Chair of IT Security
School of Computation, Information and Technology
Technical University of Munich

Proceedings of the Seminar Cyber-Resilient Systems

Summer Semester 2024

Munich, April 18, 2024 – July 4, 2024

Editor: Lukas Gehrke

Proceedings of the Seminar
Cyber-Resilient Systems (CRS)
Summer Semester 2024

Editor:

Lukas Gehrke
Chair of IT Security (I20)
E-mail: gehrke@sec.in.tum.de
Internet: https://www.sec.in.tum.de/i20/people/lukas-gehrke

Cataloging-in-Publication Data

Seminar CRS SS 24
Proceedings of the Seminar Cyber-Resilient Systems (CRS)
Munich, Germany, April 18, 2024 – July 4, 2024

DOI: 10.14459/2024md1759337
Cyber-Resilient Systems (CRS)
Series Editor: Lukas Gehrke, Technical University of Munich, Germany

Preface

We are pleased to present to you the first version of the proceedings of the seminar Cyber-Resilient Systems (CRS) during the summer semester 2024. From this year onwards, the seminar will take place in both the summer and winter semesters.

The seminar investigates topics around the idea of cyber resilience, an approach to cyber security that lays a strong focus on dealing with harm in the aftermath and minimizing damage. In the seminar, students write a scientific paper about a topic they choose together with their supervisor. Supervisors are staff members from the Chair of IT Security at the Technical University of Munich. During the semester, the supervisors review the papers of the students and give them feedback. The seminar concludes with presentations by all participants.

Among the participants of the seminar, one is awarded with a *Best Paper Award*. For this semester, the award goes to Adrian Stein who wrote the paper *SoK: A Review of Cyber Resilience Frameworks for System Architectures* , congratulations!

We would be delighted if the contributions from this seminar help you. If you would like to know more about our work, please take a look at our homepage `https://sec.in.tum.de`.

This work was partially supported by the TUM Innovation Network ReTruSt: Resilient, Trustworthy, Sustainable.

Finally, we thank the Chair of Network Architectures and Services at TUM for advising the editor on the design and publication of this report.

Munich, October 2024

Claudia Eckert

Lukas Gehrke

Seminar Organization

Chair Holder

Claudia Eckert, Technical University of Munich, Germany

Technical Program Committee

Lukas Gehrke, Technical University of Munich, Germany

Advisors

Lukas Gehrke (gehrke@sec.in.tum.de)
Technical University of Munich

Seminar Homepage

`https://www.sec.in.tum.de/i20/teaching/ss-2024/cyber-resilient-systems`

Contents

Seminar

SoK: A Review of Cyber Resilience Frameworks for System Architectures

Adrian Stein
Chair of IT Security
Technical University of Munich
Garching, Germany
adrian.stein@tum.de

***Abstract*—The term *cyber resilience* has become almost ubiquitous in today's cyber security vocabulary with good reason. It is a concept that aims to ensure the continued operation of systems and gives rise to the idea that there is more to a good cyber defense than just prevention. *Cyber resilience* highlights the importance of other system properties during adverse events, like *recovery* and the ability to *adapt*. Many frameworks, like the *NIST Cybersecurity Framework*, aim to guide system designers in building *cyber resilient* architectures. However, it sometimes remains unclear which properties of *cyber resilience* are covered by these frameworks. Therefore, this paper first establishes a definition of cyber resilience that is coherent with the current literature. Additionally, we provide insight into related concepts to the idea of *cyber resilience* like the *kill-chain* by *Lockheed Martin* and the *Mitre ATT&CK* matrix. With this knowledge, we review nine frameworks that address *cyber resilience* of system architectures and categorize each framework based on which *resilience* properties are covered. We have found that the frameworks that cover most or all *resilience* properties have strong ties to organizations associated with the US government. The work of independent researchers tends to focus on one selected property. Moreover, we found that there is a lack of research assessing the real-world performance of these frameworks. These findings highlight the need for a more unified approach concerning *cyber resilience* frameworks.**

***Index Terms*—Cyber Resilience, System Architecture, NIST Cybersecurity Framework, CSF**

I. Introduction

Over the last few decades, cyber threats to governments, critical infrastructure, and the private sector have increased. One of the more notable attacks, discovered in 2010 by *VirusBlokAda*[1], was the *Stuxnet* worm. *Stuxnet* targeted uranium enrichment facilities in Natanz, Iran, by reprogramming the PLCs[2] controlling the European-made IR-1 centrifuges. The infected PLCs were instructed to operate the motors of the centrifuges outside of their intended operational specification, thus causing physical damage to the enrichment facility. *Stuxnet* is no simple attack tool. It consists of more than 50.000 lines of code and utilizes four *zero-days* in conjunction with stealth techniques to hide from system designers and engineers, thus becoming the first PLC *rootkit*. [1]–[3]

While *Stuxnet* did not spy or otherwise cause harm to other machines outside of its limited targets in Natanz [2], other malicious software is undoubtedly intended to disrupt the orderly operation of businesses, public services, and governments alike. *WannaCry*, and *Petya*, two *ransomware* tools discovered in the mid to late 2010s, caused outages of critical infrastructure and considerable financial damage to businesses. In 2017, a malware called *NotPetya* emerged, spreading around corporate networks by instrumentalizing *Microsoft Exchange Servers* as the point of entry, encrypting the data stored on any infected machine, and demanding payment to restore access to the encrypted data. However, the restoration part was a lie; the data was wiped instead (thus facilitating its name as not the malware it looked to be). *NotPetya* primarily targeted companies operating out of Ukraine. *Merck & Co.*, *Mondelēz*, *A.P. Møller-Mærsk*, and *TNT Express*, to name a few, were heavily affected by the *NotPetya* attack. Production lines, distribution of goods, and global shipping & logistics infrastructure stopped, prompting these companies to later report massive losses (in some cases, multiple hundred million dollars). [1]

However, not every company has to shut down its operations during attacks that have the potential to devastate the global economy. Sometimes, these companies do not even need to disclose any losses. Circling back to *NotPetya*, companies like *Boeing*, *Oracle*, *Johnson & Johnson*, and even *Microsoft* continued working from their offices in Ukraine. [1]

Due to this circumstance, a simple question arises. How is it that some companies (or, for that matter, any institutions) are better equipped to deal with emerging cyber threats than others? Answering that question is not trivial; comparing two companies in the first place is not simple. One explanation could be the different degrees to which companies incorporate robust *cyber resilience* techniques into their defense strategy.

For this reason, the US administration under Obama issued *Executive Order 13636* [4], named "Improving Critical Infrastructure Cybersecurity," on February 12th, 2013. *EO-13636* proposed creating a baseline framework that provides standards and guidelines for reducing cyber risks. The *National Institute of Standards and Technology* (*NIST*) was tasked to develop this framework. The first version of the aptly called "NIST Cybersecurity Framework" (*CSF*) was published a year later to meet the demands from *EO-13636*. [4], [5]

The *CSF* aims to help organizations strengthen their security and resilience of critical infrastructure regardless of size by

[1] Antivirus company based in Belarus

[2] Programmable Logic Controller

doi: 10.14459/2024md1759337_01

providing a framework of adequate standards, guidelines, and practices. The framework can be incorporated into existing processes to identify and fill the gaps concerning the current cyber risk management coverage, complementing the existing development cycles instead of replacing them. [5]

While the *CSF* provides a solid foundation for improving and unifying the cyber risk management strategies of independent actors in the private sector and governments alike, it is not the only extant framework addressing this issue. Therefore, this paper reviews and summarizes nine frameworks, one of which is the *CSF*, to better understand the currently available frameworks that task themselves to improve the *cyber resilience* of system architectures. To achieve this, we first define in section II what *cyber resilience* stands for, along with related ideas and concepts: namely the *kill-chain* [6], the *ATT&CK* matrix [7], and the *impact-wave* analogy [8]. Section III establishes our literature research process and the definition of *cyber resilience* we will use for our evaluation. In IV, we categorize and summarize the nine selected frameworks in detail. The selected frameworks are formulated generally, thus enabling them to be applied to various system architectures and businesses. Afterward, we assess our categorization in section V. Section VI then gives an overview of other papers detailing systematizations for *cyber resilience* frameworks. We also included two papers, [3] and [9], discussing technical aspects of improving *cyber resilience* within a more specified architecture. We justified their inclusion since they aim to strengthen the *cyber resilience* of *cyber-physical* systems and *authentication* & *authorization*, both common concepts within Industry 4.0. In VII, we discuss the limitations of this work (see VII-A) and propose creating a unified *cyber resilience* framework to streamline the incorporation process of such frameworks into the defense strategies of today's businesses and institutions (see VII-B). Then, VIII summarizes this work's results and contributions to the field of *cyber resilience* research.

II. Background and Definitions

First and foremost, we introduce the necessary definitions and concepts on which we base our categorization of proposed approaches and solutions to secure system architectures. II-A points out existing definitions of *cyber resilience*, which we refine through additional insight via the *impact-wave* [8] analogy model shown in II-B. Then we briefly discuss the actions and tools an attacker might utilize to compromise a system by introducing *Lockheed Martin*'s *Kill-Chain* [6] in II-C and *Mitre*'s *ATT&CK* matrix [7] in II-D. These concepts give a comprehensive overview of *cyber resilience* goals and an introduction to adversarial behavior.

A. Cyber-Resilience

The *National Academy of Sciences* (*NAS*) defines *resilience* as the ability to "prepare and plan for, absorb, recover from, and more successfully adapt to adverse events." [10] This initial definition from 2012 was based on existing literature and is, therefore, consistent with the definitions used by other institutions and committees. However, their work primarily focused on improving the response strategies to hazards and natural disasters. Here, *resilience* was used in the literature since 1999 [10], gaining traction in 2005 when widely adopted frameworks for building disaster resilience started to surface. [10]

One can find many definitions regarding *cyber resilience*. But they all have one thing in common: they split *resilience* into distinct *phases* or *functions*, each responsible for addressing one aspect of resisting an adverse event. These functions may be defined granularly or broadly. In the following, we will discuss different definitions for *cyber resilience* to arrive at a definitive one used as a baseline in this work. We begin with the *NIST Cyber Security Framework*, as it seems to be the work that introduced *cyber resilience* concepts to a broader audience in 2014 in response to *EO-13636* (see I).

The *NIST CSF* nowadays defines six core functions: *Govern*[3], *Identify*, *Protect*, *Detect*, *Respond*, and *Recover*. These functions give a general overview of best practices and protocols for organizing cybersecurity at a higher level. A concurrent and continuous application of these functions is preferable over a procedural one. Forming an operational culture under these core principles can adequately address dynamic risks. [5], [11], [12]

1) *Govern*: This function primarily aims to establish an organizational context. As already stated, no two companies are the same. Therefore, it is important to identify which strategies and mechanisms achieve and prioritize the outcome of the remaining functions. Additionally, roles, responsibilities, authorities, and policies should be implemented as a basis for the risk management strategy. [11]
2) *Identify*: The cybersecurity team needs to understand the company's assets and resources. These are not limited to data, hardware, people, and facilities but also include suppliers and related cybersecurity risks. With this understanding, already established policies and risk management strategies can be adjusted accordingly to help manifest the other functions. [5], [12]
3) *Protect*: Appropriate physical security controls and safeguards can be implemented by identifying assets and risks, thus reducing the impact of adverse conditions. Potential outcomes of this function may include identity management, access control, awareness & training, data security, and platform security. [11], [12]
4) *Detect*: Timely discovery of cyberattacks and anomalies. Detection categories include continuous monitoring, detection processes, and successful incident response. [11], [12]
5) *Respond*: Containment of anomalies and adversarial activity during a detected cybersecurity incident. Outcome categories include incident management, response planning, analysis, and mitigation. [5], [11]

[3] the initial version of the *CSF* omitted *Govern* as a core function

6) *Recover*: Restoration of normal operations after a cybersecurity event. Common activities are appropriate communication during recovery and recovery planning improvements. [11], [12]

The *NIST CSF* defines some of the core concepts involved in *cyber resilience*. However, their definition lacks proper methods for *adaptation*. While *detection*, *response*, and *recovery* can restore the system to a normal state, it is crucial that the system also appropriately adapts to adverse conditions. Cyber defense is, therefore, similar to evolution. If an organism does not adjust to a changing environment or external stresses, then none of its survival mechanisms are working correctly, thus necessitating improvements. This does not mean that *adaptation* is not within the scope of the *CSF*; instead, it's just part of the *response* function and not a category in and of itself [5].

Linkov and Kott define *cyber resilience* more concisely in [13]. They give four essential phases: *Plan*, *Absorb*, *Recover*, and *Adapt* (based on the findings of the *National Academy of Sciences*). However, while defining each phase, they also note what resilience means in other disciplines (socioecological, psychological, organizational, and engineering) to help establish a defining resilience feature for each phase. [13]

Resilience Phase	Resilience Feature	Description
Plan	Critical Functions	Identification of system functions by which system performance is measured
Absorb	Thresholds	Tolerance to stress or change. Exceeding certain thresholds due to stresses perpetuates a regime shift
Recover	Time (and scale)	Duration in which the system performance is degraded
Adapt	Memory / adaptive management	Changes in management or response handling enabled by learning outcomes from previous disruptions

TABLE I: Resilience feature of each resilience phase as defined by Linkov and Kott (see [13])

Table I links the resilience phases to their defining feature. When looking at the defining features of each phase, one can quickly identify many similarities with the previously described core functions of the *NIST CSF*. However, it becomes clear that Linkov and Kott offer a new perspective for *absorption* and *adaptation*. *Absorption* is now better defined through *thresholds*. There is a limit to what a system can withstand; exceeding these limits must entail a permanent change in policy or strategy to better adapt to the current environment [13]. We can derive from this that every system must have some core principles that must be retained for the system to *exist* meaningfully. These core principles are dependent on the system itself. However, these may include basic safety guarantees in the case of cyber-physical systems or the protection of specific security goals (like *confidentiality* or *authenticity*) at all times, to name some examples (we will discuss some more detailed examples in section IV). In the case of *adaptation*, the concept of *memory* describes the ability of a system to self-organize itself. In contrast, *adaptive management* supports learning about a system's limits and opportunities to facilitate improvement after an incident [13].

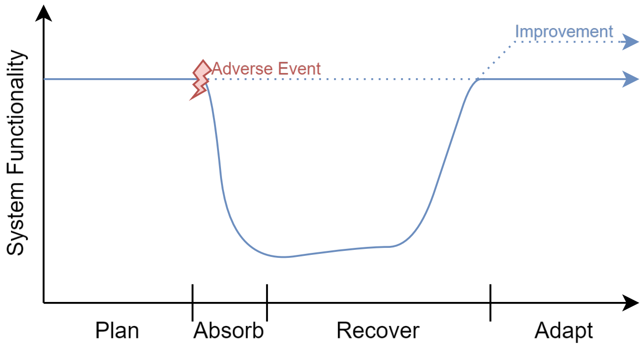

Fig. 1: System performance during each *cyber resilience* phase (see [13])

Figure 1 depicts a system's performance during the different phases of *cyber resilience* as defined by [13]. Here, it is essential to note that resilience should not only refer to the ability to return to a previous good state in the face of an adverse event but also imply coming back stronger than before [13]. Clarke and Knake also point this out when comparing the meaning of resilience to other fields of study, specifically psychology, where resilience does not simply mean forgetting a trauma but overcoming it and improving [1].

Ligo, Kott, and Linkov provide additional insight into the interpretation of figure 1 in [14]. The system's resilience for a given period is defined as the area under the curve within a timeframe Δt. Forming the integral has the advantage of considering every aspect of resilience that applies during the selected timeframe, i.e., the ability to absorb external pressure, time to recover, and the system performance after recovery. This analogy can also be adapted to remove the element of time by measuring the system's functionality against adversary effort instead of time, therefore allowing better comparisons between systems. Nevertheless, care must be taken that the measurement of resilience in the case of zero adversary effort does not read zero, as that does not reflect reality. [14]

They also remark that "cyber-resilience is not a meaningful quantity in the absence of a threat." [14] However, this statement is questionable when considering the global IT outage on July 19th, 2024 due to a faulty update of *CrowdStrike*'s endpoint detection and response (EDR) software. This outage affected millions of devices, impairing the continued operation of critical infrastructure, government agencies, and many more industries like aviation, shipping, and healthcare to name a few. Most importantly, this incident was not caused by an external threat but rather by an internal error of a service provider. It revealed critical weaknesses in contingency planning, supply chain management, and reliance on external services. [15]

The *CrowdStrike* outage therefore emphasizes the importance of developing *cyber resiliency* measures (especially for the *Recover* phase) that not only consider threats by malicious actors but also accidents and faults of trusted third parties and suppliers that may degrade system functionality.

B. Strategic and Operational Perspectives

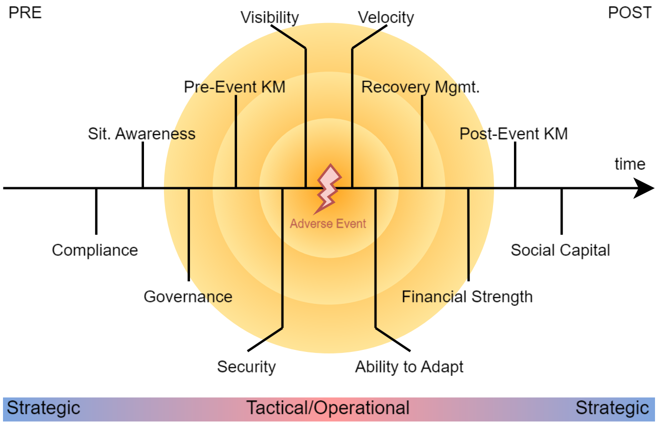

Fig. 2: Impact-Wave Analogy (see [8])

Guerra and Sepúlveda Estay proposed the *impact-wave* analogy (see figure 2), which enables the categorization of solutions addressing challenges in cybersecurity as either *strategic* or *tactical/operational* in nature. *Strategic* mechanisms apply long before or after an adverse event, while *tactical/operational* mechanisms apply during an adverse event. [8]

They identified 12 actions suitable for managing cyber risk (see figure 2). We will now describe these actions in further detail as stated in [8]:

1) *Compliance* addresses the upholding of proper standards and practices as required by the applicable legislation
2) *Situational Awareness* describes the identification of potential risks and vulnerabilities
3) *Governance* defines the deciding authority w.r.t. the IT infrastructure as well as their capabilities
4) *Pre-Event Knowledge Management* facilitates proper training and education to cultivate resilient behavior based on the current understanding of the system's risks and defenses
5) *Security* refers to the defenses and countermeasures of the system protecting its physical and digital assets
6) *Visibility* is the ability to detect an adversary's actions
7) *Velocity* describes the reaction time to disruptions
8) *Ability to Adapt* refers to the ability to adjust the allocation of system resources in response to challenges
9) *Recovery Management* focuses on preparing for previously unknown disruptions, developing contingency plans, and ensuring the availability of resources to execute those plans
10) *Financial Strength* reflects the capacity to overcome changes in cash flow
11) *Post-Event Knowledge Management* describes the ability to learn from past events
12) *Social Capital* involves the social network and relationships with other actors, allowing to learn from them

The actions five through eight are all *operational*. They depend on the specific implementation of the system that dictates which defenses are applicable in the first place. The other actions are *strategic* and thus broadly applicable. However, a clear distinction between *operational* and *strategic* might not always be possible depending on the system itself.

If an adverse event occurs, it creates wave-like ripples (hence the name) affecting every action by demanding improvement in the pre-event stages for prevention (actions 1–6) and testing the capabilities and effectiveness of post-event reactions (actions 7–12). In an ideal case, the defenses before an exploit occurs are enough to outrightly prevent it. [8]

C. The Kill-Chain

In 2011, Hutchins, Cloppert, and Amin from *Lockheed Martin* published a paper [6] that pushed the idea that a cyber attacker needs to perform multiple actions building upon each other to achieve a goal while the defender only needs to detect and interrupt any one of these actions for a successful defense [1]. Taking inspiration from the military doctrine of the US *Air Force*, they adapted their concept of the aptly named *kill-chain*, consisting of the links *Find*, *Fix*, *Track*, *Target*, *Engage*, and *Assess* (*F2T2EA*), to cyberspace. Traditionally, the *kill-chain* identified the steps for successfully locating and engaging a target. It gained its name because a disruption in any of its links would break down the entire process and thus lead to mission failure. Hutchins, Cloppert, and Amin adapted this process to cyber threat actors, defining the steps a modern-day cyber attacker would need to take to successfully engage a cyber system or resource. [6]

The adapted *kill-chain* looks as follows (as per [6]):

1) *Reconnaissance*: Research, identification, and selection of targets.
2) *Weaponization*: Coupling remote access tools or malware with an exploit into a deliverable payload.
3) *Delivery*: Delivery of the payload to the target via physical (e.g., USB flash drives) or non-physical (e.g., email attachments) media.
4) *Exploitation*: After delivery, the weapon executes the intruder's code via exploits or exploitable features in applications, operating systems, or the users.
5) *Installation*: Some means for the adversary to achieve persistence (e.g., via a backdoor).
6) *Command and Control (C2)*: The compromised hosts connect to the attacker's command and control infrastructure, providing the attacker with remote access to and control over the infected machines.
7) *Actions*: After completing the first six stages, the attacker can fulfill their original objective, like data exfiltration.

Like with the original *kill-chain* used by the US *Air Force*, any disruption in one of the links breaks the whole chain. This concept helps identify areas where a defender can act against the adversary. The defender does not need to wait until an actual intrusion occurs. Instead, precautions can be taken even before the initial exploitation. Conversely, the attacker has not won directly after the exploitation, allowing the defenders to disrupt the attacker while installing their command and control infrastructure. Thus, the defender can develop strategies to break each link of the attacker's *kill-chain*. [1]

Identifying the steps an attacker needs to take also helps develop tangible assessment strategies to categorize *cyber resilience* solutions and their effects on different parts of the *kill-chain*. One such example would be the *Course of Action Matrix* [6] or the *Mitre ATT&CK* matrix [7].

D. The Mitre ATT&CK Matrix

Mitre's ATT&CK matrix [7] is a globally accessible database of adversary techniques grounded in real-world data. The *ATT&CK* matrix lists standard attack techniques for each step of an expanded *kill-chain* (now 14 distinct steps). Thus, it offers the defenders, i.e., system designers and architects, a wealth of knowledge on how an attacker might try to compromise a system. Every attack listed also comes with recommendations for mitigation and detection. [7]

Such a knowledge base is ideal for a balanced understanding of the current security landscape if a proper mapping exists from defense mechanisms to links of the *kill-chain*. Then, these mechanisms can be tested against the listed attacks of the *ATT&CK* matrix.

III. Methodology

This section briefly discusses how we systemize and evaluate the extant literature on resilient architectures in III-A. Afterward, we present the literature research process in III-B.

A. Systematization

We must first establish how we systemize each paper to understand the current research landscape concerning *cyber resilient* system architecture in section IV.

In section II-A, we introduced different extant definitions of *cyber resilience*. Linkov and Kott [13] already provide a concise definition that captures the essential concepts of *cyber resilience* (see II-A). However, their definition does not accurately reflect the importance of the organizational context and management of responsibilities (i.e., *governance*). Strong leadership throughout a crisis with clear responsibilities for each actor undoubtedly unifies the response efforts. Therefore, we will extend their definition to include the *govern* function of the *NIST CSF*, which addresses the abovementioned issue.

Hence, we derive the following five core functions of *cyber resilience* from [5], [11], [13]:

1) *Govern*: Establishment of an organizational context with clear roles, responsibilities, authorities, and policies. Additionally, management of human and cyber resources.
2) *Plan*: Identification of critical system functions, assets, and current coverage of defense strategies.
3) *Absorb*: Tolerance to stresses provided by the in-place defenses.
4) *Recover*: Duration of degraded performance and restoration of system functionality.
5) *Adapt*: Changes in the other functions to better defend against unknown attacks enabled by learning from previous disruptions.

Our research focuses on frameworks guiding system designers to solve common problems when designing new or upgrading system architectures and infrastructures to achieve *cyber resilience*. As already stated in II-B, these frameworks can be *strategic* or *operational*. We intend to present more frameworks of the former kind because they broadly apply to various system architectures. However, we also plan to showcase a few *operational* solutions, which necessitates detailing specific implementation scenarios when needed. These *operational* ones might not be broadly applicable. Still, we feel they provide creative solutions that emphasize particular aspects of *resilience* as we have defined it, hence their inclusion.

B. Research

We conducted our research primarily via *IEEE Xplore*[4], *Google Scholar*[5], and *Semantic Scholar*[6], ensuring that the literature we included matched our interpretation of *cyber resilience* from III-A. Additionally, backtracking through the citations of the found literature helped us include relevant research not found by our initial search.

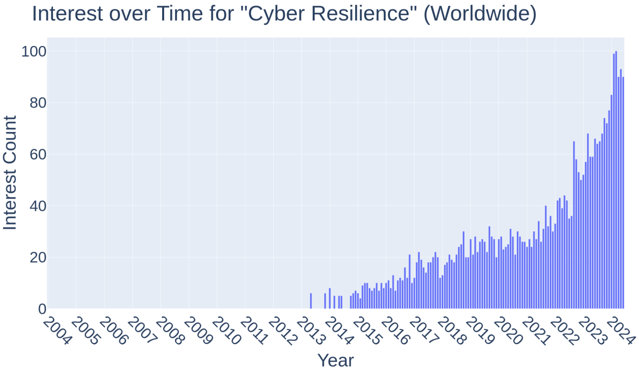

Fig. 3: Google search trends for *cyber resilience* [16]

At this point, it is essential to mention that we didn't limit our search terms to always include *cyber resilience*. We also included similar search terms like *security* and *secure* as a replacement for *resilience*. Therefore, we included research contributing to our designated area of interest, which didn't use today's common terminology. As already theorized in II-A, the term *cyber resilience* started to gain traction after *EO-13636* in 2013 [4] and the NIST *CSF* in 2014 [5]; a claim supported by figure 3, showing *Google's* search analytics for *cyber resilience* [16]. The recency of the term *cyber resilience* was also already noted in [17].

