

Article

AI-Amplification Indicator: An Actor-Level Scoring Framework for Ransomware Operations on the Dark Web

Mostafa Moallim ¹, Seokhee Lee ², Ibrahim Alzahrani ², Faisal Abdulaziz Alfouzan ³
and Kyounggon Kim ^{2,*}

¹ Center of Artificial Intelligence for Security, Naif Arab University for Security Sciences, Riyadh 14812, Saudi Arabia; mmoallim@nauss.edu.sa

² Department of Cybersecurity and Digital Forensics, Naif Arab University for Security Sciences, Riyadh 14812, Saudi Arabia; slee@nauss.edu.sa (S.L.); ialzahrani@nauss.edu.sa (I.A.)

³ Private Company, Riyadh 11564, Saudi Arabia; alfouzanfaisal@gmail.com

* Correspondence: kkim@nauss.edu.sa

Abstract

Ransomware operations have evolved from isolated malware incidents into organized ransomware-as-a-service (RaaS) ecosystems that employ coordinated tactics, techniques, and procedures and increasingly rely on automation and artificial intelligence to scale intrusions. However, most assessments remain artifact-centric, focusing on malware signatures or aggregate victim counts, which provide limited visibility into differences in actor-level behavior and operational capability. This study introduces the AI-Amplification Indicator (AIAI), an interpretable actor-level scoring framework that transforms publicly observable leak-site disclosures and verifiable open-source evidence into quantitative behavioral profiles. Using continuous monitoring of dark web leak portals, we construct a standardized dataset of ransomware disclosures for 2025 with temporal, geographic, and sector metadata. AIAI measures four complementary dimensions: GenAI-enabled social engineering, operational tempo and orchestration, targeting breadth and diversification, and temporal scaling dynamics. Indicators are computed for all observed actors, while comparative profiling focuses on the ten most active actors to ensure stable behavioral estimation. The analysis reveals substantial heterogeneity in posting cadence, targeting strategies, and scaling dynamics, as well as limited but measurable evidence of automated or AI-assisted deception. These differences are not captured by victim counts alone. The proposed framework provides a transparent and reproducible approach for actor-level ransomware intelligence, enabling systematic comparison of operational styles and supporting data-driven defensive prioritization.

Keywords: ransomware; ransomware-as-a-service (RaaS); cyber threat intelligence; threat actor profiling; tactics, techniques, and procedures (TTPs); behavioral scoring; dark-web intelligence; generative AI; operational scaling



Received: 12 February 2026

Revised: 23 March 2026

Accepted: 4 April 2026

Published: 8 April 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Ransomware has evolved from opportunistic malware campaigns into organized, profit-driven operations conducted by structured cybercriminal groups. Modern attacks increasingly follow a professionalized ransomware-as-a-service (RaaS) model and frequently employ double- or pure-extortion strategies in which victims are publicly exposed through dedicated leak portals rather than relying solely on file encryption [1,2]. This shift has transformed ransomware from isolated technical incidents into coordinated, multi-stage

campaigns that resemble persistent adversarial operations, often following structured lifecycles spanning initial access, internal activity, and eventual impact or exfiltration phases [3].

Concurrently, advances in automation and generative artificial intelligence (GenAI) have been associated with reduced operational costs and lower skill barriers for large-scale intrusion activity, enabling threat actors to scale campaigns with fewer resources and less manual effort than previously required [4]. Recent research shows that GenAI can enhance phishing and social engineering by enabling more convincing, personalized, and scalable deceptive content compared with traditional approaches [5,6]. These capabilities allow attackers to rapidly localize lures, adapt messaging, and expand victim reach with minimal human effort, effectively amplifying the operational capacity of ransomware groups.

Despite these developments, most ransomware assessments remain artifact-centric, focusing on malware signatures, vulnerability severity, or aggregate victim counts. While useful for detection, such metrics provide limited visibility into how attacker organizations differ in operational tempo, targeting breadth, or scaling behavior. In response, threat intelligence research has increasingly emphasized tactics, techniques, and procedures (TTPs) and adversary-centric behavioral modeling, often structured through frameworks such as MITRE ATT&CK [7].

Recent work has further highlighted the need for systematic and quantitative metrics that evaluate ransomware operations at the group level rather than at the level of individual malware samples [8]. However, there remains a lack of interpretable methods that directly measure actor-level behavior and explicitly account for observable signals of operational amplification associated with automation and AI adoption.

To address this gap, we introduce the AI-Amplification Indicator (AIAI), a rubric-based framework that transforms publicly observable leak-site disclosures and verifiable open-source evidence into actor-level measures of operational scaling. AIAI evaluates four complementary dimensions: GenAI-enabled social engineering, operational tempo and orchestration, targeting breadth and diversification, and cross-country and cross-sector expansion. Each dimension is derived from measurable signals and mapped to an interpretable ordinal score, enabling transparent and reproducible comparison across groups.

We apply the proposed framework to all ransomware actors observed during the 2025 study period. To ensure stable and interpretable comparisons, detailed profiling focuses on the ten most active actors, as lower-activity groups frequently provide insufficient behavioral signal for reliable estimation of cadence, diversification, and scaling metrics. This targeted comparison enables clearer differentiation of operational styles beyond simple victim counts, revealing patterns such as burst-style disclosures, sustained multi-month activity, diversified targeting, and varying degrees of social-engineering enablement. Together, these profiles provide a more structured view of how ransomware organizations organize and scale their operations.

The main contributions of this study are threefold:

- We develop a ransomware intelligence dataset derived from continuous monitoring of dark web leak portals, providing standardized temporal, geographic, and sector metadata for actor-level behavioral analysis.
- We propose the AI-Amplification Indicator (AIAI), a scoring framework that operationalizes observable signals of operational scaling and AI-assisted amplification at the ransomware-actor level.
- We apply the framework across the full set of observed actors and present a focused comparative analysis of leading groups, demonstrating how actor-level scoring reveals distinct operational styles and threat profiles that are not captured by aggregate victim counts alone.

The remainder of this paper is organized as follows: Section 2 reviews related work. Section 3 describes the dataset and coverage. Section 4 presents the AIAI methodology. Section 5 reports the analysis and results. Section 6 discusses the findings and future research directions. Section 7 presents the limitations. Section 8 concludes this paper.

2. Literature Review

To establish a theoretical foundation for the proposed scoring framework, we reviewed the existing literature across four key domains: (1) evolution of cyber threat assessment, (2) the RaaS ecosystem on the Dark Web, (3) automated threat intelligence mapping, and (4) the emerging impact of generative AI in cybercrime.

From a cyber threat intelligence (CTI) perspective, this review is critical as it transforms raw behavioral data into actionable strategic insights. We analyzed how threat evaluation methodologies have shifted from static vulnerability metrics to dynamic, adversary-centric behavioral analysis. This approach builds upon our previous research, which focused on an examination of cryptocurrency-driven ransomware syndicates operating on the darknet [8], and utilized social networks to analyze attack trends and TTPs specifically within the Arab world [9]. By extending these principles, the current study focuses on quantifying the operational gains achieved through AI adoption.

The evaluation methodologies for cyber threats have evolved from static, asset-oriented security assessments into dynamic, adversary-centric behavioral analysis. Early frameworks primarily relied on indices such as the Common Vulnerability Scoring System (CVSS) to evaluate the severity of specific vulnerabilities; however, modern threat modeling has shifted toward quantifying the tactics and techniques employed by attackers [10]. Notably, scoring mechanisms are being utilized as a pivotal methodology for supporting complex decision-making across various domains, including financial risk assessment and medical diagnostics, beyond just cybersecurity [11,12]. A model to quantify attack complexity was proposed by assigning weights to attack technologies based on the MITRE ATT&CK framework [11], while TTP-based risk assessment studies have been analyzed to emphasize the necessity of dynamic models that reflect evolving adversary behaviors over static indicators [10]. Furthermore, the importance of scoring in evaluating organized intrusions has been underscored by classifying techniques across the Advanced Persistent Threat (APT) lifecycle [12], and quantitative comparative analyses of ransomware detection performance have been provided using deep learning technologies [13].

Recent ransomware trends have rapidly evolved into “Double Extortion” schemes, where attackers encrypt data and subsequently threaten to leak sensitive corporate information via Dedicated Leak Sites (DLSs) [1]. This ecosystem has become highly professionalized through the RaaS model, with attackers actively leveraging dark web infrastructure to evade law enforcement tracking [2]. The impact of double-extortion ransomware on corporate reputation and psychological pressure has been analyzed [1], and the business models and potential expansion of the RaaS ecosystem have been forecasted [2]. Consequently, generating intelligence based on real-time data collected from the dark web is essential for identifying practical threat trends. Methodologies for monitoring the underground economy on the dark web have been surveyed [14], and a platform for the automated extraction and statistical analysis of data from TOR-based sites has been proposed [15]. Additionally, research on intelligently analyzing Dark Web threat data using a multi-agent framework has been conducted [16], alongside studies on self-configuring defense mechanisms against ransomware in Internet of Medical Things (IoMT) environments [17].

