High Stakes, Low Certainty: Evaluating the Efficacy of High-Level Indicators of Compromise in Ransomware Attribution

Horst, Max van der; Kho, Ricky; Gadyatskaya, Olga; Mollema, Michel; Eeten, Michel van; Zhauniarovich, Yury

Publication date
2025

Document Version
Final published version

Published in
USENIX Security '25

USENIX

THE ADVANCED COMPUTING
SYSTEMS ASSOCIATION

High Stakes, Low Certainty: Evaluating the Efficacy of High-Level Indicators of Compromise in Ransomware Attribution

Max van der Horst, *Delft University of Technology;* Ricky Kho, *Sogeti;*
Olga Gadyatskaya, *Leiden University;* Michel Mollema, *Northwave Cybersecurity;*
Michel Van Eeten and Yury Zhauniarovich, *Delft University of Technology*

https://www.usenix.org/conference/usenixsecurity25/presentation/van-der-horst

This paper is included in the Proceedings of the
34th USENIX Security Symposium.

August 13–15, 2025 • Seattle, WA, USA

978-1-939133-52-6

Open access to the Proceedings of the
34th USENIX Security Symposium is sponsored by USENIX.

High Stakes, Low Certainty: Evaluating the Efficacy of High-Level Indicators of Compromise in Ransomware Attribution

Max van der Horst
Delft University of Technology

Ricky Kho
Sogeti

Olga Gadyatskaya
Leiden University

Michel Mollema
Northwave Cybersecurity

Michel van Eeten
Delft University of Technology

Yury Zhauniarovich
Delft University of Technology

Abstract

As ransomware attacks grow in frequency and complexity, accurate attribution is crucial. Victim organizations often feel compelled to pay ransom, but must first attribute the attack and conduct sanction screening to ensure the threat actor receiving the payment is not a sanctioned entity, avoiding severe legal and financial risks. This cyber threat actor attribution process typically relies on Indicators of Compromise (IoCs) matching known threat profiles. However, the emergence of the Ransomware-as-a-Service (RaaS) ecosystem and rebranding behavior complicate attribution for sanction screening.

Our mixed-methods study, combining interviews with 20 experts with an analysis of ransomware incident reports, reveals significant challenges and limitations in the current attribution process. High-level IoCs, widely regarded as more reliable, lack the necessary specificity for accurate attribution, leading to potential risks of misattribution. Practitioners rely on lower-level IoCs, which provide clearer links to threat actors but are highly volatile, further complicating sanction enforcement. These challenges highlight the need for urgent improvements in the attribution and sanction processes.

To mitigate these risks, we offer recommendations aimed at enhancing data-sharing practices, improving attributions frameworks, and refining the sanction violation policy to better support sanction screening efforts. While we do not recommend paying ransomware actors, we acknowledge that some organizations may face pressures to do so in certain situations. In such cases, it is vital to ensure legal compliance, particularly regarding sanctioned entities. This work aims to help victims of ransomware shield themselves from transgressing against sanctions.

1 Introduction

Ransomware attacks are increasingly impacting the digital as well as the physical landscape [85]. For instance, the infamous ransomware attack on Colonial Pipeline in May 2021 resulted in fuel shortages at filling stations across several US states and an increase in fuel prices to the highest levels since 2014 [25]. The increased prevalence of these attacks [85] is partially due to the creation of Ransomware-as-a-Service (RaaS) models, pushed by notorious cybercriminal groups like BlackBasta [37], BlackCat [22], and LockBit [20]. Moreover, ransomware remains highly profitable. In 2024, 63% of ransom demands exceeded $1 million, with a median payment of $2 million – a fivefold increase compared to 2023 [84]. Still, the business disruption caused by a ransomware attack often forces victim organizations to accept the cybercriminal's demands and pay ransom. Sophos reports that 56% of affected organizations make payment to recover access to their data [84], while 62% of CISOs admit that their organization would consider settling ransom demand [27].

However, paying ransom introduces significant legal concerns. Policymakers, such as the U.S. Office for Foreign Assets Control, sanction specific ransomware actors in order to deter, disrupt, and diminish ransomware campaigns [60, 63, 88]. Additionally, embargoes on countries like North Korea impose strict restrictions on financial transactions with these regions [87]. These sanctions complicate the decision-making process for organizations facing ransomware attacks, as they must navigate legal restrictions alongside technical challenges. Victims risk further damages due to fines or even criminal prosecution for sanction violations [60].

To avoid such violations, victims require cybersecurity analysts to perform *Cyber Threat Actor (CTA) attribution*, linking the attack to a specific group or jurisdiction. This growing demand for attribution has led to the emergence of new services, known as *sanction screening*, offered by security companies to determine whether the perpetrator is under sanctions and making payment to it thus being a potential violation [19, 49, 62]. Attribution of attacks is based on forensic analysis that identifies *Indicators of Compromise (IoCs)* – pieces of evidence that suggest malicious activity within a network or system – and compares them to known IoCs associated with particular threat actors.

Researchers and practitioners typically classify indicators into different levels [36, 61]: *low* and *high* IoCs. Examples

of low-level indicators include file hashes, domain names, IP addresses, and registry keys. High-level indicators, such as *Tactics, Techniques, and Procedures (TTPs)*, are more abstract and require interpretation or correlation with other pieces of data to be useful in active defense efforts such as CTA attribution.

Frameworks such as the Pyramid of Pain [10], Rid and Buchanan's Q-model [72], and other threat intelligence literature [5, 12, 38, 91], emphasize that high-level IoCs contribute more substantially to threat actor investigations than low-level IoCs. For example, the Pyramid of Pain argues that "*Not all indicators are created equal,[...] some of them are far more valuable than others.*" [10] This reasoning originates from the idea that low-level indicators are transient and often easily manipulable, while high-level IoCs capture a threat actor's behavioral patterns and operational methodology, which is much more difficult ("painful") for the actor to change [10]. Although models like the Pyramid of Pain and Q-model are widely recognized in industry and academia as applicable to all threats, including ransomware [16, 40, 41, 73, 82, 94], empirical research specifically examining their assumptions in the context of ransomware remains limited. This underscores the need for further investigation, as ransomware dynamics are likely distinct from typical cybercrime due to the tendency of ransomware actors to reveal themselves.

In this paper, we present the first empirical study into the ransomware attribution process in the industry, and specifically the role of IoCs in this process. To gain a more practical insight into the attribution process, we collaborated with a cybersecurity company, which specializes in ransomware forensics and sanction screening, and conducted 15 semi-structured interviews with analysts from it. In addition, we corroborated those findings by interviewing 5 more analysts from 5 similar companies. Our interviewees confirmed that they find TTPs to be too generic for precise attribution. Instead, they rely on IoCs specific to ransomware, like the ransom note, and other lower-level IoCs. They also reported concerns about the current attribution process for sanction screening, as the legal framework for this process, especially the required level of certainty, is not clearly defined.

To further explore the role of TTPs in attribution, we complemented the interview study with the document analysis of incident reports. Using the MITRE ATT&CK framework [18], we extracted TTPs associated with individual attacks from ransomware incident reports produced by our partnering company in 2023. Next, we compared the TTPs of the attacks attributed to one ransomware threat actor (RTA) with each other and with the TTPs reported within the #StopRansomware campaign [21] of the U.S. Cybersecurity and Infrastructure Security Agency (CISA). This comparison evaluated the overlap of ransomware TTPs associated to different RTAs or reported by different organizations. Our findings reveal a substantial overlap between IoCs discovered for different RTAs, further exacerbated by a far-from-perfect overlap of TTPs observed for the same perpetrator. The report analysis demonstrates that TTPs, indeed, might not aid in attribution as much as expected.

Finally, we integrated insights from both components of the study to develop a comprehensive and practical understanding of the ransomware attribution process, including the use of high-level indicators in this process and, therefore, in sanction screening. We challenge the established frameworks that prioritize high-level IoCs, such as those suggested by the widely adopted Pyramid of Pain. While this model emphasizes that high-level IoCs are more valuable because they are harder for attackers to change, our empirical findings suggest that these indicators might not be as effective as traditionally thought for ransomware contexts. In combination with the volatility of low-level IoCs, this creates a vulnerability in the sanction screening process that prevents victims from being certain about violating sanctions. By presenting countervailing evidence to this long-standing belief, we aim to stimulate a reevaluation of attribution strategies in ransomware defense, highlight the necessity of improved attribution methodologies, and reconsideration of sanction enforcement to better protect ransomware victims and facilitate better decision-making.

Thus, the most important contributions of this work to the study of ransomware attribution and its implications on sanction enforcement are the following:

- **An empirical investigation into ransomware attribution process for sanction screening:** Our empirical study, combining 20 expert interviews with an analysis of 27 incident reports, dissects the ransomware attribution process followed by practitioners and the IoCs they rely on in this process, highlighting the pertinent challenges and complexities.

- **Challenging the value of TTPs in ransomware attribution:** Our findings challenge the belief in high-level IoCs for ransomware attribution, showing they are too ambiguous, while low-level IoCs, though preferred, are too volatile for reliability.

- **Expanding the Pyramid of Pain framework for ransomware attribution:** We propose an expansion to the Pyramid of Pain with IoCs specific for ransomware attribution.

- **Sanction policy implications:** We highlight the risks of relying on unreliable attribution methods, which can lead to disproportionate penalties in sanction enforcement, posing legal and financial risks for victims.

- **Recommendations for improved attribution:** We advocate for better intelligence sharing, standardized reporting, and revising attribution and sanction practices to meet the challenges of the ransomware ecosystem.

2 Background

Cyber sanctions are measures imposed by governments or international organizations against individuals, entities, or nations responsible for cyberattacks or cybercrime activities in order to change their behavior [58]. These sanctions can include asset freezes, travel bans, and restrictions on doing business with designated persons or entities. The objective of cyber sanctions is to deprive criminals of funds and further deter malicious cyber activities by increasing the cost of engaging in such behavior, thereby enhancing national and international cybersecurity.

Sanctions imposed on ransomware criminals may prohibit making ransom payments, thus discouraging individuals and organizations from supporting this criminal activity [88]. Governments enact these measures to stop the cycle of extortion. Breaking these sanctions is considered a serious crime in many countries, resulting in significant fines or even imprisonment. Therefore, victim organizations use sanction screening to verify whether or not a sanction has been imposed on the attacking group. Typically, this sanction screening is done by comparing with lists (such as the OFAC lists by the US Department of the Treasury [64]) that contain information about the imposed sanctions, including identifying details of the sanctioned group, individual, or state.

Although the sanctioned entities are known, establishing a direct link between them and an attack remains challenging. *CTA attribution* is the process of collecting, analyzing, and associating data from malicious cyber activities to identify and determine the identity, location, or other identifying information of an attacker or intermediary [52, 72, 73, 94, 95]. CTA attribution is often a mix of technical attack analysis and threat actor profiling [82]. Layton and Watters defined two types of attribution outcomes: *absolute* and *relative* attribution. With absolute attribution, the actual actor is identified. In the case of relative attribution, the attribution remains relative to a previous incident [41]. Examples of these in practice are the activity clustering conducted by companies such as Microsoft [56] and Mandiant [44].

For such analyses and profiling to be effective, IoCs, or *indicators*, are essential. Indicators are pieces of evidence or artifacts used to identify and attribute cyber threats to specific actors, groups, or campaigns. Researchers and practitioners typically classify IoCs into *low-level* (or *direct*) and *high-level* (or *indirect*), such as in [36, 61, 70]. Low-level indicators are straightforward, tangible, and concrete pieces of evidence that can be directly related to the activity or presence of a malicious actor. High-level indicators, as opposed to low-level indicators, are more abstract and require interpretation or correlation with other pieces of data in order to be useful in CTA attribution – but they aim to capture the key characteristics of CTA behavior. Examples of low-level indicators are hashes, domain names, IP addresses, and specific signatures, whereas high-level indicators may involve TTPs. This intuition of low and high IoCs is captured by the Pyramid of Pain model [10]. It positions different IoCs hierarchically according to the difficulty for CTAs to pivot to another IoC at the same level: while it is trivial to change the hash value of a file or switch to another IP address, it involves more effort to switch to completely new malware, and it is even more challenging to adopt new behaviors (Tactics, Techniques, and Procedures) [10]. Another important framework for the practitioners is the MITRE ATT&CK [18], which is a global reference for describing adversarial tactics, techniques, and sub-techniques used in cyberattacks [3]. These components, structured hierarchically, are key to this study.

