MANTRA: Towards a Conceptual Framework for Elevating Cybersecurity Applications Through Privacy-Preserving Cyber Threat Intelligence Sharing

Philipp Fuxen*, Murad Hachani*, Rudolf Hackenberg*, Mirko Ross†
*Dept. Informatics and Mathematics, OTH Regensburg
Regensburg, Germany
Email: {philipp.fuxen, murad.hachani, rudolf.hackenberg}@oth-regensburg.de
†asvin GmbH
Stuttgart, Germany
Email: m.ross@asvin.io

***Abstract*—In light of the escalating cyber threat landscape, this paper highlights the critical importance of Cyber Threat Intelligence while acknowledging the challenges that impede its effective dissemination, including reputational risks, technical barriers, and the existence of data silos. To address these issues, we propose the conceptual framework of the MANTRA network—a theoretical privacy-preserving Cyber Threat Intelligence sharing model intended to enhance cybersecurity measures across organizations of varying sizes and resource capacities. The MANTRA concept endeavors to overcome these dissemination challenges through the adoption of federated learning for dismantling data silos, the enhancement of data analytics for managing information overload, the application of secure protocols and peer-to-peer communication for safeguarding the confidentiality, integrity, and availability of Cyber Threat Intelligence data, and the promotion of inter-organizational collaboration via socio-economic governance models. This holistic strategy aims not only to facilitate the exchange of information on cyber threats, but also to strengthen the collective defense against the ever-evolving cyber threats. Central to this theoretical exploration are pivotal research questions: identifying the most effective data sources for the envisioned MANTRA network, discerning the methodologies and technologies critical for secure and efficient data exchange within MANTRA, and comprehending how specific application scenarios of MANTRA might impact the efficiency of cybersecurity tactics across diverse organizational contexts. In conclusion, MANTRA presents a concept that combines a hybrid peer-to-peer architecture with federated learning and offers a promising framework for privacy-preserving Cyber Threat Intelligence sharing that should be further explored and validated in future research.**

***Keywords*-Cyber Threat Intelligence; Federated Learning; Privacy-Preserving Data Sharing; Cybersecurity.**

I. INTRODUCTION

In an increasingly hyper-connected world, where the digital infrastructure of companies and organizations is constantly growing and evolving, we face a challenge: the threat of cyber-attacks. These attacks are not only on the rise in frequency, but they are also becoming more sophisticated, targeting a wide range of sectors and businesses, regardless of their size or industry. Recent analyses, including the Federal Office for Information Security (BSI) Report 2023 [1], the Google Cloud Threat Horizons from August 2023 [2], and the CrowdStrike Global Threat Report 2024 [3], underline the complexity and broad spectrum of cyber threats. They point to the diversity of cyberattacks, ranging from critical infrastructure to cloud resources, and emphasize the increasing use of Artificial Intelligence (AI) by cyber attackers. This work demonstrates the significant impact of these threats on privacy, security, and economic stability and highlights the urgent need for robust cybersecurity measures. Against this backdrop, organizations must act to protect themselves against these threats. A central pillar of this is Cyber Threat Intelligence (CTI) - the collection, analysis, and understanding of information about potential threats. CTI enables companies to detect threats at an early stage, react proactively, and continuously improve their defense strategies.

A key challenge in the field of CTI is the effective exchange of relevant threat information between organizations and actors. Despite the growing awareness of the need for CTI sharing, many organizations face several challenges. These include reputational and privacy concerns, as well as technical barriers such as incompatibility of the system, different data formats, and communication channels. Furthermore, the exchange of CTI data often takes place informally by email or telephone, with the effectiveness and scope of the exchange strongly dependent on personal relationships. These concerns and barriers lead companies to keep their CTI data in isolated environments or silos to minimize the risk of disclosure, among other things. This leads to a limited overall threat landscape, incomplete information, and increased security risks, as potential threats may not be identified or addressed early enough. Dealing with an enormous flood of information is another challenge. Companies are faced with an overwhelming amount of CTI data that needs to be captured, processed, and interpreted efficiently. The volume of information can tie up resources and affect their ability to distinguish relevant insights from irrelevant noise. This makes it difficult to identify and prioritize potential threats, slows response times, and can lead to delayed or inadequate defense against attacks. For small organizations, this problem is severe as they have limited resources and cybersecurity expertise. For them, managing and making sense of the vast amounts of CTI data can be even more challenging, increasing their vulnerability to cyber threats and impacting their overall security posture.

