



City Research Online

City, University of London Institutional Repository

Citation: Lakka, E., Hatzivasilis, G., Karagiannis, S., Alexopoulos, A., Athanatos, M., Ioannidis, S., Chatzimpyrros, M., Kalogiannis, G. & Spanoudakis, G. (2022). Incident Handling for Healthcare Organizations and Supply-Chains. 2022 IEEE Symposium on Computers and Communications (ISCC), doi: 10.1109/iscc55528.2022.9912965 ISSN 1530-1346 doi: 10.1109/iscc55528.2022.9912965

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/32111/>

Link to published version: <https://doi.org/10.1109/iscc55528.2022.9912965>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

Incident Handling for Healthcare Organizations and Supply-Chains

*Towards an automated detection, evaluation, and response to cyber-incidents

Eftychia Lakka
*Foundation for Research and
Technology-Hellas (FORTH)*
Heraklion, Greece
elakka@ics.forth.gr

George Hatzivasilis
Sphynx Technology Solutions AG
Zug, Switzerland
g.hatzivasilis@sphynx.ch

Stylianos Karagiannis
PDMFC, Ionian University
Portugal, Greece
stylianos.karagiannis@pdmfc.com

Andreas Alexopoulos
AEGIS IT RESEARCH GmbH
Braunschweig, Germany
andreas.alexopoulos@aegisresearch.eu

Manos Athanatos
*Foundation for Research and
Technology-Hellas (FORTH)*
Heraklion, Greece
athanat@ics.forth.gr

Sotiris Ioannidis
*Technical University of Crete,
Foundation for Research and
Technology-Hellas (FORTH)*
Crete, Greece
sotiris@ece.tuc.gr & sotiris@ics.forth.gr

Manolis Chatzimpyrros
Sphynx Technology Solutions AG
Zug, Switzerland
m.chatzimpyrros@sphynx.ch

Grigoris Kalogiannis
Sphynx Technology Solutions AG
Zug, Switzerland
g.kalogiannis@sphynx.ch

George Spanoudakis
Sphynx Technology Solutions AG
Zug, Switzerland
spanoudakis@sphynx.ch

Abstract—Healthcare ecosystems form a critical type of infrastructures that provide valuable services in today societies. However, the underlying sensitive information is also of interest of malicious entities around the globe, with the attack volume being continuously increasing. Safeguarding this complex computerized setting constitutes a major challenge for the involved organizations. This paper presents an incident handling system for healthcare organizations and their supply-chain. The proposed approach utilizes swarm intelligence in order to assess the current security posture in a continuous basis and respond to attacks in real-time. The overall solution is based on the related NIST 800.61 standard and implements the operations of i) preparation, ii) detection and analysis, iii) containment, eradication, and recovery, and iv) post-incident activity. The system is developed under the EU funded project AI4HEALTHSEC and is applied in the relevant healthcare pilots.

Index Terms—Healthcare sector, incident handling, incident response, response team, security, privacy

I. INTRODUCTION

The healthcare sector has undergone considerable changes in the past several years, primarily spurred by the adoption of new medical technology including IoT, Cloud Computing, and Big Data. As a result, healthcare organizations are increasingly affected by cybersecurity attacks. For that reason, appropriate response mechanisms are necessary when a security-related incident occurs [1], [2]. In fact, incident handling process is responsible to address the specific problem including the

incident response and management which is the protection of an organisation’s information by developing and implementing an incident response process (e.g. plans, defined roles, training, communications, management oversight) in order to quickly discover an attack and then effectively contain the damage, eradicate the attacker’s presence, and restore the integrity of the network and systems [3].

These incidents can have numerous devastating effects on healthcare organizations, from the inadvertent release of protected health information to disruptions in clinical care [4]. According to Ponemon institute, “healthcare organizations are in the cross hairs of cyber attackers” that grow increasingly frequent [5]. Except that, based on ENISA analysis [6], incidents handling is one of the major challenges in the healthcare security domain. This phenomenon can be explained by combining two factors: (i) the high value of healthcare facilities’ assets and (ii) the ease with which they can be compromised. Medical data is ten-twenty times more valuable than financial data for the reason that healthcare records can continue being exploited even after resolving the security breach which released them. At the same time, the healthcare industry is behind other industries in protecting its infrastructure and data.

Although that the majority of organizations implement security policies in their healthcare systems and/or infrastructures, there are incidents that can be neither anticipated nor avoided. In fact, security incidents root causes include, human errors,

natural phenomena, malicious actions (DDoS attack, MITM attacks, etc.) and system failures (including third party failure, i.e., hardware failure). It is worthwhile noting that system failures and human errors account equally for most of the incidents reported.