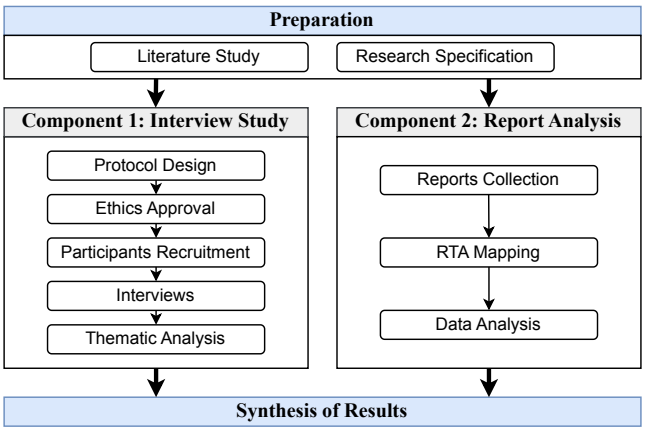

Figure 1: Overview of our research methodology.

Digital forensics plays a crucial role in CTA attribution by analyzing system artifacts like log files, registry entries, and malware samples to trace the attack's origin and determine a root cause [65, 79]. Unfortunately, attribution through digital forensics faces various limitations, particularly due to the volatility of low-level indicators, which adversaries can easily modify or obscure [77]. To address this, attribution techniques based on high-level indicators such as broader behavioral patterns are recommended [41], as suggested by models like the Pyramid of Pain [10] and Rid's Q-model [72]. Reporting TTPs for RTAs is also mandated by governing bodies [63]. Yet, though more reliable, attribution using high-level indicators requires extensive data and manual effort [77].

3 Methodology

To investigate the ransomware threat actor (RTA) attribution process and the effectiveness of high-level indicators, we used a mixed-methods approach. This included a qualitative component (semi-structured interviews with security experts involved in ransomware analysis and sanction checks) and a quantitative analysis of real ransomware incident reports. Figure 1 outlines the components of our methodology, which are detailed in this section.

Table 1: Demographics of the interviewees.

ID	Position	Experience	Org. Size
E1	Lead Digital Forensics	13 years	Medium
E2	Lead Cyber Threat Intelligence	16 years	Medium
E3	Digital Forensics	3 years	Medium
E4	Digital Forensics	3 years	Medium
E5	Digital Forensics	3 years	Medium
E6	Digital Forensics	4 years	Medium
E7	Threat Intelligence Analyst	8 years	Medium
E8	Digital Forensics	3 years	Medium
E9	Reverse Engineer	8 years	Medium
E10	Digital Forensics	4 years	Medium
E11	Lead High Tech Crime	5 years	Medium
E12	Manager	13 years	Medium
E13	Cyber Resilience Consultant	9 years	Medium
E14	Head of Threat Intelligence	12 years	Medium
E15	Digital Forensics	3 years	Medium
V1	Manager Security Operations	6 years	Small
V2	Threat Intelligence Analyst	6 years	Large
V3	Threat Intelligence Analyst	5 years	Large
V4	Digital Forensics	12 years	Large
V5	Digital Forensics	8 years	Large

Org. size in staff: small (1-50), medium (51-250), large (250+)

To study the use of indicators in ransomware investigations, we partnered with a leading international cybersecurity company in Europe specializing in ransomware forensics, recovery, and sanction checks. This collaboration provided access to industry experts and real ransomware reports, ensuring interviewees had relevant knowledge and experience.

3.1 Component 1: Interview Study

We conducted expert interviews to gain insights into the attribution process by working together with our partnering company. Fifteen interviewees were selected among its employees based on their roles and at least three years of experience in ransomware investigations, focusing on cyber threat intelligence, digital forensics, or law enforcement. Table 1 provides demographic details for these participants, marked as *E*. For the organization sizes, we used the EU definition of small, medium, or large size enterprises[1]. To confirm data saturation, five additional interviews were conducted with experts adhering to the same criteria from five different companies (marked as *V*). To mitigate duplication bias, we confirmed no overlap in the organizational history among interviewees. All organizations from which the interviewees were recruited have an international focus and are either based in Europe or have a branch office there.

Semi-structured interviews were conducted between November 2023 and July 2024, either in person or via conference calls. Participants were informed about the study's goals and the data we collect, and signed an informed consent form. No compensation was provided for participation.

Interview Structure. The interview questions were organized into three segments: the attribution process for RTAs, strengths and limitations of current practices, and suggestions for improvement. The first segment, informed by the literature review, focused on examining RTA attribution methodologies and their differences from other CTAs. The second segment explored the strengths, limitations, and unique challenges in current attribution practices, while the final segment gathered expert opinions on how to improve the attribution process. Our full interview questionnaire is provided in Appendix A.

Each interview lasted for approximately 90 minutes. The interviews were captured through audio recording to ensure accuracy and completeness and, subsequently, transcribed.

Interview Coding and Analysis. Two coders conducted a five-phase thematic analysis of the transcripts, following best practices from Braun and Clarke [13]. The first phase, *Familiarization with the Data*, involved reading the transcripts to understand their content. In the second phase, *Generation of Initial Codes*, the data were systematically analyzed and split into meaningful units or segments, each representing a distinct concept or theme relevant to the research questions. The third phase, *Searching for Themes*, identified patterns and recurring concepts. In the fourth phase, *Reviewing Themes*, the coherence and relevance of the themes were evaluated. Finally, in *Defining and Naming Themes*, the themes were refined and reported. The coders regularly met to align their observations and unify codebooks. We used the ATLAS.ti software [6] to facilitate the process of interview coding and analysis.

The 5 verification interviews did not yield new themes, and only one code was added, without affecting the meaning of existing codes. This confirms that the data from the original 15 interviews were sufficiently saturated [4, 31, 32].

To ensure reliability, the final codebook was verified through an inter-coder reliability assessment using Krippendorff's Alpha, resulting in a score of 0.872, indicating substantial agreement. An overview of the codes and themes is provided in Appendix B.

3.2 Component 2: Report Analysis

To complement the findings from the interview study, we also analyze real ransomware reports quantitatively.

Reports Collection. We used two report sources in this study: *Company reports* from our partnering cybersecurity firm and *CISA reports*. To ensure up-to-date trends and aligned analyses, we only considered reports produced by the company and published by CISA in 2023.

Company reports are based on ransomware incident investigations and include a mandatory Root Cause Analysis (RCA), detailing the techniques and tools used by the RTA in

[1] https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Glossary:Enterprise_size

Table 2: RTAs and the number of company reports (#).

Threat Actor	#	Threat Actor	#
Blackcat	5	Monti	1
Play	4	RansomHouse	1
Lockbit	3	Carver Phobos	1
Black Basta	2	Ragnar	1
Royal	2	BlackSuit	1
Mallox	1	INC	1
HelloKittycat	1	C3RB3R	1
ESXiArgs	1	8Base	1

Table 3: Mapping between RTAs as named by the partnering company and CISA.

Company RTA	CISA RTA	Company RTA	CISA RTA
Blackcat	ALPHV	Hellokittycat	N/A
Lockbit	Lockbit	Mallox	N/A
Play	Play	ESXIArgs	N/A
Blackbasta	Black Basta	INC	N/A
Carver Phobos	Phobos	Monti	N/A
Blacksuit	Royal	Ransomhouse	N/A
8base	Phobos	Ragnar	N/A

the attack. Due to the sensitive nature of the reports, only one researcher was given access to manually extract relevant information. The ransomware strain and the tactics, techniques, and procedures used were recorded, ensuring all classified data was removed for team analysis. However, confidentiality restrictions prevented independent evaluation of the data extraction process. Twelve reports in our set lacked the details necessary for this study due to incomplete log files, where adversaries had removed logs, making it impossible to reconstruct the kill chain. These reports were omitted, leaving 27 documents available for analysis. In the analyzed reports, certain ransomware actors were more prevalent than others. However, the frequency of observed groups did not align with their overall contribution to attacks in 2023, except for LockBit, Play, and ALPHV Blackcat, which matched their reported activity [1]. Eleven ransomware actors appeared only once in the dataset, likely due to a combination of activity and regional factors. Table 2 lists the identified actors and their respective case counts.

The second source is *CISA reports* [21], published by the U.S. Cybersecurity and Infrastructure Security Agency (CISA) within the *#StopRansomware* campaign. CISA reports contain aggregated data from multiple incident documents provided by authoritative companies. Appendix C lists the ransomware actors, links to their reports, and publication dates. This study uses 13 CISA reports published in 2023.

RTA Mapping. The names of threat actors in company reports often differ from those in CISA reports due to rebranding, varying naming conventions, or the use of encryption tool names. Affiliates also add complexity to the mapping. We conducted an extensive search and created a mapping shown in Table 3, aligning the RTA names in the company and CISA reports.

Data Representation. To unify the data, we mapped information from the company and CISA reports to the MITRE ATT&CK Enterprise framework (Version 14.1) [18]. We used a tactic/technique/sub-technique combination for specificity. In cases where only a technique was mentioned, we conservatively also included all corresponding sub-techniques. Each report was then encoded as a binary vector of length 887, representing the presence or absence of a specific tactic, technique, and sub-technique combination[2].

Data Analysis. We analyzed the company reports using various data analysis techniques and metrics, focusing on the similarity of TTPs used by the same RTA across reports. To do this, we calculated similarity scores between each pair of reports and averaged the results. We opted for the *overlap* similarity metric (see Equation 1) over the Jaccard one, as the latter heavily penalizes sets with different cardinalities.

$$S_o(\mathbf{A},\mathbf{B}) = \frac{\sum_{i=1}^{n}\min(A_i,B_i)}{\min(\sum_{i=1}^{n}A_i,\sum_{i=1}^{n}B_i)} \quad (1) \qquad S_c(\mathbf{A},\mathbf{B}) = \frac{\sum_{i=1}^{n}\min(A_i,B_i)}{\sum_{i=1}^{n}A_i} \quad (2)$$

First, we assessed whether RTAs could be clustered based on TTPs from company reports and if these clusters correspond to the RTAs themselves. For this, we used the *silhouette score*, a metric that measures how well data points are clustered (ranging from -1 to 1). A higher score indicates well-defined clusters, meaning distinct TTPs for different RTAs, while a lower score suggests overlapping clusters, implying the opposite. We calculated this score using the `sklearn` library [81]. Then, we evaluated the distinctiveness of TTPs employed by different RTAs in CISA reports. To achieve this, we calculated both overlap (Equation 1) and containment (Equation 2) similarities. The key distinction between these metrics lies in their treatment of the compared sets: overlap similarity measures the shared TTPs relative to the smaller of the two sets, while containment similarity evaluates shared TTPs relative to one specific set. Finally, we compared the TTP similarities for the same threat actors across both the company and CISA reports.

3.3 Synthesis

We distill the key findings from our study based on the results of both interviews and report analysis. We do this by organizing the findings into three key topics in which the insights from the interviews and reports are discussed cohesively. Our key findings are presented in the three following sections: (Section 4) RTA attribution hinges on low-level indicators; (Section 5) RTA attribution is complex; and (Section 6) data

[2]There are 887 unique tactic/technique/sub-technique combinations.

fragmentation is a key challenge.

