Degree project

UNIVERSITY OF SKÖVDE

Adaptive Network Recovery Techniques for Enhancing Cyber Resilience: A Systematic Literature Review

Bachelor Degree Project in Informatics

First Cycle 30 credits

Spring term 2026

Student: Martina Fletcher Pettersson

Examiner: Thomas Fischer

ABSTRACT

As our modern network environments become increasingly connected, adaptive techniques have emerged to help support faster response and recovery. Cyber resilience is the ability to maintain or restore acceptable functionality during and after cyber incidents. The aim of this study was to identify adaptive network recovery techniques presented in recent research, examine how they may support cyber resilience, and examine how they compare with more traditional recovery methods where reported. This systematic literature review was conducted using IEEE Xplore, ScienceDirect and Web of Science. Publications dated between 1 January 2024 and 2 March 2026 were screened using predefined inclusion, exclusion and quality criteria. The final publication set included 21 publications. Their findings were analysed through a thematic coding process adapted for this systematic literature review. The selected literature presented reinforcement learning and autonomous response, large language model supported incident response and playbook generation, self-healing and resilience oriented approaches, as well as network-level adaptive mechanisms. Reported benefits included shorter response or recovery times, less manual involvement, better decision support and improved service availability within the environments that were evaluated. Three publications provided the clearest comparisons with traditional or static methods for response or recovery related outcomes, using baselines such as a traditional control method, a rule-based response system, a deep reinforcement learning-based system or Snort. These methods were not always directly comparable, as the studies measured different response and recovery outcomes and not all of the baselines provided automatic recovery. Most of the remaining evidence instead focused on detection, containment, response planning, preventive reconfiguration or modelled resilience rather than complete service recovery. The findings therefore suggest that adaptive techniques can support different parts of network response and recovery, but they do not show that adaptive recovery is generally more effective than traditional recovery methods in real operational networks. Further research using more common recovery metrics, better comparisons with traditional methods and realistic production environments is still needed.

LIST OF ABBREVIATIONS

ABBREVIATION	FULL TEXT
AI	Artificial Intelligence
ARCS	Adaptive Reinforcement Learning for Cybersecurity Strategy
DRL	Deep Reinforcement Learning
IoT	Internet of Things
LLM	Large Language Model
MTD	Moving Target Defence
PICOC	Population, Intervention, Comparison, Outcome and Context
SDN	Software Defined Networking

Contents

1 Introduction 1
1.1 Aims 2
2 Background 3
2.1 Concepts and Techniques 3
2.1.1 Cyber Resilience 3
2.1.2 Traditional Network Recovery 3
2.1.3 Adaptive Network Recovery 5
2.1.4 Distributed Systems and Network Vulnerabilities 6
3 Problem Formulation 8
3.1 Study Aims 8
3.2 Motivation 8
3.3 Delimitations 9
4 Methodology 10
4.1 Systematic Literature Review 10
4.1.1 Review Approach 11
4.1.2 PICOC and Search Concepts 12
4.1.3 Database Selection and Search Strategy 12
4.1.4 Inclusion and Exclusion Criteria 14
4.1.5 Quality Assessment 14
4.1.6 Data Extraction Form 15
4.1.7 Study Selection 15
4.2 Thematic Analysis 17
4.3 Validity and Reliability 18
5 Results 20
5.1 Overview of Selected Studies 20
5.2 Thematic Findings 21
5.2.1 Adaptive Response and Recovery Mechanisms 21
5.2.2 Reported Resilience Functions and Outcomes 24
5.2.3 Evaluation Methods and Supporting Evidence 26
5.2.4 Operational Limitations and Safeguards 29
5.3 Summary of Key Findings 31
6 Discussion 33
6.1 Ethical and Societal Aspects 36
7 Conclusion 38
7.1 Future Work 39
References 41

Appendix A List of Included Literature
Appendix B Theme and Subtheme Mapping
Appendix C Thematic Coding Sample
Appendix D Quality Assessment Results

1 Introduction

Cyber security incidents can cause issues with availability and reliability in modern digital systems (Ross et al., 2021). As organisations connect and become reliant on connected digital infrastructure, the ability to recover swiftly after a cyber attack has become an important part of cyber security. More traditionally, cyber security has often revolved around the prevention of attacks along with being able to quickly identify intrusions, and whilst this is still an important factor, cyber resilience also concerns how systems are able to withstand, recover from and adapt after disruptions (Ross et al., 2021). The term cyber resilience does not just describe the process of protecting but also the ability to maintain or restore important functions during and after a disruption (Ross et al., 2021).