The main objective of the paper is to propose the refined concept of the privacy-preserving sharing network MANTRA [4] and to answer the following research questions:

 ISBN: 978-1-68558-156-5

RQ1: What types of data sources are suitable for enhancing the cyber threat intelligence capabilities of the MANTRA network?

RQ2: What methods and technologies enable secure and efficient sharing of data within the MANTRA network?

RQ3: What impact could the tailored use of MANTRA have on cybersecurity strategies tailored for diverse organizational contexts?

This paper is structured as follows: It starts by reviewing existing literature in Section II. The objectives of the MANTRA initiative are detailed in Section III, followed by an overview of the architecture in Section IV. Section V discusses the various data sources utilized by the network. The use and impact of MANTRA applications are explored in Section VI. Finally, the paper concludes with a summary of findings and outlines directions for future research in Section VII.

II. Related Work

In the domain of CTI sharing, contemporary research highlights a diversity of methodologies and associated challenges. Various architectures for Threat Intelligence Sharing Platform (TISP) have been developed, including centralized systems like Malware Information Sharing Platform (MISP) [5], cloud-based frameworks [6], and, increasingly noted in the literature, blockchain-enabled platforms [7]–[10]. Each architecture aims to achieve specific objectives, such as improving anonymity to lower the threshold for the sharing of organizational information or creating incentives for participation. These architectures each present a unique set of advantages and limitations. Given the ascending interest in blockchain-enabled TISPs, this section delves into an analysis of differing approaches within this category.

"Siddhi" [9] and "LUUNU" [10] represent blockchain-enhanced platforms for CTI sharing, devised to enhance organizational engagement through the robust privacy protections offered by ledger technologies, including traceability and data provenance. Siddhi, built upon Rahasak, introduces an administrative validation process to bolster trust within the network and adopts a Self-Sovereign Identity (SSI)-enabled registration for anonymity. LUUNU, while employing similar technologies, extends its functionalities with Federated Learning (FL) and improved data storage through MISP and Model Cards, going beyond mere CTI sharing as seen in Siddhi to also include cyber threat detection capabilities, such as Denial of Service (DoS) attacks, through the training of Machine Learning (ML) models leveraging MISP's off-chain repository. Although a significant focus is placed on anonymity and data integrity, the broader discourse often omits considerations related to blockchain aspects, such as consensus algorithms, which significantly affect network leadership dynamics and throughput. Zhang et al. [7] proposed a consensus algorithm tailored for a consortium blockchain that employs Proof-of Reputation (PoR), encompassing mechanisms for CTI data sanitization, the generation of sensitive information proposals, and the automation of CTI responses. Unlike these approaches focused on technology and security, Nguyen et al. [8] advocate for a blockchain framework aimed at Industrial Control Systems (ICS), predominantly emphasizing incentives to foster participation, including subscription discount strategies, thus presenting a distinct perspective focused on engagement through economic motivators.

Heo et al. [6] developed a hybrid cloud-based model for CTI sharing designed to facilitate the ease of use for individuals by addressing resource constraints such as time, capabilities, and cost. Importantly, their approach heavily relies on industry standards to ensure interoperability, security, and ease of integration. This reliance on standards is strategic to reduce barriers to entry for CTI sharing and ensure that even entities with limited cybersecurity resources can participate effectively and safely in the threat intelligence ecosystem.

Wagner et al. [11] highlight several barriers to CTI sharing, pinpointing the critical need for enhanced automation, trust, and interoperability. MANTRA, as a concept, introduces a distinct framework within this diverse ecosystem. It structures a hybrid peer-to-peer network, primarily designed to create an optimal data flow for FL. This approach not only maintains the anonymity of data sources, but also ensures that sensitive information is kept secure, addressing key concerns in cybersecurity information sharing. Moreover, the focus on Federated Learning enables MANTRA to leverage the collective intelligence of various entities while minimizing the risks associated with centralized data storage and management. To facilitate the sharing of sensitive information, in the envisioned framework of MANTRA, a systematic procedure is proposed for the cleansing and attribution of data, which is then leveraged for the training of AI models. This ensures that only attack-specific information is externalized in the form of trained AI models, addressing concerns regarding the disclosure of vulnerable information. Furthermore, MANTRA addresses time, capabilities, and cost constraints by providing dedicated models and guidelines for the direct implementation of security measures in detection, prevention, attribution, and response tasks. This approach not only streamlines the process, but also helps eliminate redundant or missing information through a guided cleanup process. In addition, MANTRA strives to address the scarcity of security information through the implementation of advanced attribution models and guidelines. This effort aims not only to increase the quantity of CTI information but also to enhance its quality, directly addressing the critical issues of CTI sharing identified in current research and development landscapes. By focusing on the sanitized model-based exchange of intelligence and attribution, MANTRA aims to overcome the limitations of current CTI sharing platforms, offering a new paradigm that emphasizes data privacy, security, and the efficient use of collective cybersecurity insights.