4 RTA Attribution Hinges on Low-Level Indicators

Our first key finding includes details on the RTA attribution process to support sanction screening and the fact that it hinges on low-level indicators. Interviewees give their opinions on the current attribution process and how they value various IoC types. The analysis of the reports corroborates the opinions of the interviewees, showing that using low-level indicators for ransomware attribution is currently the least bad alternative.

The interviewed experts consistently reported that low-level indicators, such as ransom notes, communication channels, leak sites, and network IoCs, are the primary means of attribution of attacks to a specific ransomware threat actor. The participants identified eight key indicators that are displayed in Table 4, of which seven are low-level. For example, interviewee E4 explained, "*We look at the ransom note and assume we are dealing with the group that claims responsibility*". These low-level indicators are considered reliable because they directly link the attack to the operational infrastructure of the RTA. A caveat that is mentioned for low-level IoCs is that such indicators often change and that this assumption only holds for recent attacks. In contrast, high-level indicators, such as TTPs, are regarded by experts as too generic, and thus, less useful for precise attribution. As interviewee E3 said, "*You cannot conclude that much when looking at just the TTPs, as they are too generic to really make a high-confidence distinction between groups.*" Instead, these indicators are valued for their theoretical potential to provide insights into attacker profiling, exemplified by interviewee E12: "*You can easily build a profile for ransomware groups based on the TTPs they use.*"

Table 4: Key indicators mentioned in the interviews and their description.

Indicator	Description
Leak site	A website, typically on the dark web, used to share publicly stolen data and conduct negotiations with victims.
Communication channel	A communication platform used to reach out to the ransomware adversaries.
Ransomware sample	Software that encrypts files or the operating system.
Ransomware note	A file/piece of information reporting that the system has been attacked with the data on how to get in contact with the ransomware threat actor.
TTPs	Tactics, Techniques, and Procedures employed by threat actors.
IP address	IP addresses that make up threat actor infrastructure.
Cryptowallet address	Cryptowallet related to, for example, Bitcoin payments.
Tools	Tools used by threat actors during attacks.

4.1 RTA Attribution Process

The practitioners described a structured workflow that demonstrates how low-level indicators are systematically utilized during attribution. This process includes the following steps:
1. Initial Identification: Attribution typically begins with analyzing the ransom note, which often includes the name of the group and a means of communication. This serves as the initial attribution hypothesis.
2. Verification of Identity: The communication channel referenced in the ransom note is analyzed to ensure that it aligns with the claimed group. Interviewee E1 emphasized this point, saying, "*The contact point is identifiable as only the group manages it.*" During the negotiation, attackers may prove involvement by decrypting files or using their leak site, further verifying their identity. This step acts as a secondary verification of the attackers' identity. However, to ensure sufficient confidence in the attribution for use in sanction screening, additional evidence is required, interviewee E1 continued: "*You must make a clear distinction between the quality of the attribution and the risk of a sanction violation after you have made that attribution.*" This is supported by interviewee E7, who explained that "*a single indicator is very weak, but a group of indicators pointing to the same conclusion is already much stronger when you have to be as certain as possible.*" Therefore, this necessitates a more in-depth attribution process.
3. Infrastructure Analysis: Analysts investigate network IoCs, such as IP addresses and domains, to map the infrastructure used in the attack. These indicators are compared to historical records of known infrastructure to confirm patterns or re-use. Low-level IoCs are deemed decisive evidence in this stage.
4. Malware and Tool Analysis: In some cases, ransomware samples are analyzed for unique signatures or errors in their code, but this is rarely definitive. As interviewee E9 noted, "*Ransomware code is not exceptionally complex, offering fewer unique signatures than conventional malware.*" The RaaS model and rebranding behavior further complicate this step, as multiple groups may use the same ransomware strain or core elements of its code. As interviewee E1 notes, "*It may also turn out that the group you have attributed is ultimately the same as another rebranded group.*" Groups rebrand after disruptions, adding complexity to the ransomware landscape. Participants observed that former group members often "brought TTPs" into newly formed groups, possibly making behavioral patterns useful to trace the origin of a new ransomware group but also contributing to the ambiguity of TTPs. Despite this, interviewees E2 and E7 mention that malware

Table 5: Average overlap similarity of the techniques in the reports attributed to the same RTA.

RTA	Overlap Similarity	
	Mean	Stddev
Blackcat	0.56	0.10
Lockbit	0.41	0.26
Play	0.28	0.23
Blackbasta	0.36	-
Royal	0.25	-
Average	**0.37**	**0.20**

analysis, including dynamic analysis and code similarities, helps authenticate the ransomware message and identify its type.

5. Cross-Referencing with External Data: Indicators are cross-checked against OSINT, SOC feeds, and threat intelligence reports from paid services or internal databases. Interviewees highlighted that this diversity of sources is crucial for increasing attribution confidence levels.

6. Attribution Confidence Levels: Finally, analysts assign a confidence level to their findings, typically categorized as low, medium, or high. Interviewee V2 pointed out that commercial pressures can distort confidence ratings, stating, "*Commercial companies often inadvertently pollute the threat intelligence ecosystem [by publishing inaccurate or inconsistent attributions] because they want to publish their research or think it discredits their research by giving it low or medium confidence[, and therefore mark it as high confidence even though they are far from certain].*"

7. Reporting and Sanction Screening: Attribution results are documented and checked against sanction lists to avoid payments to sanctioned entities.

4.2 Insights from Incident Reports

Our quantitative examination of incident reports supports experts' reliance on low-level indicators while confirming the challenges of using high-level indicators.

Low Overlap of TTPs for the Same RTA. Incident reports attributed to the *same* RTA shared a mean overlap similarity score of 0.37 (see Table 5), and this would be even lower if a more typical Jaccard similarity metric is used. This indicates considerable variability in the TTP sets *within* one group. Such variability stems from factors like incomplete forensic data or evolving attacker behaviors, aligning with the interviewees' remarks on the difficulty of pinning down a group based on high-level indicators alone.

Overlap Across RTAs. When aggregating all TTPs associated with each RTA and comparing different groups (Table 6a and Table 6b), the average overlap similarity was 0.21. Although lower than the 0.37 *within*-RTA overlap, it is still substantial enough to fade clear distinctions between separate RTAs. Negative silhouette scores (-0.0864 for Euclidean distance and -0.0873 for cosine distance) further highlight the lack of well-defined TTP "clusters", reinforcing the claim of interviewee E7 that "*it is very difficult to use TTPs to identify ransomware attackers because there is a lot of overlap.*"

However, despite the poorly defined clusters, some techniques might be unique enough for an RTA. A full list of unique TTPs found for the company and CISA reports can be found in our supplementary material [89]. There, a unique technique for an RTA means that any other RTA does not employ the same technique according to the same source of reports. However, for some groups, there are no unique techniques, and, moreover, the unique techniques mentioned in the company reports do not coincide with the ones provided by CISA.

Implications for Indicator Use. These quantitative results corroborate the interviews: TTP sets are neither stable nor sufficiently unique for unambiguous identification. By contrast, low-level IoCs like ransom notes and communication channels are more volatile but reliably pinpoint a specific group, reflecting the direct operational ties the interviews described. Thus, the quantitative analysis strongly supports the view that high-level indicators, though useful for long-term profiling and attack mitigation, fall short for precise, high-confidence attribution.

Key Finding: For RTA attribution, low-level IoCs are typically seen as more actionable and precise but volatile. High-level IoCs are valued for their theoretical potential in long-term RTA profiling, but are deemed too generic and inconsistent for precise attribution.

5 RTA Attribution is Complex

A second topic centers on the complexity introduced by rebranding, affiliate-based RaaS models, and nation-state false flags—all of which can lead to inaccuracies in attribution and serious legal or financial ramifications for victims.

5.1 Challenges Highlighted by Experts

The interviews highlighted numerous challenges in the attribution of ransomware, with the dynamic and fluid nature of RTAs being a recurring point of discussion. RTAs frequently rebrand to evade sanctions or regain operational capacity after disruption. As interviewee E1 explained, "*The moment a group is sanctioned, they rebrand, leading to a new group that is not sanctioned.*"

This rebranding strategy not only complicates tracking but also blurs the lines between older and newly formed groups, making attribution even more challenging. Interviewee E5 explained that "*some families have so much overlap with each other that you cannot always say based on TTPs who you are*

Table 6: Generated overlap similarity matrices based on the company and CISA reports.

(a) Company reports

RTA	Blackcat	Lockbit	Carver Phobos	Play	Blackbasta	Royal	Mallox	Ransomhouse	INC	Monti	Ragnar	Blacksuit	8base	Hellokittycat	C3RB3R	ESXIArgs
Blackcat	1.00	0.65	0.33	0.47	0.19	0.28	0.28	0.19	0.26	0.30	0.23	0.21	0.23	0.07	0.05	0.07
Lockbit		1.00	0.43	0.45	0.17	0.24	0.17	0.31	0.24	0.33	0.33	0.19	0.26	0.10	0.05	0.07
Carver Phobos			1.00	0.17	0.11	0.11	0.17	0.20	0.17	0.23	0.11	0.17	0.11	0.06	0.03	0.03
Play				1.00	0.32	0.46	0.11	0.29	0.36	0.36	0.36	0.32	0.36	0.07	0.07	0.11
Blackbasta					1.00	0.42	0.08	0.12	0.23	0.23	0.15	0.19	0.23	0.08	0.04	0.04
Royal						1.00	0.10	0.29	0.43	0.38	0.19	0.24	0.38	0.05	0.05	0.10
Mallox							1.00	0.26	0.11	0.11	0.11	0.11	0.05	0.11	0.05	0.05
Ransomhouse								1.00	0.32	0.26	0.32	0.26	0.37	0.16	0.11	0.11
INC									1.00	0.53	0.35	0.35	0.41	0.06	0.06	0.06
Monti										1.00	0.35	0.35	0.41	0.12	0.12	0.12
Ragner											1.00	0.25	0.44	0.12	0.12	0.12
Blacksuit												1.00	0.42	0.08	0.08	0.08
8base													1.00	0.09	0.09	0.18
Hellokittycat														1.00	0.29	0.29
C3RB3R															1.00	0.50
ESXIArgs																1.00

(b) CISA reports

RTA	Snatch	Lockbit	Cl0p	Phobos	Royal	Bian Lian	Rhysida	Akira	AvosLocker	Black Basta	DPRK	Play	ALPHV
Snatch	1.00	0.50	0.30	0.11	0.38	0.09	0.09	0.06	0.27	0.15	0.19	0.06	0.02
Lockbit		1.00	0.52	0.30	0.50	0.23	0.27	0.18	0.38	0.20	0.17	0.15	0.08
Cl0p			1.00	0.16	0.31	0.18	0.22	0.11	0.31	0.15	0.07	0.07	0.09
Phobos				1.00	0.24	0.39	0.27	0.41	0.12	0.22	0.04	0.27	0.12
Royal					1.00	0.26	0.33	0.22	0.09	0.26	0.22	0.20	0.13
Bian Lian						1.00	0.71	0.51	0.11	0.31	0.09	0.31	0.14
Rhysida							1.00	0.44	0.09	0.32	0.12	0.29	0.15
Akira								1.00	0.03	0.34	0.07	0.55	0.10
AvosLocker									1.00	0.07	0.04	0.07	0.07
Black Basta										1.00	0.13	0.35	0.22
DPRK											1.00	0.09	0.04
Play												1.00	0.06
ALPHV													1.00

dealing with", which reinforces E1's claims in Section 4.1 about how attributed groups could essentially be the same as another rebranded group.

As discussed in Section 4.1, the RaaS ecosystem adds complexity. Interviewee E6 underscored this challenge, stating, "*The probability of a forensic analyst encountering the same affiliate and corresponding modus operandi in multiple separate incidents is very low.*" Affiliates' roles are often indistinct, further complicating efforts to separate their activities from those of core RTA members.