Additionally, deliberate human intervention to disrupt the workflow (i.e. malicious actions) also accounts significantly for security risk. On the other hand, the impact of natural phenomena is responsible for a small only percentage of the reported security incidents. It has to be noted that human factor may also relate to malicious actions, with the prospect of causing system holes through negligence or oversights, which could lead to system failures, hence the infrastructures can be vulnerable to possible attacks. Moreover, the incorrect security practices by personnel are included in human error that can result in security incidents; thus, apart from implementing cybersecurity measures, awareness raising, and training have a significant role in building a secure system. Therefore, healthcare organisations need to have an incident response capacity, in order to timely identify incidents and restore and reconstitute systems and services in a trusted manner. Indeed, there is a vital need for the development of a healthcare specific incident reporting, classification, and alerting mechanism in pan European level. International good practices could be consulted towards this direction.

Taking the above into consideration, IT security in healthcare systems, services and applications is positioned as a major concern due to the high privacy and confidentiality requirements of sensitive healthcare data and faces many security challenges which are highlighted below [6]:

- **Systems availability:** It is about continuous accessibility of critical health information by authorized professionals in order to ensure the best healthcare services. Systems availability may relate to physical systems function (e.g., networks, storage) and affect significantly the healthcare delivery
- **Lack of interoperability:** The high-level interoperability aims to guarantee that information of healthcare infrastructures is transmitted safely through individual information systems, health service institutions, healthcare providers and patients. It is important as many diverse systems and applications interconnected at various scales i.e. a medical device collecting clinical data can be linked in the same network that a computer uses to access the Internet.
- **Access control and authentication:** Authentication is the initial stage of the users' validation in order to determine their identity, which is necessary to ensure that they are authorized to access the system, which is a key-security feature in healthcare infrastructures [7]
- **Data integrity:** It purposes to ensure the quality and integrity of the data that are stored and exchanged for clinical and administrative purposes; a crucial part of healthcare systems for the reason that errors in personal or clinical data may affect a person's medical treatment, insurance or employability [8]

- **Network Security:** It is a fundamental challenge in securing healthcare infrastructures, especially when the system is network based (e.g., EHR/PHR, cross border eHealth)
- **Security expertise and awareness:** A critical factor that includes the adequate and sufficient organisational structure and especially the role of a security officer
- **Data loss:** It is mentioned to the protection of the data from loss; it is considered a very important part of the healthcare sector as a significant amount of vital, personal, and confidential data is stored in digital format

Motivated by the above, this work aims to provide a solution that utilizes novel decision-making approaches to efficiently and effectively monitor any possible threats to healthcare infrastructure, detect the anomaly stated and handle the corresponding incidents. The below-proposed approach is based on the development and the architecture that is carried out in AI4HEALTHSEC¹, a horizon 2020 project. The remainder of this paper is organized as follows: Section II provides the background of existing incident identification approaches, security event management and relevant frameworks and decision-making approaches. Section III details the methodology and the implementation details of the proposed approach. Section IV gives a brief overview of a preliminary application and evaluation of the AI4HEALTHSEC incident response. Finally, Section V features the concluding remarks and pointers to future work.

II. BACKGROUND & RELATED WORK

A. Existing Incident Identification approaches

Incident identification and handling strategy is a crucial task to mitigate risks to the confidentiality, integrity, and availability (CIA) of organisations' assets, as well as minimising loss (e.g., financial, reputational, and legal) [9]. The emerging paradigms of attacks, challenge the enterprise cybersecurity with sophisticated custom-built tools, unpredictable patterns of exploitation, and an increasing ability to adapt to cyber defences. Therefore, they have raised the needs of specific methodologies for identifying incidents, as well as frameworks specifically tailored for this task. The 2018 Annual Cybersecurity report [10], published by Cisco, indicates that organizations and enterprises have implemented cyber-awareness programmes, including seeking for outsourcing service to strengthen defences on cybersecurity incidents. In order to apply the best practises in preventing, handling, and managing all cybersecurity activities, it is first necessary to identify cybersecurity incidents. For this reason, many specific methodologies and frameworks for incident identification have been developed in the recent years.

Some consolidated procedures for security incident identification are defined in ISO/IEC 27035-1:2016 [11] and ISO/IEC 27035-2:2016 [12] standards. ISO/IEC 27035-1:2016 is the foundation of this multipart International Standard. It presents

¹AI4HEALTHSEC: A Dynamic and Self-Organized Artificial Swarm Intelligence Solution for Security and Privacy Threats in Healthcare ICT Infrastructures <https://www.ai4healthsec.eu/>

basic concepts and phases of information security incident management and combines these concepts with principles in a structured approach for detecting, reporting, assessing, and responding to incidents, while applying the lessons learnt.

