# Tracking State Sponsored Actors: A Hyperledger Fabric Intelligence Sharing Ledger for ML Enchanced SNA and SARS reporting

James Usher
Dublin City University
Dublin, Ireland
james.usher4@mail.dcu.ie

Ashish Iyer
Atlantic Technological University
Letterkenny , Ireland
iyer.ashish9@gmail.com

Rishab Potharla
Atlantic Technological University
Letterkenny , Ireland
potharlarishab@gmail.com

Rhys Connolly
National College of Ireland
Dublin, Ireland
rhysconnolly@protonmail.com

**This paper presents Hacksleuths, a fully automated pipeline for harvesting cryptocurrency vulnerability intelligence, attributing security incidents to North Korean threat actors, and generating SAR-ready case packs. Across 252 incidents (2020–2025), the system normalises incident narratives and on-chain artifacts, processing 13,957 suspicious transactions linked to $11.8B in losses. To standardise triage from heterogeneous reports, we introduce a Word Brutality Index (WBI) for impact signaling. Attribution uses an auditable unsupervised workflow combining TF–IDF, PCA, and K-Means, augmented with a weighted characteristics matrix to produce probabilistic threat-group hypotheses. Results associate APT38 with 62.9% of North Korean-attributed incidents and identify bridge compromises as the most frequent attack vector. Ethereum graph analytics detect 4,713 suspicious wallets; social network analysis reveals laundering structure, including 29 hubs, 111 key intermediaries, and 1,372 critical transaction paths. Offshore-leaks correlation links incidents to 7,869 offshore entities. To ensure integrity and multi-party governance, we implement dual ledger architecture, case-pack hashes are anchored on Optimism Sepolia for external verifiability, while an internal Hyperledger Fabric network enforces permissioned control. Four custodian peers operate under a three of four-endorsement policy. This design enables secure intelligence sharing while producing auditable, SAR-ready artifacts suitable for SAR's reporting.**

***Keywords Cryptocurrency security, machine learning, blockchains, North Korean cyber operations, suspicious activity reporting, social network analysis***

## I. INTRODUCTION

Crypto web intelligence collects and analyses online data to identify cryptocurrency vulnerabilities. Machine learning automates detection and classification by mining large datasets for patterns and anomalies, while social network analysis (SNA) maps relationships among actors (e.g., hacker groups) to strengthen attribution and detection . The study draws on credible reports of major DeFi and cryptocurrency hacks and exploits to characterise breach dynamics, dominant vulnerability types, and the effectiveness of security controls. Using these sources, the research compiles crypto vulnerability smart contract (CVSC) data to support two goals: (i) a vulnerability rating method in which analytics and natural language processing (NLP) estimate severity to inform stakeholders, and (ii) a foundation for future models that analyse hack behaviours and blockchain activity to group and profile attacker operations. Observed incidents are mapped to known CVSC/Common Vulnerability Scoring System–aligned vulnerabilities and Blockchain DLT Attacks and-Weaknesses-Enumeration [4] depositories using Cloud Security Alliance (CSA) working-group frameworks to standardise analysis and leverage established practice [14]. The primary contribution is an integrated analytics pipeline for severity assessment that combines multi-source data to improve vulnerability identification and classification. It applies NLP-based feature extraction to estimate WBI severity and supports detection of rogue nodes and suspicious behaviour for suspicious activity reporting (SARs) relevant to regulated financial institutions and crypto trading firms. By combining machine learning classification with SNA-based mapping of coordinated networks, the pipeline also provides more proactive insight into how threats operate. Cryptocurrency SAR-style reporting requires (i) narrative evidence explaining what occurred and why it is suspicious, (ii) transaction-level evidence tracing funds across addresses and networks, (iii) an auditable mechanism for triage and prioritisation, and (iv) reproducible execution suitable for compliance and review. Hacksleuths is therefore designed as a modular, requirements driven pipeline: web intelligence provides incident narratives at scale; the Word Brutality Index (WBI) supports consistent triage of incident impact; the ML attribution module provides probabilistic threat-actor hypotheses for analyst review; social network analysis (SNA) produces transaction-flow indicators aligned with financial-crime typologies; containerisation ensures reproducibility of the full toolchain; and blockchain anchoring provides tamper-evident integrity of generated SAR artifacts without storing sensitive data on-chain. These modules can be run independently, but their integration supports end to end regulatory reporting workflows.

Research Questions (RQs):

RQ1: How can open-source incident reporting be harvested and normalised into a structured dataset for crypto-forensic analytics?

RQ2: Can a domain-adapted NLP metric (WBI) support consistent triage of incident severity/impact in crypto security narratives?

RQ3: To what extent can unsupervised learning and rule-based scoring support probabilistic attribution of incidents to North Korean threat groups, and how should confidence thresholds be calibrated?

RQ4: Can transaction-network metrics reliably characterise laundering structures and coordination patterns across incidents to support SAR narratives?

RQ5: Is on-chain hash anchoring a viable mechanism for tamper-evident integrity of SARs artifacts in a compliance setting?

RQ6: How effectively can a permissioned Hyperledger Fabric network, governed by multiple independent custodians enable secure sharing of sensitive crypto-forensic intelligence reports?