A notable challenge arises from the forensic evidence itself. Interviewee E9 emphasized that even when forensic analysts find ransomware samples or IoCs, the shared nature of RaaS tools often makes it difficult to identify the exact threat actor. They noted, "*Even if you identify the ransomware strain, it doesn't necessarily mean you've identified the actors behind it, as multiple affiliates may use the same strain.*"

Misattribution was frequently mentioned as a critical concern. Twelve out of twenty interviewees (E2, E3, E4, E5, E6, E8, E9, E12, E13, V1, V2, V4) noted that incorrect attribution could have severe legal and financial consequences, especially when sanctions compliance is involved. Public misattribution can also harm the attributing party's reputation, as inaccurate information is deemed to pollute the threat intelligence landscape.

There are also conflicting incentives of the threat actors themselves to be identified correctly. On the one hand, interviewee E14 notes that "*ransomware threat actors portray themselves as a company offering penetration testing and decryption services and have a reputation to uphold*". Thus, RTAs have no incentive for false flag operations, which would misattribute their strain. On the other hand, nation-state actors complicate attribution, as they sometimes mimic ransomware groups to obscure their motives or use destruction to obscure their objectives. Interviewee E12 remarked, "*APTs are significantly harder to profile as they have a larger variety of TTPs and sometimes mimic ransomware groups to hide their goals.*" These false-flag operations blur the distinction between financially motivated RTAs and state-sponsored groups, complicating attribution further, especially in cases where attackers deliberately use ransomware to obscure geopolitical objectives. Examples of such cases are APT41 (China) and APT45 (DPRK), which are suspected of serving state interests while simultaneously deploying ransomware on victim infrastructure [28, 45]. While the exact motivations of their attacks, and therefore whether they are false-flag operations, are unconfirmed, it shows that state-linked groups should be expected to employ ransomware to facilitate their goals.

5.1.1 Ethical and Legal Challenges

Ethical challenges were sometimes mentioned as an obstacle to attribution. For example, some participants discussed difficulties in balancing their obligations to clients, victims, and the broader threat intelligence community. For instance, interviewee E9 mentioned that "*threat actor attribution on a more fine-grained level than on group level is not the main priority and therefore cost-ineffective.*"

According to our participants, ethical and legal considerations often limit the degree to which analysts can share sensitive evidence or findings or act on them, which in turn affects the completeness of attribution reports. For example, as interviewee E5 stated, "*Understanding the infrastructure of an actor and potentially hacking or logging into their infrastructure [, aside from the legal concerns,] poses ethical questions. There may be data from the investigation that we cannot use without the client's permission.*" Some experts also mentioned

that they discovered credentials allowing them to access the attacker's systems, raising the dilemma of whether to proceed with attribution, risking evidence compromise, or hand it over to law enforcement, potentially delaying the work.

Interviewee E2 underscored the significant impact of legal and geopolitical barriers, such as restrictive extradition rules, which prevent the detention of certain attackers. They mentioned that "*the effect of attribution is very limited, as ransomware actors are often in countries without an extradition treaty with Western countries, which severely delays or even completely prevents prosecution*". As a result, the adversaries exploiting these legal loopholes often evade accountability. Moreover, it was also mentioned that some group activities, such as Evil Corp (assigned UNC2165 by Mandiant, [46]), are even welcomed at the state level.

The interviewees with a background in law enforcement stated that, as they work on more fine-grained levels of attribution, they must take into account the Code of Criminal Procedure to ensure legality. Hence, law enforcement procedures impose constraints on attribution in the service of criminal prosecution. From the legal perspective, participants also mentioned the lack of clarity regarding the attribution process for the sanctions check purposes, for example, to what extent the investigation should proceed and what confidence level is mandatory.

5.2 Incident Report Data

Our report analysis confirms some of the concerns mentioned by the interviewees.

Negative Silhouette Scores & Fluid Clusters. The negative silhouette scores indicate that TTP-based clusters are poorly defined, supporting the observation that RTAs, or their affiliates, often employ overlapping and evolving sets of techniques. This aligns with interviewee E5's assertion that certain ransomware strains exhibit so much overlap that accurate TTP-based attribution is often impossible.

Building on this, E6's observation about inconsistent modi operandi highlights how phenomena such as rebranding and affiliate turnover contribute to significant ambiguity in the use of TTPs among many ransomware groups. Rebranding and the RaaS model further exacerbate this issue: while the likelihood of encountering the same TTPs increases, the probability of encountering the same actor or affiliate decreases. This contradicts E12's earlier vision on the long-term profiling of ransomware groups based on the TTPs they use and adds significant complexity to the attribution process.

Low Intra-RTA Similarity. The 0.37 average TTP overlap within the same group (see Table 5) highlights how rebranding or affiliate diversification fragments the group's "signature". This variability supports interviewee E9's comments about the shared nature of RaaS tooling, where multiple affiliates use the same ransomware strain, creating ambiguity in attributing attacks to a specific group or affiliate. The data underscore that, if TTPs alone remain so inconsistent, it is easier for analysts to misattribute – precisely as interviewees warned.

Mimicry Is Not Difficult. While available reports in our dataset do not explicitly label state-sponsored adversaries while masquerading as RTAs, the fluid and shared nature of TTPs still aligns with the interview statements that nation-state groups can mimic ransomware. The wide technique overlap (averaging 0.21 among distinct RTAs) suggests it would be relatively straightforward for advanced actors to adopt existing TTPs, camouflaging themselves as regular ransomware.

Overall, these data confirm the interviewees' emphasis on caution: the TTPs overlap and limited intra-group separation lead to the likelihood of misattribution, especially if victims or investigators focus only on high-level indicators.

Key Finding: Frequent rebranding, affiliate-based RaaS models, ethical considerations, and possible nation-state false-flag operations make it hard to pinpoint which actors are behind ransomware attacks, sometimes leading to misattribution with serious legal and financial risks for victims.

6 Data Fragmentation is Key Concern

Finally, participants stressed that organizational hesitancy to share incident data, coupled with ambiguous sanction enforcement legislation, hinders effective attribution. This topic also ties into the need for centralized databases and better cross-organizational and cross-national collaboration.

6.1 Experts Want to Reduce Data Gaps

The interviews underscored significant data gaps and the pressing need for improved collaboration among entities involved in ransomware attribution. A *centralized database* including both low- and high-level indicators was universally acknowledged as a critical improvement among participants. Such a resource could address the fragmentation of evidence and enable analysts to identify patterns more effectively. As interviewees E5 and E15 noted, historical data play a crucial role in attribution, helping analysts identify rebranded RTAs or distinguish affiliates from core members.

We already mentioned in Section 5.1.1 that ethical and legal considerations complicate the attribution process and sometimes limit the degree to which analysts can share their findings. Participants also noted the hesitancy of organizations to disclose incident data due to concerns about reputational harm (E2, E6, V3). This reluctance creates information silos, limiting the ability of the larger community to track RTA behaviors and improve the accuracy of the attribution. The lack of shared data prevents the development of a more comprehensive understanding of ransomware operations.

Finally, the interviewees identified the lack of collaboration between cybersecurity firms and law enforcement agencies as another obstacle. Interviewee E2 emphasized the potential value of stronger partnerships, stating, "*Law enforcement often has a better overview of what is going on in a country, whereas cybersecurity companies could deliver pieces of evidence from attacks that can be used for prosecution.*" However, these collaborations are often hindered by differing priorities, extradition limitations as discussed in Section 5.1.1, and barriers to data sharing.

6.2 Coverage Differences in the Reports

Our comparative analysis of company and CISA reports reinforces the interviewees' concerns about data gaps and framework limitations. Specifically, the discrepancy in the number of TTPs attributed to each RTA and the partial overlap between the two data sources provides quantitative evidence of fragmented reporting.

Coverage Discrepancies. As shown in Table 7, CISA often registers more techniques than our partnering company for a specific RTA (such as Royal and Carver Phobos) by pooling information from multiple intelligence feeds. In other cases, the company's internal investigations list significantly more TTPs than CISA publications for the same RTA (such as Blackcat). In both scenarios, the *overlap similarity* for the same RTA across the two sources averages around 0.35. That is – each data source independently observes or documents different techniques – implying that no single dataset offers a complete picture of the RTA's behavior. These coverage discrepancies underline how partial data can lead to potential misalignment in threat intelligence. Similar issues with coverage were previously reported for general threat intelligence feeds [12,30,43], but our study confirms this for RTAs-related intelligence reports.

Fragmentation and Potential Under-Reporting. Even for RTAs with multiple incident reports available, we identified large gaps in the sets of documented techniques. Such gaps can arise from incomplete logging, limited forensic access, or legal and reputational concerns that discourage sharing full findings, which is in line with the comments of interviewees E2 and E6, who stressed that organizations might not want to provide full transparency. Consequently, both the overlap similarity (0.35) and the containment similarity (see Table 7) suggest that neither dataset likely captures the entire spectrum of a given RTA's techniques. This fragmentation directly validates the interviewees' experiences of scattered or partial forensic data, highlighting why analysts may miss critical signs of rebranding, different affiliates, or new TTPs.

Key Finding: Neither private companies nor public agencies appear to have a fully comprehensive repository of RTA's IoCs, prompting our participants to call for clearer guidelines and robust, centralized databases. Blind spots in different sources hinder the comprehensive tracking of RTAs and could lead to a victim organization inadvertently paying a sanctioned threat actor and facing legal consequences.

7 Discussion

We now discuss the findings from our mixed-methods empirical study and go into further detail on what they mean in practice.

7.1 Implications of the Key Findings

This section summarizes the study's three primary insights, highlighting their implications for both research and practice. Although low-level indicators offer actionable intelligence for immediate threat response, they tend to be volatile and are quickly rendered obsolete once adversaries modify their infrastructure. High-level indicators, such as adversarial tactics or broader behavioral signatures, promise more longevity and can theoretically aid in strategic profiling; however, interviewees noted that these remain too generic to be consistently used for precise attribution.

This nuance becomes more important in the second key finding of this work, which draws attention to the complexities of attributing ransomware attacks to distinct RTAs. Participants emphasized that the indicator shortcomings in combination with frequent rebranding, affiliate-based RaaS models, ethical considerations, and occasional nation-state false flags not only hinder accurate attribution but also increase the risk for serious legal and financial consequences. Errors in attribution can occur when investigators rely on overlapping TTPs or incomplete intelligence, sometimes leading victims to engage – knowingly or unknowingly – with sanctioned entities.

Finally, both private entities and public agencies seem to lack a comprehensive indicator repository, which is a deficiency that interviewees believe contributes significantly to the risk of misattribution. The resulting blind spots can leave victim organizations at risk of inadvertent sanction violations. Because of these blind spots, there is broad support for more robust, centralized databases alongside clearer guidelines for sanction compliance. Addressing these gaps through standardized data-sharing frameworks and multi-stakeholder collaboration could reduce the likelihood of misattribution and strengthen the response to ransomware.

7.2 Mapping Indicators to the Pyramid of Pain

The dynamic ransomware ecosystem, particularly the RaaS subtype, complicates accurate attribution as affiliates use dis-

Table 7: Comparison of the same RTAs between the company and CISA reports using containment and overlap similarity.

RTA		# of Techniques		Containment Similarity		Overlap Similarity
Company	CISA	Company	CISA	Company -> CISA	CISA -> Company	
Blackcat	**ALPHV**	43	10	0.05	0.20	0.20
Lockbit	**Lockbit**	42	60	0.21	0.15	0.21
Play	**Play**	28	18	0.36	0.56	0.56
Blackbasta	**Black Basta**	26	23	0.27	0.30	0.30
Royal	**Royal**	21	46	0.43	0.20	0.43
Carver Phobos	**Phobos**	35	49	0.23	0.16	0.23
Blacksuit	**Royal**	12	46	0.42	0.11	0.42
8base	**Phobos**	11	49	0.45	0.10	0.45
Mean		**27.25**	**37.62**	**0.30**	**0.22**	**0.35**

Table 8: Suggested positioning of different ransomware indicators in the Pyramid of Pain (PoP.)