Network recovery is a central part of cyber resilience and this term is used for the processes and techniques used to recover, maintain or restore network functionality after a cyber attack or other disruption. Traditional recovery methods are often reliant on predefined procedures, static playbooks, human intervention, backups, system rollbacks or other rule-based actions (Bartock et al., 2016; Nelson et al., 2025; Schlette et al., 2024). These methods can work reasonably well in known and predictable situations, but they may require human judgement when an incident does not match a predefined procedure (Bartock et al., 2016; Nelson et al., 2025). Recovery processes can also involve manual decisions, especially when analysts need to decide which systems should be restored first, whether services should be disconnected or when a system is safe to return to operation (Bartock et al., 2016; Nelson et al., 2025).

Newer approaches have also been developed around how networks can respond and recover in changing threat environments. Adaptive network recovery refers to recovery methods that can adjust according to changes in the system, network conditions or previously learned information. These techniques can include reinforcement learning, autonomic or self-healing systems, software defined networking (SDN), moving target defence (MTD) and large language model (LLM) supported cyber security (Arulkumaran et al., 2017; Kephart & Chess, 2003; Kreutz et al., 2015; Cho et al., 2020; Hasanov et al., 2024). These methods are relevant to cyber resilience because they provide ways for response or recovery actions to be adjusted rather than relying only on actions that were decided beforehand.

At the same time, adaptive recovery is not free from issues. AI-supported techniques may require suitable training, testing and continuous validation before they can be relied upon. They can also introduce security risks, particularly when AI-generated information is inaccurate or when systems are affected by attacks such as prompt injection or attempts to bypass safeguards (Autio et al., 2024; Hasanov et al., 2024). These issues are especially important when automated systems are used to make decisions that could affect important services or sensitive data. Human oversight, allocation of responsibility and trust are therefore still important aspects when considering adaptive network recovery.

Throughout this study, the term “network” is used as a broad term which is also consistent with its use in the reviewed literature. It refers to organisational and technical network environments where digital systems, applications and infrastructure are interconnected and need to be able to keep or recover functionality after a disruption. This includes enterprise

networks, cloud networks, IoT environments and other cyber related systems discussed in the literature. This study does not focus on national-level Internet infrastructure, telecommunication politics or global Internet routing systems. Instead, focus is placed specifically on recovery and response techniques that can support cyber resilience in network-based systems.

Even if research around adaptive recovery techniques has been increasing over the last few years, the literature is still spread over multiple different techniques and uses different methods to evaluate them. Some studies focus more on reinforcement learning, whilst others use LLMs, self-healing systems or dynamic network configuration techniques. This variation creates a challenge for this review because findings from the different techniques and evaluation methods cannot always be compared consistently. Therefore, this study uses a systematic literature review to identify, organise and examine recent research about adaptive techniques for network recovery along with its relation to cyber resilience. The limitations caused by the variation in terminology, techniques, and evaluation methods are all discussed further in Section 4.3.

1.1 Aims

The aim of this study is to identify adaptive network recovery techniques presented in recent research and analyse how they may support cyber resilience. Comparisons with traditional or static recovery methods are also examined where reported in the included publications. This study also examines how techniques were evaluated and discusses their reported benefits and limitations.

2 Background

This section of the thesis explains various concepts and techniques that are relevant to the research topic in order for the reader to better understand the technical background. The topics in this section are divided into three main concepts; cyber resilience as a whole, traditional network recovery as well as adaptive network recovery which is explained in more detail due to the importance of the subject in the research topic. There is also a section on distributed systems and network vulnerabilities.

2.1 Concepts and Techniques

2.1.1 Cyber Resilience

Cyber resilience refers to the ability of a system or organisation to anticipate, withstand, recover from and adapt to adverse conditions, attacks or compromises. It is broader than cyber security based mainly on prevention because it assumes that disruption may still occur and focuses on maintaining or restoring important functions afterwards (Ross et al., 2021). A broad view is important within cyber resilience because modern infrastructure is exposed to not only isolated incidents, but also to conditions involving multiple variables that must be managed accordingly in order to maintain systems with near perfect uptime, security and reliability.