The NIST cybersecurity framework (NISTCSF, 2018) [13], [14] offers a quantitative and measurable risk reduction guide on how organizations can incorporate cybersecurity activities as part of their risk management process, including incident identification procedures. The framework provides guidance that is useful and applicable to any organization, therefore offers a common, consistent, and comparable set of guidelines and practices.

Another approach for incident identification and management relies on Computer Security Incident Response Team (CSIRT), whose main function is to react in a timely fashion to cybersecurity threats. The CSIRT will typically be called into action by a notification or a triggered event, but can also be called into action by a relevant discovery while performing one of many passive services. The latter case may also include incident identification tasks. Frameworks for defining CSIRT services, roles, policies, standards, as well as procedures in case of incidents have been widely studied in literature [15].

The European Union Agency for Cybersecurity (ENISA) has provided a Good Practice Guide for Incident Management [16], which complements the existing set of ENISA guides that support CSIRT [17]. The guide describes good practices and provides practical information and guidelines for the management of network and information security incidents, with an emphasis on incident handling. In particular, it includes the identification of the incidents and its characteristics in the suggested procedures and handling process.

B. Security Event Management and Relevant Frameworks

Security events and evaluation approaches of the incidents are supported by Security Information and Event Management (SIEM) systems, providing a comprehensive view of the organization's security [18]. SIEM systems allow to evaluate and consolidate messages and alerts of individual components of an IT system. The main objective of SIEM systems is to provide a solid log management for security-related events [19]. There are two main types of approaches to identify an intrusion. The first is the Signature-Based Detection System (SIDS) and the second is the Anomaly-Based Intrusion Detection System (AIDS) [20]. Their combination is considered as well in hybrid settings. The first approach uses pattern-matching techniques by using signatures to identify the inconsistencies. The second approach uses machine learning and statistical-based knowledge to compare the network traffic. Therefore, any deviation from the expected behaviour is captured, and although it creates a lot of false positives, this approach can detect zero-day attacks and internal malicious activities that have not been reported in the past. There are two ways to proceed to the inspection, either by using a Host-Based Intrusion Detection System (HIDS) or by using Network-Based Intrusion Detection System (NIDS) [21]. The first approach analyses log files or syslog files that are coming

directly from a host OS, firewall, database, or logs from other services. However, this approach maintains limitations, mostly because since the HIDS must be deployed on every device. On the other side, NIDS extracts information by inspecting the network packets, and therefore, it can monitor all the devices on a certain part of the network. This approach also has limitations since some packets might be encrypted or the result could be highly influenced by the data bandwidth.

C. Decision-Making Approaches

With the huge amount of heterogeneous data gathered every day, decision-making becomes more and a more challenging task in the context of defining how to prevent, detect, and recover from information security breaches. A major part of decision-making involves the appropriate analysis of the data, which is the key success factor influencing the performance of decision makers, specifically the quality of their decisions. Decision-making has been studied in many different disciplines and various approaches have been proposed in the recent years to handle these problems and made great contributions to the decision-making. Specifically, Diesch et al. [22] proposed a comprehensive model of Management Success Factors (MSF) for information security decision makers, which were defined to help information security decision makers to apply the appropriate management decision. Another research in the area of road safety [23], Shah et al. proposed a methodology in order to analyse road safety performance. This was achieved by combining Data Envelopment Analysis (DEA) with the decision tree (DT) technique. Furthermore, Albeshri et al. [24] presented a multi-priority model as a theoretical model influenced with the network calculus and access control. Jalali et al. [25] proposed a simulation game to study the effectiveness of decision-makers to help managers in making proactive investment decisions for building cybersecurity capabilities. M'manga et al. [26] proposed in their paper a normative model providing systematic traceability to risk rationalization in order to facilitate the transparent understanding of security decision-making. Poletto et al. [27] believe that big data plays an essential role to support the decision-making, so the objective of our work aims to provide a theoretical approach about the elements necessary to apply the big data concept in the decision-making process.

D. AI4HEALTHSEC Holistic Incident Handling Approach

Considering all the above background analysis, many well-documented methodologies that describe the security incident response process, have already been proposed and applied in the healthcare domain. As mentioned before, the major aim of these strategies is to analyse a procedure for rapid detection of incidents, along with minimizing the effects, mitigating the causes, and restoring the affected resources. In fact, the popular Incident Handling recommendations pro-

posed by ENISA², NIST³, ISO/IEC 27035-1⁴, and CSIRT and CERT/CC⁵. As shown already above, all approaches share common characteristics, and it seems possible to derive a general methodology which would cover the entire procedure by the conjunction of the practices introduced by the various sources. After carefully investigating the relevant methodologies, this work proposes the solution that has been designed adapting the NIST methodology to the architecture scheme of AI4HEALTHSEC. It consists of four steps:

- Preparation (Step 1): It contains the steps that are taken before an incident occurs, such as training, writing incident response policies and procedures, and providing tools such as laptops with sniffing software, crossover cables, original OS media, removable drives, etc. In fact, preparation should include anything that may be required to handle an incident or will make incident response faster and more effective
- Detection and Analysis (Step 2): It is the phase in which events are analysed in order to determine whether these events might comprise a security incident (triage principles are included in this step)
- Containment, Eradication and Recovery (Step 3): The containment phase of incident response is the point at which the incident response team attempts to keep further damage from occurring as a result of the incident (i.e., taking a system off the network, isolating traffic, powering off the system, etc.). The eradication phase involves the process of understanding the cause of the incident, so that the system can be reliably cleaned and ultimately restored to operational status later in the recovery phase. The recovery phase involves cautiously restoring the system or systems to operational status
- Post-Incident Activity (Step 4): It includes the creation of a follow-up report, which each incident response team should evolve to reflect new threats, improved technology, and lessons learned aiming to reduce the probability of a similar incident happening again and to improve incident handling procedures.

III. METHODOLOGY AND IMPLEMENTATION DETAILS

A. Methodology – Overview

This section specifies the incident handling solution that has been proposed for AI4HEALTHSEC. In fact, it includes procedures for incident identification, security events evaluation and analysis, and decision-making. Figure 1 outlines the main tools/components that comprise the overall incident handling process, as they will be described in the following paragraphs.

An intelligent agent (IA) is anything that perceives its environment, takes actions autonomously to achieve goals (e.g., for security or safety), and may improve its performance with learning or may use knowledge. They may be

simple or complex. For AI4HEALTHSEC, two intelligent agent types are defined. The **primary agents** administrate a local subsystem of a healthcare setting. One **supervisor agent** per healthcare organization collects information from the underlying primary agents and manages the whole system. Then, the supervisory agents of different organizations can form a network and perform the swarm intelligence operations, controlling the whole ecosystem.

Swarm intelligence is defined as collective behaviour of a decentralized or self-organized system. These systems consist of numerous agents with limited intelligence interacting with each other based on simple principles. As mentioned before, the swarm intelligence for AI4HEALTHSEC is mainly performed by the supervisor agents of the involved organizations. Also, each supervisor agent orchestrates the operation of its underlying primary agents to accomplish internal collaborative tasks.

B. Implementation

In general, a Security Information and Event Management (SIEM), called Metadon, implements the main incident handling functionalities at the local/edge level. Metadon utilizes the Elasticsearch, Logstash, and Kibana (ELK) stack [28] in order to collect security incident -related data from the underlying system. Then, it implements a rule-based Artificial Intelligence (AI) operation, exploiting what is called Sigma-rules [29], in order to reason about the security posture. Metadon can also respond automatically to specific type of incidents and perform relevant actions, like i) send a message to the system administrator or ii) perform a script to tackle/mitigate the malicious side-effects. Apart from its own data gathering and evaluation capabilities, Metadon also receives information from the Cloud-based IDS. This module is based on Snort [30], which monitors the networking activity of the examined machine or network and is also utilizing rule-based reasoning to reveal potential security incidents. Snort implements more than 1900 specialized and technical rules that capture real cases of malicious/suspicious actions, while the user can also implement his/her own rules. This component is responsible to analyse information at a lower system level. The detected incidents are notified to the relevant Metadon instance. Asset Explorer and Data Fusion and Pattern Recognition module could also be utilized in order to support data unification and pre-processing prior to Machine Learning (ML) operations. All this operation resembles the main functionality of the conceptual Primary Agent, which is performing autonomous and self-adaptable operations at the local system.

At the backend of the system, Assurance Platform is deployed [31], [32], which collects information by the local Metadon agents (aka Particular Agents which have been deployed in the HCII) and perform AI processes, reasoning about the whole healthcare setting for an organization. The Assurance Platform is also utilizing the ELK stack, as well as customized Event Captors, in order to gather data and evidence. This feature is used for the collection of security-related information at the backend side, similar to Metadon.

