International Journal of Information Security (2024) 23:1821–1838
https://doi.org/10.1007/s10207-024-00820-4

REGULAR CONTRIBUTION

RAMA: a risk assessment solution for healthcare organizations

Michail Smyrlis[1,2] · Evangelos Floros[3] · Ioannis Basdekis[1] · Dumitru-Bogdan Prelipcean[4,5,6] · Aristeidis Sotiropoulos[7] · Herve Debar[8] · Apostolis Zarras[9,10] · George Spanoudakis[1]

Published online: 1 March 2024

Abstract

Recent cyber-attacks targeting healthcare organizations underscore the growing prevalence of the sector as a prime target for malicious activities. As healthcare systems manage and store sensitive personal health information, the imperative for robust cyber security and privacy protocols becomes increasingly evident. Consequently, healthcare institutions are compelled to actively address the intricate cyber security risks inherent in their digital ecosystems. In response, we present RAMA, a risk assessment solution designed to evaluate the security status of cyber systems within critical domain, such as the healthcare one. By leveraging RAMA, both local stakeholders, such as the hospital's IT personnel, and global actors, including external parties, can assess their organization's cyber risk profile. Notably, RAMA goes beyond risk quantification; it facilitates a comparative analysis by enabling organizations to measure their performance against average aggregated mean scores, fostering a culture of continuous improvement in cyber security practices. The practical efficacy of RAMA is demonstrated through its deployment across four real-world healthcare IT infrastructures. This study not only underscores the significance of addressing cyber security risks within healthcare but also highlights the value of innovative solutions like RAMA in safeguarding sensitive health information and enhancing the sector's overall cyber resilience.

Keywords Cyber security · Healthcare · Risk Assessment · Software security · Information security

1 Introduction

Cybersecurity in healthcare is absolutely essential due to the critical nature of patient data and the potential consequences of breaches. In an increasingly digitised healthcare landscape, sensitive information, such as medical records, personal identifiers, and financial data, is stored and exchanged electronically. Protecting this data is paramount, as cyber-attacks can result in identity theft, fraud, or ransomware

This work has received funding from the European Union's Horizon 2020 research and innovation programmes under grant agreements No. 883275 (HEIR) and No. 965343 (RETENTION).

✉ Michail Smyrlis
smyrlis@sphynx.ch

Evangelos Floros
efloros@pagni.gr

Ioannis Basdekis
i.basdekis@sphynx.ch

Dumitru-Bogdan Prelipcean
bprelipcean@bitdefender.com

Aristeidis Sotiropoulos
a.sotiropoulos@aegisresearch.eu

Herve Debar
herve.debar@telecom-sudparis.eu

Apostolis Zarras
zarras@ics.forth.gr

George Spanoudakis
spanoudakis@sphynx.ch

1 SPHYNX Technology Solutions AG, Zug, CH, Switzerland

2 Department of Computer Science, City University of London, London, UK

3 University General Hospital of Heraklion, Crete, Hellas

4 Bitdefender, Bucharest, Romania

5 Alexandru Ioan Cuza University, Iaşi, Romania

6 Paris-Est Créteil University, Créteil, France

7 AEGIS IT Research, Braunschweig, Germany

8 SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France

9 Foundation for Research and Technology, Crete, Hellas

10 University of Piraeus, Piraeus, Greece

incidents that may compromise patient safety and trust in the healthcare system. Moreover, with the rise of telemedicine and connected medical devices, vulnerabilities in cybersecurity could lead to life-threatening disruptions in healthcare services. By investing in robust cybersecurity measures, healthcare providers can safeguard patient privacy, maintain the integrity of medical data, and ensure the delivery of high-quality and secure healthcare services.

Most EU countries invest enormous resources into acquiring the most updated e-health tools and applications to deliver effective and efficient healthcare services to their citizens [1]. These services include sharing health information with relative ease, aiming to boost the interaction between healthcare professionals and their patients [2].

Unfortunately, the healthcare industry is not only becoming a prime target for cybercriminals but also has the most expensive data breaches for the past 13 years [3]. This concerning trend has been further exacerbated during the COVID-19 pandemic, as reported by Muthuppalaniappan [4]. The underlying reasons for this worrisome phenomenon can be attributed to two key factors: First, healthcare facilities possess highly valuable assets [5]. Moreover, a patient's aggregated data can be likened to a valuable goldmine, offering a comprehensive biography of an individual that encompasses fundamental details, health patterns, family history, and even financial information [6]. Lastly, healthcare organisations are prone to easy compromise, as pointed out by Alzahrani [7]. Notably, Kumar emphasises that medical data holds greater value than financial data, as healthcare records can be exploited long after the initial security breach that exposed them [8]. This underscores the critical importance of reinforcing cybersecurity measures in the healthcare sector.

To tackle the aforementioned issues, security and risk management leaders needed to update their cybersecurity strategies to protect modern organisations' ever-expanding digital footprints from new and emerging threats. According to Gartner [9], due to the expansion of enterprise attack surfaces, organisations tend to utilise Digital Risk Protection Services (DRPS), External Attack Surface Management (EASM) technologies, and Cyber Asset Attack Surface Management (CAASM) to assist security experts in visualising internal and external business systems, as well as automating the detection of security coverage gaps. Moreover, digital supply chain attacks are becoming increasingly popular, as the realisation of such attacks can provide a high return on investment. Thus, new mitigation strategies encompass resilience-based thinking, requests for evidence of security measures and best practices, more deliberate risk-based vendor/partner segmentation and scoring, and efforts to stay ahead of forthcoming requirements.

Although the healthcare sector has been capitalising on digital advancements to improve patient outcomes and experiences substantially, poor security practices are still heavily used [10]. According to Coventry et al. [11], poor computer and user account security, remote access and home working, and lack of encryption are critical issues in the healthcare sector. These, in combination with the lack of up-to-date risk assessment techniques and security awareness, have led the healthcare sector to be one of the most impacted in terms of average data breach cost [12], as well as be the sector that faces the most significant influx of cyber-attacks, regarding both volume (69%) and complexity (67%) [13].

This study presents the Risk Assessment for Medical Applications (RAMA), a risk assessment solution used to estimate the attack surface and resilience of medical applications and systems by incorporating several critical issues. The proposed solution (comprising of a methodology, and its implementation) is further separated into the Local and Global RAMA Scores to address the need to calculate it locally (within a single healthcare organisation) and globally (within all the assessed healthcare organisations). The former incorporates several critical issues reported by a centralised client (see Sect. 3) and estimates the attack surface and resilience of the underlying medical devices per healthcare organisation. In contrast, the latter will serve as a global benchmark against which local RAMA scores will be compared. This will allow external stakeholders and/or healthcare practitioners to understand the security status of healthcare organisations worldwide. The proposed solution has been successfully deployed in four healthcare organisations across Europe (two in Greece, one in the UK, and one in Norway). The evaluation of our solution is available in Sect. 5.

In summary, we make the following main contributions:

- We provide state-of-the-art risk assessment tools that could leverage the cyber security of a critical infrastructures, such as healthcare organisations. The involved tools not only serve as standalone solutions for addressing specific issues in different aspects of a cyber system (e.g. network, operating system, applications etc.) but also integrate their outputs in a versatile format suitable for various roles. This format includes raw findings from the tools, the local RAMA score as an independent unit, and more, providing a comprehensive and adaptable approach to cybersecurity.
- We combine the output of these tools to create the local RAMA score (and corresponding metadata) that helps an organisation understand its security posture. By provided a unique score that integrates the findings of different tools, we allow end users to (a) gain a holistic understanding of their cybersecurity landscape, (b) streamline the interpretation of diverse tool outputs into a unified and actionable format, (c) facilitate more informed decision-making by presenting a consolidated assessment, (d)

enhance the efficiency of their cybersecurity strategy by leveraging comprehensive insights, and (e) foster a more proactive approach to addressing vulnerabilities and mitigating risks within their organization.

– We construct the global RAMA score that helps organisations to compare their security status with other healthcare organisations across Europe. Through this approach, we enable organisations within the same domain (healthcare) to (a) promote a collective effort towards elevating overall cybersecurity standards within the healthcare sector, (b) strengthen the resilience of healthcare organizations by fostering a community-driven approach to cybersecurity enhancement, (c) foster a culture of benchmarking, allowing organizations to understand where they stand in comparison to their peers, and (d) encourage a collaborative environment where best practices can be shared and adopted, which can be especially useful for organisations with smaller resources.

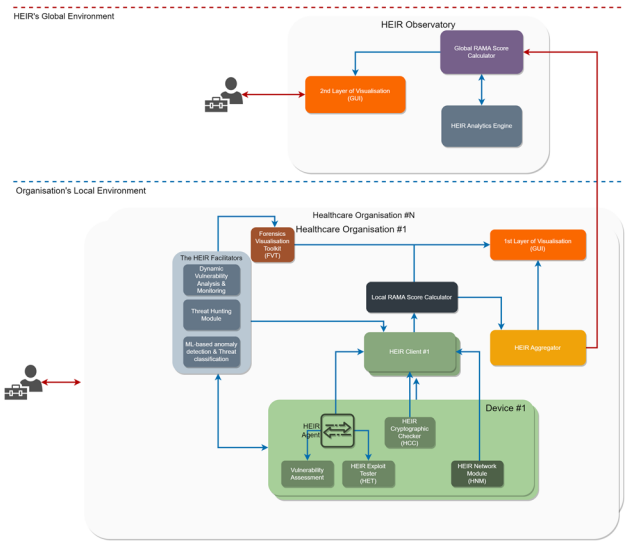

Fig. 1 High-level architecture of the proposed solution

2 Risk assessment

Risk assessment is a pivotal cornerstone for organisations, enabling a comprehensive understanding, effective control, and proactive mitigation of cyber risks. This critical process operates seamlessly across all three tiers of the risk management hierarchy: organisation level, mission/business process level, and information system level [14, 15].