In summary, MANTRA primarily aims to streamline complexity by concentrating on federated learning, steering clear of the sometimes resource-demanding or complex blockchain architectures that are inherently designed for privacy and traceability. Instead, MANTRA manifests in the domain-specific

 ISBN: 978-1-68558-156-5

model exchanges, such as Intrusion Detection System (IDS) models or attribution models, which are developed across various peers. The absence of blockchain-provided financial incentives is counterbalanced by an application-driven architecture. This architecture encourages the fortification of organizations, industry sectors, or even supply chains through models trained collaboratively that are already accessible. Ultimately, MANTRA commits to enhancing the security of training and distributing AI models, ensuring security measures are in place from the outset, even before the information traverses the network.

III. Objective of MANTRA

In this section, we outline the primary goals of the MANTRA network: Bridging data silos and managing the CTI data flood, ensuring the principles of confidentiality, integrity, and availability of shared CTI data, and developing socio-economic governance models to promote information sharing.

MANTRA is designed as a robust platform that facilitates the efficient exchange of CTI, ensuring data protection and privacy for all participating organizations. Our initiative focuses on dismantling data silos that hinder the seamless flow of information between different sectors and organizations. These silos often impede the effective dissemination of CTI and make it difficult to detect threats. To overcome these obstacles, MANTRA adopts a privacy-friendly approach through federated learning, allowing data to stay decentralized and processed locally. This strategy not only preserves the participant's privacy but also promotes effective collaboration and sharing of CTI.

Additionally, MANTRA addresses the issue of information overload by streamlining data aggregation, processing, and analysis. By enhancing the efficiency of data processing, we aim to refine the quality and relevance of shared CTI, enabling organizations to better understand the threat landscape and strengthen their defense mechanisms.

Another focus of MANTRA centers on the confidentiality, integrity, and availability of shared CTI data, employing robust security protocols to safeguard against unauthorized access. By utilizing Peer-to-Peer (P2P) communication, MANTRA enhances security and data protection, diminishing dependency on vulnerable central servers. This strategy reduces outage risks and potential attack vectors, ensuring a more secure, private, and resilient communication network. Through these initiatives, MANTRA seeks to transform CTI exchange, facilitating efficient and secure dissemination of crucial threat intelligence throughout the network.

MANTRA emphasizes developing socio-economic governance models to facilitate sensitive cybersecurity information sharing within supply chains and across organizations. These models offer a collaborative framework that uses social and economic incentives to encourage participation in data sharing. Through such incentives, organizations are encouraged to contribute their CTI, improve security and resilience within supply chains, and increase the overall effectiveness of cybersecurity.

IV. Overview of MANTRA Architecture

The general architecture of MANTRA consists of several components that work together to achieve the objectives. These components include a protocol layer, an application layer, and a federated learning layer. In the design of the MANTRA network, several foundational assumptions are introduced to outline the architecture's operational framework. Primarily, entities within the network are conceptualized as peers, with the typical participant being identified as an organization. Beyond the baseline of organizational participation, MANTRA integrates a subset of peers distinguished by their high trust level. This cathegorization facilitates a governance model wherein the generation of new global models and the oversight of the federated learning process are predominantly managed by entities of higher trust, such as governmental agencies. The network becomes a hybrid peer-to-peer model, due to its federated learning process that ensures a secure and regulated environment, protecting shared threat intelligence models.

The protocol layer serves as the core component and uses a hybrid P2P protocol that ensures secure and confidential information exchange between participants and forms the basis for reliable network communication and data transmission. This hybrid P2P model optimizes resource utilization and improves scalability by combining the efficiency of centralized management with the robust, secure framework of decentralized networks, providing a flexible architecture for secure CTI data processing.

The application layer is a key component of the MANTRA architecture that is responsible for pre-processing and training the models with CTI data. In addition to creating and training models, this layer is also responsible for data integration and data management to ensure efficient use and processing of CTI information. Moreover, the application layer implements applications that include the trained and aggregated models to provide participating organizations with relevant insights in areas, such as prevention, detection, response, and attribution.

The federated learning layer plays a central role in the MANTRA architecture by aggregating the individually trained models. This aggregation enables a global overview of the CTI data without having to share the raw data between participants. This preserves the privacy and security of the data while allowing the results to be analyzed and shared.

