International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 3, May-June-2026, ISSN (Online): 2395-566X

Beyond the Surface Web: An Analytical Study of Deep Web and Dark Web Threat Ecosystems

Deepa Barethiya, Himanshu Praveen Dethekar, Bhavesh Tembhurkar
Dept. of MCA
GHRCEM, Nagpur, Maharashtra, India

Abstract — The dark web constitutes a stratified, operationally sophisticated cybercrime ecosystem whose threat dynamics are shaped by layered anonymity infrastructure, AI-augmented criminal tooling, and resilient financial obfuscation mechanisms. While existing literature provides valuable but fragmented analysis of individual components, few studies integrate these elements within a unified analytical framework. This paper addresses that gap through a hybrid analytical survey approach, advancing four primary contributions: (1) a six-dimension taxonomic model differentiating surface web, deep web, and dark web environments; (2) a Five-Layer Dark Web Threat Ecosystem Model characterising the functional architecture of criminal infrastructure; (3) a structured capability taxonomy of AI-augmented criminal tools (Dark LLMs); and (4) a proposed Cyber Threat Intelligence (CTI) extraction pipeline for dark web environments. Drawing on peer-reviewed literature spanning 2020–2025, operational intelligence from Europol IOCTA, Chainalysis Crypto Crime Reports, and FBI IC3 data, and documented threat actor behaviour, the paper analyses ransomware-as-a-service dynamics, cryptocurrency financial obfuscation, law enforcement response limitations, and post-Tor architectural evolution. Persistent research gaps in multilingual CTI extraction, post-Tor forensic methodology, and AI-threat detection are identified, with a structured research agenda proposed.

Keywords— dark web threat intelligence; cyber threat intelligence; Dark LLMs; ransomware-as-a-service; anonymity networks; cryptocurrency laundering; threat ecosystem modelling; OSINT; post-Tor architectures; AI-augmented cybercrime.

I. INTRODUCTION

1. Motivation and Problem Context

The architecture of the internet is not uniform in accessibility, indexability, or anonymity design. Beneath the publicly accessible surface web lies a substantially larger domain of non-indexed content—the deep web—which includes private databases, passworded enterprise systems, and deliberately anonymised subsections engineered to resist conventional access and monitoring. The dark web, an intentional subset of this environment, is designed specifically to provide actor anonymity, content concealment, and resistance to both surveillance and attribution.

Despite broad acknowledgement in academic and policy discourse, the analytical treatment of the dark web in published research has remained predominantly descriptive: cataloguing what the dark web contains rather than modelling how it functions as an integrated operational system. This orientation leaves a consequential gap: an absence of analytical frameworks capable of capturing the ecosystem's structural dynamics, inter-component dependencies, and evolving threat characteristics.

The dark web of 2025 bears limited resemblance to the relatively static network of illicit marketplaces documented in early academic surveys. It constitutes a dynamically adaptive criminal infrastructure incorporating multiple anonymity network protocols, sophisticated identity management systems, cryptocurrency laundering pipelines of institutional scale, and—critically—artificial intelligence-augmented tooling that represents a qualitative escalation in criminal capability. Europol's Internet Organised Crime Threat Assessment (IOCTA) for 2024 confirmed that law enforcement takedowns have shortened the lifecycle of individual criminal sites while simultaneously accelerating the ecosystem's fragmentation and diversification [1]. This pattern—adaptive resilience rather than structural degradation under pressure—is analytically significant and inadequately theorised in existing literature.

2. Current Threat Escalation: The AI Dimension

Concurrent with this adaptive resilience, the integration of large language model (LLM) technologies into dark web criminal operations has introduced a capability escalation without precedent. Kela Cybersecurity documented that mentions of malicious AI tools on cybercriminal forums increased by 219% between 2023 and 2024 [2]. Tools including WormGPT, FraudGPT, and GhostGPT are actively marketed on dark web forums and Telegram channels without ethical constraints, enabling actors with limited technical expertise to generate convincing phishing content, functional malware code, and automated social engineering campaigns at scale [3]. This

lowering of the capability barrier has structural implications for the entire threat ecosystem, particularly at the operational layer.

3. Literature Gaps and Paper Contributions

Existing dark web survey literature exhibits several recurring limitations. Definitional frameworks remain inconsistent across publications, complicating cross-study comparison [11], [17]. Analytical models treat the dark web as a collection of discrete criminal phenomena rather than an integrated operational system [1]. The intersection of AI capability escalation with dark web criminal ecosystems has been documented in threat intelligence reports but not systematically synthesised academically [2], [4]. CTI extraction methodology from dark web sources is treated in a growing but fragmented body of work, with significant gaps in multilingual coverage and operational tooling integration [5], [6].

This paper proposes an integrated analytical framework to address these gaps, structured around four primary contributions: (1) a six-dimension taxonomy resolving existing definitional inconsistency; (2) a Five-Layer Dark Web Threat Ecosystem Model organising dark web criminal operations by functional layer; (3) a Dark LLM capability taxonomy classifying AI-augmented criminal tools; and (4) a five-stage CTI extraction pipeline designed for dark web intelligence environments. These contributions are advanced through a hybrid analytical survey methodology synthesising peer-reviewed literature, operational intelligence, and documented threat actor behaviour.

4. Paper Organisation

The remainder of this paper is organised as follows. Section II describes the research methodology. Section III proposes the six-dimension taxonomic framework. Section IV provides a comparative analysis of anonymity network architectures. Section V presents the Five-Layer Threat Ecosystem Model. Section VI characterises AI-augmented criminal ecosystems. Section VII proposes the CTI extraction pipeline. Sections VIII and IX analyse ransomware and cryptocurrency sub-ecosystems. Section X evaluates law enforcement and forensic limitations. Section XI addresses post-Tor architectures. Sections XII–XIV cover ethical considerations, research gaps, and study limitations respectively. Section XV concludes.

