNS poisoning, Blockchain Denial-of-Service—BDoS) [8]. Recent studies confirm the acuteness of these threats: nascent private networks can be compromised by 51% attacks at costs several orders of magnitude lower than established public chains [16], and real-world BDoS incidents—such as the Solana network outage of January 2022—demonstrate the disruptive potential of transaction flooding [2]. Eclipse attacks have furthermore been shown to serve as enablers for cascading attacks including selfish mining and double spending [18].

Digital forensics—the science of identifying, preserving, analyzing, and presenting digital evidence within a legal framework [5]—is well-established for conventional computing environments. Its application to blockchain, however, remains nascent. A recent systematic review [12] confirms that while blockchain's immutability protects evidence from tampering, it simultaneously prevents investigators from correcting erroneous records, introduces attribution challenges through pseudonymity, and complicates evidence localization across distributed nodes. Another recent literature survey [15] further establishes that the most explored domains are IoT and cloud forensics, while forensic investigation of the blockchain itself—especially private permissioned chains—remains critically understudied.

The existing body of work on blockchain forensics can be partitioned into two streams. The first uses blockchain as forensic infrastructure: ForensicTransMonitor [13] embeds each investigative step as an immutable ledger entry via smart contract APIs; ZAKON [23] deploys a Hyperledger Fabric-based admissibility framework for courtroom-ready evidence; and multiple chain-of-custody systems [14] have been validated on permissioned platforms. The second stream targets blockchain systems as objects of investigation: Balaskas and Franqueira [7] catalogued analytical tools (Chainalysis, Elliptic, Crystal Blockchain), while deep anomaly detection approaches [19]] have been proposed for real-time attack monitoring on public chains. Neither stream, however, addresses the specific forensic challenge of acquiring evidence following mining and network attacks within a private blockchain environment—where the investigator operates inside a closed, permissioned network with access to node infrastructure, event logs, and Docker container internals.

In this paper, we propose and empirically validate two complementary evidence acquisition models—one post-mortem, one live—specifically designed for Hyperledger Fabric private blockchain environments subjected to mining and network attacks.

The principal contributions of this work are:

- A blockchain-specific forensic evidence taxonomy distinguishing user-generated and machine-generated artefacts in private permissioned environments;
- A post-mortem acquisition model based on bit-for-bit forensic disk imaging with SHA-256 integrity verification, compliant with NIST chain-of-custody standards;
- A live acquisition model leveraging Hyperledger Explorer's REST API for real-time, targeted evidence retrieval from a running Fabric network;
- A structured tool selection framework mapping evidence type and acquisition mode to appropriate tooling within the Hyperledger Fabric ecosystem;
- Empirical validation through controlled BDoS flooding attack simulation on a Hyperledger Fabric test network, with measurable quantitative acquisition metrics (imaging time, storage overhead, API latency, and CPU impact).

The remainder of this paper is structured as follows: Section 2 reviews related work across digital forensics methodologies, blockchain attack taxonomies, and forensic frameworks. Section 3 presents background on digital investigation and blockchain attack classification. Section 4 details the proposed acquisition models and their Hyperledger Fabric implementation. Section 5 presents and discusses results, and Section 6 concludes with research perspectives.

Related Work:-

This section reviews the literature across three intersecting areas: digital forensics methodologies, blockchain security and attack taxonomies, and blockchain-assisted or blockchain-targeted forensic frameworks. We organize findings chronologically within each theme and identify the gap that the present work addresses.

Digital Forensics Methodologies and Evidence Acquisition:-

Digital forensics has historically focused on conventional computing environments, with established methodologies covering file system analysis, memory forensics, network forensics, and mobile device investigation [4]. The first formal definition emerged from the Digital Forensics Research Workshop (DFRWS) in 2001, describing it as the application of scientifically proven methods for the preservation, collection, validation, identification, analysis, interpretation, documentation, and presentation of digital evidence [5]. Since then, the discipline has expanded considerably alongside technological evolution.

Casino et al. [6] provided a comprehensive systematic review of research trends in digital forensics up to 2022, identifying blockchain as one of the most rapidly emerging investigation domains. Their work catalogued open challenges including data volume, encryption, and the fragmentation of evidence across distributed systems. Atlam et al. [12] conducted a more recent systematic literature review of blockchain forensics specifically, surveying state-of-the-art techniques, applications, and open challenges. They highlighted that while immutability protects evidence from tampering, it also prevents investigators from correcting erroneous records—a double-edged property critical for forensic methodology design.