In the MANTRA network, peers can take on different tasks depending on their skills and resources, as shown in Figure 1, which shows the types of peers and the MANTRA layers described. Three peer types are defined to manage the variety of tasks in the network and to meet the different requirements: Training Peer (TP), Aggregation Peer (AP), and Operational Peer (OP). The TP trains models locally with CTI data, starting with an initial model from the AP, and sends its updated model back for aggregation. The AP combines these local models into a comprehensive global model for network-wide distribution. The OP, lacking the resources for local training, can use the global model to leverage collective insights. This structure ensures that all peers, regardless of their model train-

 ISBN: 978-1-68558-156-5

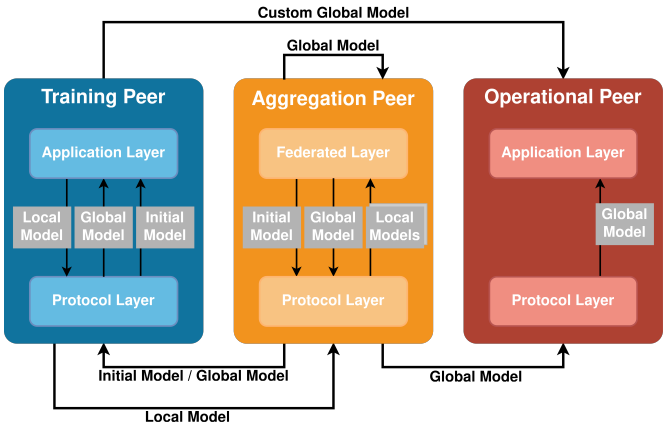

Figure 1: Peer Types of MANTRA.

ing capabilities, contribute to and benefit from the network's collective insights and cybersecurity efforts. Beyond the scope of federated learning communication, training peers also have the capability to distribute enhanced models, enriched with supplementary data, to other training or operational peers. This can particularly be exemplified by the transfer of information from a larger entity to smaller, dependent organizations.

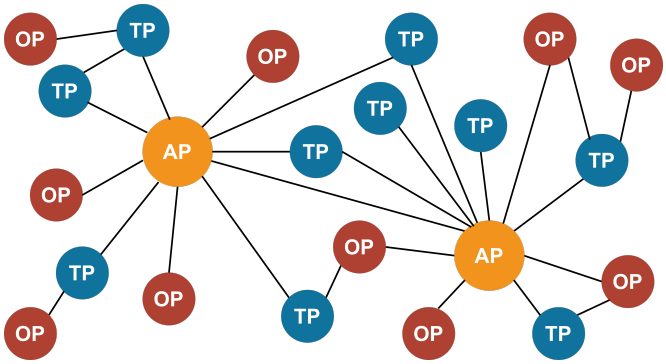

Figure 2: Topology of Peers.

Figure 2 shows a schematic representation of the structure of the MANTRA network, illustrating the roles and interconnectivity of the different peer types. Peers have full control over communication, exchange, and the training of models. To maintain the effectiveness and security of the data flow in the FL architecture, certain communication paths are essential. Thus, a TP must communicate with at least one AP to introduce new information in the form of models into the system. The 1:N relationship ensures network reliability. In addition, TPs and OPs have the option to develop their own security models outside of the central FL cycle, which could, for example, enhance the detection rate. Since smaller companies often do not manage or are unable to manage CTI, they can act as OPs within the network. This enables them to protect themselves through global MANTRA models or, by retrieving models through TPs, to map and stabilize entire supply chains. A network architecture oriented towards FL must include some form of evaluation. Since an AP only aggregates a portion of the models in the network, this represents merely a preliminary stage of aggregation. Therefore, APs must be capable of exchanging information. However, since authenticity and validation are beyond the scope of this work and are still partially under research, they are not addressed in this publication.

V. Data Sources for MANTRA

The following section examines the various data sources for the MANTRA system, especially CTI data. CTI data includes information about current threats, vulnerabilities, attack patterns, malware analysis, and other relevant security aspects that can come from different sources. This data plays a central role in detection, analysis, and defense against cyber attacks. In addition to the internal data generated by the network participants, it is also important to use external data sources to obtain a comprehensive picture of the threat landscape. One example of external data sources is Open Source Intelligence (OSINT), which is already in use for the CTI. Other possible data sources that are important for creating and improving the models in the MANTRA system are discussed in the following.