IV. Analysis

In the following, we will present our analysis of the extant literature for resilient architectures, discussing how different frameworks address the *cyber resilience* of architectures enabled by cyber resources. We do this by analyzing frameworks

[4] https://ieeexplore.ieee.org/Xplore/home.jsp
[5] https://scholar.google.com/
[6] https://www.semanticscholar.org/

General Information			Addressed Resilience Phases					Wave Analogy [8]	
Source	Name	Affiliation	Govern	Plan	Absorb	Recover	Adapt	Strategic	Tactical
[11]	CSF	NIST	☒	☒	☒	☒	☒	☒	☒
[18]	FCF	PNNL	☒	☒	☒	☒	☒	☒	☒
[19]	NICE Framework	NIST	☒	□	□	□	☒	☒	□
[20]	CREF	MITRE	□	☒	☒	☒	☒	☒	☒
[21]	Cyber Resiliency Design Principles	MITRE	□	☒	☒	☒	☒	☒	☒
[22]	Resilience Metrics for Cyber Systems	—	□	☒	□	□	☒	☒	□
[23]	Decision-Theoretic Approach to Designing Cyber Resilient Systems	—	□	☒	□	□	☒	☒	□
[14]	How to Measure Cyber-Resilience of a System with Autonomous Agents	—	□	☒	□	□	☒	☒	□
[24]	Cybersecurity: Risk Management Framework and Investment Cost Analysis	—	☒	☒	□	□	☒	☒	□

TABLE II: Categorization of extant frameworks

that guide system designers during development, deployment, and assessments of the underlying system architectures. Table II shows the results of our study, where we summarize which of our previously identified *cyber resilience* phases (see III-A) are addressed by the respective framework. Additionally, we determine if a framework is of *strategic* or *tactical* nature, or both, based on the *wave-analogy* (see II-B).

A. NIST Cyber Security Framework (CSF)

The NIST *CSF* is one of the primary catalysts for *cyber resilience* research today, as already alluded to in III. We based our initial definition of *cyber resilience* from section III-A partly on the findings of the *CSF*. Therefore, it is no surprise that the *CSF* covers all our identified resilience criteria (see table II). The main goal of the framework is to support organizations of all sizes to reduce their cybersecurity risks [11]. It consists of 3 main components:

1) *Core*: Taxonomy of 6 primary functions (*govern*, *identify*, *protect*, *detect*, *respond*, and *recover*; see II-A) and respective categories and subcategories which detail high-level cybersecurity outcomes for managing risks [11]. These categories and subcategories are updated regularly based on the current state of the cyber threat landscape [25].
2) *Organizational Profiles*: These profiles describe the prioritization of the cyber security outcomes from the framework's *core* tailored to the specific needs of the organization's objectives and risks. A profile either reflects the current state of the organization's cybersecurity posture or a desired target state. [11]
3) *Tiers*: The organizational profiles can be further described using the framework's tiers: *partial* (Tier 1), *risk-informed* (Tier 2), *repeatable* (Tier 3), and *adaptive* (Tier 4). These tiers describe the degree to which an organization has implemented its cybersecurity risk management strategy (from an informal and ad-hoc approach to a formal and risk-informed one). [11]

The *CSF* provides pointers to address specific issues like risk assessment, privacy, and even emerging threats from artificial intelligence by referencing the implementation of related frameworks tailored to these sub-problems that comply with the goals of the *CSF* [11]. As such, the *CSF* provides both solutions for the *strategic* and *tactical* domains according to the *wave-analogy*.

B. Facility Cybersecurity Framework (FCF)

Other frameworks base themselves on the standards the *CSF* sets. One such framework is the *Facility Cybersecurity Framework* (*FCF*) [18], which helps to identify security gaps and assess the current and target states of an organization's cybersecurity state via a web-based interface [18], [25]. Additionally, the framework provides cybersecurity training for different scenarios through gamification [18]. Since the framework is directly based on the *CSF*, it also inherits all its properties (see table II).

C. Workforce Framework for Cybersecurity (NICE Framework)

Training the cyber workforce for adaptability to emerging threats is an important task. This is partly addressed by the additions the previously mentioned *FCF* imposes onto the *CSF*. However, while such training possibilities are suitable for existing cyber defense teams, they do little to help organizations establish and maintain talent with training in cyber security.

That is where the *Workforce Framework for Cybersecurity* (*NICE* framework) from the *National Initiative for Cybersecurity Education* fills the gap. The *NICE* framework provides organizations with the appropriate building blocks that describe the skills and knowledge needed to perform cybersecurity tasks. The framework emphasizes that the learning outcomes should apply to current employees, students, and people searching for work. This view stresses that lifelong learning in cybersecurity is essential in addition to facilitating an environment where cybersecurity skills are continuously being improved. Additionally, employers can identify career pathways that synergize with the skillset needed for the organization's defense strategy. [19]

Thus, the *NICE* framework addresses the importance of managing cybersecurity talent, not only during recruitment but also through continuous training to adapt to emerging threats. *Governance* and *adaptation* are therefore central concepts of the *NICE* framework (see table II).

D. Cyber Resiliency Engineering Framework (CREF)

Mitre's Cyber Resiliency Engineering Framework (*CREF*) tries to improve the resilience of system architectures by helping system engineers understand which objectives best support common resilience goals. These objectives consist of *understand*, *prepare*, *prevent*, *constrain*, *continue*, *reconstitute*, *transform*, and *re-architect*. Achieving these objectives and understanding their interplay helps to improve overall resilience. In addition to the previous frameworks, this one considers *cost-effectiveness* as an important factor when implementing *cyber-resilience*. The framework determines three types of costs associated with implementing a resilience measure: the initial cost of installation, operating costs, and consequential costs, which can either be positive or negative. [20]

The *CREF* adds perspective to the existing frameworks by considering cost-effectiveness. However, it falls short of adequately addressing the issues associated with *governance*.

E. Cyber Resiliency Design Principles

Design principles are concise statements describing a fundamental concept of a research domain. For already established domains, a set of widely accepted principles exists. For the security domain, some examples would include *least privilege* or *defense in depth*. However, this is not true for *cyber resilience* due to its recency. *Mitre* now defines a set of *strategic* and *structural* design principles for *cyber resilience* based on their experiences from the *CREF*. *Strategic* design principles are applied during the engineering stages and guide analysis and programmatic decisions, while *structural* design principles affect the system architecture (this differentiation is similar to the *wave-analogy* [8] discussed in II-B). [21]

Nevertheless, these design principles are not universally applicable. They can be selected to fit the underlying architecture and system operations. These principles are generally relevant during the development and redesign phases, where making informed decisions about improving or changing the architecture is essential. Therefore, a basis for analysis and discussion is established by selecting proper *cyber resilience* design principles. [21]

F. Resilience metrics for cyber systems

Linkov et al. present in [22] a methodology that enables the development and organization of resilience metrics for systems enabled by cyber resources. Unique to their approach is the ability to link each metric against national security policy goals (like the ones derived from *EO-13636* [4] as already mentioned in I). This enables the system designer to take tangible actions in terms of *cyber resilience* by properly allocating available resources. [22]

They base their work on the *resilience matrix* [22], which plots the already known NAS resilience phases, *plan*, *absorb*, *recover*, and *adapt* (see II) against the *Network Centric Warfare* (NCW) doctrine, which focuses on creating "shared situational awareness and decentralized decision-making by distributing information across networks operating in physical, information, cognitive, and social domains." [26]

This approach unites the *cyber resilience* efforts by removing the fragmentation of resilience knowledge of separate disciplines (like engineering, environmental management, and cyber security) into one generalized, broadly applicable framework. In their mind, resilience is a property that takes the system as a whole into account (not just technological components but also human and physical components as well). This interconnectedness becomes apparent when looking at an example: During an attack, the system might report the state of every connected machine (*physical domain*). It then forwards the gathered information to the responsible decision-makers (human or non-human, see IV-H), who must decide on an appropriate response (*cognitive domain*). [22], [26]

A system designer can populate this matrix with their current defense mechanisms for each *NAS* event management cycle stage and the *NCW* domain it addresses. This allows for identifying gaps in the current defense strategy while also allowing the designer to understand the interplay between mechanisms. However, the designer should also be mindful when populating the matrix. Adding a metric can affect other event management cycle stages, even across domains. For example, adding sensors might improve the detection ability for adverse conditions, but the designer should also consider where and if this extra information can enhance overall resilience. [22]

G. Decision-Theoretic Approach to Designing Cyber Resilient Systems

While the *resilience matrix* supports system designers by revealing coverage and coupling of implemented defense mechanisms, it cannot quantitatively measure the effectiveness of each mechanism.

Mehta et al. propose a more formal method that utilizes *Markov decision processes* (*MDP*) to compare the achieved level of *cyber resilience* for various system designs and configurations. Such a decision-theoretic approach has the advantage of accounting for the inherent uncertainty in the system's dynamics. The proposed solution is illustrated by applying the developed decision models to a prototypical running example system utilizing *network intrusion detection* and *host reconstitution mechanisms*. [23]

H. How to Measure Cyber-Resilience of a System with Autonomous Agents

Another problem some architectures might face is the introduction of *autonomous agents*. These agents can manifest as a combination of hardware and software, like robots or crewless vehicles, but might also entirely exist as software. Such purely software-based agents might serve the task of modern intrusion detection systems, which may or may not use AI. While adding those agents into systems can benefit automation and reaction speed (speeds that human defenders cannot achieve), they might also pose dangers by increasing the available attack surface an attacker can leverage to gain a foothold within the system. Performance analysis of the effectiveness of these agents is critical today, as more and more systems rely on independent, self-sufficient agents like the ones used in self-driving cars. Ligo, Kott, and Linkov discuss potential methods for assessing the performance of such agents, specifically those focused on cyber defense. They discuss the benefits and drawbacks of different assessment approaches, mainly *qualitative assessments*, *expert analysis*, *simulations and modeling*, *red-teaming & pen-testing*, and *wargaming*. [14]

I. Cybersecurity: Risk Management Framework and Investment Cost Analysis

Much like *Mitre's CREF*, Lee emphasizes in [24] the importance of cost in his risk management framework. The framework has four layers, each addressing specific cyber risk management functions. The first layer is the *cyber ecosystem layer*, whose purpose is to identify the external actors of the system. These include suppliers, customers, consultants, regulatory agencies, and adversaries. A thorough understanding of these ecosystems helps to make informed decisions about the system's infrastructure. Through the *cyber infrastructure layer*, the organization should identify the roles and responsibilities of internal users and employees and critical data and data flows at risk of being targeted by cyberattacks. The framework's core is the third layer, the *cyber risk assessment layer*. Risk assessment, in this case, is achieved by a three-step process: *risk identification*, *risk quantification*, and *cyber investment analysis*. *Risk identification* and *quantification* assign a severity and likelihood of occurrence to every possible threat of the system through the use of a *cyber risk matrix*. This step is crucial for the following *cyberinvestment analysis* that provides heuristics that consider the frequency in which a specific attack type occurs, the cost of that attack succeeding, and the defense probability. This heuristic supports system designers in adding financially sound cyber defenses that act on a sensible proportional scale to an identified threat. Lastly, the *cyber performance layer* suggests the implementation of cyber defenses with, among other things, the proper associated testing, policy development, and personnel training. Additionally, monitoring and logging infrastructure should be available to collect data about the effectiveness of the deployed cyber defenses. [24]

Performing each layer's activities generates feedback for every other layer [24], thus facilitating constant improvement if this framework is adequately adapted.

V. Evaluation

When reviewing table II, we can see that only four of the nine chosen frameworks that fit our research methodology (discussed in section III) address more than three of our five core *cyber resilience* functions. Unsurprisingly, the frameworks that cover most resilience criteria strongly link with the US government through institutions like *NIST*, *PNNL*, or NPOs like *Mitre*. In the case of the *NIST CSF*, it is clear that it addresses every facet of *cyber resilience* because it is the direct response to *EO-13636*, calling for an improvement in cybersecurity for critical infrastructure (see I) [4]. *EO-13800* under the US Trump administration then mandated that US federal networks and critical infrastructure comply with the *CSF* [27]. Thus, the *CSF* must be an exhaustive framework for *cyber resilience* to effectively address the sought-after improvements in cyber security. Something similar might be the case for the frameworks provided by *PNNL* and *Mitre*.

Whereas the research conducted by independent researchers ([22], [23], [14], and [24]) seem to focus their efforts on the *plan* resilience phase. It is at this point important to note that especially the *govern* and *plan* phases almost always reciprocate with the *adapt* phase. The proper identification of assets and management of available resources either succeeded in identifying threats and properly managing defenses or failed. In the latter case, change in the identification and management processes needs to take place to more adequately defend against actual threats that are likely to happen.

In light of this, it makes sense that so much independent research focuses on assessing *cyber resilience* coverage or specific heuristics for cost-effective risk management, which we attributed to the resilience phase *plan*. It is undoubtedly one of the most critical tasks in defending against cyber threats.

However, the apparent lack of independent research for the *absorb* and *recover* phases seems concerning. That doesn't mean that no research in that regard exists. We excluded most from our analysis in section IV because they make strong assumptions about the underlying system architecture, thus making these solutions not broadly applicable. More research discussing the *absorb* and *recover* functions in a more general sense would be desirable. Nevertheless, we will share and briefly summarize some of these more technical papers in section VI since they still provide helpful insight into *cyber resilience* techniques, albeit for architectures displaying specific properties. Afterward, we will discuss our results in light of our peers' research to identify future research possibilities in section VII.

VI. Related Work

This section provides a comprehensive review of related work in the field of *cyber resilience*, encompassing a range of perspectives from state-of-the-art systematizations and conceptual reviews in VI-A to detailed technical solutions for common architectural structures in VI-B. By examining these various contributions, we aim to contextualize our findings within the current landscape of *cyber resilience* research, detailed in VI-C.

A. Systematizations and Reviews

Sepúlveda Estay, Sahay, Barfod, and Jensen comprehensively reviewed 208 articles discussing *cyber resilience* assessment. Their review not only includes a summary of the application areas and the addressed resilience features (specifically from the *wave-analogy*, see II-B) but also investigates general publication data, like country of origin, publishers, citation count, and the collaboration between authors and countries. Through this analysis, they provide information on the current state of the global research landscape in this field. [17]

The study by Lezzi, Lazoi, and Corallo analyzes how cybersecurity issues within the context of *Industry 4.0* are addressed within the current literature. They do this by systematically analyzing 40 articles. These articles are compared based on their characterization/identification of cyber threats to certain industry types/assets and their proposed solutions. Through this comparison, they arrive at a set of recommended countermeasures and guidelines tailored for systems within *Industry 4.0.* [28]

In [29], the authors discuss the evolution of *cyber resilience* frameworks for network security, documenting the progression from traditional cybersecurity solutions to resilience measures capable of addressing dynamic cyber threats. They describe that traditional defenses focus on *prevention*, whereas resilience also includes *detection*, *reaction*, and *recovery*, emphasizing *preparation* and the ability for *adaptation*. In addition to an overview of the historical evolution, they identify these frameworks' critical components and how they are integrated into cyber response strategies. Through a discussion of emerging trends complementing their research, they are able to provide helpful insights for companies seeking to improve their *cyber resilience* posture by "highlighting key components [of cyber resilience] such as risk assessment, threat intelligence, incident response, and recovery planning." [29] Furthermore, they stress the importance of aligning the overall business strategy with *cyber resilience* goals. [29]

The work of Azmi, Tibben, and Win discusses 12 extant *cyber resilience* frameworks, where they derive shared concepts from the selected frameworks. These shared concepts are then further categorized into different perspectives. Through this process, they provide a foundation for creating a more general model for *cyber resilience* frameworks. [30]

Kwon et al. noticed the lack of specific countermeasures when working with current frameworks. Thus, they propose a *Cyber Threat Dictionary* which provides a mapping from the *Mitre ATT&CK* framework (see II-D) to the *NIST CSF*. This mapping enables response teams to select effective countermeasures against attacks from the current threat landscape. [25]

Ibrahim, Valli, McAteer, and Chaudhry examine the adoption of the *NIST CSF* by an undisclosed governmental organization in Western Australia. They found that the *CSF* helped to identify gaps and shortcomings in the current threat mitigation strategy of the organization, thus creating opportunities and objectives for improvement. Implementing the *CSF* was more streamlined than other frameworks. However, they remark that the implementation process could be improved even further. [31]

B. Technical Solutions for Specific Architectures

The *Secure Simplex System Architecture* (*S3A*) proposed in [3] provides a practical solution for retaining a minimum safety level for architectures with *cyber-physical* components. They focus on ensuring that the *cyber-physical* systems operate within the given safety margins by introducing additional trusted hardware that intervenes when infected control hardware sends malicious commands to the machines [3]. This approach also provides an excellent example of the concept of thresholds introduced in section II-A. The *S3A* ensures that a *cyber-physical* system cannot operate outside its specified operational envelope, thus retaining a minimum level of safety that any intervention by an attacker cannot undercut. This further limits the maximum possible damage an attack can inflict on the system.

The research conducted by Kreutz et al. sheds light on the implementation of *cyber resilient* authentication and authorization services. These services are central to almost any IT system. They demonstrate that developing such services with a focus on resilience is achievable through standard practices in security and dependability. Resilience concerning such services primarily refers to fault tolerance and intrusion tolerance. Their findings are supported by a resilient re-implementation of the *OpenID* and *RADIUS* services. [9]

C. Research Landscape

Having introduced some of the research conducted by others, we can see that the majority ([17], [28]–[30]) focuses on the systematization and analysis of existing solutions. These papers form a basis for comparison of existing high-level resilience solutions and approaches, much like this paper does. Few add additional insight and create new opportunities by linking existing concepts to develop solutions focusing on real-world applicability (see [25]). Furthermore, there seem to exist many solutions (e.g., [3], [9]) that demonstrate the specific implementation of some resilience properties in scenarios with well-defined architectures and operational scope. An important area where we found little available research is reviews of the implementation process of the frameworks discussed in IV (see [31]).

VII. Discussion

Our findings from section V indicate that most frameworks that cover all our resilience criteria are from organizations with strong ties to US governmental institutions. The amount of research conducted by independent researchers is usually restricted to one resilience phase, mainly *plan*. However, this allows for a deeper discussion of techniques for that phase, resulting in a wealth of information on, e.g., how to accurately assess *cyber resilience* coverage.

The main contribution of our paper is the introduction, summary, and systematization of some of the most essential *cyber resilience* frameworks, as well as a comprehensive introduction to the core concepts of *cyber resilience*. Through our analysis in IV, we detail which resilience phases are addressed by each framework (see table II). Additionally, we provided an overview of the current research landscape through a brief overview of related work (see VI).

Nevertheless, our study is not without limitations. We will discuss the constraints encountered during our research and their potential impact on our conclusions in section VII-A. Having discussed these shortcomings, we can more accurately outline promising future research topics, aiming to address these limitations in section VII-B

A. Limitations

The core limitation of our research is the amount of evaluated frameworks. Many more exist that address different aspects of *cyber resilience*, e.g., the *NIST SP 800-53* [32], *NIST SP 800-221* [33], *COBIT5* [34], *ISO/IEC 27001:2022* [35], and additional frameworks that were discussed in [17], [28]–[30] to name a few.

Because of this limitation, our work only serves as an introduction to *cyber resilience* and frameworks addressing *cyber-resilient* system design and architectures.

B. Future Work

Since so many additional frameworks that fit our research topic exist (see VI and VII-A), one central question arises: Are so many frameworks even practical? Our research has already shown an overlap of addressed resilience phases between frameworks. It would be interesting to consider the similarities of these frameworks and try to incorporate the core contributions and unique approaches of all those frameworks into one, more generalized framework. This framework would work towards a common goal, achieving and assessing the *cyber resilience* of system architectures enabled by cyber resources.

Such a framework would be rather general in its design and thus be a purely *strategic* approach for managing cyber risks (see II-B). However, the framework could provide more *tactical* solutions to problems faced by specific architectural designs by maintaining an up-to-date database of the current threat landscape and possible mitigation techniques. Through such a system, the framework could provide insight into solutions for absorbing and recovering from adverse events.

To that end, the framework could utilize the already introduced *kill-chain* and *Mitre ATT&CK* matrix (see II-C and II-D). The *kill-chain* gives insight into adversarial behavior, which could be incorporated into the framework's definition of the *plan* function. The *plan* function identifies the system's assets, critical functions, and current resilience coverage (see III-A). Incorporating the *kill-chain* could extend this identification process by including an analysis of which assets are at risk of being instrumented by an attacker at each step of the *kill-chain*. This would provide more insight into each asset's importance and provide grounds for a deeper analysis of mitigation techniques. The *ATT&CK* matrix would then complement this analysis by already providing an extensive database of mitigations for common attacks that an attacker can perform at each stage of the *kill-chain*. Something similar was also already proposed with the *Cyber Threat Dictionary* from [25] (see VI).

Organizations might face the following scenarios when they improve their *cyber resilience* posture. First, they must decide which resilience framework to adopt, which seems impossible due to the many available choices. Second, they might adopt multiple frameworks simultaneously to achieve better resilience coverage. In the latter case, frictions and conflicts between the processes suggested by the used frameworks might exist, further complicating the formation of a proper resilience strategy.

A unified framework might address these issues by providing tuned and coordinated guidelines derived from many frameworks for achieving *cyber resilience*. However, such a framework might entail high complexity and, thus, limit its adoption by smaller organizations and institutions. Nevertheless, we believe its implementation is still worth considering because the research artifacts needed to create said framework would provide valuable insight in and of themselves. We therefore suggest the following 4-step process for making such a framework:

1) *Applicability of Current Frameworks*: The success of currently available frameworks should be evaluated in detail, documenting the implementation process, shortcomings, and limitations of each framework (much like [31]). This would help to identify which resilience concepts and strategies work in practice and which do not.
2) *Filtering Frameworks*: The in step 1 evaluated frameworks should be filtered based on their real-world performance. Strategies and concepts that have merit should be kept, while those found ineffective, overly complex, or not adaptable enough should be discarded.
3) *Unification*: The narrowed list of *cyber resilience* strategies and concepts from steps 1 and 2 should be aligned to form a cohesive strategy for managing *cyber resilience*. Additionally, the *kill-chain* and *ATT&CK* matrix should be incorporated to form the basis for an attack-oriented resilience strategy (similar to [25]).
4) *Deployment & Continuous Improvement*: The unified framework should then be adopted and tested if it achieved its goal, much like the other frameworks from step 1. These tests allow further analysis of the performance of the unified framework and also facilitate the possibility for further improvement.

While the third step, *unification*, might prove to be difficult or even impossible, the first two steps would determine which currently available frameworks achieve proper *cyber resilience* in reality. This alone provides valuable information on the applicability of currently available frameworks and thus conclusively informs about which frameworks are worth adopting.

VIII. Conclusion

This research examined the current state of frameworks that address *cyber-resilient* system design and architecture. To introduce this topic, we first provided a comprehensive understanding of *cyber resilience* and its importance. Additionally, we overviewed closely related concepts, such as the *kill-chain* or the *ATT&CK* matrix, which inform in depth about adversarial behavior and link the resilience concepts to tangible *attack vectors* along with their mitigations. The core of our research was summarizing and systematizing a select amount of frameworks that support the design and operation of resilient architectures. Through this systematization, we noticed that organizations with close ties to the US government create frameworks covering most *cyber resilience* criteria. This is in direct contrast to research conducted by independent entities, which tend to focus on only one or two aspects of *cyber resilience*. While our study was limited to a small subset of currently available frameworks, we noticed that far too many exist. We therefore proposed extensively examining these frameworks' implementation and success to determine their real-world viability. With this knowledge, a unified framework could be created that streamlines the implementation experience and features the *cyber resilience* concepts and strategies that work.

References

[1] R. A. Clarke and R. K. Knake, *The Fifth Domain*. New York, NY: Penguin Press, Jul. 2019.

[2] M. Baezner and P. Robin, "Stuxnet," *CSS Cyberdefense Hotspot Analyses*, 2017.

[3] S. Mohan, S. Bak, E. Betti, H. Yun, L. Sha, and M. Caccamo, "S3a: secure system simplex architecture for enhanced security and robustness of cyber-physical systems," in *Proceedings of the 2nd ACM International Conference on High Confidence Networked Systems*, ser. HiCoNS '13. New York, NY, USA: Association for Computing Machinery, 2013, pp. 65–74. [Online]. Available: https://doi.org/10.1145/2461446.2461456

[4] Executive order – improving critical infrastructure cybersecurity. Visited 2024-05-18. [Online]. Available: https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/executive-order-improving-critical-infrastructure-cybersecurity

[5] NIST. (2014) Visited 2024-05-18. [Online]. Available: https://www.nist.gov/system/files/documents/cyberframework/cybersecurity-framework-021214.pdf

[6] E. Hutchins, M. Cloppert, and R. Amin, "Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains," *Leading Issues in Information Warfare & Security Research*, vol. 1, 01 2011.

[7] MITRE. MITRE ATT&CK. Visited 2024-05-07. [Online]. Available: https://attack.mitre.org/

[8] P. J. G. Guerra and D. A. Sepúlveda Estay, "An impact-wave analogy for managing cyber risks in supply chains," in *2018 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, 2018, pp. 61–65.

[9] D. Kreutz, O. Malichevskyy, E. Feitosa, H. Cunha, R. da Rosa Righi, and D. D. de Macedo, "A cyber-resilient architecture for critical security services," *Journal of Network and Computer Applications*, vol. 63, pp. 173–189, 2016. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1084804516000539

[10] *Disaster Resilience: A National Imperative*. Washington, DC: The National Academies Press, 2012, visited 2024-05-18. [Online]. Available: http://dx.doi.org/10.17226/13457

[11] *The NIST Cybersecurity Framework (CSF) 2.0*, Feb. 2024. [Online]. Available: http://dx.doi.org/10.6028/NIST.CSWP.29

[12] IBM. Visited 2024-05-18. [Online]. Available: https://www.ibm.com/topics/nist

[13] I. Linkov and A. Kott, *Fundamental Concepts of Cyber Resilience: Introduction and Overview*. Cham: Springer International Publishing, 2019, pp. 1–25.

[14] A. K. Ligo, A. Kott, and I. Linkov, "How to measure cyber-resilience of a system with autonomous agents: Approaches and challenges," *IEEE Engineering Management Review*, vol. 49, no. 2, pp. 89–97, 2021.

[15] D. George, "When trust fails: Examining systemic risk in the digital economy from the 2024 crowdstrike outage," 2024. [Online]. Available: https://zenodo.org/doi/10.5281/zenodo.12828222

[16] Google. Visited 2024-06-09. [Online]. Available: https://trends.google.com/trends/explore?date=all&q=Cyber%20Resilience&hl=en

[17] D. A. Sepúlveda Estay, R. Sahay, M. B. Barfod, and C. D. Jensen, "A systematic review of cyber-resilience assessment frameworks," *Computers & Security*, vol. 97, p. 101996, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167404820302698

[18] F. Cybersecurity. Visited 2024-06-13. [Online]. Available: https://facilitycyber.labworks.org/

[19] R. Petersen, D. Santos, M. C. Smith, K. A. Wetzel, and G. Witte, "Workforce framework for cybersecurity (NICE framework)," National Institute of Standards and Technology, Tech. Rep., Nov. 2020. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-181r1

[20] D. Bodeau and R. Graubart, "Cyber resiliency engineering framework," *The MITRE Corporation*, 2011.

[21] ——, "Cyber resiliency design principles," *The MITRE Corporation*, 2017.

