Journal of Software Engineering and Applications, 2024, 17, 359-377
https://www.scirp.org/journal/jsea
ISSN Online: 1945-3124
ISSN Print: 1945-3116

A Comparative Analysis of Cybersecurity Threat Taxonomies for Healthcare Organizations

Mahima Jaikanth, Vijay K. Madisetti

School of Cybersecurity and Privacy, Georgia Institute of Technology, Atlanta, GA, USA
Email: vkm@gatech.edu

How to cite this paper: Jaikanth, M. and Madisetti, V.K. (2024) A Comparative Analysis of Cybersecurity Threat Taxonomies for Healthcare Organizations. *Journal of Software Engineering and Applications*, **17**, 359-377.
https://doi.org/10.4236/jsea.2024.175020

Received: April 28, 2024
Accepted: May 27, 2024
Published: May 30, 2024

Abstract

Information technology is critical in coordinating patient records, smart devices, operations, and critical infrastructure in healthcare organizations, and their constantly changing digital environment, including suppliers, doctors, insurance providers, and regulatory agencies. This dependence on interdependent systems makes this sector vulnerable to various information technology risks. Such threats include common cybersecurity risks such as data breaches and malware attacks, unique problems occurring in healthcare settings such as unauthorized access to patient records, disruptions in services provided at medical facilities, and potential harm caused to patients due to the compromise of medical devices. The threat taxonomies, such as the Open Threat Taxonomy, NIST, or ENISA, are foundational frameworks for grasping and categorizing IT threats. However, these taxonomies were not specifically designed to deal with the complexities of the healthcare industry. The problem arises from the gap between these taxonomies' general nature and the industry-specific threats and vulnerabilities that affect healthcare organizations. As a result, many healthcare institutions fail to holistically address and eliminate the unique risks related to confidentiality, integrity, and availability of patients' data as well as critical systems used in healthcare. This paper aims to narrow this gap by carefully assessing these taxonomies to determine the framework best suited for addressing the threat environment in the healthcare sector.

Keywords

Threat Taxonomies, Open Threat Taxonomy (OTT)

1. Introduction

The healthcare sector has a complex problem in the field of information tech-

nology that is essential to service delivery—digital systems integration. The complex network of interde-pendent electronic products, computer information systems, and medical infrastructure makes the healthcare industry vulnerable to cybersecurity threats. This threat encompasses all kinds of cyberattacks, including data breaches, malware infections, and healthcare-specific threats like unauthorized access to patient information and disruptions of critical medical services. In addition, healthcare technology integrated with the IoT creates new security threats that endanger patient safety and confidentiality. The growing number of ransomware infections that infect healthcare organizations indicates the need to address these unique problems because consequences extend beyond data compromise to potential harm to patients and disruptions in medical care. Custom taxonomy, specifically designed for the healthcare industry's defense against emerging cyber threats, is lacking. Healthcare information systems have specific vulnerabilities that include critical patient data, continuous contact between medical devices, and possible impact on patients' safety. A clearly defined taxonomy that reflects the complexities of healthcare is critical to accurately recognize, separate, and address these distinct industry-specific threats. A lack of customized taxonomy causes healthcare organizations to work with an uncharted cybersecurity terrain, undermining their capacity to deal with vulnerabilities productively. This shortcoming emphasizes the necessity of assessing and adapting current taxonomies or developing new ones that are more suitable for the healthcare information ecosystem.

Taxonomies

A threat taxonomy is a structured classification system that categorizes and organizes cyber threats. It provides a common language for discussing and analyzing threats. Threat taxonomies aid in standardizing threat descriptions, promote cross-domain consistency in threat analysis, and ease communication. As the risk management program within an organization matures, threat taxonomies allow the organization to move from less detail to a more detailed classification of all possible threats and events with different levels of specificity. A taxonomy is a system of classification often hierarchical in nature, where each parent tier has its own grouping of terms corresponding to the child's category. It typically includes threat categories, subcategories, and specific threat types. As seen in **Figure 1**, ENISA's Threat Taxonomy, NIST's Cyber-security Framework, and Open Threat Taxonomy (OTT) are well-known examples of threat taxonomies studied in this paper.

The ENISA Threat Taxonomy was developed by the European Union Agency for Cybersecurity (ENISA). It plays a pivotal role in categorizing cyber threats relevant to healthcare. This taxonomy spans a broad spectrum of threat types, including those specific to medical devices, electronic health records, and telemedicine [1]. The ENISA threat taxonomy is based on source, target, impact, and vector and operates at four hierarchical levels—Threat Category, Threat

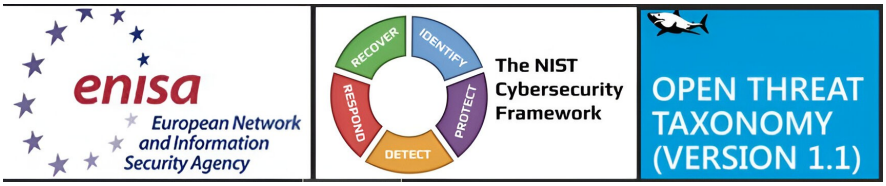

Figure 1. ENISA, NIST, and OTT taxonomies.

Event, Threat Sub-Event, and Threat Detail. It encompasses diverse threats such as malware, web-based attacks, insider risks, and physical manipulation. It aims to enhance cybersecurity across the European Union and foster collaboration among member states. ENISA's annual threat landscape reports provide valuable insights into prevalent cyber threats and emerging trends. Within the healthcare industry, the most relevant categories within ENISA are data breaches, identity theft, ransomware, and supply chain attacks. The National Institute of Standards and Technology (NIST) taxonomy provides a flexible set of standards, guidelines, and best practices for managing cybersecurity risks. Its five core functions—Identify, Protect, Detect, Respond, and Recover—form a comprehensive approach to safeguarding critical systems and data. Open Threat Taxonomy (OTT) was initiated by the SANS Institute and the Center for Internet Security (CIS) as a collaborative effort. It aims to create a common language for describing cyber threats. By classifying threats based on attributes like actors, methods, targets, and impacts, the OTT facilitates information sharing, analysis, and response among cybersecurity professionals and organizations.