The toolbox of risk assessment techniques offers a systematic approach to identifying, analysing, and evaluating security risks associated with diverse assets. In detail, a cyber security risk assessment navigates the intricate landscape of information assets, spanning hardware, systems, laptops, customer data, and intellectual property. This holistic evaluation extends to the intricate web of potential threats that could compromise these assets [16].

The paramount importance of aligning security controls with an organisation's unique risk profile underscores the utility of risk assessment tools. In this pursuit, two primary methods, qualitative and quantitative, stand as pillars of risk assessment. Qualitative risk assessment is rooted in gauging the likelihood of threats materializing, forming the bedrock of informed decision-making. Conversely, quantitative risk assessment entails a meticulous, formalised approach that quantifies the potential risks tied to the operation of an engineering process.

Amidst this intricate landscape, a range of risk assessment strategies come into play. These strategies address core organisational facets associated with the risk assessment process and encompass the definition of primary risk assessment protocols, supportive documentation, high-level process descriptions, and stakeholder involvement. While some resources offer a broad overview without delving into the intricacies of risk assessment algorithms [17, 18], others provide comprehensive methods and tools for risk assessment [19, 20].

In essence, risk assessment transcends being a mere procedural step, emerging as a fundamental paradigm that empowers organisations to navigate the intricate labyrinth of cyber threats with vigilance and strategic precision.

3 Threat hunting

Our proposed implementation for the risk assessment solution described in Sect. 4 is based on a layered architecture (see Fig. 1), communication between components, aggregation and transmission of the relevant information, and presentation through a web-based application (i.e., namely the 1st Layer GUI and the Observatory). More specifically, this implementation allows us to provide meaningful risk assessment, as we incorporate tools that are able to identify issues in (a) the data link layer (Network and Threat Detection modules), (b) the network layer (Network Module), (c) the Transport layer (Cryptographic Checker Module), (d) the Session Layer (SIEM Module), and (e) the Application Layer (Vulnerability Assessment, Exploit Tester, and SIEM modules).

This paper aims to showcase the Threat Hunting capabilities embedded in our approach. These capabilities encompass diverse methods and techniques designed for the detection of emerging threats or targeted attacks. Within our approach, the Threat Hunting Layer plays a pivotal role by consolidating various tools contributing to the identification and evaluation process.

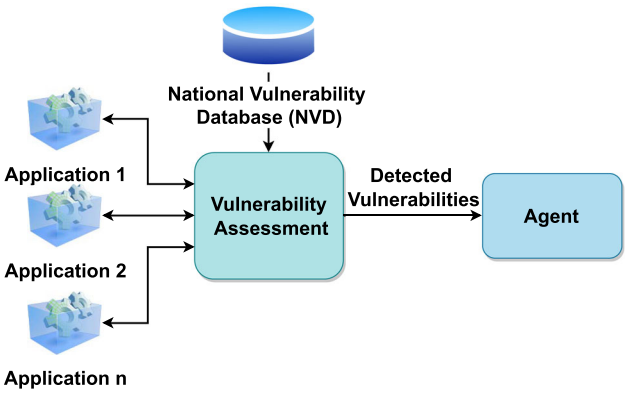

Fig. 2 Vulnerability Assessment high-level architecture

3.1 Vulnerability assessment module

A vulnerability assessment (VA) is a systematic review of security weaknesses in an information system. It determines whether the system is vulnerable to known vulnerabilities, rates their seriousness, and makes remedy or mitigation recommendations. In our approach, the VA module targets the operating system (OS) configurations and application information. These points of interest might place the endpoint and the entire medical system at risk for security breaches if they are improperly set or if the applications are outdated. The module interacts with the assessed system and collects the application's name and version. As Fig. 2 shows, VA exists within the context of an endpoint agent, i.e., a software deployed at an endpoint level that manages and collects the necessary data for further analysis.

VA's outcomes are first collected by the Local Correlation Component of the agent and then forwarded to the client through a Kafka message broker. The VA identifies vulnerabilities in the installed applications based on known CVEs as existing in NIST's National Vulnerability Database (NVD) to produce the results. One of the limitations of this module is that if software running on the platform is not indexed in the National Vulnerability Database (NVD), it can pose a challenge for conducting a comprehensive vulnerability assessment. The NVD is a well-known repository that provides a centralised source of vulnerability information for widely used software and systems. However, it may not cover less popular or custom-built applications that are unique to a specific organisation or industry. However, if the software vulnerability is also based on a misconfiguration from the operating system, the VA module will be able to identify the risk. Below is a snippet of the VA output:


```
{
"application_name": "Mozilla Firefox",
"cves": [
{
"cve": "CVE-2018-10892",
"description": "Mozilla developers reported
    memory safety bugs present in Firefox 80
    and Firefox ESR 78.2. Some of these bugs
    showed evidence of memory corruption and we
     presume that with enough effort some of
    these could have been exploited this
    vulnerability affects Firefox < 81,
    Thunderbird < 78.3, and Firefox ESR < 78.3
    .",
"publish_date": "2020-10-01 19:15:00.0",
"score": 67
}
],
"version": "69.0"
}
```


3.2 Exploit tester

The Exploit Tester (see Fig. 3) is responsible for assessing the attack surfaces for the operating system's configuration. Unlike the vulnerability assessment, which detects issues with the installed applications, the operating system's vulnerabilities and misconfigurations are discovered by the Exploit Tester. To accomplish that, it queries the OS's registry keys and configuration parameters as input, then outputs a list of the misconfigured security-related components, followed by suggestions and descriptions. An example of the Exploit Tester's output is as follows:


```
{
"availability": "None",
"confidentiality": "None",
"description": "Verifies the local group policy
     settings for User Configuration\\
    Administrative Templates\\System\\Ctrl+Alt+
    Del Options\\Remove Task Manager. When
    Remove Task Manager is enabled, the
    endpoint is vulnerable to security threats.
     Since Task Manager can list and terminate
    currently running processes, some malware
    may disable it to prevent themselves from
    being closed.",
"integrity": "High",
"name": "Task Manager",
"score": 25,
"triggered": true,
"type": "MisConfiguration"
}
```


Unlike the VA, the ET is focused on evaluating the attack surfaces based on operating system configurations. Some configurations can pose specific risks or might indicate some intrusion. For instance, the macro running enabled by default in Office applications increases significantly the attack surface, more specifically the risk of infection. Another example is disabling the task manager. It might be a legitimate action,

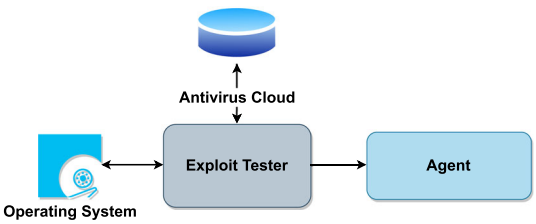

Fig. 3 Exploit Tester high-level architecture

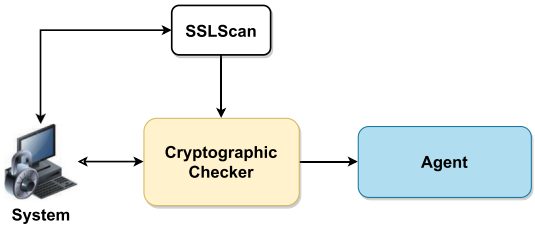

Fig. 4 Cryptographic Checker high-level architecture

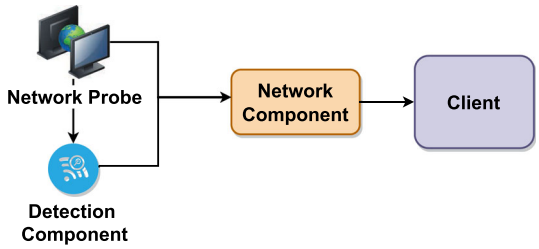

Fig. 5 Network Module high-level architecture

but in most cases, it is disabled by malware or an attacker to make the process visibility harder and mitigate as well.

3.3 Cryptographic Checker

The Cryptographic Checker (CC) detects outdated security protocols within the assessed cyber system's servers or targets the external servers that are service providers for the system itself (Fig. 4). CC is based on SSLScan [21] and can detect the used protocol and its version and the usage of vulnerable cryptographic implementations. The former can be cross-listed with the required ones per component (e.g., the latest TLS protocol). An example of the CC's output is as follows:


```
{
"description": "",
"host": [private],
"sniname": [private],
"port": "631",
"protocol": [
{
"type": "tls",
"version": "1.0",
"enabled": "1"
},
{
"type": "tls",
"version": "1.1",
"enabled": "1"
},
{
"type": "tls",
"version": "1.2",
"enabled": "1"
},
{
"type": "tls",
"version": "1.3",
"enabled": "1"
}
]
}
```


3.4 Network module

The Network Module monitors the assessed network traffic and provides security insights regarding any identified malicious activity. This module can analyze both the inbound and outbound network traffic of the system and detect (i) private information leaks, (ii) malicious content sent over the network, and (iii) ongoing attacks on the network. The output of this component provides connection information for the endpoints connected to the analyzed network and usage statistics that the anomaly detection component can use.