[22] I. Linkov, D. Eisenberg, K. Plourde, T. Seager, J. Allen, and A. Kott, "Resilience metrics for cyber systems," *Environment Systems and Decisions*, vol. 33, 12 2013.

[23] V. Mehta, P. D. Rowe, G. Lewis, A. Magalhaes, and M. Kochenderfer, "Decision-theoretic approach to designing cyber resilient systems," in *2016 IEEE 15th International Symposium on Network Computing and Applications (NCA)*, 2016, pp. 302–309.

[24] I. Lee, "Cybersecurity: Risk management framework and investment cost analysis," *Business Horizons*, vol. 64, no. 5, pp. 659–671, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0007681321000240

[25] R. Kwon, T. Ashley, J. Castleberry, P. Mckenzie, and S. N. Gupta Gourisetti, "Cyber threat dictionary using mitre att&ck matrix and nist cybersecurity framework mapping," in *2020 Resilience Week (RWS)*, 2020, pp. 106–112.

[26] I. Linkov, D. Eisenberg, M. Bates, D. Chang, M. Convertino, J. Allen, S. Flynn, and T. Seager, "Measurable resilience for actionable policy," *Environmental science & technology*, vol. 47, 09 2013.

[27] Executive order – strengthening the cybersecurity of federal networks and critical infrastructure. Visited 2024-06-14. [Online]. Available: https://trumpwhitehouse.archives.gov/presidential-actions/presidential-executive-order-strengthening-cybersecurity-federal-networks-critical-infrastructure/

[28] M. Lezzi, M. Lazoi, and A. Corallo, "Cybersecurity for industry 4.0 in the current literature: A reference framework," *Computers in Industry*, vol. 103, pp. 97–110, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0166361518303658

[29] M. O. Akinsanya, C. C. Ekechi, and C. D. Okeke, "The evolution of cyber resilience frameworks in network security: A conceptual analysis," *Computer Science &amp; IT Research Journal*, vol. 5, no. 4, pp. 926–949, Apr. 2024. [Online]. Available: https://www.fepbl.com/index.php/csitrj/article/view/1081

[30] R. Azmi, W. Tibben, and K. T. Win, "Review of cybersecurity frameworks: context and shared concepts," *J. Cyber Policy*, vol. 3, no. 2, pp. 258–283, May 2018.

[31] A. Ibrahim, C. Valli, I. McAteer, and J. Chaudhry, "A security review of local government using NIST CSF: a case study," *J. Supercomput.*, vol. 74, no. 10, pp. 5171–5186, Oct. 2018.

[32] *Security and Privacy Controls for Information Systems and Organizations*, Sep. 2020. [Online]. Available: http://dx.doi.org/10.6028/NIST.SP.800-53r5

[33] S. Quinn, N. Ivy, M. Barrett, L. Feldman, D. Topper, G. Witte, K. Scarfone, R. Gardner, and J. Chua, *Enterprise impact of information and communications technology risk: governing and managing ICT risk programs within an enterprise risk portfolio*, Nov. 2023. [Online]. Available: http://dx.doi.org/10.6028/NIST.SP.800-221

[34] ISACA. (2012) Cobit 5. Visited 2024-06-15. [Online]. Available: https://www.isaca.org/resources/cobit

[35] ISO. (2022) Iso/iec 27001:2022. Visited 2024-06-15. [Online]. Available: https://www.iso.org/standard/27001

SoK: Improving Cyber Resilience in Cloud Computing with Alternative Architectures

Lukas Hertel
School of Computation, Information and Technology
Technical University of Munich
Garching, near Munich, Germany
lukas.hertel@tum.de

***Abstract*—As society increasingly relies on cloud computing, ensuring its resilience against large-scale outages has become critical. Such outages, possibly caused by natural disasters, wars, or cyber-attacks, can disrupt daily life, leading to downtime of critical infrastructure, data loss or theft, and significant economic losses. Given the diverse domains encompassed by cloud computing — ranging from storage to networking — each with unique resiliency characteristics, a comprehensive approach to cloud resiliency is essential. This paper examines alternative cloud architectures specifically designed to enhance system resiliency. First, we define key concepts, including cyber resilience and cloud computing. Subsequently, we provide a concise overview of cyber resilience in cloud environments, emphasizing the necessity of architectural-level resilience as an all-encompassing solution. Building on this premise, we conduct an in-depth evaluation of five promising architectures, assessing their effectiveness in improving resilience and their associated complexity and overhead. Our analysis identifies diversity and replication as key drivers of resilient cloud architectures while noting that all approaches face significant intricacy, posing challenges for widespread adoption. Future work must address this issue to make such architectures feasible in production environments.**

***Index Terms*—cyber resilience, cloud computing, cloud architecture**

I. Introduction

Over the past two decades, cloud computing has undergone significant growth. In 2006, Amazon introduced Amazon Web Services (AWS) with the launch of their Elastic Compute Cloud (EC2) service, which enables organizations to rent large-scale compute capacity on demand [10], [1]. Since its inception, the public cloud computing market has surged to 478 billion U.S. dollars in 2022 and is projected to reach around 2,500 billion U.S. dollars by 2032 [5], [6].

This remarkable growth can be attributed to the numerous benefits cloud computing offers. Companies no longer have to invest in infrastructure up-front and are billed only for actual resource usage [15]. As companies grow and demand increases, resources can easily be scaled up, even while expanding to new regions, without significant hurdles. Furthermore, the maintenance is more manageable. Depending on the service model, the cloud provider abstracts away at least the underlying hardware and possibly much more.

As a result, the security and availability of services offered through the cloud can benefit. Large cloud providers, such as Google, Amazon, and Microsoft, can invest substantially more in the security of their service than most cloud users could. This allows them to employ dedicated security teams responsible for maintaining and implementing security policies, fund in-house and academic security research, or support attractive bug-bounty programs [4].

However, this does not mean that clouds are inherently secure and resilient to adverse conditions. Large cloud providers are highly lucrative targets and the homogeneous structure simplifies attacks once motivated adversaries find a vulnerability despite the investments made by such companies [14], [23].

Furthermore, many businesses, critical infrastructure, and the military are increasingly relying on clouds [8]. Datacenter infrastructure is often located in major cities, where regional catastrophic events, ranging from earthquakes to war, can lead to widespread outages [11].

Due to these reasons, it is necessary and beneficial to protect clouds beyond the usual security best practices and have cloud-specific resilience strategies deployed. Given the multidisciplinary nature of cloud computing, this paper will review and discuss various approaches aimed at enhancing the resilience of cloud computing architectures as comprehensive solutions. These approaches acknowledge the varying resilience levels of lower layers, such as storage, physical data centers, and networks. We argue that improving overall resilience requires accepting these variations and focusing on distributing and heterogenizing the cloud. Although the reviewed works share this core idea, they differ significantly in terms of complexity, performance overhead, and the guarantees they provide. By highlighting these differences, this paper aims to present an overview of this research field and identify remaining limitations to guide future work.

II. Background

This section will offer a more precise definition of cyber resilience and explore the fundamental concepts of cloud computing. We will end with a brief discussion of resiliency challenges associated with cloud environments.

A. Cyber Resilience

To discuss the impact of different architectures on the resilience of a system, we first need to define it. Cyber resiliency is often used in slightly different contexts, depending on the literature. Our work adheres to the definition given

doi: 10.14459/2024md1759337_02

in the second volume of the National Institute of Standards and Technology (NIST) Special Publication 800-160: “Cyber-resilience is the ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises on systems that use or are enabled by cyber resources” [21]. To grasp the core concept of resilience and its distinction from cybersecurity, it is essential to look at its goals in more detail.

Anticipate: In order to improve anticipation, it is crucial to identify vulnerable elements with operational significance in a system. A “state of informed preparedness for adversity” [21] must be maintained. Adversity can take the form of cyber-attacks, and natural disasters, but also unexpectedly high loads. One can achieve a state of preparedness by keeping track of new threats, conducting regular stress tests, and having contingency plans. [7]

Withstand: In the case of an attack, the essential functionality should still be guaranteed to satisfy critical mission needs. To achieve this, secure systems are needed. [7]

Recover: After an incident, normal operation should be restored quickly. This includes the technical aspects of data and functional restoration and handling of possible reputational and legal consequences. [7]

Adapt: Strategies must adapt as technology, operations, and threads evolve. As a prerequisite, changes must be analyzed to keep track of modifying attack surfaces. Concrete enhancements of security policies or risk management strategies might then be needed. [7]

B. Cloud Computing

NIST [17] defines cloud computing as “…a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics, three service models, and four deployment models.” The essential characteristics consist of on-demand self-service, broad network access, resource pooling, rapid elasticity and measured service.

Erl and Monroy [10] provide a more concise definition: “Cloud computing is a specialized form of distributed computing that introduces utilization models for remotely provisioning scalable and measured resources.”

Service Models: Service models range from Infrastructure as a Service (IaaS) over Platform as a Service (PaaS) to Software as a Service (SaaS). In an IaaS model, the consumer can run arbitrary software, such as operating systems, on the provided computing resources. The service provider supplies processing, storage, networks, and other computing resources. On the next level of abstraction, PaaS, the consumer can deploy applications written in a programming language or service supported by the provider. In SaaS, the consumer no longer manages the application or underlying infrastructure but only uses the application through a client device. An example would be web-based email or Google Docs. [17]

TABLE I
Cloud Provider Security Responsibilities According to Google [4]

	On-prem	**IaaS**	**PaaS**	**SaaS**
Content				
Access Policy				
Usage				✓
Deployment				✓
Web application security				✓
Identity			✓	✓
Operations			✓	✓
Access and authentication			✓	✓
Network security			✓	✓
Guest OS			✓	✓
Audit logging		✓	✓	✓
Network		✓	✓	✓
Storage & Encryption		✓	✓	✓
Hardened kernel & IPC		✓	✓	✓
Boot		✓	✓	✓
Hardware		✓	✓	✓

Depending on the service model, different parties are responsible for the security and resilience of system components. Table I gives an overview of various security considerations and whether the cloud customer or the provider is responsible for them.

Deployment Models: Cloud infrastructure can be provisioned in four distinct ways. In a *private cloud*, a single organization is the cloud user. Nevertheless, it can still be managed or owned by a third party and be on or off premises. The infrastructure of *community clouds* is used by a community, of which one or more organizations can also be the owner. A *public cloud* is operated by a cloud provider, such as a business or academic organization, and can be used by the general public. The fourth deployment model is the *hybrid cloud*, where two cloud infrastructures are connected, e.g., a private and public cloud. As a result, data and applications can be moved between them.

C. Resilience Concerns in Cloud Computing

Colman-Meixner et al. [9] categorize cloud service disruptions, along with their associated costs.

Causes: Intentional or unintentional human errors can lead to failures in a service; whereas mistakes can mostly be relatively easily reverted, attacks are often more targeted, leading to more significant damage. Next are software failures caused either by bugs or malicious software. Physical failures are outages of physical cloud components and spread to higher levels, at some point reaching the running application. The final category of causes includes disasters, which can be classified into natural disasters, such as floods and earthquakes, and human-instigated disasters, such as those involving weapons of mass destruction. [9]

Impact: The consequences of a service disruption can vary, ranging from masked disruption, causing no noticeable impairment in service, to instances of degraded service, unreachable service, or even corrupted service. Only in the last scenario is recovery no longer fully possible, and critical data lost. [9]

Consequences: Costs can be classified into three categories. First, repair costs may arise in the case of damaged hardware. Second are penalty costs, paid by the cloud provider to the user for not meeting availability targets. The most significant portion of costs is attributed to business revenue losses resulting from missed opportunities. [9]

III. Related Work

It is challenging to find papers discussing cyber resilience in the cloud using a definition similar to the one provided by NIST, which was published in 2021. This is because the concept of resilience predates NIST's definition, leading many papers to either use different interpretations of resilience or employ other terms for concepts that align closely with the goals highlighted by NIST.

The work by Khorshed, Ali, and Wasimi [15] provides an overview of the general security landscape of cloud computing as of 2012. Even though the authors do not directly use the term "resilience", they review techniques related to anticipation, detection, and adaption. They identify the following seven primary threats which are specific to cloud computing [15]:

1) Abuse of nefarious use of cloud computing
2) Insecure application programming interfaces
3) Malicious insiders
4) Shared technology vulnerabilities
5) Data loss/leakage
6) Account, service, and traffic hijacking
7) Unknown risk profile, due to lack of transparency

They discuss these points primarily from the perspective of a cloud customer, emphasizing that many of the threats stem from the customer's lack of control over the cloud infrastructure.

Colman-Meixner et al. [9] provide a comprehensive survey on resiliency techniques. After a detailed discussion on cloud computing concepts, they continue discussing service disruptions and categorizing techniques to improve cloud resiliency. Failure forecasting, protection — either through replication or checkpointing — and restoration are the three main strategies they identified in their review. Replication can be either active, where all replicas stay updated (such as redundant array of inexpensive disks (RAID)), or passive, where the component is replaced with a backup on failure. Furthermore, they distinguish on a high level between resiliency in cloud computing infrastructure and resiliency in cloud computing applications. For infrastructure resiliency, they group the reviewed literature into the resiliency of data centers, communication networks, and the resiliency of middleware. Our work focuses on the resiliency of the middleware, within which Colman-Meixner et al. highlight the benefit of efficient integration of multiple resiliency techniques across various layers. The primary issue identified with resiliency solutions on the middleware is their high complexity and the requirement for sophisticated measurement tools.

Based on their work, Welsh and Benkhelifa [27] provide a review of techniques aimed solely at cloud service providers to deliver high-resiliency services to their customers. They also classify the work into layers, focusing on the differences between centralized and decentralized clouds, the target cloud delivery model, and the techniques used to achieve resilience. They highlight that most resiliency approaches are based on some form of redundancy, which is inherently expensive. Due to the varying use cases of clouds, adaptive resilience is needed, which selects appropriate techniques for hardening against identified weaknesses. Other than Colman-Meixner et al. [9], they focus only on techniques applicable to a cloud environment and ignore general resiliency approaches.

Prokhorenko and Babar [20] use a slightly different definition of resilience in their survey on architectural resilience. According to them, resilience constitutes the process necessary for attaining the actual goal of high reliability. For a system to be resilient, it must be trustworthy, efficient, and have high capacity. Trustworthiness denotes the system's ability to safeguard against unauthorized workloads, whereas capacity refers to its ability to manage substantial legitimate workloads. Trustworthiness extends beyond individual nodes; it is equally critical to ensure the security of communication between nodes. Encryption and authentication are classic solutions to this problem. To support high loads and consequently ensure high capacity, techniques such as load balancing and failure handling are necessary. The final pillar identified by the authors is efficiency, aimed at overcoming the physical constraints of devices. This is achieved mainly by cyber-foraging, a technique where external computational resources are used to augment the capabilities of a mobile or resource-constrained device — or node in the case of cloud computing.

Prokhorenko and Babar distinguish between different architectures on a coordinate system, with node resources on one axis and proximity to the user on the other. Central clouds, with virtually endless resources, are on one end, whereas fog-based architectures and peer-to-peer systems are on the other. Between those two extremes are cloudlets and edge nodes. They highlight the high latency issues in central clouds due to the considerable distance to powerful nodes and the drawback of a single point of failure. According to the authors, these problems are gradually resolved when moving to an architecture with less powerful but more nodes near each other — and the user.

IV. Methods

The fundamental principle behind cloud computing is distributed computing [10], which can be loosely characterized as "...a collection of independent computers that appears to its users as a single coherent system" [26]. Distributed computing is, therefore, a considerably older concept, dating back to the beginnings of the Internet in the 1960s. Key concepts such as communication, consistency and replication, fault tolerance, and security in distributed systems apply equally to cloud systems. These topics serve as a valuable entry point into the study of resiliency in cloud environments, having been extensively covered in literature, including Tanenbaum and van Steen's "Distributed Systems: Principles and Paradigms" [26].

Based on this, we looked for contributions more specifically covering resilience in combination with parallel computing. Back in 1984, Svobodova [25] already used the term resilient distributed systems combined with atomic actions for more fine-grained recoverability, resulting in faster resumptions after a failure. She also argues for the generally enhanced reliability of distributed systems due to the physical disconnection of processes and resources and the ability to use other nodes to finish computations in the case of failing nodes.

Other reviewed literature includes the work by Motresor et al. [18], making the case for simple nodes in peer-to-peer networks as opposed to large and complex centralized distributed systems, and Kyamakya et al. [16] using the term survivability in a very similar fashion to the definition of resiliency given by NIST in [21]. They highlight the need for robust systems, secure end-to-end networking, and system validation and verification to achieve these survivability goals without going into much detail.

However, much of the available literature was either of insufficient quality or focused on overly specific solutions to narrow problems, rendering them less relevant for our subsequent survey on cloud resiliency. Therefore, we continued by focusing more heavily on cloud-specific literature.

Cyber resilience in cloud computing is the responsibility of the provider and the user. To understand the current security and resilience practices employed in state-of-the-art clouds, we reviewed the documentation provided by large cloud providers such as Google Cloud [4] and AWS [3]. Even though technical details were sparse and issues and limitations were hardly discussed, it gave insights into the methods deployed and the strengths of large public clouds.

For a more academic view, we selected papers discussing cloud computing in combination with cyber resilience by querying research search engines such as Consensus, Semantic Scholar, and Google Scholar. As cloud computing resilience affects many aspects of resilience problems, the literature only scraped the surface of many different elements to be considered in cloud computing.

One can look at the issue of cloud computing on different layers [27]. The lowest layer handles physical resilience on a data center level. On the next higher level, the abstraction of virtual resources takes place. Examples would be replicating data to achieve higher storage resilience or having multiple data center locations hidden through virtual networking. The third layer handles cloud management, trying to improve resilience on the middleware level.

Resilience improvements on lower levels also help achieve better resilience on a higher level. Nevertheless, a highly resilient storage protocol, for example, does not help much with a highly unreliable network. For this reason, we focused our search on new architectures promising high resilience despite unreliable lower layers.

To gain a comprehensive understanding of these approaches, we initially reviewed the majority of the papers analyzed by Welsh and Benkhelifa [27]. We excluded [24] due to ambiguity regarding the authors' specific contributions, and [11] because it does not introduce any specific concept and primarily focuses on making a case for community clouds in general to achieve higher resilience. We then extended our search for new approaches and updates on the ones discussed in [27]. By using the "cited by" feature on Google Scholar, we searched for all papers that cited these proposals and filtered them to include only the ones published after 2020. However, the results were very sparse and mainly consisted of other surveys with different focuses.

Furthermore, we sought to find updates on the papers we discuss in the next section, as several indicated an intention to implement their theoretical ideas. For MEERKATS, discussed in Section V-B, we found a final report [12] published by the US Air Force Research Laboratory in 2016, four years after the release of the original paper proposing MEERKATS. We discovered the report through an extensive Google search conducted shortly before the deadline for this paper rather than through academic research databases. Although this time constraint limited our ability to discuss the report in full detail, it contains valuable insights, including several modifications and proof-of-concept implementations, along with their evaluations. The surveys by Colman-Meixner et al. [9], and Welsh and Benkhelifa [27], both published after the final report's release, refer only to the original proposal. So will we if not explicitly stated otherwise. Despite our efforts, we found no new information or updates on the other technologies.

Finally, we employed the index terms from the existing data, again in combination with a filter to show only literature published after 2020 to search for new research in this field. By doing so, we encountered the work by Moura and Hutchison [19] on Resilient Edge Cloud Systems (RECS). They introduce a novel resilience approach to networks in the form of a Software Defined Network (SDN)-based system, which can handle communication failures and provides load balancing on a server and network level. Furthermore, their work includes an implementation and exhaustive evaluation thereof. However, the complexity of the approach and the requisite knowledge of computer networking presented significant challenges, given our current level of understanding. Despite our efforts to delve deeper into the material, including consulting supplementary resources, we decided not to include this approach in further discussions. This decision highlights the complexity involved in implementing resiliency solutions in the cloud and the extensive expertise required across a broad range of computer science disciplines.

Most of the additional work we identified through this method consisted either of further surveys or contributions that did not meet our inclusion criteria, due to issues such as clarity, accuracy, or relevance, and were therefore excluded from our analysis.

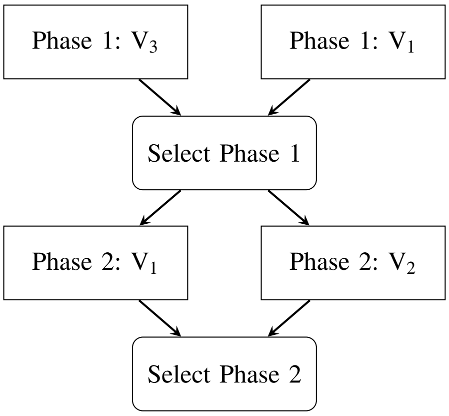

Fig. 1. The core idea behind BioRAC. Every phase of a task runs a different version. Due to additional replication and acceptance tests, attacks can be detected by the CAM. Adapted from [13].

V. Results

This section introduces the reviewed literature, of which an overview is given in Table II. The discussed approaches are ordered chronologically. A detailed comparison of their impact on resilience and feasibility follows in the next section.

A. BioRAC

Hariri et al. [13] introduce with Biologically-inspired Resilient Automatic Cloud (BioRAC) a new Cell-Oriented Architecture, which conceptually consists of organisms and cells. During execution, an organism associates itself with specific mission or application roles and enlists cells operating on one or more host machines to perform its functionality. The cells form a cluster, and the needed resources are assigned to them.

Furthermore, they present the Cooperative Autonomic Manager (CAM), which is responsible for realizing the automatic resiliency and is part of every cluster. Based on generated logs and anomalous events, it updates the underlying structure of cells.

The authors state that this approach is employed not solely for the purpose of redundancy but also to ensure diversity. Cells periodically “shuffle” functionally equivalent code modules, such as algorithms implemented in different programming languages. Figure 1 shows an example. First, two different code modules execute the first phase of a task, both with different versions. During the second phase, different versions are chosen again. The heart of the algorithm lies in the selection phases when the CAM executes an acceptance test, such as an agreement protocol, ensuring the cell is running correctly and no attack has succeeded. If a test should fail, the result provided by another cell is used and forwarded to the next phase. Cells can also be isolated immediately upon detecting discrepancies, depending on the policy. Based on the identified vulnerable elements of a system, as discussed in Section II-A, the amount of shuffling and replications can be tailored to a reasonable tradeoff between resilience and effort.

BioRAC automatically stores negative acceptance test results and other incidents. This allows the CAM to change security policies during runtime and adapt to new threads without manual intervention.

As a result, the authors show how applying BioRAC to a cloud will make it considerably more difficult for attackers to find a way into a system. Not only would an adversary need to exploit possible weaknesses in a system that can change from execution to execution, but due to the replication, the vulnerability would have to be present in multiple versions.

B. MEERKATS

The Maintaining Enterprise Resiliency via Kaleidoscopic Adaption and Transformation of Software Services (MEERKATS) [14], [2], [12] cloud architecture consists of many different components with the goal of having “an environment for cloud services that constantly changes along several dimensions, toward creating an unpredictable target for an adversary” [14]. One of the ten newly introduced components, DREME [8], will be discussed in more detail in the next chapter. In this section, we will review the overall ideas without discussing any in great detail.

The authors argue that while cloud computing providers need to be capable of handling the highest possible load, a significant portion of their resources remain idle most of the time. These resources can be utilized to enhance the resilience of critical components and adapt to evolving threats by analyzing past incidents. Furthermore, continuous service and data mutation can be employed, akin to the approach discussed in BioRAC.

At the core of the MEERKATS system sits a component responsible for gathering information from every other component, operating on multiple layers of abstraction. The goal is to detect unusual activity, which might indicate an ongoing attack. This works particularly well with another component responsible for tracking the flow of information between processes and systems on a byte level.

Another introduced component can rewrite running binaries and instrument them with hardening techniques. These can be in the form of buffer overflow protection, NULL pointer protection, or even Control Flow Integrity (CFI). To recover from bugs that lead to a crash or exploit, MEERKATS can use existing, regular code as exception handlers. Furthermore, the component maintains read and write sets for instructions to detect illegal memory accesses. Through an additional component, they show that running program replicas with enforced thread schedule variations allows for the identification of concurrency vulnerabilities.

Should an adversary be able to circumvent these proactive measures, she would be facing a rapidly mutating software and network state. With the help of lightweight virtualization and traffic rerouting, applications are migrated regularly to new machines. If the system detects or suspects a compromised machine, MEERKATS promptly and securely transfers all data to an alternative location.

Moreover, MEERKATS can generate “deceptive” information to mislead attackers, similar to honeypots. The goal is

TABLE II
Overview of Reviewed Approaches to Resilient Cloud Architectures

Approach	Short Description	Distributed	Single Point of Failure	Limitation	Implementation Provided
BioRAC [13]	Replicated, diverse cells executing a single task	No	Yes, the CAM	High complexity and overhead	No
MEERKATS [14]	Collection of components increasing resilience	Yes	-	High complexity and overhead	Partially
DREME [8]	Replication of diverse services	Unspecified	Yes, the I/O redirector	High resource overhead	Yes
DefCloud [23]	Infrastructure and process level diversity	Yes	No	High complexity	No
Cloudy [22]	Provides virtual clouds on distributed clouds	Yes	Yes, the CSM	High complexity	Yes

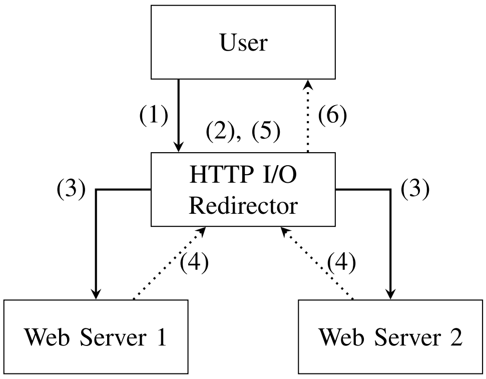

Fig. 2. In the DREME system, a user request is first sent to the HTTP I/O redirector (1). If the protocol at a higher level associates state with the user, such as a nonce in challenge-response authentication, the user request must be modified (2). In (3) and (4), the servers handle the request and send a response. After selecting one response in (5), the result is transparently sent back to the user (6). Adapted from [8].

to use the information-gathering component to automatically generate believable data while monitoring the traps in parallel.

Lastly, the authors of MEERKATS envision custom hardware accelerators, especially for moving data faster between nodes. Accelerators would also benefit other components to reduce or altogether avoid slowdowns.

C. DREME

The Diversified Replica Execution and Monitoring Environment (DREME) [8] framework is a part of the MEERKATS [14] cloud architecture and allows the execution of synchronized replicas. Even though it is an integral component in the original proposal of MEERKATS, it is no longer mentioned in its final report [12]. The replicas are based on lightweight containers, each of which can have different hardening techniques enabled, have a different environment, or even be an entirely different program. For example, the authors propose replicas of a database application might once be realized with PostgreSQL and once with MySQL.