A key part of cyber resilience is not only whether an attack can be prevented, but also how well a system is able to keep functioning and recover after a disruption. Recovery is therefore an important part of cyber resilience, as systems need to be able to restore important functions and return to an acceptable level of operation after an incident (Ross et al., 2021). This means that recovery can be seen as one part of the wider resilience process rather than something completely separate that only happens after an attack.

Cyber resilience is therefore closely connected to network recovery, as recovery affects how well and how quickly a system can return to normal operation after a disruption. Organisations may need to decide which systems should be restored first, how the impact of an incident can be limited and how important services can be kept available during recovery (Bartock et al., 2016; Nelson et al., 2025). Adaptive techniques are relevant here because they can change their response depending on the current state of the system or network instead of only following actions that were decided beforehand. Reinforcement learning, SDN and MTD are examples of techniques that can support this type of adaptation (Arulkumaran et al., 2017; Kreutz et al., 2015; Cho et al., 2020).

2.1.2 Traditional Network Recovery

Traditional network recovery often includes more manual response mechanisms. Recovery actions are usually based on specific rules, predefined procedures, static playbooks, backups, system rollbacks or other rule-based actions. However, traditional recovery does not always mean that every action has to be done manually. Scripts and security tools can still automate some parts of the process, but the rules, conditions and approvals are usually decided

beforehand rather than learned and changed by the system itself. Incident response covers the actions used to detect, respond to and recover from an incident, while recovery is more specifically focused on restoring affected systems and services (Nelson et al., 2025). These methods can work reasonably well on known and predictable threats, but they could struggle more in the face of complex, live or previously unseen attacks.

In this context, a response playbook is a predefined sequence of actions that can be followed when handling a particular type of incident. A playbook can include the actions that should be taken, the order in which they should be carried out, who is responsible for them and where further investigation or approval may be needed (Schlette et al., 2024). These approaches can make recovery more repeatable, but they may be less effective when an incident does not match the conditions that the playbook or rollback procedure was designed for. Playbooks may also need to be adjusted to fit the organisation and its systems instead of being used the same way in a different environment (Schlette et al., 2024).

Traditional recovery can also involve containing or isolating affected systems before recovery can properly begin. Containment is mainly used to stop an incident from spreading while it is still being investigated. This can include blocking network traffic, isolating an affected device, quarantining malicious files or temporarily shutting down a service. Eradication then focuses more on removing what caused the compromise, such as malware, breached accounts or exploited vulnerabilities (Nelson et al., 2025). These actions do not mean that the system has fully recovered yet, but they can help limit the damage and make the later recovery process safer.

Another major weakness in traditional recovery is the high dependence on manual intervention. Many methods of recovering from an attack with traditional techniques mean that a human has to get involved, possibly by triggering a restart, shutting down affected servers, approving recovery actions or restoring important data from backups. Human involvement can also be important when deciding which systems should be restored first, approving actions that could affect important services and making sure that a restored system is working properly before it is put back into normal use. Recovery plans therefore also need to define responsibilities, priorities and the actions that should be taken during and after an incident (Bartock et al., 2016; Nelson et al., 2025).

Backups are also an important part of traditional recovery because they can be used to restore affected data or systems. A rollback instead returns a system to an earlier working state, while rebuilding means creating a clean version of the affected system when the original can no longer be trusted (Koo & Toueg, 1987; Nelson et al., 2025).

Traditional recovery can also involve having to balance security with keeping important services available. Disconnecting an affected system may help stop an attack from spreading, but it can also cause downtime or interrupt a service that is still needed. Restoring a system too quickly could also bring it back before the cause of the compromise has been properly dealt with. This means that recovery still involves deciding when a system is ready to be restored and whether a backup, rollback or complete rebuild is the better option (Bartock et al., 2016; Nelson et al., 2025).

Traditional recovery methods therefore present an important conflict. Their predefined procedures, backups and manual controls remain necessary for providing reliable and controlled recovery, especially when responding to known incidents. This same dependence on predefined actions and manual intervention can, however, make them too slow for complex or previously unknown incidents. These limitations help to explain why adaptive recovery has become more relevant, but they do not mean that adaptive methods are automatically better.

