IDTrace: A Deterministic Digital Footprint Risk Assessment Platform

K.Tamil Selvi [1]
Assistant Professor,
Department of Computer Science & Engineering,
Adhiyamaan College of Engineering
(An Autonomous Institution),
Hosur, India

Thirunavukarasan Y [1]
Yokeshraj S[2]
Tamilarasan J[3]
UG Scholars,
Department of Computer Science & Engineering,
Adhiyamaan College of Engineering,
(An Autonomous Institution),
Hosur, India

***Abstract*—The widespread adoption of digital identities across numerous online platforms has significantly heightened individuals' vulnerability to credential exposure and large-scale data breaches. Current breach notification mechanisms remain largely siloed, offering piecemeal alerts without any meaningful interpretation of cumulative risk. This paper introduces IDTrace, a deterministic framework for quantifying digital footprint risk, which consolidates multi-source open-source intelligence (OSINT) and produces a transparent, explainable risk score on a 0–100 scale through severity-weighted deductions and a critical breach ceiling mechanism. To improve throughput and ensure resilience against failure, the system leverages concurrent intelligence retrieval through parallel execution pipelines. Empirical results demonstrate a 37.5% decrease in scan latency via asynchronous processing, linear scalability with respect to computational load, and perfectly consistent deterministic outputs with zero variance across repeated evaluations. A comparative study highlights IDTrace's advantages in interpretability and holistic exposure modeling over conventional breach detection solutions and probabilistic risk scoring methods. Collectively, IDTrace offers a structured, transparent, and user-centered foundation for cybersecurity risk assessment in an increasingly interconnected digital landscape**

Keywords—Digital Footprint; OSINT Aggregation; Deterministic Risk Scoring; Cybersecurity Analytics; Exposure Intelligence; Identity Risk Modeling; Explainable Security Systems; Parallel Processing.

I.INTRODUCTION

The ongoing digital transformation of modern society has led individuals to maintain numerous interconnected accounts spanning diverse and heterogeneous platforms. Every digital interaction contributes incrementally to the formation of a persistent and growing digital footprint. While such pervasive connectivity offers undeniable convenience and broader accessibility, it simultaneously exposes users to an expanding array of cybersecurity threats, including credential stuffing attacks, phishing campaigns, identity theft, and widespread data breaches.

Contemporary breach detection platforms predominantly rely on a lookup-centric model, in which a user's email address or username is queried against known repositories of compromised data. Although this approach succeeds in notifying users of past exposure incidents, it falls short of offering a coherent or structured understanding of cumulative risk. Not all data exposures carry equivalent consequences; for instance, the compromise of a plaintext password or a government-issued identifier poses substantially greater harm than the mere leakage of a username. Despite this disparity, conventional systems neither quantify such distinctions nor account for the amplifying effects that arise from cross-platform exposure.

This absence of structured risk quantification produces a critical gap between simple breach awareness and meaningful, actionable cybersecurity management. To bridge this divide, the present work introduces IDTrace, a deterministic digital footprint risk assessment framework engineered to consolidate heterogeneous OSINT

intelligence and transform fragmented exposure events into a measurable and interpretable risk score.

The principal contributions of this work are outlined as follows:

- A deterministic model for quantifying cumulative digital exposure risk.
- A severity-weighted, data-sensitive penalty framework for breach classification.
- A critical breach cap mechanism designed to prevent artificial score inflation.
- A parallel, multi-source OSINT aggregation architecture for comprehensive intelligence gathering.
- An AI-assisted remediation module operating within clearly defined ethical boundaries.