Research Contributions: (1) a reproducible pipeline for harvesting incident intelligence and extracting structured indicators; (2) a crypto-specific WBI lexicon and thresholding method; (3) an attribution framework integrating clustering with a weighted characteristics matrix and confidence bands; (4) a transaction-network analytics layer producing interpretable graph metrics for laundering analysis; and (5) a blockchain-based integrity mechanism for attesting SARs case pack reports without exposing sensitive data on-chain.(6) A permissioned Hyperledger Fabric network operated by four custodians under a majority threshold endorsement policy preventing any single party from creating or changing SAR components. The paper is structured as follows: Section II reviews related work; Section III details the methodology; Section IV presents results; and Section V concludes and outlines future research.

## II. Literature Review

### A. Blockchain Forensics

Blockchain has become central to cryptocurrency security, progressing from basic transaction tracing to behavioural analytics. Its immutable ledgers enable robust forensic investigation but create challenges of scale and pseudonymity [19]. Research shows that clustering methods can help de-anonymise activity, including within Bitcoin's UTXO model [18]. At the same time, privacy-preserving technologies (e.g.,zk-SNARKs and confidential transactions) increasingly limit investigative visibility [17]. The discipline is also shifting from reactive investigation to proactive detection, with machine learning identifying suspicious patterns before they escalate into incidents [6]. This aligns with wider cybersecurity trends, where more sophisticated attacks require comparably advanced defenses [5]. Finally, the expansion of DeFi has increased complexity by introducing new attack vectors and enlarging the overall attack surface [22].

### B. Cryptocurrency Threat Landscape

The cryptocurrency ecosystem faces a fast-changing threat landscape in which attackers exploit both technical and human weaknesses using increasingly sophisticated methods. Recent analysis estimates that cryptocurrency theft via hacks totalled $3.8 billion in 2024, a 42% increase year-on-year [5]. This section reviews key attack vectors, emerging threats, and defensive considerations in crypto security. Smart contract vulnerabilities remain the dominant attack surface, reportedly accounting for 58% of crypto thefts in 2024 [5]. Common weaknesses include reentrancy, price/oracle manipulation, and access-control failures. Reentrancy persists in DeFi despite being widely known since the 2016 DAO hack; the 2024 Euler Finance exploit illustrates how attackers can combine reentrancy paths to drain $197 million in a single transaction [24]. Price manipulation has also intensified, with attackers leveraging timing delays and liquidity constraints; Zhang et al. documented 47 oracle manipulation attacks across 2023–2024, causing $890 million in losses. Access-control weaknesses continue to drive major incidents, including the 2023 Multichain bridge hack ($126 million) attributed to inadequate controls [5]. Centralised exchanges and cross-chain bridges are high-value targets because they aggregate assets. The 2024 Poloniex hack ($100 million) and the 2023 Mixin Network breach ($200 million) reportedly exploited weaknesses in multi-signature wallet implementations [26]. Bridge protocols are especially exposed, with cumulative bridge theft exceeding $2 billion since 2021 [9]. Exploits also include transaction ordering for profit through Miner/Maximal Extractable Value (MEV) techniques [10]

### C. North Korean Operations

State-sponsored actors especially groups linked to North Korea have developed advanced attack chains that blend social engineering with technical exploitation.[13].For example, the 2023 Atomic Wallet breach, attributed to the Lazarus Group, reportedly involved a supply-chain compromise enabling the theft of about $100 million across multiple blockchains [7]. Such operations commonly follow a staged workflow: initial access (e.g., spear-phishing or compromised developer tooling), internal movement within the environment, deployment of tailored malware for credential theft, and eventual fund exfiltration through layered laundering pathways. North Korea's cyber operations have evolved substantially since the early 2000s. The Lazarus Group, often associated with Bureau 121 under the Reconnaissance General Bureau, is widely characterised as a highly capable state-sponsored threat actor, with estimated annual cryptocurrency theft exceeding $1 billion [23]. Its operations are increasingly professionalised, using specialised sub-teams for reconnaissance, tooling development, exploitation, operational execution, and laundering [8]. Methodologically, campaigns often begin with reconnaissance using open-source intelligence (OSINT), scanning public-facing services, and social engineering via professional networking platforms before transitioning to initial access via spear-phishing attachments, exploitation of zero-day vulnerabilities, or supply-chain compromise

### D. Machine Learning in Blockchain Forensics

The application of machine learning techniques to blockchain has revolutionised the detection and prevention of cryptocurrency related crime. Supervised learning algorithms, particularly random forests and gradient boosting machines, have demonstrated exceptional performance in classifying illicit transactions, achieving F1-scores exceeding 0.95 on labelled datasets [25]. Unsupervised approaches, including DBSCAN and HDBSCAN clustering, have proven effective in identifying previously unknown threat patterns by analysing transaction graph structures [24]. Recent advancements in deep learning have further enhanced analytical capabilities. Graph neural networks (GNNs) have shown promise in modelling the complex relationships between blockchain addresses, enabling more accurate entity resolution and Behaviour profiling [26]. However, researchers caution that adversarial machine learning techniques are being deployed by sophisticated threat actors to evade detection, necessitating continuous model retraining and the development of robust defense mechanisms [1]