Figure 2 shows how this works in practice. When a user sends a request, this request is first handled by the I/O redirector. The I/O redirector is aware of the replicas and forwards the request appropriately. Here, we have two replicas of a web server. In general, there can be arbitrarily many replicas, which might use replicated services themselves. The web server handles the request as usual and returns the result to the I/O redirector. After choosing one of the responses or having merged them, the result is sent back to the user.

The authors express concern that issues arise with this concept because the web servers might introduce different states for a single user. For example, the web server might generate a unique, random token for authenticated users, which they can send as part of their HTTP requests. Therefore, user requests must be modified and tracked by the I/O redirector, potentially introducing considerable complexity.

They show that, by monitoring the resource utilization of the replicas, DREME can detect attacks such as Denial of Service (DoS). Currently, the host on which a replica runs tracks the container's CPU usage, memory usage, and network traffic. The authors then configured MEERKATS to terminate any application that causes a metric to surpass a predefined threshold.

D. DefCloud

DefCloud's [23] contributions consist of two parts. First, they ensure infrastructure diversity in data centers and distribute them among multiple physical locations. The second part consists of program diversity, similar to the previously discussed approaches.

To mitigate a single vulnerability that affects an entire data center, DefCloud proposes a diverse data center architecture. Figure 3 shows an exemplary schematic of such a data center. The infrastructure is divided into k subtrees, whereas different subtrees use different technology. (In Figure 3 is $k = 2$.) One subtree might run Linux on AMD processors and is interconnected with Cisco switches, whereas the other subtree runs Windows machines on IBM processors with a Juniper network infrastructure. Malware is then limited to a single subtree, considerably decreasing the likelihood of it spreading.

The authors argue that despite the robust protection of data centers, they remain vulnerable to natural disasters, war, or terrorist attacks. Consequently, they advocate for distributing data centers across multiple locations to achieve high resilience.

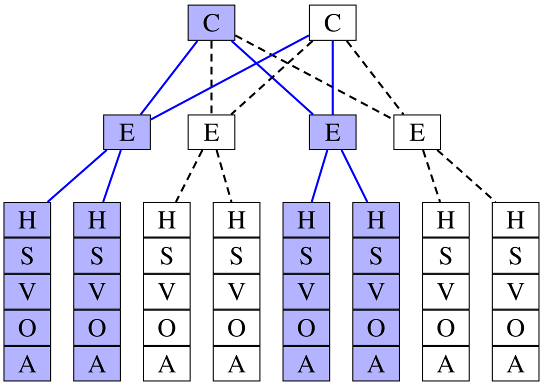

Fig. 3. Diverse data center architecture, as proposed by DefCloud. Blue nodes and edges represent a different technology stack. Core (C) switches connect to edge (E) switches and those, in turn, to hosts (H). A host consists of storage management (S), a virtual machine monitor (V), an operating system (O), and applications (A). Adapted from [23].

When choosing locations, diversity has to continue to play a significant role in reducing the chances of a problem impacting multiple sites. For instance, it is advisable to connect data centers to different Internet service providers (ISPs) and ensure that not several locations are susceptible to the same natural disaster, such as flooding.

Additionally, DefCloud introduces spatial and temporal process diversity to achieve diversity on the application layer, very similar to the approach discussed in Section V-C with DREME. In contrast to DREME, the compiler does this automatically by "slicing" the program into multiple smaller programs, computing intermediate states. During execution, the unmodified program compares its state with the states produced by the slices. These hardened slices are generally faster than running a complete binary replica and allow different program parts to be fortified to varying degrees based on their vulnerability and importance. Furthermore, a just-in-time (JIT) compiler, in combination with a binary decompiler, periodically analyzes and changes the layout of the program's text and data segment.

E. Cloudy

Cloudy [22] is a software distribution based on Debian Linux, trying to improve the resilience of community clouds. It delivers tailored decentralized solutions for managing networks and discovering services for an IaaS or PaaS model. Services run as Docker containers. Users can enable multiple additional apps through a web interface or install new containers.

The definition of a community cloud provided in Section II-B has to be narrowed down slightly in the case of Cloudy, as the infrastructure runs only on the edge of a network and, in parts, even on a community network. The devices are low-powered computers, called resource devices (RDs), in homes or offices instead of large data centers, and little traffic goes over the internet. This results in a heterogeneous cloud in terms of its network topology, networking devices, and RD performance and reliability.

According to the authors, a cloud user sends a deployment request to the Cloud Service Manager (CSM) when he wants to deploy a service. Depending on the users' needs, the degree of replication and other resilience parameters are defined in such a request. If the request is accepted, the CSM generates a *resource slice*, consisting of resources (compute power, storage, etc.) and a virtualized network. The CSM then continuously monitors resource usage and reconfigures slices dynamically when needed. For example, suppose a RD becomes overloaded. In that case, the CSM moves virtual machines to another RD — in the case of network congestion, the topology of a slice might have to be adapted to decrease the load on specific paths.

To stay operational despite link failures or outages of RDs, the CSM ensures that the RDs are sufficiently geographically distributed and replicated. This allows the surviving RDs to reconfigure links in the network to reroute traffic to new RDs, allowing the slice to stay operational.

They also propose a new distributed algorithm to support the migration of services to new locations. As the communication link is unreliable, they propose a layered controller, building upon the link a failure detector, followed by a generic consensus service and an agreement protocol to decide group membership and leader election. Upon these layers, a reconfiguration algorithm can then be built.

VI. Evaluation

The reviewed solutions exhibit significant variations in the resilience properties they offer. In this section, we evaluate these properties alongside their associated costs, specifically in terms of complexity and overhead.

A. Resilience

We individually examine the distinct cyber resilience goals to compare the resiliency of the discussed approaches. These objectives are not binary values; a system's ability to withstand attacks is not an all-or-nothing measure but rather can be assessed on a continuum [27]. Furthermore, the extent to which a resilience goal is achieved depends on the context, including users' needs and the underlying infrastructure. For instance, a community cloud for media sharing, composed of highly unreliable servers, may require a more resilient architecture to attain the same level of resilience as the media-sharing service of a big company operating on a large cloud provider's infrastructure.

1) Anticipate: Identifying vulnerable components within a system and tracking evolving threats can only be partially automated. Data centers, for instance, can hardly accurately assess the risk posed by terrorism, flooding, or advanced cyberattacks. However, it is feasible to identify critical technical components and monitor for ongoing attacks or attempted breaches. Therefore, intrusion detection systems (IDSs) can also be seen as a form of anticipation, allowing us to prepare and react to adverse conditions.

Only MEERKATS [14], and BioRAC [13] have exhaustive monitoring systems for detecting threats and attacks. MEERKATS has a complex IDS, monitoring information flow and anomaly alerts, accompanied by a component responsible for deceptive information generation for honey pots. These are then automatically monitored. The authors highlight that the system's efforts are concentrated on mission-critical applications. They do not specify how such applications are identified, leaving it unclear how one marks an application as critical.

The authors of BioRAC introduce a monitoring system as well, stating that it can use internal and "publically reported" alerts to identify new threats and attacks.

The minimal support for anticipation goals in these cloud system architectures is unsurprising, as all NIST [21] proposed improvements involve manual tasks. Activities such as conducting stress tests, creating contingency plans, and running red teaming scenarios remain essential.

2) Withstand: The main focus of all the reviewed work lies in the ability to withstand adverse conditions. Ensuring critical mission needs can be achieved through various methods. One primary approach is replication, which can be implemented at multiple levels: replicating data, computation, network components, or even entire data centers. Second, many techniques can be categorized as a form of hardening. One can instrument binaries with additional checks to detect exploitation attempts or increase the number of guards protecting a facility. The last common practice to improve the capabilities to withstand adverse conditions in a cloud system is the usage of diversification to limit the impact and spreading of attacks and errors.

All of the discussed solutions, except for Cloudy [22], run diverse and active replicas of applications. In the case of BioRAC, these have to be manually developed, whereas they can be generated automatically in the case of MEERKATS with DREME [8] and are only generated automatically in the case of DefCloud [23]. DREME runs entirely different applications, whereas BioRAC switches versions during phases of a task, and DefCloud uses diverse replicas to calculate critical state multiple times. In the case of a programming bug, different compilers or compilation flags do not automatically resolve the issue. As a result, DefCloud is not necessarily protected, as the replicas might face the same fault. The other approaches do not face this issue, as entirely new versions are unlikely to contain the same bug. In any case, exploiting the potential vulnerability is considerably more challenging, as the exploit has to work on both versions simultaneously. Even running the same binary twice will make many binary exploitation attempts unsuccessful due to the randomization techniques deployed by the operating system (OS), such as address space layout randomization (ASLR) or stack canaries.

Cloudy and DefCloud are the only two papers explicitly discussing hardware replication, although from slightly different angles. The authors of DefCloud show the importance of diverse replicas inside a data center and the replication of entire data centers. By doing so, DefCloud is the only architecture that explicitly handles resiliency in the event of regional disasters, where an entire data center might become non-operational. In the case of Cloudy, replicated hardware is required due to the very unreliable infrastructure. Consider an essential node within a home-based distributed network, where various factors can disrupt connectivity (for instance, a fuse might blow due to the simultaneous use of high-power appliances like an oven and a hairdryer). Incidents of such nature can render large portions of the network unreachable. Replicating these nodes is crucial to support and mitigate the impact of such outages.

TABLE III
OVERVIEW OF THE DIFFERENT FOCUSES TO WITHSTAND ADVERSE CONDITIONS, CATEGORIZED BY ARCHITECTURE.

	Replication	**Hardening**	**Diversification**
BioRAC [13]	++		++
MEERKATS [14]	++	+++	++
DREME [8]	++	+	++
DefCloud [23]	+++	+	+++
Cloudy [22]	++		+

MEERKATS, including DREME, employs the most sophisticated hardening techniques comprising various concepts. Besides binary rewriting, they propose continuous migration and alteration of the software, data, and network state. DefCloud follows a similar, though more constrained, approach, aiming to randomize the software layout over time.

All reviewed approaches use some sort of diversification to improve the capabilities to withstand adverse conditions. As already discussed, BioRAC, MEERKATS, DREME, and DefCloud diversify and replicate software to find discrepancies during execution and, therefore, hope to identify exploitation attempts. But even in the absence of replication, DefCloud and Cloudy show that diversification enhances resiliency. This is because adverse conditions are likely to affect only a portion of the system rather than spreading to all distinct components. For large cloud providers with a monoculture infrastructure, this issue is particularly critical, as a single point of failure can lead to widespread outages.

Table III provides an overview of the different approaches taken by the discussed architectures. The more plus symbols a technology has, the more it relies on this general approach.

3) Recover: Restoring regular operation after an incident is not explicitly mentioned in any of the reviewed papers. All previously discussed techniques on the withstand capabilities of a system try to avoid large system breaches. Through a high level of replication, one arguably hopes that there will always be a working replica, from which, in the case of an incident, new replicas can be created again.

On a smaller scale, MEERKATS features a self-healing software component that supports error recovery by utilizing existing code as exception handlers. However, this is also more of a hardening technique to stop vulnerabilities from being exploited as a technique to recover from an actual incident.

Cloudy aims to migrate services after the failure of infrastructure, which can be seen as a form of recovery, as the service might experience a period of degraded service until the migration has been completed. It should be noted that in the case of the failure of the main component, the CSM, the system can no longer recover automatically.

4) Adapt: Similar to the objectives of anticipation, this process can only be partially automated. Adapting strategies in response to evolving technologies and threats remains a manual task, necessitating the creation of specific business processes. The architecture is capable of adapting only to ongoing and measurable adverse conditions. Consequently, automatic anticipation and adaptation are feasible only within a short time frame.

The authors of BioRAC state that their system is capable of "self-adaption", based on gathered information by their monitoring systems. Based on this information, the number of phases can be increased or decreased — resulting in more or less checkpoints respectively, or the number of versions can be altered.

Similarly, MEERKATS also responds to the findings of its comprehensive monitoring. It then alters its architecture, as we have discussed when talking about MEERKATS' withstand strategies.

In Cloudy, the architecture adapts to underlying infrastructure changes but cannot respond to threats or ongoing attacks. DefCloud has no automatic adaption capabilities.

B. Complexity

All of the approaches suffer from a high degree of complexity. Unfortunately, the more complex components are typically only superficially explained, leaving the precise implementation unclear. In this section, we nevertheless try to identify the problematic parts in all systems and compare them.

BioRAC's complexity comes mainly from its CAM, which manages the replication based on gathered information. The versions and acceptance tests between phases are created manually and are not part of the system's inherent complexity. However, managing these versions, especially deciding what needs to be replicated based on logs and "publically" reported information, is not trivial.

MEERKATS builds on these ideas and includes many more components over multiple layers, all communicating and influencing each other. Each of those introduces considerable complexity itself. For example, the system is supposed to generate "believable deceptive information" for honeypots automatically. The authors give as an example the automatic generation of documents with fake credit card numbers, account numbers, usernames, passwords, and URLs to sites, which are then automatically monitored for input of such automatically generated data. Other components of the MEERKATS system face similar complexity.

For DREME, the authors provide a limited example implementation, which allows a more accurate discussion of its complexity. The critical part of the system is the I/O redirector, whose complexity closely correlates to the complexity of the application it redirects to. The I/O redirector is relatively straightforward for a stateless application, but the complexity increases as the managed state grows.

DefCloud's complexity can be divided into additional complexity due to infrastructure and program diversity. We estimate a reasonable complexity overhead for the infrastructure, as it only introduces the need to manage two different technology stacks, which might be very expensive but doable. The complexity of the spatial diversity component of the program diversity is arguably similar to the approaches seen with DREME and BioRAC. Conversely, the temporal diversity for programs looks considerably more involved. A binary decompiler is supposed to recreate high-level semantic information of programs on the fly. Such decompilers are highly complex and often inaccurate. A JIT compiler is then supposed, based on the decompiled binary, to introduce new entropy in the program. Similar to decompilers, JIT compilers are highly complex and often produce suboptimal code, as they cannot spend the same amount of time on optimization as the ahead-of-time compiler did.

Running cloud services on unreliable infrastructure fundamentally requires a certain degree of complexity, as many edge cases in other systems are common in such a scenario. The approach chosen by Cloudy handles that complexity in a reasonable fashion, abstracting complexity with the help of a layered architecture.

C. Overhead

The overhead that occurs when deploying a specific architecture is twofold. The first is the overhead at runtime, which can be increased memory usage, CPU usage, or slower execution times. On the other hand, when looking at these overheads, we identified approaches that use similar techniques but pay the price for those upfront. They reduce the overhead at runtime by compiling multiple versions of an application before production rather than applying hardening techniques dynamically during execution. Although these approaches will run faster when released, neglecting considerably longer compile times or the need to develop multiple versions would be an unfair comparison.

The resource usage overhead at runtime for BioRAC, DREME, and DefCloud depends on the degree of replication. Nevertheless, BioRAC and DefCloud have the advantage that they do not need to replicate the entire application and can harden different parts — phases in the case of BioRAC and slices in the case of DefCloud — more than others. This will make them more resource-efficient and flexible, at the cost of running more acceptance tests instead of only one when the final result is produced. In DefCloud, the continuous operation of a JIT compiler, and occasionally a decompiler, incur additional costs due to the support of real-time hardening. DREME is the only approach providing actual numbers where the authors show an increased execution time of around 30% for a webserver instrumented with an I/O redirector. This number fluctuates considerably depending on the size of a request.

MEERKATS' overhead is arguably the highest because of its numerous and complex components. It combines multiple previously discussed approaches and many more, leaving unclear how they should be realized and how high their overhead is.

BioRAC and — at least when not depending on the binary rewriting capabilities of MEERKATS — DREME need multiple program versions beforehand. This requires the development of such, or at least the compilation with different flags, compilers, or linked libraries.

Looking at the overhead of Cloudy is difficult, as there is no natural baseline to which we could compare it. Making the system reliable on highly unreliable infrastructure will inherently come with a need for replication; without it, the system would not be usable.

VII. Discussion

Even though different architectures can provide considerable improvements to the resiliency of a system, their main focus lies on the ability to withstand adverse conditions as good as possible. To achieve anticipation and adaption goals will continue to require manual intervention, and for a successful automated recovery, the system has to be at least partially operational — hence also depending on good withstand capabilities. In this section we identify the limitations of the current architectures and what future work is needed to resolve them.

A. Limitations

Many aspects of resiliency cannot be solved by a new architecture alone. Nevertheless, the current proposals exhibit distinct weaknesses that hinder their widespread adoption, for which improvements are possible.

The fundamental complexity is usually barely addressed, and solutions are offered that are unlikely to be feasible for general application. Such solutions range from writing acceptance tests for every phase in a task in BioRAC over the automatic generation of arbitrary "bait" documents in MEERKATS, a task state-of-the-art artificial intelligence (AI) systems clearly could not handle, to on-the-fly decompilation in DefCloud. Most of the architectures are only partially realizable, as with the research implementation of DREME, where only parts of the original proposal are implemented, and this also only with a lot of manual work that cannot be automated.

Another significant concern is the substantial resource overhead. Our estimates indicate that these solutions could incur an overhead of well above 100%, which would severely impede the adoption of such architectures.

B. Future Work

Given these limitations, future research should focus on implementing novel resilient architectures capable of resolving high complexity. Expanding upon foundational efforts like the DREME research implementation is crucial for developing and evaluating effective solutions. All-encompassing approaches to cloud resiliency, such as MEERKATS, sound promising, and their impact on production environments needs further examination.

VIII. Conclusion

Resiliency is fundamental in a cloud environment and is a shared responsibility between the cloud provider and the customer, depending on the service model. Enhancing resiliency can be achieved on multiple layers, from replicated storage and virtualized computing to virtual machine (VM) checkpointing techniques. Each approach improves resiliency only within a specific aspect of cloud computing. This paper reviewed comprehensive solutions to this issue by exploring novel cloud architectures prioritizing resiliency at their core. All of them are designed to withstand adverse conditions, emphasizing replication and diversity. These concepts can be applied to the applications running on the infrastructure, the technology stack, and the infrastructure itself. Consequently, the complexity and overhead of these systems increase significantly, which may account for the limited number of proof-of-concept implementations. Future research must continue there, developing feasible solutions and investigating their practical impact on resiliency. This work does not include a comparison of cloud architectures with other cyber resilience approaches for cloud environments. Additionally, due to time constraints, it does not address the extent to which the same properties could potentially be more efficiently and effectively achieved at lower layers.

The use and importance of cloud technology in all areas of daily life are expected to continue rising. Therefore, achieving a high level of resiliency against targeted attacks, natural or human-induced disasters, and other adverse conditions on all levels of cloud computing is crucial.

References

[1] "About AWS." [Online]. Available: https://aws.amazon.com/about-aws/

[2] "Angelos D. Keromytis Curriculum Vitae." [Online]. Available: https://angelosk.github.io/cv.html

[3] "Cloud Security – Amazon Web Services (AWS)." [Online]. Available: https://aws.amazon.com/security/

[4] "Google security overview | Documentation." [Online]. Available: https://cloud.google.com/docs/security/overview/whitepaper

[5] "Public cloud computing market size 2024." [Online]. Available: https://www.statista.com/statistics/273818/global-revenue-generated-with-cloud-computing-since-2009/

[6] "Cloud Computing Market to be Worth USD 2,495.2 Billion by 2032, growing at 17.8% CAGR," Oct. 2023. [Online]. Available: https://finance.yahoo.com/news/cloud-computing-market-worth-usd-130600292.html

[7] O. Bakalynskyi and F. Korobeynikov, "Establishing Goals in the Creation of Cyber-Resilient Systems per NIST," in *2023 13th International Conference on Dependable Systems, Services and Technologies (DESSERT)*. Athens, Greece: IEEE, Oct. 2023, pp. 1–4. [Online]. Available: https://ieeexplore.ieee.org/document/10416540/

[8] A. Benameur, N. S. Evans, and M. C. Elder, "Cloud resiliency and security via diversified replica execution and monitoring," in *2013 6th International Symposium on Resilient Control Systems (ISRCS)*, Aug. 2013, pp. 150–155. [Online]. Available: https://ieeexplore.ieee.org/document/6623768

[9] C. Colman-Meixner, C. Develder, M. Tornatore, and B. Mukherjee, "A Survey on Resiliency Techniques in Cloud Computing Infrastructures and Applications," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 3, pp. 2244–2281, 2016, conference Name: IEEE Communications Surveys & Tutorials. [Online]. Available: https://ieeexplore.ieee.org/document/7409914

[10] T. Erl and E. B. Monroy, *Cloud Computing: Concepts, Technology, Security, and Architecture*, 2nd ed. Pearson, Aug. 2023. [Online]. Available: https://learning.oreilly.com/library/view/cloud-computing-concepts/9780138052287/

[11] G. Garlick, "Improving Resilience with Community Cloud Computing," in *2011 Sixth International Conference on Availability, Reliability and Security*, Aug. 2011, pp. 650–655. [Online]. Available: https://ieeexplore.ieee.org/document/6046040

[12] R. Geambasu, D. Mitropoulos, S. Sethumadhavan, J. Yang, A. Stravrou, D. Fleck, M. Elder, and A. Benameur, "Maintaining Enterprise Resiliency via Kaleidoscopic Adaption and Transformation of Software Services (MEERKATS):," Defense Technical Information Center, Fort Belvoir, VA, Tech. Rep., Apr. 2016. [Online]. Available: https://apps.dtic.mil/sti/citations/AD1007307

[13] S. Hariri, M. Eltoweissy, and Y. Al-Nashif, "BioRAC: biologically inspired resilient autonomic cloud," in *Proceedings of the Seventh Annual Workshop on Cyber Security and Information Intelligence Research*, ser. CSIIRW '11. New York, NY, USA: Association for Computing Machinery, Oct. 2011, p. 1. [Online]. Available: https://dl.acm.org/doi/10.1145/2179298.2179389

[14] A. D. Keromytis, R. Geambasu, S. Sethumadhavan, S. J. Stolfo, J. Yang, A. Benameur, M. Dacier, M. Elder, D. Kienzle, and A. Stavrou, "The MEERKATS Cloud Security Architecture," in *2012 32nd International Conference on Distributed Computing Systems Workshops*, Jun. 2012, pp. 446–450, iSSN: 2332-5666. [Online]. Available: https://ieeexplore.ieee.org/document/6258191

[15] M. T. Khorshed, A. B. M. S. Ali, and S. A. Wasimi, "A survey on gaps, threat remediation challenges and some thoughts for proactive attack detection in cloud computing," *Future Generation Computer Systems*, vol. 28, no. 6, pp. 833–851, Jun. 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167739X12000180

[16] K. Kyamakya, K. Jobman, and M. Meincke, "Security and survivability of distributed systems: an overview," in *MILCOM 2000 Proceedings. 21st Century Military Communications. Architectures and Technologies for Information Superiority (Cat. No.00CH37155)*, vol. 1, Oct. 2000, pp. 449–454 vol.1. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/904993

[17] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," Sep. 2011.

[18] A. Montresor, H. Meling, and O. Babaoglu, "Toward Self-Organizing, Self-Repairing and Resilient Distributed Systems," in *Future Directions in Distributed Computing: Research and Position Papers*, ser. Lecture Notes in Computer Science, A. Schiper, A. A. Shvartsman, H. Weatherspoon, and B. Y. Zhao, Eds. Berlin, Heidelberg: Springer, 2003, pp. 119–123. [Online]. Available: https://doi.org/10.1007/3-540-37795-6_22

[19] J. Moura and D. Hutchison, "Resilience Enhancement at Edge Cloud Systems," *IEEE Access*, vol. 10, pp. 45 190–45 206, 2022, conference Name: IEEE Access. [Online]. Available: https://ieeexplore.ieee.org/document/9751692

[20] V. Prokhorenko and M. Ali Babar, "Architectural Resilience in Cloud, Fog and Edge Systems: A Survey," *IEEE Access*, vol. 8, pp. 28 078–28 095, 2020, conference Name: IEEE Access. [Online]. Available: https://ieeexplore.ieee.org/document/8978538

[21] R. Ross, V. Pillitteri, R. Graubart, D. Bodeau, and R. McQuaid, "Developing cyber-resilient systems : a systems security engineering approach," National Institute of Standards and Technology (U.S.), Gaithersburg, MD, Tech. Rep. NIST SP 800-160v2r1, Dec. 2021. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-160v2r1.pdf

[22] A. Sathiaseelan, M. Selimi, C. Molina, A. Lertsinsrubtavee, L. Navarro, F. Freitag, F. Ramos, and R. Baig, "Towards decentralised resilient community clouds," in *Proceedings of the 2nd Workshop on Middleware for Edge Clouds & Cloudlets*, ser. MECC '17. New York, NY, USA: Association for Computing Machinery, Dec. 2017, pp. 1–6. [Online]. Available: https://dl.acm.org/doi/10.1145/3152360.3152363

[23] J. P. G. Sterbenz and P. Kulkarni, "Diverse Infrastructure and Architecture for Datacenter and Cloud Resilience," in *2013 22nd International Conference on Computer Communication and Networks (ICCCN)*, Jul. 2013, pp. 1–7, iSSN: 1095-2055. [Online]. Available: https://ieeexplore.ieee.org/document/6614125

[24] G. Suciu, C. Cernat, G. Todoran, V. Suciu, V. Poenaru, T. Militaru, and S. Halunga, "A solution for implementing resilience in open source Cloud platforms," in *2012 9th International Conference on Communications (COMM)*, Jun. 2012, pp. 335–338. [Online]. Available: https://ieeexplore.ieee.org/document/6262565

[25] L. Svobodova, "Resilient Distributed Computing," *IEEE Transactions on Software Engineering*, vol. SE-10, no. 3, pp. 257–268, May 1984, conference Name: IEEE Transactions on Software Engineering. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/5010234

[26] A. S. Tanenbaum and M. v. Steen, *Distributed systems: principles and paradigms*, 2nd ed., ser. Always learning. Harlow, Essex: Pearson Education, 2014.