2.1.3 Adaptive Network Recovery

Adaptive network recovery can be seen as a response to the limitations of traditional methods in modern cyber security environments. One potential benefit of adaptive recovery is that response actions can be adjusted while an incident is still being handled. Instead of only restoring systems after an incident has already caused disruption, adaptive recovery can support changes to the system or network based on its current condition. Autonomic computing and self-healing systems are based on systems being able to manage and adjust parts of their own behaviour, while techniques such as software-defined networking and moving target defence can allow network configurations to be changed more dynamically (Kephart & Chess, 2003; Kreutz et al., 2015; Cho et al., 2020). This helps to support the idea that adaptive recovery is not just about eventually restoring service, but also about being able to respond and adjust while the disruption is still taking place.

This method of network recovery is also more aware of context than more traditional approaches. An indicator of compromise is a technical artifact or observable that suggests that an attack may be underway or that a compromise may already have occurred (Nelson et al., 2025). Adaptive or AI-supported approaches can use information about the current state of a system or network when deciding how to respond. Reinforcement learning, for example, allows an agent to learn which actions to take based on its interaction with an environment and the results of earlier decisions (Arulkumaran et al., 2017). This differs from more traditional static methods, where actions are more likely to be decided beforehand. Adaptive techniques such as reinforcement learning, LLM-supported response, self-healing systems and SDN-based network reconfiguration are therefore relevant because they can support response or recovery actions that change according to the current system state or previously learned information (Arulkumaran et al., 2017; Hasanov et al., 2024; Kephart & Chess, 2003; Kreutz et al., 2015).

Reinforcement learning is a type of machine learning where an agent learns which actions to take by interacting with an environment and receiving rewards based on its decisions. In adaptive network recovery, the environment can represent the network, the state can represent its current condition and the actions can represent different defensive measures (Arulkumaran et al., 2017). LLMs are models trained on large collections of text to process and generate language. In cyber security, they have been investigated for tasks such as analysing security information, supporting threat detection and assisting with incident response (Hasanov et al., 2024).

Another important part of adaptive network recovery is that adaptation can happen at different levels. Some methods focus on decision-making, where an agent chooses the most suitable response action. Other methods focus on automating incident response workflows, such as generating playbooks or mitigation steps. SDN separates the part of the network that makes decisions from the devices that actually forward traffic. This allows network behaviour to be changed through software and controlled more dynamically (Kreutz et al., 2015). MTD is based on changing parts of a system or network over time in order to make it less predictable to an attacker. Address hopping is one example of this, where network addresses are changed dynamically (Cho et al., 2020). Self-healing systems represent another direction and are designed to detect faults or attacks and restore functionality with less human intervention. This is connected to the idea of autonomic computing, where systems are designed to manage parts of their own behaviour through functions such as self-configuration, self-protection and self-healing (Kephart & Chess, 2003).

Adaptive and AI-supported techniques also introduce their own limitations. LLM-based systems can produce inaccurate information and can be affected by risks such as prompt injection or attempts to bypass safeguards (Autio et al., 2024; Hasanov et al., 2024). These issues are especially important in cyber security because an incorrect automated response could affect the availability or security of a system. Earlier research on machine learning for network security has also shown that results from controlled environments may not always transfer well into real operational networks (Sommer & Paxson, 2010). This means that adaptive recovery still needs suitable validation and human control before more serious actions are carried out.

Human involvement can also still be important when deciding how much of the recovery process should be automated. Parasuraman et al. (2000) explain that automation can be used at different levels instead of everything being either completely manual or completely automatic. This means that the amount of automation used can depend on what kind of action is being carried out and what could happen if the system makes the wrong decision.

Therefore, adaptive network recovery differentiates from a more traditional cyber security perspective and shows why researching adaptive network recovery is relevant for the future of network security and reliability. While conventional cyber security primarily focuses on preventing, detecting and responding to attacks, adaptive network recovery also emphasises the ability of systems to continue functioning during disruption and recover effectively afterwards. This makes adaptive recovery especially relevant in modern network environments, as expecting to prevent every single attack is unrealistic and recovery capabilities are therefore important for maintaining availability, limiting disruption and supporting cyber resilience.

2.1.4 Distributed Systems and Network Vulnerabilities

Modern network environments are often distributed across multiple systems and locations. Distributed systems can consist of large numbers of networked devices or nodes that work together toward a common goal (Ahmed & Boutaba, 2011). This can make recovery more complicated, as a cyber attack may affect more than one isolated system and can also spread to other connected parts of the infrastructure. In distributed environments, one single

interference in a service can affect many other systems and services that are dependent on the same communication routes, control systems or digital services. This means that recovery does not only mean being able to get systems back to their original state and therefore recovering one singular service, but also being able to understand how different parts of a network affect and rely on each other.