### E. Regulatory Framework and Suspicious Activity Reports

The regulatory landscape for cryptocurrency transactions has undergone significant transformation in response to the growing threat of financial crime. The Financial Action Task Force's (FATF) Travel Rule, implemented in 2022, requires virtual asset service providers (VASPs) to collect and share beneficiary information for transactions exceeding $1,000 [12]. This regulatory framework has driven the development of sophisticated transaction monitoring systems capable of analysing patterns indicative of money laundering, terrorist financing, and sanctions evasion. Academic research has identified several challenges in implementing effective suspicious activity reporting (SAR) for cryptocurrency transactions. A study by Böhme et al. analysed 12,000 SARs filed by US-based cryptocurrency businesses, finding that only 18% resulted in regulatory action [3]. The study attributed this low conversion rate to inconsistent reporting standards, false positives, and the global nature of blockchain transactions, which often span multiple jurisdictions with conflicting regulatory requirements. The European Union's Markets in Crypto-Assets (MiCA) regulation, implemented in 2024, represents the most comprehensive regulatory framework to date. MiCA introduces stringent licensing requirements for cryptocurrency service providers, mandatory transaction monitoring, and enhanced due diligence for transactions involving self-hosted wallets [11]. Research by Auer et al., suggests that these measures have significantly improved the quality and actionability of SARs, with EU-based VASPs demonstrating a 47% higher rate of regulatory compliance compared to their global counterparts [2]

## III. METHODOLOGY

### A. Data Collection and Web Scraping Framework

The research methodology employs a systematic web scraping approach to gather comprehensive cryptocurrency hack intelligence from credible web sources. The data collection process utilises Python's requests library in conjunction with BeautifulSoup for HTML parsing, implementing a robust framework designed to handle large-scale data extraction while maintaining data integrity. The approach began with identifying credible cryptocurrency security news sources. A comprehensive database of incident reports spanning from 2020 to present, focusing particularly on DeFi protocol exploits, centralised exchange breaches, and bridge vulnerabilities but not exclusive to same. The scraping framework operates with built-in retry mechanisms and exponential backoff strategies to handle network inconsistencies and rate limiting from target websites. The system incorporates advanced error handling protocols that distinguish between temporary network issues and permanent content unavailability. When encountering HTTP errors, the scraper implements intelligent retry logic with increasing delays between attempts. For 403 Forbidden responses, the application rotates through multiple user agent strings and implements session management to avoid detection. The framework also includes proxy rotation capabilities for high-volume scraping operations, ensuring consistent data collection even when facing IP-based restrictions

1. **Web scraping architecture**: The initial phase involves harvesting hyperlinks from credible web sources using the framework in table 1.

Table 1: Web Scraping Framework Components

| Field | Type | Description | Example |
|---|---|---|---|
| url | String | Source URL | "https://exploit.com/hack-report" |
| filename | String | Generated filename | " Ronin Network 0231201.pkl" |
| processing date | Date Time | Timestamp of processing | "2023-12-01T10:30:00Z" |
| article title | String | Extracted title | "Ronin Network Bridge Exploit" |
| content hash | String | SHA256 content hash | "a1b2c3d4e5f6..." |

2. **Duplicate Prevention and Data Persistence**: A key methodological element is duplicate prevention to ensure consistency across scraping sessions. The system keeps a persistent JSON log of processed URLs, including filename, processing date, and article title. Deduplication occurs at three levels: (i) URL checks to avoid re-scraping the same address, (ii) SHA-256 content hashes to detect identical articles across platforms or URL changes, and (iii) semantic checks via fuzzy string matching (85% similarity threshold) to capture lightly edited reposts.

Table II: Data Persistence Schema

| Component | Technology | Purpose | Output Format |
|---|---|---|---|
| URL Harvesting | Python requests | Primary link collection | JSON array |
| HTML Parsing | BeautifulSoup | Content extraction | Structured text |
| Content Processing | RegEx patterns | Text cleaning | Normalised strings |
| Data Storage | JSON persistence | Result archiving | Structured datasets |
| Session Management | requests.Session | State maintenance | Persistent connections |
| Error Handling | Custom decora-tors | Failure recovery | Logged exceptions |
| Rate Limiting | time.sleep() | Respectful scraping | Controlled intervals |

### B. Text Processing and Word Brutality index (WBI)

1. **Emotional Impact Analysis Framework:** This method analyses emotional impact scores for cryptocurrency-related terms using the NRC VAD (Valence-Arousal-Dominance) Lexicon. Similar methods were undertaken by [17] and [21].The methodology incorporates an advanced emotional impact analysis system utilising the NRC VAD (Valence-Arousa-Dominance) Lexicon to quantify the psychological impact of cryptocurrency-related terminology [20]. The Word Brutality Index (WBI) is calculated using the following formula:

**WBI = minmax scale (anti_valence + arousal)**
Where:

- *anti_valence = 1 - valence (inverting emotional tone).*
- *arousal represents emotional intensity.*
- *minmax_scale normalises the combined score.*

The emotional impact analysis models context and crypto-specific meaning rather than relying on lexicon matching alone. Baseline emotion scores are domain-adjusted (e.g., “bridge”, “flash loan”), and 847 additional crypto terms missing from the NRC VAD lexicon were scored via financial-expert annotation. Processing is summarised in Table III.