²European Network and Information Security Agency

³National Institute of Standard and Technology

⁴ISO/IEC 27035-1:2016 Information technology — Security techniques — Information security incident management — Part 1

⁵CERT Coordination Center

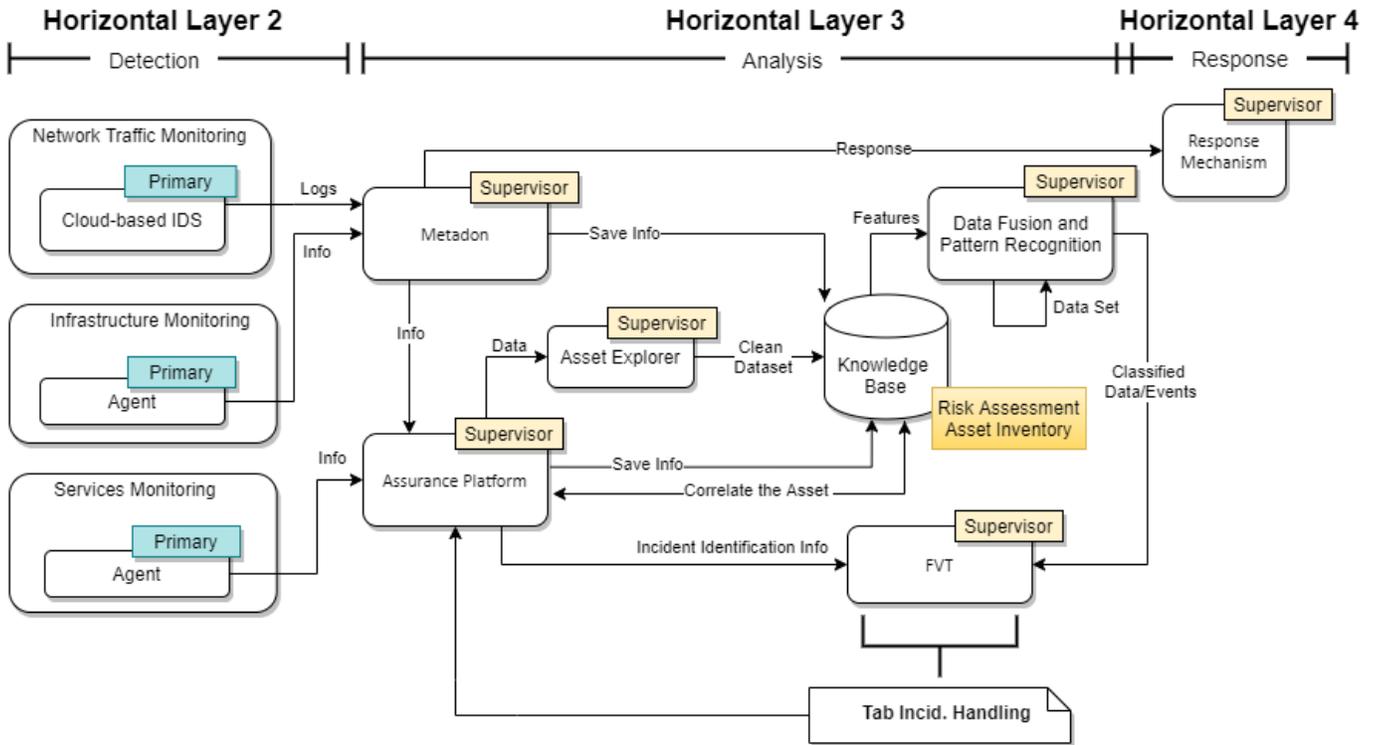


Fig. 1. Outline of the Incident Handling Process

The Assurance Platform also supports the use of Automated ML (AutoML) capabilities to further enhance the reasoning capabilities (e.g., perform User and Entity Behaviour Analytics (UEBA) to further analyse the access of users to the backend services). A Knowledge Base is established, maintaining all pieces of knowledge for an organization. This also includes the identification of the underlying assets based on the Common Platform Enumeration (CPE) [33]. Again, Asset Explorer and Data Fusion and Pattern Recognition module could be utilized for data transformation and unification purposes. Then, all information from the incident handling are accessible to the user via a Unified Dashboard. The asset explorer provides a Graphical User Interface (GUI), allowing the user to analyse the information, receive recommendations for potential response or restorative actions. The actions include, among others, human-driven decision-making for higher transparency of the AI functionality. The user is also enabled to execute the post-incident actions and share cyber threat intelligence (CTI) with other external stakeholders. All this operation resembles the main functionality of the conceptual Supervisory Agent, which administrates the system of a single healthcare organization and exchanges knowledge with the rest ecosystem.

IV. DEPLOYMENT AND DEMONSTRATION

A. System Implementation and Deployment

To evaluate the performance and functionalities of the described incident handling process, a server with Proxmox⁶

⁶<https://www.proxmox.com/>

is used to create Virtual Machines (VMs) for the different involved components. All the involved components are instantiated as different VMs running Linux as operating system inside Proxmox (Figure 2). Linux systems were deployed to execute test scenarios and to demonstrate the AI4HEALTHSEC. A subset of systems is used to execute the cyberattacks between the deployed hosts. The interaction generates numerous security events on the potential victim hosts/clients. Three types of AI4HEALTHSEC agents are deployed to distribute and store the events responsible for the following actions, a) Collect alerts related to network anomalies generated by the Cloud-based IDS, b) Collect customized alerts generated by other devices, c) Use ELK Beats to retrieve various events from the systems. The agents are the data sensors being setup to collect and distribute security events and alerts.