II. RESEARCH METHODOLOGY

1. Survey Approach and Analytical Synthesis

This paper employs a hybrid analytical survey methodology, combining systematic literature review with original framework development. Unlike a purely descriptive survey, this approach treats the synthesised literature as evidence for constructing analytical models that extend beyond the sum of their constituent sources. The methodological rationale is grounded in the recognised value of framework-synthesis approaches in cybersecurity research, where empirical data collection is constrained by the operational nature of the subject matter [11].

2. Source Selection and Literature Review Strategy

Academic literature was sourced from IEEE Xplore, ACM Digital Library, SpringerLink, Elsevier ScienceDirect, and arXiv, with a primary timeframe of 2020–2025. Search terms included combinations of: dark web, darknet, threat intelligence, cyber threat intelligence, ransomware-as-a-service, anonymity networks, cryptocurrency laundering, Dark LLMs, dark web forensics, and post-Tor architectures. Foundational technical papers published before 2020 (notably [8] on Tor architecture) were retained where no more recent primary source supersedes them.

Inclusion criteria required: (1) publication in a peer-reviewed venue or recognised operational intelligence body; (2) primary focus on technically relevant dark web, cybercrime, or CTI topics; (3) sufficient methodological transparency to permit quality assessment. Exclusion criteria eliminated non-peer-reviewed commercial blog posts, Wikipedia-derived sources, and studies whose primary framing was sociological or criminological without technical content. Threat intelligence reports from Europol [1], Chainalysis [12], [19], and FBI IC3 were treated as operational intelligence sources supporting analytical claims rather than as primary academic evidence.

3. Limitations of Threat Intelligence Evidence

Operational intelligence reports from commercial and law enforcement bodies provide uniquely valuable real-world evidence but carry inherent limitations: reporting bias toward detected and investigated incidents; potential for strategic framing by reporting organisations; and temporal lag between events and publication. Where statistics from such sources are cited, they are presented as documented observations rather than population-representative measurements.

III. TAXONOMY AND DEFINITIONAL FRAMEWORK

1. The Problem of Definitional Fragmentation

A persistent methodological challenge in dark web research is the absence of a stable, consistently applied definitional framework. Terms including ‘dark web,’ ‘darknet,’ ‘deep web,’ and ‘dark internet’ are used interchangeably across the literature with partially overlapping meanings. Ngo et al. define

the dark web as 'a subsection of the deep web that conventional search engines cannot index' [17]—a formulation shared broadly but operationalised inconsistently. Alahmari et al.'s 2024 systematic review identifies 'dark net' as sometimes referring specifically to encrypted overlay network infrastructure independent of the content layer, compounding the ambiguity [11]. This fragmentation affects the comparability of findings across independent studies and the design of monitoring systems that operationalise these distinctions.

2. Proposed Six-Dimension Taxonomy

This paper proposes a six-dimension classification matrix that differentiates surface web, deep web, and dark web environments across operationally relevant axes. The six dimensions are: (1) access mechanism; (2) content indexability; (3) user anonymity level; (4) primary content type; (5) dominant threat class; and (6) forensic traceability. Table I presents the complete taxonomy.

Table I. Six Dimension Classification Of Internet Tiers

Dimension	Surface Web	Deep Web	Dark Web
Access Mechanism	Standard browser; no authentication required	Standard browser + auth; direct URL or API credentials	Specialised client required (Tor, I2P, Freenet, Lokinet)
Content Indexability	Fully indexed by standard search engines	Not indexed; behind access controls or auth walls	Not indexed; hidden service or content-hash addressing
User Anonymity Level	Low; IP visible to server; ISP monitoring possible	Moderate; varies by service and institutional context	High by design; multi-hop routing; onion/garlic encryption
Primary Content Type	Public information, commercial services, social media	Private data, enterprise/academic proprietary content	Mixed: privacy-preserving use + criminal markets/forums
Dominant Threat Class	Phishing, credential harvesting, surface malware distribution	Database exposure, misconfigured APIs, insider threats	RaaS ecosystems, exploit markets, AI-augmented tooling
Forensic Traceability	High; standard digital forensic tools applicable	Moderate; requires authorised backend system access	Low by design; specialised protocols required [10]

3. Terminological Clarifications

For precision throughout this paper: 'Darknet' denotes the technical overlay network infrastructure (Tor, I2P, Freenet, Lokinet). 'Dark web' refers to the content layer accessible through darknet infrastructure, encompassing both legitimate privacy-preserving uses and criminal applications. 'Dark web ecosystem' denotes the structured system of interacting functional layers through which criminal operations are sustained on the dark web. 'Dark LLM' denotes a large language model—purpose-built or jailbroken—deployed for criminal purposes without ethical constraints [3], [4].

IV. COMPARATIVE ANALYSIS OF ANONYMITY NETWORK ARCHITECTURES

1. Tor: Architecture and Forensic Vulnerabilities

Tor (The Onion Router) remains the dominant anonymity network by user base and criminal adoption. It employs layered encryption across three-hop circuits of volunteer-operated relays, with traffic decrypted incrementally at each hop so no single relay possesses both the origin and destination of a communication [8]. Hidden service (".onion") addressing enables server-side anonymity. Documented vulnerabilities include exit node monitoring for clearnet-destined traffic, browser fingerprinting and JavaScript-based deanonymisation, and traffic correlation attacks by global passive adversaries. Law enforcement competence with Tor-based network analysis has increased substantially, motivating criminal migration toward alternative platforms [1].

2. I2P: Garlic Routing and Hidden Service De-anonymisation

I2P (Invisible Internet Project) employs garlic routing within a distributed hash table (DHT), bundling multiple encrypted streams into "garlic messages" and routing them peer-to-peer rather than through centralised relay directories. Eepsites (I2P-hosted services) provide persistent internal communication infrastructure preferred by criminal actors for backend operations. Wang et al. (2024) demonstrated a de-anonymisation attack against I2P hidden services via live behaviour alignment, revealing that I2P's anonymity guarantees are not absolute under active adversarial conditions [9]. This finding has direct forensic implications for law enforcement investigations targeting I2P-based criminal infrastructure.