More recently, the application of blockchain technology as a forensic infrastructure tool—rather than a target of investigation—has attracted significant research attention. ForensicTransMonitor by Syed et al. [13] proposed a generic methodology embedding each forensic transaction as an immutable blockchain entry, using smart contracts as API connectors between forensic applications and the ledger. The framework demonstrated applicability across IoT and cloud domains with low computational overhead. Complementing these infrastructure-oriented approaches, recent research has explored AI-driven techniques for automated forensic evidence identification. Hasan et al. [25] proposed a machine learning framework integrating explainable AI techniques for anomaly detection in blockchain transactions, demonstrating that ensemble classifiers (XGBoost, stacking) outperform single classifiers with high accuracy in detecting anomalous Bitcoin transactions. Similarly, Mohammed and Ngom [26] proposed a forensic-friendly AI framework integrating blockchain with federated learning to enhance AI trustworthiness, enabling traceable and auditable model training across distributed nodes without centralising raw data. These AI-augmented forensic approaches represent a complementary direction to the rule-based acquisition models proposed in the present work, and are discussed further in Section 5. The work in [15] systematically compiled blockchain applications in digital forensics up to early 2025, confirming that IoT forensics and cloud forensics are the most explored domains, while forensic investigation of the blockchain itself—particularly private chains—remains understudied.

Blockchain Attack Taxonomy and Security Analysis:-

The security landscape of blockchain systems has been extensively studied. Saad et al. [8] presented a comprehensive survey of the blockchain attack surface, systematically classifying attacks across consensus mechanisms, peer-to-peer networks, and application layers. König et al. [10] reviewed current vulnerabilities and attack vectors, with particular emphasis on PoW and PoS consensus weaknesses and PBFT fault tolerance limits.

Regarding mining attacks, a recent review on 51% attack vulnerability of nascent blockchains [16] characterized the security economics of consensus attacks across the blockchain lifecycle, finding that small private networks with fewer nodes can be compromised at costs several orders of magnitude lower than established public chains—making the private blockchain context particularly sensitive. Selfish mining detection models based on attack state-intensity-time relationships have also been proposed, providing a quantitative basis for forensic timeline analysis [17].

On network attacks, Shi et al. [18] analyzed eclipse attacks on Ethereum's peer-to-peer network, demonstrating how node isolation facilitates cascading attacks including selfish mining and double spending. A smart eclipse attack detector proposed behavioral monitoring of peer connection patterns as an early warning mechanism. For denial-of-service attacks at the blockchain level (BDoS), the literature documents real-world incidents including the Solana network outage, where transaction flooding caused a four-hour service interruption [2]. Deep anomaly detection systems leveraging machine learning for real-time blockchain attack detection—covering 51%, selfish mining, double-spending, and Sybil attacks—have been surveyed[19], establishing a computational basis for proactive forensic monitoring.

Blockchain-Targeted and Blockchain-Enabled Forensic Frameworks:-
Several recent works have addressed the integration of blockchain technology into forensic investigation pipelines, primarily for IoT and cloud environments. Xiao et al. [20] proposed a blockchain-based digital forensics scheme for Industrial IoT (IIoT), using decentralized storage for forensic data and smart contracts for evidence chain tracing, with a token-based access control mechanism. Brotsis et al. [21] developed BEvPF-IoT, a blockchain-based evidence preservation framework for IoT devices designed to prevent third-party manipulation of digital evidence until court submission.

In the Hyperledger Fabric context specifically, the authors demonstrated the use of Hyperledger Fabric for maintaining evidence integrity in containerized cloud ecosystems, integrating Docker engine audit logging with blockchain-based accountability [22]. The ZAKON framework [23] introduced a decentralized architecture for forensic evidence admissibility, deployed on Hyperledger Fabric and benchmarked using Hyperledger Caliper, addressing courtroom admissibility through multi-dimensional checking of evidence transactions and post-trial query resolution. Ali et al. [24] leveraged Hyperledger Fabric for trusted cybersecurity threat intelligence sharing using IPFS and MITRE ATT&CK-structured smart contracts, demonstrating the platform's suitability for security-critical permissioned applications.

Concerning the forensic investigation of blockchain systems themselves—as opposed to using blockchain as forensic infrastructure—Balaskas and Franqueira [7] catalogued analytical tools for blockchain investigation, distinguishing between collection tools (Bitcoin Core, Ethereum ETL), transaction analysis platforms (Chainalysis, Elliptic), and behavioral analysis tools (Crystal Blockchain, CipherTrace). However, these tools target public chains. The ZAKON framework and the Hyperledger-based chain-of-custody systems cited above address the use of private chains as forensic enablers, not as forensic targets.

The review above reveals a clear and persistent gap in the literature: while substantial work exists on (a) forensic methodologies for conventional systems, (b) blockchain attack taxonomies, (c) blockchain-assisted forensic chain-of-custody systems, and (d) forensic tools for public blockchains, no prior work has proposed a comprehensive, empirically validated evidence acquisition methodology specifically designed for forensic investigation of a private blockchain system following mining and network attacks. The present work addresses this gap by designing, implementing, and validating two complementary evidence acquisition models—post-mortem and live—within a Hyperledger Fabric deployment subjected to controlled attack scenarios.