The agent that collects alerts from network anomalies (Agent 01) was installed in the network gateway to capture the network interaction between the deployed hosts and clients. The overall deployment is scalable and sufficient for real-time operation, being able to deploy extra hosts if required. The test environment has been used to integrate the AI4HEALTHSEC components and to automate the processes required to deploy the platform elsewhere.

B. Demonstration Example

In order to demonstrate the analysis and evaluation capabilities, two use cases have been developed respectively, presenting: i) the core incident evaluation, and ii) the support

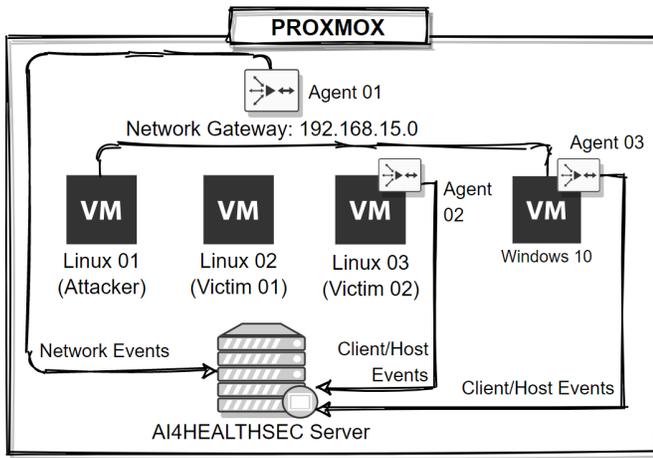


Fig. 2. Deployment and Demonstration Environment for AI4HEALTHSEC

of anomaly detection with the AutoML. Figure 3 illustrates the related UML diagram.

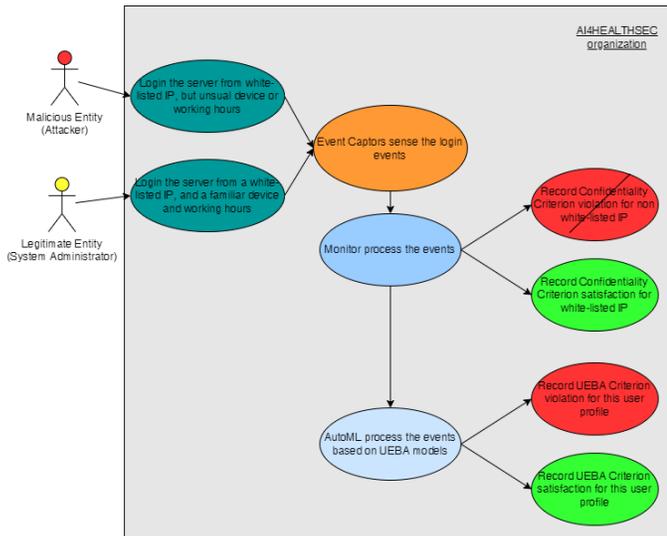


Fig. 3. UML Use Case Diagram – Incident Handling with Monitor and AutoML

1) *Core Reasoning with Predefined Ruleset*: A legitimate or a malicious actuator interacts with a monitoring asset. The sensing mechanisms (e.g., Beats, Event Captors, Snort, etc.) capture and process these events. The reasoning procedure analyzes the events and deduces wherever they are legitimate or malicious actions. The first case concerns the confidentiality property and the control that a specified user accesses a service from a set of white-listed IPs. A Filebeat or Auditbeat captures the user interaction with a service (or other resource). If a user access is recorded from a different IP, it can be due to some attack that manage to overcome the deployed defences (e.g., firewall) and infiltrate the system. For this example, we evaluate wherever the system administrator accesses the backend server from specific IPs. This is the server where the AI4HEALTHSEC Supervisory Agent is deployed and can

additionally contain piloting services as well. If the user accesses the system from a valid IP, then a success event is recorded. Otherwise, if the user accesses the system from a non-valid IP, then a violation event is recorded. For the testing of this use case, we emulate the legitimate and the malicious entity by performing two relevant events that will trigger the reasoning process. One event triggers the ruleset that captures the satisfaction of the confidentiality criterion, and the other event triggers the ruleset that discloses the violation of it.

2) *Reasoning Enhanced with ML*: Here again, a legitimate or a malicious actuator interacts with a monitoring asset. Here, the reasoning is also supported by the AutoML component. The system captures and process these events, and deduces wherever they are legitimate or malicious actions.