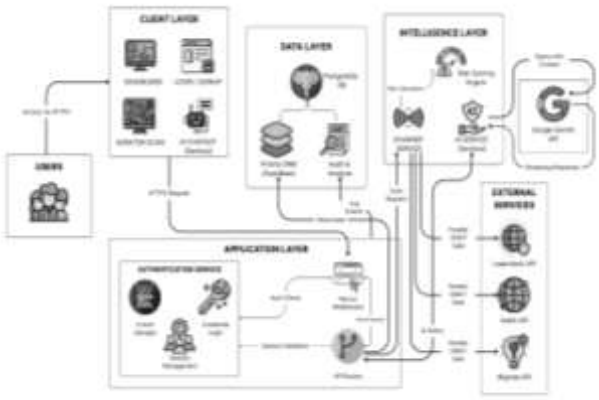

Figure 1 : IDTrace Orchestrational Workflow

II.LITERATURE REVIEW: RELATED WORK AND CRITIQUE

The design and development of IDTrace draw upon interdisciplinary research encompassing digital footprint intelligence, OSINT aggregation, structured risk modeling, AI governance, and established cybersecurity standards. The following section critically reviews relevant prior works and identifies the gaps that motivate the proposed framework.

A. OSINT Aggregation and Dark Web Intelligence

Singh, Pilli, and Laxmi (2024) presented a dark web surveillance framework leveraging multi-source OSINT correlation to uncover compromised identity evidence. Their work underscores the value of structured intelligence gathering through modular adapters capable of interfacing with heterogeneous data sources, demonstrating the effectiveness of correlating publicly

accessible intelligence streams to reconstruct identity exposure patterns.

While this approach meaningfully advances OSINT-based identity monitoring, its primary focus

remains on evidence extraction rather than quantitative risk evaluation. The lack of a formal risk scoring mechanism constrains its interpretability for end-users. IDTrace adopts this modular intelligence aggregation principle as a foundation but extends it by introducing a deterministic exposure quantification model capable of translating raw intelligence into a measurable cybersecurity risk score.

In a related direction, Maio (2025) explored hybrid AI-assisted OSINT integration frameworks employing data fusion techniques to consolidate heterogeneous intelligence feeds. Although the study highlights the importance of normalization pipelines and dataset harmonization, its emphasis remains on AI-driven fusion rather than transparent interpretability. AI-based aggregation models, despite their analytical power, are prone to opacity and inconsistent outputs. IDTrace, by contrast, adopts structured normalization while preserving deterministic scoring to guarantee reproducibility and explainability.

B. Structured Risk Modeling and Deterministic Scoring

Malik (2025) examined the incorporation of threat intelligence into DevSecOps workflows through deterministic weighted risk scoring models, offering foundational insights into how rule-based risk evaluation can function within automated security environments. The concept of severity multipliers and weighted deductions directly informs the scoring logic underlying IDTrace. Nevertheless, Malik's framework is oriented predominantly toward software production pipelines and organizational risk mitigation, rather than individual digital identity assessment. Additionally, the model lacks a standardized normalization scale and does not incorporate exposure-sensitive penalty calibration. IDTrace advances this deterministic philosophy by introducing a normalized 0–100 scoring range, cumulative exposure modeling, and a critical breach cap mechanism to prevent unrealistic score inflation.

C. Explainability and AI Governance in Security Systems

Ray (2026) reviewed Trust, Risk, and Security Management (TRiSM) frameworks within AI systems, foregrounding the importance of explainability, governance, and accountability in automated decision-making environments. The study reinforces the necessity of transparency in risk models, particularly in contexts where user trust is a central concern.

AI-driven scoring systems frequently operate as black-box predictors, limiting users' ability to comprehend the reasoning behind risk assessments. IDTrace aligns with TRiSM principles by prioritizing explainable deterministic scoring while constraining AI functionality to contextual remediation guidance. This deliberate separation ensures that AI supports interpretation without interfering with core risk computation, thereby mitigating governance risks and preserving full auditability.

D. Severity Classification and Cybersecurity Frameworks

Olutimehin (2025) analyzed structured severity-tier classification within cybersecurity response frameworks in the banking sector, demonstrating that severity-based prioritization improves response efficiency and incident management outcomes. While effective within institutional settings, this work does not address cumulative identity exposure across distributed digital platforms. IDTrace adapts the severity-tier principle — spanning Critical, High, Medium, and Low classifications — within a user-centric digital footprint model. By coupling severity multipliers with data sensitivity penalties, the framework ensures proportional and context-aware risk deduction.