Pyramid Layer	PoP Types	Ransomware IoCs
Tough!	TTPs	Ransomware Brand
Challenging	Tools	TTPs, leak site, communication channel, ransomware sample, used tools
Annoying	Network/Host Artifacts	Ransomware note linguistics
Simple	Domain Names	Leak site URL, communication channel URL
Easy	IP Addresses	C2 address, other IP addresses, Cryptowallet address
Trivial	File Hashes	Malware hashes, tool hashes

tinct infrastructure, TTPs, and strains, increasing the volatility of low-level indicators and the ambiguity of high-level ones. The Pyramid of Pain (PoP) model [10], capturing how pain is inflicted when an indicator is denied, does not fully account for the indicators mentioned in the interviews. Though less used, these indicators are taken into account by models like the Q-model [72]. Given the PoP's central role in both intelligence and defense, it is worth reconsidering how ransomware fits into this model. The mapping below reflects how difficult it is for RTAs to change indicators and their value for defense.

Refining the Pyramid of Pain for ransomware attacks introduces the following considerations:

- **Cryptowallet addresses:** Cryptowallet addresses, linked to financial activity and infrastructure tracking, could be considered a separate indicator in the model. Being easy to change, it may fit the *Easy* level.
- **Ransomware note:** Although ransomware notes are simple to rewrite, interviews noted that aspects like syntax, grammar, and cultural references are harder for adversaries to alter. The note itself fits the *Easy* level, but ransomware linguistics could be classified as *Annoying*, unless tools like Large Language Models become widely used to rewrite notes.
- **Ransomware brand:** While unconventional as an IoC, ransomware brands provide strategic value similar to TTPs, affecting reputation and operational impact. Brands may fit the *Tough!* layer, as rebranding usually disrupts operations. Although not a technical indicator, they hold strategic significance.
- **TTPs of RaaS affiliates:** To bring attention to the discovered issues with adversarial TTPs in the RaaS ecosystem, we explicitly downgrade TTPs from *Tough!* to the *Challenging* layer in the Pyramid. The affiliates in the RaaS ecosystem face lower costs to change them, relying on Initial Access Brokers and RaaS developers. This challenges the Pyramid of Pain's assumptions and suggests that targeting TTPs in RaaS is more about disrupting the ecosystem's supply chain than forcing a single entity to change.

In light of these considerations, Table 8 shows the positioning of different ransomware-related IoCs in the Pyramid of Pain. While the Pyramid of Pain still holds its value for defense against ransomware attacks, practitioners need to adjust their strategies accordingly to account for these dynamics.

7.3 Improving the Status Quo

Taking into account all the different results and findings, there are a few key problems that require a solution. These are: an improved way of (central) knowledge sharing and management, raising awareness on the pollution of the cyber threat intelligence (CTI) landscape, and refining the current policy surrounding sanction violations.

Knowledge sharing and management. One common theme among interviewees was the centralized sharing of knowledge, as attribution data is currently fragmented, adding complexity to the process. Systematizing knowledge management could provide a comprehensive, historical view, despite its vulnerabilities. This is also acknowledged by Badva et al. in [7], who also conclude from qualitative research with security analysts that centralization of knowledge will help. An example is the suggested Cyber Solidarity Act (CSA), adopted in December 2024 for EU Member States [26], which promotes cross-border Security Operation Centers and mutual

assistance in threat detection and mitigation. While the private sector is involved through the Solidarity Mechanism, the act mainly focuses on (inter)national threat detection and mitigation. Organizations that systematize intelligence sharing in a similar way already exist and are called Information Sharing and Analysis Centers (ISACs). As described in [23], ISACs play a crucial role in facilitating effective information exchange among organizations. However, ISACs are sectoral, so it would be beneficial to expand on initiatives that aim to function as cross-sectoral ISACs or to further incentivize and stimulate inter-ISAC coordination. Future research should prove to what extent such centralization would improve the situation and what risks come with this approach.

Raising awareness on threat intelligence pollution. A growing issue with centralized knowledge sharing is the risk of threat intelligence pollution. Maschmeyer et al. [47] highlighted the bias in the CTI industry, where companies focus on high-profile attackers and unique TTPs. Interviewees noted that commercial cybersecurity companies often publish reports to gain attention, even without the ability to attribute attacks accurately, expressing high confidence in specific indicators. This reflects a misunderstanding of the threat intelligence landscape, as few entities, among which law enforcement, can attribute with such high confidence. Interviewees emphasized that attributing with low or medium confidence demonstrates competence, despite companies' fears it would appear otherwise.

The problem with this behavior is that consumers of this intelligence will not have the right information they need to effectively mitigate a threat. Consequently, this worsens the already poor observability of RTAs.

Refinement of sanction violation policy. The practical value of both high- and low-level indicators presents challenges to sanctions against cyber threat actors. The found implications of their use in attribution raises concerns about the proportionality of current sanction policies. Institutions that pay ransoms cannot verify if the recipient is sanctioned, and when rebranded ransomware groups are linked to sanctioned entities, victims may become liable for indirect transactions [88]. Interviewees noted ambiguities in sanction policy and screening procedures, which is supported by legal researchers [86]. This ambiguity, combined with potential penalties, fosters Fear, Uncertainty, and Doubt (FUD) among victims, which may benefit ransomware actors, as is also suggested by Connolly and Borrion, who found that reward systems could be more effective than punishment [101].

The US Office for Foreign Assets emphasizes that both paying ransom and facilitating payments violate sanctions [88]. As a result, insurance companies hesitate to reimburse or assist with ransom payments [80]. One interviewee noted that insurers may reclaim ransom payments after discovering that a rebranded group is tied to a sanctioned entity.

Everything considered, current laws on sanction violations may disproportionately affect ransomware victims. The risk of misattribution raises ethical and legal concerns, especially given the severe consequences for victims. Ransom payments should either be banned or regulations refined to better protect victims, possibly through provisions safeguarding organizations acting on the best available evidence. There should also be procedures to review sanctions as new information emerges, ensuring fairness for victims and insurance companies.

8 Limitations

This study has some limitations that need to be acknowledged.

The majority of the interviewed experts (15 out of 20) originated from a single company specializing in ransomware incident response and negotiations. To corroborate these findings, we conducted five additional interviews with practitioners from five other relevant organizations, which demonstrated data saturation. All the organizations involved operate internationally (within the EU or even globally, including in the Global South), so while their headquarters may be Western-based, their practical experience extends beyond a single national landscape. However, we acknowledge a potential Western-centric bias tied to specific regulatory frameworks, work cultures, and threat landscapes that could stem from our participants and the documents we analyzed. This is particularly relevant in the context of global geopolitics and differing ransomware trends, especially in parts of the Global South that may face distinct challenges. Future research that includes participants from a broader array of geopolitical contexts – especially those outside of the EU/US – would further validate and expand upon these findings, offering deeper insights into how ransomware threats and sanction screening processes might vary worldwide.

The study also has certain limitations related to the report analysis. First, the company reports analyzed in this study were written by a single team. This raises issues about how representative this dataset is. While the partnering company is reputable, we cannot confirm the universal representativeness of the studied dataset due to inconsistencies with other reputable sources like CISA's #StopRansomware campaign [21]. This means that regional and cultural biases may remain in this study.

Second, the incident reports of the partnering company are considered highly sensitive. Therefore, only one researcher had access to these reports, and we were not able to measure an agreement metric for the data extracted from the reports. Thus, the study results might be skewed due to possible mistakes or subjectivity when extracting data (researcher bias).

Third, the incident reports of the partnering company gave limited insights into the procedures and tools used by the threat actors. There was not enough data to draw any conclusions from this; hence, they have not been included in this

study. It is still possible for these indicators to display patterns that can be used for attribution. Moreover, it is worthwhile to investigate the effect of combining low-level and high-level indicators on attribution accuracy.

Fourth, we employed similarity metrics to compare TTPs in ransomware reports, treating all TTPs as equally valuable. This is an approach commonly used for TTP comparisons in large datasets [15, 39], but this representation may be too coarse-grained, if, for example, the analysts assign varying weights to different TTPs or consider some combinations of TTPs to be more revealing of particular CTAs. However, we believe that our approach is appropriate, because in our study we did not find any evidence that TTPs are used in more than a binary fashion. For example, industry reports, such as the CISA reports we employed, provide only lists of TTPs common to the analyzed threat actors, and do not discuss groups of TTPs or more important TTPs. When discussing the RTA attribution process, our practitioners also did not mention any TTP groupings or some important techniques pertinent to some RTAs. Moreover, our approach is justified by the observation that inconsistent reporting undermines the ability of TTP sets to enable unambiguous identification. Consequently, assigning different weights to TTPs would not resolve this lack of differentiation. Still, we welcome future studies zooming into the TTP usage and establishing better ways to report TTPs for CTAs.

Finally, a common practice among ransomware actors and affiliates is the use of Initial Access Brokers, who break into a company and sell that access on underground forums [68]. Our research design does not account for this practice, taking along the Initial Access (sub-)techniques. As we have no way of determining whether or not an Initial Access Broker was used for the break-in, we decided to include these results. We expect that, in the worst case, this could lead to some generalization of the used Initial Access methods among ransomware actors.

9 Related Work

Ransomware. There is a rich literature on ransomware-related research, with several recent survey studies summarizing this area [8, 50, 51, 59, 67].

Research on sanctions in ransomware incidents predominantly focuses on the political and economic effects of imposing sanctions on RTAs and those who facilitate ransom payments. To the best of our knowledge, none of the studies focused on enabling victims to verify whether their attacker has been sanctioned by the jurisdiction the victim falls under. Abely [2] writes about what seems to be the general perspective: ransomware should be battled by prohibiting ransom payments, as the system for imposing sanctions on an RTA suffers from timing issues, as payments to criminals are (often) not prohibited until they are sanctioned by the proper authorities. However, some studies argue against the blanket ransom payment ban due to challenges in enforcement and the drastic implications for critical infrastructures [11, 60, 63].

While the ransomware type itself may be easy to identify, the RaaS model introduces a form of ambiguity with affiliates bringing in more noise. Sanction screening is currently based on lower-level indicators such as cryptocurrency addresses [33]. However, if these indicators span attacks from multiple ransomware types, as our research and Cable et al. [14] show independently, the current form of RTA attribution for sanction screening is ineffective.

Bendiek and Schulze [9] discuss that the EU attribution policy is incoherent due to the asymmetry in technical and intelligence capabilities. The study mentions that attribution takes a long time and advocates for better international cooperation and data sharing for overcoming the jurisdictional and operational challenges that currently hinder effective sanction implementation. Our participants also strongly voice the need for better information sharing as the current fragmented data landscape hinders precise RTA attribution.

While the attribution problem in cyber threat intelligence is still an unsolved problem [76], research on ransomware detection largely focuses on the (automatic) classification of ransomware samples with machine learning algorithms [8, 57, 67, 90], often focusing on code similarity analysis of either the malware itself or, when available, the source code [24, 61, 90]. As mentioned by our interviewees, such sample analysis is helpful and is part of the attribution process they follow, but alone it is not sufficient for a confident RTA attribution due to the frequent code reuse among groups and the relatively small size of ransomware code.

Many studies focus on tracking ransomware payments and analyzing RTAs via cryptocurrency transactions recorded in blockchains [14, 17, 29, 35, 66, 69, 92]. Notably, Cable et al. [14] demonstrate that RTAs are interconnected (due to rebranding and evolution of the groups) also via the blockchain transactions. For example, the Conti group had a substantial overlap in terms of the transferred funds over a set of addresses with several other RTAs, e.g., Royal and Black Basta, revealing the interconnections among the remnants of this group. Thus, cryptocurrency transactions as an indicator also show that a precise attribution of RTAs is challenging, and there is an overlap among indicators of different RTAs.

