

Augmenting threat intelligence: A framework for integrating LLMs, AI Agents, and RAG in cybersecurity analysis

Bhanu Prakash Reddy Mettu *

Independent Researcher, USA.

World Journal of Advanced Engineering Technology and Sciences, 2025, 15(03), 1523-1530

Publication history: Received on 02 May 2025; revised on 14 June 2025; accepted on 16 June 2025

Article DOI: <https://doi.org/10.30574/wjaets.2025.15.3.1032>

Abstract

This comprehensive framework for integrating advanced artificial intelligence technologies into threat intelligence workflows addresses the increasing volume and complexity of cybersecurity data. The strategic deployment of Large Language Models (LLMs), AI agents, and Retrieval-Augmented Generation (RAG) across the threat intelligence lifecycle—from data collection and processing to analysis and dissemination—demonstrates significant potential for automating routine tasks, enhancing analytical capabilities, extracting actionable insights from vast datasets, and improving the timeliness of intelligence reporting. Through detailed examination of implementation strategies and technical considerations, the transformative impact on traditional threat intelligence practices becomes evident while complementing human analyst expertise. The practical methodologies presented enable security teams to leverage generative AI in identifying and responding to emerging threats more effectively.

Keywords: Threat intelligence; Large language models; AI agents; Retrieval-augmented generation; Cybersecurity automation

1. Introduction

1.1. Current Challenges in Threat Intelligence

The landscape of cybersecurity threat intelligence faces unprecedented challenges due to the exponential growth in data volume and complexity. Modern security teams must process an overwhelming amount of information from diverse sources including threat reports, data leaks, network logs, and social media discourse. This information overload makes it increasingly difficult to identify and respond to genuine threats in a timely manner. The challenge is further compounded by the sophistication of modern threats that operate across multiple vectors and employ advanced evasion techniques. Previous research has highlighted how these challenges significantly impact the effectiveness of security operations, creating bottlenecks in analysis and response capabilities [1].

1.2. Manual Nature of Threat Analysis Workflows

Traditional threat intelligence workflows remain largely manual and resource-intensive processes. Analysts typically spend considerable time sifting through extensive reports, correlating information across disparate sources, and attempting to extract actionable insights. These manual approaches are not only time-consuming but also prone to human error and cognitive biases. The reliance on manual processes creates significant operational inefficiencies and can lead to missed threats or delayed responses to critical security incidents. Furthermore, the technical complexity of modern threats often requires specialized expertise that may not be readily available within security teams [2].

* Corresponding author: Bhanu Prakash Reddy Mettu

1.3. Introduction to AI Technologies for Threat Intelligence

Recent advancements in artificial intelligence technologies offer promising solutions to address these challenges. Large Language Models (LLMs) demonstrate remarkable capabilities in understanding context, identifying key elements within complex texts, and generating coherent analyses. AI agents built upon these models can analyze situations, reason through complex scenarios, plan responses, and take actions using various tools. Retrieval-Augmented Generation (RAG) models combine powerful retrieval mechanisms with generative capabilities to provide context-relevant responses based on external data sources. These technologies present opportunities to augment human analysts and streamline threat intelligence workflows through automation of routine tasks and enhancement of analytical capabilities [2].

1.4. Thesis and Scope

This paper proposes that the strategic integration of LLMs, AI agents, and RAG models can fundamentally transform the threat intelligence lifecycle by addressing key challenges identified in recent literature. By leveraging these technologies across the processes of data collection, processing, analysis, and dissemination, security teams can significantly enhance their capacity to identify, understand, and respond to emerging threats. The remainder of this article explores the current state of threat intelligence practices, examines the specific capabilities of these AI technologies in security contexts, details their application across the threat intelligence lifecycle, provides practical implementation guidance, and concludes with recommendations for security practitioners.