3. Freenet, Lokinet, and Nym: Emerging and Specialist Protocols

Freenet operates as a decentralised, content-addressed distributed storage network providing strong resistance to content removal but exhibiting high latency and limited criminal adoption for interactive services. Lokinet implements a Low Latency Anonymous Routing Protocol (LLARP) that hybridises Tor's circuit-based architecture with I2P's DHT-based node discovery, incorporating blockchain-based service node incentives for routing bandwidth. Nym implements a mix network architecture employing packet shuffling and deliberate timing delays at intermediate mix nodes, providing strong resistance to traffic analysis attacks that Tor and I2P remain vulnerable to; however, Nym does not currently support hidden services, limiting its direct applicability as dark web hosting infrastructure.

Table Ii. Anonymity Protocol Comparison: Operational And Forensic Characteristics

Dimension	Tor	I2P	Freenet	Lokinet	Nym
Routing Model	Circuit/onion	Garlic + DHT	Content-hash storage	Hybrid LLARP	Mix network
Hidden Service Support	Yes (onion)	Yes (eepsites)	Yes (distributed)	Yes	No
Clearnet Access	Yes (exit nodes)	Limited	No	Yes	No
Latency Profile	Low-medium	Medium	High	Low-medium	High (deliberate)
Traffic Analysis Resistance	Moderate	Moderate	High (passive)	Moderate	High
Criminal Adoption	Primary platform	Specialist infra.	Minimal (legacy)	Emerging	Nascent
LE Capability Maturity	High	Medium	Low	Low	Minimal

The strategic implication is that criminal actors operate across multiple protocols, selecting them by operational requirement: Tor for user-facing services, I2P for persistent backend communications, and emerging platforms for actors seeking additional metadata protection against maturing law enforcement capabilities.

V. FIVE-LAYER DARK WEB THREAT ECOSYSTEM MODEL

1. Framework Rationale

Existing analytical frameworks organise dark web criminal phenomena categorically—by crime type, actor type, or technical mechanism. While each lens provides partial insight, none models the ecosystem's functional structure or accounts for its documented resilience under law enforcement pressure. The Five-Layer Dark Web Threat Ecosystem Model organises the ecosystem by operational function—what each component does in sustaining criminal operations—rather than by crime category. This functional orientation enables more precise analysis of disruption propagation and intelligence collection opportunity identification. The model is empirically grounded in the observed post-Operation Cronos LockBit disruption: seizure of the dominant RaaS platform produced affiliate migration rather than ecosystem collapse [1], a dynamic explicable through the model's inter-layer dependency structure.

2. Layer Definitions

Layer 1 — Infrastructure: Technical hosting and connectivity resources including bulletproof hosting providers (jurisdictions shielding them from law enforcement cooperation), mirror network architectures ensuring content resilience against individual platform seizures, and clearnet-to-darknet bridge services. This layer is the most capital-intensive ecosystem component, making it a natural concentration point and high-value law enforcement target [16].

Layer 2 — Identity: The system governing pseudonymous identity establishment, management, and verification. On dark web markets and forums, identity constitutes a structured reputation asset—PGP key fingerprints, transaction history, reputation scores, and vouching relationships—that carries substantial economic value and persists across platform disruptions [16]. Identity asset preservation is a primary mechanism of ecosystem resilience.

Layer 3 — Marketplace: Transactional infrastructure for commodity and service exchange, including drug markets, weapons markets, stolen credential markets, exploit markets, access broker markets, and criminal-as-a-service platforms (DDoS-for-hire, Malware-as-a-Service). The marketplace layer is the most externally visible and consequently the most frequently targeted by law enforcement; however, its disruption without simultaneous Layer 1 and 2 disruption produces displacement rather than elimination.

Layer 4 — Financial: Cryptocurrency selection, transaction obfuscation, and fiat conversion mechanisms constituting the economic foundation of all other layers. The evolution of this layer—from Bitcoin-dominated transactions toward stablecoin intermediation, privacy coin adoption, and DeFi protocol exploitation [12]—represents one of the most technically dynamic dimensions of the ecosystem.

Layer 5 — Operational: Human and organisational structures through which criminal actors coordinate and execute operations. This includes RaaS provider-affiliate hierarchies, initial access broker markets, dark web forum-based knowledge transfer, and the AI-augmented tooling ecosystem documented in Section VI. The operational layer's capability ceiling has been qualitatively elevated by Dark LLM integration.

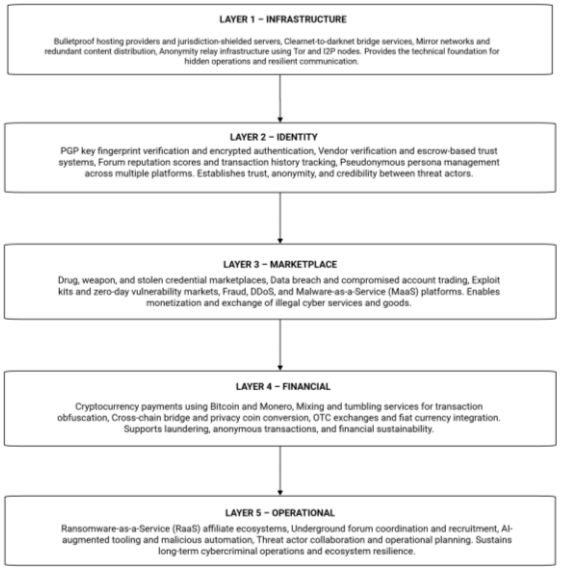

Fig. 1. Five-Layer Dark Web Threat Ecosystem Model.

3. Inter-Layer Dependencies and Resilience Mechanisms

The model exposes the inter-layer dependencies accounting for the ecosystem's documented resilience. Layer 3 (Marketplace) depends on Layers 1 and 2 for hosting and trust infrastructure; when a marketplace is seized, infrastructure and identity assets remain intact, enabling rapid reconstitution on alternative platforms. Layer 4 (Financial) integrity is a precondition for the economic viability of all other layers; sanctions on mixing services (Tornado Cash, 2022; Huione Guarantee, 2024) produce stress but are countered through technical adaptation. The operational implication for law enforcement is that multi-layer simultaneous disruption is necessary for structural ecosystem impact—a conclusion supported by the post-Operation Cronos evidence [1].