A. Internal Data

The use of internal data is an essential part of the MANTRA system and enables an understanding of the individual security landscape of the participants. Internal data includes participant-generated information, such as log files, event data, network traffic analysis, system configurations, and other internal security data. This data provide unique insights into specific security threats and vulnerabilities that an organization is exposed to. Important internal data sources include:

- **Network and security logs:** These include firewall logs with details of blocked or suspicious connections, IDS/Intrusion Prevention System (IPS) logs that reveal unusual patterns in network traffic, and VPN logs that provide indications of unusual login attempts. They are crucial for detecting and responding to security threats.
- **System and application logs:** Operating system logs provide information about unauthorized access or system errors. Web server and application logs provide information about suspicious requests and security incidents that are essential for the security of applications and systems.
- **Security Event and Information Management (SIEM) data:** SIEM systems collect event data and generate alerts that identify suspicious activity on the network, which is essential for comprehensive security monitoring.
- **Endpoint Detection and Response data:** Endpoint Detection and Response (EDR) data provides insights into behavior-based threat detection and forensic information about endpoints that are important for detecting and responding to advanced threats.
- **Threat intelligence feeds and incident reports:** These provide information on current threat trends, Indicator of

 ISBN: 978-1-68558-156-5

Compromise (IoC)s and the attackers' TTPs and help to improve preventive security measures.

The use of internal data within the MANTRA system presents several challenges that must be carefully addressed to ensure the security, effectiveness, and reliability of the sharing of cyber threat intelligence. In particular, these challenges include protecting the privacy and confidentiality of sensitive information, ensuring data quality and integrity, overcoming data integration and compatibility issues, scaling the system to handle growing volumes of data, and complying with legal and regulatory requirements. These challenges require the use of technologies and methods, such as encryption, access controls, data validation mechanisms, efficient data integration tools, and scalable architectures to create a secure and effective platform for the exchange of threat data. Additionally, continuous adaptation to changing regulatory frameworks is essential to ensure compliance and increase the confidence of participants in the MANTRA system.

B. External Data

External data sources complement internal data by providing a broader perspective on the global cyber threat landscape. These sources are essential for the MANTRA system to obtain a complete landscape of threats and attack tactics that go beyond the immediate experience of the participating organizations. Important external data sources include:

- **OSINT:** OSINT includes data from publicly available sources that contain information on new and existing threats. These sources include news reports, articles, security blogs, and public vulnerability databases that provide information on current cyber threat trends.
- **Threat Intelligence Feeds:** Specialized services provide real-time information and data feeds on detected threats and vulnerabilities. These feeds provide valuable data that can be integrated directly into the MANTRA system to improve threat detection and response capabilities.
- **Sector-specific security reports:** Reports and analyses published by security companies and industry associations provide deep insight into specific threat vectors and attack patterns within specific sectors. This information is particularly valuable for companies operating in high-risk areas.
- **Government and authority notifications:** Information from national and international security agencies provides authoritative information on cyber threats, warnings, and recommendations. Integrating these data helps the MANTRA system adapt to the evolving security landscape and strengthen defense strategies against state-sponsored cyberattacks.
- **Other CTI sharing platforms:** Community platforms and networks that promote the exchange of threat intelligence between organizations are a valuable source of up-to-date and relevant information about threats and attacks.

The use of external data sources enables the MANTRA system to create a more comprehensive and up-to-date basis for generating Threat Intelligence. By integrating data from different sources, MANTRA can develop more accurate models for threat detection, making a valuable contribution to strengthening the cyber resilience of participating organizations. The challenge lies in continuously evaluating the credibility and quality of external data and ensuring that this information is used effectively to improve understanding and responsiveness to cyber threats.

VI. Applications of MANTRA

This section focuses on the components of the MANTRA application layer. It includes two sections: "Data Integration and Management", which looks at the processes and technologies used to efficiently integrate and manage CTI data on the network, and "Use Cases", which looks at practical applications and tangible benefits of MANTRA for cybersecurity operations. It discusses how MANTRA could transform operational performance and collective security standards of participating organizations.

A. Data Integration and Management

The integration and management of data in MANTRA is crucial to achieving CTI, as it determines the quality and effectiveness of subsequent analyses. Data integration is carried out by the participants in the MANTRA network and involves various input methods. To ensure that all relevant information is efficiently captured and prepared for further processing, several steps are carried out before the data is integrated.

