BLACK KITE

Large Language Model Use Cases in Cyber Threat Intelligence

Enhancing Third-Party Risk Management Through AI

Authors: Müzeyyen Gökçen Tapkan, Gizem Toprak

LLMs in CTI: A New Era for Third-Party Risk Management

The security of your organization is only as strong as its weakest link – and increasingly, those links are external. While internal cybersecurity measures are essential, they offer limited visibility into the security postures of third-party vendors and partners. This report is specifically designed to bridge that gap, empowering Third-Party Risk Management (TPRM) professionals to harness the vast potential of open-source intelligence (OSINT).

Today's challenge is that while Cyber Threat Intelligence (CTI) is crucial for identifying potential risks, traditional methods struggle to efficiently process the vast amounts of unstructured data relevant to third-party security—before adversaries can exploit vulnerabilities.

This is where Large Language Models (LLMs) emerge as a game-changer. This report will demystify the application of LLMs in CTI, demonstrating how these powerful tools can be practically applied to strengthen your third-party risk management program.

We will show how LLMs can enable you to:

- **Streamline OSINT Analysis**
- **Enhance Supply Chain Visibility**
- **Improve Risk Prioritization**
- **Facilitate Informed Decision-Making**

Ultimately, by transforming the often overwhelming volume of publicly available data into clear, actionable intelligence, LLMs empower you to move from reactive threat response to proactive risk mitigation.

What are LLM and CTI?

LLMs are trained on extensive datasets that allow them to understand, generate, and summarize human-like text with remarkable accuracy. By leveraging advanced deep learning techniques, LLMs can sift through vast amounts of information to extract meaningful insights and patterns that might otherwise go unnoticed. Their ability to interpret detailed language, context, and emerging trends makes them particularly valuable in cybersecurity, where timely and precise intelligence can be the difference between thwarting an attack and suffering a breach.

At its core, CTI is about transforming raw data into actionable intelligence. It involves a structured process of gathering information from multiple sources, analyzing that data for indicators of compromise, and disseminating the findings to relevant stakeholders. CTI not only helps with real-time threat detection, but also informs strategic decision-making and proactive defensive measures.

Integrating LLMs into CTI workflows promises to automate repetitive analytical tasks, enhance data correlation, and ultimately improve the speed and accuracy of threat assessments.

The connection between LLMs and CTIs is poised to redefine the cybersecurity landscape.

By leveraging the advanced text processing capabilities of LLMs, CTI systems can accelerate the discovery of critical threat patterns, automate the creation of intelligence reports, and even predict emerging attack vectors. This integration not only increases the efficiency of threat analysis, but also provides security teams with a powerful tool to stay ahead of sophisticated adversaries in an ever-evolving digital battlefield.