This is especially clear in critical infrastructure and cyber-physical systems. A cyber-physical system combines digital systems and communication with physical processes, where the digital parts can monitor or control physical operations and changes in the physical system can also affect the digital side (Lee, 2008). Rinaldi et al. (2001) describe how critical infrastructures can be highly interconnected and dependent on each other, meaning that disruption in one part can also affect other connected systems and services. This is important from a recovery perspective, as problems in one part of an infrastructure may therefore have wider effects on the systems that depend on it. This shows that recovery in distributed systems must take into account both the various technical components and the wider functionality of the systems.

As mentioned above, distributed systems can also create a larger attack surface. When multiple devices, servers, services and communication routes are all connected, there can be more potential points that need to be protected. An attack surface describes the parts of a system through which it may be exposed to attacks, and reducing this attack surface can help to reduce security risk (Manadhata & Wing, 2011). The dependencies between different systems can also mean that attacks or disruptions against one part may affect other connected parts of the infrastructure (Rinaldi et al., 2001). This can make recovery more complicated, as predefined rules or static recovery routines may not always account for every dependency or condition that exists within a larger interconnected system.

Adaptive network-level techniques can be relevant in environments such as these. SDN allows network behaviour to be controlled and changed through software, while MTD is based on changing parts of the network over time in order to make it less predictable to an attacker (Kreutz et al., 2015; Cho et al., 2020). Altogether, distributed systems make network recovery a more complex task as recovery decisions need to take into account different dependencies, service continuity, attack surfaces and ever-changing network circumstances. This reinforces the connection between network recovery and cyber resilience, as more resilient systems do not just need to be able to resist attacks, but also limit the impact, recover functionality and adjust according to new threats.

3 Problem Formulation

Recovery from cyber incidents is an important part of cyber resilience, as modern organisations and digital infrastructures are increasingly dependent on connected systems. This creates a need to examine whether adaptive techniques may support stronger recovery and resilience. Recent publications have examined adaptive recovery in several technical areas and have used different methods, metrics and settings to evaluate it. However, the research is spread across these technical areas, which makes it difficult to form a clear picture of how adaptive network recovery techniques may contribute to cyber resilience, what benefits they may provide and how they compare with traditional recovery methods. There is therefore a need for a systematic review in order to bring together, organise, and examine the available research, while also identifying its limitations and areas where further research is needed.

3.1 Study Aims

The aim of this study is to identify and examine adaptive network recovery techniques presented in recent research and analyse how they may support cyber resilience.

The study aims to answer the following research question:

> *What adaptive network recovery techniques are presented in recent research, and how may they contribute to cyber resilience?*

Where reported in the selected publications, comparisons with traditional recovery techniques and methods are also examined.

3.2 Motivation

The importance of recovering from cyber attacks in modern networked systems motivated this study of cyber resilience. Organisations depend on networked systems to support their operations, making fast and effective recovery important for limiting disruption and supporting resilience. Traditional recovery methods often rely on manual intervention, which can make them less suitable for complex network environments.

Newer adaptive methods aim to improve recovery through learning-based, context-driven and partly autonomous techniques. However, the research is still new and often focuses on specific technical solutions rather than wider comparisons. By reviewing current publications, this study aims to provide a clearer understanding of how adaptive recovery techniques are developing, what benefits and limitations they have, and how they compare with traditional recovery methods. The study contributes academically by bringing together a growing research area and practically by highlighting methods that may support cyber resilience in real network environments.

3.3 Delimitations

This study is limited to research concerning adaptive network recovery techniques and traditional or static recovery methods when they are mentioned or used as baselines in the included publications. It focuses specifically on recent developments within this area and therefore does not aim to go into historical depths regarding cyber security or recovery research. Instead, the purpose is to examine newer approaches and how these developments could influence and change the future of cyber security practice.

This study also primarily focuses on the recovery and resilience related aspects of cyber security. It does not examine prevention or detection in depth unless they are directly connected to recovery processes. Similarly, the study does not aim to evaluate cyber security as a whole, but rather focuses on how systems and networks can adapt, recover, and maintain resilience following disruptions or attacks.