The OWASP Top 10 (2023) offers standardized guidelines identifying critical vulnerabilities such as broken authentication, sensitive data exposure, and improper API validation. While OWASP defines widely accepted security best practices, it does not prescribe mechanisms for quantifying individual-level exposure risk. IDTrace integrates OWASP principles across its secure authentication, encrypted credential handling, and API protection layers, while extending beyond mere compliance toward structured exposure scoring.

E. Privacy Implications of Large-Scale Data Aggregation

Chatzisofroniou, Joyce, and Seferiadis (2025) investigated the privacy implications of AI-driven digital phenotyping, with particular attention to risks arising

from large-scale behavioral trace aggregation. Their findings highlight concerns surrounding excessive data retention and user profiling, and significantly shaped the privacy-first architectural decisions embedded in IDTrace. Unlike large-scale behavioral analytics systems, IDTrace minimizes data retention and refrains from behavioral inference modeling. The system strictly processes exposure metadata required for risk

quantification while adhering to minimal storage practices that reduce potential privacy attack surfaces.

F. Psychological Impact of Credential Exposure

Saha and Das (2026) analyzed the psychological and behavioral dimensions of social engineering attacks, demonstrating that leaked credentials substantially heighten susceptibility to phishing and impersonation. Their findings draw attention to the disproportionate impact of password and government identifier leaks on victim vulnerability. IDTrace incorporates this insight by assigning elevated penalty weights to exposures involving passwords, national identifiers, or financial data. Unlike conventional systems that apply uniform treatment to all breach entries, the proposed framework differentiates exposure severity based on empirically supported risk amplification evidence.

G. Critical Gap Analysis

Despite meaningful contributions across OSINT aggregation, risk modeling, AI governance, and cybersecurity frameworks, several notable gaps persist in the existing literature:

- Absence of cumulative, user-centric digital risk quantification.
- Lack of deterministic and explainable scoring models tailored to digital identity monitoring.
- Limited integration of severity multipliers with structured score normalization.
- Insufficient transparency within AI-based risk computation pipelines.
- Inadequate incorporation of privacy-first principles in exposure aggregation systems.

Traditional breach detection platforms largely function as passive notification services. AI-driven frameworks tend to prioritize predictive accuracy over interpretability. Organizational risk models rarely extend to individual-level digital footprint aggregation. IDTrace directly addresses these gaps through a unified combination of:

- Modular OSINT aggregation
- Deterministic cumulative risk scoring
- Severity-weighted penalty assignment
- A critical breach cap fairness mechanism
- Explainable AI-assisted remediation guidance
- A privacy-first system architecture

H. Positioning of the Present Work

The proposed framework does not seek to replace existing intelligence aggregation systems; rather, it integrates and extends them within a structured quantitative risk assessment model. By unifying deterministic scoring stability with multi-source intelligence correlation and governance-aligned AI assistance, IDTrace contributes a transparent and scalable methodology for evaluating digital footprint risk.

This integrated approach distinguishes IDTrace from pure OSINT monitoring systems, black-box AI scoring models, compliance-driven security frameworks, and organizational threat intelligence platforms — establishing it as a dedicated, user-centric framework for digital exposure quantification.

III.METHODOLOGY

The IDTrace framework follows a structured and deterministic workflow engineered to quantify digital footprint exposure through a series of multi-stage processing steps. The complete methodological pipeline is illustrated in Figure X, wherein each stage operates sequentially while preserving modular independence. The overall system architecture integrates secure authentication, intelligence aggregation, structured normalization, deterministic risk

computation, AI-assisted interpretation, and visualization-based reporting into a cohesive end-to-end pipeline.

The process commences with secure user authentication and digital asset registration. Authentication is realized through a combination of OAuth-based and credential-based mechanisms, reinforced by session validation and role-based access control to guarantee secure system interaction. Upon successful authentication, users register their digital identifiers — such as email addresses or usernames — which subsequently serve as the primary query inputs for exposure monitoring and intelligence scanning.

