Cyber Threat Intelligence Sharing: A Privacy-Preserving Framework for Cross-Organization Collaboration

Gustavo Nunes
Department of Computer Science, Federal University of Campina Grande, Brazil

Abstract:

This paper explores a privacy-preserving framework for cross-organization collaboration in cyber threat intelligence sharing. The proposed framework aims to enhance collective defense against cyber threats by facilitating information exchange among organizations while ensuring the privacy and security of shared data. By integrating advanced cryptographic techniques and privacy-preserving mechanisms, the framework addresses key challenges in threat intelligence sharing, including data confidentiality, integrity, and trust.

Keywords: Cyber Threat Intelligence, Privacy-Preserving Framework, Cross-Organization Collaboration, Homomorphic Encryption, Secure Multi-Party Computation, Differential Privacy.

1. Introduction:

In the rapidly evolving landscape of cybersecurity, organizations face increasingly sophisticated and diverse cyber threats. Effective defense against these threats often requires timely and accurate threat intelligence (CTI) sharing among organizations. CTI encompasses data and insights about emerging threats, vulnerabilities, and attack techniques, which are crucial for proactive security measures. However, the collaborative nature of CTI sharing introduces significant challenges, particularly in maintaining data privacy and protecting sensitive information from unauthorized access or misuse[1].

The traditional models of CTI sharing often involve exchanging raw data or detailed threat reports, which can expose sensitive details about an organization's internal systems, vulnerabilities, and security posture. This exposure raises critical concerns regarding data privacy and confidentiality. As organizations collaborate to strengthen their collective cybersecurity defenses, they must also address the risks associated with sharing potentially sensitive information, ensuring that their privacy is not compromised. The primary objective of this research is to develop a privacy-preserving framework for cross-organization collaboration in CTI sharing. The proposed framework aims to facilitate the exchange of threat intelligence while safeguarding the privacy of participating organizations. By integrating advanced cryptographic techniques and privacy-preserving mechanisms, this framework seeks to address the core challenges of maintaining data confidentiality, integrity, and trust during information exchange[2].

In summary, this paper explores the development of a novel framework designed to enhance the security and efficacy of CTI sharing among organizations. The framework will provide a comprehensive approach to balancing the need for collaborative cybersecurity efforts with the imperative of protecting sensitive data, thereby contributing to more robust and resilient cybersecurity practices.

2. Literature Review:

The exchange of cyber threat intelligence (CTI) among organizations is vital for effective cybersecurity. Traditional models of CTI sharing include Information Sharing and Analysis Centers (ISACs), industry-specific sharing groups, and ad-hoc information exchanges facilitated through informal networks. ISACs, for example, provide a structured approach to sharing threat data within specific sectors, enabling organizations to receive timely alerts and collaborate on threat mitigation strategies. However, these models often face limitations in terms of scalability, interoperability, and data granularity. Additionally, the reliance on centralized entities can create bottlenecks and single points of failure[3].

The primary concern in CTI sharing is the preservation of data privacy and security. Sharing threat intelligence often involves exchanging detailed information about an organization’s security environment, which can include sensitive data about network vulnerabilities, attack vectors, and incident response strategies. This exposure raises significant risks related to data confidentiality and integrity. Current practices, such as anonymization and aggregation, offer some protection but may not be sufficient to address all privacy concerns. There is a need for more advanced methods that ensure privacy without compromising the utility of the shared intelligence.

Several privacy-preserving techniques have been explored to address the challenges associated with sensitive data sharing. Homomorphic encryption, for example, allows computations to be performed on encrypted data, ensuring that the data remains confidential even during processing. This technique has shown promise in secure data analysis but faces challenges related to computational efficiency and practicality. Secure multi-party computation (MPC) enables multiple parties to jointly compute functions over their inputs while keeping those inputs private. MPC has been used in various applications, including privacy-preserving data sharing, but its implementation can be complex and resource-intensive. Differential privacy offers a framework for sharing aggregated data in a way that prevents the disclosure of individual data points, providing strong privacy guarantees. However, achieving a balance between privacy and data utility remains a significant challenge[4].

Recent research has focused on integrating these privacy-preserving techniques to enhance the security and effectiveness of CTI sharing frameworks. Combining homomorphic encryption with secure multi-party computation can provide robust privacy protection while enabling useful data analysis. Differential privacy techniques can be incorporated to offer additional layers of privacy guarantees, particularly in scenarios involving aggregated threat data. The integration of these mechanisms into a cohesive framework requires addressing various challenges, including system design, scalability, and performance optimization.