2. Background

Recognizing and controlling cybersecurity risks is critical in the quickly changing healthcare IT environment. Healthcare companies may better prioritize risks, manage resources, and deploy focused security solutions using threat taxonomies.

The healthcare industry needs robust threat taxonomies that can help it better manage its systems' inherent complexities. Healthcare organizations often operate within intricate ecosystems with multiple stakeholders, interconnected devices, and sensitive patient data. Threat taxonomies offer an organized way to recognize, classify, and analyze these challenging situations related to cybersecurity. This approach is vital for effective risk management because it enables healthcare organizations to prioritize risks, allocate resources efficiently, and put security measures specifically tailored toward them. Furthermore, using standardized threat descriptions made possible by these threat taxonomies enhances interoperability and integration between various health information system instances. Adopting appropriate threat taxonomies should enable healthcare providers to enhance their resiliency against cyber threats, secure patient privacy, and maintain operations continuity.

The existing approaches to threat taxonomies within the healthcare industry have several limitations. Firstly, the healthcare environments differ significantly:

hospitals and clinics, telemedicine, and medical devices. Generalized taxonomies do not adequately capture risks specific to the context of the environment. Secondly, there are challenges in striking the right balance between granularity and scope of taxonomies because overly broad taxonomies can miss serious hazards, while overly specific ones may overload users. The dynamic and ever-changing threat landscape causes taxonomies to become outdated quickly and miss new and emerging threats. It is important to note that no taxonomy can fully address every potential danger. Several risks could go unnoticed because of changing attack vectors or a lack of awareness. Inconsistency in taxonomies' structures and terminologies makes it difficult for taxonomy designers and companies to map threats between frameworks. Thirdly, human bias during the taxonomy-development process can result in a biased classification of threats. Lastly, taxonomies mainly focus heavily on technical threats while overlooking non-technical threats like social engineering attacks that are crucial in the healthcare sector. The limited resources available to healthcare organizations for threat management make it hard to achieve the perfect taxonomy that involves finding a balance between the use of resources and effectiveness. Taxonomies must also comply with privacy laws such as GDPR and HIPAA.

Existing Work

Existing threat modeling tools in the industry include STRIDE, PASTA, VAST, OWASP Application Threat Model, MITRE ATTCK, and threat ranking tools like DREAD. These models can mitigate healthcare information systems (HIS) risks and protect protected health information (PHI). While many papers about Threat Modeling in Healthcare specialize in domains like Mobile Health Systems, Electronic Health Systems [2], Medical Cyber-Physical Systems, and Cloud Computing, very few papers have information about threat taxonomies within the healthcare industry [3].

STRIDE is a well-known threat modeling system that is useful for identifying many kinds of security risks. It addresses denial of service, spoofing, tampering, repudiation, information disclosure, and elevation of privilege. It provides an organized method for comprehending and resolving risks to patient information, medical equipment, and healthcare systems. DREAD is also currently used as a risk rating model [4].

The paper by Alhassan *et al.* uses a methodology for modeling threats in Electronic Health Systems that can be extended to other healthcare domains [5]. The threat modeling steps include identifying access points and assets, identifying threats, and rating the identified threats.

Other valuable sources include Verizon Data Breach Investigations Reports (DBIR), HIPAA Journal, and the OCR portal within the HHS. Examining the DBIR's healthcare section offers insightful information on the unique threat environment specific to the healthcare industry. The research examines trends in data breaches, including threat sources, attack types, exploited vulnerabilities,

and targeted assets. The evidence from DBIR reports highlights the need for a systematic threat taxonomy approach that addresses the wide range of potential threats to healthcare systems, devices, and patient data.

By analyzing the risks to healthcare over time using the Verizon DBIR's longitudinal perspective, it will be possible to identify new dangers and conduct more efficient, evidencebased risk assessments [6]. The HIPAA Journal offers helpful content about cybersecurity risks, compliance challenges, and breaches that affect the healthcare sector [7]. Healthcare businesses may enhance their threat modeling efforts by including the analysis of trends and information from these HIPAA-focused sources [8]. These sources emphasize security flaws, compliance gaps, and regulatory issues that need to be addressed through threat modeling assessments in the present healthcare context.

Current research articles studying different frameworks include a paper: "A Novel Trust Taxonomy for Shared Cyber Threat Intelligence" by Thomas D. Wagner *et al.* that provides insights that can be used to indirectly assess and compare the trustworthiness of frameworks when sharing cyber threat intelligence. However, it does not directly compare OTT, ENISA, and NIST based on trust parameters [9]. The article's primary emphasis lies in gauging threat intelligence's credibility instead of discussing threat taxonomies. It highlights the main criteria required when verifying the trustworthiness of the entities providing this framework. These include the origin of the intelligence (it may originate from the government source, private companies, or individuals), the motivation of the intelligence generator (whether it is carried out for the public benefit or profit), and the methodology (the approach in gathering and analyzing the threat intelligence). The article underlines the necessity of looking beyond taxonomies and instead looking at the holistic aspects to analyze the reliability and trust of intelligence. This holistic evaluation can assist organizations in making informed decisions on leveraging threat intelligence for their cybersecurity posture and risk management strategies.

An insightful study delves into the complexity of taxonomies, identifying commonalities and differences among them [10]. The ENISA Threat Taxonomy has a detailed classification system for cyber threats. It spans various threat categories, subcategories, and types with specific relevance to healthcare. Additionally, ENISA offers the latest version of its Threat Taxonomy in an accessible Excel format on its official website. Furthermore, an article presents the results of a comprehensive evaluation of taxonomies for information technology threats, including ENISA, NIST, and OTT. The assessment evaluates based on parameters such as completeness, complexity, and clarity [11].