The MITRE ATT&CK framework is currently recognized as the most advanced industry standard for systematizing cyber attacker behavior patterns [7]. The design philosophy of the framework for classifying adversary tactics and techniques has been presented,

which continues to be refined annually to reflect new threat scenarios [7]. As ransomware is a subset of cyberattacks, active research is underway to define and classify its TTPs through this framework. Techniques for automatically mapping threat reports to the ATT&CK matrix using Natural Language Processing (NLP) have been researched [18,19], and the TTPDrill model for extracting tactical steps from unstructured intelligence data has been introduced [10]. Furthermore, a dataset for classifying public threat reports into ATT&CK tactics has been established [20], and adversary behavior patterns in Industrial Control System (ICS) environments have been profiled in a matrix format [21]. The academic applications of the ATT&CK framework have also been synthesized through the analysis of over 400 publications [22].

Finally, the emergence of generative AI has introduced an “amplification effect,” allowing attackers to strike a broader range of targets with fewer resources. The impact of Large Language Models (LLMs) on social engineering techniques and malware production has been analyzed [4], and concepts for measuring the operational gains resulting from AI adoption in modern cybercrime have been presented [23]. Technical sophistication was confirmed through the analysis of leaked Conti ransomware source code [24], and the utility of the TRAM tool for the automated mapping of threat reports has been demonstrated [25]. This study incorporates the intelligence generation techniques and classification systems of these prior works to quantitatively evaluate the AI amplification effect of ransomware groups based on empirical dark web data.

Gap Analysis

The synthesis of the existing literature reveals several critical research gaps that this study aims to address:

- Integration of multi-domain scoring logic: moving beyond traditional security metrics, this study introduces scoring techniques validated in other industries, such as finance and healthcare, to objectively quantify the capabilities of attack groups.
- Empirical intelligence via dark web analysis: extending previous monitoring methodologies, we utilize timestamps and metadata extracted from actual Dark Web leak sites to quantify the real-world operational tempo of ransomware groups.
- Measurement of AI amplification via MITRE ATT&CK: by combining technology-centric scoring with the proposed AIAIs, we introduce an actor-level behavioral scoring framework that operationalizes observable proxy signals of AI-enabled amplification in ransomware operations, enabling systematic, interpretable, and reproducible comparison of operational scaling patterns across threat actors.

These gaps necessitate a shift toward more empirical, behavior-based quantification that accounts for the operational realities of modern threat actors. In response, this study focuses on a recent and complete observation window (January–December 2025), enabling consistent, cross-actor comparison using contemporaneous operational data rather than aggregated or retrospective reports. While dark web data is inherently subject to selective disclosure by adversaries, it remains the most empirical foundation for generating real-time threat intelligence. By acknowledging these inherent data constraints, this study focuses on quantifying a conservative lower bound of cyber threats based on observable publication behavior, ensuring that the resulting AIAI profiles represent verifiable operational output rather than speculative estimates.

3. Data and Coverage

This study analyzes ransomware activity observed between 1 January and 31 December 2025 across Arab countries. The unit of analysis is a distinct victim organization publicly

named on a ransomware leak site. Each observation represents a unique victim–actor disclosure event rather than the number of technical intrusions or ransom payments.

3.1. Data Sources and Collection Procedure

Victim disclosures were collected through continuous monitoring of actor-operated leak portals on the dark web maintained by NAUSS, which served as the primary source of actor attribution and victim identification. Collection combined automated scraping of leak-site posts with analyst review to verify timestamps and remove reposts or mirrored entries.

To assess monitoring coverage and mitigate source bias, the dataset was cross-validated against supplementary OSINT sources including vendor threat reports, regional cybersecurity advisories, and media coverage of confirmed ransomware incidents during 2025. Disclosures identified through supplementary sources were used for validation and consistency checking of actor attribution and timing within the primary monitoring feed. This multi-source validation approach aligns with established practice in dark web ransomware intelligence research [8] and improves coverage robustness relative to single-source monitoring while maintaining a consistent primary data collection protocol.

Duplicate and mirrored entries were identified by cross-checking four criteria simultaneously: victim organization name, actor attribution, publication timestamp, and source portal. Entries sharing the same victim name and actor attribution with identical timestamps were flagged and removed following manual confirmation. Where the same organization appeared under different country entries on the same actor’s portal, records were retained as separate disclosure events if the victim operated across multiple jurisdictions, as confirmed through open-source investigation. When a victim organization appeared under multiple actors within a short temporal window, entries were treated as separate disclosure events attributed to the respective actors rather than merged, reflecting documented re-extortion behavior and affiliate overlap in the RaaS ecosystem. For example, the same victim (Hôpital La Rabta) was claimed by two different actors (Devman and Qilin) within a three-day window; both entries were therefore retained as distinct actor-attributed disclosures.

All records were reviewed and validated by the NAUSS CTI team to ensure consistency of actor attribution and temporal alignment.

3.2. Dataset Composition and Inclusion Criteria

After validation and deduplication, the final dataset contains 147 distinct victim organizations across 14 countries and activity attributed to 48 ransomware actors. Each row represents a unique victim–actor disclosure event. Duplicate postings referring to the same organization or repeated extortion attempts were removed, while all actors with at least one confirmed disclosure during the observation window were retained. No additional filtering or minimum-volume thresholding was applied in order to preserve the full observable ecosystem.

Figure 1 illustrates the distribution of publicly disclosed victims across ransomware actors. The activity is highly skewed, with a small number of groups accounting for a disproportionate share of disclosures, notably *CLOP* and *Qilin* (nine victims each), followed by *INC Ransom* and *Kill Security* (eight victims each), and several mid-tier actors with comparable but slightly lower volumes. This uneven distribution indicates substantial heterogeneity in operational intensity across groups rather than uniform activity levels. Consequently, the AIAI framework adopts an actor-centric perspective, evaluating tempo of attacks, targeting breadth, scaling dynamics, and social-engineering enablement at the group level instead of relying solely on aggregate victim counts.




Figure 1. Number of publicly disclosed ransomware victims per actor in 2025 across Arab countries. Raw victim counts are shown here for completeness; normalized activity rates accounting for group lifespan (victims per active month) are presented as part of the AIAI-2 operational tempo analysis in Section 4.3.

3.3. Metadata and Derived Attributes

Country, sector, and publication timestamps were extracted directly from leak-site disclosures when provided. The sector labels were assigned to the Critical Infrastructure Taxonomy of the Cybersecurity and Infrastructure Security Agency (CISA) to ensure consistent categorization between actors [26].

When sector information was explicitly stated in the leak post, the stated label was accepted as primary and mapped directly to the nearest CISA category. When sector information was absent or ambiguous, classification was inferred through open-source investigation, including review of the victim organization’s official website, public business registries, and professional network profiles. Organizations were then assigned to the most specific applicable CISA critical infrastructure sector. Where classification remained uncertain after this process, the entry was assigned to the most general applicable category. All sector classifications were reviewed by a second CTI specialist from the NAUSS CTI team and verified by the supervising author to ensure consistency and reduce labeling error.

These fields were used to compute derived metrics for the AIAIs, including temporal cadence (AIAI-2), geographic and sector diversification (AIAI-3), and scaling dynamics over time (AIAI-4).

3.4. Dataset Structure

Each observation corresponds to a single ransomware leak-site disclosure event. The analytical dataset stores standardized fields extracted directly from leak portals and validated through analyst review. Table 1 summarizes the primary variables retained for analysis. AIAI-2–AIAI-4 are computed directly from these structured fields through aggregation and transformation. AIAI-1 differs methodologically and is derived from actor-attributed campaign artifacts and external reporting, as it relies on evidence not observable in leak-site metadata.

Table 1. Primary fields in the ransomware disclosure dataset.

Field	Type	Description
Victim	text	Organization name listed on the leak site
Ransomware Gang	categorical	Actor or group self-identified on the portal
Detection Date (UTC + 0)	datetime	First observed publication timestamp
Ransomware URL	text	Link to the leak or disclosure page
Victim Site	text	Official website or domain of the victim
Victim Country	categorical	Country extracted from the disclosure
Industrial Sector	categorical	Sector label as stated or inferred
CISA Industry	categorical	Standardized mapping to CISA critical infrastructure taxonomy

3.5. Coverage Limitations

Leak-site disclosures reflect only publicly revealed victims and do not represent the full universe of compromises. Actors may selectively publish victims or delay disclosure, and some incidents may never appear publicly. Consequently, the dataset measures observable publication behavior rather than total ransomware activity and should be interpreted as a conservative lower bound of regional impact.

In addition to these visibility constraints, a further consideration relates to the RaaS affiliate model. Leak-site disclosures are attributed to the ransomware group operating the portal (e.g., Qilin or LockBit) without distinguishing whether the underlying intrusion was conducted by the core development team or an affiliated operator. Because dark web monitoring platforms capture victim publications at the group-brand level, AIAI metrics reflect the aggregate observable output of the ransomware ecosystem associated with that brand.

Accordingly, cadence, tempo, and scaling indicators represent group-level behavioral profiles that capture the collective operational footprint of the ransomware brand, including affiliated activity. This aggregation aligns with how ransomware campaigns are publicly disclosed and analyzed in threat intelligence practice, where the group identity is defined by the leak-site presence rather than individual operators.