2. The Evolving Landscape of Threat Intelligence

2.1. Current State of Threat Intelligence Practices

The threat intelligence landscape has undergone significant evolution in recent years, transitioning from isolated security practices to becoming an integral component of comprehensive cybersecurity frameworks. Contemporary threat intelligence operations encompass a diverse range of activities, including the monitoring of external threats, analysis of attack patterns, and the development of proactive defense strategies. Organizations increasingly rely on multiple intelligence sources, including commercial feeds, open-source intelligence, dark web monitoring, and industry sharing communities. Despite these advancements, the field continues to face fundamental challenges in effectively transforming raw data into actionable intelligence that can drive security decisions. Research indicates that many organizations struggle to fully operationalize threat intelligence despite recognizing its strategic importance [3].

2.2. Challenges Faced by Threat Analysts

Threat intelligence analysts confront numerous challenges that impede their effectiveness. The phenomenon of data overload represents perhaps the most significant obstacle, as analysts must process an ever-expanding volume of information from disparate sources. This challenge is compounded by strict time constraints, particularly when responding to emerging threats that require immediate attention. The combination of information overload and time pressure frequently results in missed threats or delayed identification of critical security issues. Additionally, analysts often struggle with the problem of signal-to-noise ratio, where genuinely important threats are obscured by vast amounts of less relevant data. These challenges are further exacerbated by the growing sophistication of threat actors who continuously adapt their tactics to evade detection [4].

2.3. Limitations of Traditional Threat Intelligence Tools

Traditional threat intelligence tools and methodologies exhibit significant limitations in addressing modern cybersecurity challenges. Many existing platforms excel at data collection but provide insufficient capabilities for automated analysis and correlation. Manual processing requirements create bottlenecks in intelligence workflows, particularly when dealing with large datasets. Furthermore, traditional tools often operate in isolation rather than integrating seamlessly with broader security ecosystems, creating information silos that hinder comprehensive threat visibility. The static nature of many threat intelligence platforms also limits their ability to adapt to rapidly evolving threat landscapes. These technical limitations are frequently compounded by process-related challenges, including inadequate standardization of intelligence formats and inconsistent methodologies for threat assessment [3].

2.4. The Need for Automation and Augmentation

The growing complexity of the threat landscape, combined with the limitations of traditional approaches, underscores the pressing need for automation and augmentation in threat intelligence workflows. Without technological advancement, security teams will continue to face resource constraints that prevent them from fully leveraging

available intelligence. Automation offers the potential to significantly reduce the time required for routine tasks such as data collection, initial triage, and correlation analysis. Augmentation technologies can enhance analyst capabilities by surfacing relevant connections, identifying patterns, and providing decision support. Together, these approaches can enable more proactive security postures and facilitate the shift from reactive threat response to anticipatory defense. Recent industry reports highlight this transition as essential for organizations seeking to maintain effective security operations in increasingly complex digital environments [4].

3. AI Technologies for Threat Intelligence

3.1. Large Language Models (LLMs) in Security Contexts

Large Language Models represent a significant advancement in artificial intelligence with particular relevance to cybersecurity applications. These sophisticated neural network architectures are trained on vast corpora of text data, enabling them to understand and generate human-like text across diverse domains. In security contexts, LLMs demonstrate remarkable capabilities for understanding technical documentation, analyzing threat reports, identifying patterns in security incidents, and generating coherent summaries of complex security events. Their ability to process natural language allows security analysts to interact with them using conversational queries rather than specialized query languages. However, LLMs also present certain limitations when applied to cybersecurity tasks, including potential knowledge boundaries, challenges with temporal awareness, and possibilities of generating incorrect information when presented with unfamiliar scenarios. Additionally, security teams must carefully consider data privacy implications when utilizing these models for sensitive threat intelligence activities [5].

3.2. AI Agents: Architecture and Capabilities

AI agents build upon the foundation of LLMs by incorporating planning and execution capabilities that enable autonomous or semi-autonomous operation within defined domains. Their architecture typically includes components for perception (understanding inputs), reasoning (analyzing situations and determining appropriate responses), planning (developing sequences of actions), and execution (carrying out selected actions through integrated tools). In threat intelligence contexts, these agents can be designed to perform routine tasks such as monitoring threat feeds, enriching indicators of compromise with contextual information, and generating preliminary analyses of security events. Their reasoning capabilities allow them to assess the relevance and severity of potential threats based on organizational context, while their tool utilization features enable integration with existing security infrastructure. This architecture creates possibilities for significant workflow automation while maintaining human oversight for critical decisions [6].