[27] T. Welsh and E. Benkhelifa, "On Resilience in Cloud Computing: A Survey of Techniques across the Cloud Domain," *ACM Comput. Surv.*, vol. 53, no. 3, pp. 59:1–59:36, May 2020. [Online]. Available: https://dl.acm.org/doi/10.1145/3388922

SoK: Different ways to increase Resilience of Virtual Machines

1st Carl König
IT Security
Technical University of Munich
Munich, Germany
carl.koenig@tum.de

***Abstract*—Hypervisor technology allows for multiple operating systems to be run securely, safely, and synchronously on the same hardware. This makes it the critical technology to enable cloud computing and, as such, Infrastructure as a Service (IaaS). Consequently, significant parts of modern software infrastructure like healthcare, banking, and even the entire Internet rely upon the security and safety guarantees offered by hypervisors. Due to their central role in everyday life, they are an attractive target for malicious actors who plan on causing extensive social and economic damage. As Virtual Machine Monitors (VMMs) are highly complex pieces of technology, so are the different approaches that try to protect them. In this paper, we categorize and systematize different approaches that try to increase the resilience of Virtual Machines (VMs), which refers to the ability of a VM to maintain its functionality and performance in the face of various challenges or threats. We look at systems that allow VMs to continue running even if the underlying hypervisor crashes, different techniques to fuzz VMMs, and attempts to reduce the potential attack surface of virtualization technology. Then, we will compare the different approaches, discuss their applicability, and argue their contribution towards making VMs more resilient. We managed to identify a new research direction for replication systems. We analyze why creating a system capable of differentiating crashes between internal faults and external exploits is highly relevant. In addition, we argue for the real-world relevance of replication systems and against the real-world relevance of attack surface reduction.**

I. Introduction

Hypervisors and their ability to host multiple VMs on the same hardware allowed for the rise of cloud computing. Thus, VMMs resemble the backbone of modern IT infrastructure. Most everyday services, websites, critical infrastructure, and more depend on working VMM. Consequently, these systems also rely on the security guarantees offered by hypervisors. Here, the most important one is the isolation between the different VMs. Due to their incredibly central role in modern computing, they are a highly interesting target for malicious actors. As VMMs are highly complex pieces of software that are interconnected with underlying hardware, it makes them extremely prone to contain vulnerabilities. Attackers can then leverage these to break these security guarantees. This is why it is highly relevant to discuss methods to not only increase the security of hypervisors but also the different ways to make them more resilient.

This paper aims to explore the different techniques and approaches and derive new research questions regarding increased cyber resilience. We are unique in that we create an overview of all the different technologies enhancing VM security and not only focus on a singular one like replication [14], [17] or discuss one security concept and mention its application to VMs like the fuzzing roadmap [30]. This paper serves as a synopsis of what has been achieved in hypervisor cyber resilience so far and inspects where further research is possible.

We systematize different approaches that aim to increase the resilience of VMMs and discuss their effectiveness. In particular, we take a look at:

- **Replication** of VMs, which involves the continuous synchronization of the inner state of the running VM with a backup VM so in case the hypervisor crashes, the VM can resume execution on a different instance [7]. We then discuss how this availability-enhancing feature can be improved and identify open research questions that would further the development of such a system to make it withstand outside attacks. In addition to that, we also discuss VM migration, which serves more as an administrative tool to move VMs between different hardware or VMMs.
- **Fuzzing** of Hypervisors or, more precisely, their virtual device drivers. Here we look at the different techniques to search for vulnerabilities and analyse if it is useful to develop new strategies to systematically search for bugs. In addition, we consider the challenges of fuzzing VMMs.
- Last but not least, we discuss **attack surface reduction** approaches to the common goal of achieving a reduced trusted code base. Similarly, we look at how VMMs use hardware features to protect and question the applicability of attack surface reduction techniques.

II. Background

This section discusses the technical knowledge needed to understand this paper. First, we will examine the difference between hypervisors/VMMs and VMs, the two types of hypervisors, and what security promises they deliver. Subsequently, we will discuss how hypervisors use hardware like network cards or memory. After that, the focus shifts to CPU features and instructions like protection rings, which are necessary for understanding this paper. Finally, we will define how the term cyber resilience is used in this paper.

doi: 10.14459/2024md1759337_03

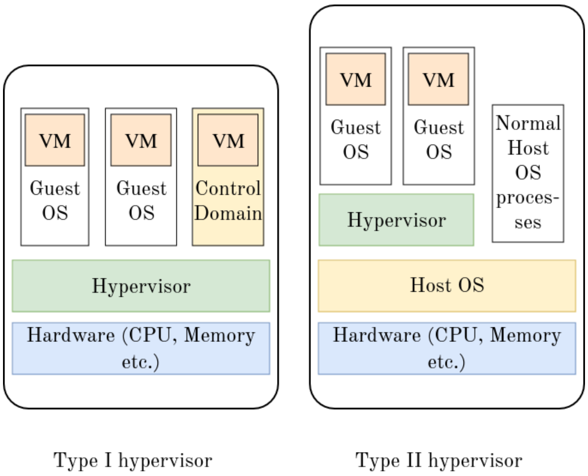

Fig. 1. The two different Types of hypervisors visualized.

A. Classification and guarantees of hypervisors

The idea of virtualization is to allow for multiple VMs to run securely on the same hardware. The hypervisor or VMM is responsible for isolating the different VMs and creating the illusion that they have the entire hardware for themself. The VM will often run an OS, which is then called a guest OS. There are three different types of VMs: System Virtual Machines, Process Virtual Machines, and Containers. The focus of this paper lies solely on System VMs. Two types of hypervisors run System VMs:

- **Type 1 hypervisor**, also often called **bare-metal hypervisor**, runs, as the name suggests, directly on the hardware. This means that the hypervisor has to perform tasks usually done by an OS, such as managing page tables or interrupt handling. Often, a *control domain* is used, which is an OS running in a VM. The VMM then communicates with the *control domain* and lets it perform most tasks OS has to handle, like memory management.
- **Type 2 hypervisors** instead run inside a so-called host operating system, e.g. Linux or Windows. Consequently, the hypervisor does not need to worry about typical OS tasks as the host OS already provides a layer of abstraction between hardware and hypervisor.

fig. 1 visualizes the how basic layout of both types of hypervisors. An example of a commonly used Type 1 hypervisor is Xen, and a Type 2 hypervisor is QEMU.

The main security guarantee of nearly all hypervisors, and the only one that this paper focuses on, is isolation between the different VMs. Thus, if one VM executes malicious code, crashes, or does something else, the other VMs are neither compromised nor affected [18], [28].

B. Isolation from a software side

This isolation is, among other things, provided by virtual memory as well as virtual device drivers. This isolation is often implemented by using *multilevel page tables* and potentially *shadow page tables*, usually with the help of hardware support offered by the CPU. This leads to each VM thinking they have their physical memory even though they share the same physical RAM-Stick with many other VMs. Usually, literature distinguishes between the OS in the VM knowing (paravirtualization) and not knowing about being virtualized. However, this difference is not important for the contents of this paper and thus will not be explained in further detail. In addition to virtualizing memory, devices like network cards and USB ports must also be virtualized to keep the isolation of VMs in tact. Usually, the OS will initialize those devices on boot. If a device is detected, the OS configures and prepares it for later usage. However, when running on a VMM, these devices which are actual hardware. Thus, the VM should not be allowed direct access as that would break the isolation. Consequently, the hypervisor has to virtualize these resources and provide virtual device drivers who communicate between software and hardware. An example of that would be the hard drive. Multiple VMs believe they have sole ownership of the same drive, meaning the hard disk must also be virtualized. To achieve this, the hypervisor uses a virtual device driver, which interfaces with the VM and pretends to be a hard drive, when in reality, it just creates another layer of abstraction between hardware and VM [28].

C. Isolation from a hardware side

While hypervisors provide isolation via software, this is only possible because that software utilizes hardware features supplied by the CPU Some hardware features are crucial for virtualization and allow for it to be possible in the first place, while others just provide enhanced performance on virtualization. The first concept, CPU protection or privilege rings, predates hypervisors and is of utmost importance for an OS. There are four different privilege levels ranging from zero (the highest) to three (the lowest) on x86_64. The lower levels are used to limit access to the CPU and other resources like memory. Usually, the kernel, as well as the device drivers, run in ring 0 while uesrspace programs run in ring 3. Thus, they do not have access to privileged instructions with which they could, e.g., modify the page table. This level of hardware-aided protection is crucial for the security of operating systems, and while their use case for hypervisors is limited, they are important to note here because all common Type 1 hypervisors run the control domain and virtual device drivers directly alongside the VMM in ring 0 [13].

However, for hypervisors, the ability of the CPU to switch between *root-mode* and *non-root-mode* is far more crucial. These modes allow the CPU to differentiate if a VM or the VMM is running. This means that if inside of a VM (so in *non-root-mode*) a privileged instruction is run that requires hypervisor intervention, the CPU traps into the VMM. This way, the VMM can ensure that the VM cannot do something that would break the isolation or be malicious towards the hypervisor. The switches between *root-mode* and *non-root-mode* are referred to as *VM Entrys* and *VM Exits*. The CPU provides something called virtual-machine control data structures (VMCSs) that are used to manage the transitions between *root-mode* and *non-root-mode*. The hypervisor can perform instructions like reading or writing from VMCSs via privileged instructions, e.g., *VMREAD* or *VMWRITE*. It is important to note that the terminology above is the one Intel uses and can differ for other x86_64 architectures. However, the functionality remains unchanged regardless of the name [13].

D. Cyber Resilience

This paper uses the definition of cyber resilience given by the NIST. This definition splits cyber resilience into the ability to **anticipate**, **withstand**, **recover**, and **adapt** [21]. In this paper, anticipation describes not only the direct forecast of an attack but also the design philosophy of a system in which something is created with inherent protections against malicious actors. The ability to withstand describes ways of keeping a running system even in the face of internal faults and external attacks. Recovery means that a system can rebound into a working state after it has stopped working. Last but not least, adaption describes the process of quickly reshaping and changing a system in the face of a threat.

III. Methodology of the research

This section outlines our approach to finding and systematizing the knowledge related to cyber resilience in VMs. We first looked at the NIST definition of cyber resilience to understand which direction we should orient our research. After identifying and reading through papers regarding VM security, we sorted them into the four categories of cyber resilience (according to NIST). The next step consisted of pinpointing the most relevant and promising research directions. This was done by, among others, checking if real-world applications also use the systems presented in the papers. For obvious reasons, we only looked at open-source projects like Xen. After identifying an exciting research direction, we used the papers. We further analyzed the papers we had already found, research work mentioned in their related work section, and papers citing the initial papers.

With the knowledge of which research directions we want to write about, we asked the following two key questions:

- 1) Which technical aspects need to be mentioned so a reader can understand the research direction?
- 2) Which papers are the most relevant and are responsible for the most significant accomplishments?

After answering those questions, we had a good enough overview to write this paper. If it became evident during the writing process that some systems needed to be explained better or some knowledge gaps were left open, we just reiterated these steps.

IV. VM migration and replication

In this section, we delve into the concept of virtual machine migration and replication and the difference between the two. Furthermore, we will explore the types of attacks and faults that VM replication can protect against. These include hardware failures, software bugs, and even localized outages. However, it is important to understand that VM replication, while a powerful tool, is not a panacea for all possible attacks and faults. In specific scenarios, that we will discuss later, live replication and migration cannot provide adequate defense.

A. Migration and replication

Migration describes the process of moving something, like a Linux process or operating system, from one platform or environment to another.

Clark et al. were the first to implement live migration at the magnitude of operating systems. They showed that with the help of VMs, it was possible to migrate running operating systems between different physical hosts with a reasonably low downtime. [6]. While their system allows fast and seamless migration between different physical hosts, their tool aims to purposefully move operating systems. Thus, it cannot make the VM more resilient towards internal or external faults.

Replication describes the act of copying or duplicating. Hereinafter, we will differentiate between migration and replication of VMs as follows: Migration describes the process of intentionally moving a VM from one physical host to another. In contrast, replication describes the process of continually sending the internal state of one VM to another and thus allowing the contents of the hypervisor to run on a different host in case of an internal or external fault.

The *Remus* system, created by Cully et al., does precisely that. Their tool runs two hosts at the same time. The first host (active or primary host) runs speculatively and regularly sends snapshots of its entire internal state to the second host (backup host). Running two hypervisors in that way allows the second host to resume execution if a fault inside the first host occurs. The backup host can replicate the inner state due to the snapshots sent by the primary host [7]. Their paper created the basis upon which most of the recent VM asynchronous replication research builds upon.

B. Functionality of live replication

Since nearly all modern live replication systems work similar to *Remus*, we will use it and its design goals to explain the technical details of live replication systems. Fig. 2 also shows the process of one relicatoin step on the left and the handeling of a fault on the right. In their paper Cully et al. identified three core objectives:

- **Generality:** Fault tolerance should be provided as a low-level service that does not require retrospective tailoring towards the applications running inside the replicated hypervisor and the hardware it is running on.
- **Transparency:** The operating systems and applications running inside the VMM should not require any adjustments to allow for state recovery and replication. In addition, the system should not need any additional hardware besides what the (average) hypervisor is already running on.
- **Seamless fault recovery:** The fault recovery should happen without any external state being either lost or re-

peated. The occurrence of a fault should not be noticeable from an outside perspective.

Remus replicates a VM by taking a snapshot of the entire inner state of an active host and sending it to a backup host. For hypervisors, this usually entails the memory (or at least all the dirty pages since the last backup), the state of external device drivers, CPU registers, and so on. Considering that all those aforementioned resources are virtualized accessible for the hypervisor, the goal of **generality** is achieved by replicating the entire VM. By using the abstraction provided by the hypervisor, the system is entirely independent from the hardware of the active or backup host. In addition, replicating the entire state of the VM makes it also independent of the software running inside the Monitor [7]. Furthermore, the abstraction of the hypervisor allows for the replication and, in case of a fault, the switch to the backup VM, to be completely transparent and unnoticeable for any software inside the hypervisor.

Important to note is that *Remus* performs the execution of the active host speculatively. The replication works because all the output (e.g., everything in the virtual network controller) is buffered up until the inner-state snapshot has been taken. After that, the buffer is flushed. Because all modern CPUs have multiple cores and there can be non-deterministic events in the operating system and application in the hypervisor, running the same snapshot multiple times or running the same snapshot on the active and backup VM is likely to yield different results. However, even though the output of the VM is not deterministic, this type of execution has the advantage of being much faster than synchronizing the active and backup host step by step. This also means that if an internal fault occurs and the backup host has to take over, only the computations after the last snapshot are lost. Because the output of the faulting hypervisor was not released, it does not matter that the results of the backup host might differ from the results of the active host because it is not noticeable from an outside perspective. With that, the goal of **seamless fault recovery** is also achieved [7].

In addition, *Remus* performs all its replication asynchronously. This means that the active host takes the snapshot and sends it to the backup host, not waiting for confirmation if the backup host received the snapshot. Implementing the replication in such a way leads to a substantial increase in performance, as the active host only has to pause execution for a concise amount of time [7].

It is worth noting that having one backup VM for each actively running one is unnecessary. Internal faults in hypervisors are rare, so an N-to-1 configuration with one backup VM for multiple active hosts is usually preferable [7].

The live replication for VMs, as described in the *Remus* paper, increases the resilience of hypervisors. In case of an internal fault, the live replication allows for a fast and effective recovery of the VM back into a working state, meaning that systems with similar functionality to *Remus* can make hypervisors more resilient in the face of internal bugs, hardware failures, or localized outages. In systems where constant availability is imperative, the prospect of live replications is highly relevant. Suppose the hypervisor backups are performed at the scale of data centers. In that case, they can guarantee that a system will stay online even in the face of regional power outages or natural disasters. While discussing the technical details for such a system is outside this paper's scope, there are papers tackling these issues and implementing solutions [4], [11].

C. Protection against zero-day-hypervisor-DoS exploits

In this section, we explain how live replication upholds availability in the face of internal faults, hardware failure, etc. However, live replication systems have an Achilles heel. They cannot protect the hypervisor against zero-day Denial of Service (DoS) exploits. Suppose a malicious actor leverages such an exploit and manages to crash the primary host. In that case, even though the backup will continue the execution, the virtualization will still be done by the same VMM, meaning the attacker can use the same vulnerability to crash the backup. Protecting against such zero-day DoS attacks is especially relevant as most hypervisor vulnerabilities affect the availability. An analysis of the different impacts on exploits done by Decourcelle et al. shows that for KVM, QEMU and XEN over 90% of found attacks affect the availability of the system and nearly 80% for both ESXi and Hyper-V [9].

Ngoc et al. propose using hypervisor transplantation to combat this vulnerability towards DoS attacks. They show that their system (HyperTP) can quickly and seamlessly migrate software inside a Xen VM to a KVM (Kernel-based Virtual Machine - a Linux kernel module aiding virtualization) and vice versa. Because their approach is not replication-based, it only protects against disclosed but unpatched zero-day exploits. Nevertheless, in all cases where the exploit becomes publicly known, HyperTP can protect the infrastructure until the patch to fix the vulnerability is released. Additionally, they show that there is only minimal overlap between vulnerabilities in different hypervisors, meaning that, e.g., if Xen is exploitable, KVM will most likely be unaffected [20].

Decourcelle et al. propose HERE, a system similar to HyperTP, also leverages the benefits of software diversity; however, in this case, to implement live replication. HERE uses Xen for the active VM and KVM as the backup VM. Furthermore, Decourcelle et al. propose two further enhancements for live replication:

- **Dynamic control of the checkpointing period.** Live replication introduces overhead, which reduces the execution speed of applications inside the VMM. Thus, live replication systems are often only relevant for software with a more significant need for availability than speed. The constant replication intervals used in most systems are ill-suited for applications that can only tolerate minimal downtime. As a countermeasure, HERE dynamically controls the intervals of inner-state snapshots based on the maximum tolerable downtime and desired replication overhead.

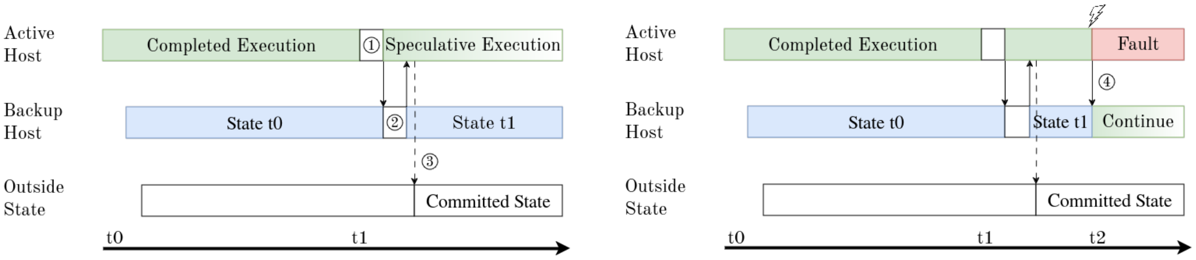

Fig. 2. One pass through of live replication on each side respectively. The left figure illustrates the functionality of the system with no fault occurring in the primary host and the right figure displays the timeline in case of a fault after a snapshot has been send to the backup host. The design of this figure is inspired by the one of the original *Remus* paper [7].
1) Checkpoint where the primary host pauses execution and gathers all changes.
2) Changes are transferred to the backup host.
3) Any outgoing IO is released.
4) The fault occurs, the backup host gets notified, and subsequently starts running with his latest knowledge of the current state.

- **Optimized multithreaded replication.** Modern network card throughput speed is in the hundreds of gigabits, and using only a single thread to send dirty pages to the backup host does not utilize the hardware's full potential. It risks turning the serialization time of the backup data into an unnecessary bottleneck.

With the aforementioned optimizations, HERE outperforms *Remus* in most benchmarks while itself having negligible overhead [9].

V. Fuzzing of hypervisors

So far, the focus has been placed on ways to keep a system available, allowing it to withstand power outages, natural disasters, and attacks using zero-day DoS attacks. While this increase in fault tolerance is highly relevant for cyber resilience, so is the ability to anticipate attacks and thwart them by patching the vulnerability. Over the last years **fuzzing** has been shown to be highly effective when it comes to finding bugs in complex pieces of software [1], [10], [23], [30]. This section will discuss applying fuzzing to hypervisors and the techniques for finding possible attack vectors.

A. Fuzzing and hypervisor Fuzzing

Virtual device drivers used by VMMs are highly complex. They combine hardware and software, often with multiple developers creating them and working on different levels of abstraction. In addition, many device drivers run in the most privileged CPU ring. Because of all these reasons, virtual device drivers are especially prone to errors and are the cause of most hypervisor exploits in the past few years [22]. Hence, there is an inherent need to test them thoroughly and efficiently. The golden standard for such kinds of testing in userspace programs are evolutionary graybox fuzzers like American Fuzzy Lop (AFL) [1], or its successor AFL++ [10]. They are code coverage guided, meaning they receive information about which code branches are executed. Graybox, in this case, means that while the fuzzer does not know about every execution step, the coverage guidance gives it a rough sketch of what was executed and what was not. This information and other metrics about the tested program influence the generation of new test cases. This information and other metrics about the program being tested are then used to adapt and evolve the test cases. Consequently, this allows unexplored code branches to be executed and fuzzed by the next set of test cases [10].

Henderson et al. apply evolutionary fuzzing to Hypervisosrs with their system VDF (Virtual Device Fuzzer). VDF uses AFL as a basis, meaning that it is also evolutionary. In this case, the fuzzing is also low dimensional, meaning that it tests only one virtual device via memory interactions and I/O. However, implementing an AFL-based fuzzer for hypervisors presents unique challenges. Henderson et al. laid out three core issues. Firstly, the testing must only target one specific virtual device, and thus, fuzzing needs to be focused on a tiny part of the hypervisor codebase. Secondly, the virtual device driver constantly interacts with the hypervisor, so VDF needs to be directly embedded into the VMM. Thirdly, the testing framework must be stateful as device drivers must be adequately initialized before being tested [12]. The way Henderson et al. solved these problems is outside the scope of this paper. However, with their system VDF, they managed to find 1,014 crashes or hangs in virtual device drivers.

B. Different types of hypervisor fuzzers

While VDF is one of the major fuzzers for hypervisors, it has one major flaw. As mentioned earlier, it performs low-dimensional fuzzing at a slow rate. However, the reality of VMMs is that different device drivers constantly work together and influence each other's results. For that reason, a high-dimensional fuzzer that tests many interfaces simultaneously would be able to check a broader range of code and potentially discover new bugs.

Schumio et al. [25] created a high-dimensional fuzzer called HYPER-CUBE. It has a high test case throughput, interacts with all interfaces simultaneously, and provides stable and deterministic tests for different VMMs. The novel system runs inside a custom operating system and a custom Bytecode interpreter.

By wrapping the fuzzer inside the custom operating system, Schumio et al. can run it on all hypervisors (open source and proprietary VMMs) that can boot commercial of-the-self operating systems and fuzz them. Because of their novel approach, they found 54 bugs and obtained 43 CVEs. Due to HYPER-CUBE's high throughput, the system was magnitudes faster at finding bugs than VDF. Their paper shows that rediscovering bugs originally found by VDF took their system only up to 10 minutes, while VDF took over 60 days. However, these differences in speed are only the case for device drivers with a smaller codebase. HYPER-CUBE achieves its high test throughput speed only because it uses no coverage guidance. This difference between fuzzing simple and complex virtual device drivers makes sense because random tests will automatically cover most possible branches if the codebase is small and straightforward. Consequently, when fuzzing complex device drivers, HYPER-CUBE struggles to test them thoroughly and detect bugs [25].

To counter this issue, Schumilo et al. [26] created another system called Nyx, which uses the same costume Operating System and bytecode interpreter as HYPER-CUBE but instead utilizes gray box fuzzing, similar to AFL, to search for bugs. Nyx utilizes KVM-PT (Processor Trace) and QEMU-PT, which use the Intel-PT hardware features to obtain the needed code coverage. Intel-PT collects execution information like control flow or privilege level of the code running on the CPU and provides this data to a process. However, due to the complexity of some device drivers, coverage guidance is often insufficient to effectively fuzz, and a specification is needed. These specifications give the fuzzer additional insides into the test code but have to be created by hand. Thereupon, a human needs to understand the inner workings of the test subject and create a specification describing them. This strongly impedes the scalability of the system. For example, understanding the structures performing VirtIO and creating a specification for them took Schumilo et al. two days. In addition, the setup of Nyx is more complex due to the hypervisor needing to run inside of KVM-PT. However, this has the upside of allowing for faster recovery in case of a crash because KVM-PT needs to reload an old snapshot of the VMM, and the fuzzer can continue testing. Nevertheless, Schumilo et al. used Nyx to uncover 44 Bugs, and they requested 22 CVEs [26].

Pan et al. [22] propose a different and scalable approach to fuzzing virtual device drivers that do not rely on the creation specifications. They focus on direct memory access, which most drivers use to transfer data structures. Part of their research entailed studying and analyzing these types of memory accesses and their code. Pan et al. used their gathered knowledge of direct memory access to create V-Shuttle, a fully automatic semantics-aware fuzzer. By organizing the different types of direct memory access into categories and fuzzing based on that information, V-Shuttle gains a deep understanding of the underlying protocol, allowing for the better and purposeful creation of test cases with potentially interesting results. Due to not needing any specifications, V-Shuttle scales a lot better than Nyx while retaining the advantage of having an understanding of the inner workings of the virtual device driver. In addition, Pan et al. show that their system has higher code coverage than Nyx and HYPER-CUBE. With their new approach, they discovered 35 new vulnerabilities, 17 of them being accepted as CVEs at the time of the paper release [22].

VI. Attack surface reduction

Another way of increasing the security of VMMs is by reducing their potential attack surface. This reduction is primarily achieved by implementing the principle of least privilege, which effectively decreases the trusted code base (TCB), mirroring the approach most microkernels take. The most common method of deprivileging hypervisor code involves the relocation of code that doesn't necessarily require privileged instructions from protection ring 0 (or in modern CPUs *root-mode*) to a lower one (or *non-root-mode*).

Murray et al. [18] highlight create an application for reducing the TCB in the Xen hypervisor. In their paper, they move the userspace of Dom0 (domain zero) from ring 0 into ring 1, where it runs alongside all the guest VMs. Dom0 is the control domain of Xen, which acts similar to how the host OS would on a type-2 hypervisor, providing the guest VMs with access to the hardware. Because all of Dom0 is running in the highest protection ring, Xen needs to trust everyone accessing Dom0, even if it is just userspace access. However, the totally missing protection between VMM and dom0 is not ideal which is why Murray et al. propose a disaggregated domain builder (domB) which runs in ring 0 and communicates with Dom0. These changes manage to reduce the TCB by an order of magnitude. These proposed changes, however, have not been integrated into Xen. It is important to note that Xen has a Dom0less mode where dom0 is not run, meaning that all the VMM resources are statically partitioned.

DeHype has a similar approach but is applied to KVM, the virtualization module of the Linux kernel. It manages to deprivilege 93.2% of the existing code while adding only a marginal amount. This is achieved by decoupling the dependencies of the KVM. This is done by moving the memory access API, scheduling-related operations, and other functions into the userspace. However, privileged instructions like VMREAD and VMWRITE, among others, are essential for virtualization; thus, some functionality has to stay inside the kernel. [29] However, because the memory API is now running in ring 3, it does not know the physical addresses for allocated memory. Consequently, a form of memory rebasing needs to be implemented. In this case by letting the privileged module allocate pages in the kernel space and then map them into the userspace. If the kernel then announces the physical addresses of the pages to the userspace program, it can efficiently translate virtual addresses to physical addresses.