In conclusion, while existing CTI sharing models provide valuable insights and facilitate collaboration, they often fall short in addressing privacy concerns effectively. Privacy-preserving techniques offer promising solutions but need to be integrated into practical and scalable frameworks. This paper aims to build upon these insights to propose a novel privacy-preserving framework that addresses the limitations of current models and enhances the collaborative potential of CTI sharing.

3. Framework Design:

The proposed privacy-preserving framework for cyber threat intelligence (CTI) sharing is designed to facilitate secure and efficient information exchange among organizations while preserving the confidentiality of sensitive data. The framework architecture consists of several key components: the Data Provider, the Privacy-Preserving Engine, the Data Aggregator, and the Data Consumer. The Data Provider is responsible for submitting threat intelligence data to the system, while the Data Consumer queries and retrieves the necessary information. The Privacy-Preserving Engine is the core component, implementing cryptographic techniques to ensure that data privacy is maintained throughout the process. The Data Aggregator consolidates threat intelligence from multiple sources, ensuring that the aggregated data is both useful and privacy-preserving[5].

To address the privacy concerns inherent in CTI sharing, the framework integrates several advanced privacy-preserving mechanisms:

Homomorphic Encryption: This technique allows computations to be performed on encrypted data, enabling the Privacy-Preserving Engine to process threat intelligence without decrypting it. This ensures that the data remains confidential throughout the analysis and query process. Although homomorphic encryption provides strong privacy guarantees, its computational overhead is a consideration, which the framework addresses by optimizing encryption schemes and reducing the complexity of operations. Secure Multi-Party Computation (MPC): MPC enables multiple organizations to collaboratively compute functions over their inputs while keeping their data private. In the context of CTI sharing, MPC can be used to perform joint analysis of threat intelligence data without revealing individual contributions. The framework utilizes efficient MPC protocols to balance privacy with computational efficiency, ensuring that the collaborative analysis is both secure and feasible[6]. Differential Privacy: This technique is employed to provide privacy guarantees when sharing aggregated threat intelligence. Differential privacy ensures that individual data points cannot be discerned from the aggregated results, even if an attacker has access to the aggregated data. The framework integrates differential privacy mechanisms to protect the privacy of organizations' data when providing high-level insights and trends. Ensuring the integrity and authenticity of threat intelligence data is crucial for the framework's effectiveness. The proposed framework incorporates mechanisms for data integrity verification and authentication. Cryptographic hashing and digital signatures are used to verify the integrity of the data submitted by the Data Providers. This ensures that the data has not been tampered with during transmission or storage. Additionally, access controls and authentication protocols are implemented to verify the identity of Data Providers and Consumers, preventing unauthorized access to sensitive information. Scalability is a key consideration in the design of the privacy-preserving framework. The architecture is designed to handle large volumes of threat intelligence data and support a growing number of participating organizations. Performance optimization techniques, such as efficient encryption algorithms, parallel processing, and data compression, are employed to minimize the impact of privacy-preserving mechanisms on the system's overall performance. By leveraging these techniques, the framework aims to provide timely and actionable threat intelligence while maintaining strong privacy protections[7].

In summary, the framework design integrates advanced privacy-preserving techniques within a structured architecture to facilitate secure and efficient CTI sharing. By addressing privacy, data integrity, and performance concerns, the proposed framework aims to enhance collaborative cybersecurity efforts while safeguarding sensitive information.

4. Implementation:

To bring the proposed privacy-preserving framework to life, a prototype implementation was developed to demonstrate its functionality and effectiveness. The prototype was built using a modular architecture, allowing for easy integration of various privacy-preserving techniques and components. The core modules of the prototype include the Privacy-Preserving Engine, the Data Aggregator, and the Data Consumer Interface. The Privacy-Preserving Engine integrates homomorphic encryption, secure multi-party computation (MPC), and differential privacy mechanisms to ensure data confidentiality and privacy. The Data Aggregator module is responsible for consolidating and aggregating threat intelligence from multiple sources, while the Data Consumer

Interface provides a user-friendly way for organizations to query and retrieve the information they need[8].

Integrating the prototype with existing CTI platforms and tools is a crucial step in validating its practicality and effectiveness. The framework was designed with compatibility in mind, supporting standard data formats and protocols commonly used in CTI systems. To facilitate integration, a set of application programming interfaces (APIs) was developed, allowing seamless communication between the framework and existing cybersecurity tools. The prototype was tested with several widely-used CTI platforms to ensure that it can be incorporated into existing workflows without disrupting current operations. Additionally, the integration process included ensuring that the privacy-preserving mechanisms did not compromise the functionality or performance of the existing systems[9].

