DLINE JOURNALS

Journal of Information Security Research

Print ISSN: 0976-4143
Online ISSN: 0976-4151

JISR 2026: 17 (1)
https://doi.org/10.6025/jisr/2026/17/1/1-22

Sector-Aware Cyber Threat Intelligence: A Security-Enhanced RAG Framework for Precision Threat Analysis

Maleerat Maliyaem
Faculty of Information Technology, King Mongkut's University of Technology North Bangkok, Bangkok. Thailand
maleerat.m@itd.kmutnb.ac.th

ABSTRACT

This paper presents a comprehensive framework for enhancing Cyber Threat Intelligence (CTI) analysis through a domain specific, security aware Retrieval Augmented Generation (RAG) architecture. It begins by defining CTI as the process of transforming raw cyber data into actionable intelligence, emphasizing its role in understanding adversarial Tactics, Techniques, and Procedures (TTPs) via structured knowledge bases like MITRE ATT&CK. The paper reviews the evolution of CTI analysis from traditional rule based and machine-learning systems to modern Large Language Model (LLM) driven approaches highlighting persistent challenges such as hallucinations, a lack of contextual grounding, and vulnerability to adversarial manipulation in retrieval pipelines.

Drawing on 2025 CrowdStrike threat data, the analysis reveals critical trends: a 27% year over year rise in interactive intrusions, 73.4% of which are financially motivated (eCrime), while 26.5% stem from nation-state actors targeting strategic sectors like Government (+126%), Telecommunications (+130%), and Industrials & Engineering (+185%). Opportunistic attacks declined by 12%, indicating a shift toward high-value, precision targeting.

These insights inform a proposed CTI RAG system that integrates real world threat statistics to enable sector aware retrieval, sector weighted ATT&CK alignment, and evidence grounded reasoning. Unlike generic or LLM-only baselines, this approach reduces the risk of hallucination, improves TTP accuracy, shortens analyst triage time, and enhances explainability by linking outputs directly to verified CTI sources and sector-specific tactics.

The architecture is designed as an end to end, modular pipeline supporting ingestion, semantic indexing, secure retrieval, grounded reasoning, and analyst feedback making it suitable for operational Security Operations Centers (SOCs). By unifying empirical threat trends, structured knowledge, and robust AI reasoning, the framework offers a scalable, trustworthy solution for next generation CTI that balances tactical eCrime defence with strategic awareness of nation state threats.

Keywords: Cyber Threat Intelligence (CTI), Retrieval-Augmented Generation (RAG), Large Language Models (LLMs), MITRE ATT&CK, Sector-Specific Threat Modeling, eCrime, Nation-State Actors, Hallucination Mitigation, Security-Aware AI

Received: 18 September 2025, Revised 4 December 2025, Accepted 21 December 2025

1. Introduction

Cyber Threat Intelligence (CTI) refers to the systematic process of collecting, analyzing, and disseminating actionable information regarding existing, emerging, or potential cyber threats that may compromise an organization's digital assets, operational continuity, or reputation. Unlike raw security data, CTI transforms fragmented technical signals into meaningful intelligence that supports informed decision making and proactive defense strategies. CTI outputs may range from low level indicators such as Indicators of Compromise (IoCs), attack vectors, and malware signatures, to higher level contextual intelligence describing adversarial intent, capabilities, and operational patterns.

CTI plays a critical role in understanding adversarial Tactics, Techniques, and Procedures (TTPs**)**, which describe how attackers plan, execute, and sustain cyber operations. By mapping observed behaviors to structured knowledge bases such as MITRE ATT&CK, CTI enables defenders to anticipate future attack stages, strengthen detection capabilities, and deploy targeted countermeasures. CTI data is typically aggregated from heterogeneous sources, including Open Source Intelligence (OSINT), industry reports, vulnerability disclosures (e.g., CVE), threat feeds, malware analyses, and discussions from underground cybercrime forums.

However, the value of CTI is tightly coupled with the effectiveness of its analysis. The diversity, scale, and unstructured nature of CTI data introduce significant analytical challenges, motivating the development of automated and intelligent analysis frameworks.

2. Background

2.1 Attack Investigation and Attribution Using CTI

Attack investigation is a core CTI task that aims to reconstruct the sequence of events surrounding a cyber incident and infer critical characteristics of the attack. These include the attacker's tools, exploited vulnerabilities, operational techniques, and strategic objectives. A closely related and particularly challenging problem is attack attribution, which seeks to identify the responsible individual, group, or nation state actor behind an attack.

Attribution is inherently difficult due to deliberate adversarial obfuscation, the reuse of tools by multiple threat actors, false flag operations, and incomplete or noisy evidence. As highlighted by Alhuzali et al. [1], attackers intentionally manipulate artefacts to mislead analysts, making definitive attribution elusive even with substantial technical evidence.

The investigation process is further complicated by the heterogeneity and temporal dependency of CTI data. Information may arrive incrementally over time, span multiple formats (textual reports, logs, indicators), and originate from sources with varying credibility and completeness. Manual investigation under such conditions is time consuming, error prone, and does not scale to the volume of modern cyber threat data. Consequently, the research community has explored automated and semi automated CTI analysis methods to

improve scalability, consistency, and analytical depth.

3. Related literature

3.1 Traditional Automated CTI Analysis Approaches

Early automated CTI analysis methods primarily relied on rule based systems and classical machine learning techniques**.** These approaches aimed to assist analysts by automating tasks such as attack pattern recognition, incident clustering, and basic attribution support. Studies by Liao et al. [2], Gao et al. [3], and Arikkat et al.[4] demonstrate that automation can partially alleviate the burden of manual analysis and improve consistency across investigations.

Despite their benefits, traditional approaches exhibit several limitations. They often require extensive feature engineering, handcrafted rules, and deep domain expertise to remain effective. [5, 6, 7]Fengxiao Tang, Santos, P, Pedro Santos Rafael,] As adversaries continuously evolve their tactics, static rules and manually engineered features quickly become outdated. [8, 9] McCall, Bassog,] Furthermore, classical machine learning models struggle to capture complex semantic relationships and contextual dependencies present in unstructured threat reports.