3. Methodology

This paper uses qualitative and quantitative metrics to evaluate ENISA, NIST, and OTT. **Figure 2** shows the block diagram of the implementation steps. The first step towards this solution was to identify the assets within the healthcare

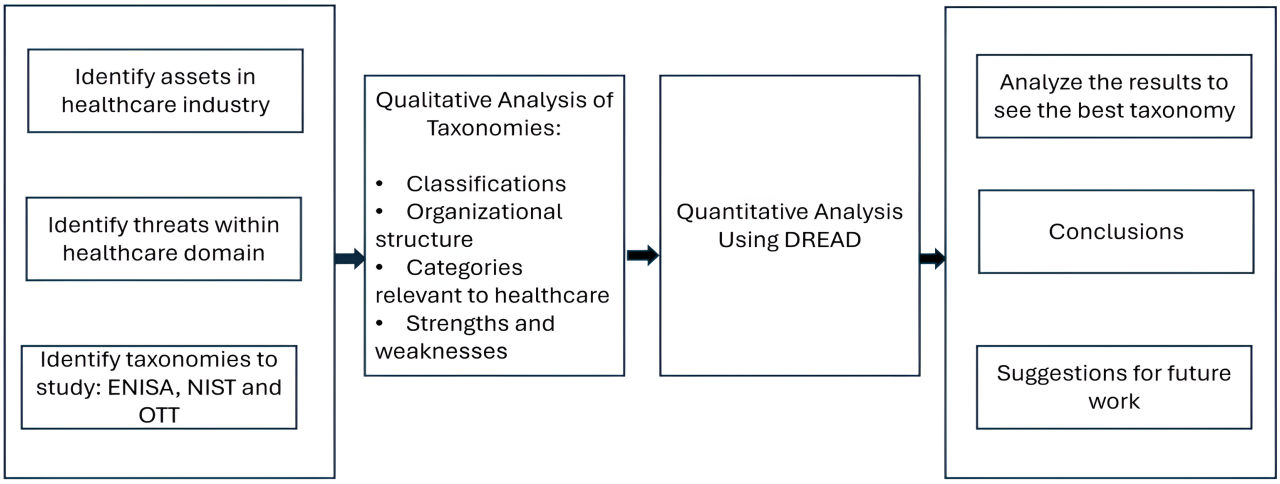

Figure 2. Block diagram of the implementation.

industry and healthcare-specific threats. Next, the threats specific to a healthcare organization in different domains were identified. Following the threat identification, a qualitative analysis of taxonomies was performed. The qualitative analysis includes research of the various methodologies used for threat modeling, finding security and privacy issues in the healthcare domain, listing the threat actors, and documenting the implications of loss of confidentiality, integrity, and availability of PHI (Protected Health Information). Each taxonomy's strengths, weaknesses, and organizational structure were identified in the analysis. Qualitative analysis also involved looking at the classification of each taxonomy and identifying the categories relevant to healthcare within each taxonomy. Following qualitative analysis, a threat ranking tool named DREAD was used to perform quantitative analysis across the ENISA, NIST, and OTT taxonomies and rate the identified threats across various dimensions: Damage (D), Reproducibility (R), Exploitability (E), Affected Users (A), Discoverability (D). Overall Rating and Severity were determined as the result of the analysis. This paper presents the results of the qualitative and quantitative analysis to determine the taxonomy that best fits the healthcare industry.

Improvement over the Baseline

This paper presents the results of a qualitative and quantitative analysis of threat taxonomies using the DREAD model for the healthcare sector. The progress achieved is a significant improvement over the existing work because while quantitative analysis has been performed for threat modeling tools and frameworks in industries such as finance and management, similar research is lacking in the healthcare domain.

Existing research emphasizes the challenges of comparing the quantitative performance of threat modeling tools because of the need for manual configuration and customization, the scarcity of information related to threat taxonomies, and the need for customization of threat modeling tools. Threat modeling has been researched using tools like OWASP and the MITRE ATTCK model; how-

ever, quantitative analysis and assessment of threat taxonomies is lacking, particularly in the healthcare domain.

The research in this paper offers a systematic method to assess the effectiveness of various threat taxonomies for healthcare organizations by performing a quantitative analysis of these taxonomies using the DREAD model. Furthermore, the research highlights the importance of quantitative analysis in the healthcare sector, which traditionally relies on qualitative assessments and subjective evaluations. By introducing a quantitative methodology, the research contributes to a more objective and measurable approach to threat taxonomy selection and implementation in healthcare environments. By using quantitative measurements within particular healthcare requirements, this study can assist companies in selecting the most suitable threat taxonomy and help with their decisionmaking processes. This novel approach can reinforce risk management and enhance cybersecurity procedures, which in turn could help safeguard confidential healthcare data and critical infrastructure.

4. Qualitative Analysis of Threat Taxonomies: ENISA, NIST, and OTT

Qualitative analysis is a critical factor in cybersecurity for creating a better threat categorization. Table 1 summarizes the results of the quantitative analysis. The qualitative analysis descriptions of the three taxonomies are as follows.

Table 1. Summary of qualitative analysis parameters for ENISA, NIST, and OTT.

Parameters	ENISA	NIST	OTT
Scope	Broad coverage of cybersecurity threats, including healthcare-specific ones.	General framework applicable to various sectors, including healthcare	Comprehensive coverage, especially for industrial automation and control systems
Granularity	Detailed threat descriptions, suitable for in-depth analysis.	Balanced approach, neither too detailed nor too high-level	Focus on technical threats, may lack granularity in non-technical aspects
Emerging Threats	May not cover the latest emerging threats comprehensively	Provides guidelines for managing evolving risks	May need updates to address emerging healthcare threats
Applicability	Suitable for healthcare but requires tailoring	Relevant for healthcare but not specific to the industry	May need customization for healthcare context
Alignment with Standards	May align with ISO/IEC 27001 and other standards	Provides overall direction for securing Operational Technology systems	Specific focus on industrial automation security
User-Friendliness	May be complex due to extensive categorization	Balanced and practical for various users	Depends on user familiarity with industrial control systems
Privacy considerations	May need additional privacy-specific criteria	Addresses privacy but not healthcare-specific privacy concerns	Privacy aspects may need customization
Resource requirements	May require significant resources for implementation	Balanced approach considering resource constraints	Resource-efficient but may need adaptation