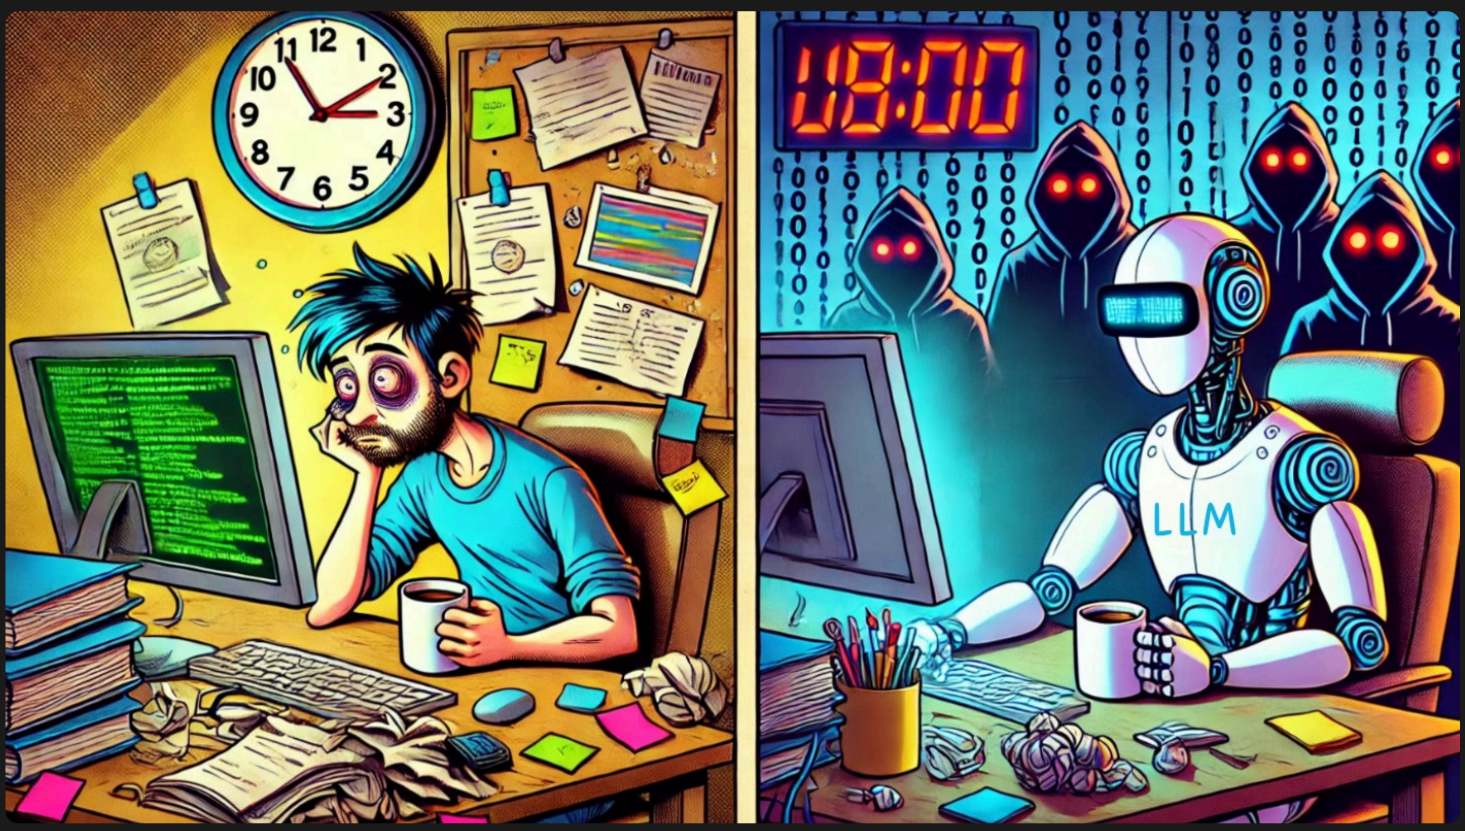

CTI workflows involve repetitive analytical tasks. LLMs improve speed and accuracy. (Photo created with ChatGPT-4.0)

Why Supply-Chain CTI is Critical Today

Supply chains today are no longer simple linear systems. Instead, they form complex networks involving numerous vendors and service providers. Each entity represents a potential entry point for attackers, creating a broad attack surface. Incidents such as the recent vulnerabilities uncovered in DrayTek Vigor routers, VMware ESXi, Apache Tomcat, and Axios HTTP Client illustrate how a single vendor's weakness can lead to widespread disruption, exposing sensitive data and damaging countless organizations' reputations.

At the same time, as regulations become stricter, companies are being held accountable for the security of their entire supply chain. This increases the need for Cyber Threat Intelligence (CTI) programs. As data volumes grow and threats evolve more rapidly, traditional security methods are falling short. Advanced technologies like AI and machine learning are now vital to processing large amounts of information and providing real-time, actionable intelligence.

Today, Supply chain CTI isn't just about preventing attacks – it's about building resilience. Organizations that invest in strong CTI practices can protect themselves, maintain trust, ensure continuity, and develop more secure partnerships throughout the supply chain.

ACCURACY IN SUPPLY CHAIN

In supply chain management, traditional CTI relied on simple keyword and company name searches to identify threats in dark web data. This approach often generated high false positive rates, mistakenly flagging harmless data as threats and potentially misallocating resources.

Integrating LLMs into the CTI process radically improved performance. By leveraging LLM's advanced text analysis capabilities, the system was able to accurately distinguish between real cyber threats and innocuous phrases.

As a result:

- False positive rates dropped from 60% to 20%
- Accuracy increased from 40% to 88%

This striking development demonstrated how LLMs can significantly reduce noise and provide more reliable intelligence, allowing CTI analysts to derive actionable insights without the burden of sifting through countless false alarms.

Accuracy in Supply Chain:

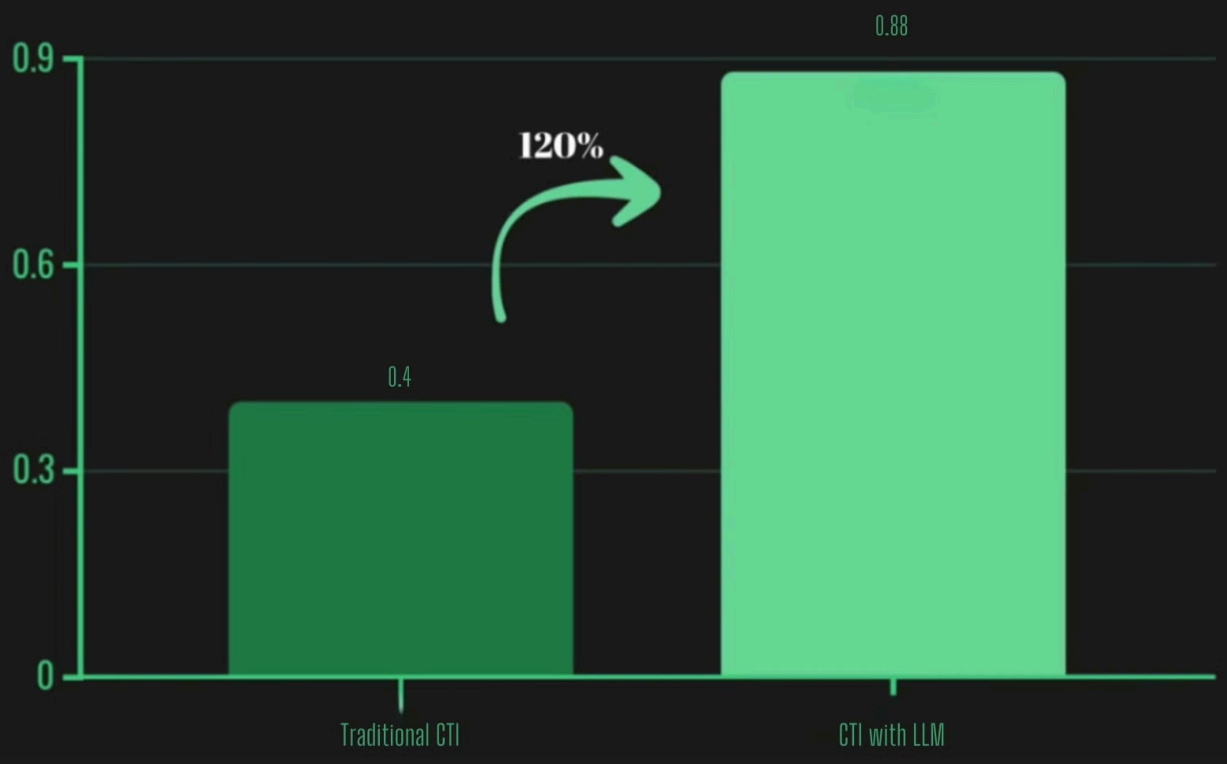

False Positive Rate in Supply Chain:

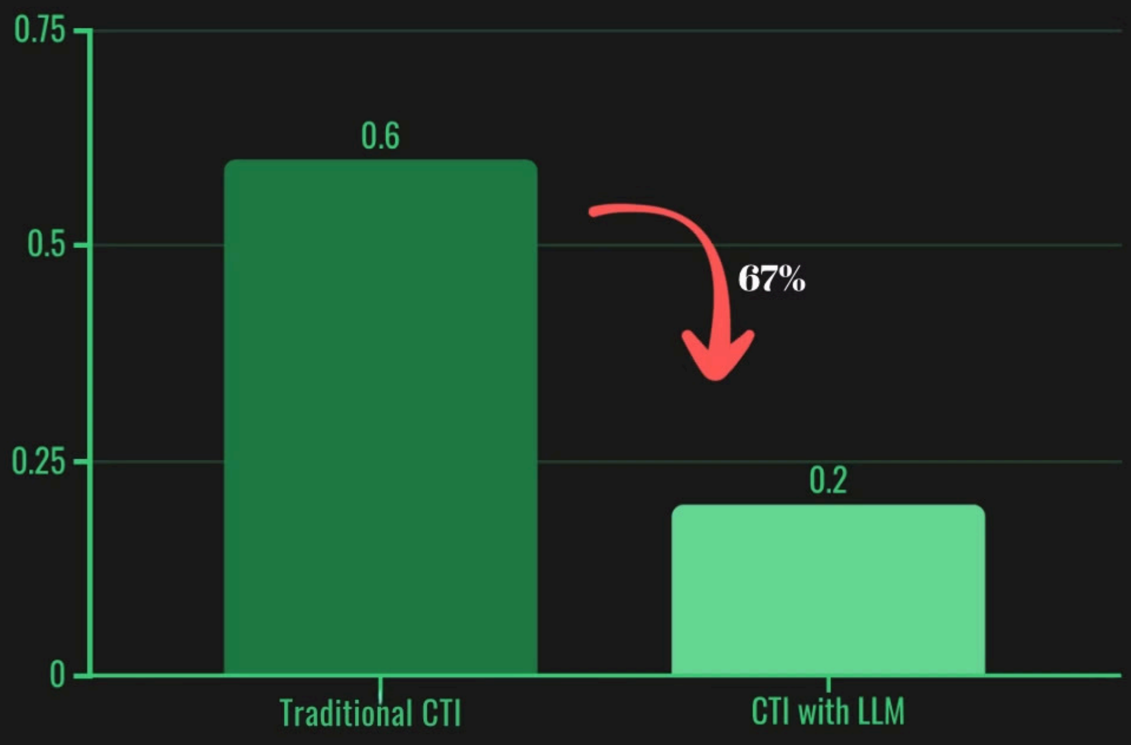

It is usually impossible to sort out large and complex intelligence data and make an inference from it with traditional CTI techniques, both in terms of human power and cost. It should not be forgotten that these texts are difficult to understand even for a cyber threat analyst. Thanks to LLM, very large and complex data was made meaningful and inferred. Companies had very little confidence in the data coming from here because with classical word searches, it was not possible to reach correct information from a platform where cyber threat keywords were already frequently used. However, thanks to LLM techniques, meaningful and accurate results could be extracted from this complex data. Not only that, the attack method and risk level could also be learned.