It is important to note that while Murray et al.'s proposed changes reduce Xen's TCB still remains quite sizable. Steinberg et al. take the minimization of TCB for a hypervisor with NOVA (a microhypervisor) to an extreme. Nova provides a very thin virtualization layer with an extremely small TCB by utilizing hardware features such as nested paging and I/O virtualization, and keeping the VMM functionality as minimal as possible. These design choices allow NOVA to only consist of 36 thousand LOC [27].

Due to the similarities between hypervisors with a small TCB and microkernels they can also be combined usefully. For this, the seL4 microkernel is particularly interesting. This is because of its security promises and its frequent use in critical systems due to its implementation being formally proven [2]. For that reason de Matos et al. propose an approach to leverage the security of seL4 in using it as a type 1 hypervisor and running QEMU as well as VirtIO (paravirtualized drivers from the Linux kernel) on top of the kernel [8].

VII. Evaluation

In this section, we take a direct look at which aspects of cyber resilience, as defined by NIST [21], the different presented techniques enhance. In addition, we discuss differences between the approaches to increase resilience and identify exciting and new research directions.

The relation between the techniques presented and the abilities of cyber resilience is also visualized in table I. As presented by Clark et al. [6], live migration only marginally increases hypervisors' resilience. This is because it only allows moving VMs between different physical hosts. While this is extremely useful for updating VMMs seamlessly and unnoticed to both the outside world and the program running inside the VM, it only affects the ability to keep systems available. However, security updates are arguably a form of adaption, meaning the system increases cyber resilience in this specific use case.

Nevertheless, live migration systems like HyperTP [20] are designed to make VMs more adaptable. The main objective of this tool is to allow a seamless transition of VMs from one type of hypervisor to another. At the very core of that idea is the improvement of VM adaptability in case of a potential threat. It also improves in the category of anticipation by providing support infrastructure in case an attack on a specific type of hypervisor is anticipated. Nonetheless, systems of this type have a significant disadvantage. They can only offer their protection if it is known that the hypervisor in use is potentially compromised.

While from a technical aspect, there are lots of similarities between live migration and live replication, they are divergent when looking at them from a cyber resilience perspective. For replication, the main focus lies on recovery and, thus, protection from internal faults. Consequently, these systems also help VMs to withstand stress and adverse conditions. While in the case of Remus [7], the system is not designed to protect against a malicious outside actor, it still improves cyber resilience in these two aspects. However, the relatively recently developed system HERE [9] focuses on making VMs even more robust. If a malicious actor attacks the VMM, the system nullifies the security vulnerability exploited by switching hypervisors. This enormously increases the ability to withstand any exploits launched at the virtualisation software.

Nevertheless, the question arises of what would happen if an attacker found a vulnerability in the primary hypervisor and the backup hypervisor. It could be highly relevant for live replication systems to find out if malicious actors exploiting one vulnerability also have more in stock. As discussed with fuzzers, they all discovered bugs on multiple VMMs, and thus, hypervisor transplant can potentially only lull the users into a false sense of security. Another Problem with systems like HyperTP is that due to the switch in hypervisors, there is a potential loss of performance (especially if the VM is moved from a Type-1 to a Type-2 hypervisor). A system that can differentiate between an internal fault and an external attack could help utilize systems similar to HyperTP most efficiently, as the recovery and ability to withstand faults and attacks can be paired with efficiency.

While fuzzing systems do not directly increase the resilience of VMs, their mention and explanation in this paper are well deserved as they have proven over time to be the most efficient way to discover significant security vulnerabilities in hypervisors. They play a role in adaption as they often provide the need to update a VMM. In addition, they are used as a form of anticipation. It is known that the virtualization code contains bugs and vulnerabilities, but the idea is to find them before malicious actors do and thus, in a way, anticipate the attack vector and fix the vulnerability before it can be exploited.

When talking about fuzzers, it is important to add that with each new fuzzer, new bugs are discovered. This can best be seen in table II, where we show the amounts of vulnerabilities and CVEs discovered by the different fuzzing systems. Even though they are all based on the same idea, all of them managed to detect quite a substantial amount of bugs only because they searched for them in different ways. This further emphasizes how complex hypervisors and their virtual device drivers are and also shows that fuzzing them with a slightly different technique can lead to publishable results.

TABLE II
Comparison of the number of vulnerabilities found and CVEs assigned by the different fuzzers*

Fuzzer	Amount Vulns.	Amount CVEs
VDF**	unknown	unknown
HYPER-CUBE	54	43
Nyx	44	22
V-Shuttle	35	17
Total	133	82

∗ All numbers are taken from the corresponding papers. Vulnerabilities that the fuzzers found by third parties are not listed either. The amount of Vulnerabilities only entails those that have not been known before the fuzzer discovered them.

** In the VDF paper, there is no appendix where security critical bugs are listed.

Like fuzzers, attack surface reduction does not actively anticipate potential attacks, stress, or adverse conditions [21]. However, admitting that attacks will happen and preparing against them from the ground up is effectively anticipating exploitation. Using the seL4 microkernel as a base for the

TABLE I
Different systems and approaches and the areas in which they make hypervisors more resilient.

Category	System	Anticipation	Ability to Withstand	Recovery	Adaption
Live migration					
	Live migration*	✗	✗	✗	(✓)
	HyperTP	(✓)	✗	✗	✓
Live Replication					
	Remus	✗	✓	✓	✗
	HERE	✗	✓✓	✓	✗
Fuzzer		(✓)	✗	✗	(✓)
Attack surface reduction		✓	✗	✗	✗

✗ : Does not increase resilience in this aspect
(✓) : Only marginally increases resilience/increases resilience only in special use cases
✓ : Increases resilience in that aspect
✓✓ : Strongly increases resilience in that aspect

* by Clark et al.. Their system is nameless.

hypervisor or moving the control domain out of ring 0, all these approaches try to make exploiting the VMM as hard as possible from the beginning. Nevertheless, the practicality of such systems is a valid concern. While reducing the attack surface sounds very good on paper, the system's efficiency is often ignored, and the practicality not discussed. It could be interesting to analyze the real-world usage of such attack-surfaced reduced systems.

VIII. Discussion

In this paper, we do not discuss many aspects of hardware security. Not only did we not mention hardware virtualization support in much detail, but we also just assumed that it could not contain any bugs [13]. This is, of course, not true, and there have been many side-channel-based attacks like Meltdown and Spectre and other attacks like RowHammer [15], [16], [19] that use hardware bugs to jeopardize VM security. Thus, the hardware aspects of hypervisor cyber resilience could also be of great interest. Additionally, many attack surface reduction techniques are obsolete due to hardware features like *root-mode* implementing better security between hypervisor and VMM.

We also did not look at any other definitions of cyber resilience. It could be of interest to compare the differences between them and check if the discussed technologies also increase resilience according to the other interpretations. In the same breath there could have been a discussion on which definition is best suited for hypervisors.

Something that was also not discussed is the possibility of system misconfiguration that could lead to vulnerabilities unrelated to hardware or software errors.

IX. Related Work

The amount of work focussing on systematizing knowledge or surveys related to cyber resilience for VMs is slim. However, the work related to our paper is the survey from Cinque et al. [5]. Their paper focuses on the current industry trends in virtualizing *mixed-critical systems*. Clinique et al. discuss the security and isolation of VMs as well as the usage of microkernels and TCB reduction but do not mention live replication or any other kind of ability of a VM to withstand attacks. In addition, they only mention fuzzing once while listing different ways to test hypervisors, not going into depth on how they work and what they fuzz [5].

Contrary to that, Zhu et al., with their survey paper on fuzzing [30], details VMM testing. Just as in this paper, they also talk about the fuzzing frameworks HYPER-CUBE and Nyx; however, they are in less detail than we do. Unlike us, Zhu et al. also discussed many other approaches to fuzzing hypervisors but not in great detail as the context they mention them in is only examples of applications of techniques explained above [30].

Jeba et al. [14] analyze live VM migration, replication, load balancing, and energy management migration techniques. They talk about ways to increase the resilience of virtual machines but do so on a more abstract level than this paper [14]. Live replication and the history of systems intending to increase the availability of VMs are discussed at great length in this survey paper by Medina et al. [17]. Similar to this paper, they also confer on *Remus* and its functionality.

As VMs run on the same hardware, the security of, e.g., the CPU, is essential to guarantee isolation and security. In their Survey paper, Anwar et al. [3] studied and classified different side-channel attacks that exploited bugs related to the L2 or L3 cache and their countermeasures. In addition, they compared the attacks and fixes against those of typical side-channel attacks. Their work is quite different from ours as we did not go into hardware security in this paper at all [3]. Another paper that focuses on side-channel attacks is from

Riddle et al. [24]. Their survey focuses on common side-channel attacks like timed attacks and defenses against them.

X. Conclusion

This paper looks at the most important technologies that increase the resilience of hypervisors. It classifies the discussed ideas into the four cyber resilience categories defined by NIST [21]. In addition, the presented systems are critically analyzed.

We question the real-world applicability of attack surface reduction approaches as the core idea of the deprivileging software tends to add much overhead. Furthermore, we identify potentially exciting areas for further research. While live replication that changes the hypervisor if an internal fault occurs does protect against DoS attacks targeted at the specific VMM, it adds potential overhead. This is especially true if a switch from Type-1 to Type-2 hypervisor occurs, as presented in HyperTP [20]. Here, a system that could detect the difference between an internal fault and an external attack could be of great use. It would not only save computing power but could also act as a early warning system. Such technology can potentially be the most promising research regarding the withstanding attacks part of cyber resilience.

References

[1] American fuzzy loop. Accessed: 27-6-2024.

[2] Virtualisation on seL4 | seL4 docs. Accessed: 20-6-2024.

[3] Shahid Anwar, Zakira Inayat, Mohamad Fadli Zolkipli, Jasni Mohamad Zain, Abdullah Gani, Nor Badrul Anuar, Muhammad Khurram Khan, and Victor Chang. Cross-vm cache-based side channel attacks and proposed prevention mechanisms: A survey, 2017. *Journal of Network and Computer Applications*, 93:259–279, 2017.

[4] Omran Ayoub, Obed Huamani, Francesco Musumeci, and Massimo Tornatore. Efficient online virtual machines migration for alert-based disaster resilience, 2019. In *2019 15th International Conference on the Design of Reliable Communication Networks (DRCN)*, pages 146–153.

[5] Marcello Cinque, Domenico Cotroneo, Luigi De Simone, and Stefano Rosiello. Virtualizing mixed-criticality systems: A survey on industrial trends and issues, 2022. 129:315–330.

[6] Christopher Clark, Keir Fraser, Steven Hand, Jacob Gorm Hansen, Eric Jul, Christian Limpach, Ian Pratt, and Andrew Warfield. Live migration of virtual machines, 2005. In *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation - Volume 2*, NSDI'05, pages 273–286. USENIX Association.

[7] Brendan Cully, Geoffrey Lefebvre, Dutch T. Meyer, M. Feeley, N. Hutchinson, and A. Warfield. Remus: High availability via asynchronous virtual machine replication. (best paper), 2024.

[8] Everton de Matos, Conor Lennon, Eduardo K. Viegas, Markku Ahvenjärvi, Hannu Lyytinen, Ivan Kuznetsov, Joonas Onatsu, and Anh Huy Bui. Integrating VirtIO and QEMU on seL4 for enhanced devices virtualization support, 2023. In *2023 IEEE 22nd International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 1076–1085. ISSN: 2324-9013.

[9] Jean-Baptiste Decourcelle, Tu Dinh Ngoc, Boris Teabe, and Daniel Hagimont. Fast VM replication on heterogeneous hypervisors for robust fault tolerance, 2023. pages 15–28. Conference Name: Middleware '23: 24th International Middleware Conference ISBN: 9798400701771 Publisher: ACM.

[10] Andrea Fioraldi, Dominik Christian Maier, Heiko Eißfeldt, and Marc Heuse. Afl++ : Combining incremental steps of fuzzing research, 2020. In *WOOT @ USENIX Security Symposium*, 2020.

[11] Andreas Fischer, Ali Fessi, Georg Carle, and Hermann de Meer. Wide-area virtual machine migration as resilience mechanism, 2011. In *2011 IEEE 30th Symposium on Reliable Distributed Systems Workshops*, pages 72–77.

[12] Andrew Henderson, Heng Yin, Guang Jin, Hao Han, and Hongmei Deng. VDF: Targeted evolutionary fuzz testing of virtual devices, 2017. In Marc Dacier, Michael Bailey, Michalis Polychronakis, and Manos Antonakakis, editors, *Research in Attacks, Intrusions, and Defenses*, pages 3–25. Springer International Publishing.

[13] Intel. Intel® 64 and ia-32 architectures software developer's manual. https://www.intel.com/content/www/us/en/developer/articles/technical/intel-sdm.html.

[14] P. Getzi Jeba Leelipushpam and J. Sharmila. Live VM migration techniques in cloud environment — a survey, 2013. In *2013 IEEE Conference on Information & Communication Technologies*, pages 408–413.

[15] Paul C. Kocher, Daniel Genkin, Daniel Gruss, Werner Haas, Michael Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. Spectre attacks: Exploiting speculative execution, 2018. *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1–19, 2018.

[16] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul C. Kocher, Daniel Genkin, Yuval Yarom, and Michael Hamburg. Meltdown: Reading kernel memory from user space, 2018. In *USENIX Security Symposium*, 2018.

[17] Violeta Medina and Juan Manuel García. A survey of migration mechanisms of virtual machines, 2014. *ACM Comput. Surv.*, 46(3), jan 2014.

[18] Derek Gordon Murray, Grzegorz Milos, and Steven Hand. Improving xen security through disaggregation, 2008. In *Proceedings of the fourth ACM SIGPLAN/SIGOPS international conference on Virtual execution environments*, VEE '08, pages 151–160. Association for Computing Machinery.

[19] Onur Mutlu. Retrospective: Flipping bits in memory without accessing them: An experimental study of dram disturbance errors, 2023. *ArXiv*, abs/2306.16093, 2023.

[20] Tu Dinh Ngoc, Boris Teabe, Alain Tchana, Gilles Muller, and Daniel Hagimont. Mitigating vulnerability windows with hypervisor transplant, 2021. pages 162–177. Conference Name: EuroSys '21: Sixteenth European Conference on Computer Systems ISBN: 9781450383349 Publisher: ACM.

[21] National Institute of Standards and Technology. Accessed: 28-6-2024.

[22] Gaoning Pan, Xingwei Lin, Xuhong Zhang, Yongkang Jia, Shouling Ji, Chunming Wu, Xinlei Ying, Jiashui Wang, and Yanjun Wu. V-shuttle: Scalable and semantics-aware hypervisor virtual device fuzzing, 2021. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 2197–2213. ACM.

[23] Sanjay Rawat, Vivek Jain, Ashish Kumar, Lucian Cojocar, Cristiano Giuffrida, and Herbert Bos. Vuzzer: Application-aware evolutionary fuzzing, 2017. In *Network and Distributed System Security Symposium*, 2017.

[24] Andrew R. Riddle and Soon M. Chung. A survey on the security of hypervisors in cloud computing, 2015. pages 100–104, 2015.

[25] Sergej Schumilo, Cornelius Aschermann, Ali Abbasi, Simon Worner, and Thorsten Holz. HYPER-CUBE: High-dimensional hypervisor fuzzing, 2020. Conference Name: Network and Distributed System Security Symposium ISBN: 9781891562617 Publisher: Internet Society.

[26] Sergej Schumilo, Cornelius Aschermann, Ali Abbasi, Simon Wörner, and Thorsten Holz. Nyx: Greybox hypervisor fuzzing using fast snapshots and affine types, 2021. pages 2597–2614.

[27] Udo Steinberg and Bernhard Kauer. NOVA: a microhypervisor-based secure virtualization architecture, 2010. In *Proceedings of the 5th European conference on Computer systems*, EuroSys '10, pages 209–222. Association for Computing Machinery.

[28] Andrew Tanenbaum and Herbert Bos. *Modern Operating Systems, Global Edition, 2014*. Pearson Higher Education & Professional Group.

[29] Chiachih Wu, Zhi Wang, and Xuxian Jiang. Taming hosted hypervisors with (mostly) deprivileged execution, 2013.

[30] Xiaogang Zhu, Sheng Wen, Seyit Camtepe, and 2022 Xiang, Yang. Fuzzing: A survey for roadmap. *ACM Comput. Surv.*, 54(11s), sep 2022.

doi: 10.14459/2024md1759337_03

SoK: Cyber Resilient Systems in different market sectors

1st Hartung Hanne
Munich, Germany
hanne.hartung@outlook.com

***Abstract*—Cyber resilience has become a major concern for both public and private sector due to the increasing number of cyber threats caused by the growing complexity of modern technology systems. As organizations rely more on digital infrastructure, their exposure to cyber-attacks grows, requiring robust measures to protect assets and ensure business continuity. In this paper, we explore the current state of art of cyber-resilient applications in the field, drawing upon existing research to provide a comprehensive overview. Our analysis begins with a detailed examination of different sectors such as finance, healthcare, transportation, energy, supply chain and communication, where we explore the state of cyber resilience application in this area. A critical observation from our study is the lack of research on other sectors such as the Consumer Staples, Real Estate and Materials sector, reflecting an overall lack of comprehensive studies in the field of cyber resilience. We then present cyber resilience products and technologies currently available on the market and judge their effectiveness, incorporating insights from an interview with a security expert from a large German enterprise who explained the company's cyber strategy and recommended several products to us. By combining the sector analysis and the products, we provide an overview of where each product is applied. Following this, we analyze which sectors have exemplary cyber resilience practices, assessing how well the sectors have implemented these products and technologies. We identify the finance sector as a prime example and find several areas for improvement in the healthcare, transportation, and communication sectors. This study aims not only to provide a clearer understanding of how different sectors are tackling the challenges of cyber resilience, but also to help developers choose the right products for their application. It also emphasizes the need for increased research in the field of cyber resilience applications, noting that it has received less attention compared to other fields despite its growing importance. To meet these challenges, both the cybersecurity and cyber resilience frameworks need to be constantly updated to keep up with the fast development of technology.**

***Index Terms*—Cyber resilience, resilience, public sector, private sector, market sector**

I. Introduction

Cyber resilience is defined as the ability to anticipate, withstand, recover from, and adapt to various cyber-attacks [1]. Another frequently referenced definition of resilience is by the National Academies of Science (NAS): "the capacity to prepare for, absorb, recover from, and more effectively adapt to adverse events." [2]. The term "cyber resilience" emerged in the early 2000s when the need to develop systems capable of surviving and recovering from cyber incidents was recognized [3]. In 2005, the UK Cabinet Office introduced the concept of cyber resilience and emphasized how important it is for companies to adapt to evolving threats and maintain critical operations. They were the first government to bring attention to this issue [4]. The foundation of inaugural Australia International Cyber Resilience Conference in 2010, was an important milestone in the development of cyber resilience. It paved the way for future advancements and discussions in this field. This was the first time the concept recieved significant attention from the academic community [3]. In light of numerous successful cyber-attacks such as the WannaCry attack in 2017 [5] or the shutdown of the Ukrainian power grid in 2015 [6], the topic of resilience has become increasingly important for IT experts in recent years. Cybersecurity aims to protect IT assets such as data [7], where cyber resilience is the ability to defend against cyber-attacks and return to a system's original state when cybersecurity fails to protect it. Cyber resilience offers the opportunity of business continuity when cyber-attacks are missed by the deployed cybersecurity solutions [8]. The focus of cyber resilience is not just on the existing cyber threats. Rather, it is about learning from them and continuously adapting the system to ensure the sustainability of its services. This concept, gaining recognition among IT experts around the world [3], is of significant importance from an information security perspective. As a result, architects and engineers are actively exploring strategies to integrate resilience principles into designs and frameworks and support them with advanced technologies. To achieve cyber resilience, best practices for business continuity, IT security, and other disciplines are combined to form a strategy addressing the current needs and goals. A company or organization can effectively adapt to cyber-attacks if it can maintain its business operations at least partially during the attack, which is very important in the areas including critical Infrastructures.

This research focuses on the application of cyber resilience in different market sectors including critical Infrastructures. The research's goal is to give an overview of the current state of art of resilience 'in the field'. It not only assesses the current landscape, but also challenges existing practices and encourages innovation in the field of cyber resilience. By critically reviewing existing tools and technologies, we aim to enable cybersecurity professionals to rethink their strategies and find more effective solutions in the fight against cyber threats.

We begin with an overview over the research that has been done in the sectors finance, healthcare, transportation, energy, Supply chain and communication. This introduction sets the

doi: 10.14459/2024md1759337_04

stage for a more in-depth analysis of the practical applications of cyber resilience strategies in these sectors. To understand the significance of our focus, it is important to recognize that the integration of resilience measures varies greatly across different sectors. By exploring the current resilience practices in these areas, we can identify both common challenges and sector-specific barriers.

Following this sector analysis, the next phase of our research involves an evaluation of various cyber resilience products and technologies currently available on the market. Our aim is to assess their effectiveness and applicability in the context of the previously discussed sectors. This will include a review of the technologies' ability to withstand cyber threats and their adaptability to ever-evolving cyber risks. Our assessment is also supported by insights gained from the interview with the expert. He recommended to us the use of Multi-Factor Authentication, Vulnerabilities Assessment, Penetration Testing and Red Team Testing. He also mentioned potential disadvantages of these approaches.

In addition, our evaluation will not only focus on the effectiveness of these products but will also consider their deployment and operational effects in practice. By examining where and how these products have been implemented, we assess their suitability and effectiveness for each sector. This approach allows us to highlight best practices for implementing resilience solutions effectively across diverse environments. This practical guide is designed to help industry members make informed decisions about the implementation and integration of cyber resilience technologies to protect critical infrastructure.

II. Methodology

A. Study retrieval

A literature search was conducted in IEEE Xplore Digital Library, ACM Digital Library, Science Direct, Vectoral Publishing and Researchgate for studies published between 2010 and 2024. Following up, we searched for studies published at the International Conference on Electronics, Communications and Control Engineering, the IEEE International Conference on Cyber Security and Resilience, the International Conference on Reliability and Quality in Power Supply and the International Conference on the EU Cyber Security and Resilience Acts. A comprehensive search strategy was developed in which search terms were combined and used in two different sets (set 1: cyber resilience, recovery, security, critical infrastructure, set 2: supply chain, finance, healthcare, communication, transportation, energy) to find the studies. Most of the studies we found were published between 2020 and 2024, indicating that this is an emerging topic that has only recently started to gain recognition.

B. Study selection

The titles and abstracts of the papers were then manually reviewed against the inclusion and exclusion criteria. Inclusion criteria were review articles, conference papers and original articles published in English between January 2010 and May 2024, focusing on the application of cyber resilience in different sectors. Studies reporting on the state of art and recent advances in cyber resilience were also included and used to obtain additional sources. Based on the sectors different other studies and surveys researched [9], [10], we chose to analyze the Healthcare, Energy, Transportation, Finance, Supply chain and Communication sector. We limited our selection to papers that address these sectors. Exclusion criteria included studies related to very specific smaller areas of cyber resilience, for example Cyber resilience in Australian small businesses [11], given the scope of the paper. Singular papers researching one of the sectors we excluded were also not used, based on the decision that an in-depth analysis of a sector requires several sources. From both academic and industry perspectives, we looked for studies on currently on the market available cyber resilience products. Additionally, we searched for publications by firms and organizations stating which of these products they use and what their resilience strategy looks like. We found almost no papers with actual technical details published by a company itself, after searching Websites and company blogs from several big companies, for example Microsoft and Telekom. The companies tend to publish guides for achieving Cyber resilience [12], [13], but not their own strategy. To gain insight on an actual cyber resilience strategy in a company, we conducted an interview with the Head of Security at a large German enterprise. He disclosed the company's resilience strategy to us and explained in detail how this strategy was developed after falling victim to a cyber-attack recently. Some of the products analyzed in the paper were recommended and explained to us by him.

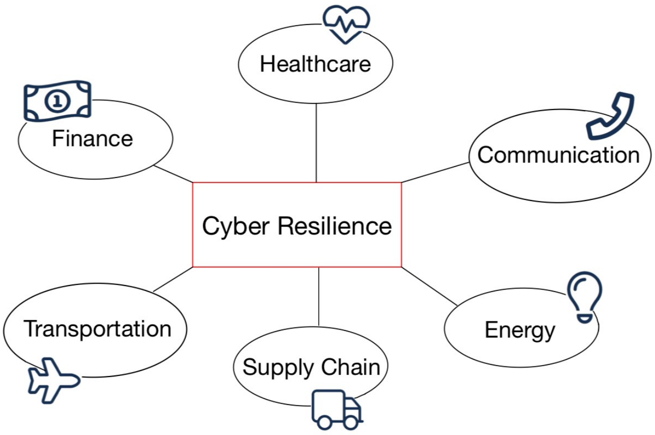

Fig. 1. Overview of the different sectors

C. Main themes

We found a lack of research on other sectors, such as the Consumer Staples, Real Estate and Materials sectors and the chemical and pharmaceutical industry, which is the reason for the limitation of the sectors to Healthcare, Energy, Transportation, Finance, Supply chain and Communication sector. The sector that has been studied most intensively is finance, a sector whose cybersecurity and resilience received a lot of attention as it lost approximately 20 billion USD due to cyber-attacks in the past 20 years [14]. The lack of research on other sectors can be explained by the lack of research on the topic in general. Most papers on the application of cyber resilience have only recently been published. Many organizations may not yet have developed a solid cyber resilience strategy, let alone made it available to the public. This could be due to the risk that by publishing such information, they expose themselves to a greater risk of cyber-attacks.

III. Applications of cyber resilience

A. Healthcare sector

Garcia-Perez et al. [15] use PLS-SEM to analyze cyber resilience in healthcare. In this research, digital resilience is considered as part of the overall performance of healthcare institutions and not as an independent category. Digital transformation efforts are explored as significant for both same-level digital resilience and higher-level institutional resilience. The security aspect of digital transformation in healthcare is modeled using three constructs from the Blanchet et al. [16] framework: Knowledge and Resources, Awareness of Risk, and Partnerships & Supply Chain. These constructs focus on understanding and integrating cybersecurity measures, managing uncertainties, and considering interdependencies within the healthcare sector, emphasizing the relationship between digital transformation and healthcare resilience.

Boddy et al. [17] discussed improving cyber resilience in healthcare systems by using advanced visualization techniques and data analytics to detect unusual data behaviors. They explored a sophisticated system employing machine learning algorithms to analyze data patterns and user profiles, focusing on three primary services:

1. The Active Directory Domain Services (AD DS) server, which manages access to organizational infrastructure, including security group user accounts and passwords.