Background:-
Digital Forensics
Digital forensics encompasses the application of scientifically proven methods for the systematic investigation of digital systems. The standard investigation process comprises five phases: Identification, Collection (Acquisition), Examination, Analysis, and Presentation—all governed by an overarching chain of custody requirement [3, 4].
The acquisition phase is of particular importance, as it determines the evidentiary value of all subsequent analysis.

Two main acquisition modes exist:
- Post-mortem (cold) acquisition: The target system is powered down before imaging. Bit-for-bit copies of storage media are created in a controlled environment, ensuring data authenticity and reproducibility.
- Live (warm) system acquisition: The target system remains operational, enabling capture of both static and volatile data including active processes, network connections, memory contents, and event logs.

Chain of custody (chain of possession) is a fundamental principle ensuring that evidence integrity is maintained from acquisition to courtroom presentation. Best practices include rigorous documentation, cryptographic hashing for integrity verification, secure storage, and use of industry-standard tools such as EnCase or FTK [3]. Digital evidence can be classified into two categories: user-generated data (documents, messages, account details, web pages) and machine/network-generated data (system logs, router logs, IP addresses, configuration files, temporary files). In the context of blockchain forensics, this taxonomy requires extension to accommodate blockchain-specific artefacts.

Blockchain Technology and Attack Taxonomy:-

A blockchain is a decentralized, distributed, and immutable ledger where data is organized into cryptographically linked blocks. Each block contains a block hash, the previous block's hash, a Merkle root, a timestamp, a nonce, and transaction data [1, 11]. Blockchain systems are classified into four types: public (permissionless), private (permissioned), consortium, and hybrid. This study focuses on private blockchains, which are controlled by a single organization and restrict participation to authorized nodes.

As shown in Table 1, blockchain attacks can be organized into three principal categories [8]:

Table 1. Classification of blockchain attacks by type.

Attack Category	Attack Type	Primary Impact
Mining Attacks	51% Attack	Consensus integrity compromise
	Selfish Mining	Block withholding, unfair rewards
	Cryptojacking	Unauthorized compute resource use
	Timejacking	Block timestamp manipulation
Network Attacks	Eclipse Attack	Node isolation, information manipulation
	Sybil Attack	Identity spoofing, vote manipulation
	DNS Attack	Routing hijack to counterfeit network
	BDoS Attack	Transaction throughput denial
Application Attacks	Smart Contract DoS	Fund lock, auction manipulation
	Re-entrancy Attack	Recursive fund withdrawal
	Replay Attack	Transaction replay across chains

Mining attacks primarily affect the computational power balance of the network. In PoW-based systems, the 51% attack requires an adversary to control over half the total hash rate, enabling double-spending, block suppression, and chain forking. In private PBFT-based chains (as used in Hyperledger Fabric), the equivalent compromise requires controlling the primary node, with a failure tolerance threshold of only 33% of nodes [8, 10].

Network attacks exploit the peer-to-peer communication substrate. Eclipse attacks isolate target nodes by monopolizing their inbound and outbound connections with malicious peers, enabling information poisoning. Sybil attacks create multiple false identities to overwhelm honest node votes. BDoS attacks flood the transaction pool with illegitimate transactions, degrading throughput—as demonstrated in the Solana network [2].

Blockchain Forensics: Specific Challenges:-

Forensic investigation of blockchain environments is complicated by several inherent characteristics:

- Data immutability: Once recorded, blockchain data cannot be modified or deleted, preserving evidence integrity but complicating error correction ;
- Pseudonymity: Blockchain addresses are not directly linked to real-world identities, complicating attribution ;
- Data volume and distribution: Evidence is distributed across all network nodes, requiring coordinated multi-node acquisition ;
- Technical complexity: Investigators must understand consensus mechanisms, smart contract execution, and cryptographic data structures ;
- Evidence encryption: Sensitive blockchain data is cryptographically protected, limiting direct inspection without appropriate credentials.

Methodology:-

Experimental Setup:-

The experimental environment consisted of a workstation running Kali Linux with 16 GB RAM and 500 GB storage, supplemented by a 500 GB external drive for forensic image storage. The private blockchain platform was

Hyperledger Fabric v2.5, selected for its enterprise-grade modularity, PBFT-compatible consensus (Raft), permissioned access model, and open-source availability.

The test network was deployed using the fabric-samples reference architecture, comprising two peer organizations (Org1, Org2), one node per organization (peer0.org1.example.com, peer0.org2.example.com), and one ordering node (orderer.example.com). Communication channels were established using the createChannel script, and a basic asset-transfer chaincode was deployed for transaction simulation.