3.3. Retrieval-Augmented Generation (RAG)

Retrieval-Augmented Generation represents a hybrid approach that combines the knowledge retrieval capabilities of traditional information systems with the generative capabilities of language models. This technology addresses certain limitations of standalone LLMs by incorporating external knowledge bases that can be updated independently of the model itself. In threat intelligence applications, RAG systems can access current threat databases, security bulletins, and organizational knowledge repositories to provide context-aware responses grounded in accurate, up-to-date information. By retrieving relevant documents or data points before generating responses, these systems reduce the likelihood of generating incorrect information and enhance the specificity of security recommendations. The retrieval component also improves transparency by allowing analysts to review the sources that informed generated outputs, which proves particularly valuable in security contexts where decision justification is essential [5].

3.4. Integration Possibilities

The integration of LLMs, AI agents, and RAG technologies creates opportunities for developing comprehensive threat intelligence systems that exceed the capabilities of any individual technology. Potential integration approaches include utilizing RAG to provide LLMs with access to specialized security knowledge bases, embedding LLMs within agent frameworks to enable natural language interaction with security tools, and creating multi-agent systems where specialized agents collaborate to analyze different aspects of security events. These integrated systems could potentially support end-to-end threat intelligence workflows, from initial data collection and triage through comprehensive analysis and report generation. By combining strengths of different technologies, integrated systems can address limitations of individual components while providing security teams with enhanced analytical capabilities [6].

Table 1 Comparison of AI Technologies for Threat Intelligence [5, 6]

Technology	Key Capabilities	Primary Applications in TI	Limitations
Large Language Models	Natural language understanding, Pattern recognition	Report summarization, IOC extraction	Knowledge boundaries, Temporal limitations
AI Agents	Autonomous operation, Tool integration	Workflow automation, Continuous monitoring	Governance requirements, Integration complexity
RAG	Grounding in external knowledge, Source transparency	Context-specific intelligence retrieval	Data quality dependencies, Implementation complexity

3.5. Technical Implementation Considerations

The implementation of AI technologies in cybersecurity environments requires careful consideration of various technical factors. Security teams must evaluate infrastructure requirements, including computational resources needed to deploy and operate these systems effectively. Data management considerations include establishing processes for securely storing and accessing the information these systems require. Integration with existing security tools and workflows represents another critical factor, as AI technologies must complement rather than disrupt established security operations. Additionally, organizations must implement appropriate governance mechanisms to ensure these technologies operate within defined ethical and operational boundaries. Performance monitoring frameworks are also essential to track system effectiveness and identify areas for improvement. Each of these considerations plays a crucial role in successful deployment of AI for threat intelligence purposes [5].

4. Application Across the Threat Intelligence Lifecycle

4.1. Data Collection

The initial phase of the threat intelligence lifecycle involves gathering relevant information from diverse sources, a process that can benefit significantly from AI technologies. AI-assisted identification of sources helps security teams discover and prioritize the most valuable intelligence feeds based on organizational needs and threat profiles. Machine learning algorithms can evaluate source reliability and relevance, enabling more targeted collection efforts. Additionally, AI systems can orchestrate collection activities across multiple channels, including surface web, dark web, technical feeds, and industry-specific resources. These technologies enable continuous monitoring of selected sources with minimal human intervention, adaptively adjusting collection parameters based on emerging trends. When implemented effectively, AI-driven collection approaches provide comprehensive coverage while reducing the manual effort traditionally required to maintain awareness across numerous information sources [7].