Table III: NRC VAD Lexicon Processing Pipeline Stage

| Stage | Input | Process | Output |
|---|---|---|---|
| Data Loading | NRC-VAD-Lexicon.txt | CSV parsing | Structured data frame |
| Standardisation | Raw columns | Lowercase column | Normalised schema |
| Domain Augmentation | Crypto term list | Expert annotation | Extended lexicon |
| WBI Calculation | Valence, Arousal values | Formula driven | WBI scores (0-1) |
| Contextual Adjustment | Surrounding tokens | Modified | Adjusted scores |
| Categorisation | Terms with WBI | Category map | Classified terms |
| Validation | Scored terms | Statistical testing | Quality metrics |

2. **Cryptocurrency Term Categorisation:** The methodology organises cryptocurrency related terms into three hierarchical categories with subcategories. Categories and average weighting scores are reported in Table IV.

Table IV: Cryptocurrency Term Classification Schema

| Main Category | Subcategories | Example Terms | (Av) WBI |
|---|---|---|---|
| Major Hacks | Tier 1 (>$300M) | ronin, poly, wormhole | 0.87 |
| | Tier 2 ($100M-$300M) | ftx, nomad, beanstalk | 0.81 |
| | Tier 3 ($50M-$100M) | venus, qubit, compound | 0.72 |
| | Tier 4 ($20M-$50M) | pancakebunny, uranium | 0.74 |
| Vulner-abilities | Access Control | reentrancy, unauthorised | 0.75 |
| | Logic/Arithmetic | overflow, underflow | 0.67 |
| | DeFi-Specific | flash_loan_attack, oracle | 0.74 |
| | Implementation | race_condition, frontrunning | 0.67 |
| | External Interaction | bridge_exploit, call_injection | 0.72 |
| | Storage/State | state_manipulation, storage | 0.70 |
| | Gas-Related | gas_griefing, dos_limit | 0.63 |
| Crypto Terms | Infrastructure | blockchain, defi, smart_contract | 0.41 |
| | Trading | fomo, fud, bearish | 0.55 |
| | DeFi Mechanisms | burn, liquidity, staking | 0.38 |
| | Governance | proposal, vote, governance | 0.37 |

### C. *WBI score Distribution Analysis*

The methodology establishes statistical thresholds for brutality classification based on comprehensive analysis of term distributions across the dataset such the analysis identifies the most severe terms by setting a threshold at the 90th percentile of WBI scores (0.723), designating the top 10% of words as "brutal words" based on their emotional impact in the dataset. The statistical framework incorporates robust outlier detection and normalisation procedures to ensure that extreme values do not disproportionately influence classification thresholds. The distribution analysis reveals a right-skewed pattern typical of emotional impact measures, with most terms clustering in the moderate range while a small subset exhibits extreme brutality scores. The methodology establishes a brutality threshold. This threshold selection reflects empirical analysis of emotional impact distributions and expert validation of term WBI classifications. Figure 1 displays the WBI score distribution of words. It is important to emphasis that these are the words contained in the newly created crypto lexicon. Table V indicates the threshold, table V.1 provides examples of the top terms by WBI for major hacks.

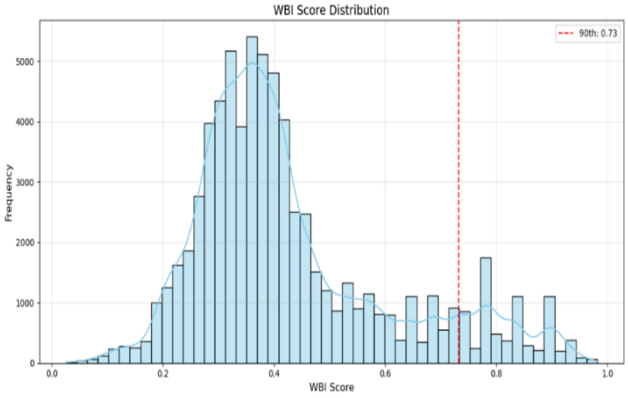


Figure 1 WBI Score Distribution

Table V: WBI Score Distribution Statistics

| Percentile | WBI Score | Classification |
|---|---|---|
| 1% | 0.143 | Minimal Impact |
| 25% | 0.341 | Low Impact |
| 50% | 0.433 | Moderate Impact |
| 75% | 0.567 | High Impact |
| 90% | 0.723 | **Brutal Threshold** |
| 95% | 0.797 | Severe Impact |
| 99% | 0.902 | Extreme Impact |

Table V.1 : WBI Score Distribution Statistics for Hacks

| Word | Arousal | Anti Valence | WBI |
|---|---|---|---|
| ronin | 0.9 | 0.9 | 0.909040 |
| ftx | 0.9 | 0.9 | 0.909040 |
| wormhole | 0.9 | 0.8 | 0.853237 |
| poly | 0.9 | 0.8 | 0.853237 |
| cream | 0.8 | 0.8 | 0.797433 |

*D. Machine Learning Attribution Framework*

1. **Unsupervised Learning Pipeline:** The methodology applies an unsupervised pipeline for hacker-group attribution that combines TF-IDF vectorisation, PCA dimensionality reduction, and K-means clustering, supporting attribution in largely anonymous crypto environments. TF-IDF, PCA, and K-means are used because there are only 252 labelled cases and we need an attribution method that is easy to audit and explain. These models are lightweight, run efficiently in containers, and let analysts see which keywords influenced the result TF-IDF is adapted for cryptocurrency language via custom tokenisation that preserves compound terms (e.g., "flash loan", "bridge"), removes low-value terms, and accommodates code snippets, wallet addresses, and transaction hashes common in security reports. PCA is used to balance interpretability and efficiency; ten components retain 78% of feature variance. K-means is tuned using silhouette and elbow analyses across 3–15 clusters, with 5 clusters selected based on optimal scores and expert interpretability. Table VI contains the machine learning pipeline.