While affiliate heterogeneity can influence posting cadence across campaigns, this variation is inherently captured within the AIAI framework, as tempo and orchestration metrics are derived from the full observable output of the ransomware group. These indicators therefore reflect aggregate operational dynamics at the group level rather than assuming uniform behavior across individual operators.

4. Methodology

4.1. Overview of the AIAI Scoring Framework

This section defines the scoring methodology applied to the dataset described in Section 3. The proposed framework implements actor-level indicators designed to quantify observable patterns consistent with AI-assisted operational scaling by capturing changes in operational efficiency, coordination, and campaign scalability that are not readily achievable under purely manual workflows.

The AIAI framework does not directly measure the use of AI systems by ransomware actors. Instead, it captures observable behavioral patterns, such as rapid scaling, sustained activity, and diversified targeting, that are consistent with the operational effects of automation and AI-assisted workflows. AI and automation reduce the marginal cost of generating phishing content, coordinating campaigns, and adapting targeting strategies, which, in turn, enables higher operational tempo and broader campaign scaling. These indicators therefore function as proxy signals of AI-enabled amplification rather than direct evidence of AI adoption, capturing changes in operational capability rather than specific underlying technologies.

Specifically, the adversarial use of AI to scale ransomware operations reflects a broader technological arms race in cybersecurity, where defenders increasingly deploy advanced AI-based detection systems, including self-attention and sequence modeling architectures, to identify complex and evolving attack patterns [27]. This parallel highlights the dual role of AI amplification in both offensive and defensive contexts and motivates the development of structured behavioral indicators capable of systematically capturing attacker-side scaling dynamics, as formalized in the AIAI framework.

We introduce the AI-Amplification Indicators (AIAI), computed at the ransomware-actor level for the 2025 observation window, as illustrated in Figure 2. AIAI consists of four complementary components (AIAI-1 to AIAI-4), each capturing a distinct measur-

able dimension of actor behavior derived from leak-site disclosures and, where required, verifiable actor-linked open-source reporting.




Figure 2. Overview of the AI-Amplification Indicator (AIAI) framework, illustrating the transformation of leak-site disclosures into actor-level indicators and ordinal AIAI profiles.

The framework is rubric-based and interpretable. Rather than training a predictive model, observable indicators are mapped to discrete ordinal sub-scores and aggregated into a 0–3 scale, where higher values indicate stronger and more consistent signals. Explicit scoring thresholds are defined to support transparency, reproducibility, and robustness under the partial-visibility constraints inherent to leak-site data.

The 2025 dataset contains 48 distinct ransomware actors. Activity is highly skewed, with many actors responsible for only one or two disclosures during the observation window. Such sparse activity provides limited behavioral signal for cadence, diversification, or scaling metrics and may yield unstable or noisy estimates.

Accordingly, the automated indicators (AIAI-2–AIAI-4), which rely exclusively on structured leak-site metadata, are computed for the full set of actors to preserve completeness and avoid selection bias. In contrast, AIAI-1 relies on manual, evidence-gated artifact analysis and therefore requires actor-attributed campaign materials that are not uniformly available across all actors. As a result, AIAI-1 is evaluated only for the most active actors with sufficient reporting coverage.

For interpretability and comparative clarity, the result tables and figures report only the ten highest-scoring actors per indicator. This reporting choice reduces visual clutter while ensuring that computation itself remains population-wide for automated metrics.

AIAI-1 (GenAI social-engineering enablement) evaluates documented evidence of GenAI-enabled social engineering and phishing enablement and is computed using an evidence-gated rubric due to uneven artifact availability. AIAI-2 (operational tempo and orchestration) measures posting cadence, burst dynamics, and cross-country spread speed. AIAI-3 (targeting breadth and adaptation) captures sectoral and geographic diversity as proxies for operational reach. AIAI-4 (temporal scaling dynamics) characterizes scale-up dynamics from initial observation to peak activity.

4.2. AIAI-1: GenAI Social-Engineering Enablement

AIAI-1 evaluates whether a ransomware actor’s social-engineering operations exhibit credible, documented signals consistent with generative AI (GenAI)-enabled content generation. Recent industry reporting highlights a growing shift toward AI-assisted phishing, impersonation, and collaboration-tool abuse in ransomware intrusions, motivating the need for an indicator that captures scalable deception workflows [5].

Unlike AIAI-2–4, which rely on systematically observable leak-site behavior, AIAI-1 is evidence-gated and derived from actor-attributed lure artifacts and public incident reporting. The construct is operationalized using primary campaign materials (e.g., phishing emails, landing pages, scripts, or related infrastructure) together with measurable indicators of scalable content production, including rapid localization, contextual fluency, template or

variant diversity, and repeatable personalization. The unit of analysis is an actor-linked campaign or lure set for which artifacts are available for inspection.

Actors are assigned an ordinal score (0–3) using an evidence-gated rubric (Table 2), where higher scores require verified artifacts and repeated measurable indicators across campaigns. Because AIAI-1 depends on the availability of publicly documented, actor-attributed artifacts, evidence availability varies across ransomware groups and introduces asymmetry in observable signals. Actors with greater public visibility, for example, through vendor investigation or high-profile campaigns, are more likely to have attributable artifacts available for evaluation. To address this variability, each actor’s evidence category is recorded alongside the assigned score, distinguishing evaluations based on verified artifacts from those based on reporting alone. The evidence-gated rubric ensures that higher scores are assigned only when verifiable artifacts provide measurable indicators of scalable or automated content generation, thereby reducing the risk that differences in reporting volume inflate actor capability assessments.

Table 2. AIAI-1 evidence-gated scoring rubric for GenAI social-engineering enablement.

Score	Evidence Criteria
0	No credible reporting or artifacts link the actor to social-engineering activity.
1	Reporting links the actor to phishing or social engineering, but no actor-attributed artifacts are available for evaluation.
2	At least one actor-attributed campaign with verified lure artifacts demonstrating multiple observable indicators of scalable or automated content generation.
3	Multiple actor-attributed campaigns across time with repeatable artifacts showing consistent scaling indicators.

Scores reflect the best available public evidence at the time of evaluation and should be interpreted as conservative lower bounds on observable GenAI-enabled social-engineering activity.

To assess scoring reliability, AIAI-1 evaluations were conducted by the lead researcher and independently reviewed by a CTI specialist from the NAUSS using identical evidence sources and rubric criteria. Discrepancies, where present, were resolved through discussion to reach a final consensus among the authoring team.

The rubric-based design constrains evaluator discretion by requiring documented, verifiable evidence for each score level; in particular, actors cannot receive score 2 or higher without confirmed actor-attributed artifacts. This reduces subjectivity inherent to evidence-gated assessment and ensures that scores reflect observable evidence rather than absence of reporting. The evidence category classifications reported in Section 5 provide transparency about the evidentiary basis underlying each assigned score.

4.3. AIAI-2: Operational Tempo and Orchestration

AIAI-2 measures the observable operational tempo of ransomware actors based on leak-site publishing behavior. The indicator captures how frequently, how rapidly, and how widely an actor releases victim disclosures. Rather than estimating the true number of intrusions, AIAI-2 focuses on externally visible operational output derived from dated leak-site posts.

Four dimensions are considered: posting throughput during active periods, burst-style release behavior, geographic spread across affected countries, and the temporal spacing between consecutive publications. Each dimension is converted to an ordinal subscore (0–3) using fixed thresholds (Table 3).

Table 3. AIAI-2 operational tempo and orchestration subscore rubric (0–3).

Metric (Input)	Score 0	Score 1	Score 2	Score 3
Victims per active month (VPM)	<1.2	1.2 to <2.0	2.0 to <4.0	≥4.0
Burst index (BI)	≤1.0	>1.0 to <1.5	1.5 to <2.0	≥2.0
Spread velocity (SV)	<0.75	0.75 to <1.25	1.25 to <2.0	≥2.0
Cadence	>60	>21 to ≤60	>7 to ≤21	≤7

For each ransomware actor observed in 2025, the following quantities are computed from leak-site records:

- Total number of posted victim organizations (V);
- Number of months in which at least one victim was published (active months, M_{active});
- Victims per active month ($VPM = V / M_{\text{active}}$);
- Burst index ($BI = \max(\text{monthly}) / \text{mean}(\text{monthly})$, over active months);
- Spread velocity ($SV = \text{unique_countries} / M_{\text{active}}$);
- Cadence (median inter-posting gap in days, computed from publication timestamps at day-level granularity).

All inputs are derived directly from publication timestamps and victim metadata.

Cadence is computed as the median inter-posting gap in days. Because timestamps are recorded at day-level granularity, multiple disclosures within the same day may produce zero-day gaps, reducing the median interval. These near-zero values reflect burst-style publication behavior, where multiple victims are disclosed within concentrated time windows, rather than measurement error. To ensure stability, cadence is interpreted jointly with the burst index, which captures within-period concentration at the monthly level. This complementary design prevents same-day posting from dominating the tempo signal. Alternative statistics (e.g., quantile-based gaps or intensity models) may provide additional robustness and are left for future work.