The high-level architecture of this component is depicted in Fig. 5. More specifically, the Network Probe intercepts the traffic for the configured Ethernet device and (sub)network. The Probe submits data to the detection module and directly to the network component (mainly for network telemetry). It comprises threat detection signatures and heuristic rules and evaluates the attack indications. It also exhibits patterns for network-wide information leaks. Lastly, as depicted in the high-level architecture, the Network component is the only component that communicates with the Detection one. An example of the Network Module's output is as follows:


```
{
"DestinationMAC": "[private]",
"DestinationPort": 50975,
"AlertType": "ATTACK",
"GMID": "984a2797-190b-4d28-a5b8-d97597a5bb11",
"Description": "Network Probe has prevented a
    suspicious DNS request to a public server
    that could contain private data. This is a
    potential data exfiltration marker. Data
    exfiltration is a form of a security breach
     that occurs when an individual's or
    company's data is copied, transferred, or
    retrieved from a computer or server without
     authorization.",
"DestinationIp": "[private]",
"SourceIp": "[private]",
"event_name": "detection",
"AlertName": "Exploit.DNS.ExfiltrationQuery",
"TimeCreated": 1637750665615,
"SourceMAC": "00:0c:29:a3:01:b7",
"SourcePort": 445
}
```


3.5 Threat detection module

The Threat detection module (TDM), much like the Network Module (NM), shares information with the client in a consistent manner, adhering to the same interpretational context. The TDM's capabilities extend to the meticulous scanning of files and processes situated within the confines of an endpoint machine. This process involves scrutinizing potential malicious content as it is executed. In essence, the TDM serves as an advanced guard, primed to identify and isolate any malicious activity that may transpire within an endpoint environment. In essence, the TDM functions as a sentinel, reinforcing the system's ability to identify and respond to

potential threats with heightened efficacy. The TDM is based on a lightweight approach of scanning technology from Bitdefender and is specially tailored in order to accommodate the low impact on healthcare environments but at the same time to provide a high rate of detection.

```
{
"ScannedObject": "C:/test/samples.",
"ObjectType": "File",
"AlertType": "Malware",
"event_name": "detection",
"AlertName": "Trojan.NG.Test.1",
"TimeCreated": 1670944872
}
```

As opposed to the NM that analyzes the network traffic and provides detection for malicious traffic and private data possible leaks, the TDM is a content-based scanner that is able to detect malicious files that are present on disk or executed files.

3.6 Security information and event management (SIEM)

The SIEM component feeds the interactive Forensics Module with various security-related data across all agent-installer endpoints. It is built on the open-source Wazuh [22], which offers a variety of security-related services that continuously monitor an IT infrastructure. The Wazuh Manager, where data is gathered, processed, indexed, and stored, receives events from lightweight agents that run on the monitored systems and collect all data. As a result, the server level is the only location where security intelligence and data analysis are carried out, ensuring that the resources required at the client level are kept to a minimum. Wazuh clients can be used with a variety of operating systems, such as Windows, Linux, Mac OS X, AIX, Solaris, and HP-UX. The events that the Wazuh agents report is the result of a variety of tasks, including (i) inventory of running processes and installed applications, (ii) collection of log and events data, (iii) monitoring of file and registry key integrity, (iii) monitoring of open ports and network configuration, and (iv) configuration assessment and policy monitoring.

The Wazuh server receives these events and processes them using a series of decoders and rules while employing threat intelligence to search for well-known Indicators Of Compromise (IOCs). All occurrences are given a severity level as a result of this analysis, allowing administrators to concentrate on pressing problems that must be solved. Additional delivery of this is made possible by way of tailored alerts sent to an Elastic Stack, which also offers a strong interface for data visualization and analysis address[] to its integration with Kibana.

Moreover, Wazuh can gather and combine OS-derived logs as well as logs from network devices like routers and firewalls. This can be done by either monitoring the log files directly or by transmitting logs via Rsyslog [23]. This could facilitate the collection of logs from medical equipment that needs to be monitored in hospital use-case situations. Furthermore, the Wazuh manager can communicate with web browsers, command-line tools like cURL, or other scripts or programs that can perform web requests address[] to the rich RESTful API that Wazuh provides. This, together with the RESTful APIs that ElasticSearch offers, feeds the Client that will then notify the Local RAMA calculator. An example of SIEM's output is as follows:

```
{
"description":"Windows Defender: ERROR: BAD
    INPUT DATA",
"severity": "12",
},
{
"description":"Short-time multiple Windows
    Defender error events",
"severity": "14"
}
```

3.7 Aggregator

The Aggregator is a component that collects the individual Local RAMA scores and corresponding metadata and forwards them to the Global RAMA Score calculator and the local hospital's environment. For the former, since data are transferred over the Internet, the Aggregator is responsible for anonymizing all the hospital-related information and submitting it to a TLS-enabled Kafka. Its task is to compute and report an aggregated score for all the involved organizations to the Global RAMA Score calculator and compute and report a weighted aggregated score per healthcare organization. The latter considers the severity per department, as communicated by the organization. In summary, the primary metadata that the Aggregator reports are the number of (i) critical events, (ii) benign/malicious findings, (iii) OS vulnerabilities, (iv) vulnerabilities at the application level, (v) misconfigurations, (vi) identified heartbleeds [24], (vii) attacks, and $(viii)$ exploits.

4 Risk assessment for medical applications

The Risk Assessment for Medical Applications (RAMA) is a comprehensive solution that empowers healthcare organisations to identify potential hazards and evaluate their potential for causing harm. In the proposed solution, scores are not merely numerical representations; they encapsulate a comprehensive evaluation of risks based on a range of factors specific to the cyber system. These factors may include but are not limited to the likelihood of occurrence, potential impact, historical data, and contextual relevance. Unlike some existing solutions that might adopt a one-size-fits-all approach, RAMA allows for the customisation of scoring

criteria (e.g. by including one or more of the aforementioned tools) to align with the unique characteristics and priorities of the assessed organisation. This flexibility ensures that the scores generated by RAMA are not only meaningful but also directly applicable to the organisation's specific risk landscape.

The aforementioned can be achieved through the RAMA calculator, the implementation of the proposed solution, which provides a score and metadata that can help the assessed healthcare organisation understand its security posture. The score can be calculated either for the organisation's local infrastructure (Local RAMA) or as act as a score (Global RAMA) against which the local RAMA scores of individual organisations are compared. It is important to clarify that the term "global" in this context indicates that the score is calculated from various healthcare organisations that already utilise the proposed solution, encompassing a wide-ranging perspective. Below, the Local and Global RAMA scores are presented.

4.1 Local RAMA score

The Local RAMA score is based on the weighted aggregation of the Vulnerability Assessment (VA), Exploit Tester (ET), Cryptographic Checker (CC), Network Module (NM), Threat Detection Module (TDM), and SIEM Module sub-scores. Each score is equally important, as it reveals potential issues in different layers of a cyber system, e.g., network, presentation and application, and, subsequently, provides a metric that would allow the end user to understand the security posture of its organisation better. The final local RAMA Score is a composite calculation of two sub-scores, the base and the temporal. The former acts as a static risk assessment score, i.e., it reveals the risk that is evaluated at a given time but not updated frequently and is based on the ET, CC, and VA sub-scores. On the contrary, the temporal score is a dynamic risk assessment metric based on the outcome of the NM, TDM, and SIEM modules. The temporal score is calculated based on a continuous process of identifying hazards and assessing risk. The main difference with the base part is that the temporal one measures the current state of the cyber system, meaning that the score can fluctuate more easily over time. The calculation of the Local RAMA Score is based on the formula below:

$$LRS = 0.7 * \text{Base}_{\text{score}} + 0.3 * \text{Temporal}_{\text{score}} \quad (1)$$

Vulnerability assessment sub-score The calculation of the VA score (normalised from 0 to 100) takes into account the overall number of recognised applications and the number of vulnerabilities per application, as shown below:

$$V_a = \sum_{i=1}^{n} A_i * V_s \quad (2)$$

where n is the total identified applications, A_i is the application's severity and V_s is the vulnerability score per application. The vulnerability score takes into account each vulnerability's severity (as supplied by the VA), as shown below.

$$V_s = \sum_{i=1}^{n} VSS_i \quad (3)$$

where n represents the overall number of vulnerabilities and VSS_i represents the severity-based vulnerability severity score. The severity per vulnerability is calculated as follows: $0-20=1$; $21-50=3$; $51-80=5$; $81-100=7$. The metadata for the VA score includes (i) the total number of vulnerabilities (across all applications), (ii) the total number of vulnerabilities (by application), as well as (iii) the top 10 identified vulnerabilities (per severity and frequency).

Exploit Tester Sub-Score. The formula for calculating the ET sub-score is a weighted average between the OS vulnerabilities and the identified misconfigurations. As the former is more severe, its weight is higher than the misconfiguration one. The formula is as follows:

$$\begin{aligned} ET_{score} = {} & 0.85 * \text{vulnerability}_{\text{score}} \\ & + 0.25 * \text{misconfiguration}_{\text{score}} \end{aligned} \quad (4)$$

The formula used to determine the vulnerability and misconfiguration scores (normalized from 0 to 100) takes into account both the triggered value and the impact of each recognized security property, particularly confidentiality, integrity, and availability (CIA). More specifically, the following weight is added whenever a certain vulnerability or configuration is exploited:

$$\text{Vulnerability}_{\text{score}} = \sum_{i=1}^{n} CIS_i + IIS_i + AIS_i \quad (5)$$

$$Misconfiguration_{score} = \sum_{i=1}^{n} CIS_i + IIS_i + AIS_i \quad (6)$$

where n is the total number of vulnerabilities/misconfigurations, CIS_i is the Confidentiality, IIS_i the Integrity, and AIS_i the Availability impact score. The score per impact is calculated based on the impact value: $none = 0$; $low = 2$; $medium = 7$; $high = 10$. The metadata for the ET sub-score includes (i) the number and percentage of malicious findings, (ii) the number of OS vulnerabilities, (iii) the number of misconfigurations, (iv) the number and percentage of benign findings, and (v) the ET vector, which depicts the

qualitative impact per security attribute, and the number and percentage of benign results (CIA).

Cryptographic checker sub-score The heartbleeds reported by the component are taken into account in the procedure for computing the CC sub-score (normalized from 0 to 100). The sub-score is raised by 10 if a certain TLS version is both enabled (enabled = 1 in the protocols) and vulnerable (vulnerable = 1). (as it is considered a major issue). The formula reads as follows:

$$CC_{score} = \sum_{i=1}^{n} H_i \tag{7}$$

where n is the total number of identified heartbleeds and H_i always equals 10. The number of identified vulnerable TLS protocols and their list are included in the metadata created for the CC.