DETECTION OF 4 CRITICAL COMPONENTS

LLM can help companies to identify the critical components in a CTI task , the threat actors, their targets, methodology and tools.

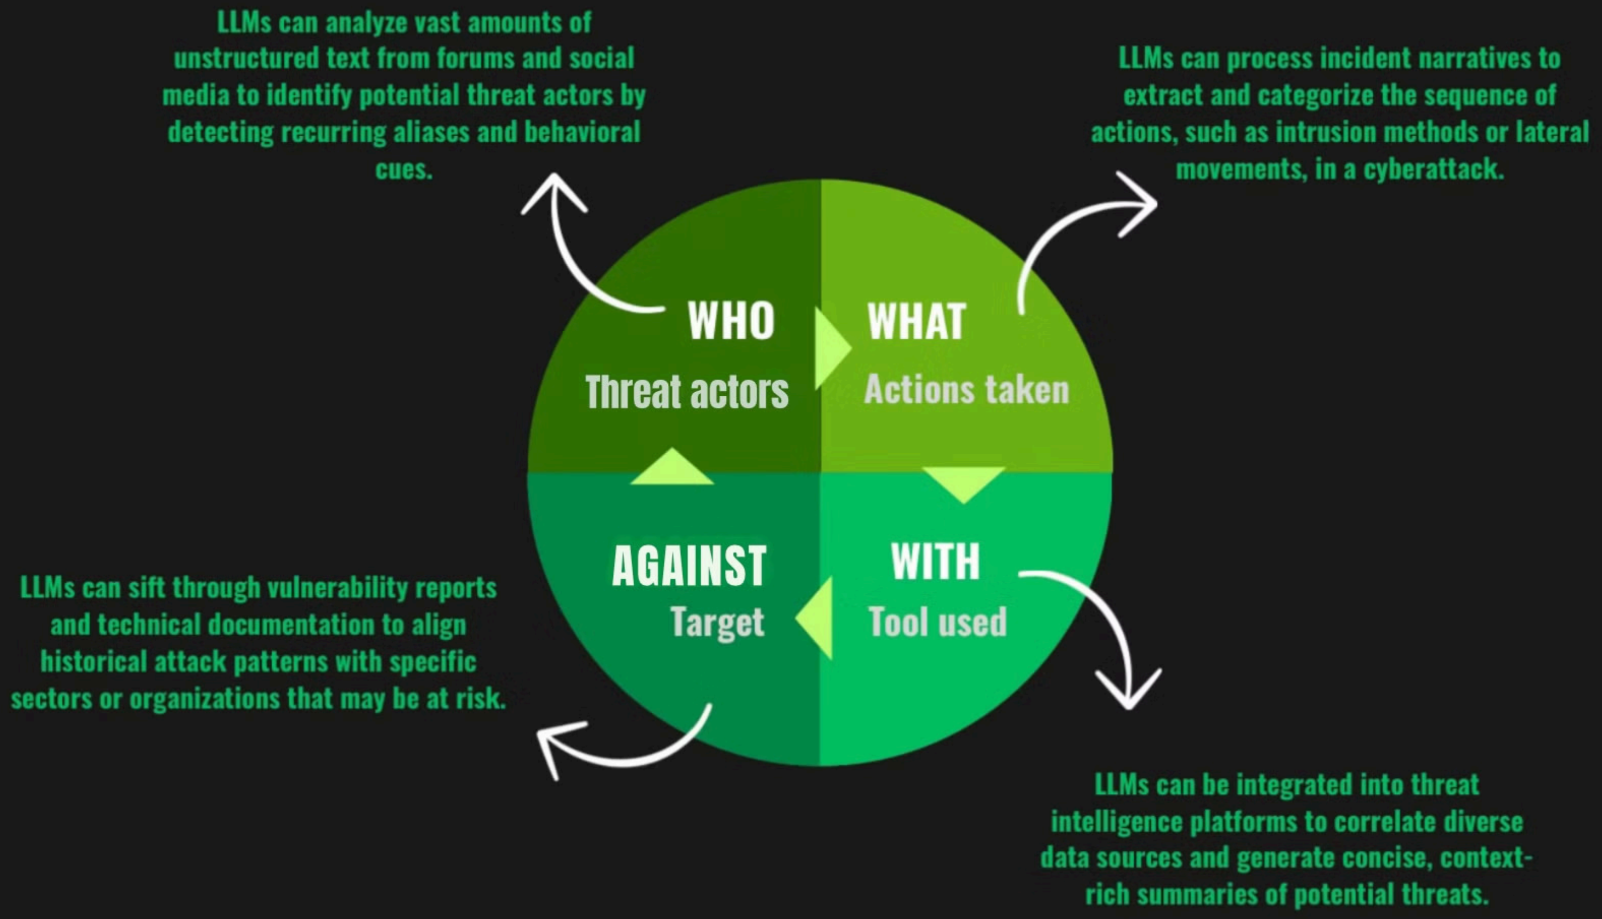

LLMs can analyze vast amounts of unstructured text from forums and social media to identify potential threat actors by detecting recurring aliases and behavioral cues.