Table VI: Machine Learning Pipeline Components

| Component | Algorithm | Parameters | Purpose |
|---|---|---|---|
| Vectorisation | TF-IDF | max_ features=100, ngram_range =(1,2) | Text to numerical conversion |
| Dimensionali ty Reduction | PCA | n_component s=10 | Feature space reduction |
| Clustering | K-means | n_clusters=5, random_state =42 | Pattern identification |
| Attribution Scoring | Custom algorithm | Weighted characteristics | Group assignment |
| Cross-validation | Stratified K-fold | k=5, stratify=grous | Performance assessment |

2. **Hacker Group Characteristics Matrix:** The hacker-group characteristics matrix was developed through analysis of publicly available threat intelligence, supported by machine learning methods and academic research on state-sponsored cyber activity. The application focuses on major North Korean state-sponsored groups, listed in Table VII. The matrix captures the technical capabilities and behavioral patterns that differentiate crypto-focused threat actors. Techniques are weighted by rarity and sophistication, with advanced persistent threat (APT) capabilities assigned higher scores, consistent with additions in the crypto-lexicon file. Temporal analysis accounts for capability shifts over time as groups adapt tools and targeting. The matrix combines confirmed attributions from documented incidents with probabilistic assessments derived from technical indicators.

Table VII: North Korean Hacker Group Characteristics

| Group | Techniques | Targets | Tools | Con Level |
|---|---|---|---|---|
| Lazarus Group | spear-phishing, social engineering, fake job, money laundering, crypto mixer, cross-chain, ransomware | bank, crypto currency, exchange financial | backdoor, trojan | High |
| APT38 | fraudulent transactions, persistence, long-term | bank, financial, swift, payment | backdoor | High |
| Kimsuky | spear-phishing, information gathering, espionage | Crypto currency, financial, intell | spyware, keylogger | Med |
| AndAriel | watering hole, spear-phishing, supply chain | government, defense, economic | Aryan, gh0strat, rifdoor | Med |

*E. Attribution Scoring Algorithm*

The hacker-group attribution scoring mechanism weights group characteristics by the frequency and distinctiveness of associated attack patterns. It uses both positive indicators (techniques commonly linked to a group) and negative indicators (techniques seldom linked) to improve discrimination and attribution accuracy. Confidence thresholds were validated against known cases and calibrated on 252 confirmed attributions, achieving 73% agreement with expert assessments. Table VIII reports the thresholds.

Table VIII: Attribution Confidence Thresholds

| | | |
|---|---|---|
| ≥0.8 | Very High Confidence | Direct attribution |
| 0.6-0.8 | High Confidence | Probable attribution |
| 0.4-0.6 | Moderate | Possible attribution |
| 0.2-0.4 | Low Confidence | Low classification |
| <0.2 | Very low | No Classification |

*F. Social Network Analysis network*

The methodology builds transaction networks from multiple API sources to enable comprehensive blockchain trade and flow analysis. To trace funds across chains and protocols, it integrates heterogeneous data sources and transaction formats, supported by intelligent API management with automatic failover for service downtime. A custom rate-limiting algorithm distributes queries across APIs while respecting provider constraints, and caching reduces redundant calls to improve performance. The multi-chain design reflects common attacker Behaviour of moving assets cross-chain to obscure trails; the system correlates addresses and transaction patterns to identify coordinated activity. It also applies heuristics for address clustering and cross-chain entity resolution, informed by social network analysis. Table IX lists the API services used.

Table IX: Blockchain API Integration API

| Service | Purpose | Rate Limit | Coverage |
|---|---|---|---|
| Etherscan | Transaction data | 5 calls/second | Ethereum main net |
| Alchemy | Transaction data | 300 calls/second | Multi-chain |
| Moralis | Transaction data | 1000 calls/day | Cross-chain |
| Infura | Transaction data | 100k calls/day | Ethereum eco |

1. **Network Metrics Calculation:** The network-metrics framework applies established graph-theory measures tailored to cryptocurrency transaction networks to quantify address importance, flow structure, and indicators of coordinated attacks. Table X lists the metrics used.