Network module sub-score The alert type is taken into account in the calculation for computing the Network Module sub-score (normalized from 0 to 100), as this indicates the seriousness of the identified problem. The alert type may be (i)none, (ii)info, (iii)suspicious, (iv)malware, (v)attack, or (vi)exploit. This leads to the following formula:

$$NM_{\text{score}} = \sum_{i=1}^{n} NIS_i \tag{8}$$

where n is the total number of identified network issues (alerts or detection) and NIS_i is the network impact score. The latter is calculated as follows: $none = 0; info = 2; suspicious = 4; malware = 6; attack = 8; exploit = 10$. The metadata for this sub-score includes the number of exploits, attacks, and findings, the destination port and IP, the source port and IP, and a brief description per network issue.

Threat detection module sub-score The alert type is taken into account in the calculation for computing the Network Module sub-score (normalized from 0 to 100), as this indicates the seriousness of the identified problem. The alert type may be (i)none, (ii)info, (iii)suspicious, (iv)malware, (v)attack, or (vi)exploit. This leads to the following formula:

$$TDM_{score} = \sum_{i=1}^{n} TDMI_i \tag{9}$$

where n is the total number of malicious issues (alerts or detection) and $TDMI_i$ is the threat impact score. The latter, just like the network's component, is calculated as follows: $none = 0; info = 2; suspicious = 4; malware = 6; attack = 8; exploit = 10$. The metadata for this sub-score includes the number of exploits, attacks, and findings, the destination port and IP, the source port and IP, and a brief description per network issue.

SIEM sub-score The formula for calculating the SIEM sub-score (normalized from 0 to 100) takes the severity as reported through the SIEM component. Since the severity calculation is based on Wazuh's ruleset, no further reasoning is applied through the calculator. The SIEM's formula is as follows:

$$SIEM_{score} = \sum_{i=1}^{n} SIS_i \tag{10}$$

where n is the total number of identified issues and SIS_i is the SIEM impact score. The latter is calculated as follows and is based on Wazuh's rules classification (as denoted within the parentheses): $ignored(0) = 0; low(2-4) = 2; medium(5-8) = 3; high(9-12) = 5; critical(13-14) = 8$. The metadata for this sub-score includes the number of issues as well as the description as reported from the SIEM component.

4.2 Global RAMA score

The concept of the Local RAMA score was introduced in our solution as an effective way to communicate the overall security status of individual clinical sites to IT personnel, security experts, and other stakeholders. Similarly, the Global RAMA score serves as a benchmark for a group of clinical sites, enabling its comparison to local RAMA ones. Having said that, a fluctuation in the local part will always affect the Global part. The Global RAMA Score calculator acts as a mediator between the Local RAMA Score Calculator of a single healthcare organisation and the Observatory and creates a single Global Score that incorporates Local RAMA Scores between all the healthcare organisations that have a computed local RAMA score. This allows interested parties to either compare their Local RAMA score to the global one or to identify the status of attack surface and resilience from a more global perspective. The calculation of the Global RAMA score relies on input from the Aggregator and subsequently, the Local RAMA Score calculator. The Global RAMA Score calculator is deployed outside the premises of the hospital and communicates with the Aggregator. Communication takes place via a message broker, over a TLS-secured communication channel. Prior to the data transfer from the Aggregator to the Global RAMA Score calculator, the former anonymizes data to be sent (mostly metadata fed by the Local RAMA Score calculator) that might expose personal information from a specific hospital. The definition of the Global RAMA Score is a weighted sum of all the Local RAMA aggregated scores as depicted in the equation below.

$$GlobalRAMA_{score} = \sum_{i=1}^{n} LRA_i \quad (11)$$

where N is the number of the available Local RAMA aggregated scores and LRAi is the Local RAMA Aggregated score for clinical site i as provided through its local HEIR Aggregator. The Global Rama Score can also be translated in a qualitative form, as mentioned below:

- 100—None
- 80–99—Low
- 50–79—Medium
- 10–40—High
- 0–9—Critical

The Global RAMA score is accessible to interested parties through the Observatory. The Observatory is a web-based platform responsible to collect, analyse and present the results of all the deployed clients in order to provide global insights on the level of security in healthcare environments. The collected information is stored in the Observatory database and will be analysed by the Analytics Engine in order to produce statistics, historical analysis and trends. Furthermore, data is collected from the Aggregators deployed in each hospital. The Global RAMA Score Calculator consumes this data, generates the Global RAMA score, and provides relevant metadata. The results are presented in the Observatory GUI, which represents the 2nd Layer of visualizations. The Global RAMA Score also comes with metadata provided by the Aggregator, initially produced from the Local RAMA calculator and the Global RAMA calculator, such as the top 10, most severe and most frequent vulnerabilities in all the involved healthcare facilities. This fact enables hospitals to identify if the locally identified vulnerabilities are also present in other hospitals of the HEIR ecosystem. The Observatory is meant to be accessible to stakeholders, policymakers or legislators. It comprises intelligent knowledge-base and interactive visualisation tools and its focus is on depicting the landscape of cyber threats for electronic medical devices, detailed cybersecurity assurance statuses, and their evolution over time.

5 Real-world deployment and use cases

To evaluate the RAMA score, we deployed it in four hospitals (two in Greece, one in the UK, and one in Norway). To realize the setup, we created an ecosystem of servers and workstations. Specifically, we provided hospitals with a Virtual Machine (VM) as the central server to host all the components. Further, the hospitals provided additional VMs as workstations replicating workstations at various clinics. In addition, we utilized productive workstations from different hospital departments where the components are deployed. The workstations are distributed to the hospital's departments and belong to various clinics. The initial use cases we examined are described below.

5.1 Outdated software

Outdated software can pose a number of problems and risks in a hospital system. For one, it may no longer be supported by the manufacturer, which means that it may not receive important security updates or bug fixes. This can make it more vulnerable to security breaches, malware, and other online threats. To tackle this issue, the deployed security mechanisms could detect outdated software on the servers and workstations that belong to the hospital and inform the IT department of the issues and the actions needed. To do so, the administrators check the base part of the Local RAMA score and identify if any connected client (a clinic workstation) has malicious findings. The system alerts individuals that the client has an outdated version of specified software when they examine the vulnerability details for the particular client in this situation. Once the issue has been resolved, the client in concern has no more vulnerabilities, and the Local RAMA score goes up. This way, the hospital's security is hardened as no attackers will exploit outdated software vulnerabilities.

5.2 Threat detection

For the smooth operation of a hospital, a threat detection system that acts quickly and efficiently is crucial. Security mechanisms can identify risks to the hospital's servers and workstations, eliminate them, and notify the IT department of the problems and the necessary steps. Here, the stage is set with an external storage device and a piece of malware designed to execute privilege escalation and employ lateral movement strategies. Within this expanded context, the malicious software is transplanted from the external storage onto the target system and subsequently set into motion, all while operating under the unwitting engagement of the end user. The target of this use case is to demonstrate how the temporal part of the local RAMA score will be recalibrated, based on the findings of the TDM and SIEM modules.

5.3 Cryptographic protocols

The objective of this scenario is to showcase the functionality of a key component within our solution, namely the Cryptographic Checker (CC). The spotlight is on the CC as it undertakes a pivotal role in this showcase. More specifically, the CC is responsible to establish connections and meticulously evaluate the cryptographic capabilities of each

interconnected system that bolsters the IT infrastructure. Within this scenario, the CC assumes the responsibility of assessing the vulnerability of devices or servers to potential cryptographic attacks. Once this assessment is conducted, it seamlessly recalculates the base part of the local RAMA score, taking into account the newfound insights. These recalibrations play a crucial role in offering a comprehensive overview of the overall security status. The end result is an enriched RAMA score, equipped with a finely-tuned accuracy reflecting the cryptographic health of the IT infrastructure.

6 Results and discussion

To evaluate the efficacy of our solution, and more specifically how the calculated RAMA score, along with the risk assessment tools integral to the computation of our proposed solution, will influence the actions of a healthcare organisation, we built upon the use cases presented in Sect. 5. More specifically, in this section, we will delve into an in-depth examination of the responsive actions taken by the evaluated healthcare organisation in response to variations in the RAMA score.

Time to detect an attack The time it takes to detect an attack varies depending on many factors. Some of these include the attack type, the efficiency of the systems in place to detect the attack, and the overall level of security the target employs. More specifically, in our study, the application layer's risks, attack surfaces, and vulnerabilities can be detected as soon as their definition is available in the databases. The evaluation is made periodically as a scanning task, and the periodicity can be configured by the administrator depending on the module and focus. Considering this, the detection time can be computed as the delta time between definition availability and the next task time. For detecting attacks, private information leaks, and malware we provide real-time availability of detection since, for the Network, Threat Detection, and SIEM modules, the traffic is continuously monitored. Of course, the latest definition depends on the update process of the module, but there is no delta time for detection from the moment of the latest availability. So in case of an attack, detection will be provided in real-time from the moment the inbound malicious content is transmitted over the inbound traffic or the leaks are submitted in the outbound traffic. This is possible due to the architecture of the NM by continuously extraction of traffic by the network probe and submitting to the detection component that is able the identify the problematic traffic on a stream-based paradigm.