4.1. ENISA

- **Scope:** Its scope extends broadly across cybersecurity and includes details specific to healthcare. It covers various threats, making it suitable for various sectors.
- **Granularity:** The ENISA taxonomy performs a comprehensive threat analysis, which may be considered broad. It provides a holistic view of threat analysis but lacks fine-grained details.
- **Emerging Threats:** While ENISA captures existing threats effectively, its coverage of emerging threats may not be exhaustive. Continuous monitoring and adaptation is essential to use ENISA to capture new threats.
- **Applicability:** Guidelines are suitable for healthcare but often require tailoring. Although they align with ISO/IEC 27001 standards, they require customization to address healthcare-specific threats [12].
- **Alignment with Standards:** It provides comprehensive guidance, but its alignment with specific standards needs adaptation.
- **User-Friendliness:** Due to extensive threat categorization, threat practitioners may find it complex to adopt ENISA taxonomy.
- **Privacy Considerations:** It emphasizes privacy issues but may need additional criteria specific to healthcare.
- **Resource Requirements:** Implementing ENISA recommendations may demand significant resources.

4.2. NIST

- **Scope:** NIST's framework [13] is not specific to any industry but offers a balanced approach that may apply to healthcare.
- **Granularity:** The methodical and customized strategy employed by NIST addresses high-level hazards while considering changing threat environments. It successfully balances practicality with depth.
- **Emerging Threats:** Its focus on risk management allows it to adapt to emerging threats. However, it does not explicitly highlight healthcare-specific emerging risks.
- **Applicability:** It offers a balanced approach relevant to healthcare but is not specific to every industry. However, it only provides an overall direction for securing operational technology systems.
- **Alignment with Standards:** NIST aligns with general security standards and best practices.
- **User-Friendliness:** It considers resource constraints and ease of implementation, making it very user-friendly.
- **Privacy Considerations:** It does not explicitly address privacy concerns.
- **Resource Requirements:** While NIST is resource-efficient, it requires adaptation within healthcare scenarios.

4.3. OTT

- **Scope:** It has comprehensive coverage for industrial automation and control

systems.

- **Granularity:** OTT [14] assumes user familiarity with industrial control systems (ICS). It lacks granularity in nontechnical aspects and requires customization.
- **Emerging Threats:** It acknowledges many healthcare threats but needs many updates to address emerging cybersecurity risks.
- **Applicability:** OTT requires more customization to fit within the healthcare industry.
- **Alignment with Standards:** It is specifically focused on Industrial Control Systems and lacks specific alignment with established security standards.
- **User-Friendliness:** Its usability depends on the user's familiarity with industrial control systems. It is not very user-friendly as compared to the NIST taxonomy.
- **Privacy Considerations:** OTT needs customization to address privacy aspects in healthcare.
- **Resource Requirements:** Resource implications for OTT in healthcare have not been explored in depth. Existing research indicates that OTT may be resource-efficient but needs further adaptation

Threat events can also be categorized hierarchically by ENISA's threat taxonomy according to their source, target, impact, and vectors. Threat categorization and organizational schemes differ between NIST and Open Threat Taxonomy. While Open Threat Taxonomy employs a four-level hierarchy: threat categories, threat subcategories, threat types, and threat variants, NIST employs a three-level hierarchy: threat sources, threat events, and threat actors [15]. NIST categories relevant to healthcare include natural, human, and environmental threats [16]. Open Threat Taxonomy also has some healthcare-related categories, such as physical, environmental, and human threats. Malicious code, web-based attacks, identity thefts, and physical damage are some of the categories relevant to healthcare in ENISA.

Trustworthiness is another factor that can be used to compare threat taxonomies. While the paper by Thomas D. Wagner *et al.* does not compare OTT, NIST, and ENISA according to the trustworthiness parameter, this research compares the threat taxonomies using the parameters source, motivation, and reputation. The NIST taxonomy is developed and maintained by the National Institute of Standards and Technology, a reputable U.S. government agency. ENISA taxonomy is developed and maintained by the European Union Agency for Cybersecurity, a well-recognized European cybersecurity agency. The OTT taxonomy is a community-driven, opensource project with varying contributors, which may impact some users' trust.

While NIST and ENISA establish standardized frameworks for a range of stakeholders to serve the public good, OTT's motives can differ among community members. Furthermore, while OTT depends on community contributions, NIST and ENISA follow defined methods and rigorous review processes, strengthening confidence in their methodological soundness. The community contri-

butions within OTT may result in lower levels of methodological rigor across different categories. Additionally, NIST and ENISA are well-known in the cybersecurity space, which promotes confidence among users and organizations. However, OTT is a relatively young open-source initiative with a developing reputation.

Based on these trust factors, NIST and ENISA may offer higher indirect trust than OTT due to their established sources, motivations, methodologies, and reputations. However, the actual trustworthiness of each taxonomy depends on the specific user or organization's needs and risk tolerance.

5. Quantitative Analysis of Taxonomies

5.1. DREAD Tool

The DREAD tool is used to perform quantitative risk analysis by rating the severity of a cyber threat on a scale of 1 to 10 for each of the following categories: Damage potential (D), Reproducibility (R), Exploitability (E), Affected users (A), and Discoverability (D).

Damage potential is a category used to evaluate the possible harm that a threat in the taxonomy could bring about. This category can be ranked depending on the seriousness of the damage—such as data loss, monetary loss, or reputational damage. Reproducibility assesses the ease with which a threat can be replicated or replayed. Highly replicable threats pose a greater risk since they can be exploited more frequently. They can be ranked depending on the difficulty of replicating the threat within a taxonomy.

Exploitability is a metric used to gauge how simple it is to take advantage of the threat's associated vulnerability. Higher risk is associated with threats that are simpler to exploit. Risks can be ranked according to the difficulty of the exploitation, the level of expertise needed, and the accessibility of the exploit code. The number of people or systems the threat may impact can be determined by investigating the affected users' category.