System prerequisites included Git, cURL, Docker, Go (v1.19.3), and the Hyperledger Fabric binaries. Network deployment followed the standard fabric-samples installation procedure via the install-fabric.sh script.

Evidence Taxonomy in Private Blockchain:-

Based on the standard digital forensics evidence classification framework [4] and the specific characteristics of Hyperledger Fabric, we define the following evidence taxonomy applicable to post-attack investigation(Table 2):

Table 2. Digital evidence taxonomy in private blockchain environments.

Evidence Category	Evidence Type	Forensic Relevance
User-generated	Cryptographic transactions & blocks	Proof of asset transfers, double-spend detection
	Smart contract interactions	User behavior, unauthorized chaincode calls
	Digital signatures	Attribution, identity verification
	DApp-level data	Application-layer activity reconstruction
Machine-generated	Node event logs (peers, orderer)	Attack timeline, anomaly detection
	Network traffic metadata	Eclipse/Sybil/BDoS pattern identification
	Transaction endorsement records	Consensus flow analysis
	CouchDB state database entries	Ledger state at time of attack
	Docker container logs	Infrastructure-level error traces

Proposed Evidence Acquisition Models:-

General Acquisition Workflow:-

The proposed acquisition workflow (Figure 1) is applicable to both post-mortem and live acquisition contexts. It proceeds through the following stages:

- Phase 1 – Identification: Secure the crime scene; determine the incident type (mining or network attack); answer the 5W1H investigative questions (Who, What, Where, When, How, Why); identify evidence types (user vs. machine) ;
- Phase 2 – Evidence Type Determination: Classify target evidence as user-generated or machine-generated (Table 2), and determine whether post-mortem or live acquisition is appropriate ;
- Phase 3 – Tool Selection and Blockchain Access: Select appropriate tooling based on evidence type and acquisition mode; authenticate to the Hyperledger Fabric network ;
- Phase 4 – Collection, Hashing, and Copying: Collect all relevant data; generate SHA-256 hash of collected data; create a working copy; verify integrity of the copy ;
- Phase 5 – Transmission and Storage: Transmit one copy to investigators for analysis; securely store another copy for future reference throughout the investigation ;
- Phase 6 – Chain of Custody Documentation: Document all steps following the NIST Evidence Chain-of-Custody Tracking Form standard.