Figures 6, 7, and 8 shows the evolution of the RAMA score, metrics and findings of the deployed risk assessment tools for a healthcare organisation for the past 9 months. Since this figure presents a more concentrated view of the

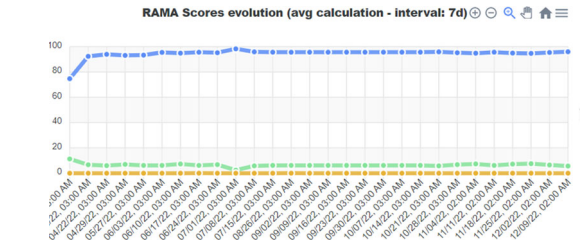

Fig. 6 Local RAMA fluctuation over time

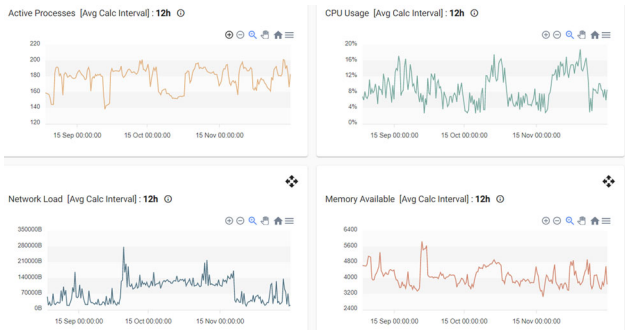

Fig. 7 Performance metrics over time

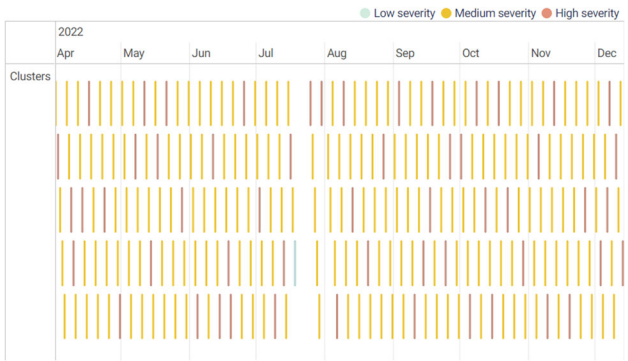

Fig. 8 SIEM tool findings over time

results so far, below we present the results of the execution of the use cases described in Sect. 5 that presents (a) how our solution contributed to the security of the examined organisation, and (b) how the different integrated tools contributed to the different parts of the calculation of the RAMA score.

6.1 Identification of outdated software

In this scenario, the hospital's administrator initiates the process by inspecting the local RAMA score, effectively highlighting the local environment's cybersecurity status. During this inspection, the administrator discerns a noteworthy revelation-specifically, a connected client, in this instance, a workstation situated within a clinic, has brought to light certain vulnerabilities. Prompted by this revelation, the administrator seamlessly progresses towards a more in-depth analysis. This involves a meticulous exploration of the

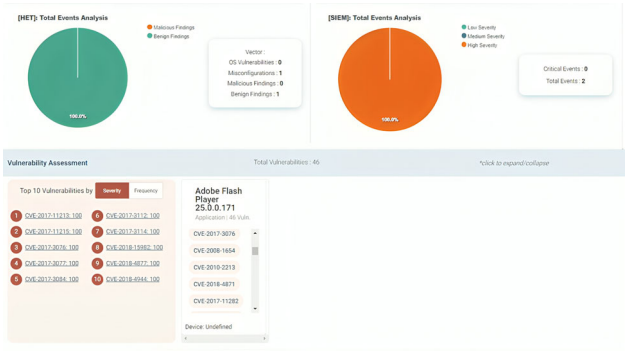

Fig. 9 Detected outdated software and constructed metadata

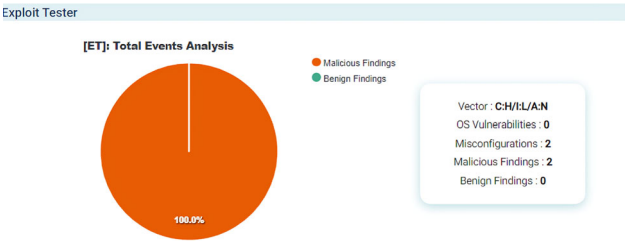

Fig. 10 Detected misconfiguration from the Exploit Tester

Fig. 11 Cryptographic Checker's report (JSON format), detecting the Heartbleed[2]

metadata attributed to the identified client, a task seamlessly facilitated by HEIR's integrated Vulnerability Assessment feature. The metadata constructed through the calculation of the base part of the Local RAMA Score (see Fig. 9) provides a wealth of valuable information, allowing the administrator to pinpoint a specific concern: the client in question currently operates on an outdated version of the Mozilla Firefox browser.

Besides the outdated software, we also have deployed the Exploit Tester which is able to detect operating system configurations that may pose a security risk. Such configurations are also reported to the RAMA score to indicate the possible risks. The misconfigurations are based on the most common features of operating systems that are exploited by attackers and malware to obtain initial compromise or further collect data. For instance, in this use case, we considered the default enablement of macro in Office applications. Even if this is disabled by default (with a clean installation) when enabled it increases the infection risk significantly. Let us consider that a VBA script or a macro is used for processing medical data or just for scheduling purposes. With everyday use, an administrator might be tempted to enable this by default. The ET is able to detect and report this (see Fig. 10) as a bad configuration from a security perspective.

In essence, this scenario illustrates how the administrator, empowered by the base part of the Local RAMA score, undertakes a methodical journey of assessment. By first assessing the local RAMA score and subsequently delving into the associated metadata, the administrator is adeptly equipped to identify and address the presence of outdated software-a pivotal step in upholding robust cybersecurity within the healthcare organization.

6.2 Cryptographic protocol issue detection

The primary objective of the depicted use case revolves around the timely detection and resolution of devices or servers susceptible to cryptographic attacks. In this scenario, the central actor is the Backend System Administrator. Their pivotal task is to identify vulnerabilities and subsequently utilise the insights provided by the Cryptographic Checker to rectify the security weaknesses. The journey commences with the Backend System Administrator accessing the 1st Layer GUI. Their primary focus is to check the local RAMA score for any significant reduction.

[1] Should a substantial reduction in the RAMA score be identified, the Administrator, still operating within the 1st Layer GUI, proceeds to uncover the source of concern. Through this graphical interface, they pinpoint a specific client that has raised a vulnerability flag. The Administrator then proceeds to delve into the details of this vulnerability. Within this investigative phase, the Administrator seamlessly navigates through the CC reports, extracting crucial information (see Figs. 11 and 12). Among the valuable insights garnered are the impacted devices, along with a clear understanding of the severity (as reported through the CC module).

[1] https://heartbleed.com/.

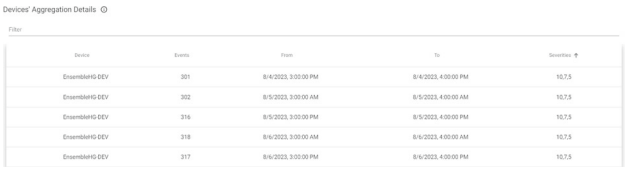

Fig. 12 Cryptographic Checker's metadata

In a determined effort to bolster security, the Administrator takes action beyond the proposed environment. Depending on whether the vulnerability stems from the OpenSSL library or the application utilising the vulnerable port, they undertake the necessary updates. This proactive measure serves to mitigate the identified vulnerability. Following the implementation of these updates, the Administrator circles back to the local RAMA score assessment. They await the occurrence of a fresh CC scan, effectively marking the completion of the remediation process. The ultimate confirmation lies in the RAMA score, which, upon reassessment, should reflect the resolution of the issue and the absence of any reported vulnerabilities. In essence, this use case showcases a comprehensive and dynamic cycle of vulnerability detection and resolution. The role of the Backend System Administrator emerges as pivotal in maintaining the robustness of the healthcare ecosystem's cybersecurity.

6.3 Threat detection

In this depicted scenario, we encounter a compelling interplay between an external storage device and insidious malware that adeptly employs privilege escalation and lateral movement tactics. The malicious application is copied from the external storage and unwittingly executed by an unsuspecting end-user. Our threat detection module steps forward as the vigilant guardian. More specifically, it is the malware, that promptly communicates this discovery to the HEIR platform. Meanwhile, our Security Information and Event Management (SIEM) component assumes its role as a sentinel, attuned to every critical event unfurling within the digital domain. This vigilance extends to all entities, whether users, attackers, or processes, casting a watchful eye on their every move. More specifically, the malware—after its execution—modifies the user groups to obtain some privilege escalation. Additionally, it tries several predefined logins to obtain a lateral movement. As the logins do not work, they will generate failed login events that will be considered an indicator of a compromise, from a SIEM point of view (see Fig. 13). This type of events, depending on their severity will be submitted to the HEIR platform to be visible and the RAMA score will be decreased. The SIEM and Threat Detection Module work in a complementary way.

Let us presume that the detection for the malware was not available, the SIEM high-severity events would be reported

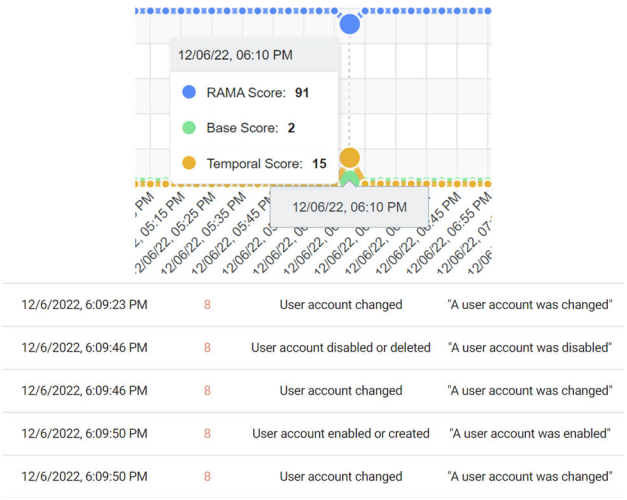

Fig. 13 Failed logins captured by the SIEM component

providing a clear indication that an attack or malicious action is active in the indicated endpoint. On the other hand, the reports from Threat Detection Module to the SIEM can also increase the importance of maybe not important action but they may provide context into the administrator analysis. In conclusion, this scenario demonstrates how two of our components work together to increase the temporal part of the local RAMA score in case a malicious action happens.

7 Related work

In the subsequent sections, we expound upon the pertinent research in the context of this paper. Our examination of the field is organised into three primary categories: (i) real-time threat monitoring, (ii) threat intelligence, and (iii) risk assessment. Elaboration on these categories is provided as follows. In our work, the related studies serve as a valuable reference point, that allowed us to: (a) draw upon the methodologies employed in the related work to design and structure our research approach (e.g. TVRA, CVSS), as, by understanding how previous studies addressed similar challenges guided our methodological decisions, (b) create concepts and theoretical foundations explored in the related work contribute to the development of our conceptual framework, and (c) emphasised how our work extends beyond the existing literature by introducing innovative approaches, addressing limitations identified in prior studies, or presenting novel insights that contribute to the advancement of the healthcare domain.