Greater risk is associated with threats that impact vital systems or a more extensive user base [17]. Affected users can be ranked using the estimated number of affected users or systems. Discoverability evaluates the possibility that the threat will be discovered and exploited. Less danger is associated with threats that are harder to find or identify. Discoverability can be rated using the difficulty of identifying the threat and the easy availability of detection tools [18].

Users use the DREAD tool not only to identify and prioritize critical threats within a system or a network but also to communicate and justify risk assessment and mitigation strategies to stakeholders within the company's management. DREAD's structured approach allows users to assess and quantify the risks posed by different threats, which helps with better decision-making and resource allocation for threat mitigation and risk management [19].

Table 2 presents a comparative assessment of threat actors across the ENISA, NIST, and OTT threat taxonomies, evaluated across various dimensions using

Table 2. Threat mapping using the DREAD tool for ENISA, NIST, and OTT.

ENISA	Threat Actor	Damage (D)	Reproducibility (R)	Exploitability (E)	Affected Users (A)	Discoverability (D)	Overall Rating	Severity
	Cybercriminals	8	7	9	9	6	7.8	High
	Nation-states	9	6	8	9	3	7	Medium
	Insider Threats	7	5	6	6	9	6.6	Medium
	Supply Chain	8	7	8	9	6	7.6	High
NIST								
	Cybercriminals	8	5	5	8	5	6.2	Medium
	Nation-states	9	4	8	8	4	6.6	Medium
	Insider Threats	8	4	6	7	8	6.6	Medium
	Supply Chain	7	8	8	8	5	7.2	High
OTT								
	Healthcare-specific Threats	7	7.5	6	6	9	7.1	High
	Industrial Control Systems	8	7.5	8	9	6	7.7	High

DREAD: Damage (D), Reproducibility (R), Exploitability (E), Affected Users (A), Discoverability (D), Overall Rating, and Severity. Within the ENISA taxonomy, Cybercriminals stand out as the most severe threat actor, with a high overall rating of 7.8 because of their great potential for harm and exploitability. Supply-chain attacks are rated 4.6 and assigned a high overall rating. Insider threats and nation-state threat actors both have medium severity ratings.

In contrast, the NIST taxonomy assigns a lower overall grade to nation-states, cybercriminals, and insider threats, classifying them as moderately severe threats. On the other hand, supply chain threats are ranked higher, highlighting their significant impact potential.

The OTT taxonomy focuses on specific threat actors, such as Healthcare-specific Threats and Industrial Control Systems. Table 2 highlights some threat actors affecting each taxonomy, scored quantitatively using the DREAD model (rating 0 - 10), offering valuable insights for cybersecurity practitioners and healthcare organizations to prioritize risk mitigation efforts effectively.

The overall ranking for each of the taxonomies ENISA, NIST, and OTT was determined using the DREAD model. None of the current research papers have the numbers for the affected users (A) and discoverability (D) parameters within the DREAD model. Because of the qualitative nature of these parameters, they are assigned labels such as High, Medium, or Low.

For each taxonomy, the following score ranges were assigned as part of the quantitative analysis ranking process: Low: 0 - 4, Medium: 5 - 7, High: 8 - 10. Within the range of high, medium, and low, the values are selected based on extensive research and understanding of the taxonomies. Patients, healthcare pro-

fessionals, and healthcare institutions were considered targets while ranking each taxonomy.

5.2. Scoring Using DREAD

Damage Potential: A higher score in this category indicates that the threat has the potential to cause significant harm or damage. A higher score here is desirable because it highlights the taxonomy's ability to address a broad spectrum of threats. ENISA's taxonomy covers a wide range of threats, including potential damage to health data, systems, and services, and hence, it is assigned a score of 8.NIST is given a high score of 9 because it includes critical infrastructure in addition to sensitive and non-sensitive user data. In the case of the OTT, a high score of 8 is assigned for Damage Potential because it covers a wide range of threat categories, including sensitive and non-sensitive user data compromise.

Reproducibility: A higher score in this category implies the threat is easily reproducible or exploited. Conversely, a lower score indicates that the threat is difficult to reproduce. Thus, a higher reproducibility score is undesirable because it means the threat can be easily exploited, NIST and ENISA are rated 5 for reproducibility because of their complex taxonomy structure, which makes them difficult to reproduce. OTT, on the other hand, has a score of 7.5 because its clear organization and structure of taxonomy facilitate the consistent identification and classification of threats, contributing to its higher reproducibility.

Exploitability: A higher exploitability score shows that the threat is readily exploitable due to vulnerabilities or weaknesses. Conversely, a lower score indicates that the threat is more complex to exploit. A higher exploitability score is undesirable in most cases because it signifies a higher risk of successful attacks. ENISA and NIST are assigned a value of 5 for exploitability because they acknowledge that attackers can exploit common vulnerabilities with available attack tools. OTT achieves a high score of 9 for exploitability because of its thorough threat descriptions that contain details about web application proxies or other exploit vectors. In other words, OTT is more easily exploitable than ENISA or NIST.

Affected Users: A higher score indicates that a more significant number of users or systems are vulnerable to threats. Conversely, a lower score means that fewer users or systems are affected. A higher score is undesirable for Affected Users because it implies a broader impact on users or systems. The threat ranking is 6 for ENISA because the threat affects a few users, specifically healthcare professionals and patients. NIST scores 8 for Affected Users because the threats affect administrative users, including healthcare administrators and IT staff. According to NIST, the threat's impact is concentrated on those responsible for managing and securing healthcare systems. OTT is assigned a median score of 6 as it acknowledges that the number of users affected by a particular threat can vary depending on the specific context and environment in which the threat is present.

Discoverability: A higher Discoverability score signifies that the threat is easily discoverable or detectable. Conversely, a lower score implies that the threat can remain hidden or undetected. A higher discoverability score is preferable because it allows for timely detection and mitigation of threats. ENISA is assigned a score of 8 for Discoverability because its vulnerabilities are publicly known within the cybersecurity community. They could be disclosed through research papers and public databases. NIST has a lower score of 6 because its vulnerabilities are not as widely known as ENISA's, but they are still accessible to security practitioners and attackers.OTT receives the highest score of 10 for Discoverability due to its comprehensive coverage of threats.