These limitations constrain the adaptability of traditional methods and motivate the exploration of more flexible, data driven models that can reason across diverse CTI sources.

3.2 Emergence of Large Language Models for CTI Reasoning

Recent advances in Large Language Models (LLMs) have introduced new opportunities for CTI analysis, particularly in tasks involving unstructured text, semantic understanding, and reasoning. LLMs can process large volumes of natural language data, making them well suited for analyzing threat reports, security advisories, and cybercrime forum discussions.

Empirical studies show that LLMs can achieve high accuracy in extracting CTI variables, summarizing threat narratives, and identifying attack characteristics. For instance, Clairoux Trepanier et al. [10] evaluated an LLM based system using GPT 3.5 turbo on more than 700 conversations from underground forums such as XSS, Exploit_in, and RAMP. The system achieved an average accuracy of 96.23%, with strong precision and recall when extracting key CTI variables, demonstrating the feasibility of LLM driven CTI extraction.

Similarly, Divakaran and Peddinti [11] argue that LLMs significantly enhance attack investigation by improving the processing and reasoning over massive, heterogeneous CTI datasets. By leveraging pretrained knowledge and contextual understanding, LLMs enable analysts to synthesize insights from previously inaccessible sources.

However, the adoption of LLMs in cybersecurity is not without risks. A major concern is hallucination, where models generate fluent but incorrect or unverified outputs. Martino et al. [12] and Perkoviæ et al. [13] highlight that hallucinations pose serious threats in security critical applications, as inaccurate intelligence can lead to flawed defensive decisions. These limitations underscore the need for mechanisms that ground LLM outputs in verifiable evidence.

3.3 Retrieval-Augmented Generation for Grounded CTI Analysis

Retrieval-Augmented Generation (RAG) has emerged as a promising solution to mitigate hallucinations and improve factual grounding in LLM outputs. RAG combines information retrieval with generative modeling, allowing LLMs to generate responses conditioned on retrieved external documents rather than relying solely on internal model knowledge.

In CTI contexts, RAG enables models to ground their reasoning in authoritative sources such as threat reports, vulnerability databases, and structured knowledge bases. Divakaran and Peddinti (2024) note that RAG improves factual correctness and domain relevance, particularly for knowledge intensive tasks.

Despite these advantages, RAG introduces new challenges. Bang et al. [14] conduct a comparative analysis of RAG and non RAG frameworks across eleven LLMs and find that RAG can significantly alter a model's safety profile. Even when combining safe models with benign documents, RAG systems may produce unsafe outputs. Furthermore, conventional red teaming techniques prove less effective in RAG settings, revealing gaps in current AI safety methodologies.

These findings indicate that while RAG improves grounding, it also introduces new attack surfaces, including corpus poisoning and inter context conflicts. As a result, RAG systems require specialized safety evaluation and defense mechanisms tailored to their unique architecture.

3.4 Domain-Specific RAG for Adversarial Technique Identification

Generic RAG frameworks often rely on off the shelf retrievers trained on broad corpora, which may retrieve noisy or weakly relevant documents for specialized domains like cybersecurity. This limitation motivates the development of domain specific RAG approaches.

Lekssays et al. [15] propose TechniqueRAG, a domain specific RAG framework designed to accurately identify adversarial techniques from security texts. Technique RAG addresses a fundamental trade off in existing methods: generic models lack domain precision, while specialized pipelines require extensive labeled data and computational resources.

The proposed framework integrates off the shelf retrievers with instruction tuned LLMs and minimal text technique pairs. To improve retrieval relevance, TechniqueRAG employs zero shot LLM based re ranking that explicitly aligns retrieved documents with adversarial techniques. Experimental evaluations across multiple security benchmarks demonstrate state of the art performance without extensive task specific optimisations or large labelled datasets. [16, 17, 18]

A series of advanced RAG techniques have been proposed recently, including the ITER RETGEN method proposed by Shao et al. [19]. FLARE method by Jiang et al. [20] IRCoT, by Trivedi et al. [21], Selective Context, proposed by Li et al. [22] and the SuRe method by Kim et al.

TechniqueRAG highlights the importance of tailoring RAG pipelines to domain specific semantics, particularly in CTI applications where precision and contextual alignment are critical.

3.5 LLM-Based Interpretation and Summarization of TTPs

Interpreting TTPs within the MITRE ATT & CK framework is a cognitively demanding task, often requiring substantial expertise to understand procedural dependencies and attacker intent. Recent studies investigate how LLMs can support TTP comprehension and summarization.

Fayyazi et al. [24] compare Supervised Fine Tuning (SFT) of encoder only models (e.g., RoBERTa) with RAG-enhanced decoder only LLMs (e.g., GPT-3.5) for TTP interpretation. Their findings show that RAG based approaches consistently outperform SFT, particularly in low resource settings where labeled data is scarce.

The results suggest that retrieving relevant procedural context is more effective than fine tuning alone, reinforcing the role of RAG as a flexible and scalable method for enhancing LLM based cybersecurity analysis.

3.6 Security and Robustness of RAG Systems

While RAG improves grounding, its reliance on external knowledge introduces vulnerabilities. Attackers may manipulate retrieved content to influence model outputs, thereby introducing misinformation or degrading system reliability.

Yao et al. [25] propose EcoSafeRAG, a defense framework that detects malicious or poisoned documents through sentence level context diversity analysis. Unlike prior methods, EcoSafeRAG does not rely on the LLM's internal knowledge, aligning more closely with RAG's design philosophy. Experiments show that EcoSafeRAG improves security while reducing latency and token usage compared to vanilla RAG pipelines.

Complementarily, Liang et al. [26] introduce SafeRAG, a benchmark designed to systematically evaluate RAG security. SafeRAG categorizes attack scenarios into silver noise, inter context conflict, soft advertising, and white Denial of Service attacks. Evaluations across fourteen RAG components reveal that even advanced retrievers and LLMs remain vulnerable, highlighting the need for robust, standardized evaluation frameworks.

3.7 End-to-End RAG Architectures for CTI

Beyond isolated tasks, recent research explores end-to-end RAG architectures for comprehensive CTI analysis. Gandhi et al. [27] propose a hybrid pipeline that integrates dense vector retrieval, semantic indexing, and knowledge graph reasoning. The system ingests heterogeneous CTI sources, applies semantic chunking and entity linking, and indexes embeddings alongside a cybersecurity knowledge graph.