Actors active during only one month necessarily exhibit neutral burst statistics ($BI = 1.0$) even when multiple victims are released within a short period, because the burst index is defined as the ratio of peak to mean monthly victims and collapses to 1.0 when only one month is active. This creates a structural floor that would systematically under-represent concentrated single-month campaigns. To ensure that concentrated single-month campaigns are not structurally penalized by this mathematical property, actors with a single active month and four or more posted victims are assigned the maximum burst subscore. The threshold is aligned with the subscore rubric in Table 3, where $VPM \geq 4.0$ already qualifies as maximum throughput; an actor simultaneously meeting this volume criterion within a single month provides convergent evidence of burst-style campaign behavior. This adjustment affects only a small subset of actors in the 2025 dataset (3 out of 48). Sensitivity checks conducted without the adjustment produced the same relative ranking among the highest-tempo actors, indicating that the rule does not materially change the comparative results but prevents systematic underestimation of concentrated disclosure bursts.

The four subscores are summed to produce a total orchestration value ranging from 0 to 12, which is then mapped into the final AIAI-2 indicator using four ordinal bins: a sum of 0–3 maps to score 0, 4–6 to score 1, 7–9 to score 2, and 10–12 to score 3. This binning divides the achievable range into four equal intervals, preserving interpretability while ensuring consistency with the other AIAI components.

4.4. AIAI-3: Targeting Breadth and Portfolio Diversification

AIAI-3 measures the breadth of a ransomware actor's victim portfolio using the leak-site disclosures contained in the dataset described in Section 3. The indicator captures whether an actor concentrates activity on a narrow set of countries or sectors, or instead demonstrates diversified targeting across geographies and industry verticals. AIAI-3 is designed as an observable proxy for scalable target selection and multi-market reach, and it reflects published victim exposure rather than the true underlying compromise footprint.

Two dimensions are considered: sectoral diversification (how evenly victims are distributed across sectors) and geographic diversification (how evenly victims are distributed across countries). Each dimension is converted to an ordinal subscore (0–3) using fixed thresholds (Table 4).

Table 4. AIAI-3 targeting breadth and diversification subscore rubric (0–3).

Metric (Input)	Score 0	Score 1	Score 2	Score 3
Sector entropy (H_{sec})	<0.7	0.7 to <1.2	1.2 to <1.6	≥ 1.6
Country entropy (H_{ctry})	<0.7	0.7 to <1.2	1.2 to <1.6	≥ 1.6
Sector concentration ($p_1^{(sec)}$)	≥ 0.70	0.50 to <0.70	0.35 to <0.50	<0.35
Country concentration ($p_1^{(ctry)}$)	≥ 0.70	0.50 to <0.70	0.35 to <0.50	<0.35

For each ransomware actor observed in 2025, the following quantities are computed:

- Sector entropy (H_{sec}): Shannon entropy of the victim distribution across sectors:

$$H_{sec} = - \sum_i p_i \log p_i$$

where p_i is the proportion of victims in sector i . Higher values indicate more even diversification across industries.

- Country entropy (H_{ctry}): Shannon entropy of the victim distribution across countries:

$$H_{ctry} = - \sum_j q_j \log q_j$$

where q_j is the proportion of victims in country j . Higher values indicate broader geographic spread.

- Sector evenness (J_{sec}): Pielou's evenness index for the sector dimension:

$$J_{sec} = \frac{H_{sec}}{\log N_{sec}}$$

where N_{sec} is the number of distinct sectors observed for that actor. J_{sec} ranges from 0 (complete concentration in one sector) to 1 (perfect evenness across all sectors) and is independent of category count.

- Country evenness (J_{ctry}): Pielou's evenness index for the country dimension:

$$J_{ctry} = \frac{H_{ctry}}{\log N_{ctry}}$$

where N_{ctry} is the number of distinct countries observed for that actor. J_{ctry} ranges from 0 (complete concentration in one country) to 1 (perfect evenness across all countries) and is independent of category count, enabling direct comparison of geographic diversification across actors with different portfolio sizes.

- Sector concentration ($p_1^{(sec)}$): dominance of the most frequent sector:

$$p_1^{(sec)} = \max_i \left(\frac{v_i}{V} \right)$$

where v_i is the number of victims in sector i . Higher values indicate specialization.

- Country concentration ($p_1^{(ctry)}$): dominance of the most frequent country:

$$p_1^{(ctry)} = \max_j \left(\frac{c_j}{V} \right)$$

where c_j is the number of victims in country j . Higher values indicate geographic concentration.

These four quantities provide complementary views of diversification: entropy captures distributional evenness, whereas concentration captures dominance by a single sector or country. Sectors follow the CISA taxonomy and countries follow the victim country named in the leak post, consistent with the counting rules defined in Section 3. Sectoral diversification is scored jointly using entropy and concentration. Actors receive higher subscores when victims are distributed across multiple sectors without heavy dominance by a single sector. In cases where entropy is moderate but one sector dominates, the concentration rule governs to avoid overstating diversification.

Geographic diversification is scored analogously, using country entropy and country concentration. Higher subscores indicate broader multi-country targeting with lower dominance by a single country.

The sector and country subscores are summed (range: 0–6) and then mapped into the final AIAI-3 score using four ordinal bins: a sum of 0–1 maps to score 0, 2–3 to score 1, 4–5 to score 2, and 6 to score 3.

4.5. AIAI-4: Temporal Scaling Dynamics of Ransomware Activity

AIAI-4 measures the observable scaling dynamics of ransomware actors within the monitored observation window. Specifically, it captures how quickly disclosure activity ramps up following first appearance, how long it is sustained, and how rapidly peak activity is reached. The indicator distinguishes actors that exhibit short, localized bursts from those demonstrating more sustained and structured expansion patterns.

Importantly, AIAI-4 metrics are computed relative to each actor's first observed appearance in the dataset rather than against a fixed calendar reference. As a result, scores reflect within-window scaling behavior rather than absolute expansion speed. Actors appearing late in the observation window may exhibit compressed ramp-up and persistence metrics due to limited exposure time, and scores should therefore be interpreted as observable-publication proxies rather than definitive measures of true operational velocity.

Three dimensions are considered: (i) spread velocity across countries and sectors during active months, (ii) time-to-spread, defined as how quickly an actor reaches broader victim coverage following first appearance, and (iii) campaign span, used to distinguish short burst-style expansions from sustained multi-month scaling. AIAI-4 operationalizes scaling as a composite indicator by combining these observable dimensions into a single score. Each dimension is converted to an ordinal subscore (0–3) using fixed thresholds (Table 5).

Table 5. AIAI-4 temporal scaling dynamics (0–3).

Metric (Input)	Score 0	Score 1	Score 2	Score 3
Ramp-up speed (t_5 days)	>120 or $V < 5$	46–120	15–45	≤ 14
Sustained presence (span_days , M_a)	$\text{span} < 30$ or $M_a < 2$	$\text{span} \geq 30$ and $M_a \geq 2$	$\text{span} \geq 90$ and $M_a \geq 3$	$\text{span} \geq 180$ and $M_a \geq 4$
Acceleration (peak lag)	≥ 5 months	3–4 months	2 months	0–1 months

For each actor, three subscores are computed from dated leak-site disclosures after deduplication:

- Ramp-up speed (s_{ramp}): measured using the time elapsed in days from the first observed disclosure to the fifth victim disclosure (t_5). This captures early-expansion velocity. Actors with fewer than five observed victims are assigned $s_{\text{ramp}} = 0$.
- Sustained presence (s_{persist}): measured using a joint criterion combining the total activity span in days (span_days) and the number of active months (M_a). This combined logic ensures that brief high-volume bursts are distinguished from sustained multi-month operations, requiring actors to satisfy both duration and activity thresholds to receive higher subscores.
- Acceleration (s_{accel}): measured using the peak lag, defined as the number of months between the first observed disclosure and the month of highest disclosure volume. Shorter lags indicate faster early acceleration toward peak operational intensity.

The three subscores are summed to produce a total scaling value ranging from 0 to 9. This total is then mapped into the final AIAI-4 indicator using four ordinal bins: a sum of 0–2 maps to score 0, 3–5 to score 1, 6–7 to score 2, and 8–9 to score 3. This binning preserves interpretability while maintaining consistency with the ordinal scoring structure used across the other AIAIs.

To improve reproducibility, the AIAIs follow a deterministic rubric-based scoring procedure. The general computation is summarized in Algorithm 1 as follows:

Algorithm 1 AIAI Indicator Scoring Procedure (AIAI-2 through AIAI-4)

Require: Indicator-specific input metrics computed from leak-site records (see Sections 4.3–4.5)

Ensure: AIAI score $\in \{0, 1, 2, 3\}$

- 1: Compute ordinal subscores using predefined thresholds (Tables 3–5)
 - 2: Sum subscores to obtain Total
 - 3: Map Total to final score using ordinal bins
 - 4: **return** AIAI score
-

4.6. TTP-Oriented Interpretation and ATT&CK Alignment

Although the AIAIs are computed directly from observable leak-site behavior and open-source evidence, the framework is grounded in a TTP-oriented threat intelligence perspective. Following indicator computation and ordinal scoring, AIAI profiles are interpreted in relation to higher-level tactics and techniques as defined by the MITRE ATT&CK framework.