Figure 3 illustrates the data flow from input to integration into the MANTRA system. MANTRA participants enter data from a Threat Intelligence Platform (TIP), via a web form or with predefined templates. This data is converted into the Structured Threat Information eXpression (STIX) format, which provides a standardized structure for the exchange of CTI. During the classification phase, the CTI data is assigned metadata and tags such as CTI type, trustworthiness, and relevance. The data then undergoes various processing steps: it is cleansed, aggregated, and anonymized to ensure that it is used in compliance with data protection regulations. The color code of the Traffic Light Protocol (TLP) indicates the degree of sensitivity of the information and guides the data sharing. After this processing, the data is filtered and integrated into the MANTRA system.

B. Use Cases

Within the architecture of MANTRA, the applications play a crucial role in the practical implementation and benefits of the system. This section shows how the CTI provided by MANTRA could contribute to the protection and resilience of organizations in different use cases. MANTRA is designed to provide organizations of all sizes - from small businesses to large corporations - with tools that enable them to obtain accurate and timely threat intelligence. By applying the global models created through the interaction of training peers and aggregator peers, network participants should be able to improve their security in the areas of prevention,

 ISBN: 978-1-68558-156-5

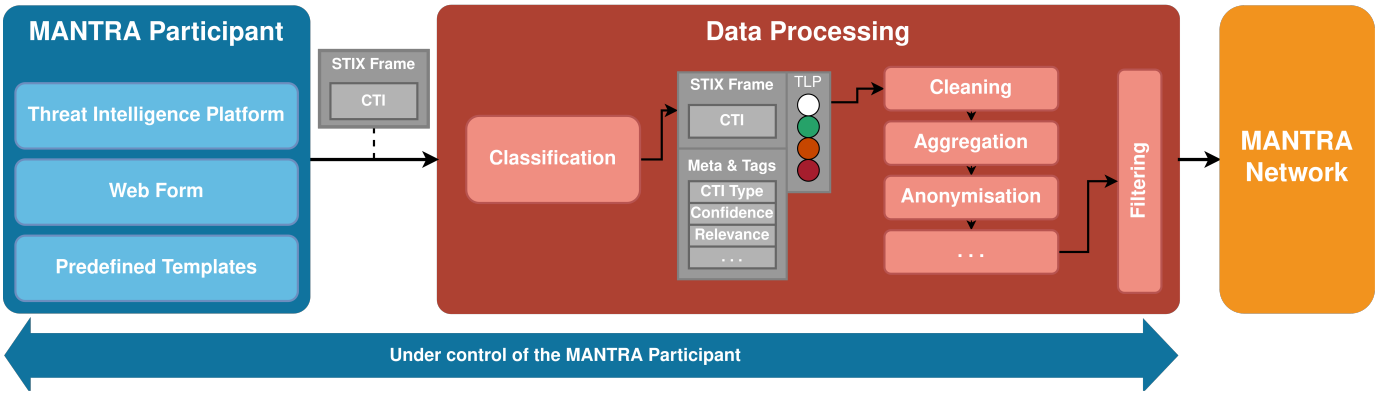

Figure 3: Data Integration of MANTRA.

detection, response, and attribution. In the following, ideas for the application of MANTRA are described, which will be researched and evaluated in further publications.

1) Prevention: In the domain of prevention, a practical application of MANTRA could be to enhance the functionalities of Security Information and Event Management (SIEM) systems within a Security Operations Center (SOC). Leveraging the global models derived from MANTRA's federated learning process, SIEM tools could be enhanced to detect potential threats more precisely, and promptly. This integration enables organizations to proactively identify and mitigate vulnerabilities or suspicious activity before they can be exploited. Additionally, the global models should provide a diversified threat landscape for the organization. By integrating the advanced threat intelligence provided by MANTRA, SOCs could refine security policies and alert thresholds, leading to more proactive and effective defense strategies.

2) Detection: In the realm of detection, MANTRA could use the CTI provided to increase the effectiveness of SIEM, EDR/Extended Detection and Response (XDR), and IDS. By integrating threat data from the MANTRA network, these tools are expected to detect potential security threats earlier and with better accuracy. Integrating MANTRA with SIEM systems could improve their ability to detect suspicious patterns and anomalies in network traffic and log data by enabling comparison with current global threat models. EDR/XDR platforms could benefit from MANTRA by improving the detection of malware and attempted attacks on endpoints based on the latest intelligence on threat actors and their Tactics, Techniques, and Procedures (TTP). IDS systems could strengthened by updating their detection signatures with MANTRA-powered data, enabling better detection of intrusion attempts and unusual activity.