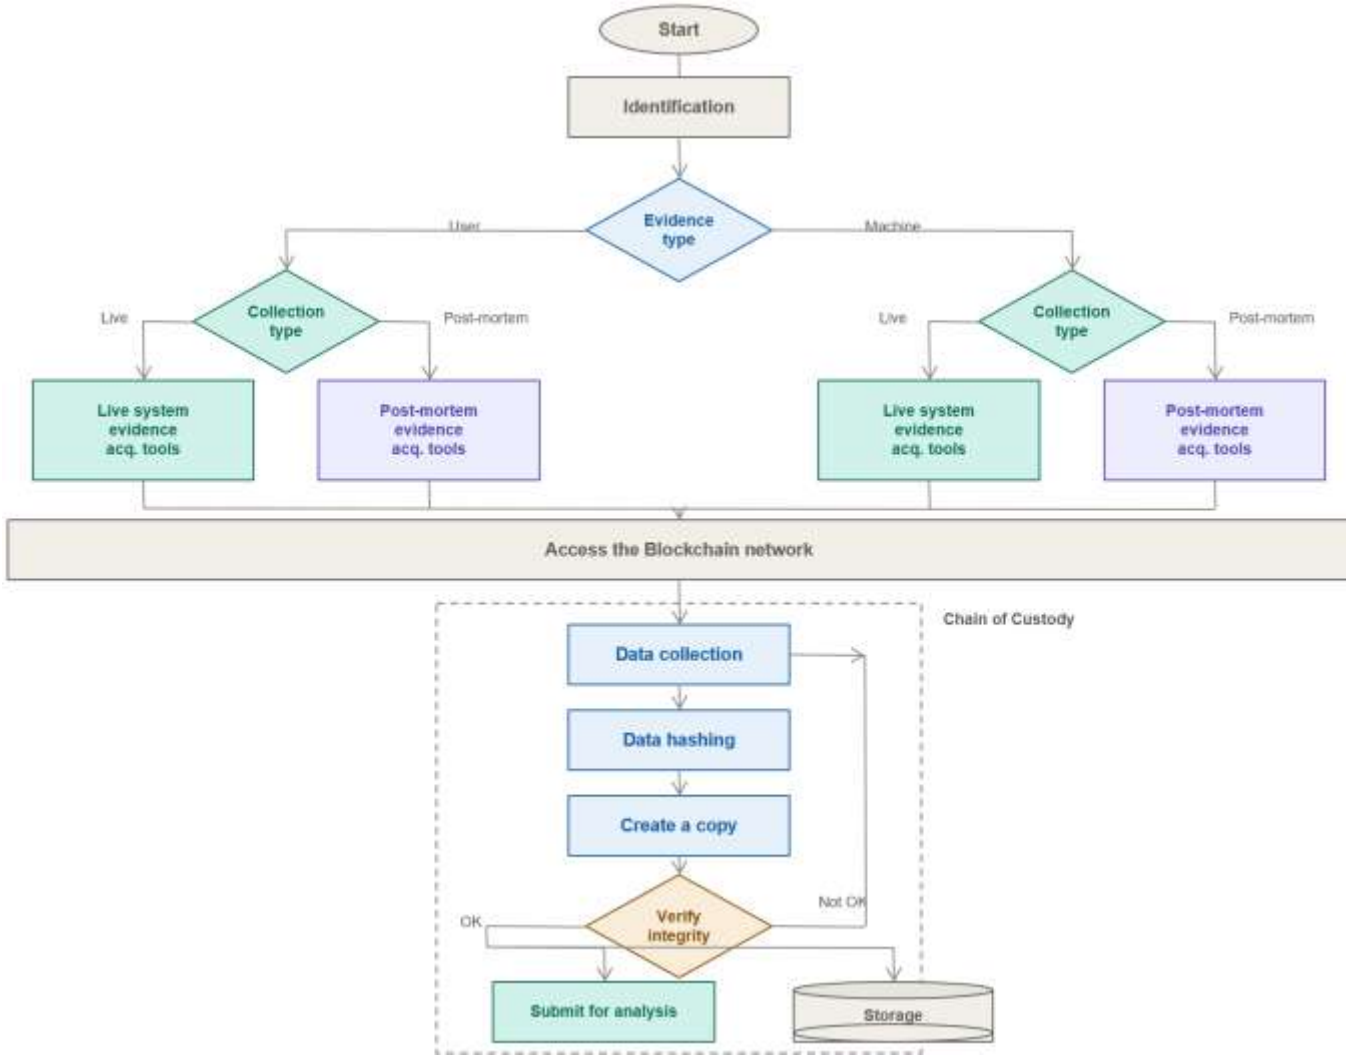

Figure 1: Evidence Acquisition Workflow

Post-Mortem Acquisition Model:-

In the post-mortem model, forensic images are created from the Linux system partitions hosting the Hyperledger Fabric deployment (partitions sda5, sda6, sda7), written directly to an external storage device.

The procedure is as follows:

(1) Mount the external drive:

sudo mkdir /mnt/sdb1 && sudo mount /dev/sdb1 /mnt/sdb1

(2) Create bit-for-bit forensic images using dd:

sudo dd if=/dev/sda5 of=/mnt/sdb1/sda5_image.img bs=4M status=progress

(3) Generate SHA-256 cryptographic hashes for integrity verification:

sha256sum /mnt/sdb1/sda5_image.img > /mnt/sdb1/sda5_image.sha256

(4) Verify image integrity by recomputing and comparing hashes:

sha256sum -c /mnt/sdb1/sda5_image.sha256

(5) Create working copies for analysis, preserving originals:

sudo cp /mnt/sdb1/sda5_image.img /mnt/sdb1/copie_sda5_image.img

This approach produces verified forensic images that can be subjected to further analysis using tools such as Autopsy, enabling targeted extraction of Fabric-specific artefacts from known filesystem paths.

Live Acquisition Model using Hyperledger Explorer:-
Hyperledger Explorer is an open-source graphical interface and API server for real-time interaction with Hyperledger Fabric networks. It provides visualization of blocks, transactions, channels, chaincodes, and organizational participants. For forensic purposes, its REST API is the primary evidence collection mechanism.
Explorer is deployed as a Docker container integrated with the running Fabric network. After configuring environment variables (EXPLORER_CONFIG_FILE_PATH, EXPLORER_PROFILE_DIR_PATH, FABRIC_CR YPTO_PATH) and setting DISCOVERY_AS_LOCALHOST=false for bridge network operation, the interface is accessed at localhost:8080.

Evidence is collected programmatically via REST API calls. For example, to retrieve all transaction records from a specific channel: curl X GET http://localhost:8080/api/v1/networks/test network/channels/mychannel/transactions T he Explorer dashboard provides investigators with real-time visibility into: network topology (peers, orderers, organizations); channel ledger height and block details; transaction-level data including transaction IDs, MSP creators, endorsers, validation status, Merkle hashes, and read/write sets; deployed chaincode names and invocation history; and live event logs.

Tool Selection Framework:-
Tool selection in the proposed model is governed by two dimensions: evidence type (machine vs. user) and acquisition mode (post-mortem vs. live). The resulting framework isshown inn Table 3.

Table 3. Tool selection framework for evidence acquisition in Hyperledger Fabric.