LLMs can process incident narratives to extract and categorize the sequence of actions, such as intrusion methods or lateral movements, in a cyberattack.

LLMs can sift through vulnerability reports and technical documentation to align historical attack patterns with specific sectors or organizations that may be at risk.

LLMs can be integrated into threat intelligence platforms to correlate diverse data sources and generate concise, context-rich summaries of potential threats.

BLACK KITE

Key Use Cases LLM in CTI

Summarization

Extracting summaries and information from complex incident reports.

NER

Name Entity Recognition (NER): Identifying affected companies, sensitive data, threat actors, etc.

Q&A

Extracting further intelligence from unstructured text, e.g, hacker forums, news.

TTP Tagging

Extracting Tactics, Techniques, and Procedures (TTP).

Graph Relationship Extraction

Extracting the graph of who did what with which tools against whom (the "w" questions).

Risk Level Analysis

Rating the extent of cyber threats.

Patch Confidence

Providing CVE and CPE analysis and detailing.

OSINT Analysis

Quickly identifying emerging threats or shifts in attacker narratives, providing early warnings.

Challenges for LLM

What LLMs excel at:

- Speed
- Pattern recognition
- Summarization

Where LLMs fall short:

- Contextual depth
- Domain-specific nuances

Although LLMs seem like a dizzying development in the world of cyber threats, as with any technology, they have both strengths and weaknesses.

Speed and Scalability

LLMs Excel at processing massive amounts of unstructured data at unprecedented speeds. For example, approximately 1.5 million texts are pulled from the dark web monthly. Manually parsing this data would be nearly impossible, but LLMs can sift through this deluge very quickly to extract relevant information, enabling real-time threat detection.

Advanced Pattern Recognition

Deep Learning Techniques LLMs use deep learning techniques to recognize complex patterns and relationships in text. This means they can distinguish between benign phrases or text that only mentions payment methods, and real threats. For example, while traditional methods may flag a company as affected by a cyberthreat simply because of the presence of a company name, an LLM can distinguish whether the context implies a real cyberthreat.

Effective Summarization

Concise Summaries LLMs can condense long, complex threat reports into concise summaries that highlight key details such as the nature of the threat, the methods used, and the level of risk. This capability not only saves time, but also helps CTI analysts quickly understand the situation without getting lost in irrelevant details.

Limited Contextual Depth

Despite their strengths, LLMs sometimes struggle with deeply contextual or highly granular data. While they are adept at capturing general patterns, the subtleties of specific cyber threat narratives (especially those that require a deep understanding of evolving attack methodologies) may not always be fully captured.

Domain-Specific Nuances

CTI often contains specialized language and terminology that can vary significantly across industries or cyber threat types. Unless specifically fine-tuned on domain-specific datasets, LLMs can miss these nuanced distinctions. This can sometimes lead to misinterpretations where benign activities are confused with malicious behavior or vice versa.

Everyone can tap into their power, but only experts can transform raw data into real security—distinguishing truth from phantom threats!

Examples of How Bad Inputs Lead To Bad Outputs

In CTI, LLMs are only as effective as the data they receive. Poor quality input or incorrect prompts can lead LLMs astray. Not only incorrect prompts, but incorrect configurations can lead to misinterpretation of input and misleading, incorrect outputs that can compromise security efforts. Companies can lose reputation and trust because of this. Here are some detailed examples of how bad input can result in bad outputs:

Misleading Request Leading to Incorrect Inference

Example:
Imagine an analyst is reviewing a dark web post that ambiguously mentions "intrusion" near a company's name. The prompt given to the LLM is:

"Does this post confirm a ransomware attack on Company X?"

Because the prompt is overly specific and assumes a particular type of threat, the LLM may force an interpretation that aligns with ransomware—even though the original text only hints at a generic intrusion. This results in a false alarm, prompting unnecessary escalation.

Over-Creativity Introducing Nonexistent Details

Example:

An analyst sets the temperature parameter too high in the LLM configuration to encourage creative text generation. When processing a simple report that states, "Multiple login attempts detected," the LLM, influenced by the high temperature, might embellish the summary to include an invented narrative like:

A coordinated cyber espionage campaign targeting Company X was underway, involving sophisticated credential stuffing tactics.

This output infuses details and threats not present in the original data, leading to a misleading risk assessment.

Incomplete Extraction of Key Information

Example:
When tasked with summarizing a lengthy technical threat report, the input data provided is cluttered and lacks clear markers for critical details. As a result, the LLM produces a summary such as:

The report discusses some cyber threat activity.