3) Reaction: MANTRA could enhance reaction capabilities by facilitating better response management, achieved through integration with Security Orchestration, Automation and Response (SOAR) platforms and other incident management tools. By integrating MANTRA's CTI with these systems, security teams can create automated workflows for efficient and rapid response to detected threats. MANTRA can provide SOAR solutions with up-to-date and contextualized threat intelligence to accelerate security incident decision-making. Based on this information, SOAR platforms can prioritize specific alerts, orchestrate investigations, and initiate automated response actions based on the severity and relevance of the threat. Incident response management tools could also benefit from integration with MANTRA, as they gain access to detailed data on attack patterns and tactics. This not only improves the analysis and investigation of security incidents, but also helps to develop more effective response strategies and shorten response times.

4) Attribution: MANTRA is set to implement multifaceted attribution methodologies across its architecture, as elaborated in section VI-A. This involves conducting attribution during the initial data processing stage, which inherently enhances the overall quality of the local CTI repository within the TIP. By eliminating duplicates and enriching the dataset through correlation and supplementation of potentially missing information, MANTRA could improve the integrity and comprehensiveness of its CTI data.

To achieve this, MANTRA leverages contemporary frameworks and AI strategies. These are designed to identify and assimilate IoCs and TTPs, facilitating the establishment of connections between them. This advanced approach not only streamlines the attribution process but also ensures a more robust and actionable CTI repository, empowering stakeholders with enhanced capabilities for threat detection and response [12]–[15].

VII. Conclusion

Facing an intensifying cyber threat landscape, this paper underscores the essential role of CTI in safeguarding organizations across diverse sectors. It explores the MANTRA network as a solution to overcome obstacles, such as reputational risks, technical barriers, and data silos that impede effective CTI sharing. MANTRA promotes privacy-preserving CTI exchange, particularly benefiting organizations with limited resources, and underscores the need for improved sharing mechanisms to bolster collective cyber defense.

The objectives set forth by MANTRA address fragmented data silos through federated learning, tackle information over-

 ISBN: 978-1-68558-156-5

load with enhanced data analysis, and ensure CTI data's confidentiality, integrity, and availability via secure protocols and peer-to-peer communication. Additionally, it introduces socio-economic governance models to foster cross-organizational information sharing for heightened security and resilience.

MANTRA's architecture, featuring a protocol layer for secure data exchange, an application layer for CTI model processing, training, and use, and a federated learning layer for model aggregation, supports a collaborative ecosystem of training, aggregator, and operational peers. This structure allows for effective CTI utilization while ensuring data privacy and facilitating broad cybersecurity intelligence sharing.

Key to MANTRA's functionality are both internal and external data sources, which collectively provide a rich intelligence base for precise threat detection and robust organizational cybersecurity. Despite challenges in data privacy, quality, and compliance, MANTRA emphasizes the need for secure intelligence-sharing mechanisms.

Finally, the application layer's focus on "data integration and management" and "use cases" showcases MANTRA's capability to deliver actionable insights for prevention, detection, response, and attribution. By integrating with SIEM, EDR/XDR, IDS, and SOAR systems, MANTRA could enhance early threat detection, attack identification, and incident response. What is expected to have an impact on cybersecurity strategies and support in various security areas.

We plan to fully implement the MANTRA framework by creating the necessary infrastructure and technology. In this phase, the MANTRA concepts will be translated into functional modules for effective CTI operations. In addition to creating the foundational framework, we will enhance the application layer and focus on developing prevention, detection, response, and attribution tools that leverage global network models. These tools will be designed to provide organizations with robust capabilities to combat cyber threats amidst data proliferation by leveraging advanced analytics for actionable insights. To ensure the long-term impact of MANTRA, we will continuously evaluate and refine the framework and its applications to adapt them to evolving threats, incorporate the latest research, and optimize their efficiency and scalability.

ACKNOWLEDGEMENT

This study has been supported by funding from the Agentur für Innovation in der Cybersicherheit GmbH (Cyberagentur). The Agentur für Innovation in der Cybersicherheit GmbH did not interfere in the research process and its results.

REFERENCES

[1] "The state of it security in germany in 2023," Bundesamt für Sicherheit in der Informationstechnik (BSI), Threat Report, Nov. 2023.

[2] "Threat horizons: August 2023 threat horizons report," Google Cloud, Threat Report, Apr. 2023.

[3] "Global threat report," CrowdStrike, Threat Report, 2024.

[4] P. Fuxen *et al.*, "Mantra: A graph-based unified information aggregation foundation for enhancing cybersecurity management in critical infrastructures," in *Open Identity Summit 2023*, Bonn: Gesellschaft für Informatik e.V., 2023, pp. 123–128, ISBN: 978-3-88579-729-6. DOI: 10.18420/OID2023_10.