The second case concerns the confidentiality property and the anomaly detection based on User and Entity Behavior Analytics (UEBA). Here, we have trained a model to recognize a user based on his/her past interaction with the system. it is considered the protection SSH login service of the users to the backend. Usually, the users access the backend server from specific devices, locations, and/or working hours. Based on his/her routine, if a user logins the server from another region or country (i.e., Maldives, China), this could be a suspicious event. At initialization, the ML parses a set of testing logfiles and discloses the usual IPs, devices (e.g., based on the MAC details), working hours, and other pieces of information that are utilized by some test users, as well as other UEBA-related information. At runtime, the Monitor will examine every successful service login (as explained in the first use case) and request the AutoML module to check if the login action for the specific user is complying with the related UEBA profile.

For this example, we further evaluate wherever a successful system administrator login to the backend server (i.e., from white-listed IPs) is also complying with the specific user's UEBA profile. If the user login matches with his/her UEBA model, then a success event is recorded. Otherwise, if the user login does not match with his/her UEBA model, then a violation event is recorded.

For the testing of this use case, we emulate the legitimate and the malicious entity by performing two relevant events that will trigger the reasoning process. One event triggers the ruleset that captures the satisfaction of the UEBA criterion, and the other event triggers the ruleset that discloses the violation of it. Figure 4 shows the web interface that is available to the system administrator.

V. CONCLUSION AND FUTURE WORK

This work presented a solution of an incident handling approach for the event detection and analysis, collection of security-and privacy-related data, automated monitoring for potential events, and the analysis of the incidents via a user-driven investigation of the identified events from the lower levels of the Health Care Information Infrastructures. The strong point of our work lies in the fact that it investigates how to address one of the major challenges of the healthcare sector,

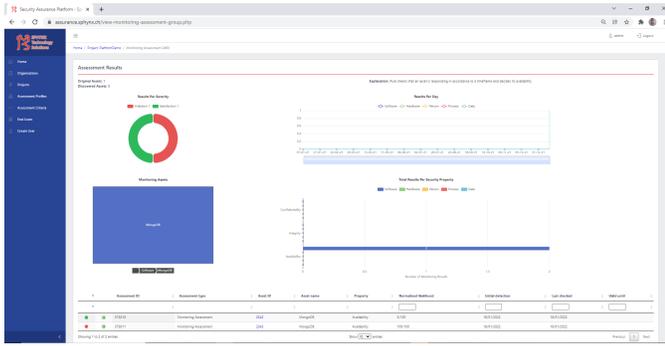


Fig. 4. Assurance Platform GUI – Evaluation of ongoing results

which is to be prepared for an unexpected attack. As reported above, based on ENISA's Threat Landscape report, "The attention paid by cybercriminals to health targets has increased considerably due to financial motives and the importance of the sector during the COVID-19 pandemic". This means, that now, even more than before, the health sector needs to be prepared against a variety of attacks. Hence, this paper provides insights on how to detect, handle, and recover from such unexpected incidents. As for future work, our proposal will be extended in order to offer the incident correlation, management of data from medical devices, data unification, and pre-processing prior to ML operations.

VI. ACKNOWLEDGMENTS

This work has received funding from the European Union's Horizon 2020 research and innovation programmes under grant agreements No. 883273 (AI4HEALTHSEC), No. 883275 (HEIR), No. 830927 (CONCORDIA), and No. 952644 (FISHY).

REFERENCES

- [1] ENISA, Sectoral/thematic threat analysis, ENISA Threat Landscape 2019-2020", [Online]. Available: <https://www.enisa.europa.eu/publications/sectoral-thematic-threat-analysis> [Accessed: 25-February-2022].
- [2] Hatzivasilis, George, et al., "Review of Security and Privacy for the Internet of Medical Things (IoMT)." International Workshop on Smart Circular Economy (SmaCE), Santorini Island, Greece, 30 May, 2019, IEEE, pp. 457-464.
- [3] ENISA "Strategies for incident response and cyber crisis cooperation", [Online]. Available: <https://www.enisa.europa.eu/publications/strategies-for-incident-response-and-cyber-crisis-cooperation> [Accessed: 25-February-2022].
- [4] Jalali, Mohammad S., et al. "EARS to cyber incidents in health care." Journal of the American Medical Informatics Association 26.1, pp: 81-90, 2019.
- [5] "Healthcare organizations are in the cross hairs of cyber attackers". [Online]. Available: <https://ponemonsullivanreport.com/2016/04/>. [Accessed: 25-February-2022].
- [6] ENISA, "Security and Resilience in eHealth Infrastructures and Services — ENISA." [Online]. Available: <https://www.enisa.europa.eu/publications/security-and-resilience-in-ehealth-infrastructures-and-services>. [Accessed: 25-February-2022].
- [7] Bell, Greg, and Michael Ebert. "Health care and cyber security: increasing threats require increased capabilities." KPMG, 2015.
- [8] Zender, A. "Ensuring Data Integrity In Health Information Exchange." AHIMA. American Health Information Management Association, 2012.