Table 2 Threat Intelligence Lifecycle with AI Enhancement [7, 8]

Lifecycle Phase	Traditional Challenges	AI Enhancement Opportunities	Key Technologies
Data Collection	Source identification, Coverage gaps	Automated discovery, Continuous monitoring	LLMs, AI agents
Data Processing	Format inconsistencies, Integration issues	Automated cleansing, Entity extraction	LLMs, Machine learning
Analysis	Information overload, Manual correlation	Pattern recognition, Relationship mapping	LLMs, RAG, AI agents
Dissemination	Time-consuming reporting, Distribution issues	Automated reporting, Targeted distribution	LLMs, AI agents

4.2. Data Processing

Once collected, raw threat data requires extensive processing before it becomes suitable for analysis. AI technologies offer significant advantages in automating transformation, cleansing, and integration techniques necessary for effective data preparation. Natural language processing capabilities can standardize information from diverse textual sources, while machine learning algorithms identify and resolve inconsistencies across datasets. Entity recognition systems

extract and normalize key elements such as IP addresses, domain names, and attack techniques, creating structured representations that facilitate subsequent analysis. These technologies also excel at integrating information across multiple sources, establishing connections between seemingly disparate data points. By automating these labor-intensive processing tasks, AI enables security teams to focus their expertise on higher-value analytical activities rather than manual data manipulation [8].

4.3. Analysis and Extraction

The analysis phase represents perhaps the most promising application area for AI in threat intelligence. Machine learning and natural language processing techniques demonstrate particular effectiveness in extracting Indicators of Compromise (IOCs), Tactics, Techniques, and Procedures (TTPs), and other security-relevant information from complex datasets. Large Language Models excel at summarizing extensive threat reports into concise, actionable intelligence briefs focused on organizational priorities. Pattern recognition algorithms identify correlations and trends across historical security events that might elude human analysis. These capabilities enable security teams to process significantly larger volumes of intelligence data while maintaining or improving analytical depth. Furthermore, AI systems can maintain continuous awareness across the threat landscape, identifying emerging attack patterns and notifying analysts when developments warrant human attention [7].

4.4. Few-Shot Learning for Specialized Tasks

Few-shot learning approaches offer particular value for specialized extraction tasks in threat intelligence contexts. This technique enables AI systems to perform new intelligence-related tasks with minimal training examples, addressing the challenge of limited labeled data in cybersecurity domains. By providing a small number of examples demonstrating desired outputs, analysts can guide language models to extract specific types of information from security reports, such as particular attack patterns or emergent threats. This approach proves especially valuable when analyzing novel attack techniques or adapting to evolving threat actor behaviors. Few-shot learning capabilities allow security teams to rapidly adjust analytical focus as threat landscapes change, without requiring extensive model retraining or development of new extraction rules [8].

4.5. Implementation with Security Libraries

Practical implementation of AI for threat intelligence benefits significantly from specialized security libraries and frameworks. Tools like MSTICpy provide extensive capabilities for security data acquisition, enrichment, analysis, and visualization within Python environments. These libraries offer pre-built connectors to common security data sources, specialized data structures for representing security information, and visualization components designed specifically for threat analysis. When combined with AI technologies, these tools enable rapid development of customized intelligence workflows tailored to organizational requirements. Implementation examples demonstrate how these libraries can be integrated with language models and machine learning components to create end-to-end intelligence processing pipelines that enhance analyst productivity while maintaining necessary security controls [7].

4.6. Multi-Agent Systems for Complex Workflows

The most sophisticated applications of AI in threat intelligence involve multi-agent systems where specialized components collaborate to perform complex analysis workflows. These systems typically incorporate agents with distinct roles and capabilities, such as data collection agents, enrichment agents, analytical agents, and reporting agents. By distributing intelligence tasks across specialized components, these architectures can process information at scale while maintaining analytical depth. Agent coordination frameworks manage workflow sequencing and information sharing between components, ensuring coherent end-to-end processing. This approach proves particularly valuable for comprehensive threat investigations that span multiple data sources and analytical techniques. As these systems mature, they increasingly demonstrate capabilities for autonomous investigation of potential threats, presenting human analysts with comprehensive assessments rather than raw data [8].

5. Practical Implementation and Case Studies

5.1. Framework for AI-Enhanced Threat Intelligence Systems

Implementing AI-enhanced threat intelligence systems requires a structured approach that accounts for organizational needs, existing security infrastructure, and desired operational outcomes. A comprehensive implementation framework typically includes distinct phases for assessment, design, development, deployment, and continuous improvement. The assessment phase involves evaluating current threat intelligence capabilities, identifying gaps that AI technologies could address, and establishing clear objectives for enhanced systems. Design considerations include selecting appropriate AI

technologies based on specific use cases, determining data requirements, and establishing integration points with existing security workflows. Development approaches range from utilizing commercial AI-enhanced platforms to building custom solutions that address organization-specific requirements. Throughout implementation, security teams should maintain focus on practical operational outcomes rather than technological sophistication for its own sake [9].