This phase does not introduce additional computed metrics. Instead, it provides analytical validation by relating each AIAI dimension to established tactical constructs. For example, AIAI-1 aligns with *Initial Access* and *Resource Development* through scalable social engineering, AIAI-2 reflects coordinated execution associated with *Impact*, AIAI-3 captures systematic targeting behavior consistent with *Reconnaissance*, and AIAI-4 reflects campaign operationalization and expansion over time, consistent with generalized adversary lifecycle models that describe progression from initial access through internal activity and eventual

impact or exfiltration [3]. This perspective complements the ATT&CK mapping by situating observed scaling behavior within a broader campaign lifecycle context.

This alignment ensures that AIAI scores remain interpretable within established CTI taxonomies while preserving transparency and reproducibility under partial visibility constraints.

4.7. Threshold Sensitivity Analysis

To assess sensitivity to threshold selection, each key parameter was independently perturbed in both directions. For AIAI-2, VPM thresholds were shifted by ± 0.5 and cadence thresholds by ± 3 days. For AIAI-3, entropy thresholds were shifted by ± 0.1 and concentration thresholds by ± 0.05 . For AIAI-4, time-to-5-victims and span thresholds were shifted by ± 15 days.

Across all perturbations, top-tier actor membership remained stable for all indicators. For AIAI-2 and AIAI-3, both membership and ordering of the top 3 actors remained unchanged across all tested perturbations. For AIAI-4, the same three actors consistently occupied the top tier, with minor within-tier reordering observed under one ramp-up perturbation.

Lower-ranked actors exhibited occasional reordering across all indicators, reflecting proximity to threshold boundaries rather than structural instability. These results indicate that the principal findings of the AIAI framework are robust to reasonable threshold perturbations and are not artifacts of specific boundary choices.

Future work may further refine threshold calibration using percentile-based binning, bootstrapping, or stability-based optimization.

5. Analysis and Results

This section reports the results of applying the AIAI framework to ransomware actors observed during the 2025 study period.

5.1. AIAI-1 Results: GenAI Social-Engineering Enablement

Evidence availability varies substantially across actors, as detailed in Section 4.2; scores should be interpreted as conservative lower bounds on GenAI-enabled activity rather than definitive capability assessments. Table 6 reports AIAI-1 scores for the ten most active actors evaluated in the 2025 dataset. Consistent with the evidence-gated design of AIAI-1, scores are assigned only when actor-attributed phishing or social-engineering artifacts are publicly documented and sufficient to evaluate measurable indicators of scalable or automated content generation.

Among the evaluated actors, only *Qilin* attains score 2. Vendor reporting provides actor-attributed phishing campaigns with verified lure artifacts exhibiting multiple observable indicators of scalable content production, including contextual pretexting and infrastructure supporting repeatable variant deployment across targets [28]. These characteristics meet the rubric criteria for documented, repeatable social-engineering workflows consistent with automation.

To illustrate what constitutes a verified lure artifact under the AIAI-1 rubric, the *Qilin*-affiliated campaign documented by Sophos provides a representative example based on vendor-analyzed campaign materials [28]. In this campaign, targets received spear-phishing emails that closely mimicked legitimate ScreenConnect authentication alerts, incorporating contextually accurate account details and domain references. The campaign directed victims to lookalike domains designed to replicate the legitimate login interface, leveraging adversary-in-the-middle infrastructure to capture credentials and authentication tokens.

Table 6. AIAI-1 GenAI social-engineering enablement scores (2025).

Actor	Score	Evidence Category	Evidence Summary
Qilin	2	Verified artifacts	Actor-attributed phishing lure artifacts with contextual pretexting and infrastructure enabling repeatable, scalable variant deployment.
CL0P	1	Reporting only	Social-engineering or extortion activity reported; no attributable lure artifacts available for evaluating scalable or automated content generation.
INC Ransom	1	Reporting only	Reporting references phishing or intrusion activity without verifiable, actor-attributed lure materials.
Kill Security	1	Reporting only	Capability indications only; no documented phishing artifacts suitable for rubric-based evaluation.
DevMan	1	Reporting only	Technical ransomware activity documented, but no actor-attributed phishing or social-engineering artifacts observed.
Everest	1	Reporting only	Data-theft and extortion activity reported; no attributable lure materials demonstrating scalable or automated workflows.
NightSpire	1	Reporting only	Limited reporting on intrusion methods; no verifiable phishing artifacts available.
DragonForce	1	Reporting only	Malware/tooling lineage documented; no actor-attributed social-engineering campaign materials identified.
Kazu	1	Reporting only	Leak-site activity present; no attributable phishing or lure artifacts documented.
LockBit	1	Reporting only	Extensive payload automation and affiliate tooling documented, but no actor-attributed phishing artifacts or scalable content-generation evidence.

Importantly, Sophos links this activity to a persistent threat cluster active since 2022, characterized by consistent reuse of domain naming patterns, site structure, and phishing infrastructure across campaigns. This repeated deployment of similar infrastructure and techniques across multiple victims demonstrates systematic, scalable campaign execution rather than isolated activity. These characteristics of consistent infrastructure reuse, cross-campaign repeatability, and contextually tailored phishing lures exemplify the measurable indicators of automated or GenAI-enabled content generation required for an AIAI-1 score of 2.

All remaining actors (*CL0P*, *INC Ransom*, *Kill Security*, *DevMan*, *Everest*, *NightSpire*, *DragonForce*, *Kazu*, and *LockBit*) receive score 1. Although open-source and vendor reporting links these groups to phishing, credential abuse, or broader social-engineering activity, available sources do not provide actor-attributed lure artifacts sufficient to evaluate scalable or GenAI-style content generation under the AIAI-1 rubric. Industry surveys confirm that phishing remains a prevalent intrusion vector in 2025, but do not supply attributable campaign materials suitable for actor-level assessment [29–31].

Overall, attributable and publicly documented signals consistent with GenAI-enabled or automated social-engineering workflows remain limited within the 2025 observation period, with most actors demonstrating conventional delivery methods rather than clearly scalable content-generation mechanisms. This distribution suggests that, while ransomware operations exhibit increasing technical automation at the payload and infrastructure levels (AIAI-2–AIAI-4), verifiable evidence of GenAI-enabled scaling in initial-access social engineering remains comparatively rare.

Evidence classifications in Table 6 are derived from the cited threat intelligence and industry reports referenced in this section.

5.2. AIAI-2 Results: Operational Tempo and Orchestration

Table 7 and Figure 3 present the highest AIAI-2 scores observed in 2025. The positioning plot maps victims per active month (VPM) on the horizontal axis and spread velocity (SV) on the vertical axis, such that actors located toward the upper-right combine rapid disclosure throughput with broader multi-country exposure. Actors nearer the lower-left exhibit slower or more concentrated posting behavior.

Table 7. AIAI-2 operational tempo and orchestration score (2025), ordered by highest score.

Rank	Actor	V	M_{active}	VPM	BI	SV	G (Days)	AIAI-2
1	CL0P	9	1	9.0	1.00	4.0	0.04	3
2	Ransomhub	4	1	4.0	1.00	4.0	5.15	3
3	The Gentleman	4	1	4.0	1.00	4.0	0.00	3
4	Kill Security	8	4	2.0	2.50	1.0	6.14	2
5	Kazu	6	2	3.0	1.67	1.5	0.00	2
6	Everest	7	3	2.0	2.00	1.0	6.25	2
7	Qilin	9	4	2.0	1.50	1.5	16.60	2
8	BlackShrantac	3	1	3.0	1.00	3.0	1.70	2
9	RALord	3	1	3.0	1.00	2.0	6.52	2
10	Dragon Ransomware	2	1	2.0	1.00	2.0	1.73	2

**Figure 3.** AIAI-2 operational tempo and orchestration scores for the top ransomware actors (2025).

Three actors, *CL0P*, *Ransomhub*, and *The Gentleman*, attain the maximum AIAI-2 score of 3 and occupy this upper-right region. Although each was active during a limited number of months, they released multiple victims within highly compressed time windows, producing extreme victims-per-month values and very short inter-posting gaps. Consistent with the indicator design, these patterns are interpreted as coordinated disclosure bursts rather than isolated events, reflecting intensive short-duration campaigns.

Actors ranked fourth through tenth receive score 2 and display more distributed operational profiles. Groups including *Kill Security*, *Kazu*, *Everest*, and *Qilin* maintain activity across multiple months, combining moderate posting throughput with recurring

release surges. In the figure, these actors cluster in the central region, indicating sustained but less concentrated activity compared with the highest-tier burst operators.

Overall, AIAI-2 differentiates short-duration, high-density disclosure campaigns from steadier multi-month operations with moderate throughput. By focusing exclusively on observable publication tempo and orchestration rather than total victim counts, the indicator captures structural differences in operational intensity that are independent of targeting breadth or growth.