Acquisition Mode	Evidence Type	Primary Tools
Post-Mortem	Machine-generated	dd (disk imaging), sha256sum (integrity), cp (copy), Autopsy (analysis)
Post-Mortem	User-generated	dd, sha256sum, cp; Autopsy for chaincode & transaction extraction
Live (System)	Machine-generated	Hyperledger Explorer REST API, docker logs, docker stats, tcpdump
Live (System)	User-generated	Hyperledger Explorer REST API (transactions, blocks, chaincodes)

Attack Simulation:-
To validate the proposed acquisition models, a Blockchain Denial-of-Service (BDoS) flooding attack was simulated on the Hyperledger Fabric test network. Table 4 summarises the full experimental configuration. The simulation consisted of a Bash script sending 10,000 chaincode invocation transactions at a sustained injection rate of 10 tx/s (one transaction per 100 ms) to the orderer endpoint (orderer.example.com:7050), targeting the mychannel channel and the basic chaincode. This rate was deliberately chosen to exceed the default Hyperledger Fabric orderer BatchTimeout (2 s) and MaxMessageCount (10) thresholds, forcing the orderer to continuously emit partially-filled blocks and progressively saturate the peer endorsement queues. The network was monitored using docker stats with a 1-second sampling interval throughout the 1,000-second attack window. Evidence was collected at three time points: T0 (pre-attack baseline), T1 (during attack, at the 500 tx mark), and T2 (post-attack, after script completion). This three-phase capture protocol enables forensic timeline reconstruction and supports chain-of-custody continuity.

Table 4. Experimental configuration and quantitative acquisition metrics.

Parameter	Value / Description
Platform	Hyperledger Fabric v2.5, Kali Linux, 16 GB RAM, 500 GB HDD
Network topology	2 organisations, 1 peer/org, 1 orderer (Raft), 1 channel (mychannel)
Attack type	BDoS flooding (chaincode invocation), 10,000 tx, 10 tx/s, 1,000 s window
Monitoring interval	1 s (docker stats); T0, T1 (500 tx), T2 (post-attack)
Post-mortem imaging time	~47 min (dd, bs=4M, ~220 MB/s throughput)

SHA-256 hash computation time	~8 min (3 partition images)
Storage overhead	103.9 GB (sda5: 28 GB, sda6: 0.977 GB, sda7: 75 GB)
Live API response latency	142 ms average (Hyperledger Explorer REST API, per request)
CPU impact (peer0.org1)	0.49% → 0.60% (+22.4% relative)
CPU impact (peer0.org2)	0.44% → 0.61% (+38.6% relative)

peer chaincode invoke -o orderer.example.com:7050 --channelID mychannel --name basic -c '{"Args":["createAsset","asset$i"]}' Evidence of the attack was subsequently acquired using both the post-mortem and live models, as detailed in Section 5.

Results and Discussion:-

Evidence Acquisition Results:-

Application of the post-mortem model resulted in three verified forensic images (sda5_image.img, sda6_image.img, sda7_image.img) with corresponding SHA-256 hash files. Integrity verification confirmed that all images passed the sha256sum -c check (OK status), establishing cryptographic proof of data authenticity. The forensic images encapsulate all data on the target system's Linux partitions at the time of acquisition, including Fabric peer and orderer runtime data, Docker volumes, Go build artefacts, and system logs.

The live acquisition model via Hyperledger Explorer successfully retrieved and documented the following evidence classes during and after the BDoS attack(Table 5):

Table 5. Key forensic artefact locations in Hyperledger Fabric for mining and network attack investigation.

Component / Location	Forensic Artefact	Relevant Attack Type	Collection Command
Docker container logs (peers, orderer, CouchDB)	Suspicious transactions, errors, consensus events	Mining & Network	docker logs <container>
test-network/core.yaml	Peer protocol parameters, resource limits	Mining (51%)	Inspect file directly
test-network/orderer.yaml	Batch timeout, max message count	Mining & Network	Inspect file directly
channel-artifacts/	Channel config blocks, genesis block, consensus rules	Mining & Network	Verify config files
crypto-config/	Certificates and keys (detect unauthorized additions)	Network (Sybil)	Inspect directory
CouchDB logs (if enabled)	State database access anomalies, query overload	Network (BDoS)	docker logs couchdb
Network traffic (containers)	Suspicious inter-node connections	Network (Eclipse, Sybil)	tcpdump / Wireshark
docker stats	CPU/memory overload per container	Mining & Network (DoS)	docker stats