VI. AI-AUGMENTED CRIMINAL ECOSYSTEMS AND DARK LLMS

1. Emergence and Categories of Dark LLM Tools

The integration of large language model technologies into dark web criminal operations represents one of the most consequential developments in the ecosystem's capability trajectory since 2023. Unlike prior generations of commoditised criminal tooling—which reduced the expertise barrier for specific discrete tasks—AI-augmented tools reduce the barrier across multiple simultaneous operational requirements, from victim targeting and social engineering to malware generation and operational security guidance.

Dark LLMs comprise three analytically distinct sub-categories. First, purpose-built criminal AI tools: WormGPT, first documented by SlashNext researchers in mid-2023, is based on the open-source GPTJ language model and advertised explicitly without ethical constraints for business email compromise (BEC) template generation, malware code production, and phishing content creation [3]. FraudGPT, appearing in approximately the same period, operates on a subscription commercial model (approximately €60–€200 per month) and advertises capabilities including spear-phishing generation, scam page development, and CVE identification [3]. Second, jailbroken open-source LLMs: open-source models including Llama 2, Mistral, and derivative architectures modified through fine-tuning or adversarial prompt engineering to remove alignment constraints. Finkelstein and Rokach (2025) documented a universal jailbreak attack effective against multiple state-of-the-art models, demonstrating that alignment safety in current models is insufficiently robust against adversarial circumvention [4]. Third, repurposed research models: DarkBERT, developed by S2W for legitimate dark web cybercrime research, is a BERT-family model trained on dark web corpus data; its competency in criminal argot and darknet-specific semantics has attracted criminal interest, with claims of integration into FraudGPT derivatives documented in threat intelligence reports [7].

2. Adoption Scale and Jailbreak Methodology

Kela Cybersecurity documented that mentions of malicious AI tools on cybercriminal forums increased by 219% between 2023 and 2024 [2]. This figure reflects not merely adoption but discourse normalisation—AI tooling integrated into criminal operational planning as an expected rather than exceptional capability. The proliferation of over 15,800 publicly available

LLM model checkpoints on platforms such as Hugging Face by mid-2023 means that base models for criminal AI tool development are broadly accessible [4]. Jailbreak techniques documented by Finkelstein and Rokach include: many-shot jailbreaking (providing multiple harmful examples to condition model outputs), system prompt injection, role-playing persona manipulation, and fine-tuning on harmful datasets to remove alignment during training [4]. The practical accessibility of these techniques to criminal actors with limited machine learning expertise represents a structural change in the threat landscape.

Table III. Dark Llm Capability Matrix

Tool	Base Model	Primary Capabilities	Distribution	Pricing	Current Status
WormGPT	GPTJ (open-source)	BEC emails, phishing, malware code, payload creation	Dark web forums, Telegram	€200–€5,000/yr	Operational; iterating versions
FraudGPT	Proprietary (undisclosed)	Spear-phishing, scam pages, CVE exploit, carding	Dark web forums, Telegram	€60–€200/mo	Operational; active forum discussion
GhostGPT	Jailbroken Llama/Mistral	Malware generation, exploit scripting, obfuscation	Telegram, underground forums	Low/freemium	Active; ephemeral variants
EvilGPT / WolfGPT	Jailbroken / Python-based	Harmful content, attack scripts, phishing automation	Telegram, underground forums	Variable	Limited documented use; emerging
DarkBERT (misuse)	BERT (dark web corpus)	Criminal argot NLU, darknet content classification	Research leak/adaptation (claimed)	N/A	Academic concern; limited direct ops.
OSS Jailbroken Variants	Llama 2 / Mistral / others	Variable per jailbreak sophistication	Hugging Face, direct model repos.	Free (base) + technique cost	Distributed; difficult to enumerate

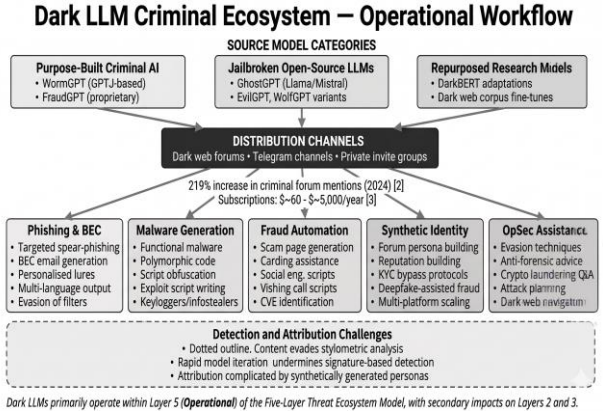

Fig. 2. Dark LLM Criminal Ecosystem Workflow

3. Detection and Attribution Challenges

AI-generated content complicates traditional detection and attribution methodologies. AI-produced phishing emails exhibit reduced stylometric signatures that distinguish non-native speaker output, undermining language-based filtering. AI-generated malware code may produce functionally novel variants that diverge from known malware families, reducing signature-based detection efficacy. AI-constructed synthetic identities complicate cross-platform threat actor correlation by providing multiple AI-coherent but unrelated persona sets [4]. These challenges represent active research problems at the intersection of adversarial machine learning and threat intelligence.

VII. CYBER THREAT INTELLIGENCE EXTRACTION FRAMEWORK

1. Existing CTI Methodology and Dark Web-Specific Limitations

Cyber Threat Intelligence is operationally defined as evidence-based knowledge about threats—including tactics, techniques, and procedures (TTPs), indicators of compromise (IoCs), threat actor characteristics, and attack prerequisites—enabling informed security decision-making [15]. The dark web is a significant CTI source whose forums contain discussions of emerging vulnerabilities, tool development, and TTPs that frequently precede observable attacks on monitored infrastructure.