5.3. AIAI-3 Results: Targeting Breadth and Portfolio Diversification

Table 8 and Figure 4 summarize the highest AIAI-3 scores observed in 2025. The indicator positions actors using sector entropy (H_{sec}) and country entropy (H_{ctry}), where larger values on both axes correspond to more evenly distributed victim portfolios across industries and geographies. Actors located toward the upper-right of the plot therefore exhibit the broadest and most diversified targeting behavior.

Table 8. AIAI-3 targeting breadth and portfolio diversification scores (2025), ordered by highest score.

Rank	Actor	V	Countries	Sectors	H_{ctry}	H_{sec}	J_{ctry}	J_{sec}	$p_1^{(\text{ctry})}$	$p_1^{(\text{sec})}$	AIAI-3
1	Qilin	9	6	6	1.733	1.667	0.967	0.930	0.250	0.375	3
2	Devman	7	6	5	1.748	1.550	0.976	0.963	0.286	0.286	3
3	CL0P	9	4	6	1.149	1.677	0.829	0.936	0.556	0.333	2
4	Kill Security	8	4	4	1.255	1.213	0.905	0.875	0.375	0.500	2
5	DragonForce	6	4	4	1.330	1.330	0.959	0.959	0.333	0.333	2
6	Everest	7	3	5	1.011	1.561	0.920	0.970	0.500	0.333	2
7	NightSpire	7	3	4	1.011	1.330	0.920	0.959	0.500	0.333	2
8	Kazu	6	3	3	1.011	1.011	0.920	0.920	0.500	0.500	1
9	Medusa	5	3	4	0.950	1.332	0.865	0.961	0.600	0.400	1
10	INC Ransom	8	6	3	1.748	0.796	0.976	0.725	0.286	0.714	1

Pielou's evenness indices J_{ctry} and J_{sec} are reported alongside the raw entropy values in Table 8 as a normalization reference. The consistently high evenness values observed across actors ($J > 0.85$ for most actors) indicate that the entropy differences reflect genuine distributional variation rather than artefacts of differing category counts, supporting the comparability of fixed ordinal thresholds across actors with different numbers of observed sectors or countries.

Two actors, *Qilin* and *Devman*, attain the maximum AIAI-3 score of 3 and occupy this upper-right region. Both demonstrate high entropy across sectors and countries alongside low concentration in any single category, indicating balanced multi-market activity rather than specialization. *Qilin* distributes victims across six countries and six sectors with limited dominance (country top-1 = 0.25; sector top-1 = 0.375), while *Devman* shows similarly even exposure across six countries and five sectors. These profiles reflect systematic diversification in both geographic and sectoral targeting.

Five additional actors receive score 2, representing moderate diversification. These actors typically span multiple countries and sectors but retain noticeable concentration in one or two dominant categories. For example, *CL0P* exhibits relatively high sector entropy but lower geographic entropy, indicating broad industry coverage concentrated within fewer countries. Others, such as *Kill Security* and *DragonForce*, display balanced yet intermediate entropy levels across both dimensions, consistent with partially diversified portfolios.

Actors assigned score 1 show more concentrated strategies. Their entropy values cluster toward the lower-left of the figure, reflecting narrower targeting footprints where a small number of sectors or countries account for a substantial share of victims. For instance, *INC Ransom* demonstrates high geographic spread but low sector entropy, suggesting

cross-border activity within a limited set of industries, while *Kazu* and *Medusa* exhibit constrained diversification across both dimensions.




Figure 4. AIAI-3 targeting breadth and portfolio diversification scores for the top ransomware actors (2025).

Overall, AIAI-3 clearly separates broadly diversified campaigns from specialized or localized operations. By focusing on entropy-based measures of portfolio dispersion rather than activity volume or tempo, the indicator captures structural differences in targeting strategy that are independent of operational intensity.

5.4. AIAI-4 Results: Temporal Scaling Dynamics of Ransomware Activity

Table 9 and Figure 5 present the highest AIAI-4 scores observed in 2025. The indicator captures the temporal scaling dynamics of ransomware activity, combining early-expansion speed, sustained activity, and acceleration toward peak disclosure.

No actor attains the maximum score of 3 during the observation period. The highest score observed is 2, achieved only by *Kazu* and *Everest*. Both actors exhibit comparatively fast early expansion together with sustained activity across multiple months. *Kazu* reaches five victims immediately ($t_5 = 0$ days) and remains active for two months, indicating an immediate burst of disclosures followed by short-term persistence. *Everest* similarly demonstrates rapid scaling ($t_5 = 27.3$ days) combined with continued activity across three months. In Figure 5, these actors occupy the upper portion of the distribution, reflecting comparatively fast ramp-up and sustained operational presence.

Table 9. AIAI-4 temporal scaling dynamics scores (2025), ordered by highest score.

Rank	Actor	V	t ₅ (Days)	Span (Days)	M _a	Peak Lag (Months)	AIAI-4
1	Kazu	6	0.0	48	2	0	2
2	Everest	7	27.3	102	3	1	2
3	CL0P	9	15.7	16	1	0	1
4	Kill Security	8	204.8	277	4	7	1
5	Devman	7	218.2	266	4	3	1
6	INC Ransom	8	160.5	205	6	5	1
7	NightSpire	7	224.7	245	4	3	1
8	DragonForce	6	273.4	288	5	9	1
9	Qilin	9	145.9	176	4	5	0
10	Medusa	5	297.9	298	3	9	0

t₅: days from first observed disclosure to the fifth victim; Span: days between first and last observed disclosure; M_a: number of months with at least one disclosure; Peak lag: months between first activity and peak disclosure month.




Figure 5. AIAI-4 temporal scaling dynamic scores for the top ransomware actors (2025).

Most remaining actors receive a score of 1, indicating slower or less consistent temporal scaling dynamics. Although several groups operate across multiple months, their time-to-spread values are substantially longer, often exceeding 150–200 days, suggesting gradual accumulation rather than rapid early expansion. Actors such as *INC Ransom*, *NightSpire*, and *DragonForce* display extended campaign spans but limited early acceleration toward peak disclosure activity.

Two actors, *Qilin* and *Medusa*, receive a score of 0. These campaigns exhibit both slow initial expansion and weaker persistence relative to other actors in the dataset.

Overall, AIAI-4 suggests that rapid temporal scaling of ransomware activity is relatively uncommon in 2025. Most actors expand incrementally over extended periods, while only a small subset demonstrate the faster early-growth patterns associated with highly coordinated campaigns.

5.5. Cross-Indicator Behavioral Profiles of Actor Activity

To synthesize the multi-dimensional results, we compare actors across all four AIAI components to identify recurring behavioral patterns rather than produce a single composite ranking. Because each indicator captures a distinct construct (social-engineering enablement, operational tempo, targeting breadth, and scaling dynamics), scores are interpreted jointly as behavioral profiles rather than summed into a predictive or weighted index.

Table 10 and Figure 6 present the individual AIAI-1–AIAI-4 scores for the ten actors with the highest overall cross-indicator activity. Scores are interpreted jointly to characterize multi-dimensional behavioral profiles rather than to derive a composite ranking. The “Sum” column in Table 10 is included solely to provide a stable display ordering that groups actors with higher cross-indicator activity for comparative readability. It does not represent a weighted composite score, and higher totals should not be interpreted as proportionally greater threat levels, as strong signals in a single indicator may reflect more operationally significant behavior than moderate scores across multiple dimensions.

Table 10. Cross-indicator behavioral profiles of ransomware actors (top 10 by combined AIAI activity, 2025).

Actor	AIAI-1	AIAI-2	AIAI-3	AIAI-4	Sum (Ordering Only)
Qilin	2	2	3	0	7
CL0P	1	3	2	1	7
Everest	1	2	2	2	7
Kill Security	1	2	2	1	6
Kazu	1	2	1	2	6
DevMan	1	0	3	1	5
NightSpire	1	0	2	1	4
DragonForce	1	0	2	1	4
INC Ransom	1	0	1	1	3
Ransomhub	0	3	0	0	3




Figure 6. Heatmap of cross-indicator AIAI behavioral profiles for the top ransomware actors (2025).

Several distinct patterns are observable. *CL0P* exhibits the strongest operational tempo signal (AIAI-2 = 3), combined with moderate diversification (AIAI-3 = 2) and limited social-engineering and scaling evidence (AIAI-1 = 1; AIAI-4 = 1). This configuration is consistent with burst-style disclosure campaigns concentrated within short time windows.

Qilin demonstrates the most balanced and socially enabled profile, combining the highest targeting breadth (AIAI-3 = 3) with the only elevated social-engineering score (AIAI-1 = 2) alongside sustained tempo (AIAI-2 = 2). This combination suggests systematic, repeatable operations spanning multiple sectors and geographies rather than isolated release bursts.

Everest and *Kill Security* display comparatively even scores across multiple dimensions. Both show moderate tempo and diversification (AIAI-2 = 2; AIAI-3 = 2), with *Everest* additionally exhibiting stronger cross-sector scaling (AIAI-4 = 2). These profiles indicate sustained and steady activity without the extreme burst behavior observed for *CL0P*.

A second group, including *DevMan*, *NightSpire*, and *DragonForce*, shows higher diversification signals (AIAI-3 of 2 or higher) but weaker tempo and scaling metrics (AIAI-2 of 1 or lower; AIAI-4 = 1), suggesting broader targeting breadth without comparable operational intensity.