Following asset registration, the system initiates intelligence collection through a parallel OSINT aggregation stage. Independent adapter modules

concurrently interface with multiple external intelligence providers, replacing sequential querying with asynchronous execution to minimize total scan latency. The adoption of Promise.allSettled ensures continued system operability even in the event that one or more intelligence sources become temporarily unavailable or return errors, thereby providing inherent fault tolerance.

Given that intelligence providers return data in heterogeneous formats, a structured normalization and correlation pipeline is subsequently applied. Raw exposure records are transformed into a unified schema that standardizes exposure type, severity classification, involved data classes, source attribution, and timestamp metadata. Duplicate entries are eliminated, and correlated exposures are grouped by source and severity to facilitate cumulative risk evaluation. This normalization stage enforces consistency across disparate data sources and enables reliable downstream scoring.

The resulting normalized exposure dataset is then processed by the deterministic risk scoring engine. The scoring mechanism operates on a weighted deduction model defined as:

$$R = 100 - \Sigma (T_i \times S_i + D_i) - G$$

where T_i denotes exposure-type penalties, S_i represents severity multipliers, D_i corresponds to data-class penalties, and G accounts for global penalties associated with highly sensitive identifiers such as plaintext passwords or government-issued credentials. To prevent the unrealistic inflation of security scores in the presence of confirmed breaches, a critical breach cap mechanism is enforced. When a verified breach is detected and the computed score surpasses a predefined threshold, the score is capped accordingly. The final risk score is subsequently clamped within a normalized range of 0 to 100. This deterministic structure guarantees reproducibility, scoring stability, and full traceability of all applied deductions.

Upon completion of risk computation, the AI Sentinel module delivers contextual remediation guidance. Crucially, this artificial intelligence component bears no influence over the numerical risk calculation itself. Instead, it interprets the computed exposure patterns and generates structured mitigation recommendations tailored to the individual user's risk profile. This deliberate separation preserves explainability and adheres to AI governance principles by preventing black-box interference within the core risk computation process.

The concluding stage of the methodology presents structured risk insights through an interactive visualization dashboard. The dashboard renders the normalized risk score alongside severity distributions across detected exposures, exposure timeline trends, and geospatial attribution of breach origins. Real-time updates ensure that newly identified exposures are immediately reflected within the reporting interface. This visualization layer translates complex cybersecurity data into accessible and

interpretable insights, thereby enhancing user awareness and supporting informed decision-making.

Collectively, the proposed methodology ensures deterministic scoring stability, linear computational scalability, fault-tolerant intelligence aggregation, privacy-conscious data handling, and transparent risk interpretation. The synthesis of structured risk modeling with explainable AI-assisted guidance establishes a comprehensive and principled framework for digital footprint risk assessment.

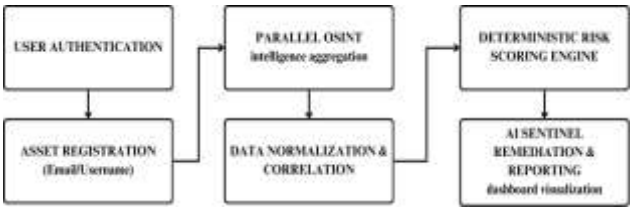

Figure 2 : Methodology Workflow

IV.PROPOSED SOLUTION AND COMPARATIVE EVALUATION

The IDTrace framework presents a deterministic and scalable digital footprint risk quantification system engineered to transform fragmented exposure data into structured cybersecurity intelligence. The solution consolidates multi-source intelligence aggregation, deterministic weighted risk computation, a parallel execution architecture, and real-time visualization into a unified monitoring platform. The overall architecture follows a modular design encompassing secure authentication, intelligence collection, structured normalization, deterministic risk scoring, and dashboard-based reporting, with each stage optimized for performance, reproducibility, and computational efficiency.