Existing CTI extraction methodologies exhibit several documented limitations in dark web contexts. Kühn et al. (2024), in the most comprehensive published evaluation of dark web CTI extraction, identified that current approaches are limited in scalability, restricted primarily to English-language Tor sources, and lack integration with live security tooling [5]. The MAD-CTI framework (IEEE Access, 2025) advances multi-agent LLM-based topic modelling for dark web forum analysis but explicitly acknowledges limitations in cross-lingual generalisation [6]. DarkBERT (S2W, 2023) demonstrates that transformer models trained specifically on dark web corpus data substantially outperform general-purpose NLP models on criminal argot classification and darknet-specific semantic understanding [7]—but published multilingual dark web datasets remain scarce.

2. Proposed Five-Stage CTI Extraction Pipeline

This paper proposes a five-stage CTI extraction pipeline specifically designed for dark web intelligence environments, addressing identified gaps in multilingual coverage, IoC-to-threat-actor correlation, and operational tooling integration.

Stage 1 — Authenticated Dark Web Collection: Tor-aware crawler deployment in isolated virtualised environments with probabilistic inter-request timing to avoid fingerprinting, structured coverage of known forum and leak site addresses, and authentication management for registered forum access. Ethical and legal constraints apply (Section XII).

Stage 2 — Preprocessing and Deduplication: Language identification, encoding normalisation, OCR extraction for image-embedded text, deduplication against previously collected corpora, and initial relevance filtering via keyword and regular expression pattern matching.

Stage 3 — NLP Classification and Entity Extraction: Fine-tuned transformer models (DarkBERT adaptations, BERT variants) for content categorisation across threat-relevant classes; named entity recognition for IoCs (IP addresses, domains, file hashes, cryptocurrency wallet addresses), CVE references, and threat actor handles; multilingual processing pipeline for non-English forum coverage.

Stage 4 — Threat Actor Profiling and Correlation: Cross-platform handle correlation through stylometric analysis, PGP key matching, and operational pattern recognition; cryptocurrency address clustering; threat actor profile construction mapping handles to capabilities, infrastructure preferences, and targeting patterns.

Stage 5 — Structured Output and Dissemination: STIX 2.1 formatting for interoperability with SIEM platforms and threat intelligence platforms (MISP, OpenCTI, Anomali) [14]; TAXII 2.1 transport for automated machine-to-machine intelligence sharing; mapping of extracted TTPs to MITRE ATT&CK framework taxonomy [15].

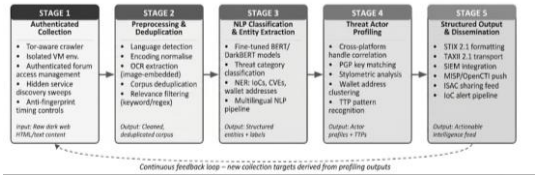

Fig. 3. CTI Extraction Pipeline Architecture. Five-stage workflow from authenticated dark web collection through structured STIX/TAXII output, with continuous feedback from Stage 4 profiling to Stage 1 collection targeting. Designed to address documented gaps in multilingual coverage and operational tooling integration.

VIII. RANSOMWARE-AS-A-SERVICE ECOSYSTEM DYNAMICS

1. The RaaS Operational Model

Ransomware-as-a-Service represents the institutionalisation of ransomware as a commercial platform, separating the functions of ransomware development, affiliate recruitment, and attack execution into a structured commercial relationship. RaaS developers maintain the ransomware codebase, payment and negotiation infrastructure, and data leak site; affiliates rent access in exchange for a percentage of ransom proceeds, typically 20–30% of received payments [1]. This model produces a criminal economy with franchise characteristics: low affiliate entry barriers, brand value for established operators, and functional specialisation increasing overall ecosystem efficiency.

The double extortion model—encrypting victim data while simultaneously exfiltrating it and threatening publication on dark web leak sites—has become the operational baseline. Triple extortion adds targeted contact with victims' customers, partners, or regulators, maximising coercive pressure beyond the direct victim. These tactical evolutions reflect the RaaS ecosystem's capacity for operational innovation within its commercial structure.

2. Fragmentation Dynamics and 2024–2025 Evidence

Operation Cronos (February 2024) seized LockBit's administrative infrastructure and unmasked key operators, disrupting what had been the most dominant RaaS operation globally [1]. The anticipated systemic impact—significant reduction in ransomware activity—did not materialise. Instead, disruption produced rapid reallocation: LockBit affiliates migrated primarily to RansomHub (founded February 2024, days before the Cronos disruption) and to Akira and Qilin groups. RansomHub became the most prolific leak group operator of 2024 despite its recent founding, exemplifying the ecosystem's adaptive capacity under pressure.

By 2025, 138 distinct ransomware groups claimed victims on dark web leak sites, compared to 98 in 2024—a 41% increase in active operations [13]. Total ransomware payments in 2024 declined 35% year-over-year from the 2023 record to approximately $813.55 million [19], primarily because growing victim backup capability reduced payment willingness rather than because attack volume diminished. The ecosystem produced 7,307 documented victims in 2025 across these 138 groups [13]. This fragmentation pattern—more groups, lower individual market share, dominant operators disrupted and replaced—is precisely the trajectory that the IOCTA 2024 analysis anticipated [1] and that the Five-Layer Model accounts for through its inter-layer dependency analysis.

IX. CRYPTOCURRENCY LAUNDERING AND FINANCIAL OBFUSCATION

1. Currency Selection and Ecosystem Composition

The financial layer of the dark web threat ecosystem has undergone significant structural evolution between 2020 and 2025. Chainalysis's 2025 Crypto Crime Report documented a seismic shift in illicit cryptocurrency composition: stablecoins (primarily USDT and USDC) accounted for 63% of all illicit crypto transactions by 2024, driven by their liquidity, USD-equivalence, and faster transaction settlement [12]. However, this generalisation requires qualification: ransomware payments and dark web marketplace transactions specifically remain Bitcoin-dominated, reflecting the liquidity requirements and established infrastructure of these use cases. Monero (XMR) has gained documented traction in dark net market ecosystems, where its ring signature, stealth address, and RingCT features provide privacy guarantees that Bitcoin's transparent ledger cannot.