The effectiveness and reliability of the privacy-preserving framework were assessed through a series of rigorous tests and validations. Testing focused on several key areas: data privacy, system performance, and usability. To evaluate data privacy, the framework was subjected to various privacy attacks and adversarial scenarios to ensure that it met its privacy guarantees. Performance testing involved measuring the system's response times, throughput, and computational overhead associated with the privacy-preserving mechanisms. The framework's usability was assessed by conducting user feedback sessions with cybersecurity professionals, who provided insights into the practicality and effectiveness of the framework in real-world scenarios[10].

The results from the testing and validation phase demonstrated that the prototype effectively maintained data privacy while providing useful threat intelligence. The integration with existing systems was successful, with no significant disruptions reported. Performance tests indicated that the privacy-preserving mechanisms, while introducing some computational overhead, did not adversely affect the overall system performance. User feedback highlighted the framework's potential for enhancing collaborative cybersecurity efforts while addressing privacy concerns. However, some areas for improvement were identified, including optimization of encryption schemes and refinement of the user interface for better usability[11].

In summary, the implementation of the privacy-preserving framework involved developing a functional prototype, integrating it with existing CTI systems, and conducting thorough testing and validation. The results indicate that the framework effectively balances privacy and performance, providing a practical solution for secure CTI sharing. Future work will focus on addressing identified areas for improvement and expanding the framework's capabilities to further enhance its utility and effectiveness in collaborative cybersecurity efforts[12].

5. Discussion:

The proposed privacy-preserving framework for cyber threat intelligence (CTI) sharing offers several advancements over traditional models. Unlike conventional models, which often rely on centralized data repositories and standard anonymization techniques, this framework incorporates advanced cryptographic methods, such as homomorphic encryption, secure multi-party computation (MPC), and differential privacy. These techniques provide a more robust approach to maintaining data confidentiality and integrity while enabling effective collaboration. Traditional models can be limited by their reliance on trust between parties and the potential for sensitive data exposure, whereas the proposed framework ensures that even during data processing and analysis, the information remains secure. The integration of these privacy-preserving techniques addresses many of the gaps and vulnerabilities present in current CTI sharing methods. One of the strengths of the proposed framework is its scalability and adaptability to various organizational contexts. The modular architecture allows for flexible deployment, making it suitable for both small organizations and large enterprises. The framework's design accommodates a growing number of participants and high volumes of threat intelligence data without compromising performance. Moreover, the framework's use of standardized data formats and APIs facilitates integration with a wide range of existing CTI platforms and tools[13]. This adaptability ensures that organizations can implement the framework in a manner that best fits their specific needs and cybersecurity environments. However, challenges related to the scalability of privacy-preserving techniques, such as the computational overhead of homomorphic encryption, remain an area for ongoing research and optimization. While the proposed framework addresses many of the current challenges in CTI sharing, there are several areas for future research and development. One key area is the optimization of cryptographic techniques to reduce computational overhead and improve efficiency. Advanced homomorphic encryption schemes and more efficient MPC protocols could further enhance the framework's performance. Additionally, exploring the integration of emerging technologies, such as blockchain for immutable data logs and decentralized trust models, could provide additional layers of security and transparency[14]. Expanding the framework's capabilities to handle diverse types of threat intelligence data, including unstructured data and real-time streaming information, would also enhance its applicability and effectiveness. The adoption of the proposed privacy-preserving framework has significant implications for collaborative cybersecurity efforts. By enabling organizations to share threat intelligence securely and privately, the framework fosters greater cooperation and information exchange across the cybersecurity community. This collective approach can lead to more comprehensive threat detection and response strategies, ultimately strengthening the overall security posture of participating organizations. The framework's focus on privacy protection also addresses concerns about data misuse and unauthorized access, building trust among collaborators and encouraging wider participation in CTI sharing initiatives.

In summary, the proposed privacy-preserving framework represents a significant advancement in secure and effective CTI sharing. Its comparison with existing models highlights its strengths in addressing privacy concerns, scalability, and adaptability. Future research and development will focus on optimizing performance, exploring new technologies, and expanding the framework's capabilities. The framework's potential to enhance collaborative cybersecurity efforts underscores its importance in advancing collective defense against cyber threats[15].