- [9] Ab Rahman, Nurul Hidayah, and Kim-Kwang Raymond Choo. "A survey of information security incident handling in the cloud." Computers & Security 49, pp: 81-90, 2015.
- [10] CISCO (2018), "2018 Annual Cyber Security report".
- [11] ISO/IEC (2016). "ISO/IEC 27035-1:2016" [Online]. Available: <https://www.iso.org/standard/60803.html>. [Accessed: 25-February-2022].
- [12] ISO/IEC (2016). "ISO/IEC 27035-2:2016" [Online]. Available: <https://www.iso.org/standard/62071.html>. [Accessed: 25-February-2022].
- [13] Barrett, Matthew P. "Framework for improving critical infrastructure cybersecurity." National Institute of Standards and Technology, Gaithersburg, MD, USA, Tech. Rep, 2018.
- [14] Scarfone, Karen, Tim Grance, and Kelly Masone. "Computer security incident handling guide." NIST Special Publication 800.61, 2008.
- [15] West-Brown, Molra J., et al. Handbook for computer security incident response teams (CSIRTs). CARNEGIE-MELLON UNIV PITTSBURGH PA SOFTWARE ENGINEERING INST, 2003.
- [16] ENISA (2010), "The European Union Agency for Cybersecurity (ENISA) have provided a Good Practice Guide for Incident Management", [Online]. Available: <https://www.enisa.europa.eu/publications/good-practice-guide-for-incident-management> [Accessed: 25-February-2022].
- [17] Sadoddin, Reza, and Ali Ghorbani. "Alert correlation survey: framework and techniques." Proceedings of the 2006 international conference on privacy, security and trust: bridge the gap between PST technologies and business services, 2006.
- [18] Aguirre, Idoia, and Sergio Alonso. "Improving the automation of security information management: A collaborative approach." IEEE Security & Privacy 10.1, 2011.
- [19] Nabil, Moukafih, et al. "SIEM selection criteria for an efficient contextual security." 2017 International Symposium on Networks, Computers and Communications (ISNCC). IEEE, 2017.
- [20] Ukwandu, Elochukwu, et al. "A review of cyber-ranges and test-beds: Current and future trends." Sensors 20.24, 2020.
- [21] Teixeira, Diogo, et al. "OSSEC IDS Extension to Improve Log Analysis and Override False Positive or Negative Detections." Journal of Sensor and Actuator Networks 8.3, 2019.
- [22] Diesch, Rainer, Matthias Pfaff, and Helmut Krömer. "A comprehensive model of information security factors for decision-makers." Computers & Security 92, 2020.
- [23] Shah, Syed Adnan Raheel, et al. "Road safety risk assessment: an analysis of transport policy and management for low-, middle-, and high-income Asian countries." Sustainability 10.2, 2018.
- [24] Albeshri, Aiiad, and Vijey Thayanathan. "Analytical techniques for decision making on information security for big data breaches." International Journal of Information Technology & Decision Making 17.02, 2018.
- [25] Jalali, Mohammad S., Michael Siegel, and Stuart Madnick. "Decision-making and biases in cybersecurity capability development: Evidence from a simulation game experiment." The Journal of Strategic Information Systems 28.1, 2019.
- [26] M'manga, Andrew, et al. "A normative decision-making model for cyber security." Information & Computer Security, 2019.
- [27] Poleto, Thiago, et al. "The full knowledge of big data in the integration of inter-organizational information: An approach focused on decision making." International Journal of Decision Support System Technology (IJDSST) 9.1, 2017.
- [28] ELK (2020), "Elastic stack", [Online]. Available: <https://www.elastic.co/> [Accessed: 25-February-2022].
- [29] SigmaHQ, "Sigma Rules", [Online]. Available: <https://github.com/SigmaHQ/sigma> [Accessed: 25-February-2022].
- [30] Snort (2021), "Snort IPS tool", [Online]. Available: <https://www.snort.org/> [Accessed: 25-February-2022].
- [31] Smyrlis, Michail, et al. "Cyra: A model-driven cyber range assurance platform." Applied Sciences 11.11, 2021.
- [32] Hatzivasilis, George, et al. "The THREAT-ARREST cyber ranges platform." IEEE CRST, IEEE, Virtual, Greece, 26 July, 2021, pp. 1-6.
- [33] CPE, "Common Platform Enumeration (MITRE)", [Online]. Available: <https://cpe.mitre.org/dictionary/> [Accessed: 25-February-2022].