2. Layering Mechanisms: Chain-Hopping, DeFi, and Forensic Limitations

The layering stage of cryptocurrency money laundering has grown substantially more complex. Traditional mixing services faced significant regulatory pressure following the US Treasury's sanctioning of Tornado Cash in August 2022; however, their closure diversified rather than eliminated mixing activity. TRM Labs documented chain-hopping—rapidly moving assets across multiple blockchains through cross-chain bridge protocols—in 68% of laundering schemes analysed in 2024 [18]. Decentralised Finance (DeFi) protocols enable token swaps without centralised exchange KYC requirements and are increasingly integrated into layering workflows. North Korean state-sponsored actors (Lazarus Group) stole $1.34 billion in cryptocurrency in 2024, representing 61% of all cryptocurrency stolen globally, employing multi-stage chain-hopping combining mixer services, DeFi protocols, and OTC brokers [12].

The forensic implications are substantial. TRM Labs reported 30% fewer successful tracings when privacy coins and mixers were combined in 2024, despite improved analytical tooling [18]. Full de-anonymisation of Monero transactions—given ring signature obfuscation of transaction inputs and stealth address obfuscation of outputs—remains beyond the capability of current blockchain forensic tools, representing a persistent asymmetry between criminal financial sophistication and investigative capability.

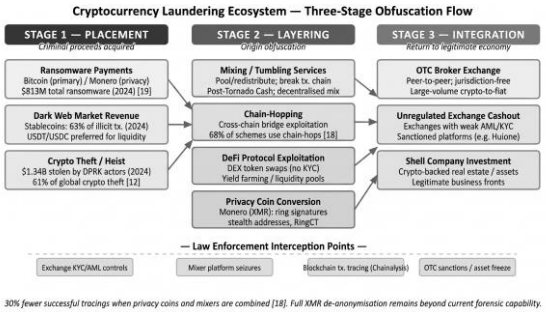

Fig. 4. Cryptocurrency Laundering and Obfuscation Flow. Three-stage process (Placement → Layering → Integration) showing specific techniques at each stage, with law enforcement interception points annotated. Chain-hopping documented in 68% of schemes; 30% fewer successful tracings when privacy coins and mixers are combined [18].

X. LAW ENFORCEMENT RESPONSE AND DIGITAL FORENSIC LIMITATIONS

1. Operational Record: Disruption Without Structural Dismantlement

The 2020–2025 period produced the most sustained and internationally coordinated law enforcement action against dark web criminal infrastructure on record. Table IV summarises key operations and their documented outcomes.

TABLE IV. Selected Law Enforcement Operations Against Dark Web Criminal Infrastructure, 2023 – 2025

Operation	Year	Primary Target	Key Outcomes	Displacement Effect
SpecTor	2023	Dark web drug markets	288 arrests; $53.4M seized	Rapid migration to remaining markets
Genesis Market Takedown	2023	Credential market	119 arrests; 200+ searches	Competitor markets absorbed traffic
Operation Cronos	2024	LockBit RaaS infrastructure	Infrastructure seized; Khoroshev unmasked	RansomHub absorbed LockBit affiliates
Operation Endgame	2024	Malware droppers / botnets	4 arrests; 100+ servers; €69M seized	Partial disruption; operator rebranding
Operation RapTOR	2025	Dark web drug vendor networks	270 arrests; $200M+ seized; 2t drugs	Investigation ongoing; long-term TBD [20]

The recurring pattern across these operations is displacement rather than elimination. Individual operators are disrupted, assets are seized, and prosecutorial outcomes are produced; however, structural degradation of the ecosystem does not follow. This paradox—where successful individual operations may accelerate ecosystem diversification rather than contracting it—is explicable through the Five-Layer Model: disruption confined to Layer 3 components without simultaneous Layer 1 and 2 disruption enables reconstitution around alternative platforms while preserving underlying infrastructure, identity assets, and financial flows [1].

2. Jurisdictional Fragmentation and Forensic Protocol Limitations

A structural challenge in law enforcement response is jurisdictional fragmentation: operators, infrastructure, affiliates, and victims are frequently distributed across multiple legal jurisdictions, requiring mutual legal assistance treaty (MLAT) processes operating on timescales mismatched to criminal operational agility. Bulletproof hosting providers—a Layer 1 component—are frequently located in jurisdictions with limited cooperation frameworks, creating persistent safe harbour for infrastructure components.

The D2WFP (Deep and Dark Web Forensic Protocol), proposed by Ghanem et al. (2023), provides a structured protocol for identifying, extracting, and analysing dark web browsing activities from end-user devices, combining Tor-specific artefact extraction with conventional digital forensic methodology [10]. While D2WFP represents a methodological advance over ad hoc approaches, it is explicitly scoped to host-device investigation and does not address network-level dark web forensics, IPFS-hosted service investigation, or the forensic challenges of decentralised anonymity architectures.

XI. POST-TOR AND DECENTRALIZED DARKNET ARCHITECTURES

1. Evidence of Criminal Infrastructure Migration

There is documented evidence of criminal infrastructure migration toward alternative and complementary anonymity platforms, motivated by concern about Tor's forensic vulnerabilities and law enforcement's demonstrated analytical competence [1]. I2P has been increasingly adopted for persistent backend criminal communication infrastructure, with its garlic routing and DHT architecture providing reduced exposure to the exit node monitoring techniques most effective against Tor. Wang et al.'s (2024) de-anonymisation research [9] represents the state of the art in I2P forensic investigation and suggests a growing but still limited law enforcement capability for this platform.