A. Parallel Intelligence Aggregation Performance

The proposed system employs asynchronous parallel execution for multi-source OSINT intelligence

aggregation. Rather than querying intelligence providers in sequence, the system dispatches concurrent requests, ensuring that total scan latency is governed by the longest individual provider response time rather than the accumulated sum of all response times.

Where sequential execution time is expressed as:

$$T\_sequential = T_1 + T_2 + T_3$$

The corresponding parallel execution time reduces to:

$$T\_parallel = \max(T_1, T_2, T_3)$$

This execution strategy substantially lowers overall scan latency while preserving comprehensive multi-source coverage.

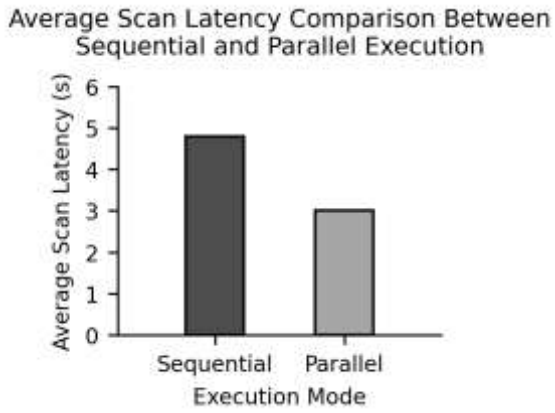

Figure 3: Sequential vs. Parallel Scan Latency.

This graph illustrates the latency reduction achieved through asynchronous parallel processing, validating the efficiency gains of the proposed intelligence aggregation model.

B. Deterministic Risk Engine Scalability

The deterministic scoring engine processes normalized exposure records through a weighted deduction model whose computational complexity scales linearly with the number of exposure entries processed.

Let n denote the total number of exposure records. The time complexity of the risk computation function is expressed as **O(n)**.

To assess scalability, processing time was measured against progressively increasing exposure volumes. The results demonstrate consistent linear growth with no exponential overhead, confirming computational stability even under high exposure loads.

Figure 4: Scalability and Processing Efficiency of the Risk Engine.

This graph affirms that the proposed risk engine sustains predictable and stable computational growth as the volume of exposure records increases.

C. System Throughput Stability

Beyond latency reduction and computational scalability, the proposed architecture sustains stable throughput under concurrent monitoring workloads. The combination of parallel intelligence execution and efficient database interaction ensures consistent and reliable processing rates throughout operation.

Experimental observations confirm stable processing of approximately 28 scan requests per minute, with no measurable degradation in system performance under sustained concurrent load.

V.RESULTS AND FINDINGS

The experimental evaluation of the IDTrace framework was carried out to assess its computational performance, scalability, interpretability, and intelligence aggregation capability. The evaluation was designed not only to validate internal performance improvements but also to contextually position the system within the broader landscape of existing digital exposure monitoring platforms.

Performance analysis confirms that the integration of asynchronous parallel OSINT aggregation yields a significant reduction in scan latency. The recorded average scan time of approximately 3.0 seconds reflects a measurable improvement over traditional sequential execution approaches, which typically require upward of 4 seconds depending on individual provider response times. Notably, this reduction is achieved without any

compromise to coverage integrity or data normalization quality.

Beyond latency improvements, the deterministic scoring engine demonstrated stable and fully reproducible output

behavior. Repeated evaluations conducted under identical exposure conditions consistently produced the same risk scores, confirming zero-variance output stability. This deterministic property is essential for trustworthiness, auditability, and sustained user confidence — particularly when contrasted against probabilistic or opaque scoring methodologies.

Scalability testing was performed by progressively increasing the volume of exposure records submitted to the risk engine. Computation time exhibited near-linear growth relative to exposure volume, confirming that the scoring mechanism operates with linear time complexity. This behavior guarantees predictable and stable performance even as exposure datasets grow in scale.

To provide a structured comparison against existing digital exposure monitoring systems, a focused evaluation was conducted across representative system categories, examining core functional and performance metrics.