2. The Patient Administration System server, which provides access to patient data for viewing or modification.

3. The Electronic Prescribing server, which could allow attackers to monitor medication doses and prescriptions. Monitoring port mapping servers, which are essential in hospital networks, is said to be a difficult task that requires significant resources. By cleansing and preparing data, cybersecurity analysts can better identify anomalous activities and minimize threats. The integration of ML algorithms helps IT departments in hospitals or other organizations to detect potential cyber-attacks within their large data infrastructures.

The UK National Health Services (NHS) introduced a program to advance cyber resilience after the WannaCry attack in 2017, which disrupted the functionality of the NHS. The strategy is based on five key pillars: Identifying Critical Areas, Unified Defense, Training their staff, Secure Technology, Response and Recovery. Ghafur et al. [18] point out several problems with the program: no clearly defined responsibility in case of an attack, no full catalogue of the software and hardware the NHS uses - leading to a lack of awareness of vulnerabilities - and chronic underinvestment in healthcare IT. Limited budget is a frequent problem most of the healthcare papers name [15], [16], [18].

Porter et al. [19] describe the methodology used to map the Health Insurance Portability and Accountability Act (HIPAA) Security Rule requirements to the CERT® Cyber Resilience Review (CRR) practice questions. This mapping allows healthcare and public health organizations to use CRR results to evaluate their cyber resilience and baseline compliance with the HIPAA Security Rule and the NIST Cybersecurity Framework (CSF). The CRR and HIPAA Security Rule were both independently mapped to the NIST CSF. The authors concluded that the CRR covers all aspects of the HIPAA Security Rule. It enables organizations to use the CRR as an indicator of compliance with the Security Rule.

B. Energy sector

The electricity sector is considered critical infrastructure because it powers the life-line infrastructure sectors such as transport, water supply and sanitation, communication, and hospitals. The electric sector is therefore subject to mandatory and enforceable cybersecurity standards according to Ram [20]. Using examples of past cyberattacks on electricity utilities, they presented best practices for electricity infrastructure resilience. They propose using the NIST Cybersecurity Framework as an electric utility/organization. They suggest implementing an Intrusion Detection System (IDS) and an Intrusion Prevention System for Advanced Metering Infrastructure, used in electrical utilities to connect with their customers.

The NIST Cybersecurity Framework was applied in the context of power systems by Pöyhönen et al. [21] as well by applying CRR to a single electricity company. They discuss the importance of cyber resilience in maintaining the reliability of the power system, especially in Finland, where electricity generation is highly distributed. A SWOT analysis is used to evaluate and improve the cybersecurity level of an individual electricity company. The Resilience metrics framework proposed by Linkov et al. [22] is applied to extend the prepared planning beyond the cyber structure.

Nguyen et al. [23] discuss the complexity and vulnerabilities of modern smart grids, emphasizing the need for robust cyber defense strategies. They point out a common problem with resilience in power systems: Most methods for detecting cyber intrusions in smart grids rely on outdated IT techniques and focus on traditional attacks such as DoS. This makes real-time cyber vulnerability assessment a challenge. The conventional, computationally intensive techniques for detecting attacks must be applied in real time in power systems. The study reviews existing detection, protection, and mitigation techniques to enhance grid resilience, emphasizing the importance of both structural and operational resilience measures. Machine learning and artificial intelligence are mentioned as promising approaches for improving attack detection and response. They conclude with the need for continuous advancements in monitoring, protection, and resilience strategies to protect against evolving cyber threats.

Most of the papers on the power sector mention the power failure in the Ukraine in 2015 [6]. They were all published after the year 2015, suggesting that this attack was the first to draw attention to resilience in this sector.

C. Transportation Sector

The cyber resilience of autonomous mobility systems has been studied quite extensively in the literature. Zou et al. [24] find that despite a growing number of internet-connected vehicles, there is a lack of awareness and protection of autonomous vehicles from cyber threats. They identify different types of cyber-attacks at both vehicle and system levels and review current practices and strategies to improve cybersecurity in Autonomous Mobility Systems (AMS). Key strategies at the vehicle level include creating layers and separation in cyber components and deploying independent data collection procedures. At the system level, maintaining redundancy in transportation capacity, providing alternative non-cyber dependent modes, and creating different subsystems to minimize attacks are recommended. The paper concludes with a call for more modeling-based research to quantitatively evaluate the benefits of these strategies and improve understanding of cyber resilience in AMS.

Another popular research topic is cyber resilience in airports. This interest is likely due to the large amount of confidential data and the strict security standards required in this environment.

Lykou et al. [25] investigate the implementation of cybersecurity measures in airports to improve cyber resilience. They find that the integration of Industrial IoT in smart airports and the increasing use of 'Bring your own device' by travelers bring new challenges. Smart airports are described as an airport solution that, unlike traditional airports, enables the control and monitoring of numerous systems from a remote location [26]. Results show that smart airports have a higher implementation rate of cybersecurity practices compared to 'basic' airports. Best practices for smart airports are introduced. To conclude, they summarize significant security gaps - including poor implementation of intrusion detection systems and BYOD controls - and point to the need for airport trust frameworks and increased security awareness. They propose a comprehensive cybersecurity framework tailored to the unique needs of smart airports. The authors emphasize the importance of collaboration among airlines, airports, vendors, and regulators to address these cybersecurity challenges effectively.

Mathew [27] explored controls for cybersecurity and cyber resilience within airports. The study included an examination of airport intelligence classifications and an analysis of cybersecurity threats. The Internet of Things (IoT) is identified as a critical technology in airports, as in the other papers on resilience in airports, improving communication among intelligent systems and devices. This technology has significantly increased cyber resilience and operational efficiency. However, the growing integration of airport services and facilities with the IoT also heightens vulnerabilities to network attacks, underscoring the critical importance of robust cyber resilience measures in airports.

D. Financial sector

Cyber resilience in the financial sector gained a lot of attention over the past few years. This increased focus can be attributed to the sector's significant role in managing large amounts of sensitive data and its critical importance to the global economic stability [28]. The financial sector often handles the most money compared to other sectors due to its core function of managing investments, assets, and financial transactions worldwide.

Asset management firms such as BlackRock manage trillions of dollars in assets. As of 2023, BlackRock alone managed approximately \$9.4 trillion in assets [29], illustrating the immense scale of capital within the financial sector. The huge value of the assets managed by financial institutions makes them prime targets for cyber-attacks, necessitating robust cybersecurity measures.

Dupont [30] examines the urgent need for cyber resilience in financial institutions. He argues that the current "prevent and protect" approach is insufficient and that incorporating a cyber resilience strategy into the risk management framework is essential. He briefly traces the scientific history of cyber resilience and outlines the five key dimensions of organizational resilience: networked, adaptive, dynamic, practiced, contested. The author analyzes three types of institutional approaches: The first is marketing cyber resilience through consulting firms and security firms promoting cyber-resilience. They do so through reports that highlight the benefits of resilience, often linking these benefits to the firm's own products and services. Second is standardizing cyber-resilience through its embedding in cybersecurity standards, such as the NIST Cybersecurity framework. Third is regulation through the development and enforcement of compliance tools and standards to improve cyber resilience by regulators. He notes that there is a lack of conceptual clarity in all of these areas as they are still at an early stage of development and application.

Overall, the author identifies a gap in research on cyber-resilience metrics and calls for more recent studies to effectively assess and improve cyber resilience.

Pinckard et al. [31] detailed the methodology and observations from mapping the declarative statements in the Federal Financial Institutions Examination Council and Cybersecurity Assessment Tool (CAT) to the best practice questions in the Cyber Resilience Review (CRR). This mapping allows financial organizations to use CRR results to measure their cyber resilience and assess their current baseline against the NIST Cybersecurity Framework. The results show that while CAT and CRR aim to improve cyber resilience, there are gaps and overlaps that need addressing for more effective implementation. The paper concludes with recommendations for using this mapping to improve organizational resilience.

Gallagher et al. [32] emphasize that cyber-attacks pose a risk to Canada's financial system by disrupting key participants' operations, such as large financial institutions or financial market infrastructures (FMIs). The attackers targeting these elements vary in motives, ranging from financial theft to business disruption. Canadian financial institutions and FMIs proactively build defenses and collaborate with each other and the federal government to minimize these threats. Key initiatives include the Public-Private Partnerships (PPP): the Canadian Cyber Incident Response Centre (CCIRC) and the Joint Operational Resilience Management (JORM) program, which enhance information sharing and crisis response. The authors point to the progress made in improving the resilience through those PPPs. Despite significant investments in cybersecurity, the financial sector must constantly evolve to maintain its resilience.

Crisanto et al. [33] published several key findings related to cyber resilience in the financial sector, focusing on banks. First, international regulatory initiatives have emphasized the importance of global cooperation, leading to widely accepted guidelines such as the FSB cyber Lexicon and the CPMI-IOSCO guidance. Second, regulatory approaches to cyber resilience are evolving to align specific regulations with broader principles and create flexible but effective frameworks to address ever-changing cyber threats. Third, regulatory frameworks now commonly incorporate vulnerability assessments, penetration testing, and red team testing to evaluate banks' cybersecurity capabilities, with an increasing focus on real-time detection and response measures. Fourth, the emphasis on cybersecurity awareness and training within banks is growing, recognizing the critical role of human factors alongside technical solutions. Finally, increased cross-border cooperation and information sharing between government regulators and private financial institutions is being sought. The aim is to better manage systemic risks and improve the overall resilience. This paper once again suggests Public-Private Partnerships for achieving financial cyber resilience.

Fedotova et al. [34] investigated the cyber resilience of credit organizations against the backdrop of the digitalization and automation of financial services. The authors point to the steady increase in data leaks over the past decade, with the financial sector being a significant target. The study underscores the importance of both planned and adaptive cyber resilience strategies. The Central Bank of Russia's implementation of smart technologies, such as continuous monitoring of unauthorized transactions, is presented as a key development trend for improving cybersecurity. FinCERT ASOI, an initiative by the Central Bank of Russia, improves cyber resilience within financial organizations by facilitating data exchange between 826 participants. These include credit institutions, telecom operators, and government agencies. To further strengthen cyber defenses, international cooperation is being promoted, with four national banks from the Eurasian Economic Union (EAEU) joining forces to create a unified information space for cyber resilience in the financial sector.

E. Supply chain

The main supply chain resilience theories were proposed in the early 2000s following the attacks of 11 September 2001, the publications in recent years have been sparse. Khan et al. [35] published one of the first journals on cyber resilience in supply chains, pointing out that this critical topic has not been the focus of academia despite growing interest from the business community. They name the convergence of information technology and supply chains as a potential reason for this gap, as these two disciplines naturally overlap in the context of supply chain cyber-risk and cyber resilience. They propose a research agenda to integrate cyber-risk into existing supply chain resilience frameworks. Key recommendations for academia and industry include fostering collaboration, implementing strategic risk management practices, and promoting a proactive culture towards cyber threats.

Davis [36] recognizes similar to Khan et al. [35] the significant efforts already made to ensure security and resilience in the physical aspects of supply chains. They also mention a lack of spending on cyber resilience in supply chains. They emphasize the urgent need for companies to protect sensitive information, as in the previous papers. The author advocates for an information-centric approach. Companies must identify, classify and protect their data throughout the procurement process and beyond. They outline five key steps to improving cyber resilience: supply chain mapping, capability building, sharing information and expertise, using common standards and frameworks, and continuous measurement and auditing. Cyber resilience is seen as an evolving concept that requires comprehensive risk assessment and collaborative efforts across the entire supply chain.

Boyes [37] examines the cyber-resilience of supply chains delivering physical products and services. This paper also emphasizes the need to address cybersecurity issues beyond just technical aspects, including personnel, process, and physical components. They introduce a cybersecurity model based on the Parkerian hexad [38], particularly relevant to complex, time-critical, and cyber-physical systems, and its application in the construction industry supply chain in the UK. Its applicability extends to other supply chains as well. They advise supply chain managers to pay attention to external vulnerabilities when selecting technologies, as the use of the cloud or remote storage of data, for example, can open up major security gaps.

F. Communication sector

Similar to the supply chain research, the research results for cyber resilience in communication networks were rather sparse, the topic hasn't received much attention from the academic world. Buinevich and Vladyko [39] investigated the concept of cyber resilience within wireless communication network technologies, focusing on applications for Intelligent Transportation Systems (ITS). Their research centered on cyber resilience in motor transport systems, including Vehicular Adhoc Networks (VANETs). The authors conducted a comprehensive analysis of cyber-attacks targeting VANETs and ITS and identified the top 10 cyber threats. They recommend 3 main options for constructing cyber-resilience telecommunication components for ITS: "standardized" (DSRC/802.11p), "mobile" (LTE-V) and their "communication mix" are allocated. In addition, several unresolved issues and future research directions were identified, such as threat formalization, vulnerability mitigation, network management integration, and prediction and modelling of cyber resilience in VANETs and ITS.

Bellini et al. [29] focus on railway communication networks and point out that transport communication systems like ITS and railway communication networks are receiving increasing attention. This research emphasizes the importance of cyber resilience in ensuring the safety and efficiency of railway operations, addressing similar challenges and using comparable methodologies as those found in ITS studies. By formalizing the Cyber Resilience Ontology using a UML Profile and Bayesian Networks, it offers a quantifiable method for resilience assessment. The approach is validated through the case study on railway communication systems to demonstrate its applicability. The main results highlight the role of Functional Dumping Capacity (FDC) as a resilience indicator, which includes Buffer Capacities, Flexibility, Margin, and Tolerance. Future work includes integrating the methodology into industrial risk management and exploring alternative formal analysis methods. Similar challenges include the integration of different technologies and protocols and balancing safety and performance.

Both papers discuss countermeasures such as strong encryption, secure communication channels, and resilience engineering techniques to mitigate cyber threats.

G. Products and Technologies

This section reviews products and technologies developed to evaluate and improve cyber resilience. It describes various tools and technologies used for cyber resilience, including the NIST Cybersecurity Framework, the Cyber Resilience Review (CRR), Vulnerability assessment, Penetration testing, Red team testing, Multi-Factor Authentication (MFA), Cybersecurity Assessment Tool (CAT) and the Cyber Resilience Assessment Tool (CRAT). Each tool and technology is described individually, followed by a comparison based on several criteria as demonstrated in Figure 2. The comparison criteria are strengths and weaknesses as well as the areas in which they are most commonly used. While all the products can be used in different sectors, some are particularly suitable for certain industries.

1) NIST Cybersecurity Framework: The NIST Cybersecurity Framework (CSF) offers comprehensive guidance for industry, government agencies, and other organizations to effectively manage cybersecurity risks. It provides a taxonomy of high-level cybersecurity outcomes that are applicable to any organization, regardless of size, sector, or maturity level. This framework helps organizations understand, assess, prioritize, and communicate their cybersecurity efforts. The framework does not prescribe specific methods for achieving these outcomes, but instead points to online resources that provide additional guidance on practices and controls that can be deployed [40].

2) Cyber Resilience Review: The Cyber Resilience Review (CRR) is a voluntary assessment designed to evaluate an organization's operational resilience and cybersecurity practices, providing recommendations for improvement. It aligns with the NIST Cybersecurity Framework [41].

3) Vulnerability assessment: Vulnerability assessment is the systematic examination of an information system and its controls and processes. The goal is to determine the adequacy of security measures, identify security deficiencies, provide data for predicting the effectiveness of proposed security measures and confirm the adequacy of such measures after implementation [42].

4) Penetration testing: A test methodology in which assessors, using all available documentation (for example, system design, source code, manuals) and working under specific constraints, attempt to circumvent the security properties of an information system [43].

5) Red team testing: A controlled attempt to compromise the cyber resilience of an entity by simulating the tactics, techniques and procedures of real threat actors. It is based on targeted threat intelligence and focuses on an entity's people, processes and technology, with minimal foreknowledge and impact on operations [44].

6) Multi-Factor Authentication: Multi-Factor Authentication (MFA) enhances security by requiring multiple verification factors to confirm a user's identity. It evolved from Single-Factor Authentication (SFA), which relies on just one factor like a password, to Two-Factor Authentication (2FA), adding another layer such as a smart card or phone. MFA also includes biometric factors like fingerprints or facial recognition to improve security [45].

7) Cyber Security Assessment Tool: The Cyber Security Assessment Tool (CSAT) is a software product to quickly evaluate the status of an organization's security and provide fact-based recommendations for improvement. The tool gathers relevant security data from the hybrid IT environment by monitoring endpoints, automated scanning of applications or active directories and employing a questionnaire. Additionally, CSAT collects data on organizational controls, policies, and other key indicators [46], [47].

8) Cyber Resilience Assessment Tool: A software solution designed to measure an organization's current cyber preparedness. It involves analyzing various aspects of the organization's IT infrastructure, policies, and procedures to identify weaknesses and areas for improvement. The tool provides insights and recommendations to enhance the organization's overall cyber resilience [48]. Many offerings are based on the NIST Cybersecurity Framework.

IV. Evaluation

This study explored the current landscape of cyber-resilient applications in different market sectors, with a focus on the healthcare, energy, transportation, finance, supply chain, and communication sectors. These are the sector specific results:

(Healthcare) There is extensive research on cyber resilience with practical applications in healthcare systems. The use of

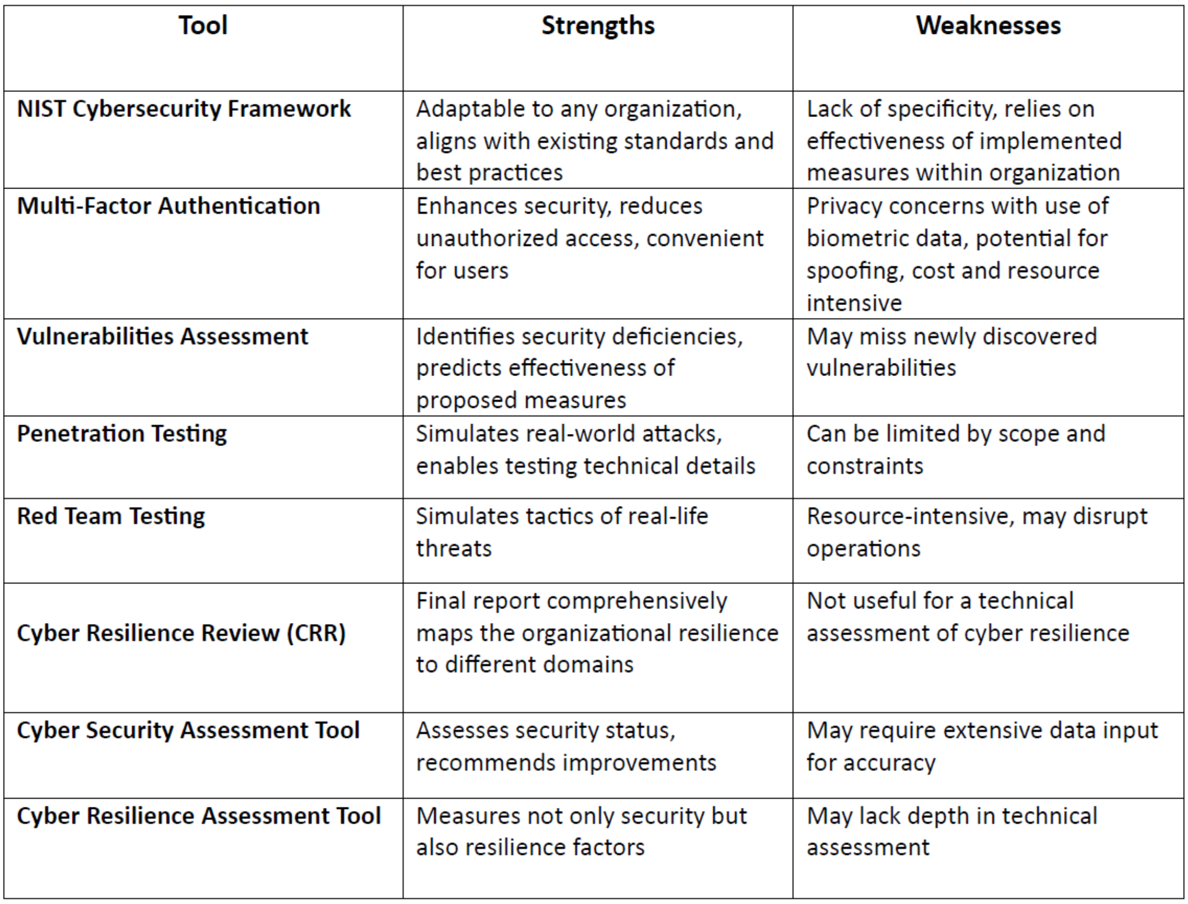

Tool	Strengths	Weaknesses
NIST Cybersecurity Framework	Adaptable to any organization, aligns with existing standards and best practices	Lack of specificity, relies on effectiveness of implemented measures within organization
Multi-Factor Authentication	Enhances security, reduces unauthorized access, convenient for users	Privacy concerns with use of biometric data, potential for spoofing, cost and resource intensive
Vulnerabilities Assessment	Identifies security deficiencies, predicts effectiveness of proposed measures	May miss newly discovered vulnerabilities
Penetration Testing	Simulates real-world attacks, enables testing technical details	Can be limited by scope and constraints
Red Team Testing	Simulates tactics of real-life threats	Resource-intensive, may disrupt operations
Cyber Resilience Review (CRR)	Final report comprehensively maps the organizational resilience to different domains	Not useful for a technical assessment of cyber resilience
Cyber Security Assessment Tool	Assesses security status, recommends improvements	May require extensive data input for accuracy
Cyber Resilience Assessment Tool	Measures not only security but also resilience factors	May lack depth in technical assessment

Fig. 2. Products Overview

machine learning and data analytics to detect cyber threats is promising. Chronic underinvestment hinders cyber resilience efforts. There is also a need for better-defined strategies for integrating cyber resilience measures in hospitals.

(Energy) The sector has begun to adopt best practices for cyber resilience, including NIST Cybersecurity Framework and advanced intrusion detection systems. Many existing detection techniques are outdated and not suitable for real-time applications in modern smart grids, so innovation is needed. There is also a need for continuous development of new monitoring and protection strategies.

(Transportation) Despite extensive research on autonomous vehicles and smart airports, other industries such as marine or road and rail have not received as much academic attention. Studies by Zou et al. [24] and Lykou et al. [25] emphasize the implementation of layered cybersecurity measures and the integration of IoT. Despite these advancements, there is a need for more modeling-based research to quantitatively assess cyber resilience strategies. The transportation sector is facing rapid development, especially with the integration of renewable energy sources. This evolution opens up numerous research opportunities and potential new dangers at the same time. It also underscores the need for ongoing updates and advancements in resilience strategies. Ensuring robust cybersecurity and resilience in this sector is important to support its future growth and innovation.

(Finance) The financial sector has made significant progress in integrating cyber resilience into its risk management frameworks. Research by Dupont [30], Pinckard et al. [31], and Gallagher et al. [32] highlights the development of Public-Private Partnerships and regulatory initiatives as key strengths. However, the sector faces challenges related to conceptual clarity and the development of effective metrics for assessing cyber resilience. Given the continued interest of hackers in the financial sector, continuous adaptation to evolving cyber threats remains essential. This sector will always be a prime target because of the substantial financial gains attackers can achieve.

(Supply Chain)

In the supply chain sector, theories of resilience to cyber-attacks were integrated early on, but recent academic research is sparse. Studies by Khan et al. [35] and Boyes [37] emphasize the need for strategic risk management and collaboration across the supply chain. There is a notable gap in research on integrating cyber-risk management into supply chain frameworks, underscoring the importance of fostering

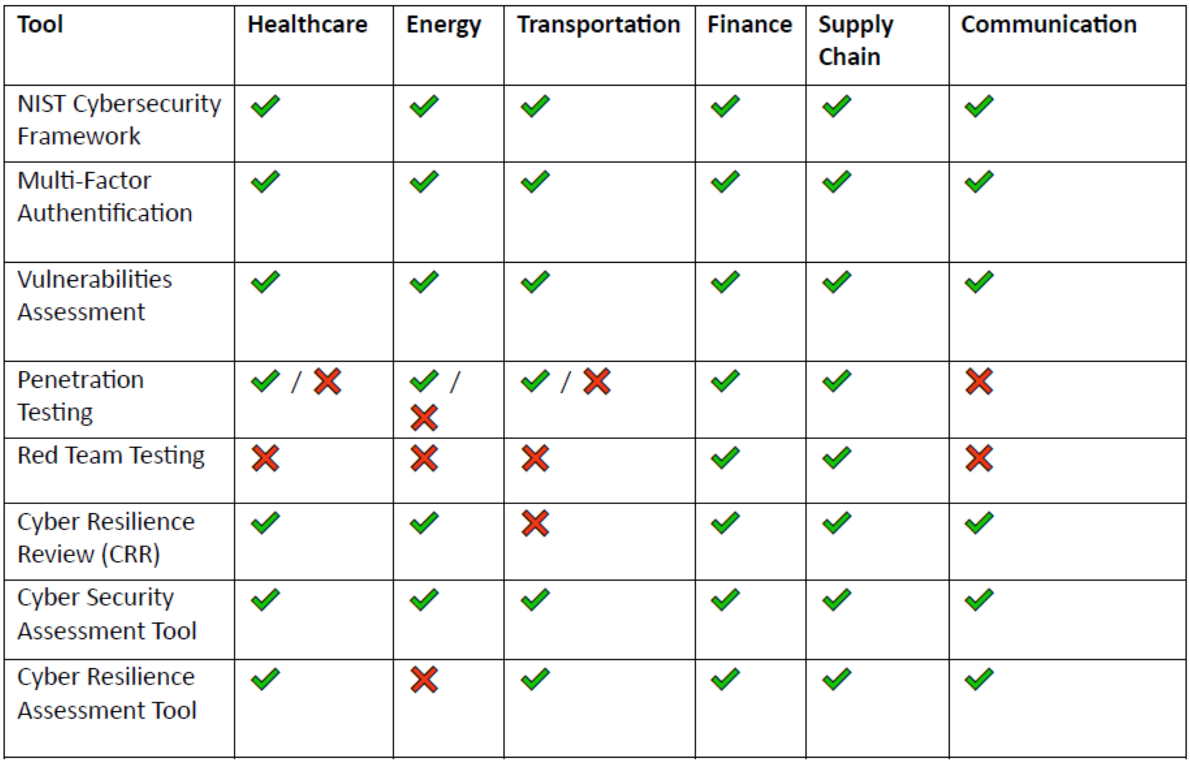

Tool	Healthcare	Energy	Transportation	Finance	Supply Chain	Communication
NIST Cybersecurity Framework	✓	✓	✓	✓	✓	✓
Multi-Factor Authentification	✓	✓	✓	✓	✓	✓
Vulnerabilities Assessment	✓	✓	✓	✓	✓	✓
Penetration Testing	✓ / ✗	✓ / ✗	✓ / ✗	✓	✓	✗
Red Team Testing	✗	✗	✗	✓	✓	✗
Cyber Resilience Review (CRR)	✓	✓	✗	✓	✓	✓
Cyber Security Assessment Tool	✓	✓	✓	✓	✓	✓
Cyber Resilience Assessment Tool	✓	✗	✓	✓	✓	✓

Fig. 3. Product Application in sectors

collaboration between IT and supply chain disciplines.