IPFS (InterPlanetary File System), a content-addressed distributed file storage protocol, has emerged as a hosting platform for criminal actors seeking takedown-resilience. Because IPFS content is addressed by cryptographic hash and stored redundantly across participating nodes, traditional law enforcement takedown procedures—requiring identification and seizure of hosting infrastructure—are not directly applicable. Documented criminal uses include hosting of phishing sites resilient to takedown, ransomware payload distribution, and experimental dark web marketplace front-ends.

2. Forensic and Intelligence Implications

Post-Tor architectural evolution has direct implications for the CTI pipeline proposed in Section VII. Extension of the pipeline to I2P eepsites, IPFS-hosted content, and Lokinet services requires modifications at the Collection stage: I2P-aware crawlers, IPFS gateway access mechanisms, and content

discovery approaches calibrated to different addressing systems. The forensic challenge is more fundamental: for IPFS-hosted criminal infrastructure, forensic approaches must rely on identifying clearnet or darknet access points through which content is served rather than the hosting infrastructure itself. The portfolio diversification model—criminal actors maintaining Tor for user-facing services, I2P for backend communications, and IPFS for resilient hosting—requires multi-platform monitoring approaches that no current published framework fully addresses.

XII. ETHICAL AND LEGAL CONSIDERATIONS

1. Passive Observation Versus Active Participation

Dark web research necessitates careful ethical boundary-setting. Passive observation—monitoring publicly accessible dark web content without authentication or interaction—is generally considered ethically and legally permissible across most jurisdictions, although applicable legal frameworks vary. Active participation—creating accounts on criminal markets, engaging in transactions, or providing content—constitutes a fundamental ethical and legal boundary that responsible research must not cross. This paper's analysis is grounded exclusively in published academic literature, public operational intelligence reports, and documented threat actor behaviour; no primary dark web data collection was performed.

2. Jurisdictional Legal Variation and Data Handling

The legal status of dark web access for research purposes varies across jurisdictions. In the United Kingdom, the Computer Misuse Act 1990 creates potential liability for unauthorised access to computer systems even where the research intent is benign. In the United States, the Computer Fraud and Abuse Act creates similar uncertainties. Researchers engaging in dark web data collection must obtain appropriate institutional ethical approvals and legal guidance specific to their jurisdiction and methodology. GDPR and equivalent privacy regulations impose additional constraints on the handling, storage, and retention of personal data collected from dark web sources, including data concerning victims whose information may appear in breach datasets.

3. Handling of Illicit Data and Non-Engagement Principles

Research involving dark web content may encounter data derived from criminal acts: stolen credentials, breach datasets, exploited payment card information. Responsible handling requires: anonymisation or hashing of personal identifiers before any processing; non-retention beyond analytical necessity; non-dissemination in forms that could enable further harm; and coordination with relevant law enforcement bodies where serious ongoing criminal activity is discovered. The principle of non-engagement—researchers must not purchase, utilise, or facilitate any criminal products or services encountered during dark web research—is a foundational ethical constraint.

XIII. RESEARCH GAPS AND FUTURE DIRECTIONS

1. Multilingual CTI Extraction

The majority of published dark web CTI research concentrates on English-language Tor sources. Russian, Chinese, Arabic, and Portuguese-language cybercriminal forums—which represent substantial shares of criminal intelligence activity—receive substantially less analytical coverage [5]. The development of multilingual dark web corpora, fine-tuned multilingual NER models, and cross-lingual IoC extraction pipelines represents an open and high-value research problem. Kühn et al. explicitly identify multilingual support as a primary limitation of current methodologies [5].

2. Post-Tor Forensic Methodology

The D2WFP protocol [10] represents the state of the art for host-device dark web forensics but does not address investigation of I2P eepsites, IPFS-hosted content, or Lokinet-based infrastructure. The development of comprehensive forensic methodology for these platforms—including network-level investigation techniques, content discovery approaches, and attribution methodologies not dependent on server identification—represents an open research problem of growing practical urgency as criminal infrastructure migration accelerates.

3. AI-Generated Threat Content Detection

The detection of AI-generated malicious content—phishing emails, malware code, synthetic personas—is an active and underspecified research problem. Current detection approaches rely on stylometric analysis and statistical anomaly detection that may not generalise to rapidly evolving AI-generated content. Research into robust, adversarially-stable detection mechanisms for AI-generated criminal content represents a high-priority future direction, given the 219% increase in Dark LLM forum mentions in 2024 [2] and the documented capability escalation documented in [4].

4. Automated Cross-Platform Threat Actor Correlation

Cross-platform threat actor correlation—linking the same criminal actor across multiple dark web forums, marketplaces, and communication channels—currently depends on manually

intensive analysis. The MAD-CTI framework [6] represents progress toward automated multi-agent correlation, but scalable, accurate automated correlation remains an open problem, particularly in the presence of AI-generated synthetic personas that can produce multiple statistically distinct but coherent identity sets.

5. Blockchain Intelligence Integration and RaaS Economic Modelling

The integration of blockchain forensic intelligence with dark web CTI pipelines—enabling automated correlation of dark web threat actor identities with known cryptocurrency wallet clusters—has been advocated but not implemented at scale [12], [18]. Additionally, economic modelling of RaaS ecosystem dynamics following major law enforcement disruptions—predicting which disruption strategies produce the greatest long-term structural impact—remains insufficiently developed to guide strategic law enforcement resource allocation.

Limitations

This study exhibits several limitations that bound the generalisability of its findings. First, private and invitation-only dark web forums—which may host the most operationally significant criminal intelligence activity—are inaccessible to public academic research, creating a systematic underrepresentation of high-sophistication criminal communications in the literature underpinning this analysis. Second, the English-language bias of published academic literature introduces a coverage gap for non-English dark web activity; Russian and Chinese criminal forums, in particular, receive substantially less analytical treatment despite their documented significance. Third, the dark web ecosystem evolves rapidly; statistics and operational intelligence cited from 2024–2025 sources reflect a snapshot that may be superseded by the time of reading. Fourth, reliance on public operational intelligence reports from commercial and law enforcement bodies introduces reporting bias: documented incidents are not population-representative, and organisations may strategically frame disclosures. Fifth, the Five-Layer Model and CTI pipeline proposed here are conceptual frameworks; their empirical validation against controlled datasets or operational deployments remains future work.