Real-time threat monitoring The categorisation of vast datasets at high velocities employs three primary distributed processing platforms: Apache Spark [25], Apache Storm [26], and Apache Flink [27]. The fundamental distinc-

tion between these platforms lies in Spark's execution of batch processing, while Storm and Flink excel in native flow processing. Derived from the Apache Spark Platform, two notable projects are Apache Spot [28] and Hogzilla [29]. Apache Spot leverages telemetry and machine learning techniques to scrutinize data packets for threat detection. Conversely, the Hogzilla tool supports Snort, SFlows, GrayLog, Apache Spark, HBase, and libnDPI, facilitating network anomaly detection. Hogzilla also enables the visualisation of network traffic, packet capture through Snort, and feature extraction via deep packet inspection.

Real-time threat monitoring encompasses the continuous surveillance of digital systems and networks, enabling prompt identification and response to potential threats as they emerge [30, 31]. Given the escalating frequency and intricacy of cyber-attacks, organisations prioritise investments in real-time threat monitoring solutions. Several key tools have garnered well-established reputations by effectively promoting their cybersecurity products. This has led to high levels of customer satisfaction and trust. Among these tools are Splunk's Enterprise Security [32], IBM's QRadar [33], and Palo Alto Networks' WildFire [34]. These tools facilitate real-time security monitoring, threat detection, log analysis, incident response automation, and data correlation from diverse sources.

Concurrently, the scientific community has introduced various methodologies within the real-time threat monitoring domain. For instance, Guimarães et al. [35] proposed TeMIA-NT, an innovative real-time flow analysis system employing parallel flow processing. Their devised system leverages data frames and a structured streaming engine, enabling real-time threat detection and swift attack response. In a similar vein, Krishnan et al. [36] formulated a comprehensive threat monitoring and security framework for multi-access edge computing (MEC) infrastructure. Among their contributions, the authors implemented anomaly detection, intelligent anti-DDoS applications, and a first-level mitigation mechanism in the data plane. This tactical deployment significantly diminishes the controller's load, leading to expedited attack detection and response times. In a separate endeavour, Cui et al. [37] introduced SD-Anti-DDoS, a system meticulously comparing diverse attack triggers utilized by detection mechanisms. Furthermore, they demonstrated the implementation of Botnet detection/traceback functions for SDN-enabled data centres, adding a layer of robust security. Kalkan et al. [38] delivered an exhaustive overview of filter-based security mechanisms aimed at combating DDoS attacks. Their review provides valuable insights into this critical aspect of threat mitigation. Lastly, Hsieh et al. [39] have made a noteworthy contribution by proposing a deep-learning-based classification technique for DDoS mitigation. Their work was showcased within the Apache Spark analytics platform, underscoring its practical application and potential significance.

Threat intelligence Threat intelligence can be described as the collection, analysis, and enrichment of threat information to deliver the necessary context to assist decision-making [40]. Over the years, researchers have proposed to utilise system logs for forensic analysis; causality analysis is vital in recognising the root causes [41]. There also attempts to minimise the dependency explosion problem by conducting fine-grained causality analysis [42], prioritising dependencies [43], and reducing data size [44].

Robertson et al. [45] proposed a comprehensive system composed of a crawler, parser, and classifier to pinpoint sites where security analysts can amass valuable information. Additionally, they devised a game theory-based framework to simulate the interactions between attackers and defenders, transforming the cyber threat intelligence process into a security game. This intricate framework incorporates historical attacks and the expertise of security professionals. In parallel, Tounsi et al. [46] categorized existing threat intelligence into three essential types: strategic, operational, and tactical. This classification offers a structured framework for understanding the various dimensions of threat intelligence.

Amid the burgeoning domain of Artificial Intelligence, Ibrahim et al. [47] embarked on a succinct exploration of how AI and machine learning methodologies can be harnessed to amplify the efficacy of threat intelligence, thereby thwarting data breaches. Concurrently, Rahman et al. [48] engaged in an in-depth discourse encompassing multiple facets of ML and Natural Language Processing, specifically concerning the automatic extraction of threat intelligence from textual descriptions. Their comprehensive discussion underscores the technological underpinnings of this critical process.

Recognising the pivotal role of threat intelligence utilisation, Wagner et al. [49] meticulously examined the state-of-the-art strategies for sharing threat intelligence. They delved into technical and non-technical challenges in automating the sharing process, providing a well-rounded perspective. Abu et al. [50] conducted a comprehensive survey, offering a panoramic overview of threat intelligence. Their work encapsulates this vital domain's definition, challenges, and key issues. Meanwhile, Ramsdale et al. [51] conducted a comparative analysis, summarizing the current landscape of formats and languages employed for sharing cyber threat intelligence. They further scrutinised a sample of cyber threat intelligence feeds, shedding light on the data they contain and the intricate challenges associated with their aggregation and dissemination.

In a parallel trajectory, Poirot [52] emerges as a sophisticated system for threat hunting. It adeptly unveils the aligned system provenance sub-graph within an input query graph. Conflict of interest [53] harnesses statistical attributes

extracted from NetFlow logs to identify botnet Command and Control (C&C) channels. The discerning use of DNS logs has also demonstrated efficacy in detecting malicious domains [54, 55]. Further amplifying the arsenal of threat detection, Hercule [56] employs community detection techniques, effectively piecing together attacks by correlating logs from diverse origins. Lastly, a cohort of research endeavours capitalises on system audit logs to execute forensic analysis and reconstruct intricate attack scenarios [43, 57, 58]. These works contribute to the arsenal of techniques to enhance our understanding of cybersecurity threats.

Risk assessment Cyber risk assessment has become paramount for organisations in recent years, driven by the escalating prevalence of cyber-attacks and data breaches. The scholarly discourse on this subject has surged substantially [59, 60], encompassing a myriad of studies that delve into various facets. These investigations span a wide spectrum, encompassing examinations of its impact on diverse organisational types [61–65], meticulous analyses of challenges and implementation constraints [66, 67], and the formulation of efficacious strategies to tackle cyber risks head-on [68, 69].

Within this expansive realm, a focal point of research has illuminated the pivotal role played by organisational culture in cyber risk management. The literature review reveal a recurring theme—organisations that foster robust risk management cultures are more adept at navigating the complex terrain of cyber risks [70, 71]. These entities promote a profound comprehension of their cyber risks and implement well-defined policies and protocols to confront them effectively [72]. Moreover, they often allocate dedicated teams or individuals to oversee cyber risk management, exhibiting proactive measures in implementing safeguards and vigilant monitoring to thwart potential threats.

Scholarly exploration has also delved into the challenges and limitations tied to the implementation of robust cyber risk assessment techniques. Among these hurdles, the rapid pace of technological evolution stands out as a significant impediment, rendering it a Herculean task for organisations to stay abreast of the latest threats and vulnerabilities [70]. Moreover, financial limitations and personnel shortages often constrain an organisation's capacity to establish and maintain resilient risk management protocols [73]. Lastly, Ganji et al. [74] reveals a notable absence of a comprehensive framework assisting organisations in both designing and adhering to the ISO/IEC 27001 standard, encompassing all 22 requirements. Its findings indicate that few studies have delved into the Information Security Management System (ISMS) at a proficient level, meeting at least half of the standard's requirements. Notably, no study was identified that covers all aspects of the ISMS.

Notwithstanding these challenges, researchers have highlighted a repertoire of effective strategies poised to navigate the complex landscape of cyber risks [75–77]. One potent approach involves the adoption of comprehensive risk management frameworks, exemplified by NIST's cybersecurity framework. This blueprint offers guidance, steering organisations toward identifying, assessing, and mitigating cyber risks [78]. Another relevant methodology is ETSI's Threat, Vulnerability, Risk Analysis (TVRA) [79]. TVRA is positioned as a method to identify risk to the system based upon the product of the likelihood of an attack, and the impact that such an attack will have on the system. Additional strategies encompass the deployment of robust security controls like firewalls and intrusion detection systems, the regular conduct of risk assessments, and consistent engagement in employee training and awareness programs [80].

Overall, the literature underscores the indelible significance of robust cyber risk assessment in today's digitised milieu [81, 82]. Through the astute deployment of resilient risk assessment frameworks and strategies, organisations fortify their defence against the perils of cyber-attacks and data breaches, safeguarding critical assets and ensuring uninterrupted operations. Distinguished by its specific focus on the healthcare sector, RAMA represents a distinctive departure from prior research endeavours. Unlike its predecessors, RAMA takes a predominantly bottom-up approach, rooted in an intricate understanding of the healthcare environment. This unique perspective allows RAMA to hone in on observed vulnerabilities intrinsic to the healthcare sector, making it particularly well-suited to address the distinctive challenges of this domain. One of the key strengths of RAMA lies in its grounding in empirically identified vulnerabilities. This factuality ensures the approach remains closely tethered to real-world risks and vulnerabilities, lending it more practicality and relevance. RAMA is designed to be a living framework poised for continuous evolution. The framework accommodates the ever-evolving threat landscape by readily incorporating newly discovered vulnerabilities as they emerge. This dynamic responsiveness ensures that RAMA remains on the cutting edge of cybersecurity defence, effectively adapting to counter emerging threats.

8 Conclusion, limitations, and future work

8.1 Overview

In conclusion, this paper introduced a comprehensive risk assessment solution that has been specially crafted for healthcare organisations, though its application is not restricted solely to this domain. The presented approach harnesses the power of cutting-edge risk assessment tools, and when inte-

grated within our suggested architecture, it provides a unique scoring system tailored to healthcare organisations. This scoring system serves two essential purposes: firstly, it furnishes a local RAMA score, enabling the healthcare organisation to gauge its own security posture accurately. Secondly, it offers a Global RAMA score, facilitating a comparative analysis with other healthcare organisations operating within our ecosystem. By utilising this amalgamated approach, healthcare organisations can gain a profound understanding of their security strengths and vulnerabilities in the context of the broader healthcare landscape.