Finally, actors such as *INC Ransom* and *Ransomhub* exhibit strength primarily along a single dimension. *Ransomhub* scores highly only in tempo (AIAI-2 = 3), while *INC Ransom* shows uniformly low-to-moderate values across indicators, reflecting narrower or lower-intensity observable behavior.

Overall, the cross-indicator comparison reinforces that ransomware capability is heterogeneous and multi-dimensional. Different actors achieve similar apparent impact through distinct combinations of tempo, diversification, scaling, and social-engineering characteristics, which would not be visible from victim counts or single metrics alone.

5.6. External Consistency and Validation

To assess whether AIAI scores reflect meaningful operational patterns, the results were compared against observable dataset proxies and external threat intelligence reporting.

Actors with higher AIAI scores consistently exhibit higher victim counts, broader geographic reach, and greater sectoral diversification. Notably, actors such as *CL0P*, *Qilin*, and *Kill Security* are both highly ranked in AIAI and identified as dominant groups in independent regional analysis [32].

This alignment is consistent with broader threat intelligence findings [29,30], where ransomware activity is concentrated among a small number of high-impact actors.

These results indicate that the AIAI framework captures meaningful differences in operational scale and behavior rather than reflecting arbitrary threshold effects. The framework should be interpreted as an observational proxy based on publicly disclosed activity rather than a direct measure of total intrusion volume.

6. Discussion

This study applied the AIAI, an actor-level behavioral scoring framework, to ransomware groups active during 2025, enabling systematic comparison of their operational characteristics and relative threat levels. This section interprets the results in light of prior research and discusses their broader implications for threat intelligence and future work.

6.1. Overall Interpretation of AIAI Scores

The scoring results of AIAI-1 through AIAI-4 reveal clear differentiation in ransomware actor behavior and operational sophistication. Across the four indicators, actors

separate into distinct tiers ranging from highly coordinated and scalable operations to more limited or episodic activity.

Actors with higher overall AIAI profiles consistently combine elevated disclosure tempo, broader cross-country and cross-sector targeting, and stronger scaling dynamics. These patterns reflect greater organizational maturity and more structured campaign execution. In contrast, lower-scoring actors tend to exhibit narrower targeting, slower operational cadence, or reduced persistence, consistent with more opportunistic or constrained operations.

Importantly, these differences are not explained by victim counts alone. Several actors with similar disclosure volumes display markedly different AIAI profiles, demonstrating that activity volume is an insufficient proxy for operational capability. By integrating tempo, diversification, scaling, and social-engineering enablement, AIAI provides a concise, behavior-driven basis for comparing ransomware actors and supporting threat intelligence prioritization.

6.2. Implications of AI-Enhanced TTP Evolution

Across the evaluated actors, the results suggest an emerging convergence between traditional ransomware tactics and AI-enabled operational capabilities. While documented evidence of GenAI-enabled activity was uneven across groups, the presence of scalable social-engineering workflows captured by AIAI-1 indicates that automation and AI are beginning to shape how certain actors generate, adapt, and deploy deceptive content.

Although AI adoption was most clearly observable in the initial access phase, particularly in phishing and impersonation activity, the broader AIAI results suggest that similar amplification effects may extend to other stages of ransomware operations, including victim selection, campaign orchestration, and cross-sector expansion. As these capabilities mature, ransomware TTPs are likely to become increasingly adaptive, personalized, and scalable.

These findings have important implications for existing threat assessment approaches, which often assume relatively static or repetitive attacker behavior. As operational workflows become more flexible and data-driven, assessment frameworks that rely solely on historical indicators or isolated technical artifacts may lose effectiveness. The results therefore support incorporating AI-related amplification signals as a core dimension of ransomware threat intelligence rather than treating them as ancillary or speculative features.

6.3. Broader Implications and Future Research Directions

By proposing a TTP-informed, actor-level scoring framework, this study contributes a structured approach for evaluating ransomware groups as organized cyber threat actors rather than isolated malware instances. The inclusion of AI-associated behavioral signals enables more nuanced comparisons of operational scaling beyond victim counts alone.

From a defensive governance perspective, AIAI profiles are intended to support prioritization and monitoring decisions rather than trigger automated responses. Actors scoring 2 or higher across multiple indicators should be treated as high-priority threats and prioritized for continuous monitoring and resource allocation. Actors exhibiting extreme values in individual indicators (e.g., AIAI-2 = 3) should trigger event-driven alerting and short-term escalation during periods of concentrated activity. Specifically, elevated AIAI-1 scores should inform phishing-resilience planning, AIAI-3 should guide sector- and region-specific defensive priorities, and AIAI-4 should trigger early-warning monitoring for emerging regional surges. Actors with consistently low scores may be monitored at a baseline level without escalated intervention. These thresholds provide a practical starting point for operationalizing AIAI within CTI workflows and should be adapted to organizational risk tolerance and threat context.

Future research should expand the dataset across regions and time periods and explore dynamic scoring models that capture temporal changes in attacker behavior. Future studies may also explore how AI-driven defensive techniques can be systematically aligned with attacker-side AI capabilities, particularly in areas such as adaptive detection, automated threat intelligence generation, and real-time response prioritization.

7. Limitations

This study identifies several limitations that should be considered when interpreting the AIAI framework and its results.

- Selective disclosure bias: The dataset is based on leak-site disclosures and reflects only publicly revealed victims. As actors may selectively publish or delay disclosure, the data captures observable publication behavior rather than the full scope of ransomware activity. AIAI scores should therefore be interpreted as conservative lower bounds.
- Geographic and temporal scope: The analysis is limited to a single year (2025) and region (Arab countries), which constrains generalizability. Longitudinal and multi-region analysis is needed to assess stability across contexts.
- Observation timing confound: AIAI-4 metrics depend on first observed activity within the window, potentially compressing metrics for late-appearing actors.
- Evidence availability bias: AIAI-1 relies on publicly attributed artifacts. Uneven OSINT availability may lead to underestimation of actors with lower public visibility.
- Group-level attribution: Leak-site disclosures are attributed at the ransomware group level. In RaaS ecosystems, observed behavior reflects aggregate affiliate activity and cannot be disentangled using public data.
- Coverage and missingness: Changes in monitored leak sites (e.g., migration or take-down) may result in incomplete capture of activity. Cross-validation with OSINT improves coverage but does not eliminate missingness.
- External validation: The framework is not validated against independent ground-truth measures such as confirmed intrusion volumes. Validation using external datasets remains an important direction for future work.

8. Conclusions

This study proposed an actor-level scoring framework for evaluating ransomware operations based on observable TTPs and indicators of AI-assisted operational amplification. By introducing AIAI, this research moves beyond artifact-centric or victim-count-based assessments and enables systematic comparison of ransomware groups in terms of operational complexity, scaling behavior, and targeting breadth.

Applying the proposed scoring model across all observed actors, with comparative profiling of the ten most active groups, revealed measurable differences in operational behavior. Higher-scoring actors exhibited faster cross-country and cross-sector spread, higher disclosure tempo, and stronger indicators of scalable orchestration. In contrast, groups with lower AIAI scores exhibited more limited operational scope and reduced observable impact. These results suggest that the proposed framework can effectively distinguish between higher-tier and lower-tier ransomware operations based on empirically observable behavior rather than isolated technical artifacts.

Importantly, the findings indicate that ransomware groups with higher AIAI scores tend to operate across multiple organizations and national contexts within the observation period, reflecting greater organizational maturity and operational reach. From a CTI perspective, this actor-centric view provides researchers and practitioners with a more intuitive and actionable understanding of which groups employ more sophisticated techniques and

exert broader operational influence, supporting prioritization, monitoring, and strategic response planning.

While this study focuses on ransomware actively observed within the Arab region, the proposed framework is not region-specific and can be extended to analyze global ransomware ecosystems. Future work may expand the dataset to include additional regions and longer observation periods, enabling longitudinal analysis of how ransomware operations evolve and how AI-assisted amplification shapes global threat dynamics.

Author Contributions: The authors confirm the contribution to this paper as follows: Conceptualization, K.K., M.M. and I.A.; methodology, K.K., M.M., I.A., F.A.A. and S.L.; validation, K.K., M.M., F.A.A. and S.L.; formal analysis, M.M.; investigation, K.K. and M.M.; data curation, M.M. and I.A.; writing—original draft preparation, M.M. and S.L.; writing—review and editing, K.K. and F.A.A.; visualization, M.M.; supervision, K.K.; project administration, K.K.; funding acquisition, K.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Naif Arab University for Security Sciences under grant NAUSS-25-R04.