Here, essential specifics—like the methods used, affected systems, and recommended mitigation strategies—are omitted, leaving CTI analysts with an overly generic and unhelpful overview that fails to guide actionable responses.

Extracting the Wrong Companies

Example:

"👉PAYPAL✅XooM✅PaysenD✅WesterN-UniuN✅IBAN👈 | Carding Forum for Professional Carders Menu Home Forums New posts Find Leaks" If the correct prompts are not written by non-specific or expert people for this text, or if incorrect configurations are made, we may get the following results.

Prompt: If a company name is mentioned in texts that may be a cyber threat, such as leak or breach, tag those companies as having experienced a cyber threat.

In this example, we can see how an unspecified, misspelled prompt can mislabel a company like PayPal, which is used only as a payment method. Unfortunately, such examples can lead to bad results, such as companies applying incorrect sanctions by increasing the false positive alert rate and causing reputation losses.

As we can see in these examples, although LLMs are very useful in CTI, it is very important to use them carefully. Although correct use can produce positive results for companies, excessive reliance on LLMs or their incorrect use can lead to material losses and loss of trust. It should not be forgotten that although the classic old method CTI tactics cannot examine all the data and catch every threat, if LLMs are not used correctly, they can also find threats that do not actually exist.

The Role of Chain of Thought and LLM Chains in CTI

Chain-of-thought prompting involves breaking down complex tasks into sequential, manageable reasoning steps. This technique is especially crucial in Cyber Threat Intelligence (CTI) projects, where the accuracy of threat detection and analysis can have significant consequences.

How It Works:
Instead of asking an LLM to provide a single, comprehensive answer, chain-of-thought prompting instructs the model to articulate intermediate reasoning steps. For instance, when analyzing a suspicious dark web post, an LLM can be guided to:

1 **Extract Key Indicators**

The data used in CTI can be confusing and disorganized, so first identify the keywords, IP addresses, or suspicious behaviors in the text.

2 **Analyze Context**

Evaluate the context surrounding these indicators to determine if they hint at an actual threat.

3 **Infer Relationships**

Connect disparate pieces of information to hypothesize about potential threat actors or attack vectors.

4 **Summarize Findings**

Provide a concise summary with actionable insights.

5 **Reasoning**

Asking for reasoning can enable the LLM to provide more accurate results by explaining why a decision was made on a critical issue such as cybersecurity.

By decomposing the task, each step becomes verifiable, allowing analysts to trace how a conclusion was reached. This transparency not only builds trust in the automated process but also makes it easier to adjust and refine prompts.

Benefits in CTI:

Improved Accuracy: Breaking tasks into smaller reasoning steps helps prevent the model from making overly broad assumptions. This method reduces the risk of misclassifications, ensuring that only genuinely relevant threats are flagged.

Enhanced Interpretability: Each step of the reasoning process is documented, offering clear insights into how conclusions were derived. This is particularly valuable when an analyst needs to understand or challenge an automated decision.

Error Isolation: If the final output is incorrect or misleading, the chain-of-thought approach allows pinpointing which step went awry. This facilitates targeted improvements in prompt design and model configuration.

Real-World Example On CoT

Consider a CTI system monitoring dark web chatter. Instead of a single prompt that asks, "What threats are mentioned in this text?", a chain-of-thought approach might break it down as follows:

1. Indicator Extraction:

"List all IP addresses and suspicious keywords from the post."

2. Contextual Analysis:

"Analyze the surrounding text for any references to specific attack methods or known threat actors."

3. Threat Correlation:

"Determine if the extracted indicators align with known cyber threat patterns."

4. Summary Generation:

"Provide a concise summary of the potential threat, including risk level and possible impact."

This multi-step process yields a more robust analysis than a one-shot prompt, reducing false positives and ensuring that key details are not overlooked.

Challenges and Opportunities for Cyber Security Experts: Upskilling to Adapt to LLM Tools, and Ethical and Accuracy Considerations

As cyber threats become more complex and data volumes increase, cybersecurity professionals are increasingly turning to Large Language Model (LLM) tools to improve operational efficiency in Cyber Threat Intelligence (CTI) and beyond. However, while LLMs offer transformative potential, integrating these tools into security workflows presents both significant challenges and exciting opportunities.

Technical Challenges and Upskilling for Cyber Security Experts:

Mastering Advanced AI Techniques

Cybersecurity professionals are traditionally well-versed in incident response, threat analysis, and vulnerability management. However, LLMs work on complex deep learning and natural language processing frameworks that require a different skillset. To get the most out of LLMs, cybersecurity professionals need to understand Model Architecture, Prompt Engineering, and even Fine-Tuning in some cases.

Data Quality and Integration

LLMs thrive on high-quality, structured data. Cybersecurity professionals must be involved in procedures to clean, normalize, and label large amounts of threat intelligence data, often originating from unstructured text such as dark web posts or social media chatter.