Additionally, some works investigate which victim and RTA characteristics influence the attack incidence and the ransom paid [55, 101, 102], while others look into analyzing ransomware communications, investigating how different RTAs function [29, 74], how dark web fora users discuss ransomware [53, 93], or which interventions are most effective against RTAs [54].

CTA attribution. Several studies specifically focus on extracting, investigating, and comparing TTPs in the context of ransomware or cyberattacks, and there are many proposed

systems and models that rely on TTPs and tools for CTA attribution [24, 34, 36, 61, 75, 82, 94, 96, 97, 100]. TTPs are widely mentioned as the more robust indicators that aid attribution [5, 12, 38, 82, 91, 96], a dynamic that is also included in various widely-accepted models like the Pyramid of Pain and Q-model [16, 40, 41, 73, 82, 94]. Our results indicate that for RTA attribution, TTPs might be limited, as they are not unique per RTA, reporting is inadequate, and our practitioners reported that they rely on lower-level indicators for a confident attribution. Closest to our study, Song et al. [83] evaluated the similarity of 12 ransomware attacks (from Conti, Lockbit, and Hive) expressed as ATT&CK tactics and techniques and specific keywords extracted from threat intelligence reports. They observed results similar to our findings from the report analysis: while some tactics are shared across different RTAs, not all tactics are equally present for any attack of a given RTA. Some overlap of the TTPs from different CTAs was also reported in [61, 71, 78]. Compared to [61, 71, 78, 83], we study this phenomenon more systematically, for a larger number of incident reports and in the context of RTA attribution. At the same time, Kim et al. relatively accurately attribute 12 CTAs in the mobile domain using vectorized TTP representations (similar to our data representation) [39]. This demonstrates, again, that RTA attribution can be inherently different from attributing CTAs, as there seems to be more reuse of TTPs. Chen et al. [15] propose a clustering approach for APTs integrating ATT&CK techniques, tools, targeted industry, and location, showing that coupling the techniques with the targeted industry has the best clustering performance.

Regarding lower-level indicators, Leite et al. demonstrate promising results for clustering attacks, including RaaS, based on DNS patterns [42], while Xiao et al. show that an integrated representation of high- and low-level indicators can achieve good performance for CTA attribution [98]. These methods highlight the usefulness of lower-level indicators in identifying CTAs, but the study [42] also emphasizes that these promising patterns have a shorter lifespan.

Qualitative research into malware analysis and threat hunting. Yong Wong et al. conducted an interview study with malware analysts [99] who reported that they extract both lower-level IoCs and the higher-level TTPs from the studied samples. However, this study did not investigate the attribution process from the obtained indicators, as this was not in the scope. Badva et al. [7] and Maxam and Davis [48] qualitatively studied the work of threat hunters. Similarly to our findings, both these studies highlighted the challenges related to data and the need for collaboration in the threat-hunting process.

10 Conclusion

This study explored the use and value of high-level indicators in ransomware attribution, particularly their role in sanction screening. Through an empirical analysis of incident reports and expert interviews, we found significant limitations in both low- and high-level IoCs for precise attribution. While low-level indicators offer direct operational links, their volatility undermines long-term reliability. Conversely, high-level indicators, though theoretically robust, proved too generic and inconsistent, often failing to uniquely identify Ransomware Threat Actors (RTAs).

These challenges are compounded by the complexities of the RaaS ecosystem and the frequent rebranding of ransomware groups. The fragmentation of this threat hinders effective attribution, increasing the risk of misattribution with legal exposure for victims during sanction enforcement as a result. Moreover, our analysis uncovers the lack of a comprehensive indicator repository among both private and public entities, resulting in blind spots that can leave victim organizations at risk of inadvertent sanction violations. These findings not only challenge the conventional reliance on high-level IoCs, but they also expose the limitations of current sanction enforcement policies in the context of ransomware. Ultimately, the study highlights that reliable ransomware attribution cannot rely on any single class of IoCs alone, necessitating a multilayered approach and informed policy to mitigate both technical and legal risks.

To address these gaps, we propose actionable steps and advocate for adjustments to address ransomware-specific nuances. By enhanced intelligence-sharing practices, the reevaluation of widely used frameworks like the Pyramid of Pain, and clearer sanction compliance guidelines, the available means for ransomware attribution can be improved. In addition, policymakers must refine sanction enforcement policies to better protect victims acting in good faith, ensuring proportionate enforcement while maintaining deterrence.

This work underscores the urgent need for interdisciplinary collaboration to develop resilient attribution methodologies, balancing technical, legal, and ethical dimensions. By advancing these efforts, we aim to provide stakeholders with the tools to navigate ransomware threats more effectively while mitigating unintended consequences of current attribution practices.

Acknowledgments

We extend our gratitude to all the experts who participated in this study. We thank the anonymous reviewers and our shepherd for their valuable and insightful feedback.

This research has been partially supported by the Dutch Research Council (NWO) under the project “THESEUS” (NWA.1215.18.006).

Ethical Considerations

For this work, we collaborated closely with the TU Delft research ethics department to craft the research protocol for our interview study component (see Appendix A). The planning and implementation of the interviews rigorously followed the applicable ethical guidelines governing research with human subjects. We also considered possible unintended consequences for our participants and the partner company that could arise from publishing this work and tried to minimize potential harm to them. All participants were informed about the goals and the process of the study, and gave their consent to participate. They also had the opportunity to review a complete draft of the paper, and, particularly, check the quotes that would be used. Multiple representatives of the partner company participated in drafting this paper and/or reviewing the content. Ethical approval for this work was obtained from the relevant TU Delft Institutional Review Board, ensuring compliance with the established standards.

Our research topic touches on an important ethical question about the morality of payments to cybercriminals. We do not advocate that ransomware payments should be allowed or facilitated. While paying ransom perpetuates the threat and is therefore considered immoral by many, the reality of many organizations is that they feel forced to do that – whether by multiple extortion schemes or because their business continuity is so severely disrupted that recovery is impossible or will be very challenging [101]. Thus, our study aims to understand and explicate the attribution process for sanction checks so that it can be improved. For example, our findings indicate that the current legal framework of attribution for sanction checks is not clear to the practitioners. Moreover, the findings indicate that imposing sanctions will only trigger RTAs to rebrand, further blurring the line for victims to determine sanction compliance. This means that enforcing sanctions punishes victims more than it does adversaries. Hence, with this work, we aim to push for a revision of sanction violation policies. Policymakers can work on developing clearer legal guidelines that will support victim organizations, insurers, and security service companies in their decision-making regarding ransomware payments.

Open Science Policy Compliance

What we share. To make our study as reproducible as possible, we share data and additional details in the supplementary material [89]. We share the data representing the studied RTA groups (the representation described in Section 3.2). For the RTAs extracted from reports provided by the company, we anonymize the adversary names, as requested by our partner company. We share the code to analyze these representations and verify our results. We also share our codebook (Appendix B).

What we do not share. Our partner company has to protect their reputation and know-how. This is why, from the start, we agreed with them that the full reports and the interview transcripts (even anonymized) would not be shared publicly. We used the same consent form for all interviewees, which explicitly mentioned our data policy (that we will not share the transcripts outside the research team). Thus, we also cannot share the transcripts from the interviews done for verification purposes.

References

[1] Ransomware recap 2023. https://e.cyberint.com/hubfs/Cyberint_Research_Report_2023_Ransomware_Recap.pdf, January 2023.

[2] Christine Abely. Ransomware, cyber sanctions, and the problem of timing. *BCL Rev. E. Supp. I-*, 63:47, 2022.

[3] Bader Al-Sada, Alireza Sadighian, and Gabriele Oligeri. MITRE ATT&CK: State of the art and way forward. *ACM Comput. Surv.*, 57(1), October 2024.

[4] Khaldoun Aldiabat and Carole-Lynne Le Navenec. Data saturation: The mysterious step in grounded theory method. *The Qualitative Report*, 23(1):245–261, January 2018.

[5] Mohammed Asiri, Neetesh Saxena, Rigel Gjomemo, and Pete Burnap. Understanding indicators of compromise against cyber-attacks in industrial control systems: A security perspective. *ACM Transactions on Cyber-Physical Systems*, 7(2), April 2023.

[6] Atlas.Ti. ATLAS.ti | The #1 Software for Qualitative Data Analysis. https://atlasti.com.

[7] Priyanka Badva, Kopo M Ramokapane, Eleonora Pantano, and Awais Rashid. Unveiling the Hunter-Gatherers: Exploring threat hunting practices and challenges in cyber defense. In *USENIX Security Symposium*, pages 3313–3330, 2024.

[8] Craig Beaman, Ashley Barkworth, Toluwalope David Akande, Saqib Hakak, and Muhammad Khurram Khan. Ransomware: Recent advances, analysis, challenges and future research directions. *Computers & Security*, 111:102490, 2021.

[9] Annegret Bendiek and Matthias Schulze. Attribution: a major challenge for EU cyber sanctions. *SWP Research Paper*, 2021.

[10] David Bianco. The pyramid of pain. https://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html, 2013.

[11] Jenny Blessing, Jules Drean, Sarah Radway, P Whartenby, and K McDermott. Survey and analysis of US policies to address ransomware. *MIT Sci. Policy Rev*, 3:38–46, 2022.

[12] Xander Bouwman, Harm Griffioen, Jelle Egbers, Christian Doerr, Bram Klievink, and Michel van Eeten. A different cup of TI? The added value of commercial threat intelligence. In *USENIX Security Symposium*, pages 433–450, August 2020.

[13] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.

[14] Jack Cable, Ian W Gray, and Damon McCoy. Showing the receipts: Understanding the modern ransomware ecosystem. *arXiv preprint arXiv:2408.15420*, 2024.

[15] Zheng-Shao Chen, R. Vaitheeshwari, Eric Hsiao-Kuang Wu, Ying-Dar Lin, Ren-Hung Hwang, Po-Ching Lin, Yuan-Cheng Lai, and Asad Ali. Clustering APT groups through cyber threat intelligence by weighted similarity measurement. *IEEE Access*, 12:141851–141865, 2024.

[16] Robert Andrew Chetwyn, Martin Eian, and Audun Jøsang. Modelling indicators of behaviour for cyber threat hunting via sysmon. In *European Interdisciplinary Cybersecurity Conference*, pages 95–104, 2024.

[17] Mauro Conti, Ankit Gangwal, and Sushmita Ruj. On the economic significance of ransomware campaigns: A bitcoin transactions perspective. *Computers & Security*, 79:162–189, 2018.

[18] The MITRE Corporation. MITRE ATT&CK v.14. https://attack.mitre.org/, 2023.

[19] CyberClan. Ransomware advisories: Is your cybersecurity firm ready? https://cyberclan.com/us/knowledge/ransomware-advisories-is-your-cybersecurity-firm-ready/, November 2020.

[20] Cybersecurity and Infrastructure Security Agency. Understanding ransomware threat actors: LockBit. https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-165a, June 2023. Cybersecurity Advisory.

[21] Cybersecurity and Infrastructure Security Agency. #StopRansomware. https://www.cisa.gov/stopransomware, 2024.

[22] Cybersecurity and Infrastructure Security Agency. #stopransomware: Alphv blackcat. https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-353a, February 2024. Cybersecurity Advisory.

[23] Josiah Dykstra, Lawrence A Gordon, Martin P Loeb, and Lei Zhou. Maximizing the benefits from sharing cyber threat intelligence by government agencies and departments. *Journal of Cybersecurity*, 9(1):tyad003, 04 2023.

[24] Kelsie Edie, Cole Mckee, and Adam Duby. Extending threat playbooks for cyber threat intelligence: A novel approach for APT attribution. In *International Symposium on Digital Forensics and Security*, pages 1–6, 2023.

[25] Will Englund, Ellen Nakashima, Gerrit De Vynck, Hannah Denham, Hamza Shaban, Katherine Shaver, Justin Wm. Moyer, Taylor Telford, Brittany Shammas, and Michael Laris. Panic buying strikes Southeastern United States as shuttered pipeline resumes operations. https://www.washingtonpost.com/business/2021/05/12/gas-shortage-colonial-pipeline-live-updates/, May 2021.