Evaluated on 2,400 CTI incident reports and synthetic network alerts, the architecture achieves a 25.7% improvement in detection F1-score, a 31% reduction in analyst triage time, and a 45% reduction in hallucinated outputs compared to keyword based baselines. These results demonstrate the effectiveness of combining RAG with structured reasoning mechanisms for scalable and reliable CTI analysis.

3.8 Research Gaps and Motivation

Across the literature, RAG-based approaches consistently enhance CTI extraction, reasoning, TTP interpretation, and analyst productivity. However, several open challenges remain. These include securing retrieval pipelines against adversarial manipulation, mitigating hallucinations under complex threat scenarios, and ensuring trustworthy deployment in operational environments.

Addressing these gaps requires security aware, domain specific RAG architectures that balance performance,

robustness, and safety. Such systems are essential for the next generation of reliable, AI-driven cyber threat intelligence platforms.

Category	Method / Study	Primary Task	Dataset / Source	Key Contributions & Findings	Limitations
Traditional Automated CTI	Liao et al. (2016)	Attack investigation and correlation	Heterogeneous CTI sources	Demonstrated feasibility of automated CTI analysis using classical ML and rule-based systems	Requires extensive feature engineering; limited adaptability to evolving threats
	Gao et al. (2023a)	Threat detection and attribution support	OSINT and structured CTI feeds	Improved scalability over manual analysis; supports structured threat inference	Performance degrades with unstructured data; high domain-dependence
	Arikkat et al. (2024)	Automated attack investigation	Multi-source CTI repositories	Highlights need for automation due to CTI scale and heterogeneity	Limited semantic reasoning capability
	Alhuzali et al. (2021)	Attack attribution	Incident-level CTI evidence	Identified attribution challenges caused by adversarial obfuscation	Attribution remains inconclusive under sparse or deceptive evidence
LLM-based CTI Reasoning	Clairoux-Trepanier et al. (2023)	CTI extraction and summarization	700+ cybercrime forum conversations	Achieved high accuracy (96.23%) in CTI variable extraction using GPT-3.5	Sensitive to prompt design; vulnerable to hallucinations
	Divakaran & Peddinti (2024)	CTI reasoning and investigation	Unstructured threat reports	Demonstrated LLM effectiveness in large-scale CTI reasoning	Lack of grounding leads to hallucination risks

	Martino et al. (2023); Perkoviæ et al. (2024)	Hallucination analysis	General LLM benchmarks	Identified hallucinations as a critical risk in safety-critical domains	No built-in mitigation mechanism
Generic RAG Frame works	Bang et al. (2024)	Safety evaluation of RAG	External document corpora	Showed RAG can alter model safety profiles and introduce unsafe outputs	Standard red-teaming ineffective for RAG
	Divakaran & Peddinti (2024)	Grounded reasoning	Knowledge bases + LLMs	Improved factual correctness and relevance via retrieval grounding	Introduces new attack surfaces
Domain -Specific RAG	Lekssays et al. (2024) – *Technique RAG*	Adversarial technique identification	Cybersecurity benchmark datasets	Achieved state-of-the-art results with minimal labeled data using zero-shot LLM re-ranking	Dependent on retriever quality
TTP Interpretation & Summarization	Fayyazi et al. (2024)	MITRE ATT&CK TTP comprehension	ATT&CK procedures	RAG-based decoder-only LLMs outperform SFT encoder models	Requires reliable retrieval context
RAG Security & Robustness	Yao et al. (2024) – *EcoSafeRAG*	Poisoning detection and defense	Synthetic adversarial corpora	Improved RAG security with reduced latency and token usage	Focuses primarily on retrieval-layer threats
	Liang et al. (2025) – *SafeRAG*	RAG vulnerability benchmarking	Manually curated attack datasets	Demonstrated widespread RAG vulnerability across attack types	Evaluation-focused; no mitigation
End-to-End CTI RAG Systems	Gandhi et al. (2024)	CTI ingestion, correlation, detection	2,400 CTI reports + alerts	Improved detection F1 by 25.7%, reduced analyst workload and hallucinations	Reliance on corpus quality; privacy risks

Table 1. Existing Methods for Cyber Threat Intelligence (CTI) Analysis

Proposed Method vs. Prior Work

Aspect	**Prior Work**	**Proposed Method**
CTI Coverage	Focused on isolated tasks (extraction, attribution, TTP mapping)	Unified end-to-end CTI analysis pipeline
Use of LLMs	LLM-only or generic RAG	Domain-specific, security-aware RAG
Retrieval Strategy	Generic dense or keyword retrieval	Semantically filtered, CTI-aligned retrieval
Hallucination Control	Limited or implicit	Explicit grounding with verified CTI sources
Security Awareness	Mostly evaluation-focused (e.g., SafeRAG)	Built-in mitigation against retrieval poisoning and inter-context conflicts
Domain Adaptation	Requires extensive labeled data or fine-tuning	Minimal supervision with domain-aligned prompts and retrieval
Explainability	Often opaque	Evidence-linked threat reasoning and traceable outputs
Analyst Support	Partial automation	Actionable insights with reduced cognitive load
Operational Readiness	Prototype-level or task-specific	Designed for real-world CTI workflows

Table 2. Proposed Method vs. Prior Work

4.1 Comparison with the State-of-the-Art

This section compares the proposed method with state-of-the-art approaches to Cyber Threat Intelligence (CTI) analysis, highlighting methodological advances, performance implications, and practical relevance. The comparison spans traditional automated CTI systems, Large Language Model (LLM)–based approaches, generic Retrieval-Augmented Generation (RAG) frameworks, and recent domain specific and security-aware RAG solutions.

4.2 Comparison with Traditional Automated CTI Methods

Traditional automated CTI systems primarily rely on rule based reasoning and classical machine learning techniques to assist with attack investigation and attribution [2, 3, 4]. These methods offer improved scalability over manual analysis but require extensive feature engineering and domain expertise. Their effectiveness is further limited when confronted with unstructured and rapidly evolving threat data.