The DREAD model calculates the overall score by adding the scores for each category and dividing them by the total number of categories. A higher overall DREAD score signifies a higher risk rating of the taxonomy [20].

The equation used in this paper to calculate the risk of each of the taxonomies is shown below where D, R, E, A, and D stand for Damage Potential, Reproducibility, Exploitability, Affected Users, and Discoverability, respectively:

$$\text{Risk value} = \left(\text{D} + \text{R} + \text{E} + \text{A} + \text{D}\right)/5 \tag{1}$$

Equation (1) determines the overall risk rating. As seen in **Table 3**, the overall threat rating for ENISA is determined by adding the values for Damage Potential, Reproducibility, Exploitability, Affected Users, and Discoverability and dividing by 5 to determine the overall threat rating:

$$\text{ENISA Risk} = \left(8 + 5 + 5 + 6 + 8\right)/5 = 6.4 \tag{2}$$

Similarly, the values for NIST and OTT are calculated.

The quantitative analysis results from **Figure 3** reveal that the ENISA and NIST taxonomies exhibit comparable overall ratings of 6.4, categorized as "Medium" severity. Although both taxonomies score relatively high in terms of Damage (D) and Affected Users (A), their lower ratings in Reproducibility (R), Exploitability (E), and Discoverability (D) contribute to their medium severity classification. In contrast, the OTT taxonomy stands out with a significantly higher overall rating of 8.1, indicating a "High" severity level. This elevated rating is driven by its higher scores in Reproducibility, Exploitability, and Discoverability, combined with a substantial potential for Damage.

Table 3. Quantitative analysis results for the 3 taxonomies.

Taxonomy	Damage Potential	Reproducibility	Exploitability	Affected users	Discoverability	Overall risk rating	Rating
ENISA	8	5	5	6	8	6.4	Medium
NIST	9	5	5	8	5	6.4	Medium
OTT	8	7.5	9	6	10	8.1	High

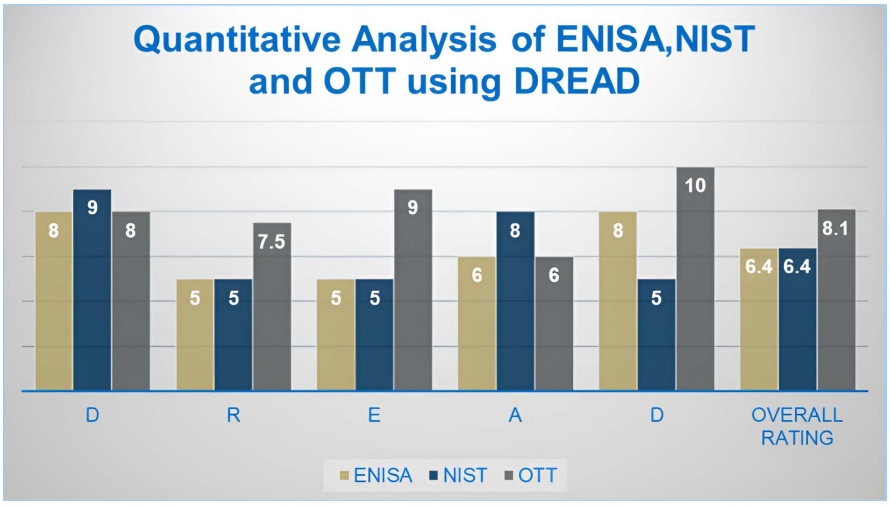

Figure 3. Quantitative analysis graph of ENISA, NIST, and OTT using DREAD.

5.3. Risks and Challenges

Several risks are associated with using the DREAD model for quantitative analysis of taxonomies. Firstly, the DREAD model is subjective and depends upon the judgment of the analyst to assign numerical values for various categories of risks. Analysts might have different opinions on the same threat, resulting in inaccurate or inconsistent results. Secondly, the DREAD model has not been validated or tested by empirical studies or data; it is purely based on the assumptions and experiences of Microsoft developers. Thus, it might not apply to different domains or scenarios. The third challenge associated with the DREAD model is the difficulty of scaling and automating the model for large systems with multiple threats and vulnerabilities. It requires manual input calibration for threats, making it time-consuming and prone to errors. The qualitative nature of the model makes it hard to assign specific values for each taxonomy. Lastly, quantitative analysis using the DREAD model considers only a subset of threat actors as part of this research because of the restricted timeline and scope. With the ever-increasing nature of cyber threats and the vast number of existing threat actors in the healthcare industry, conducting a comprehensive qualitative and quantitative analysis that encompasses all possible threat actors is necessary, which was not feasible within the scope of this research. This paper focuses on a representative sample of relevant threat actors to demonstrate the applicability of the quantitative approach while acknowledging the dynamic and evolving landscape of threats in the healthcare sector.

6. Comparison of the Qualitative and Quantitative Analysis

The Open Threat Taxonomy (OTT) adopts the high-level structure of the NIST CSF Core, which includes five functions: Identify, Protect, Detect, Respond, and Recover. The classifications of threats found in NIST publications, including "Adversary," "Malware," and "Attack Vector," are incorporated into OTT's threat taxonomy.

In contrast, the ENISA Threat Taxonomy aligns with the NIST Cybersecurity Framework's Core Functions and Categories, providing a standardized method for classifying and arranging threats. NIST rules have an impact on ENISA's incident handling and response procedures. ENISA's risk management methodologies and frameworks partially stem from NIST's Risk Management Framework (RMF).

In summary, OTT and ENISA incorporate NIST's approach to risk management, threat taxonomy, and incident response. NIST's tiered approach to cybersecurity is mirrored in OTT and ENISA. Both frameworks draw from NIST's best practices for cybersecurity, tailoring them to their specific needs and contexts to promote consistency, interoperability, and alignment with widely accepted standards and best practices. They adopt NIST's recommended security controls for protecting information and information systems.

Qualitative analysis reveals that ENISA is suited for comprehensive threat analysis. It aligns with ISO standards. However, its broad categorization lacks granularity, is complex to users for users, and requires tailoring for healthcare-specific needs.