Cost-Effective Scaling

Instead of hiring large teams to manually sift through data or constantly update legacy systems, cyber threat experts can leverage cloud-based LLM solutions that offer continuous updates and improvements, enabling high returns even with limited initial resources.

Opportunities for Cyber Security Experts

Streamlined Threat Analysis

When effectively integrated, LLMs can automate the extraction of critical insights from vast datasets, such as identifying indicators of compromise or extracting entity relationships from unstructured text. This efficiency allows cyber security teams to focus on higher-level strategic decision-making.

Examples:

- Extracting Indicators of Compromise (IOCs)
- Entity Relationship Extraction
- Summarization and TTP Tagging

Proactive Threat Forecasting

Advanced LLM configurations, supported by robust chain-of-thought reasoning and LLM chains, enable predictive analytics in CTI. By analyzing historical trends and real-time data, these systems can forecast emerging threats, allowing organizations to preemptively strengthen their defenses.

- Forecasting Based on Dark Web Trends
- Predicting Ransomware Campaign Trends
- Anticipating Advanced Persistent Threat (APT) Activities

Improved Reporting and Decision Support

Automated report generation and summarization features enable the rapid synthesis of threat intelligence reports. These tools not only reduce the manual workload but also enhance situational awareness by presenting complex threat data in an accessible and actionable format.

- Executive Summary Generation
- Comprehensive Incident Report
- Weekly Threat Digest

Challenges for Cyber Security Experts

Ethical and Accuracy Considerations

Managing Bias and Ensuring Fairness, LLMs can inadvertently perpetuate biases present in their training data. In the context of CTI, this might mean overestimating threats in certain scenarios or misclassifying benign activities as malicious.

Cybersecurity experts must:

- Regularly review AI decisions to identify and reduce bias.
- Use thought chain guidance and transparent reasoning steps to provide interpretable outputs, so automated threat assessments can be validated and corrected by analysts.

Reducing False Positives and False Negatives

The reliability of LLM outputs is paramount in CTI. High false positive rates can lead to alert fatigue and resource wastage, while false negatives can leave organizations vulnerable.

Cybersecurity experts need to focus on:

- Calibration of Model Parameters
- Combining automated assessments with human-in-the-loop reviews to cross-check and verify threat intelligence findings.

Privacy and Data Security Concerns

LLMs often process sensitive data that could include personal information or proprietary details.

Cybersecurity experts must adhere to strict data governance and compliance frameworks:

- Data Anonymization
- Secure Deployment

Open-Source Tools

Apart from these, you can also speed up your CTI processes with LLM by using open-source CTI applications. For example; to integrate these advanced prediction capabilities into your CTI workflows, open-source CTI programs such as OpenCTI, for example, can be used. OpenCTI enables organizations to effectively manage and analyze threat intelligence data. MISP (Malware Information Sharing Platform) is a widely adopted open-source tool that enables sharing, storing and correlating threat intelligence. Its capabilities include automatic report generation and integration with various data sources, making it a complement to LLM-based summarization and reporting workflows.

What's Next: The Future of AI in CTI

All AI systems are converging to an agentic structure, and AI-augmented CTI is no exception. But what is an AI agent? An **AI agent** is an autonomous or semi-autonomous system that perceives its environment, makes decisions, and takes actions to achieve specific goals. AI agents leverage **machine learning (ML), natural language processing (NLP), computer vision, and other AI techniques** to interact intelligently with data, systems, and users.

Autonomy

Operates independently with minimal human intervention and in a timely manner, especially in complex attacks.

Perception

Gathers data from its environment (e.g., logs, user input, network traffic).

Decision-Making

Uses ML algorithms to analyze data and determine the best course of action during an attack.

Action Execution

Performs tasks such as alerting, recommending actions, or automating responses.

Learning & Adaptation

Improves over time by learning from logs and past decisions.

Agentic AI can be used in cybersecurity systems to **detect and respond to threats in real-time**. For example, these agents can identify unusual network behavior and isolate impacted devices autonomously to prevent a compromise without human intervention.