XIV. CONCLUSION

This paper has presented an integrated analytical framework for the dark web as a structured cybercrime ecosystem, organised around four primary contributions. The six-dimension taxonomic model (Section III) addresses a persistent definitional fragmentation in the literature, providing a stable analytical foundation. The Five-Layer Threat Ecosystem Model (Section V) organises dark web criminal operations by functional layer, enabling structural analysis of ecosystem resilience, disruption propagation, and intelligence collection opportunity identification. The Dark LLM capability taxonomy (Section VI) classifies AI-augmented criminal tools into an analytically coherent structure, grounded in documented threat intelligence. The five-stage CTI extraction pipeline (Section VII) proposes an operationally relevant methodology for dark web intelligence, addressing identified gaps in multilingual coverage and operational tooling integration.

The central analytical thesis—that the dark web constitutes a stratified, operationally sophisticated ecosystem whose threat dynamics are increasingly shaped by AI-augmented tooling, resilient financial obfuscation, and architectural diversification beyond conventional Tor infrastructure—is supported by convergent evidence from peer-reviewed literature, Europol IOCTA operational intelligence, Chainalysis financial data, and documented threat actor behaviour across 2020–2025. The ecosystem's adaptive resilience under law enforcement pressure, exemplified by the post-Operation Cronos fragmentation dynamics, reflects the inter-layer dependency structure modelled in Section V and argues for multi-layer, coordinated disruption strategies rather than platform-level interventions alone.

For practitioners, the CTI pipeline and threat ecosystem model provide actionable analytical structures for dark web monitoring programmes. For researchers, the identified gaps in multilingual CTI extraction, post-Tor forensic methodology, AI-threat detection, and automated threat actor correlation represent a structured agenda for work that would substantively advance both academic understanding and operational capability.

REFERENCES

1. Europol, "Internet Organised Crime Threat Assessment (IOCTA) 2024," European Union Agency for Law Enforcement Cooperation, The Hague, Netherlands, 2024.
2. Kela Cybersecurity, "AI in the Underground: Malicious LLM Adoption Trends 2024," Kela Research Report, Tel Aviv, Israel, 2025.
3. LevelBlue (AT&T Cybersecurity) / SlashNext, "WormGPT and FraudGPT: The Rise of Malicious LLMs," SpiderLabs Research Blog, 2023.
4. M. Finkelstein and L. Rokach, "Dark LLMs: The Growing Threat of Unaligned AI Models," arXiv preprint

arXiv:2505.10066, Ben-Gurion University of the Negev, 2025.
5. P. Kühn, K. Wittorf, and C. Reuter, "Navigating the Shadows: Manual and Semi-Automated Evaluation of the Dark Web for Cyber Threat Intelligence," IEEE Access, vol. 12, pp. 118903–118922, 2024.
6. A. Sharma, P. Gupta, R. Sharma, and A. K. Singh, "MAD-CTI: Cyber Threat Intelligence Analysis of the Dark Web Using a Multi-Agent Framework," IEEE Access, 2025. doi: 10.1109/ACCESS.2025.10908603.
7. S2W Inc., "DarkBERT: A Language Model for the Dark Side of the Internet," in Proc. 61st Annual Meeting of the Association for Computational Linguistics (ACL), Toronto, Canada, 2023.
8. R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The Second-Generation Onion Router," in Proc. 13th USENIX Security Symposium, San Diego, CA, USA, 2004, pp. 303–320.
9. H. Wang, Z. Chen, J. Li, Y. Liu, and X. Wang, "Time Will Tell: Large-Scale De-Anonymization of Hidden I2P Services via Live Behavior Alignment," arXiv preprint arXiv:2512.15510, 2024.
10. M. C. Ghanem, P. Mulvihill, K. Ouazzane, R. Djemai, and D. Dunsin, "D2WFP: A Novel Protocol for Forensically Identifying, Extracting, and Analysing Deep and Dark Web Browsing Activities," arXiv preprint arXiv:2309.05537, 2023.
11. M. Alahmari, F. Alahmari, and A. Alotaibi, "Weaponization of the Growing Cybercrimes Inside the Dark Net: The Question of Detection and Application," Big Data and Cognitive Computing, vol. 8, no. 8, p. 91, MDPI, 2024.
12. Chainalysis, "2025 Crypto Crime Report," Chainalysis Inc., New York, NY, USA, 2025.
13. BreachSense, "Ransomware in 2025: 7,307 Victims Across 138 Groups," Annual Ransomware Report, 2026.
14. OASIS Open, "STIX Version 2.1 Specification," OASIS Committee Specification, Jun. 2021. [Online]. Available: https://docs.oasis-open.org/cti/stix/v2.1/
15. MITRE Corporation, "MITRE ATT&CK Framework for Enterprise," MITRE, McLean, VA, USA, 2024. [Online]. Available: https://attack.mitre.org
16. J. Décary-Hétu and B. Dupont, "Reputation in a Dark Network of Online Criminals," Global Crime, vol. 14, no. 2–3, pp. 175–196, Taylor & Francis, 2013.
17. F. T. Ngo, C. Marcum, and S. Belshaw, "The Dark Web: What Is It, How to Access It, and Why We Need to Study It," Criminal Justice Review, vol. 48, no. 3, pp. 305–321, SAGE Publications, 2023.
18. TRM Labs, "Blockchain Forensics and Illicit Transactions: 2024 Annual Review," TRM Labs Research, San Francisco, CA, USA, 2024.
19. Chainalysis, "Crypto Ransomware 2025: 35.82% YoY Decrease in Ransomware Payments," Chainalysis Blog, 2025. [Online]. Available: https://www.chainalysis.com
20. ICE Homeland Security Investigations, "Operation RapTOR: Global Darknet Crackdown Results in 270 Arrests, $200M+ Seized," Press Release, U.S. Department of Homeland Security, May 2025.