[5] C. Wagner, A. Dulaunoy, G. Wagener, and A. Iklody, "MISP: The Design and Implementation of a Collaborative Threat Intelligence Sharing Platform," in *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security*, ser. WISCS '16, New York, NY, USA: Association for Computing Machinery, Oct. 2016, pp. 49–56. DOI: 10.1145/2994539.2994542. (retrieved: 2024-03-07).

[6] J. Heo, Y. E. Gebremariam, H. Park, B. Kim, and I. You, "Study on Hybrid Cloud-based Cyber Threat Intelligence Sharing Model Requirements Analysis," in *Proceedings of the 2020 ACM International Conference on Intelligent Computing and Its Emerging Applications*, ser. ACM ICEA '20, New York, NY, USA: Association for Computing Machinery, Sep. 2021, pp. 1–6. DOI: 10.1145/3440943.3444737. (retrieved: 2024-03-06).

[7] X. Zhang, X. Miao, and M. Xue, "A Reputation-Based Approach Using Consortium Blockchain for Cyber Threat Intelligence Sharing," *Security and Communication Networks*, vol. 2022, e7760509, Aug. 2022, ISSN: 1939-0114. DOI: 10.1155/2022/7760509. (retrieved: 2024-03-05).

[8] K. Nguyen, S. Pal, Z. Jadidi, A. Dorri, and R. Jurdak, "A Blockchain-Enabled Incentivised Framework for Cyber Threat Intelligence Sharing in ICS," in *2022 IEEE International Conference on Pervasive Computing and Communications Workshops and Other Affiliated Events (PerCom Workshops)*, Mar. 2022, pp. 261–266. DOI: 10.1109/PerComWorkshops53856.2022.9767226. (retrieved: 2024-03-05).

[9] E. Bandara, X. Liang, P. Foytik, and S. Shetty, "Blockchain and Self-Sovereign Identity Empowered Cyber Threat Information Sharing Platform," in *2021 IEEE International Conference on Smart Computing (SMARTCOMP)*, Irvine, CA, USA: IEEE, Aug. 2021, pp. 258–263, ISBN: 978-1-66541-252-0. DOI: 10.1109/SMARTCOMP52413.2021.00057. (retrieved: 2024-03-07).

[10] E. Bandara, S. Shetty, R. Mukkamala, A. Rahaman, and X. Liang, "LUUNU — Blockchain, MISP, Model Cards and Federated Learning Enabled Cyber Threat Intelligence Sharing Platform," in *2022 Annual Modeling and Simulation Conference (ANNSIM)*, Jul. 2022, pp. 235–245. DOI: 10.23919/ANNSIM55834.2022.9859355. (retrieved: 2024-03-07).

[11] T. D. Wagner, K. Mahbub, E. Palomar, and A. E. Abdallah, "Cyber threat intelligence sharing: Survey and research directions," *Computers & Security*, vol. 87, p. 101 589, Nov. 2019, ISSN: 0167-4048. DOI: 10.1016/j.cose.2019.101589. (retrieved: 2024-03-01).

 ISBN: 978-1-68558-156-5

[12] U. Noor, Z. Anwar, T. Amjad, and K.-K. R. Choo, "A machine learning-based FinTech cyber threat attribution framework using high-level indicators of compromise," *Future Generation Computer Systems*, vol. 96, pp. 227–242, Jul. 2019, ISSN: 0167-739X. DOI: 10.1016/j.future.2019.02.013. (retrieved: 2024-03-07).

[13] M. Parmar and A. Domingo, "On the Use of Cyber Threat Intelligence (CTI) in Support of Developing the Commander's Understanding of the Adversary," in *MILCOM 2019 - 2019 IEEE Military Communications Conference (MILCOM)*, Nov. 2019, pp. 1–6. DOI: 10.1109/MILCOM47813.2019.9020852. (retrieved: 2024-03-07).

[14] M. Sahrom, S. R. Selamat, Y. Robiah, and A. Ariffin, "An Attribution of Cyberattack using Association Rule Mining (ARM)," *International Journal of Advanced Computer Science and Applications*, vol. 11, no. 2, pp. 352–358, Mar. 2020. DOI: 10.14569/IJACSA.2020.0110246.

[15] A. Nisioti, G. Loukas, A. Laszka, and E. Panaousis, "Data-Driven Decision Support for Optimizing Cyber Forensic Investigations," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 2397–2412, 2021, ISSN: 1556-6021. DOI: 10.1109/TIFS.2021.3054966. (retrieved: 2024-03-07).

 ISBN: 978-1-68558-156-5