AI Agents in Cybersecurity

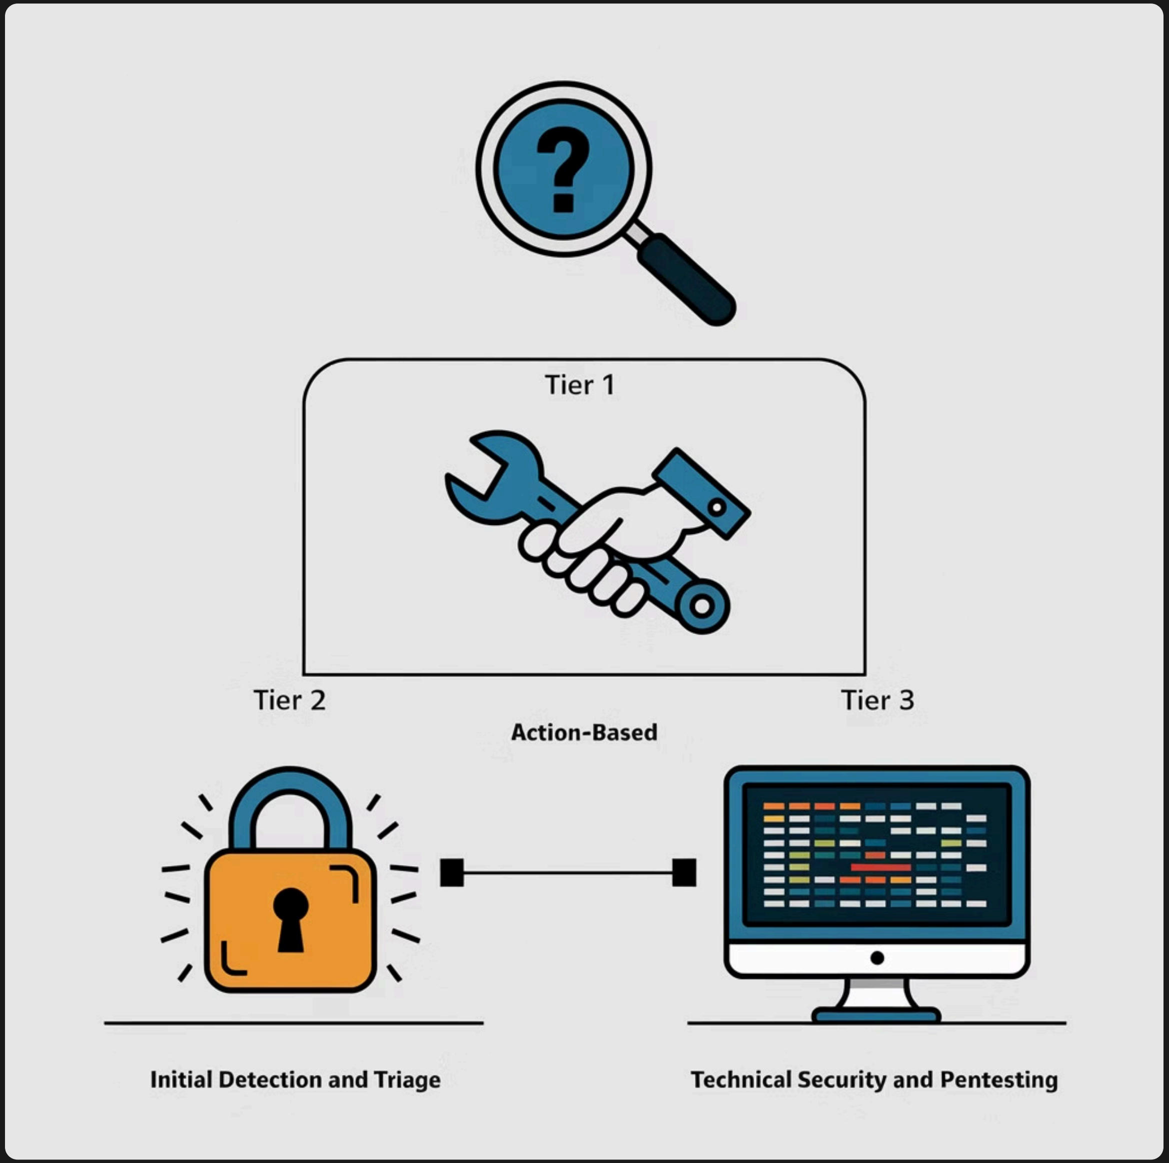
Tier 1
Tier 2
Tier 3
Action-Based
Initial Detection and Triage
Technical Security and Pentesting

- Tier 1 AI agents **aggregate** threat data from various sources
- AI can **de-duplicate** and filter noise to highlight relevant threats
- AI agents can **automate threat triage** and **rank alerts** by severity as well as **detecting false positives**, reducing analyst workload
- AI agents in Tier 2 can isolate affected systems, remove malware, and patch vulnerabilities and compromised data
- AI agents in Tier 3 can do proactive threat hunting, correlation and mapping, and can even write custom scripts for Attacks & Defense

Next Steps: Enhancing Your TPRM with Black Kite

The shift from manual, time-consuming analysis to AI-powered automation is not just a trend; it's the future of effective TPRM. To further explore how this future can become a reality for your organization, we invite you to learn more about Black Kite's capabilities in supply chain threat intelligence. Black Kite leverages AI and LLMs to provide comprehensive insights, including FocusTagsTM with detailed and action-oriented risk intelligence, ransomware susceptibility intelligence throughout your ecosystem, and nth-party supply chain visibility, empowering you to effectively manage and mitigate third-party risks.

Discover how Black Kite's solutions can transform your approach to TPRM and strengthen your supply chain security. Visit our **Supply Chain solution page** to see how we can help you harness the power of AI-driven CTI.

Learn More