Furthermore, it is worth noting that the versatility of the proposed methodology extends beyond the healthcare sector. The foundational tools forming the backbone of the risk assessment system are not only designed to be non-resource-intensive but are also not exclusive to healthcare environments. As a result, the proposed approach can be effectively employed across various infrastructures, allowing organisations from diverse sectors to benefit from its robust risk assessment capabilities. Whether in finance, education, or any other industry, the adaptability of these baseline tools empowers organisations to fortify their security measures and proactively manage risks, thereby enhancing their overall resilience in the face of potential threats. As organisations continually strive to bolster their security posture, this risk assessment methodology offers a valuable solution that combines state-of-the-art tools with a flexible architecture, paving the way for proactive risk management and better-informed decision-making.

Moreover, this paper provides a comprehensive exploration of the practical implementation and efficacy of our proposed solution in real-world environments. Through the introduction of a diverse set of compelling use cases, we illustrate the versatility and applicability of our solution across various scenarios. These use cases offer concrete examples of how our approach addresses and resolves critical challenges faced by organisations, reinforcing its practical value and relevance.

As part of our comprehensive analysis, we delve into the intricacies of each use case, showcasing the tailored implementation of our solution to address specific issues unique to each scenario. By presenting these real-world applications, we demonstrate the adaptability and robustness of our methodology in handling a wide spectrum of challenges that organisations may encounter in different domains.

Lastly, we elaborate on the successful deployment of our solution in real-world scenarios, shedding light on the process and steps taken to integrate our risk assessment tools within existing infrastructures seamlessly. Our case studies offer valuable insights into how organisations can leverage our solution to assess their security posture effectively. We highlight the quantifiable benefits and improvements observed by adopting our approach, underlining its potential to enhance resilience and bolster cybersecurity measures in practice.

Through these real-world examples and practical demonstrations, we aim to empower readers with a clear understanding of the tangible advantages our proposed solution brings to organisations across various industries. By bridging the gap between theory and practice, we provide a road map for successful implementation and encourage organisations to embrace a proactive risk management approach. As cybersecurity threats continue to evolve, our solution stands as a robust, versatile, and field-tested option for organisations seeking to safeguard their assets, protect their data, and stay ahead in an ever-changing threat landscape.

8.2 Limitations

Despite the numerous notable benefits offered by the proposed solution, it is essential to acknowledge certain limitations that warrant consideration. One primary limitation pertains to the current design of the Risk Assessment for Medical Applications (RAMA) algorithm, which relies on input from specific risk assessment tools. While these tools are undoubtedly state-of-the-art and contribute to the robustness of the RAMA score, their exclusivity could pose a barrier for certain organisations lacking access to the same set of tools. Nevertheless, this constraint can be addressed by leveraging alternative tools with similarities to those highlighted in Sect. 7. Such an approach would not only overcome the limitation but also enhance the overall solution's capability to support a broader array of tools.

Moreover, another limitation arises from the inherent nature of risk assessment methodologies in general. Risk assessments are based on historical data, known vulnerabilities, and existing threat intelligence. As a result, they may not account for emerging or previously unknown threats that have not yet manifested in the historical data. This limitation underscores the importance of continuously updating and augmenting risk assessment methodologies with up-to-date threat intelligence and adapting to evolving threat landscapes. By establishing a modular architecture that facilitates seamless integration of new risk assessment methodologies and tools, it will enable RAMA to evolve and incorporate the latest advancements in risk assessment methodologies.

Additionally, the accuracy and reliability of risk assessments can be influenced by the quality and completeness of the data used as inputs. Incomplete or inaccurate data can lead to biased risk assessments, potentially overlooking critical vulnerabilities or overemphasising certain risks. To tackle such issues, the implementation of robust data validation mechanisms within the overall solution would help to identify and rectify incomplete or inaccurate data.

Furthermore, the scope of the proposed solution might be limited to specific types of medical applications or health-

care organisations. While the RAMA score and its associated risk assessment tools are tailored to healthcare environments, their applicability to other sectors or industries may require careful evaluation and customisation.

Despite these limitations, the proposed solution remains a valuable asset for enhancing risk assessment practices in medical applications and healthcare organisations. By acknowledging these constraints and proactively addressing them, researchers and practitioners can continually improve and refine the RAMA algorithm and its accompanying risk assessment tools, making them more adaptable, inclusive, and effective for a broader range of organisations and use cases.

8.3 Future work

Moving forward, there are several avenues for future work that promise to advance our understanding and application of the proposed solution. A paramount objective is to delve deeper into the multifaceted characteristics of the RAMA score algorithm. By conducting in-depth research and analysis, we can uncover new insights, optimise its performance, and fine-tune its parameters, making it even more effective in evaluating cybersecurity in healthcare organisations.

To broaden the scope and impact of our solution, the next step involves deploying the RAMA score to additional healthcare, and other critical organisations across Europe. This expansion would provide an opportunity to garner a more comprehensive and nuanced view of the cybersecurity landscape within the healthcare sector. By analysing data from multiple organisations, we can gain valuable insights into common weak points and prevalent vulnerabilities, enabling the development of targeted strategies to bolster cybersecurity preparedness.

Furthermore, to ensure the robustness and practicality of the proposed algorithm, we envision conducting rigorous evaluations utilising focused groups. By involving cybersecurity experts, and relevant stakeholders, we can gather valuable feedback, validate the algorithm's effectiveness, and identify areas for refinement. This collaborative and iterative approach will help bolster the algorithm's credibility and applicability in real-world settings.

Beyond the immediate scope of the RAMA score algorithm, our overarching approach can be significantly enhanced by incorporating machine learning-based anomaly detection techniques. By integrating such advanced methods into our risk assessment methodology, we can effectively identify and address anomalous behaviours and potential threats that traditional approaches may overlook. Machine learning models have the potential to augment the RAMA score's precision and predictive capabilities, thereby elevating the overall efficacy of our approach in safeguarding healthcare organisations from emerging and sophisticated cyber threats.

In conclusion, the future prospects for our research encompass a comprehensive exploration of the RAMA score algorithm's characteristics, broader deployment across healthcare organisations in Europe, rigorous evaluation with focused groups, and the incorporation of cutting-edge machine learning-based anomaly detection methods. By pursuing these avenues of investigation, we aspire to strengthen our solution's relevance, effectiveness, and potential impact in safeguarding the critical infrastructure and data within the healthcare domain and beyond.

Funding The research leading to these results received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 883275 (HEIR)

Data availability No datasets were generated or analysed during the current study

Declarations

Conflict of interest The authors declare that they have no conflicts of interest.

Informed consent Informed consent was obtained from all individual participants included in the study.

References

1. Greer, S.L., et al.: Everything you Always Wanted to Know About European Union Health Policies but Were Afraid to Ask. World Health Organization, Regional Office for Europe (2022)
2. Pang, C.E., et al.: Technology preferences and routines for sharing health information during the treatment of a chronic illness. In: SIGCHI Conference on Human Factors in Computing Systems (2013)
3. Cost of a data breach report 2022 (2022)
4. Muthuppalaniappan, M., Stevenson, K.: Healthcare cyber-attacks and the COVID-19 pandemic: an urgent threat to global health. Int. J. Qual. Health C. **33**(1), mzaa117 (2021)
5. Czeschik, C.: Black market value of patient data. In: Digital Marketplaces Unleashed, pp. 883–893, Springer, (2018)
6. Javaid, M., Haleem, A., Singh, R.P., Suman, R.: Towards insighting cybersecurity for healthcare domains: a comprehensive review of recent practices and trends. Cyber Secur. Appl. **1**, 100016 (2023)