In contrast, the proposed method eliminates the need for handcrafted rules by leveraging LLM-driven semantic reasoning grounded through retrieval. Unlike traditional pipelines that operate on predefined features, the proposed architecture directly reasons over heterogeneous CTI sources, including narrative threat reports and adversarial discussions. This enables greater adaptability to emerging threats and reduces maintenance overhead, addressing a fundamental limitation of classical CTI automation.

4.3 Comparison with LLM-Only CTI Approaches

Recent studies demonstrate the effectiveness of LLMs in CTI extraction and reasoning tasks, particularly for unstructured data such as cybercrime forums and threat narratives [10, 11]. While these approaches achieve high extraction accuracy, they remain susceptible to hallucinations and lack strong guarantees of factual grounding [12, 13].

The proposed method advances beyond LLM only approaches by explicitly grounding generation in retrieved, domain relevant CTI evidence. Rather than relying on implicit model knowledge, the system constrains reasoning to verified external sources, thereby reducing the risk of fabricated or misleading outputs. This distinction is particularly critical in cybersecurity settings, where erroneous intelligence may lead to inappropriate defensive actions or misattribution.

4.4 Comparison with Generic RAG Frameworks

Generic RAG frameworks improve factual correctness by integrating retrieval with generation, but are typically designed for broad knowledge tasks rather than domain critical applications. Safety analyses reveal that RAG can unintentionally degrade model safety, even when benign models and documents are used [14]. Additionally, generic retrievers may introduce noisy or weakly relevant context, reducing domain precision.

The proposed method differentiates itself by adopting a domain specific RAG design, in which retrieval is aligned with CTI semantics and adversarial behaviors. By filtering and prioritizing CTI relevant context, the approach minimizes inter context conflicts and improves reasoning fidelity. Furthermore, security considerations are integrated into the design rather than treated as a post hoc evaluation, thereby addressing weaknesses identified in generic RAG deployments.

4.5 Comparison with Domain-Specific RAG Approaches

TechniqueRAG [15] represents a significant advancement by tailoring RAG pipelines for adversarial technique identification using zero shot LLM re-ranking. While TechniqueRAG achieves strong performance with minimal labeled data, its focus is limited to technique identification rather than holistic CTI analysis.

The proposed method generalizes domainspecific RAG beyond a single task by supporting end to end CTI workflows, including ingestion, correlation, reasoning, and summarization. Unlike task specific RAG systems, the proposed architecture integrates multiple CTI sources and reasoning stages, enabling broader analytical

coverage while maintaining domain precision.

4.6 Comparison with TTP Interpretation and Summarization Models

Studies on TTP comprehension show that RAG enhanced decoder only LLMs outperform supervised fine-tuned encoder only models, particularly when labeled data is scarce [24]. These findings reinforce the effectiveness of retrieval based contextualization for understanding complex attack procedures.

Building on this insight, the proposed method extends RAG-based TTP understanding by embedding it within a larger CTI reasoning framework. Rather than treating TTP interpretation as an isolated task, the approach contextualizes TTPs within broader threat narratives and evidence chains, improving interpretability and operational relevance for security analysts.

4.7 Comparison with RAG Security and Robustness Frameworks

Recent work highlights the vulnerability of RAG systems to corpus poisoning and adversarial manipulation [25, 26]. EcoSafeRAG introduces effective retrieval layer defenses, while SafeRAG provides a comprehensive evaluation benchmark. However, these efforts primarily focus on detection and evaluation rather than integrated mitigation within operational CTI pipelines.

The proposed method complements these works by incorporating security awareness directly into the CTI reasoning process. Instead of relying solely on external defenses or benchmarks, the system emphasizes controlled retrieval, contextual consistency, and evidence linked generation. This integrated approach enhances robustness while preserving the analytical benefits of RAG.

4.8 Comparison with End-to-End CTI RAG Architectures

End-to-end RAG architectures augmented with knowledge graphs demonstrate substantial gains in detection accuracy and analyst efficiency (Gandhi et al., 2024). These systems highlight the value of combining structured and unstructured reasoning for CTI.

The proposed method aligns with this direction but further emphasizes trustworthiness and explainability. By explicitly linking generated insights to retrieved evidence, the system supports analyst validation and reduces cognitive load. While existing architectures report strong performance gains, they remain sensitive to corpus quality and privacy concerns; the proposed approach explicitly acknowledges these limitations and motivates future work on secure retrieval and counter adversarial training.

4.9 Summary of Advancements over the State-of-the-Art

Overall, the proposed method advances the state of the art by unifying domain-specific retrieval, grounded LLM reasoning, and security awareness within a single CTI framework. Compared with prior work, it offers broader analytical coverage, greater robustness to hallucinations and retrieval attacks, and stronger alignment with real world CTI workflows. These characteristics position the proposed approach as a practical and trustworthy solution for next generation cyber threat intelligence systems.

5. Methods and Data Sources

5.1 System Architecture and Deployment Pipeline

In practice, the proposed method can be implemented as a modular, end to end CTI analysis pipeline consisting of five key stages:

1. CTI Data Ingestion
2. Semantic Preprocessing and Indexing
3. Secure Retrieval and Context Filtering
4. Grounded LLM-Based Reasoning
5. Analyst-Facing Outputs and Feedback Loop

This modular design enables incremental deployment and integration with existing Security Operations Center (SOC) tools, such as SIEM and SOAR platforms.

5.2 CTI Data Sources and Ingestion

A major strength of the proposed method lies in its ability to ingest heterogeneous CTI data from both structured and unstructured sources. In operational settings, the following data sources are particularly suitable:

5.2.1 Open-Source and Public CTI

The possible data sources for CTI studies are below.

- MITRE ATT&CK: Tactics, techniques, procedures, and procedural descriptions
- CVE / NVD: Vulnerability identifiers, severity scores, and exploit descriptions
- Security advisories: Vendor reports (e.g., Microsoft, Cisco, Palo Alto)
- Threat reports: Industry whitepapers and incident analyses

5.2.2 OSINT and Community Intelligence

- Cybercrime forums (e.g., XSS, RAMP, Exploit_in)
- Security blogs and mailing lists
- Public incident disclosures

5.2.3 Organizational Telemetry (Optional)

- Incident reports
- Alert summaries from IDS/IPS
- Malware analysis outputs

During ingestion, documents are tagged with metadata, including source, timestamp, threat category, and confidence level. This metadata is later used for retrieval filtering and trust aware reasoning.