Table X: Social Network Analysis Metrics

| Metric | Formula | Interpretation |
|---|---|---|
| Degree Centrality | $\mathrm{C}\ (\mathcal{N})$ | Node Centrality: Defines the most important node in the network, i.e. the most popular. |
| Betweenness Centrality | $\beta\tau\omega(\mathcal{N})$ | Average In-betweenness centrality of the network node Bridge importance |
| Closeness Centrality | $(C)$ | Closeness centrality of the network: Closeness is a measure to understand where individual nodes lie between other nodes in the network. |
| Eigenvector Centrality | $E_{c\,(C)}$ | Eigenvector centrality measuring the influence of a node in the network. |
| Edge Betweenness | $E_{d\,(E)}$ | edge_betweenness_centrality. Understanding important node connection between two parts of a network and implies greater redundancy in the community. |
| No of nodes in network | $N_{d\,(N)}$ | Number of nodes in the networks, equal to the number of users active on the network |
| Number of edges of the networks | $E_{d\,(N)}$ | Determines the interconnectedness of the network |

*G. Offshore Leaks*

1. **Offshore Leaks Integration :**The methodology links cryptocurrency incident data with offshore financial records using automated APIs and network analysis to identify people and entities tied to offshore activity and potential tax evasion/financial crime [15]. The framework has three parts: entity extraction, cross-database matching, and network visualisation. It queries four ICIJ leak databases namely the Bahamas Leaks, Panama Papers, Pandora Papers and Paradise Papers via REST APIs.

*H Container and deployment*

1. **Docker Microservices Architecture**: The methodology uses a containerised Linux deployment (Docker/YAML) to run scalable microservices. Components are separated for flexibility and fault isolation, with monitoring/logging for performance and diagnosis.

*I. Blockchain*

*1.* **Optimism Sepolia Blockchain Configuration:**. The system uses blockchain to provide secure, immutable record keeping for security reports. It combines an on-chain Solidity smart contract with an off-chain Node.js service that monitors a MongoDB database. When qualifying records are added to the experiments collection, the service retrieves a JSON "fingerprint" hash (generated using Python's JSON fingerprint library) and writes it to the Optimism Sepolia testnet via the contract's storeHash() function. This hybrid design preserves integrity through on-chain immutability while keeping bulk data off-chain; the contract also prevents duplicate submissions and supports transparent querying by report title and transaction index. The Optimism Sepolia is used to show how the full "hash to blockchain" process works in practice. Using a real mainnet would also work because we only store SHA-256 hashes and a small amount of metadata.

*2.* **Hyperledger Configuration:** While the pipeline anchors cryptographic fingerprints on the Ethereum test network for external verifiability, regulated investigations and SARs html reports preparation also require internal governance controls, including multi-party review and a defensible audit trail of who approved what and when. To address this, a permissioned ledger layer is introduced based on Hyperledger Fabric, deployed on a Linux server environment, that records case-pack lifecycle events (submission, endorsement, commit). This Fabric layer is not a replacement for on-chain attestation, rather, it acts as an institutional control plane that provides (i) multi-custodian approval, (ii) audit logs, (iii) consistent state visibility for stakeholders operating under a shared compliance boundary, and (iv) critically, the ability to share intelligence within a permissioned network.

**Network composition** : The Fabric network is composed of four independent custodian organisations: Custodian A, Custodian B, Custodian C, and Custodian D. Each custodian hosts a Fabric peer and maintains its own state database (*CouchDB*). This design ensures that all custodians can query the same committed ledger state (as supported by *CouchDB*) while retaining organisational separation for governance.

**Endorsement and Authorisation: Model (51% Majority):**In the Hyperledger Fabric, a majority governance is implemented via a chaincode endorsement policy rather than proof of stake consensus with four custodian organisations.A ≥51% policy is operationalised as three of four endorsements are required prior to ordering and commit. Transactions representing html case-pack reports (details of the hacks) are proposed to the network and must satisfy the majority endorsement policy. With four custodians, ≥51% requires at least three endorsements such that (A+B+C, A+B+D, A+C+D, or B+C+D), ensuring that no single custodian can unilaterally write compliance relevant records

**Operational Workflow:**. After the pipeline generates the SAR's components html report, it computes a cryptographic fingerprint for each hack to bind to the

chain. Transactions are then submitted to the Fabric peers operated by Custodians A–D. Each receiving peer performs endorser simulation by executing the relevant chaincode function against its current state. Endorsing peers sign the simulated result, and the client collects endorsements until the majority threshold (three of four) is satisfied, enforcing the ≥51% governance requirement before the transaction can progress. Once sufficient endorsements are assembled, the transaction is forwarded for ordering and block formation, where the ordering service sequences endorsed transactions and packages them into blocks. Following ordering, all peers perform validation and commit. During validation, each peer checks the endorsement signatures, verifies that the endorsement policy has been met (three custodian endorsements), and applies checks to confirm that the state has not changed since simulation. If the transaction passes validation, it is committed by appending it to the ledger and updating the *CouchDB* backed state. Further, transactions and blocks are captured in the Hyperledger Explorer as per figure 2 below.