4 Methodology

This section provides an explanation of the methodology that was used in order to conduct this study. The methodology in this study was chosen by considering what kind of information was needed in order to be able to answer the research question. Since the aim was to bring together and examine findings from existing publications across a range of techniques and settings, a method that could provide both a clear and structured process for searching, selecting and analysing earlier research was needed. A systematic literature review was therefore considered to be the most suitable methodology, instead of a study that focused on collecting new data through an experiment, interview or case study. Overall, this section explains this choice in more detail and describes the review process, including search strategy, database selection, inclusion and exclusion criteria, quality assessment, and the study selection process as a whole. This is followed by an explanation of how thematic analysis was used to code and analyse the findings. Finally, this section discusses the validity and reliability of the study, any methodological limitations, and ethical aspects.

4.1 Systematic Literature Review

A systematic literature review was considered appropriate for this study because the aim was to identify, organise and examine the available findings across the field. However, adaptive network recovery is still a developing field, and the differences in the terminology used could have affected the search terms and caused potentially relevant publications to have been missed. This validity threat is discussed further in Section 4.3.

The method presented by Carrera-Rivera et al. (2022) was used to guide the planning and execution of the study as it provides a process designed specifically for computer science research. Their method centres around two main phases: planning and conducting the review. Their guide then explains how the extracted evidence should be analysed and reported in research. The method was adapted in this study to suit the research question, limited publication period, and the qualitative thematic analysis. Extracted findings were analysed using the thematic coding procedure proposed by Naeem et al. (2023).

A few other methodological approaches have been considered that could have been conducted instead of a systematic literature review. One such approach could have been an experiment that provided comparable measurements by evaluating both adaptive and traditional techniques under the same conditions in an attack setting (Wohlin et al., 2012). This would have required access to several different recovery systems, creating multiple attack scenarios, and defining evaluation metrics for each. However, the amount of time, resources, and scale of infrastructure needed in order to conduct such an experiment would have made it difficult to execute. A smaller amount of techniques would be more viable, but may still not have been enough to answer the research question in full.

Another approach could have been a case study able to provide evidence about adaptive recovery techniques in a live environment. Case studies are especially suitable for investigating technical developments in real scenarios and contexts (Runeson & Höst, 2009).

The downside to this approach is that it would have required access to organisations currently using adaptive recovery systems, and also access to potentially sensitive incident and recovery data which may be seldom granted. It would have also only given results surrounding the few organisations chosen for such a study, and therefore not give the full picture in every context.

A third considered approach could have been conducting a systematic mapping study as described by Petersen et al. (2015). They explain that these types of studies can categorise publications and provide an overview of the structure of a research field. However, because systematic mapping studies are mainly used for organising and categorising research, it would have been less suitable for analysing the reported evidence in enough detail to make the comparison required by the research question.

For these reasons, a systematic literature review was considered to be the most suitable method for this study. The procedure based around computer science and presented by Carrera-Rivera et al. (2022) was chosen to guide the review process. As the included publications had differing research designs and evaluation methods, thematic analysis was used to group together the recurring techniques, comparisons and limitations. This review cannot verify the truth or effectiveness made by the publications on its own, but it can give insight into the available evidence.

4.1.1 Review Approach

The approach described by Carrera-Rivera et al. (2022) starts with a planning phase where the research question, search strategy, inclusion and exclusion criteria, quality assessment, and data extraction process are defined. The Population, Intervention, Comparison, Outcome and Context (PICOC) framework can be used during this phase to help structure the research question and identify the main concepts that need to be represented in the search strategy. Carrera-Rivera et al. (2022) also explain that these decisions should be documented in a review protocol to improve the replicability of the review.

The next phase concerns conducting the systematic literature review. This includes applying the search strategy to the selected databases, collecting the identified publications, removing duplicates and screening the remaining publications according to the predefined criteria. Publications that meet the inclusion criteria are then examined in more detail and any relevant information is recorded through a data extraction process.

In the last phase, extracted information is analysed in order to answer the research question. Carrera-Rivera et al. (2022) explain that the analysis can include both quantitative descriptions and qualitative interpretation, and that the selected studies can be grouped according to relevant themes, characteristics or subthemes. The study selection process can also be presented visually through a flow diagram, such as the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) diagram, in order to make the selection process easier to follow (Page et al., 2021).