On the other hand, NIST is user-friendly, resource-efficient, and has a balanced approach but lacks healthcare-specific customizations and does not address privacy issues. Finally, OTT is comprehensive for industrial control systems. However, it requires users to be familiar with Industrial Control Systems, lacks granularity in non-technical threats, lacks alignment with established standards, and needs more customization for healthcare.

There is no "best" taxonomy based only on qualitative analysis. The choice depends on the specific needs: ENISA is more suitable for gaining a broad understanding of threats in healthcare and incorporating some customization. NIST is better for an organization aiming for a balanced, adaptable approach applicable to healthcare with resource constraints.

OTT is the best option if the focus is on industrial control systems within healthcare and the organization is willing to adapt significantly.

For organizations considering standardized frameworks, established reputation, and trustiness as crucial factors in deciding on a taxonomy, both NIST and ENISA would be suitable. Quantitative analysis underscores the importance of considering multiple dimensions of threat severity when evaluating threat taxonomies. By considering factors such as reproducibility, exploitability, discoverability, damage potential, and affected users, organizations can prioritize and address cybersecurity risks within their specific contexts more effectively.

The quantitative analysis indicates that both ENISA and NIST taxonomies exhibit comparable overall ratings of 6.4, categorizing them as having "Medium" severity. In contrast, the OTT taxonomy stands out with a significantly higher overall rating of 8.1, indicating a "High" severity level. Given these findings, organizations seeking lower risk should consider adopting the ENISA and NIST taxonomies, which offer more favorable risk ratings than the OTT taxonomy.

7. Conclusion and Future Work

While no one taxonomy fits the needs of all healthcare organizations, the research in this paper suggests that a healthcare organization with threats similar to those mapped in this paper should pick ENISA or NIST taxonomies over OTT to minimize organizational risks. Integrating the quantitative factor helps provide a comprehensive and transparent approach to defining the most appropriate threat taxonomy for use within the healthcare industry. This methodology guarantees that the selected taxonomy is consistent with healthcare information systems specificity and exceeds the standards defined for efficient threat handling.

Overall, the research in this paper offers a structured approach for evaluating the suitability of threat taxonomies in healthcare. The implications of this project go beyond the enhancement of cybersecurity measures. The tailored threat taxonomy will create a paradigm shift in health care, leading to better data security by ensuring increased patient privacy and system reliability. A customized taxonomy focusing on the unique healthcare environment helps increase privacy protection and information security. Patients can trust the confidentiality of their personal data and sensitive medical information.

Adopting a custom taxonomy will be a reliable guiding tool that enables healthcare institutions to navigate the complex threat landscape effectively. Such a targeted approach ensures that vulnerabilities unique to healthcare information systems are identified and addressed orderly, thus strengthening the industry's defenses against a multitude of cyber threats. Achieving success in threat taxonomy implementation requires precision in threat mapping, a blend of quantitative and qualitative analysis, and continuous adaptation to the evolving threat landscape. The application of threat taxonomy in decision-making involves three steps. First, we need to know how taxonomy is used in risk assessments, which involves mapping the taxonomy to applicable security controls or requirements. Then, we should know where and how that mapping can go wrong. The taxonomy mapping process is intricate, and misalignment may occur if the taxonomy lacks granularity or fails to address specific industry nuances.

A quantitative approach could offer valuable insights since most existing risk frameworks are based on qualitative analysis methods. Comparative feedback on different threat taxonomies using quantitative metrics could help identify nuances that might be missed by qualitative analysis. For example, a comparative case study could address how different threat taxonomies perform in different threat scenarios. Another approach would be to employ the same risk-assessment framework to identify different types of threat taxonomies.

Generating more empirical data is crucial to mitigate subjectivity associated with using the DREAD model. As a part of future work, producing objective measurements such as historical incident data, impact severity, and real-world consequences can help provide a more accurate assessment using DREAD.

Taxonomies must be constantly updated in a quickly shifting threat environment to keep them robust enough to support worthwhile decisions.

Acknowledgements

We thank Prof. Sergio Caltagirone and the reviewers for their comments that improved the paper.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

[1] Olaniyi, O., Alhassan, J., Abba, E. and Waziri, V. (2016) Threat Modeling of Electronic Health Systems and Mitigating Countermeasures.

[2] Yeng, P.K., Stephen, D. and Yang, B. (2020) Comparative Analysis of Threat Modeling Methods for Cloud Computing towards Healthcare Security Practice. *International Journal of Advanced Computer Science and Applications*, **11**, 772-784. https://doi.org/10.14569/IJACSA.2020.0111194

[3] Maze, T. (2023) How to Use Dread Analysis with Fair.

[4] Wells, B. (2022) Threat Modeling in Healthcare.

[5] Alhassan, J.K., Abba, E., Olaniyi, O.M. and Waziri, O.V. (2016) Threat Modeling of Electronic Health Systems and Mitigating Countermeasures.

[6] Verizon (2023) 2023 Data Breach Investigations Report. https://www.verizon.com/business/resources/T31c/reports/2023-data-breach-investigations-report-dbir.pdf

[7] HIPAA Journal. Healthcare Data Breach Statistics. https://www.hipaajournal.com/healthcare-data-breach-statistics/

[8] U.S. Department of Health & Human Services, Office for Civil Rights. https://ocrportal.hhs.gov/ocr/breach/breach_report.jsf

[9] Wagner, T.D., Palomar, E., Mahbub, K., Abdallah, A.E., *et al.* (2018) A Novel Trust Taxonomy for Shared Cyber Threat Intelligence. *Security and Communication Networks*, **2018**, Article ID: 9634507. https://doi.org/10.1155/2018/9634507

[10] ENISA Publishes a Tool for the Mapping of Dependencies to International Standards. https://www.enisa.europa.eu/news/enisa-news/enisa-publishes-a-tool-for-the-mapping-of-dependencies-to-international-standards

[11] Launius, S. (2020) Evaluation of Comprehensive Taxonomies for Information Technology Threats.