5.3 Semantic Preprocessing and Indexing

To enable domain aligned retrieval, the ingested CTI data undergoes semantic preprocessing, which includes:

- Document chunking based on semantic coherence (e.g., attack phases, procedures)
- Entity extraction and linking, including:
  - Threat actors
  - Malware families
  - TTP identifiers (MITRE ATT&CK IDs)
  - Vulnerabilities (CVE IDs)
- Embedding generation using cybersecurity adapted encoders

The resulting embeddings are stored in a vector index, optionally augmented with a knowledge graph that represents relationships among actors, techniques, and incidents. This hybrid representation supports both semantic similarity search and relational reasoning.

6. Analysis

For this current study, we used the source generated by the CROWDSTRIKE 2025 THREAT HUNTING REPORT [28]. Page 8-10

Metric	**Value Reporting Period**
Increase in interactive intrusions 27%	Past 12 months

Table 3. Year-over-Year Change in Interactive Intrusions

It is clear that Interactive intrusions continue to rise sharply, indicating sustained growth in adversary activity.

Threat Motivation	**Percentage of Total Intrusions**
eCrime	73.4%
Nation-State	26.5%
Hacktivism	0.1%

Table 4. Distribution of Interactive Intrusions by Motivation (2025)

We found that financially motivated eCrime overwhelmingly dominates interactive intrusions, while hacktivism remains negligible

Sector	Percentage Change
Technology	+2%
Consulting & Professional Services	+17%
Manufacturing	+51%
Retail	+41%
Financial Services	+26%
Telecommunications	+130%
Government	+126%
Industrials & Engineering	+185%
Academic	+80%
Media	+23%
Energy	+55%
Logistics	+42%
Real Estate	+10%
Opportunistic	–12%

Table 5. Top Targeted Sectors by Intrusion Frequency (Relative Change)

Reporting Window**:** July 2023–June 2024 vs. July 2024–June 2025

Sector Category	Observed Trend
Technology	The most targeted sector for 8 consecutive years
Government	Significant increase driven by Russia-nexus activity
Telecommunications	Significant increase driven by China-nexus activity
Opportunistic Attacks	Overall decline (–12%)

Table 6. Sector-Level Intrusion Trend Summary

Aspect	Nation-State Activity	eCrime Activity
Primary Motivation	Strategic / Espionage	Financial gain
Dominant Target Sectors	Government, Telecommunications	Technology, Financial Services
Share of Interactive Intrusions	26.5%	73.4%
Growth Pattern	Sector-specific spikes	Persistent and pervasive

Table 7. Nation-State vs. eCrime Targeting Characteristics

Inferences

- 27% YoY growth confirms escalation in hands on keyboard intrusions
- Nearly 3 out of 4 intrusions are financially motivated
- Government (+126%) and Telecommunications (+130%) show nation state driven surges
- Industrials & Engineering (+185%) emerges as a rapidly escalating target class
- Opportunistic activity decline suggests adversaries are becoming more selective and strategic

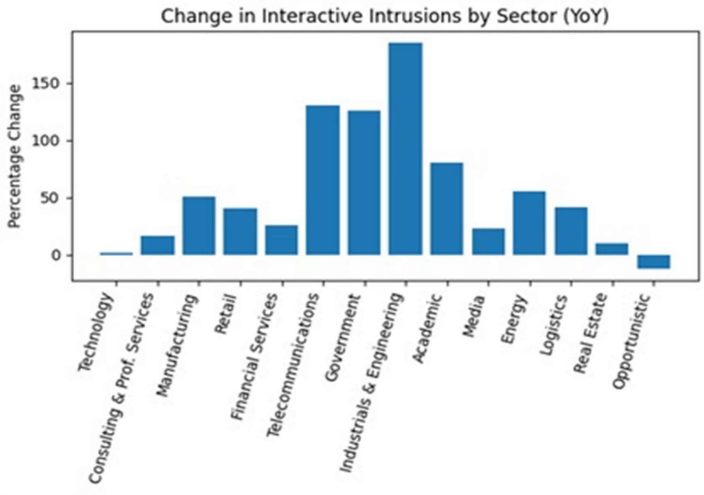

Figure 1. Changes in the Interactive Intrusions

Below, I integrate **(**i) trend charts, (ii) sector-to–MITRE ATT&CK tactic mapping, and (iii) a concrete CTI-RAG evaluation integration, so that we can directly use this material in the work or system description.

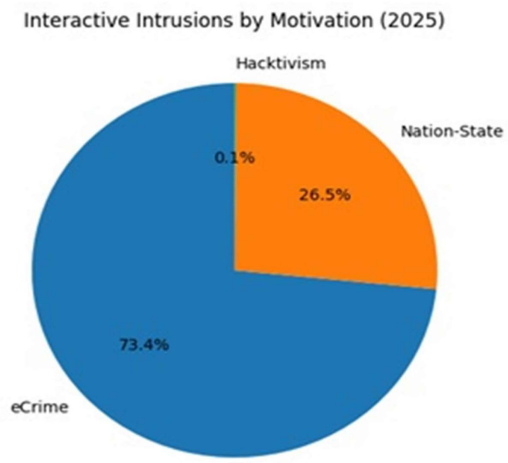

Figure 2. Intrusion Motivations

6.1 Trend Charts

Two trend charts have been generated from the extracted numerical data:

6.1.1 Sector-wise Change in Interactive Intrusions

Figure 1 shows the year over year percentage change in interactive intrusions by sector.

• Accelerated Targeting of Strategic Sectors:

The most pronounced year over year growth in cyber incidents is observed in Industrials & Engineering (+185%), Telecommunications (+130%), and Government (+126%). This sharp upward trend signals a deliberate pivot by threat actors toward sectors that underpin national infrastructure, supply chain integrity, and geopolitical influence. These industries are increasingly viewed as high-value targets due to their roles in critical operations, access to sensitive data, and potential for cascading disruption across interconnected systems. The targeting aligns with both state sponsored espionage campaigns and financially motivated actors seeking long term access or ransom leverage.