6. Conclusion:

In conclusion, the proposed privacy-preserving framework for cyber threat intelligence (CTI) sharing represents a significant advancement in addressing the critical challenge of maintaining data privacy while fostering effective cross-organization collaboration. By integrating sophisticated cryptographic techniques such as homomorphic encryption, secure multi-party computation, and differential privacy, the framework ensures that sensitive threat intelligence remains confidential and secure throughout the sharing and analysis processes. The successful implementation and validation of the framework demonstrate its ability to balance privacy concerns with the need for actionable intelligence, offering a practical solution for enhancing collective cybersecurity efforts. As organizations increasingly recognize the value of collaborative threat intelligence, the framework provides a robust foundation for secure information exchange, paving the way for more resilient and effective defenses against evolving cyber threats. Future research will continue to refine and expand upon this framework, addressing performance optimization and exploring emerging technologies to further enhance its capabilities and impact.

References:

[1] A. K. Y. Yanamala, "Secure and Private AI: Implementing Advanced Data Protection Techniques in Machine Learning Models," *International Journal of Machine Learning Research in Cybersecurity and Artificial Intelligence,* vol. 14, no. 1, pp. 105-132, 2023.

[2] L. M. d. F. C. Guerra, "Proactive Cybersecurity tailoring through deception techniques," 2023.

[3] A. K. Y. Yanamala, S. Suryadevara, and V. D. R. Kalli, "Evaluating the Impact of Data Protection Regulations on AI Development and Deployment," *International Journal of Advanced Engineering Technologies and Innovations,* vol. 1, no. 01, pp. 319-353, 2023.

[4] V. M. Reddy and L. N. Nalla, "The Future of E-commerce: How Big Data and AI are Shaping the Industry," *International Journal of Advanced Engineering Technologies and Innovations,* vol. 1, no. 03, pp. 264-281, 2023.

[5] A. K. Y. Yanamala, "Data-driven and artificial intelligence (AI) approach for modelling and analyzing healthcare security practice: a systematic review," *Revista de Inteligencia Artificial en Medicina,* vol. 14, no. 1, pp. 54-83, 2023.

[6] V. M. Reddy, "Data Privacy and Security in E-commerce: Modern Database Solutions," *International Journal of Advanced Engineering Technologies and Innovations,* vol. 1, no. 03, pp. 248-263, 2023.

[7] A. Joseph, "A Holistic Framework for Unifying Data Security and Management in Modern Enterprises," *International Journal of Social and Business Sciences,* vol. 17, no. 10, pp. 602-609, 2023.

[8] A. K. Y. Yanamala and S. Suryadevara, "Advances in Data Protection and Artificial Intelligence: Trends and Challenges," *International Journal of Advanced Engineering Technologies and Innovations,* vol. 1, no. 01, pp. 294-319, 2023.

[9] N. Pureti, "Strengthening Authentication: Best Practices for Secure Logins," *International Journal of Advanced Engineering Technologies and Innovations,* vol. 1, no. 01, pp. 271-293, 2023.

[10] B. R. Maddireddy and B. R. Maddireddy, "Adaptive Cyber Defense: Using Machine Learning to Counter Advanced Persistent Threats," *International Journal of Advanced Engineering Technologies and Innovations,* vol. 1, no. 03, pp. 305-324, 2023.

[11] N. Pureti, "Responding to Data Breaches: Steps to Take When Your Data is Compromised," *International Journal of Machine Learning Research in Cybersecurity and Artificial Intelligence,* vol. 14, no. 1, pp. 27-50, 2023.

[12] B. R. Maddireddy and B. R. Maddireddy, "Automating Malware Detection: A Study on the Efficacy of AI-Driven Solutions," *Journal Environmental Sciences And Technology,* vol. 2, no. 2, pp. 111-124, 2023.

[13] N. Pureti, "Encryption 101: How to Safeguard Your Sensitive Information," *International Journal of Advanced Engineering Technologies and Innovations,* vol. 1, no. 01, pp. 242-270, 2023.

[14] B. R. Maddireddy and B. R. Maddireddy, "Enhancing Network Security through AI-Powered Automated Incident Response Systems," *International Journal of Advanced Engineering Technologies and Innovations,* vol. 1, no. 02, pp. 282-304, 2023.

[15] N. Pureti, "Anatomy of a Cyber Attack: How Hackers Infiltrate Systems," *Revista de Inteligencia Artificial en Medicina,* vol. 14, no. 1, pp. 22-53, 2023.