[12] Logix Consulting (2019) Managed It Support Services Seattle. https://logixconsulting.com/2019/12/18/what-is-the-dread-cybersecurity-model/

[13] National Institute of Standards and Technology (2021) White Paper. https://www.nist.gov/system/files/documents/2021/10/29/34-ENISA%20MCS%20resp%20NIST%20IoT%20White%20Paper%20-%2017.10.2021fin.pdf

[14] Open Threat Taxonomy v1.1a. https://www.auditscripts.com/resources/open_threat_taxonomy_v1.1a.pdf

[15] INTECH Automation Intelligence (2022) Iso 27001, isa/iec 62443, and nist csf: Se-

lecting the Right Standard/Framework for Your OT Cybersecurity Program.

[16] Zhang, L., Taal, A., Cushing, R., de Laat, C. and Grosso, P. (2022) A Risk-Level Assessment System Based on the Stride/Dread Model for Digital Data Marketplaces. *International Journal of Information Security*, **21**, 509-525.

[17] EC-Council (2023) Dread Threat Modeling: An Introduction to Qualitative Risk Analysis.

[18] Thevarmannil, M. (2024) 10 Types of Threat Modeling Methodology to Use in 2024.

[19] Kirtley, N. (2023) Dread Threat Modeling.

[20] DREAD (2023) Risk Assessment Model. https://en.wikipedia.org/wiki/DREAD_(risk_assessment_model)

Appendix

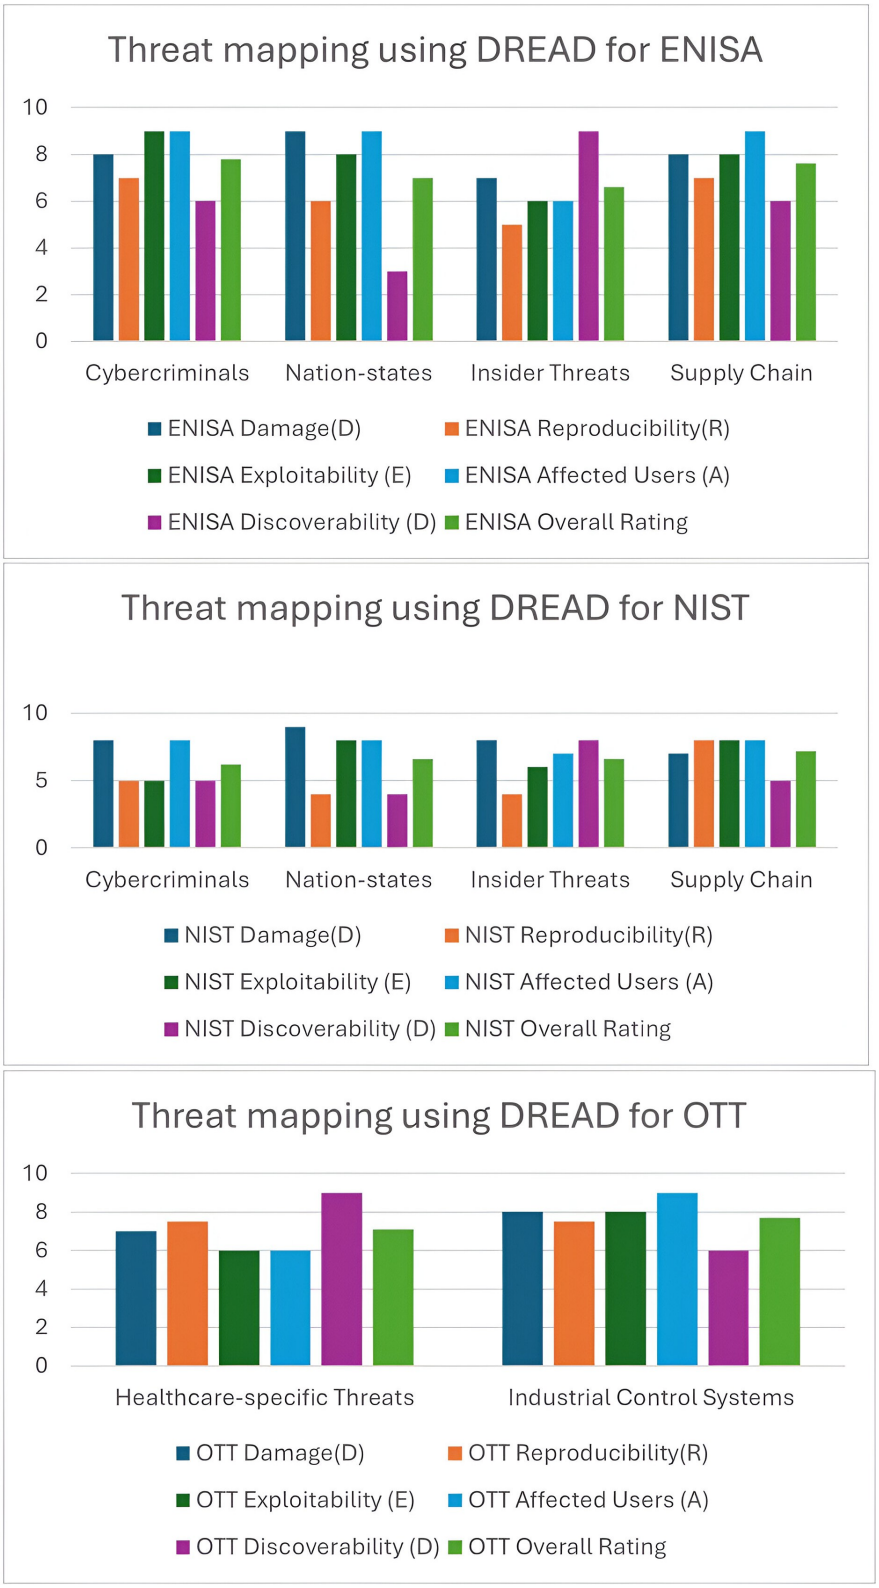

Figure A1. Individual threat mapping graphs for ENISA, NIST and OTT taxonomies.