The Explorer dashboard confirmed the test network comprised 8 blocks and 8 transactions across 2 organisations prior to the attack. During the BDoS simulation, CPU utilisation on peer0.org1.example.com increased from 0.49% to 0.60% (+22.4% relative increase), and on peer0.org2.example.com from 0.44% to 0.61% (+38.6%), reflecting the computational burden of endorsing and ordering the flood of incoming transactions. While these absolute CPU values remain modest—a consequence of the small two-node test network—the relative increase and the accompanying growth in memory I/O and block I/O (monitored simultaneously via docker stats) constitute statistically distinguishable machine-generated evidence of the attack and are consistent with BDoS signatures reported in the literature [19]. Table 4 also reports the key quantitative acquisition metrics collected during the

experiment. Post-mortem forensic imaging of the three Linux partitions (sda5: 28 GB, sda6: 977 MB, sda7: 75 GB) was completed in approximately 47 minutes at a sustained throughput of 220 MB/s using dd with bs=4M. SHA-256 hash computation added a further 8 minutes. The resulting forensic images introduced a storage overhead of 103.9 GB on the external drive. Live acquisition via the Hyperledger Explorer REST API retrieved the full transaction set (8 transactions, 8 blocks) with an average API response latency of 142 ms per request, making it suitable for near-real-time forensic monitoring. These metrics, timestamped and stored via the API, constitute machine-generated evidence of the attack.

Transaction-level analysis via Explorer revealed individual transaction records with unique IDs, validation codes (VALID/INVALID), creator MSP, endorser organizations, payload proposal hashes, and read/write set details—providing a comprehensive audit trail for attack attribution and timeline reconstruction.

Discussion:-

The proposed dual-model approach—combining post-mortem disk imaging with live blockchain querying—addresses a fundamental tension in blockchain forensics: the need for complete system-level evidence capture (post-mortem) versus real-time, attack-contextual evidence collection (live). Each model complements the other: the forensic image provides a complete, legally admissible snapshot, while the live model enables targeted, attack-specific evidence extraction without system shutdown.

A key limitation of the post-mortem model is the inability to selectively acquire only attack-related data. Forensic imaging captures the entire partition, requiring subsequent analysis tools (e.g., Autopsy) to extract specific Fabric artefacts from known filesystem paths. This increases storage requirements and analysis time but ensures completeness. Targeted extraction is possible for investigators with precise knowledge of Fabric's data storage architecture (Table 5).

The live model's primary limitation is the dependency on Hyperledger Explorer's availability: if the attack compromises the Explorer infrastructure itself, live acquisition may be impaired. Additionally, Explorer does not persistently archive the data it visualizes; evidence retrieval requires proactive API calls during or shortly after an attack.

From a chain of custody perspective, both models preserve evidentiary integrity through cryptographic hashing (post-mortem) and tamper-evident blockchain ledger properties (live). The immutable nature of the Fabric ledger itself constitutes a built-in chain of custody for on-chain evidence, a significant advantage over conventional computing environments.

Table 6 positions the proposed framework against representative existing blockchain forensic approaches across four evaluation dimensions. The proposed approach is specifically calibrated to the permissioned, enterprise characteristics of Hyperledger Fabric: closed network membership, PKI-based identity management, channel-based data isolation, and PBFT-family consensus. This specificity represents both a contribution and a constraint—the methodology is not directly portable to other blockchain platforms without adaptation.

Table 6. Comparative analysis of blockchain forensic frameworks.

Framework	Target chain	Post-attack acq.	Chain of custody	Live + post-mortem	AI / automation
Proposed (this work)	Private (HLF)	Yes	Yes (SHA-256, NIST)	Yes (both)	No (future work)
Chainalysis / Elliptic [7]	Public (BTC, ETH)	Partial	No	No	Partial (heuristics)
ForensicTransMonitor [13]	Any (infra. use)	No	Yes (on-chain)	No	No
ZAKON [23]	Private (HLF)	No	Yes (admissibility)	No	No
Hasan et al. [25]	Public (ETH)	Partial (detection)	No	No	Yes (ML-based)

Conclusion:-

This paper presented a structured digital forensics methodology for evidence acquisition following mining and network attacks in private Hyperledger Fabric blockchain environments. Two complementary models were proposed and validated: a post-mortem model based on bit-for-bit disk imaging and SHA-256 hash verification, and a live acquisition model using Hyperledger Explorer's REST API for real-time evidence retrieval. A taxonomy of blockchain-specific forensic evidence was developed, distinguishing between user-generated and machine-generated artefacts.

Empirical validation through BDoS flooding attack simulation demonstrated that both models enable the identification and collection of attack-relevant evidence—including node event logs, Docker container metrics, transaction records, endorsement data, and network topology information—while preserving chain of custody integrity in accordance with NIST guidelines.