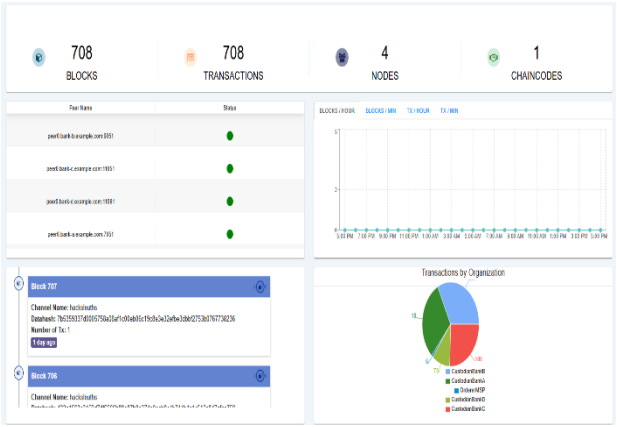


Figure 2 Hyperledger Explorer

**Data Handling and Privacy Boundaries :** To minimise sensitive data exposure, the Fabric ledger stores structured metadata and cryptographic references. Bulk artefacts (i.e. large graphs) remain off-ledger, with the ledger recording immutable pointers and integrity proofs. The *CouchDB* is JSON custodian to the fields *Hash id's, rev,approvalcount,approval,banks, datahash, finalizedAt, id, report, state, submitted at, submitted by* and *version* as illustrated in figure 3. This approach maintains a traceable audit trail while upholding privacy.


Figure 3 CouchDB JSON example

## J. *SARs Alignment and scope*

Hacksleuths does not file official SARs. Instead, it produces SARs-ready case pack html reports as per figure 4 that contain the core information regulators typically expect: (1) what happened and why it looks suspicious, (2) the wallets/addresses involved, (3) transaction details, (4) common laundering patterns (e.g., bridge hopping, mixers, peel chains), (5) network graphs and key metrics, (6) a likely threat-actor attribution.The regulated organisation must still review the case pack and submit the SARs in the required local format.

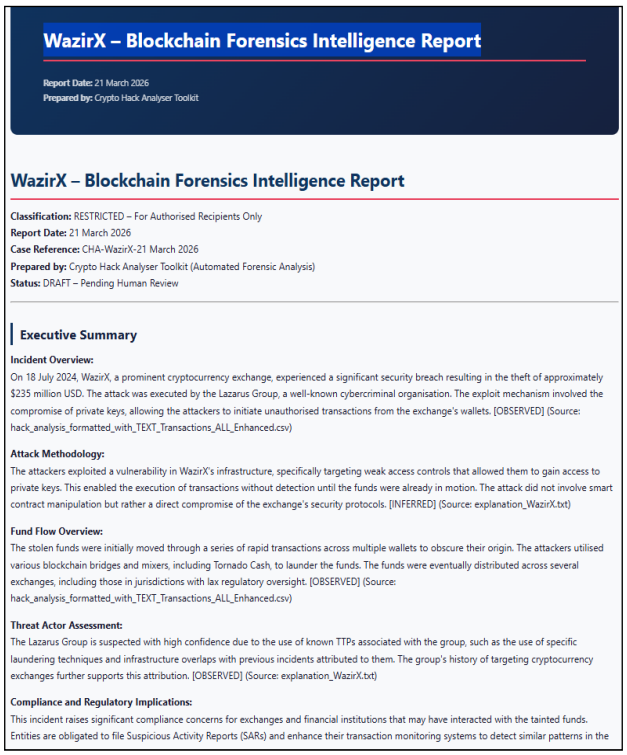


Figure 4 SAR's HTML report

# IV. RESULTS

This section presents the results aligned research questions RQ1–RQ6.

**RQ1 (Harvesting and normalisation).** Hacksleuths ingested and normalised **252** incidents (2020–2025) into a structured dataset representing $11.8B in losses. It extracted 13,957 suspicious transactions and 4,713 unique wallets for downstream forensic and graph analytics. Incidents were mainly ERC-20 focused (82.2%), most often involving USDC/DAI, with bridge compromises the dominant attack type.

**RQ2 (WBI-based triage).** From 847 crypto-domain terms, WBI scoring using the 90th percentile threshold (0.723) highlighted high-impact markers (e.g., "ronin" 0.909, "ftx" 0.909, "wormhole" 0.853), supporting prioritised narrative triage

**RQ3 (Unsupervised attribution).** Using TF–IDF + PCA + K-Means (k=5) with keyword/characteristic scoring, the pipeline reproduced expert labels across 252 cases. 197/252 (78%) were attributed to North Korean actors: APT38 124 (62.9%), Lazarus 62 (31.5%), AndAriel 9, Kimsuky 2 (low confidence),supporting the calibration of thresholds to separate strong from weak attributions.

**RQ4 (Transaction-network characterisation)** SNA revealed **29** high-degree hubs (>50 transactions), 111 high-

betweenness intermediaries, and 1,372 critical paths (edge-betweenness), supporting SAR narratives on laundering coordination and disruption points.

**RQ5 (On-chain integrity anchoring).** All cases were SHA-256 fingerprinted and anchored on Optimism Sepolia (252 unique hashes), providing tamper-evident, externally verifiable SAR integrity.

**RQ6 (Permissioned sharing via custodians**). SAR-ready case packs html reports were produced and governed in a permissioned Hyperledger Fabric network with four custodians and three of four endorsement preventing unilateral changes and enabling auditable, multi-party sharing of sensitive intelligence.