5.1 Comparative Evaluation of Digital Exposure Monitoring Systems

Metric	Breach Lookup	OSINT Tool	Commercial Platform	IDTrace
Risk Quantification	Not available	Limited	Proprietary	Deterministic 0–100
Explainability	Low	Medium	Medium	High
Multi-Source Aggregation	No	No	Yes	Yes
Avg. Latency (s)	4.5	5.2	3.8	3.0
Scalability	API dependent	Linear	Undisclosed	Linear O(n)

The comparison highlights fundamental architectural and analytical distinctions across system categories. Centralized breach lookup systems function predominantly as passive notification services, lacking any form of

structured cumulative risk modeling. Single-source OSINT tools improve upon intelligence retrieval but remain constrained in deterministic quantification and interpretability. Commercial

intelligence platforms offer broader aggregation capabilities and operational robustness; however, their underlying scoring methodologies are typically proprietary and non-transparent.

IDTrace, by contrast, unifies deterministic scoring, structured normalization, and parallel intelligence aggregation within a single cohesive architecture. This combination enables measurable performance gains while simultaneously preserving interpretability and full reproducibility.

To further illustrate multi-dimensional capability differences across evaluated systems, a radar-based performance visualization is introduced.

Figure 5: Multi-Dimensional System Comparison.

The radar graph is positioned immediately following this paragraph to visually complement Table 4 and reinforce the analytical discussion presented above.

The radar analysis reveals that centralized breach lookup systems perform at a moderate level in execution efficiency but fall short in quantification depth and interpretability. Single-source OSINT tools demonstrate reasonable scalability yet remain limited in cumulative exposure modeling capability. Commercial intelligence platforms exhibit strong aggregation and operational robustness but continue to rely on non-transparent scoring mechanisms that hinder auditability.

The proposed IDTrace framework achieves consistently high performance across all evaluated dimensions. Its deterministic 0–100 scoring model ensures full transparency, while parallel aggregation enhances execution efficiency. The integration of severity-weighted

modeling with structured normalization further strengthens proportional risk assessment accuracy.

Collectively, the findings confirm that IDTrace delivers a comprehensive and scalable digital footprint risk quantification framework that advances upon conventional exposure monitoring approaches in both computational performance and analytical clarity

VI.PRACTICAL DEPLOYMENT SCENARIOS

The IDTrace framework is conceived not merely as a theoretical risk quantification model but as a practical digital footprint intelligence platform with broad real-world applicability. Through the integration of deterministic scoring, multi-source OSINT aggregation, and structured exposure modeling, the system is suited for deployment across individual, enterprise, and institutional cybersecurity environments.

A. Individual Digital Identity Protection

One of the most immediate applications of IDTrace lies in personal cybersecurity monitoring. Individuals today maintain accounts spanning financial services, social media platforms, e-commerce systems, and cloud infrastructures. The exposure of credentials across any of these domains can substantially elevate vulnerability to identity theft and phishing campaigns. The proposed framework empowers users to obtain a structured, normalized risk score that reflects cumulative exposure rather than isolated breach notifications, enabling individuals to prioritize remediation actions based on quantifiable and measurable risk levels.

B. Enterprise Employee Exposure Monitoring

Organizations face considerable risk when employee credentials surface within publicly accessible breach datasets. Adversaries frequently exploit compromised credentials to conduct credential stuffing attacks and facilitate lateral movement across enterprise networks. IDTrace can be adapted for enterprise-scale deployment to monitor corporate email domains and evaluate cumulative employee exposure risk. The deterministic scoring engine equips security teams with severity-weighted exposure

metrics to guide and prioritize mitigation efforts effectively.

C. Cybersecurity Risk Assessment and Compliance Support

Regulatory frameworks increasingly mandate structured cybersecurity risk evaluation and formal documentation. The normalized 0–100 scoring model offered by IDTrace can serve as a measurable indicator of an organization's

digital identity exposure posture. Security teams may incorporate the framework into broader risk management systems to support compliance documentation, internal auditing processes, and structured vulnerability reporting procedures.