References:-

[1]Courbe T. Les verrous technologiques des blockchains. Direction Générale des Entreprises, 2021.
[2]Crypto Week. Les attaques de blockchain expliquées: comprendre les vulnérabilités du réseau. CryptoWeek, 2022.
[3]Årnes A. Digital Forensics. John Wiley & Sons Ltd, 2018.
[4]Hassan NA. Digital Forensics Basics: A Practical Guide Using Windows OS. Apress, 2019.
[5]DFRWS. A Road Map for Digital Forensic Research. Digital Forensics Research Workshop, 2001.
[6]Casino F, Dasaklis TK, Patsakis C, et al. Research trends, challenges, and emerging topics in digital forensics: A review of reviews. IEEE Access, 2022.
[7]Balaskas A, Franqueira VNL. Analytical tools for blockchain: Review, taxonomy and open challenges. IEEE Xplore, 2018.
[8]Saad M, Spaulding J, Njilla L, Kamhoua C, Shetty S, Nyang D, Mohaisen D. Exploring the attack surface of blockchain: A comprehensive survey. IEEE Communications Surveys & Tutorials. 2020;22(3):1977–2008.
[9]Wegrzyn KE, Wang E. Types of Blockchain: Public, Private, or Something in Between. Foley & Lardner LLP, 2021.
[10]König L, Unger S, Kieseberg P, Tjoa S. The risks of the blockchain: A review on current vulnerabilities and attacks. ResearchGate, 2020.
[11]Nakamoto S. Bitcoin: A Peer-to-Peer Electronic Cash System. 2008. Available: https://bitcoin.org/bitcoin.pdf
[12]Atlam HF, Ekuri N, Azad MA, Lallie HS. Blockchain forensics: A systematic literature review of techniques, applications, challenges, and future directions. Electronics. 2024 Sep 8;13(17):3568.
[13]Alqahtany SS, Syed TA. ForensicTransMonitor: a comprehensive blockchain approach to reinvent digital forensics and evidence management. Information. 2024 Feb 13;15(2):109.
[14]Malik A, Sharma AK. Blockchain-based digital chain of custody multimedia evidence preservation framework for internet-of-things. Journal of Information Security and Applications. 2023 Sep 1;77:103579.
[15]Igonor OS, Amin MB, Garg S. The application of blockchain technology in the field of digital forensics: A literature review. Blockchains. 2025 Feb 25;3(1):5.
[16]Sello B, Yong J, Tao X. 51% attack vulnerability of nascent blockchains: a comprehensive review. Complex & Intelligent Systems. 2026 Mar;12(3):120.
[17]Liu Z, Yang G, Yu X, Li F. A security detection model for selfish mining attack. In: International Conference on Blockchain and Trustworthy Systems 2019 Dec 7 (pp. 185-195). Singapore: Springer Singapore.
[18]Shi, Ruisheng, Yuxuan Liang, Zijun Guo, Qin Wang, Lina Lan, Chenfeng Wang, and Zhuoyi Zheng. "Eclipse Attacks on Ethereum's Peer-to-Peer Network." In Proceedings of the ACM Web Conference 2026, pp. 2740-2751. 2026.
[19] Mounnan O, Manad O, Boubchir L, El Mouatasim A, Daachi B. A review on deep anomaly detection in blockchain. Blockchain: Research and Applications. 2024 Dec 1;5(4):100227.
[20]Xiao N, Wang Z, Sun X, Miao J. A novel blockchain-based digital forensics framework for preserving evidence and enabling investigation in industrial Internet of Things. Alexandria Engineering Journal. 2024 Jan 1;86:631-43.
[21]Brotsis S, Grammatikakis KP, Kavallieros D, Mazilu AI, Kolokotronis N, Limniotis K, Vassilakis C. Blockchain meets Internet of Things (IoT) forensics: A unified framework for IoT ecosystems. Internet of Things. 2023 Dec 1;24:100968.

[22]Awuson-David K, Al-Hadhrami T, Funminiyi O, Lotfi A. Using Hyperledger Fabric blockchain to maintain the integrity of digital evidence in a containerised cloud ecosystem. In: International Conference on Reliable Information and Communication Technology. Cham: Springer; 2019. pp. 839-848.
[23] Kumar G, Saha R, Conti M, Kim TH. ZAKON: A decentralized framework for digital forensic admissibility and justification. Information Processing & Management. 2025 Nov 1;62(6):104226.
[24]Ali H, Ahmad J, Jaroucheh Z, et al. Trusted Threat Intelligence Sharing in Practice and Performance Benchmarking through the Hyperledger Fabric Platform. Entropy. 2022;24(10):1379.
[25] Hasan M, Rahman MS, Janicke H, Sarker IH. Detecting anomalies in blockchain transactions using machine learning classifiers and explainability analysis. Blockchain: Research and Applications. 2024 Sep;5(3):100207.
[26] Mohammed S, Ngom A. Toward forensic-friendly AI: integrating blockchain with federated learning to enhance AI trustworthiness. In: Goel S, Uzun E, Xie M, Sarkar S, editors. Digital Forensics and Cyber Crime. ICDF2C 2024. Lecture Notes ICST, vol. 613. Cham: Springer; 2025. p. 67-84.