[26] European Commission. The EU Cyber Solidarity Act. https://digital-strategy.ec.europa.eu/en/policies/cyber-solidarity, 2024.

[27] Erik Frank. To pay or not to pay: CISOs weigh in on the ransomware dilemma. https://www.csoonline.com/article/3488842/to-pay-or-not-to-pay-cisos-weigh-in-on-the-ransomware-dilemma.html, August 2024.

[28] Google Cloud. Apt45: North korea's digital military machine. https://cloud.google.com/blog/topics/threat-intelligence/apt45-north-korea-digital-military-machine, 2024.

[29] Ian W. Gray, Jack Cable, Benjamin Brown, Vlad Cuiujuclu, and Damon McCoy. Money over morals: A business analysis of Conti ransomware. In *APWG Symposium on Electronic Crime Research*, pages 1–12, 2022.

[30] Harm Griffioen, Tim Booij, and Christian Doerr. Quality evaluation of cyber threat intelligence feeds. In *Applied Cryptography and Network Security*, page 277–296. Springer-Verlag, 2020.

[31] Monique Hennink and Bonnie N Kaiser. Sample sizes for saturation in qualitative research: A systematic review of empirical tests. *Social science & medicine*, 292:114523, 2022.

[32] Monique M Hennink, Bonnie N Kaiser, and Vincent C Marconi. Code saturation versus meaning saturation: how many interviews are enough? *Qualitative health research*, 27(4):591–608, 2017.

[33] HM TOFSI. Ransomware and sanctions: Guidance on ransomware and financial sanctions. https://assets.publishing.service.gov.uk/media/65ca0d7c14b83c000ea716bd/Financial_sanctions_guidance_for_ransomware.pdf, 2024.

[34] Kevin Hobert, Charles Lim, and Eka Budiarto. Enhancing cyber attribution through behavior similarity detection on Linux shell honeypots with ATT&CK framework. In *IEEE International Conference on Cryptography, Informatics, and Cybersecurity*, pages 139–144, 2023.

[35] Danny Yuxing Huang, Maxwell Matthaios Aliapoulios, Vector Guo Li, Luca Invernizzi, Elie Bursztein, Kylie McRoberts, Jonathan Levin, Kirill Levchenko, Alex C Snoeren, and Damon McCoy. Tracking ransomware end-to-end. In *IEEE Symposium on Security and Privacy*, pages 618–631, 2018.

[36] Ehtsham Irshad and Abdul Basit Siddiqui. Cyber threat attribution using unstructured reports in cyber threat intelligence. *Egyptian Informatics Journal*, 24(1):43–59, 2023.

[37] Luke Irwin. Capita admits that its 'cyber incident' was ransomware and that customer data was breached. https://www.itgovernance.co.uk/blog/capita-admits-that-its-cyber-incident-was-ransomware-and-that-customer-data-was-breached, April 2023.

[38] Beomjin Jin, Eunsoo Kim, Hyunwoo Lee, Elisa Bertino, Doowon Kim, and Hyoungshick Kim. Sharing cyber threat intelligence: Does it really help? In *Network and Distributed System Security Symposium*, 2024.

[39] Kyoungmin Kim, Youngsup Shin, Justin Lee, and Kyungho Lee. Automatically attributing mobile threat actors by vectorized ATT&CK matrix and paired indicator. *Sensors*, 21(19):6522, 2021.

[40] Virendra Kumar Yadav, Shelendra Pal, Raghavendra R., Aishwary Awasthi, Laxmi Bewoor, Adapa Gopi, and Sabyasachi

Pramanik. Identification of advancing persistent risks: Expanding the MICTIC model. *Risk Assessment and Countermeasures for Cybersecurity*, pages 20–38, May 2024.

[41] Robert Layton and Paul A Watters. Indirect attribution in cyberspace. *Handbook of Research on Digital Crime, Cyberspace Security, and Information Assurance*, pages 246–262, 2015.

[42] Cristoffer Leite, Jerry Den Hartog, and Daniel Ricardo dos Santos. Using DNS patterns for automated cyber threat attribution. In *International Conference on Availability, Reliability and Security*, pages 1–11, 2024.

[43] Vector Guo Li, Matthew Dunn, Paul Pearce, Damon McCoy, Geoffrey M Voelker, and Stefan Savage. Reading the tea leaves: A comparative analysis of threat intelligence. In *USENIX Security Symposium*, pages 851–867, 2019.

[44] Mandiant. Uncategorized (UNC) threat groups. https://www.mandiant.com/resources/insights/uncategorized-unc-threat-groups.

[45] Mandiant. Apt41, a dual espionage and cyber crime operation. https://www.mandiant.com/sites/default/files/2022-02/rt-apt41-dual-operation.pdf, 2022.

[46] Mandiant. To hades and back: Unc2165 shifts to lockbit to evade sanctions. https://cloud.google.com/blog/topics/threat-intelligence/unc2165-shifts-to-evade-sanctions, 2022.

[47] Lennart Maschmeyer, Ronald J Deibert, and Jon R Lindsay. A tale of two cybers-how threat reporting by cybersecurity firms systematically underrepresents threats to civil society. *Journal of Information Technology & Politics*, 18(1):1–20, 2021.

[48] William P Maxam III and James C Davis. An interview study on third-party cyber threat hunting processes in the US department of homeland security. In *USENIX Security Symposium*, 2024.

[49] Andrew McCoomb, Ailsa Robertson, John Cassell, and Imran Ahmad. Season 1, episode 2: Ransomware. https://www.nortonrosefulbright.com/en-ca/knowledge/podcasts/disputed-podcast/2021/q4/episode-2-ransomware, July 2021.

[50] Timothy McIntosh, ASM Kayes, Yi-Ping Phoebe Chen, Alex Ng, and Paul Watters. Ransomware mitigation in the modern era: A comprehensive review, research challenges, and future directions. *ACM Computing Surveys*, 54(9):1–36, 2021.

[51] Timothy McIntosh, Teo Susnjak, Tong Liu, Dan Xu, Paul Watters, Dongwei Liu, Yaqi Hao, Alex Ng, and Malka Halgamuge. Ransomware reloaded: Re-examining its trend, research and mitigation in the era of data exfiltration. *ACM Computing Surveys*, 57(1), October 2024.

[52] Yangyang Mei, Weihong Han, Shudong Li, Xiaobo Wu, Kai-Han Lin, and Yulu Qi. A review of attribution technical for APT attacks. In *IEEE International Conference on Data Science in Cyberspace*, pages 512–518, 2022.

[53] Per Håkon Meland, Yara Fareed Fahmy Bayoumy, and Guttorm Sindre. The Ransomware-as-a-Service economy within the darknet. *Computers & Security*, 92:101762, 2020.

[54] Tom Meurs, Raphael Hoheisel, Marianne Junger, Abhishta Abhishta, and Damon McCoy. What to do against ransomware? evaluating law enforcement interventions. In *Symposium on Electronic Crime Research, eCrime 2024*, 2024.

[55] Tom Meurs, Marianne Junger, Erik Tews, and Abhishta Abhishta. Ransomware: How attacker's effort, victim characteristics and context influence ransom requested, payment and financial loss. In *APWG symposium on electronic crime research*, pages 1–13, 2022.

[56] Microsoft. How Microsoft names threat actors. https://learn.microsoft.com/en-us/defender-xdr/microsoft-threat-actor-naming.

[57] Ricardo Misael Ayala Molina, Sadegh Torabi, Khaled Sarieddine, Elias Bou-Harb, Nizar Bouguila, and Chadi Assi. On ransomware family attribution using pre-attack paranoia activities. *IEEE Transactions on Network and Service Management*, 19:19–36, 2022.

[58] Erica Moret and Patryk Pawlak. The EU cyber diplomacy toolbox: towards a cyber sanctions regime? Technical report, European Union Institute for Security Studies, 2017.

[59] Routa Moussaileb, Nora Cuppens, Jean-Louis Lanet, and Hélène Le Bouder. A survey on windows-based ransomware taxonomy and detection mechanisms. *ACM Computing Surveys*, 54(6):1–36, 2021.

[60] Karina Nad. Ransomware warfare: Exploring global and private negotiations to help US victims respond to the threat. *Cardozo Journal of Conflict Resolution*, 23:257, 2022.

[61] Umara Noor, Zahid Anwar, Tehmina Amjad, and Kim-Kwang Raymond Choo. A machine learning-based FinTech cyber threat attribution framework using high-level indicators of compromise. *Future Generation Computer Systems*, 96:227–242, 2019.

[62] Northdoor. Sanctions checker | financial sanctions check & screening. https://www.northdoor.co.uk/sanctions-checker/.

[63] Sean O'Connell. To ban ransomware payments or not to ban ransomware payments: The problems of drafting legislation in response to ransomware. *Journal of International Business & Law*, 22:151, 2023.

[64] Office of Foreign Assets Control. OFAC sanction list search. https://sanctionssearch.ofac.treas.gov.

[65] Ayodeji Ogundiran, Hongmei Chi, Jie Yan, and Ruth Agada. A framework to reconstruct digital forensics evidence via goal-oriented modeling. In *IEEE International Conference on AI in Cybersecurity*, pages 1–11, 2023.

[66] Kris Oosthoek, Jack Cable, and Georgios Smaragdakis. A tale of two markets: Investigating the ransomware payments economy. *Communications of the ACM*, 66(8):74–83, 2023.

[67] Harun Oz, Ahmet Aris, Albert Levi, and A Selcuk Uluagac. A survey on ransomware: Evolution, taxonomy, and defense solutions. *ACM Computing Surveys*, 54(11s):1–37, 2022.

[68] Constantinos Patsakis, David Arroyo, and Fran Casino. The malware as a service ecosystem. *arXiv preprint arXiv:2405.04109*, 2024.

[69] Constantinos Patsakis, Eugenia Politou, Efthimios Alepis, and Julio Hernandez-Castro. Cashing out crypto: state of practice in ransom payments. *International Journal of Information Security*, 23(2):699–712, 2024.

[70] Md Rayhanur Rahman, Rezvan Mahdavi Hezaveh, and Laurie Williams. What are the attackers doing now? Automating cyberthreat intelligence extraction from text on pace with the changing threat landscape: A survey. *ACM Computing Surveys*, 55(12):1–36, 2023.

[71] Md Rayhanur Rahman and Laurie Williams. Investigating co-occurrences of MITRE ATT&CK techniques. *arXiv preprint arXiv:2211.06495*, 2022.

[72] Thomas Rid and Ben Buchanan. Attributing cyber attacks. *Journal of Strategic Studies*, 38(1-2):4–37, 2015.

[73] Thomas Roccia. *Visual Threat Intelligence: An Illustrated Guide For Threat Researchers*. 2023.

[74] Estelle Ruellan, Masarah Paquet-Clouston, and Sebastian Garcia. Conti Inc.: Understanding the internal discussions of a large ransomware-as-a-service operator with machine learning. *Crime Science*, 13(1):16, 2024.

[75] Vinay Sachidananda, Rajendra Patil, Akshay Sachdeva, Kwok-Yan Lam, and Liu Yang. APTer: Towards the investigation of APT attribution. In *IEEE Conference on Dependable and Secure Computing*, pages 1–10, 2023.

[76] Jake Sepich. The evolution of cyber attribution. https://www.american.edu/sis/centers/security-technology/the-evolution-of-cyber-attribution.cfm, April 2023.

[77] Jawwad A Shamsi, Sherali Zeadally, Fareha Sheikh, and Angelyn Flowers. Attribution in cyberspace: techniques and legal implications. *Security and Communication Networks*, 9(15):2886–2900, 2016.