5.2. Technical Architecture Considerations

The technical architecture of AI-enhanced threat intelligence systems must address numerous considerations to ensure effective operation within existing security environments. Data architecture elements include ingestion mechanisms for diverse information sources, storage solutions for both structured and unstructured intelligence data, and processing pipelines that transform raw information into analysis-ready formats. Integration with existing security tools represents another critical consideration, with potential connection points including security information and event management (SIEM) systems, security orchestration and automation platforms, and endpoint detection and response solutions. Computational requirements must be evaluated to ensure adequate resources for AI model operation, particularly when implementing resource-intensive technologies like large language models. Additional architectural considerations include scaling capabilities to accommodate growing data volumes and ensuring appropriate access controls to maintain intelligence confidentiality [10].

Table 3 Implementation Considerations for AI-Enhanced Threat Intelligence [9, 10]

Category	Key Considerations	Potential Solutions
Technical	Infrastructure needs, Integration architecture	Cloud deployment, API-based integration
Operational	Workflow redesign, Staff training	Hybrid workflows, Skills development
Data	Source selection, Quality assurance	Data governance, Preprocessing pipelines
Ethical and Compliance	Privacy protection, Bias mitigation	Privacy-by-design, Compliance frameworks

5.3. Case Studies in Specific Threat Scenarios

Examining case studies of AI application in specific threat scenarios provides valuable insights into practical implementation approaches and potential benefits. Organizations across sectors have deployed AI-enhanced intelligence capabilities to address various security challenges, including advanced persistent threat detection, ransomware prevention, insider threat monitoring, and supply chain risk assessment. These case studies reveal how AI technologies augment human analysts in scenarios requiring processing of large data volumes, recognition of subtle attack patterns, rapid response to emerging threats, and comprehensive situational awareness. They also illustrate different implementation approaches, from targeted applications addressing specific intelligence needs to comprehensive platforms supporting end-to-end threat intelligence lifecycles. While specific implementations vary based on organizational context, common success factors include clear use case definition, phased deployment approaches, and close collaboration between security and data science teams [9].

5.4. Performance Metrics and Evaluation

Evaluating the effectiveness of AI-enhanced threat intelligence systems requires appropriate performance metrics aligned with security objectives. Technical metrics assess model performance through measures such as precision, recall, and F1 scores for detection capabilities, while operational metrics evaluate system impact on security operations through indicators like mean time to detect, mean time to respond, and false positive rates. Workflow efficiency metrics measure productivity improvements through reduced analyst time per investigation and increased intelligence processing capacity. Strategic metrics assess broader impact through improved threat visibility, enhanced risk management capabilities, and prevention of security incidents. Comprehensive evaluation methodologies incorporate multiple metric categories and establish baseline measurements before implementation to enable meaningful assessment of AI-driven improvements. Regular evaluation against these metrics supports continuous refinement of AI-enhanced intelligence capabilities [10].

5.5. Implementation Challenges and Solutions

Organizations implementing AI for threat intelligence encounter various challenges requiring thoughtful solutions. Data-related challenges include insufficient training data for specialized security use cases, inconsistent data formats

across intelligence sources, and difficulties maintaining data quality at scale. Technical challenges involve integrating AI systems with legacy security infrastructure, managing computational requirements for resource-intensive models, and ensuring system reliability during critical security events. Operational challenges include developing appropriate analyst interfaces, maintaining effective human oversight of automated systems, and establishing clear procedures for handling AI-generated intelligence. Solutions to these challenges include developing synthetic data generation approaches for training, implementing flexible integration architectures, establishing hybrid human-AI workflows, and providing comprehensive analyst training on effective collaboration with AI systems [9].