6.2 Persistent Pressure on Commercial Verticals:

Manufacturing (+51%), Retail (+41%), and Financial Services (+26%) continue to experience elevated attack volumes, reinforcing the enduring appeal of these sectors to eCrime groups. Their rich repositories of payment data, personally identifiable information (PII), and intellectual property make them lucrative targets. Additionally, the digital transformation of retail and manufacturing through IoT adoption, cloud migration, and just in time logistics has expanded the attack surface, enabling more sophisticated and scalable criminal operations such as Magecart skimming, business email compromise (BEC), and ransomware as a service (RaaS) deployments.

6.3 Decline in Opportunistic Activity:

A notable 12% decline in opportunistic attacks such as broad spectrum phishing, untargeted malware distribution, or exploitation of widely known vulnerabilities without reconnaissance suggests a strategic recalibration among adversaries. Rather than casting wide nets, attackers are investing greater resources in reconnaissance, custom tooling, and supply chain compromises to infiltrate mission critical environments with higher payoff potential. This shift reflects maturation in adversary tradecraft and an increasing focus on operational security, persistence, and impact over volume.

6.4 Cyber Threat Intelligence (CTI) Implication:

These trends underscore the inadequacy of one size fits all threat models. Organizations can no longer rely solely on generic indicators of compromise (IOCs) or broad TTPs (Tactics, Techniques, and Procedures) drawn from cross sector averages. Instead, sector-specific threat modeling informed by vertical aligned intelligence, regulatory landscapes, and ecosystem interdependencies is now essential. CTI programs must evolve to incorporate industry tailored adversary profiles, scenario based playbooks, and proactive monitoring of third-party risk vectors unique to each sector's operational environment. This precision enables more effective detection, response, and resilience planning aligned with actual threat actor behavior.

6.4.1 Interactive Intrusions by Motivation (2025)

Figure 2 provides a clear quantitative breakdown of the primary motivations driving cyber intrusions observed over the reporting period:

• eCrime Dominance (73.4%):

• The overwhelming majority of detected intrusions are attributed to financially motivated cybercriminal activity. This includes ransomware deployments, data exfiltration for sale on dark web markets, credential harvesting, business email compromise (BEC), and point of sale (POS) skimming. These operations are typically executed by organized eCrime syndicates leveraging modular malware, phishing as a service platforms, and affiliate driven ransomware models (e.g., Ransomware as a Service). Their tactics prioritize speed, automation, and monetization often exploiting known vulnerabilities, weak identity hygiene, or misconfigured cloud assets with minimal customization.

• Significant Nation State Footprint (26.5%):

• Despite being outnumbered by criminal actors, nation state or state aligned threat groups represent more than a quarter of all intrusion activity a proportion far exceeding their typical visibility in public reporting. These operations are generally intelligence driven, focused on long term espionage, intellectual property theft, pre positioning for potential disruptive actions, or influence operations. Targets often include defense contractors, critical infrastructure operators, government agencies, and high tech research institutions. Unlike eCrime actors, nation state campaigns emphasize stealth, operational security, and custom tooling, frequently employing zero day exploits, supply chain compromises, or living off the land techniques (LOTL) to evade detection.

• Negligible Hacktivist Activity (0.1%):

• Hacktivism motivated by ideological or political causes accounts for a statistically insignificant fraction of intrusions. While such campaigns can generate high profile disruptions (e.g., website defacements or DDoS attacks), they rarely involve deep network penetration or sustained access. Their low prevalence suggests that modern hacktivist efforts have either diminished in technical sophistication or shifted toward lower risk, high visibility tactics outside traditional intrusion detection scopes.

6.5 Cyber Threat Intelligence (CTI) Implication:

The dual layered threat landscape revealed by Figure B demands a tiered and context aware CTI strategy:

1. Operational Priority on eCrime TTPs:
Given that nearly three quarters of incidents stem from financially motivated actors, defensive programs must prioritize detection and mitigation of common eCrime tactics such as phishing lures, lateral movement via compromised credentials, Cobalt Strike beaconing, and data staging for exfiltration. Automated threat feeds, behavioral analytics, and rapid patching of internet facing systems remain foundational.

2. Strategic Integration of Geopolitical Context:

3. However, the substantial 26.5% nation state presence necessitates more than just tactical defenses. Organizations especially those in strategically sensitive sectors must embed geopolitical awareness into their threat modeling. This includes tracking adversary alignment with national interests (e.g., APT29 with Russia, APT41 with China), monitoring regional tensions that may trigger cyber escalation, and correlating internal anomalies with global incident timelines. Strategic CTI should inform executive level risk decisions, third-party due diligence, and incident response playbooks tailored to advanced persistent threats (APTs).

In essence, while volume driven defenses address the bulk of daily threats, context enriched intelligence is critical to identifying and neutralizing the minority but high impact nation state campaigns that could inflict strategic or existential damage. A resilient security posture must therefore balance both layers: automating responses to commodity crime while maintaining human led analysis for sophisticated, intent driven intrusions.

6.6 Mapping Sector Statistics to MITRE ATT&CK Tactics

The following table maps high growth sectors from the report to dominant MITRE ATT&CK tactics, based on observed intrusion motivations and sector characteristics.

Sector	**Growth Indicator**	**Dominant Adversary Type**	**Likely MITRE ATT&CK Tactics**
Technology	Most targeted (8th year)	eCrime + Nation-State	Initial Access, Credential Access, Persistence, Lateral Movement
Government	+126%	Nation-State (Russia-nexus)	Reconnaissance, Initial Access, Defense Evasion, Command and Control
Telecommunications	+130%	Nation-State (China-nexus)	Initial Access, Discovery, Collection, Command and Control
Industrials & Engineering	+185%	Nation-State + Advanced eCrime	Initial Access, Privilege Escalation, Lateral Movement, Impact
Financial Services	+26%	eCrime	Credential Access, Privilege Escalation, Exfiltration
Manufacturing	+51%	eCrime + Espionage	Initial Access, Lateral Movement, Impact

Retail	+41%	eCrime	Initial Access, Credential Access, Exfiltration
Energy	+55%	Nation-State	Reconnaissance, Persistence, Impact
Academic	+80%	Nation-State	Reconnaissance, Collection, Credential Access

Table 8 Sector-to–MITRE ATT&CK Tactic Mapping

This mapping matters because ot enables sector-aware ATT&CK filtering inside the RAG retrieval stage, reduces noise by prioritizing tactics statistically aligned with sector risk and provides explainability for analysts by linking sector trends to attacker behavior.