Data Availability Statement: The raw dataset is not publicly available due to privacy and legal constraints. Actor-level aggregate statistics are reported in Tables 6–10. Further data may be made available upon reasonable request to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest. Author Faisal Abdulaziz Alfouzan was employed by the Private Company, Riyadh 11564, Saudi Arabia. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. The funders had no role in the design of this study; in the collection, analyses, or interpretation of data; in the writing of this manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

AIAI	AI-Amplification Indicator
APT	Advanced Persistent Threat
CISA	Cybersecurity and Infrastructure Security Agency
CTI	Cyber Threat Intelligence
CVSS	Common Vulnerability Scoring System
DLS	Dedicated Leak Sites
GenAI	Generative Artificial Intelligence
ICS	Industrial Control Systems
LLMs	Large Language Models
NLP	Natural Language Processing
RaaS	Ransomware as a Service
TRAM	Threat Report ATT&CK Mapper
TTPs	Tactics, Techniques, and Procedures

References

1. Kerns, Q.; Payne, B.; Abegaz, T. Double-extortion ransomware: A technical analysis of maze ransomware. In *Proceedings of the Future Technologies Conference*; Springer: Berlin/Heidelberg, Germany, 2021; pp. 82–94. [\[CrossRef\]](#)
2. Patsakis, C.; Arroyo, D.; Casino, F. The Malware as a Service Ecosystem. In *Malware: Handbook of Prevention and Detection*; Gritzalis, D., Choo, K.K.R., Patsakis, C., Eds.; Springer Nature: Cham, Switzerland, 2025; pp. 371–394. [\[CrossRef\]](#)
3. Kasula, V.K.; Yenugula, M.; Konda, B. Advanced Cybersecurity Detection Techniques. In *Law and Regulation of AI, Blockchain, and Digital Rights*; IGI Global Scientific Publishing: Hershey, PA, USA, 2026; pp. 251–290.
4. Gupta, M.; Akiri, C.; Aryal, K.; Parker, E.; Praharaj, L. From ChatGPT to ThreatGPT: Impact of generative AI in cybersecurity and privacy. *IEEE Access* **2023**, *11*, 80218–80245. [\[CrossRef\]](#)

5. Jabir, R.; Le, J.; Nguyen, C. Phishing attacks in the age of generative artificial intelligence: A systematic review of human factors. *AI* **2025**, *6*, 174. [\[CrossRef\]](#)
6. Marchal, N.; Xu, R.; Elasmr, R.; Gabriel, I.; Goldberg, B.; Isaac, W. Generative AI misuse: A taxonomy of tactics and insights from real-world data. *arXiv* **2024**, arXiv:2406.13843. [\[CrossRef\]](#)
7. Strom, B.E.; Applebaum, A.; Miller, D.P.; Nickels, K.C.; Pennington, A.G.; Thomas, C.B. *MITRE ATT&CK: Design and Philosophy*; Technical Report; The MITRE Corporation: McLean, VA, USA, 2018.
8. Kim, K.; Lee, S.; Ramachandran, S.; Alzahrani, I. Cryptocurrency-driven ransomware syndicates operating on the darknet: A focused examination of the Arab world. *Egypt. Inform. J.* **2025**, *30*, 100665. [\[CrossRef\]](#)
9. Lee, S.; Mujammami, A.A.H.; Kim, K. Leveraging Social Networks for Cyber Threat Intelligence: Analyzing Attack Trends and TTPs in the Arab World. *IEEE Access* **2025**, *13*, 5679–5693. [\[CrossRef\]](#)
10. Husari, G.; Al-Shaer, E.; Ahmed, M.; Chu, B.; Niu, X. Ttpdrill: Automatic and accurate extraction of threat actions from unstructured text of cti sources. In *Proceedings of the 33rd Annual Computer Security Applications Conference*; Association for Computing Machinery: New York, NY, USA, 2017; pp. 103–115. [\[CrossRef\]](#)
11. Kim, K.; Alfouzan, F.A.; Kim, H. Cyber-attack scoring model based on the offensive cybersecurity framework. *Appl. Sci.* **2021**, *11*, 7738. [\[CrossRef\]](#)
12. Alshamrani, A.; Myneni, S.; Chowdhary, A.; Huang, D. A survey on advanced persistent threats: Techniques, solutions, challenges, and research opportunities. *IEEE Commun. Surv. Tutor.* **2019**, *21*, 1851–1877. [\[CrossRef\]](#)
13. Alraizza, A.; Algarni, A. Ransomware Detection Using Machine Learning: A Survey. *Big Data Cogn. Comput.* **2023**, *7*, 143. [\[CrossRef\]](#)
14. Lusthaus, J. Beneath the dark web: Excavating the layers of cybercrime’s underground economy. In *Proceedings of the 2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*; IEEE: New York, NY, USA, 2019; pp. 474–480. [\[CrossRef\]](#)
15. Martín Gallego, J. Data Extraction from TOR-Based Ransomware Websites and Statistical Analysis Platform. Master’s Thesis, University of Valladolid (UVaDOC), Valladolid, Spain, 2025.
16. Shah, S.; Madiseti, V.K. MAD-CTI: Cyber Threat Intelligence Analysis of the Dark Web Using a Multi-Agent Framework. *IEEE Access* **2025**, *13*, 40158–40168. [\[CrossRef\]](#)
17. Tariq, U.; Ullah, I.; Yousuf Uddin, M.; Kwon, S.J. An effective self-configurable ransomware prevention technique for IOMT. *Sensors* **2022**, *22*, 8516. [\[CrossRef\]](#) [\[PubMed\]](#)
18. Marinho, R.; Holanda, R. Automated emerging cyber threat identification and profiling based on natural language processing. *IEEE Access* **2023**, *11*, 58915–58936. [\[CrossRef\]](#)
19. Zhang, H.; Hong, H.; Shen, B.; Rojas, M.E.B.; Tahoun, E.; Zakia, U. NLP-based Cyber Threat Data Analysis with MITRE ATT&CK Techniques. In *Proceedings of the 2025 IEEE 16th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*; IEEE: New York, NY, USA, 2025; pp. 389–395.
20. Al-Sada, B.; Sadighian, A.; Oligeri, G. Analysis and characterization of cyber threats leveraging the MITRE ATT&CK database. *IEEE Access* **2023**, *12*, 1217–1234. [\[CrossRef\]](#)
21. Anthi, E.; Williams, L.; Rhode, M.; Burnap, P.; Wedgbury, A. Adversarial attacks on machine learning cybersecurity defences in industrial control systems. *J. Inf. Secur. Appl.* **2021**, *58*, 102717. [\[CrossRef\]](#)
22. Jiang, Y.; Meng, Q.; Shang, F.; Oo, N.; Minh, L.T.H.; Lim, H.W.; Sikdar, B. MITRE ATT&CK Applications in Cybersecurity and The Way Forward. *arXiv* **2025**, arXiv:2502.10825. [\[CrossRef\]](#)
23. Bukhari, M.; Sattar, S.; Saleem, S.; Khan, K.Z.; Khan, A. The Impact of Cybercrime Incidents and Artificial Intelligence adoption on Organizational Performance: A Mediation and moderation model. *J. Excell. Soc. Sci.* **2024**, *3*, 191–210. [\[CrossRef\]](#)
24. Alzahrani, S.; Xiao, Y.; Sun, W. An analysis of conti ransomware leaked source codes. *IEEE Access* **2022**, *10*, 100178–100193. [\[CrossRef\]](#)
25. MITRE Center for Threat-Informed Defense. Threat Report ATT&CK Mapper (TRAM). Official Open-Source Project for Automated TTP Mapping. 2021. Available online: <https://github.com/center-for-threat-informed-defense/tram> (accessed on 9 February 2026).
26. Cybersecurity and Infrastructure Security Agency (CISA). Critical Infrastructure Sectors. 2024. Available online: <https://www.cisa.gov/topics/critical-infrastructure-security-and-resilience/critical-infrastructure-sectors> (accessed on 3 February 2026).
27. Dontu, S.; Addula, S.R.; Pareek, P.K.; Vallabhaneni, R.; Adnan, M.M. Attack detection from Internet of Things using TPE based self-attention based bidirectional long-short term memory. In *Proceedings of the 2024 International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS)*; IEEE: New York, NY, USA, 2024; pp. 1–6.
28. Sophos. Sophos MDR Tracks Ongoing Campaign by Qilin Affiliates Targeting ScreenConnect. Threat Research Blog. 2024. Available online: <https://www.sophos.com/en-us/blog/sophos-mdr-tracks-ongoing-campaign-by-qilin-affiliates-targeting-screenconnect> (accessed on 31 January 2026).
29. Verizon. 2025 Data Breach Investigations Report. 2025. Available online: <https://www.verizon.com/business/resources/reports/dbir/> (accessed on 31 January 2026).
30. European Union Agency for Cybersecurity (ENISA). *ENISA Threat Landscape 2025*; ENISA: Athens, Greece, 2025. Available online: <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025> (accessed on 8 February 2026).

31. Flashpoint Intelligence Team. *LockBit 5.0 Analysis: Technical Deep Dive into the RaaS Giant's Latest Upgrade*; Flashpoint: New York, NY, USA, 2025. Available online: <https://flashpoint.io/blog/lockbit-5-0-analysis-technical-deep-dive-into-the-raas-giants-latest-upgrade/> (accessed on 9 February 2026).
32. Al-Morjan, A.; Kim, K.; Lee, S.; Moallim, M.; Alzahrani, I. *Ransomware Trends Report in Arab Countries 2025*; Naif University Press: Riyadh, Saudi Arabia, 2026. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.