7. Alzahrani, A., et al.: NFC Security Analysis and Vulnerabilities in Healthcare Applications. In: IEEE Pacific Rim Conference on Communications, Computers and Signal Processing (2013)
8. Kumar, C.: New dangers in the new world: cyber attacks in the healthcare industry. Intersect Stanford J. Sci. Technol. Soc. **10**(3), (2017)
9. Gartner Identifies Top Security and Risk Management Trends for 2022 https://www.gartner.com/en/newsroom/press-releases/2022-03-07-gartner-identifies-top-security-and-risk-management-trends-for-2022
10. Spanakis, E.G., et al.: Cyber-attacks and threats for healthcare: a multi-layer thread analysis. In: Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC) (2020)
11. Coventry, L., et al.: Cyber-risk in healthcare: exploring facilitators and barriers to secure behaviour. In: International Conference on Human-Computer Interaction (HCI) (2020)
12. Statista, Global average cost of a data breach by industry (2022)
13. Sophos The state of ransomware in healthcare (2022)
14. Initiative, J.T. F.T.: Guide for conducting risk assessments (NIST SP 800-30r1), National Institute of Standards and Technology (2012)
15. Dubois, É., Heymans, P., Mayer, N., Matulevičius, R.: A systematic approach to define the domain of information system security risk management. In: Intentional Perspectives on Information Systems Engineering. pp 289–306 (2010)
16. Cyber Security Risk Assessments
17. Clusif, M.: Processing guide for risk analysis and management. Club De La Securite De L'Information Francias (2011)
18. Stolen, K., et al.: Model-based risk assessment: the CORAS approach. In: iTrust Workshop (2002)
19. Amutio, M., et al.: MAGERIT-Methodology for Information Systems Risk Analysis and Management. Ministry of Finance and Public Administration, Madrid, Spain (2014)
20. Den Braber, F., et al.: Model-based security analysis in seven steps: a guided tour to the CORAS method. BT Technol. J. **25**(1), 101–117 (2007)
21. "sslscan."
22. wazuh: The Open Source Security Platform
23. rsyslog: The rocket-fast Syslog Server
24. Durumeric Z, et al.: The matter of heartbleed. In: Internet Measurement Conference (IMC), (2014)
25. Apache Spark https://spark.apache.org/
26. Apache Storm https://storm.apache.org/
27. Apache Flink https://flink.apache.org/
28. Apache Spot https://incubator.apache.org/projects/spot.html
29. Hogzilla https://ids-hogzilla.org/
30. Kebande, V.R., Karie, N.M., Ikuesan, R.A.: Real-time monitoring as a supplementary security component of vigilantism in modern network environments. Int. J. Inf. Technol. **13**, 5–17 (2021)
31. Baykara, M., Gurturk, U., Das, R.: An overview of monitoring tools for real-time cyber-attacks. In: 2018 6th International Symposium on Digital Forensic and Security (ISDFS), pp. 1–6, IEEE (2018)
32. Splunk Enteprise Security https://www.splunk.com/en_us/products/enterprise-security.html
33. IBM's QRadar https://www.ibm.com/qradar
34. Palo Alto Networks WildFire https://www.paloaltonetworks.com/network-security/wildfire
35. Guimaraes, L.C., Rebello, G.A.F., Camilo, G.F., de Souza, L.A.C., Duarte, O.C.M.: A threat monitoring system for intelligent data analytics of network traffic. Ann Telecommun. pp 1–16 (2021)
36. Krishnan, P., Duttagupta, S., Achuthan, K.: Sdnfv based threat monitoring and security framework for multi-access edge computing infrastructure. Mobile Netw. Appl. **24**, 1896–1923 (2019)
37. Cui, Y., Yan, L., Li, S., Xing, H., Pan, W., Zhu, J., Zheng, X.: SD-anti-DDoS: fast and efficient DDoS defense in software-defined networks. J. Netw. Comput. Appl. **68**, 65–79 (2016)
38. Kalkan, K., Gür, G., Alagöz, F.: Filtering-based defense mechanisms against DDoS attacks: a survey. IEEE Syst. J. **11**(4), 2761–2773 (2016)
39. Hsieh, C.-J., Chan, T.-Y.: Detection DDoS attacks based on neural-network using apache spark. In: 2016 international conference on applied system innovation (ICASI), pp 1–4, IEEE (2016)
40. Johnson, C., et al.: Guide to Cyber Threat Information Sharing. NIST Special Publication, Gaithersburg (2016)
41. King, S.T., Chen, P.M.: Backtracking intrusions. In: ACM Symposium on Operating Systems Principles (SOSP) (2003)
42. Lee, K.H., Zhang, X., Xu, D.: High accuracy attack provenance via binary-based execution partition. In: Network and Distributed System Security Symposium (NDSS), vol. **16** (2013)
43. Liu, Y., et al.: Towards a timely causality analysis for enterprise security. In: Network and Distributed System Security Symposium (NDSS) (2018)
44. Xu, Z., et al.: High fidelity data reduction for big data security dependency analyses. In: ACM SIGSAC Conference on Computer and Communications Security (CCS) (2016)
45. Robertson, J., Diab, A., Marin, E., Nunes, E., Paliath, V., Shakarian, J., Shakarian, P.: Darkweb Cyber Threat Intelligence Mining. Cambridge University Press, Cambridge (2017)
46. Tounsi, W., Rais, H.: A survey on technical threat intelligence in the age of sophisticated cyber attacks. Comput. Secur. **72**, 212–233 (2018)
47. Ibrahim, A., Thiruvady, D., Schneider, J.-G., Abdelrazek, M.: The challenges of leveraging threat intelligence to stop data breaches. Front. Comput. Sci. **2**, 36 (2020)
48. Rahman, M.R., Mahdavi-Hezaveh, R., Williams, L.: A literature review on mining cyberthreat intelligence from unstructured texts. In: 2020 International Conference on Data Mining Workshops (ICDMW), pp. 516–525, IEEE (2020)
49. Wagner, T.D., Mahbub, K., Palomar, E., Abdallah, A.E.: Cyber threat intelligence sharing: survey and research directions. Comput. Secur. **87**, 101589 (2019)
50. Abu, M.S., Selamat, S.R., Ariffin, A., Yusof, R.: Cyber threat intelligence-issue and challenges. Indones. J. Electr. Eng. Comput. Sci. **10**(1), 371–379 (2018)
51. Ramsdale, A., Shiaeles, S., Kolokotronis, N.: A comparative analysis of cyber-threat intelligence sources, formats and languages. Electronics **9**(5), 824 (2020)
52. Milajerdi, S.M., et al.: Poirot: aligning attack behavior with Kernel audit records for cyber threat hunting. In: ACM Conference on Computer and Communications Security (2019)
53. Bilge, L., et al.: Disclosure: detecting botnet command and control servers through large-scale netflow analysis. In: Annual Computer Security Applications Conference (2012)
54. Antonakakis, M., et al.: Detecting malware domains at the upper DNS hierarchy. In: USENIX Security Symposium (2011)
55. Antonakakis, M., et al.: From throw-away traffic to bots: detecting the rise of DGA-based malware. In: USENIX Security Symposium (2012)
56. Pei, K., et al.: Hercule: attack story reconstruction via community discovery on correlated log graph. In: Annual Conference on Computer Security Applications (2016)
57. Goel, A., et al.: Forensix: a robust, high-performance reconstruction system. In: IEEE International Conference on Distributed Computing Systems Workshops (2005)
58. Pohly, D.J., et al.: Hi-Fi: collecting high-fidelity whole-system provenance. In: Annual Computer Security Applications Conference (2012)
59. Lee, I.: Cybersecurity: risk management framework and investment cost analysis. Bus. Horiz. **64**(5), 659–671 (2021)

60. Amro, A., Gkioulos, V., Katsikas, S.: Assessing cyber risk in cyber-physical systems using the attack framework. ACM Trans. Priv. Secur. **26**(2), 1–33 (2023)
61. Kure, H.I., Islam, S., Razzaque, M.A.: An integrated cyber security risk management approach for a cyber-physical system. Appl. Sci. **8**(6), 898 (2018)
62. Kotenko, I., Chechulin, A.: A cyber attack modeling and impact assessment framework. In: 2013 5th International Conference on Cyber Conflict (CYCON 2013), pp. 1–24, IEEE (2013)
63. Paté-Cornell, M.-E., Kuypers, M., Smith, M., Keller, P.: Cyber risk management for critical infrastructure: a risk analysis model and three case studies. Risk Anal. **38**(2), 226–241 (2018)
64. Lee, I.: Internet of Things (IoT) cybersecurity: literature review and IoT cyber risk management. Future Internet **12**(9), 157 (2020)
65. Svilicic, B., Kamahara, J., Rooks, M., Yano, Y.: Maritime cyber risk management: an experimental ship assessment. J. Navig. **72**(5), 1108–1120 (2019)
66. Cherdantseva, Y., Burnap, P., Blyth, A., Eden, P., Jones, K., Soulsby, H., Stoddart, K.: A review of cyber security risk assessment methods for SCADA systems. Comput. Secur. **56**, 1–27 (2016)
67. Akinrolabu, O., Nurse, J.R., Martin, A., New, S.: Cyber risk assessment in cloud provider environments: current models and future needs. Comput. Secur. **87**, 101600 (2019)
68. Ganin, A.A., Quach, P., Panwar, M., Collier, Z.A., Keisler, J.M., Marchese, D., Linkov, I.: Multicriteria decision framework for cybersecurity risk assessment and management. Risk Anal. **40**(1), 183–199 (2020)
69. Silva, F., Jacob, P.: Mission-centric risk assessment to improve cyber situational awareness. In: Proceedings of the 13th International Conference on Availability, Reliability and Security, pp. 1–8 (2018)
70. Kosub, T.: Components and challenges of integrated cyber risk management. Zeitschrift für die gesamte Versicherungswissenschaft **104**, 615–634 (2015)
71. Eling, M., McShane, M., Nguyen, T.: Cyber risk management: history and future research directions. Risk Manag. Insur. Rev. **24**(1), 93–125 (2021)
72. Gatzert, N., Schubert, M.: Cyber risk management in the us banking and insurance industry: a textual and empirical analysis of determinants and value. J. Risk Insur. **89**(3), 725–763 (2022)
73. McKinsey & Company, The risk-based approach to cybersecurity. https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/the-risk-based-approach-to-cybersecurity (2019)
74. Ganji, D., Kalloniatis, C., Mouratidis, H., Gheytassi, S.M.: Approaches to develop and implement iso/iec 27001 standard-information security management systems: a systematic literature review. Int. J. Adv. Softw., vol. **12**, no. 3 (2019)
75. Amin, Z.: A practical road map for assessing cyber risk. J. Risk Res. **22**(1), 32–43 (2019)
76. Antonucci, D.: The Cyber Risk Handbook: Creating and Measuring Effective Cybersecurity Capabilities. Wiley, New York (2017)
77. Ahmed, M., Panda, S., Xenakis, C., Panaousis, E.: Mitre att&ck-driven cyber risk assessment. In: Proceedings of the 17th International Conference on Availability, Reliability and Security. pp. 1–10 (2022)
78. National Institute of Standards and Technology, NIST Cybersecurity Framework. https://www.nist.gov/cyberframework (2018)
79. Intelligent Transport Systems (ITS) Security: threat, vulnerability and risk analysis (TVRA) tech. rep., ETSI (2010)
80. Center for Internet Security , CIS RAM (Center for Internet Security Risk Assessment Method). https://www.cisecurity.org/insights/white-papers/cis-ram-risk-assessment-method (2021)
81. Alahmari, A., Duncan, B.: Cybersecurity risk management in small and medium-sized enterprises: a systematic review of recent evidence. In: 2020 international conference on cyber situational awareness, data analytics and assessment (CyberSA), pp. 1–5, IEEE (2020)
82. Ghadge, A., Weiß, M., Caldwell, N.D., Wilding, R.: Managing cyber risk in supply chains: a review and research agenda. Supply Chain Manag. Int. J. **25**(2), 223–240 (2020)