6.7 Integration into CTI RAG Evaluation

The extracted tables and trend analyses can be directly integrated into the CTI RAG evaluation framework in three concrete ways.

6.7.1 Retrieval Evaluation (Context Relevance)

We can use the data when a query targets a specific sector (e.g., *telecommunications breach*), the RAG retriever is constrained to:

- CTI documents associated with that sector
- ATT&CK tactics mapped in Table X

The evaluation metric contains

- Context Precision@k: proportion of retrieved documents matching sector-aligned tactics
- Reduction in irrelevant context (%) compared to generic retrieval

6.8 Reasoning Evaluation (TTP Accuracy)

Ground truth construction

- Use the sector growth + motivation statistics as weak supervision signals.
- Example:
- Telecommunications → Nation-state → C2, Discovery
- Financial Services → eCrime → Credential Access, Exfiltration

6.9 Evaluation metrics

- TTP identification accuracy
- Correct tactic ranking (Top-1 / Top-3)
- Hallucination rate (claims not supported by retrieved sector evidence)

6.10 Analyst-Centric Impact Evaluation

The following table shows how your numerical findings support evaluation claims.

Integration of CrowdStrike Statistics into CTI RAG Evaluation

Evaluation Aspect	**Baseline (LLM / Generic RAG)**	**Proposed CTI RAG**
Sector awareness	Implicit or absent	Explicit (trend-driven)
ATT&CK alignment	Generic	Sector-weighted
Hallucination risk	High	Reduced via evidence grounding
Analyst triage time	High	Reduced using ranked hypotheses
Explainability	Limited	Evidence-linked to sector & tactics

The table compares a conventional (baseline) approach (using standard LLMs or generic Retrieval Augmented Generation systems) with a proposed, domain specific CTI RAG system that leverages real world threat intelligence specifically, the 2025 CrowdStrike intrusion statistics.

In summary, the proposed CTI RAG system transforms raw threat statistics into an operational advantage: it makes AI-driven threat analysis more accurate, relevant, trustworthy, and actionable by embedding real-world empirical trends directly into the retrieval and reasoning pipeline.

6. Conclusion

In conclusion, the analysis of current cyber threat trends particularly the 27% year over year rise in interactive intrusions and the stark sectoral targeting shifts reveals a maturing, increasingly strategic adversary landscape. The dominance of eCrime (73.4%) underscores the persistent financial incentives driving ransomware, data theft, and credential based attacks across commercial sectors like Manufacturing, Retail, and Financial Services. Simultaneously, the significant nation state footprint (26.5%), especially in Government, Telecommunications, and Industrials & Engineering, reflects a parallel track of geopolitically motivated espionage and infrastructure targeting. The 12% decline in opportunistic attacks further signals that adversaries are prioritizing high value, mission critical environments over low effort exploits.

These dynamics necessitate a fundamental evolution in Cyber Threat Intelligence (CTI) practices. Generic, one size fits all threat models are no longer sufficient. Instead, organizations must adopt sector specific, context aware CTI frameworks that integrate both tactical eCrime patterns and strategic nation state behaviors. This includes aligning detection strategies with MITRE ATT&CK tactics most relevant to each industry's risk profile and embedding geopolitical awareness into threat assessment workflows.

To operationalize this precision intelligence, the proposed domain specific, security aware Retrieval Augmented Generation (RAG) architecture offers a promising path forward. By grounding LLM reasoning in verified, sector aligned CTI sources and filtering retrieval through adversarial behavior mappings the system reduces hallucination risks, improves TTP accuracy, and delivers explainable, analyst ready insights. Unlike prior approaches limited to isolated tasks or vulnerable to retrieval poisoning, this end to end pipeline supports real world CTI workflows with built in robustness and minimal supervision.

Ultimately, as cyber threats grow more targeted and sophisticated, defensive strategies must become equally nuanced. The convergence of empirical threat data, structured knowledge (e.g., ATT&CK), and secure AI-driven reasoning enables a proactive, adaptive defense posture one capable of addressing both the volume of criminal activity and the stealth of state sponsored campaigns. This integrated approach positions organisations not only to react but also to anticipate and neutralise emerging threats with confidence.

References

[1] Alhuzali, A. (2025). LLM-powered threat intelligence: a retrieval-augmented generation approach for cyber attack investigation. *PeerJ Computer Science* 11:e3371 https://doi.org/10.7717/peerj-cs.3371.

[2] Liao, X., Yuan, K., Wang, X, F., Xing, Z., Li, Beyah., R. (2016). Acing the ioc game: Toward automatic discovery and analysis of open source cyber threat intelligence, *In: Proceedings of the 2016 ACM SIGSAC conference on computer and communications. NY. ACM.*

[3] Gao, P., Liu, X., Choi, E., Ma, S., Yang, X., Ji, Z., Zhang, Z., Song, D. (2022). Threatkg: A threat knowledge graph for automated open source cyber threat intelligence gathering and management. arXiv preprint arXiv:2212.10388, 2022•arxiv.org.

[4] Arikkat, D. R., Abhinav, M., Binu, N., Parvathi, M., Biju, N., Arunima, K. S., Vinod, P., Rrka, M., Conti. (2024). IntellBot: Retrieval Augmented LLM Chatbot for Cyber Threat Knowledge Delivery, *IEEE 16th International Conference on Computational* •ieeexplore.ieee.org.

[5] Tang, Fengxiao., Huan, Li., Zhao, Ming., zong, Zong.,Wu. (2025). LRCTI: A Large Language Model Based Framework for Multi-Step Evidence Retrieval and Reasoning in Cyber Threat Intelligence Credibility Verification. arXiv:2507.11310v1 [cs.CR] 15 Jul.