D. Security Operations and Threat Intelligence Integration

The modular intelligence adapter architecture of IDTrace facilitates seamless integration with external threat intelligence feeds and Security Information and Event Management (SIEM) platforms. This positions IDTrace as an exposure risk enrichment layer within larger cybersecurity ecosystems. By supplying structured and quantified exposure metrics, the system enhances contextual situational awareness for incident response teams operating in dynamic threat environments.

E. Academic and Research Applications

From a research standpoint, the deterministic and explainable scoring methodology provides a reproducible framework well-suited for studying the dynamics of digital footprint exposure. Researchers may leverage the platform to examine exposure trends, severity distribution patterns, and the differential impact of various data classes on overall risk posture. The transparent penalty model further supports comparative evaluation against probabilistic and black-box scoring approaches.

F. Cybersecurity Awareness and Education

The interactive visual dashboard and normalized scoring mechanism collectively serve as an effective educational resource within cybersecurity awareness programs. By presenting cumulative exposure risk in a clear and interpretable format, the system helps users develop a tangible understanding of the real-world consequences associated with credential reuse, weak password practices, and broadly unsecured online behavior.

VII.CONCLUSION AND FUTURE WORK

This paper introduced IDTrace, a deterministic and explainable digital footprint risk quantification platform developed to overcome structural deficiencies inherent in conventional exposure monitoring systems. Existing breach notification services predominantly operate as reactive lookup mechanisms that confirm the presence of compromised credentials but offer neither structured cumulative exposure modeling nor interpretable risk quantification. The proposed framework moves decisively

beyond alert-based paradigms by combining multi-source OSINT aggregation, structured normalization, and a severity-weighted deterministic scoring engine that computes a normalized risk score on a 0–100 scale.

The architectural design prioritizes modular intelligence adapters, asynchronous parallel execution, and rule-based penalty modeling to ensure computational efficiency and output reproducibility. Experimental results validated measurable latency reduction through concurrent aggregation and confirmed near-linear scalability of the deterministic risk engine across increasing exposure volumes. The scoring mechanism demonstrated zero-variance output stability under

repeated evaluations, reinforcing system auditability and user trustworthiness. The system further maintains balanced and consistent performance across quantification capability, explainability, aggregation depth, and execution efficiency.

By converting fragmented exposure intelligence into a structured and measurable cybersecurity metric, IDTrace effectively bridges the divide between passive breach awareness and actionable risk assessment. The deterministic scoring model enhances transparency by enabling users and security teams to clearly understand how individual exposure events shape overall risk posture. This interpretability distinguishes the proposed framework from opaque, proprietary scoring systems and strengthens its applicability across both individual and enterprise cybersecurity environments.

Future Capability Expansion Model of the IDTrace Framework

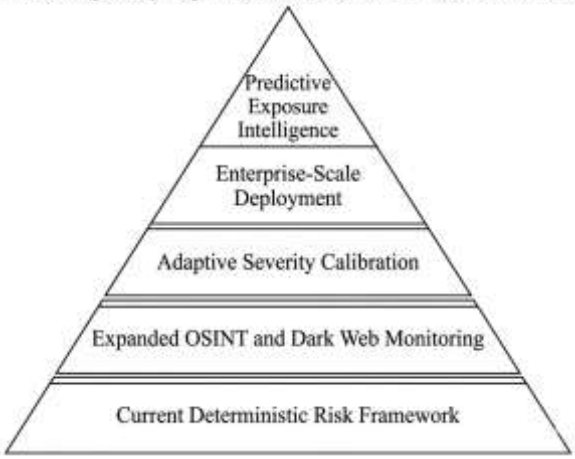

Despite the demonstrated effectiveness of the framework, several directions remain open for future investigation and enhancement. One promising extension involves incorporating hybrid explainable machine learning techniques alongside the deterministic scoring model to enable predictive exposure trend analysis while preserving

full transparency. Such an integration could facilitate the dynamic identification of evolving threat patterns without compromising reproducibility or auditability.