(Communication) The communication sector's research, particularly in wireless communication networks and railway communication systems, is still in its early stages. Studies by Buinevich and Vladyko [39], and Bellini et al. [29] highlight innovative approaches, such as the use of cyber resilience ontologies and formalized threat models. However, the sector faces challenges related to standardization and balancing security with performance. Future research should focus on solving these problems.

Overall, we noticed a lack of research on a lot of subtopics of the respective sectors and especially on the other market sectors this study excluded. Going forward, it is important to conduct research on these other sectors and increase the overall amount of research in the field of cyber resilience.

The results of the tool and technology research, which were limited due to the scope of the study: Figure 2 shows it is important to be critical when evaluating these tools. Many of the common assessment offerings provide an overview of the current cyber resilience state of the organization, but do not address specific technical problems. The interview yielded that Red teaming and penetration testing, in contrast, are more effective methods for identifying actual security gaps and testing resilience. The best planning and strategies are ineffective if they are not tested against real-world use cases. By simulating real-life attack tactics, these methods provide insights into the robustness of an organization's resilience measures and help identify areas for improvement that theoretical assessments might miss. The expert warned us that firms and organizations run into the danger of assuming they are secure based on these assessments, only to find out they are not adequately prepared when an actual cyber incident occurs [49]. The company the expert works for was the victim of a cyber-attack recently after relying solely on the results of risk assessments. This false sense of security can lead to major vulnerabilities being overlooked. From the interview, we learned that it is crucial for organizations to start Penetration Testing and Red Team Testing as soon as possible [50].

Figure 3 shows commonly used tools, highlights sectors that serve as good examples of cyber resilience and points out areas that need improvement. A tick in the respective box means that we have found more than one source - from the academic papers we used - for the implementation of the respective technology. A cross means that we found none. A tick / a cross means that we found sources that show that although there are offerings for this technology in the respective industry, it does not yet appear to be common practice [51]–[55].

Across the board, the NIST Cybersecurity Framework stands out as a universally adopted tool. Its flexibility and alignment with standards such as ISO/IEC 27001 and ISO/IEC 27002 make it a cornerstone of cyber resilience strategies across all sectors, including healthcare, energy, transportation, finance, supply chain, and communication [19]–[21], [30], [31], [40]. This widespread adoption underscores the framework's effectiveness in providing a structured approach to managing and minimizing cyber risks. Its use is recommended.

Similarly, Multi-Factor Authentication (MFA) is implemented across all sectors. MFA's ability to improve security through multiple layers of verification makes it an important tool in preventing unauthorized access and protecting sensitive information, an essential factor in critical Infrastructures. Its broad application reflects its key role in strengthening security postures universally. MFA is useful in every sector [40].

Vulnerability and Cyber Security Assessments are also commonly used across all sectors. These assessments are important for finding security weaknesses and evaluating the effectiveness of existing measures. By systematically investigating potential vulnerabilities, companies can proactively address security gaps, making this tool essential for maintaining robust cyber defenses. Every market sector needs Vulnerability and Cyber Security Assessments to uncover risks and ensure that their cybersecurity measures are both effective and up to date [19], [21], [23], [31], [33].

Penetration Testing is another tool that sees universal application. It is worth noting, however, that Penetration Testing is standard practice and particularly common in the financial and supply chain sectors. In the finance sector, this is due to regulatory requirements that mandate annual penetration tests to ensure robust security measures [56]. In the supply chain sector, the extensive use of penetration testing could be attributed to the availability of sufficient funding to enable companies to invest in comprehensive security testing.

However, the other sectors such as healthcare, energy, transportation, and communication also need more penetration testing [50]. The healthcare sector, dealing with sensitive patient data, requires penetration testing to protect against unauthorized access and to comply with healthcare data protection regulations. The energy sector, which operates critical infrastructure, needs penetration testing to secure SCADA systems and prevent disruptions to essential services. In transportation, securing connected vehicles and ensuring operational continuity through penetration testing is important to prevent potential disruptions. The communication sector, managing large amounts of data and infrastructure, also benefits from penetration testing to maintain the integrity and security of communication networks against potential cyber threats. While we found approaches in these sectors [51]–[55], it is still far too little, and the implementation of penetration testing in these areas needs to be improved. The same applies to Red Team Testing, a tool that is generally underused, perhaps due to its higher costs.

While these tools are widely adopted, sector-specific observations show varying levels of cyber resilience. The financial sector and supply chain sector are exemplary models of cyber resilience. These sectors not only use the universally adopted tools but also integrate Red Team Testing and the Cyber Resilience Review (CRR) into their strategies. Red Team Testing, which simulates tactics of real threats, and the comprehensive CRR framework significantly improve these sectors' ability to withstand and recover from cyber-attacks. Their proactive and thorough approach sets the standard for other industries.

In contrast, there is still considerable room for improvement in the healthcare and communications sectors.

In the healthcare sector, Red Team Testing and Penetration Testing are notably absent. As evaluated earlier, the sector would profit from implementing those as standard practices. It is surprising that those practices are not part of the requirements of the Health Insurance Portability and Accountability Act, although they should be [51].

The same applies to the transportation sector. We found some Penetration Test offers for logistics [52], aviation [53] and vehicles [54], using DuckDuckGo as a search engine and the query "penetration testing" combined with the respective industries. Many of the regulations for the transportation sector recommend it, such as the European Union Aviation Safety Agency, the International Civil Aviation Organization or the American Public Transportation Association, but it is not yet a standard practice.

The communication sector also lacks Red Team Testing and Penetration Testing entirely, which needs to change.

Overall, while the adoption of key cyber resilience tools in various sectors is commendable, the healthcare, transportation and communication sectors need to increase their strategies by incorporating more advanced and comprehensive testing and assessment tools. We see approaches in the sector specific regulations, such as the HIPAA [57] or the Communications and Cyber Resiliency Toolkit from CISA [58], but the implementation is often insufficient. Furthermore, these regulations are not as strict as those in the financial sector.

All other sectors, including the supply chain sector, should take the financial sector as an example to improve their own cyber resilience measures. By adopting similar standards and regulations, these sectors can better protect themselves against cyber threats and ensure the security and continuity of their critical operations.

Another observation is the lack of adoption of resilience-specific tools compared to general cybersecurity measures. Products known for cybersecurity measures - such as MFA - are more widely used than resilience technologies. This might be due to the fact that the concept of resilience has only recently gained recognition.

V. Related Work

To our knowledge, there are almost no studies that comprehensively explore the application of cyber resilience across multiple market sectors as we have done in this paper. This gap in the literature is likely due to the broad scope of the topic and the lack of in-depth examination of the individual aspects of cyber resilience within each sector. Most existing research tends to focus on specific sectors or aspects of cyber resilience, rather than providing a broad, comparative analysis across different industries. This study aims to fill that gap by offering a holistic overview and evaluation of cyber resilience practices and technologies in diverse market sectors.

Hidaifi et al. [9] conducted a comparative analysis of various popular cyber resilience tools, giving valuable insights to researchers, practitioners, and organizations in selecting the best practices for enhancing cyber resilience. They shared key findings, identified limitations and problems, and suggested future research directions. To our knowledge, this survey is the only paper that offers an analysis across several market sectors, making it a significant contribution to the field. However, it differs from our approach because it is a broad survey that compiles a large amount of information without focusing on one point. In contrast, our study focuses explicitly on particular sectors and tools, providing a detailed, sector-specific analysis and evaluation.

Safitra et al. [10] analyzed research opportunities in the field of cyber resilience. They identified a significant gap in the research, noting that the topic has received less attention compared to other disciplines.

VI. Discussion

A. Limitations

This study faced several limitations that need to be addressed in future research. Firstly, there is a notable lack of academic work on cyber resilience in a lot of market sectors, such as the Consumer Staples, Real Estate and Materials sectors. We found only one paper mentioning the state of art in the Consumer Staples sector [59]. We encountered similar difficulties when searching for relevant studies in the other sectors. This study therefore focused on the few well-researched sectors Healthcare, Energy, Transportation, Finance, Supply Chain, and Communication, excluding others due to limited available research. This lack of research poses a challenge as it limits our ability to formulate comprehensive and sector-specific findings on the current state of cyber resilience. Additionally, relying on existing literature means that the findings are limited by the quality and scope of the available studies. We found only one comprehensive related work, the survey by Hidaifi et al. [9] that tried to cover the application of cyber resilience across several market sectors. However, this survey was broad in scope and lacked detailed, sector-specific insights, making it insufficient for a thorough understanding of the subject. Another limitation is the lack of publicly available, technically detailed information on the cyber resilience strategies adopted by individual companies. This lack of transparency makes it difficult to assess the effectiveness of current practices and identify best practices across different industries. In addition, the expert preferred to remain anonymous for security reasons, so we chose not to include a protocol of the interview in the appendix. To verify the accuracy of their statements, we have provided additional sources where the information about the products can be found as well [40]–[48]. Despite these difficulties, the interview was an enrichment for our study.

B. Future Work

Future work should focus on gaining deeper insights into the cyber resilience strategies of different companies. Detailed case studies and direct collaboration with businesses can provide valuable information on the practical implementation of cyber resilience measures. In addition, there is a need for continuous and expanded research efforts in the field of cyber resilience. Researchers should aim to include a wider range of sectors, especially those that are underrepresented in the current literature, such as marine transportation, public administration, defense, education, Consumer Staples, Real Estate and the chemical and pharmaceutical industries. Given the gaps in the other market sectors, future research should prioritize exploring cyber resilience within these under-researched sectors. The paper [59] showed that sectors such as Consumer Staples, Consumer Discretionary, Technology, Materials, Industrials, and Real Estate react differently to cyber policy news. However, we currently lack comprehensive research on how these sectors specifically respond to cyber policy news. This applies in general: while it is known that some of the other sectors improve their security, the technical details and the exact measures they adopt remain under-researched. We also need more studies that comprehensively cover multiple market sectors. This approach will help to understand cross-sector impacts and shared challenges, which will facilitate the development of more effective and universally applicable resilience strategies. As the threats evolve, ongoing studies are essential to develop and refine strategies that can effectively mitigate new threats. Collaborating with industry participants and encouraging the sharing of resilience strategies plays an important role in advancing the field and ensuring that organizations can protect themselves against complex cyber-attacks.

VII. Conclusion

In this SoK, we present the current state of the art in the Healthcare, Energy, Transportation, Finance, Supply Chain, and Communication market sectors. Our analysis reveals a significant variation in the maturity and scope of cyber resilience practices across these sectors. We also find a notable research gap in other market sectors, leading to their exclusion from this paper. Furthermore, we present and evaluate the most common tools and technologies used across these sectors to enhance cyber resilience. Our evaluation highlights both the strengths and limitations of these tools, emphasizing the need for continuous improvement and field-testing to ensure their effectiveness. We then illustrate the application and prevalence of these tools across the different sectors, showing which tools are universally used. We identify areas where certain sectors need to improve their strategies, recognizing the finance sector as a prime example for implementing cyber resilience practices. Overall, this study underscores the importance of cyber resilience for protecting critical infrastructure and business operations. It also calls for increased research and development efforts to address the evolving cyber threat landscape and improve resilience across all market sectors.

A. Abbreviations and Acronyms

AMS:	Autonomous Mobility Systems

CAT:	Cybersecurity Assessment Tool
CISA:	Cybersecurity and Infrastructure Security Agency
CRAT:	Cyber Resilience Assessment Tool
CRR:	Cyber Resilience Review
CSAT:	Cyber Security Assessment Tool
CSF:	Cybersecurity Framework
HIPAA:	Health Insurance Portability and Accountability Act
IDS:	Intrusion Detection System
IoT:	Internet of Things
MFA:	Multi-Factor Authentication
NAS:	National Academies of Science
NHS:	National Health Services
NIST:	National Institute of Standards and Technology
PLS-SEM:	Partial Least Squares Structural Equation Modeling
PPP:	Public-Private Partnerships
SCADA:	Supervisory Control and Data Acquisition

References

[1] Ross, R., Pillitteri, V., Dempsey, K., Riddle, M., & Guissanie, G. "Protecting controlled unclassified information in nonfederal systems and organizations" (No. NIST Special Publication (SP) 800-171 Rev. 2 (Draft)). National Institute of Standards and Technology., 2019, doi: 10.6028/NIST.SP.800-171r2.

[2] I. Linkov and A. Kott, "Fundamental Concepts of Cyber Resilience: Introduction and Overview," in Cyber Resilience of Systems and Networks, A. Kott and I. Linkov, Eds., Cham: Springer International Publishing, 2019, pp. 1–25. doi: 10.1007/978-3-319-77492-3_1.

[3] V. Tzavara and S. Vassiliadis, "Tracing the evolution of cyber resilience: a historical and conceptual review," Int. J. Inf. Secur., vol. 23, no. 3, pp. 1695–1719, Jun. 2024, doi: 10.1007/s10207-023-00811-x.

[4] "Transformational government: enabled by technology," GOV.UK. Accessed: Jun. 16, 2024. [Online]. Available: https://www.gov.uk/government/publications/transformational-government-enabled-by-technology

[5] "WannaCry explained: A perfect ransomware storm," CSO Online. Accessed: Jun. 16, 2024. [Online]. Available: https://www.csoonline.com/article/563017/wannacry-explained-a-perfect-ransomware-storm.html

[6] "Cyber-Attack Against Ukrainian Critical Infrastructure — CISA." Accessed: Jun. 16, 2024. [Online]. Available: https://www.cisa.gov/news-events/ics-alerts/ir-alert-h-16-056-01

[7] D. Galinec and W. Steingartner, "Combining cybersecurity and cyber defense to achieve cyber resilience," in 2017 IEEE 14th International Scientific Conference on Informatics, Nov. 2017, pp. 87–93. doi: 10.1109/INFORMATICS.2017.8327227.

[8] M. Danilak, "What's the difference between cyber security and cyber resilience?," Ascentor. Accessed: May 06, 2024. [Online]. Available: https://ascentor.co.uk/cyber-security-resources/the-difference-between-cyber-security-and-cyber-resilience/.

[9] S. M. AlHidaifi, M. R. Asghar, and I. S. Ansari, "A Survey on Cyber Resilience: Key Strategies, Research Challenges, and Future Directions," ACM Comput. Surv., Feb. 2024, doi: 10.1145/3649218.

[10] M. F. Safitra, M. Lubis, and M. Kurniawan, "Cyber Resilience: Research Opportunities," Aug. 2023. doi: 10.1145/3592307.3592323.

[11] P. Williams and R. Manheke, "Small Business - A Cyber Resilience Vulnerability," International Cyber Resilience conference, Aug. 2010, [Online]. Available: https://ro.ecu.edu.au/icr/14

[12] "Cyber Resilience — Security Insider." Accessed: Jun. 21, 2024. [Online]. Available: https://www.microsoft.com/en-us/security/security-insider/intelligence-reports/5-steps-to-cyber-resilience

[13] "User Guide Cyber Resilience." Accessed: Jun. 21, 2024. [Online]. Available: https://www.telekom-mms.com/whitepaper/user-guide-cyber-resilience

[14] B. N, "Financial Sectors Lost $20 Billion Over the Past 20 Years," Cyber Security News. Accessed: May 16, 2024. [Online]. Available: https://cybersecuritynews.com/financial-sectors-lost-20-billion-over-the-past-20-years/

[15] A. Garcia-Perez, J. G. Cegarra-Navarro, M. P. Sallos, E. Martinez-Caro, and A. Chinnaswamy, "Resilience in healthcare systems: Cyber security and digital transformation," Technovation, vol. 121, p. 102583, Mar. 2023, doi: 10.1016/j.technovation.2022.102583.

[16] K. Blanchet, S. L. Nam, B. Ramalingam, and F. Pozo-Martin, "Governance and Capacity to Manage Resilience of Health Systems: Towards a New Conceptual Framework," Int J Health Policy Manag, vol. 6, no. 8, pp. 431–435, Aug. 2017, doi: 10.15171/ijhpm.2017.36.

[17] A. Boddy, W. Hurst, M. Mackay, and A. E. Rhalibi, "A study into data analysis and visualisation to increase the cyber-resilience of healthcare infrastructures," in Proceedings of the 1st International Conference on Internet of Things and Machine Learning, in IML '17. New York, NY, USA: Association for Computing Machinery, Oct. 2017, pp. 1–7. doi: 10.1145/3109761.3109793.

[18] S. Ghafur, E. Grass, N. R. Jennings, and A. Darzi, "The challenges of cybersecurity in health care: the UK National Health Service as a case study," The Lancet Digital Health, vol. 1, no. 1, pp. e10–e12, May 2019, doi: 10.1016/S2589-7500(19)30005-6.

[19] Porter, G., Trevors, M., & Vrtis, R. "Mapping of the Health Insurance Portability and Accountability Act (HIPAA) Security Rule to the Cyber Resilience Review (CRR)", CMU/SEI-2018-TN-001, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, USA, Mar. 2018.

[20] B. Ram, "Best Practices to Increase Cyber Resilience of Smart Electricity Grid with Focus on Advanced Metering and Distribution Infrastructure," presented at the International Conference on Reliability and Quality in Power Supply, 2018.

[21] J. Pöyhönen, V. Nuojua, M. Lehto, and J. Rajamäki, "Application of cyber resilience review to an electricity company," LAUREA, pp. 380–389, June 2018.

[22] "Resilience metrics for cyber systems — Environment Systems and Decisions." Accessed: May 19, 2024. [Online]. Available: https://link.springer.com/article/10.1007/s10669-013-9485-y

[23] "Electric Power Grid Resilience to Cyber Adversaries: State of the Art — IEEE Journals & Magazine — IEEE Xplore." Accessed: May 19, 2024. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9090140

[24] B. Zou, P. Choobchian, and J. Rozenberg, Cyber Resilience of Autonomous Mobility Systems: Cyber Attacks and Resilience-Enhancing Strategies. 2020. doi: 10.1596/1813-9450-9135.

[25] G. Lykou, A. Anagnostopoulou, and D. Gritzalis, "Implementing Cyber-Security Measures in Airports to Improve Cyber-Resilience," in 2018 Global Internet of Things Summit (GIoTS), Jun. 2018, pp. 1–6. doi: 10.1109/GIOTS.2018.8534523.

[26] R. AlMashari, G. AlJurbua, L. AlHoshan, N. S. Al Saud, O. BinSaeed, and N. Nasser, "IoT-based Smart Airport Solution," in 2018 International Conference on Smart Communications and Networking (SmartNets), Nov. 2018, pp. 1–6. doi: 10.1109/SMARTNETS.2018.8707393.

[27] A. R. Mathew, "Airport Cyber Security and Cyber Resilience Controls." arXiv, Aug. 26, 2019. doi: 10.48550/arXiv.1908.09894.

[28] B. N, "Financial Sectors Lost $20 Billion Over the Past 20 Years," Cyber Security News. Accessed: May 16, 2024. [Online]. Available: https://cybersecuritynews.com/financial-sectors-lost-20-billion-over-the-past-20-years/

[29] Statista, "BlackRock - statistics & facts," . Accessed: May 12, 2024. [Online]. Available: https://www.statista.com/topics/8295/blackrock/

[30] B. Dupont, "The cyber-resilience of financial institutions: significance and applicability," Journal of Cybersecurity, vol. 5, no. 1, Jan. 2019, doi: 10.1093/cybsec/tyz013.

[31] J. L. Pinckard, M. Rattigan, and R. A. Vrtis, "A Mapping of the Federal Financial Institutions Examination Council (FFIEC) Cybersecurity Assessment Tool (CAT) to the Cyber Resilience Review (CRR)," Carnegie Mellon University, Software Engineering Institute, Tech. Note CMU/SEI-2016-TN-008, Oct. 2016.

[32] H. Gallagher, W. McMahon, and R. Morrow, "Cyber Security: Protecting the Resilience of Canada's Financial System," Bank of Canada, Financial System Review, Dec. 2014, pp. 47-53. [Online]. Available: https://www.bankofcanada.ca/core-functions/financial-system/oversight-designated-clearing-settlement-systems/

[33] J. C. Crisanto and J. Prenio, "Emerging Prudential Approaches to Enhance Banks' Cyber Resilience," in The Palgrave Handbook of FinTech and Blockchain, M. Pompella and R. Matousek, Eds., Cham: Springer International Publishing, 2021, pp. 285–306. doi: 10.1007/978-3-030-66433-6_13.

[34] G. V. Fedotova, Y. V. Kuznetsov, L. A. Kargina, S. L. Lebedeva, and D. A. Kurazova, “Smart Cyber Resilience Technologies of Credit Organizations,” in “Smart Technologies” for Society, State and Economy, E. G. Popkova and B. S. Sergi, Eds., Cham: Springer International Publishing, 2021, pp. 1649–1658. doi: 10.1007/978-3-030-59126-7_180.

[35] O. Khan and D. A. Sepúlveda Estay, “Supply Chain Cyber-Resilience: Creating an Agenda for Future Research,” Technology Innovation Management Review, no. April, pp. 6–12, 2015.

[36] A. Davis, ”Building Cyber-Resilience into Supply Chains,” Technology Innovation Management Review, vol. 5, no. 4, pp. 19-27, Apr. 2015.

[37] H. Boyes, ”Cybersecurity and Cyber-Resilient Supply Chains,” Technology Innovation Management Review, vol. 5, no. 4, pp. 28-34, Apr. 2015.

[38] D. B. Parker, “Toward a New Framework for Information Security?,” in Computer Security Handbook, John Wiley & Sons, Ltd, 2012, p. 3.1-3.23. doi: 10.1002/9781118851678.ch3.

[39] M. Buinevich and A. Vladyko, “Forecasting Issues of Wireless Communication Networks’ Cyber Resilience for An Intelligent Transportation System: An Overview of Cyber Attacks,” Information, vol. 10, no. 1, Art. no. 1, Jan. 2019, doi: 10.3390/info10010027.

[40] N. I. of S. and Technology, “The NIST Cybersecurity Framework (CSF) 2.0,” U.S. Department of Commerce, NIST CSWP 29, Feb. 2024. doi: 10.6028/NIST.CSWP.29.

[41] “Cyber Resilience Review (CRR) — CISA.” Accessed: Jun. 08, 2024. [Online]. Available: https://www.cisa.gov/resources-tools/services/cyber-resilience-review-crr

[42] “What Is Vulnerability Assessment? Benefits, Tools, and Process — HackerOne.” Accessed: Jun. 08, 2024. [Online]. Available: https://www.hackerone.com/knowledge-center/what-vulnerability-assessment-benefits-tools-and-process

[43] C. C. Editor, “penetration testing - Glossary — CSRC.” Accessed: Jun. 08, 2024. [Online]. Available: https://csrc.nist.gov/glossary/term/penetration_testing

[44] “G7 Fundamental Elements for Threat-LED Penetration Testing,” GOV.UK. Accessed: Jun. 08, 2024. [Online]. Available: https://www.gov.uk/government/publications/g7-fundamental-elements-for-threat-led-penetration-testing

[45] A. Ometov, S. Bezzateev, N. Mäkitalo, S. Andreev, T. Mikkonen, and Y. Koucheryavy, “Multi-Factor Authentication: A Survey,” Cryptography, vol. 2, no. 1, Art. no. 1, Mar. 2018, doi: 10.3390/cryptography2010001.

[46] ”Cyber Security Assessment Tool (CSAT),” Cybersecurityassessmenttool.com. [Online]. Accessed: May 8, 2024. Available: https://cybersecurityassessmenttool.com/.

[47] “What Is a Cybersecurity Assessment Tool (CSAT)?,” ThreatDotMedia — Cyber Explained in Simple Terms. Accessed: Jun. 21, 2024. [Online]. Available: https://threat.media/definition/what-is-a-cybersecurity-assessment-tool-csat/

[48] “Cyber Resilience Assessment.” Accessed: May 8, 2024. [Online]. Available: https://www.dell.com/en-us/dt/data-protection/cyber-resiliency-assessment.htm

[49] F. S. Inc, “What’s the difference between a security risk assessment and a penetration test?,” F1 Solutions. Accessed: Jun. 21, 2024. [Online]. Available: https://www.f1networks.com/blog/whats-difference-security-risk-assessment-penetration-test/

[50] J. W. S. Parker, “What industries benefit the most from penetration testing?,” James Parker. Accessed: Jun. 08, 2024. [Online]. Available: https://www.jamesparker.dev/what-industries-benefit-the-most-from-penetration-testing/

[51] “HIPAA Penetration Testing - The Guide To Staying Compliant.” Accessed: Jun. 08, 2024. [Online]. Available: https://www.blazeinfosec.com/post/hipaa-penetration-testing-guide/

[52] “Penetration Testing in der Logistik,” ProSec GmbH. Accessed: Jun. 21, 2024. [Online]. Available: https://www.prosec-networks.com/pentest-logistik/

[53] “Aviation Cyber Security Testing — Pen Test Partners.” Accessed: Jun. 21, 2024. [Online]. Available: https://www.pentestpartners.com/penetration-testing-services/aviation-cyber-security-testing/

[54] “Elevating Automotive Security: Penetration Testing for Vehicles,” Deloitte Hungary. Accessed: Jun. 21, 2024. [Online]. Available: https://www2.deloitte.com/hu/en/pages/risk/articles/elevating-automotive-security-penetration-testing-for-vehicles.html

[55] “Energy Penetration Testing.” Accessed: Jun. 21, 2024. [Online]. Available: https://www.pentestpeople.com/sectors/energy

[56] New York Department of Financial Services, ”Cybersecurity Requirements for Financial Services Companies (23 NYCRR 500),” 2023. Accessed: June 8, 2024. [Online]. Available: https://www.dfs.ny.gov/system/files/documents/2023/03/23NYCRR500_0.pdf.

[57] O. for C. Rights (OCR), “Summary of the HIPAA Security Rule.” Accessed: Jun. 08, 2024. [Online]. Available: https://www.hhs.gov/hipaa/for-professionals/security/laws-regulations/index.html

[58] “Communications and Cyber Resiliency Toolkit — CISA.” Accessed: Jun. 08, 2024. [Online]. Available: https://www.cisa.gov/resources-tools/resources/communications-and-cyber-resiliency-toolkit

[59] Z. Tao, L. Zong, J. Zhai, and S. Shi, “Global Cyber Policy Developments for Cyber Security and Financial Market Stability: A Quantitative Investigation.” Rochester, NY, Sep. 26, 2022. Accessed: Jun. 19, 2024. [Online]. Available: https://papers.ssrn.com/abstract=4229989