5.6. Ethical and Privacy Considerations

The implementation of AI in threat intelligence contexts necessitates careful attention to ethical and privacy considerations. Privacy concerns arise regarding the handling of potentially sensitive information during intelligence collection and analysis, particularly when monitoring communications or analyzing user behavior patterns. Bias considerations include potential imbalances in training data that could lead to differential system performance across different threat types or actors. Transparency requirements involve ensuring analysts understand the basis for AI-generated recommendations and maintaining appropriate human oversight of automated intelligence processes. Compliance considerations include adherence to relevant data protection regulations and establishing appropriate governance frameworks for AI-enhanced security operations. Addressing these considerations requires implementing privacy-preserving architectures, establishing clear oversight mechanisms, and developing ethical guidelines specific to AI use in security contexts [10].

6. Conclusion

The integration of Large Language Models, AI agents, and Retrieval-Augmented Generation into threat intelligence workflows represents a significant advancement in addressing the challenges posed by increasing data volume and complexity in cybersecurity contexts. These technologies offer considerable potential to transform each phase of the threat intelligence lifecycle, from enhancing data collection processes to revolutionizing analysis capabilities and streamlining dissemination activities. While implementation challenges exist—including technical integration complexities, data quality concerns, and ethical considerations—the potential benefits in terms of increased analytical capacity, improved threat detection, and enhanced operational efficiency provide compelling justification for security teams to invest in AI-enhanced solutions. As threat landscapes continue to evolve in sophistication and scale, the strategic application of these technologies will likely become increasingly essential for maintaining effective security postures rather than merely providing incremental improvements to existing practices. Organizations that successfully implement these technologies within thoughtfully designed workflows, maintaining appropriate human oversight while leveraging AI capabilities for suitable tasks, will be positioned to develop more comprehensive threat awareness and responsive security operations in increasingly complex digital environments.

References

- [1] Amani Ibrahim, et al. "The Challenges of Leveraging Threat Intelligence to Stop Data Breaches." *Frontiers in Computer Science*, August 27, 2020. <https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2020.00036/full>
- [2] Venkoba, et al. "AI for Threat Detection and Prevention: Current Trends, Challenges, and Future Directions." *International Research Journal of Modern Engineering and Technology and Science (IRJMETS)*, October 2024. https://www.irjmets.com/uploadedfiles/paper/issue_10_october_2024/62498/final/fin_irjmets1729148900.pdf
- [3] Suleyman Ozarslan, PhD. "From Noise to Knowledge: Tackling Challenges in Cyber Threat Intelligence." *Picus Security Blog*, October 2, 2024. <https://www.picussecurity.com/resource/blog/from-noise-to-knowledge-tackling-challenges-in-cyber-threat-intelligence>
- [4] IBM Institute for Business Value. "IBM X-Force 2025 Threat Intelligence Index." e, 2025. <https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/2025-threat-intelligence-index>
- [5] IBM Think. "What are Large Language Models (LLMs)?", November 2, 2023. <https://www.ibm.com/think/topics/large-language-models>
- [6] MyCustomAI Blog. "Large Language Models: Biggest Strengths and Worst Limitations.", January 23, 2024. <https://www.mycustomai.io/blog/llms-top-strengths-and-worst-weaknesses>

- [7] World Bank Global Program for Safer Schools "Use of AI Technology to Support Data Collection for Project Preparation and Implementation.", April 2021. https://gps.worldbank.org/sites/gps/files/knowledge_products/2021/Use%20of%20AI%20technology%20to%20support%20data%20collection.pdf
- [8] Cem Dilmegani,. "AI Data Collection: Guide, Challenges and Methods in 2025." AIMultiple Research, April 4, 2025. <https://research.aimultiple.com/ai-data-collection/>
- [9] Arnolnt Spyros et al. "AI-Based Holistic Framework for Cyber Threat Intelligence Management." IEEE Access, January 23, 2025. <https://ieeexplore.ieee.org/stampPDF/getPDF.jsp?arnumber=10851288>
- [10] Lampis Alevizos Martijn Dekker. "Towards an AI-Enhanced Cyber Threat Intelligence Processing Pipeline." Electronics (MDPI), May 22, 2024. <https://www.mdpi.com/2079-9292/13/11/2021>