[78] Youngsup Shin, Kyoungmin Kim, Jemin Justin Lee, and Kyungho Lee. Focusing on the weakest link: A similarity analysis on phishing campaigns based on the ATT&CK matrix. *Security and Communication Networks*, 2022(1):1699657, 2022.

[79] Sikich. The role of digital forensics in fighting and preventing cybercrime. https://www.sikich.com/insight/the-role-of-digital-forensics-in-fighting-and-preventing-cybercrime/, April 2023.

[80] Simmons & Simmons. Cyber extortion – how should insurers respond? https://www.simmons-simmons.com/en/publications/ck5cr1wq11rr90988j6zjanx8/cyber-extortion-âĂŞ-how-should-insurers-respond, 2020.

[81] Sklearn Documentation. silhouette_score. https://scikit-learn.org/stable/modules/generated/sklearn.metrics.silhouette_score.html, 2024.

[82] Florian Skopik and Timea Pahi. Under false flag: using technical artifacts for cyber attack attribution. *Cybersecurity*, 3:1–20, 2020.

[83] Zheyu Song, Yonghong Tian, and Junjin Zhang. Similarity analysis of ransomware attacks based on ATT&CK matrix. *IEEE Access*, 2023.

[84] Sophos. The state of ransomware 2024. https://assets.sophos.com/X24WTUEQ/at/9brgj5n44hqvgsp5f5bqcps/sophos-state-of-ransomware-2024-wp.pdf, April 2024.

[85] Fabian M. Teichmann and Sonia R. Boticiu. The most impactful ransomware attacks in 2023 and their business implications. *International Cybersecurity Law Review*, April 2024.

[86] Delbert Tran. The law of attribution: Rules for attribution the source of a cyber-attack. *Yale JL & Tech.*, 20:376, 2018.

[87] United Nations. Sanction tracker. https://data.europa.eu/apps/eusanctionstracker/, 2024.

[88] US Office for Foreign Assets Control. Updated advisory on potential sanctions risks for facilitating ransomware payments. https://ofac.treasury.gov/media/912981/download?inline, 2021.

[89] Max van der Horst, Ricky Kho, Olga Gadyatskaya, Michel Mollema, Michel van Eeten, and Yury Zhauniarovich. Supplementary material and data for this study, https://doi.org/10.5281/zenodo.14732550, 2025.

[90] Aldin Vehabovic, Nasir Ghani, Elias Bou-Harb, Jorge Crichigno, and Aysegul Yayimli. Ransomware detection and classification strategies. In *IEEE International Black Sea Conference on Communications and Networking*, June 2022.

[91] Mattias Wahlen. Using kill chain analysis in ransomware attacks. https://www.truesec.com/hub/blog/using-killchain-analysis-in-ransomware-attacks, 2023.

[92] Kai Wang, Jun Pang, Dingjie Chen, Yu Zhao, Dapeng Huang, Chen Chen, and Weili Han. A large-scale empirical analysis of ransomware activities in bitcoin. *ACM Transactions on the Web*, 16(2):1–29, 2021.

[93] Yichao Wang, Sophia Roscoe, Budi Arief, Lena Connolly, Hervé Borrion, and Sanaa Kaddoura. The social and technological incentives for cybercriminals to engage in ransomware activities. In *International Symposium on Security and Privacy in Social Networks and Big Data*, pages 149–163. Springer, 2023.

[94] Arun Warikoo. The triangle model for cyber threat attribution. *Journal of Cyber Security Technology*, 5(3-4):191–208, 2021.

[95] David A Wheeler and Gregory N Larsen. Techniques for cyber attack attribution. *Institute for Defense Analysis*, page 2, 2003.

[96] Yiming Wu, Qianjun Liu, Xiaojing Liao, Shouling Ji, Peng Wang, Xiaofeng Wang, Chunming Wu, and Zhao Li. Price tag: towards semi-automatically discovery tactics, techniques and procedures of e-commerce cyber threat intelligence. *IEEE Transactions on Dependable and Secure Computing*, 2021.

[97] Yixin Wu, Cheng Huang, Xing Zhang, and Hongyi Zhou. GroupTracer: Automatic attacker TTP profile extraction and group cluster in Internet of Things. *Security and Communication Networks*, 2020(1):8842539, 2020.

[98] Nan Xiao, Bo Lang, Ting Wang, and Yikai Chen. APT-MMF: An advanced persistent threat actor attribution method based on multimodal and multilevel feature fusion. *Computers & Security*, 144, October 2024.

[99] Miuyin Yong Wong, Matthew Landen, Manos Antonakakis, Douglas M Blough, Elissa M Redmiles, and Mustaque Ahamad. An inside look into the practice of malware analysis. In *ACM SIGSAC Conference on Computer and Communications Security*, pages 3053–3069, 2021.

[100] Yizhe You, Jun Jiang, Zhengwei Jiang, Peian Yang, Baoxu Liu, Huamin Feng, Xuren Wang, and Ning Li. Tim: threat context-enhanced ttp intelligence mining on unstructured threat data. *Cybersecurity*, 5(1):3, 2022.

[101] Alena Yuryna Connolly and Hervé Borrion. Reducing ransomware crime: Analysis of victims' payment decisions. *Computers & Security*, 119, August 2022.

[102] Lena Yuryna Connolly, David S Wall, Michael Lang, and Bruce Oddson. An empirical study of ransomware attacks on organizations: an assessment of severity and salient factors affecting vulnerability. *Journal of Cybersecurity*, 6(1), 12 2020.

A Research Protocol

A.1 Research Questions

- SQ1: What processes and procedures do cybersecurity practitioners follow to attribute an attack to a ransomware group in service of sanction screening?
- SQ2: What techniques and indicators are currently used for ransomware threat actor attribution?
- SQ3: To what extent are high-level indicators reliable in the identification of ransomware groups?
- SQ4: What needs to be improved to further develop ransomware attribution standards?

Table 9 lists our interview questions.

A.2 Collection of Empirical Data

Invitation and Explanation.
You are being invited to participate in a research study titled "Leveraging High-Level Indicators: Correlating Ransomware Attacks to Threat Actors". This study is being done by <responsible researcher> from the <institution> and supervised by <company>.

The purpose of this research study is to find out whether high-level indicators can help correlate ransomware attacks to threat actors by the means of audio recording and will take you approximately 90 minutes to complete. Therefore, the participants are experts in identifying cyber-threat actors (Cyber-threat attribution). The data will be used for improving the state of art on the current cyber-threat attribution process and their current limitations. We will be asking you to provide information about the current cyber-threat attribution process and the limitations. In addition, improvement points or ideas to improve the current cyber-threat attribution process are appreciated.

As with any online activity, the risk of a breach is always possible. To the best of our ability, your answers to this study will remain confidential. We will minimize any risks by storing the data in a Project Storage at <institution>, which allows for access restrictions in such a way that only authorized members can access the data. By doing this, the risk of a data leak, which can lead to reputational risk, is minimized. The information will be anonymized and only the function and a small job description will be used. The information you provide will be synthesized in an anonymous summary. The summary will be sent to you for review and will be used for analysis purposes. The summary will be made publicly available with the final report. Should you have any concerns regarding the content of the summary, you will be welcome to oppose its publication.

Your participation in this study is entirely voluntary, and you can withdraw at any time. You are free to omit any questions. The participant has the right to request access to provided data and can demand to rectify or erase personal data.

If there are any questions before/after the interview, you can contact me with the following contact details: <details>

Table 9: Interview Questions.

Segment 1: Overview of Cyber Threat Attribution
1. How are cyber threat actors typically categorized based on their characteristics and traits?
2. Can you describe the techniques or methods that your organization currently employs for cyber threat actor attribution of ransomware groups?
3. In your experience, which indicators are typically considered when attributing a cyber threat to a ransomware actor or group?
4. Are there differences in cyber threat attribution techniques when dealing with different types of threat actors, such as state-sponsored groups, cybercriminals, or hacktivists?
5. Are there different levels of attribution? If so, do you observe differences in the level of attribution when comparing different organizations, such as law enforcement and cybersecurity companies?
6. At what level is the ransomware attacker identified?
Segment 2: Strengths and Limitations
7. What do you consider the main strengths of the attribution techniques or methods you use?
8. What limitations or challenges have you encountered when attempting to attribute cyber threats to specific actors or groups?
9. Can you provide an example of a recent or notable case of cyber threat attribution you have worked on, and walk through the process of attribution, including the techniques and indicators used?
10. In your opinion, how important is it to consider the potential risk of false attribution in the field of attributing cyber threat actors, and how do you mitigate this risk?
11. Are there specific legal or ethical considerations that impact your approach to attributing cyber threat actors, and if so, how do they influence your work?
12. Can you share examples of cases where attribution efforts did not lead to a clear identification of the threat actor? What were the main challenges in these cases?
Segment 3: Adaptation and Future Developments
13. How do you stay informed about evolving techniques and indicators in the field of attributing cyber threat actors, and how do you adjust your methods accordingly?
14. In your opinion, what are the most significant areas of improvement or development needed in the field of attributing cyber threat actors?

B Interview Codes and Themes

Table 10: Consolidated Summary of Themes and Codes.

Theme	Codes
Attribution Methods	Low-Level Indicators, Ransomware Analysis Methods, High-Level Indicators are Generic, Tools, Detection Methods, Darknet Investigation, Modus Operandi supports Attribution, Ransomware Code Patterns, High-Level Indicators are Valuable, Attribution by Activity Clustering, Ransomware groups outing themselves with ransom notes, Root Cause Analysis, High-Level Indicators, Communication Channel Analysis, Infrastructure Analysis, Attacker Profiling, Open-Source Intelligence for Attribution, Different Forms of Attribution, Pyramid of Pain in Attribution, High-level indicators as a way to conceptualize behavior, Strengths of the Attribution Process, group-level attribution
Attribution Objectives	Stopping at Group-Level Attribution during Investigations, Benefits of Attribution
Challenges with Attribution	Challenges in Individual-Level Attribution, Legal Challenges in Attribution, Confidence Levels in Attribution, Added Value of Active Ransomware Attribution, Ransomware Rebranding, Improvement Points in Attribution, Performing Country-Level Attribution, Complexities due to Data Fragmentation, Complexity in Correlating Indicators, Lack of Attribution Data, Attribution Accuracy, Attribution requires diverse knowledge and sources of information, Reputational Repercussions of Misattribution, Organizations Publishing Inaccurate Attributions
Ethical and Legal Considerations	Legal Considerations, Political Repercussions of Attribution, Ethical Considerations
Knowledge Management and Collaboration	The Need for an Centralized Knowledge Management Platform, Essence of Collaboration, Improving using Information Sharing
Operational Constraints	Resource Constraints in investigations, Differences Law Enforcement and Industry
Sanction Compliance	Attribution for Sanction Checks, Consequences of Inaccurate Sanction Screening, Verification against Sanction Lists, Legal Ambiguities within Sanction Screening
Threat Actor Motives and Behaviors	State vs. Criminal Actors, Ransomware Actor Motives, False Flag Operations, Ransomware Observed Procedures, Ransomware-as-a-Service, Ransomware Group Brand and Reputation
Challenges with Indicators	Complexity in Correlating Indicators, Lack of Attribution Data, Attribution Accuracy
Threat Actor Typologies	Nation-State Actors, Cybercriminal Groups, Hacktivists

C CISA #StopRansomware Reports

Table 11: CISA Report URLs.

RTA	URL	Published (YYYY-MM-DD)
ALPHV Blackcat	https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-353a	2023-12-19
Lockbit	https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-075a	2023-03-16
Akira	https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-109a	2023-04-18
Phobos	https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-060a	2023-02-29
Rhysida	https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-319a	2023-11-15
Black Basta	https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-131a	2023-03-10
Play	https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-352a	2023-12-18
AvosLocker	https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-284a	2023-10-11
Snatch	https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-263a	2023-09-20
BianLian	https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-136a	2023-05-16
CL0P	https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-158a	2023-06-07
DPRK	https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-040a	2023-02-09