[6] Santos, P., Abreu, R., Reis, M. J. C. S., Serôdio, C., Branco, F. (2025). A Systematic Review of Cyber Threat Intelligence: The Effectiveness of Technologies, Strategies, and Collaborations in Combating Modern Threats. *Sen sors* 25, 4272.

[7] Rafael, Pedro Santos., Manuel, Abreu Rafael Abreu., Reis, J. C. S., Branco, Frederico Branco Frederico. (2025). A Systematic Review of Cyber Threat Intelligence: The Effectiveness of Technologies, Strategies, and Collaborations in Combating Modern Threats, *Sensors* 25 (14) 4272.

[8] Mc Call, Grover, C., Jr. (2022). Exploring a Cyber Threat Intelligence (CTI) Approach in the Thwarting of Adversary Attacks: An Exploratory Case Study, Northcentral University ProQuest Dissertations Theses, 289-68146.

[9] Bassog, Louise. (2024). Identifying the Problems With the Data Analysis Stage of Cyber Threat Intelligence (CTI) Process and Possible Solutions, *Bowie State University Pro Quest Dissertations Theses*, 31148634.

[10] Clairoux Trepanier, Vanessa., Beauchamp, Isa-May., Ruellan, Estelle., Paquet Clouston, Masarah., Serge-Olivier Paquette., Clay, Eric. (2024). The Use of Large Language Models (LLM) for Cyber Threat Intelligence (CTI) in Cybercrime Forums}.

[11] Mon Divakaran, Dinil., Peddinti, Sai Teja. (2025). LLMs for Cyber Security: New Opportunities, https://doi.org/10.48550/arXiv.2404.11338.

[12] Martino, A., Iannelli, M., Truong, C. (2023). Knowledge injection to counter large language model (LLM) hallucination, *In:* European Semantic Web Conference, P*Springer*.

[13] Perkoviæ, G., Drobnjak, A., Botièki, I. (2025). Hallucinations in llms: Understanding and addressing challenges, *In:* 2024 47th MIPRO ICT and electronics convention (MIPRO), *ieeexplore*.

[14] An, Bang., Zhang, Shiyue., Dredze, Mark. (2025). RAG LLMs are Not Safer: A Safety Analysis of Retrieval-Augmented Generation for Large Language Models, https://doi.org/10.48550/arXiv.2504.18041.

[15] Lekssays, Ahmed., Shukla, Utsav., Sencar, Husrev Taha., Parvez, Md, Rizwan. (2025): Retrieval Augmented Generation for Adversarial Technique Annotation in Cyber Threat Intelligence Text. https://doi.org/10.48550/arXiv.2505.11988.

[16] Chris, M., Ward, Harguess, Josh., Harguess, Josh. (2025). Adversarial Threat Vectors and Risk Mitigation for *Retrieval Augmented Generation Systems*. May. 10.48550/arXiv.2506.00281.

[17] Luo, H., Luo, X., Zhao, W., Peng, Q., Chen, K., Liu, Y., Du, C. (2026). Domain Specific Retrieval-Augmented Generation with Adaptive Embedding and Knowledge Distillation Based Re Ranking. *Processes*.

[18] Barron, Ryan Calvin., Barron, Ryan Calvin., Grantcharov, Vesselin., Grantcharov, Vesselin., Wanna, Selma., B. S., Alexandrov, B. S., Alexandrov. (2024). Domain Specific Retrieval Augmented Generation Using Vector Stores, Knowledge Graphs, and Tensor Factorization. *October*. DOI: 10.48550/arXiv.2410.02721.

[19] Shao, Z., Gong, Y., Shen, Y., Huang, M., Duan, N., Chen, W. (2023). Enhancing Retrieval Augmented Large Language Models with iterative Retrieval Generation Synergy. *In:* Proceedings of the Findings of the Association for *Computational Linguistics, Singapore,* 6–10 December p. 9248–9274.

[20] Jiang, Z., Xu, F., Gao, L., Sun, Z., Liu, Q., Dwivedi-Yu, J., Yang, Y., Callan, J., Neubig, G. (2023). Active Retrieval Augmented Generation. *In:* Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, *Singapore*, 6–10 December. p. 7969–7992.

[21] Trivedi, H., Balasubramanian, N., Khot, T., Sabharwal, A. (2023). Interleaving Retrieval with Chain-of-Thought Reasoning for Knowledge Intensive Multi-Step Questions. *In:* Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics, *Toronto, ON, Canada*, 9–14 July 2023; p. 10014–10037.

[22] Li, Y., Dong, B., Guerin, F., Lin, C. (2023). Compressing Context to Enhance Inference Efficiency of Large Language Models. *In:* Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, *Singapore*, 6–10 December p. 6342–6353.

[23] Kim, J., Nam, J., Mo, S., Park, J., Lee, S. W., Seo, M., Ha, J. W., Shin, J. (2024). SuRe: Summarizing Retrievals using Answer Candidates for Open domain *QA of LLMs*. arXiv.

[24] Fayyazi, R., Taghdimi, R., Yang, S. J. (2024). Advancing TTP Analysis: Harnessing the Power of Large Language Models with Retrieval Augmented Generation, 2024 Annual Computer Security Applications Conference Workshops (ACSAC Workshops), *Honolulu, HI, USA,* p. 255-261.

[25] Yao, Ruobing., Zhang, Yifei., Song, Shuang., Gao, Neng., Chenyang, Tu. (2025). EcoSafeRAG: Efficient Security through Context Analysis in Retrieval Augmented Generation, https://doi.org/10.48550/arXiv.2505.13506.

[26] Liang, Xun Song. (2025). S}afe{RAG}: Benchmarking Security in Retrieval Augmented Generation of Large Language Model, *In:* Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), *Association for Computational Linguistics"*,p. 4609—4631.

[27] Gandhi, Sahaj Tushar. (2023). RAG Driven Cybersecurity Intelligence: Leveraging Semantic Search for Improved Threat Detection. *International Journal of Research and Applied Innovations,* 6(3), 8889-8897.

[28] Crowdstrike. (2025). Threat Hunting Report. page 8-10.www.crowdstrike.com.