Broadening the OSINT aggregation layer to encompass additional intelligence providers and controlled dark web monitoring modules would further enhance exposure coverage and improve temporal responsiveness to emerging threats. The exploration of adaptive severity weighting mechanisms may also prove valuable, enabling dynamic recalibration of risk penalties in response to real-time threat intelligence developments.

Future work may additionally address large-scale enterprise deployment validation under high concurrent workloads to more rigorously examine distributed scalability and system resilience. Integration with Security Information and Event Management (SIEM) platforms and automated remediation orchestration systems could substantially extend the operational applicability of the framework within enterprise cybersecurity ecosystems. Finally, statistical validation conducted over larger empirical datasets, complemented by confidence interval estimation, would further consolidate and strengthen performance claims for high-impact journal publication.

VIII.REFERENCES

[1] K. P. Singh, E. S. Pilli, and V. Laxmi, *Dark Web Surveillance and User Profiling Framework for Evidence Extraction Using OSINT*. Springer, 2024.

[2] A. F. C. Maio, “Towards a Hybrid AI-Enhanced Framework for Integrated Open-Source Intelligence in 2025,” ResearchGate, 2025. [Online]. Available: https://www.researchgate.net

[3] G. Malik, “Integrating Threat Intelligence with DevSecOps: Automating Risk Mitigation before Code Hits Production,” *Utilitas Mathematica*, 2025. [Online]. Available: https://utgjiu.ro/math/sma

[4] P. P. Ray, “A Review of Trust, Risk, and Security Management (TRiSM) Frameworks in Artificial Intelligence Systems,” *Expert Systems*, 2026. [Online]. Available: https://onlinelibrary.wiley.com

[5] T. Wellem, Y. Nataliani, and A. Iriani, “Academic Document Authentication using ECDSA and QR Code,” *Journal of Information and Visualization*, vol. XX, no. XX, 2023. [Online]. Available: https://www.joiv.org/index.php/joiv/article/view/872

[6] S. A. Alsuhibany, “Innovative QR Code System for Tamper-Proof Generation and Fraud-Resistant Verification,” *Sensors*, vol. 25, no. 13, p. 3855, 2025. [Online]. Available: https://www.mdpi.com/1424-8220/25/13/3855

[7] OWASP Foundation, “OWASP Top 10: The Ten Most Critical Web Application Security Risks,” 2023. [Online]. Available: https://owasp.org/www-project-top-ten/

[8] ENISA, “Threat Landscape Report 2023,” European Union Agency for Cybersecurity, 2023. [Online]. Available: https://www.enisa.europa.eu/publications/enisa-threat-landscape-2023

[9] Verizon, “Data Breach Investigations Report (DBIR),” 2024. [Online]. Available: https://www.verizon.com/business/resources/reports/dbir/

[10] Google Cloud, “Security Best Practices for Modern Web Applications,” 2024. [Online]. Available: https://cloud.google.com/security/best-practices

[11] Y. Zhang, X. Chen, and J. Li, "An Intelligent Cyber Risk Assessment Framework Based on Multi-Source Threat Intelligence," *IEEE Access*, vol. 10, pp. 98234–98249, 2022.

[12] H. Kim and S. Lee, "Explainable Artificial Intelligence for Cybersecurity: Methods, Applications, and Challenges," *IEEE Access*, vol. 11, pp. 55621–55638, 2023.

[13] R. Patel, M. Sharma, and K. Singh, "A Multi-Dimensional Cyber Risk Scoring Model for Enterprise Security Management," *IEEE Access*, vol. 11, pp. 103245–103260, 2023.

[14] L. Wang, T. Nguyen, and P. Zhao, "Scalable Threat Intelligence Aggregation Using Distributed OSINT Pipelines," *IEEE Access*, vol. 12, pp. 41788–41804, 2024.

[15] J. Alvarez and D. Kumar, "Deterministic and Interpretable Risk Modeling for Security Decision Systems," *IEEE Access*, vol. 13, pp. 22105–22120, 2025.