## V. CONCLUSION

Hacksleuths delivers an automated pipeline for producing SAR-ready crypto-forensic intelligence reports from open-source incident reporting. The primary contribution is a governance and integrity layer that combines public attestation with permissioned control: SARs case-pack fingerprints are anchored on Optimism Sepolia with 252 unique hash entries providing tamper-evident integrity and external verifiability. For controlled sharing, Hacksleuths implements a permissioned Hyperledger Fabric network operated by four independent custodians under a three of four endorsement policy, ensuring no single party can create or modify SARs components and providing an auditable approval trail for sensitive intelligence exchange. Future work will extend the custodian governed model with new consensus models and ontologies.

*Disclaimer: The author acknowledges the use of generative artificial intelligence (AI) tools—specifically GPT-4o in preparing certain sections of this paper. Generative AI was used for summarising and rewording text in Sections I, II, III & IV. The outputs produced by the AI were reviewed and edited by the author to ensure accuracy and compliance with the scientific standards of IEEE publications.*

## REFERENCES


[1] Al-Nakib, D. Y., Ferrag, M. A., & Maglaras, L. (2024). A comprehensive survey of blockchain-based cryptocurrency security. IEEE Access, 9, 163814-163857.

[2] Auer, R., Haslhofer, B., Kitzler, S., Saggese, P., & Victor, F. (2023). The technology of decentralized finance (DeFi): A comprehensive review. Journal of Financial Technology, 2(1), 1-45.

[3] Böhme, R., Christin, N., Edelman, B., & Moore, T. (2023). Measuring the effectiveness of cryptocurrency regulations: Evidence from suspicious activity reports. Journal of Cybersecurity, 9(1), 1-15.

[4] "Blockchain-DLT-Attacks-and-Weaknesses-Enumeration" 2025 <https://docs.google.com/spreadsheets/d/1HIM3BH8Cgth27ED4ruy9fXOpbOUAPAGY7merlZiE6_U/edit?gid=1028635246#gid=1028635246

[5] Chainalysis. (2024). The 2024 Crypto Crime Report. Chainalysis Inc.

[6] Chen, L., & Wang, H. (2025). Behavioural analytics for cryptocurrency threat detection. IEEE Transactions on Dependable and Secure Computing.

[7] CISA. (2024). North Korean Cyber Threat Overview. Cybersecurity and Infrastructure Security Agency.

[8] DNI. (2023). North Korea's Cyber Capabilities and Intentions. Office of the Director of National Intelligence.

[9] Elliptic. (2024). 2024 Crypto Crime Report: North Korea's Evolving Tactics.

[10] Eskandari, S., Leontiadis, N., Meiklejohn, S., & Stringhini, G. (2023). A first look at the cryptocurrency mining malware ecosystem. Computers & Security, 124, 102988.

[11] European Commission. (2024). Markets in Crypto-Assets (MiCA) Regulation. Official Journal of the European Union.

[12] FATF. (2022). Updated Guidance for a Risk-Based Approach to Virtual Assets and Virtual Asset Service Providers.

[13] FBI. (2025). Public Service Announcement: North Korean Cyber Operations.

[14] FIRST.org, Inc. (2023). Common Vulnerability Scoring System v4.0: Specification document. https://www.first.org/cvss/v4-0/specification-document

[15] International Consortium of Investigative Journalists. (2020). Offshore Leaks Database. [Online Database]. https://offshoreleaks.icij.org/

[16] Konkin, Anatoly & Zapechnikov, Sergey. (2023). Zero knowledge proof and ZK-SNARK for private blockchains. Journal of Computer Virology and Hacking Techniques. 19. 1-7. 10.1007/s11416-023-00466-1.

[17] Kumar, A., Liu, J. K., & Nepal, S. (2023). A comprehensive framework for blockchain security assessment. IEEE Transactions on Engineering Management.

[18] Liao, K., & Zhao, Z. (2022). Using blockchain to achieve decentralised privacy in IoT healthcare. IEEE Internet of Things Journal.

[19] Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., & Savage, S. (2013). A fistful of bitcoins: Characterising payments among men with no names. Proceedings of the 2013 conference on Internet measurement conference.

[20] NRC VAD Lexicon v2: Norms for Valence, Arousal, and Dominance for over 55k English Terms. Saif M. Mohammad. In arxiv preprint arXiv:2503.23547. 2025.

[21] Patel, V., & Zhang, Y. (2024). TRS: A novel transaction risk scoring system for blockchain analytics. IEEE Transactions on Information Forensics and Security.

[22] Qin, K., Zhou, L., Livshits, B., & Gervais, A. (2023). Attacking the DeFi ecosystem with flash loans for fun and profit. Financial Cryptography and Data Security.

[23] UN Panel of Experts. (2024). Report of the Panel of Experts established pursuant to resolution 1874 (2009).

[24] Wang, Y., Chen, K., & Zhang, F. (2024). Revisiting reentrancy attacks in Ethereum smart contracts. IEEE Symposium on Security and Privacy.

[25] Weber, M., Domeniconi, G., Chen, J., Weidele, D. K., Bellei, C., Robinson, T., & Leiserson, C. E. (2023). Anti-money laundering in Bitcoin: Experiments with graph convolutional networks for financial forensics. ACM SIGKDD Explorations Newsletter.

[26] Wu, J., Yuan, Q., Lin, D., You, W., Chen, W., Chen, Y., ... & Zhang, X. (2023). Who are the phishers? Phishing scam detection on Ethereum via network embedding. IEEE Transactions on Systems, Man and Cybernetics: System