4.1.2 PICOC and Search Concepts

The research question that was presented in Section 3 was the basis for the search strategy, study selection and data extraction. In order to make the connection between the research question and the search topics clearer, the question was organised using the PICOC framework outlined by Carrera-Rivera et al. (2022). Population was defined through the inclusion and exclusion criteria. Intervention, outcome and context became the basis for the Boolean search terms, and comparison terms were not included as publications concerning adaptive techniques do not typically use such terms in their titles or abstracts. Instead, comparisons were identified during the data extraction process.

Table 1

PICOC Criteria

	Description	Example
Population	Systems being studied	Computer networks and infrastructure systems
Intervention	Techniques being investigated	Adaptive, autonomous or AI-supported recovery and incident response
Comparison	Alternative approaches	Traditional, manual, static or rule-based recovery methods, if reported
Outcome	Results being examined	Response and recovery time, restoration, availability, resilience, effectiveness and limitations
Context	Settings	Cyber incidents, network recovery and cyber defence

4.1.3 Database Selection and Search Strategy

Potential publications were identified through searches in three academic databases: IEEE Xplore, ScienceDirect and Web of Science. IEEE Xplore was selected in order to identify technical research in computing and engineering, while ScienceDirect was used to identify journal publications in computer science and related areas. Web of Science was used to provide research from a wider range of related areas. The three databases were also available through the university library, which supported access to the full texts needed during screening and analysis. Together, the three databases were intended to include technical research on the topic, as well as research from related areas. The search was limited to these databases to keep the screening process manageable within the timeframe and scope of the study. This means that relevant publications that were indexed in other databases could have been missed, which is a limitation discussed in Section 4.3.

The chosen search terms were based on the research question and PICOC framework and were selected based on their relevance to adaptive network recovery and cyber resilience. Terms were selected based on their relevance to adaptive network recovery and cyber resilience, rather than how frequently individual words appeared in each publication. Terms such as “adaptive cyber security”, “autonomous cyber defence”, “artificial intelligence”, “machine learning”, “incident response”, “cyber resilience” and “self-healing systems” were used to represent the main concepts in the research question. These terms, along with the research question and the PICOC framework, are what formed the basis of the Boolean search query presented below:

```
("adaptive cybersecurity" OR "adaptive cyber defense" OR
"autonomous cyber defense")

AND

("artificial intelligence" OR "machine learning" OR "AI")

AND

("incident response" OR "cyber resilience" OR "self-healing")
```

The search terms were divided into three groups that were intended to represent the main concepts in the research question. The first group covered adaptive or autonomous cyber security, the second covered artificial intelligence and machine learning, and the third covered incident response, cyber resilience and self-healing. Terms within each group were connected using OR, because the terms represented alternative ways of describing the same concept, while all three groups were connected using AND because a publication needed to relate to adaptive cyber security, AI-supported techniques and response, recovery, or resilience in order to be relevant to the research question. This meant that each result had to include at least one term from every group.

The search query was intentionally kept very focused in order to avoid retrieving a very large number of publications that were related to cyber security or artificial intelligence in general but had little connection to adaptive recovery. This meant that the query relied on a relatively small number of search terms in a wide and still developing research area. Some relevant publications that made use of different terminology may therefore have been missed. This is a clear limitation of the search strategy and it is discussed further in Section 4.3.

Shorter test searches were also used to check whether individual search terms from the full Boolean query were being processed by each database. The term “self-healing” was checked specifically because it appeared at the end of the query. Each database returned at least one result with this term, which indicated that the final part of the query had indeed been included. As these shorter searches tested individual terms instead of the complete Boolean query, their results were not part of the study selection process. Any differences in the number of results between the databases does not necessarily mean that they processed the query differently, as they could also have been caused by differences in database context, indexing or search functions. These checks were only used to examine how the query was

processed and did not confirm that the selected terms covered all of the terminology used across the field. This limitation is further discussed in Section 4.3.

4.1.4 Inclusion and Exclusion Criteria

The inclusion criteria included peer-reviewed journal articles or conference publications written in English and published between 1 January 2024 and the final database search date of 2 March 2026. This time period was selected as a delimitation because the study aimed to examine the most recent developments in adaptive and AI-supported network recovery. It was not intended to suggest that earlier research is irrelevant. The limitations caused by this narrow publication window are discussed further in Section 4.3. Publications were also required to show a clear connection to adaptive network recovery, response or cyber resilience through evaluation results or relevant theoretical or technical findings.

Certain publications were included even if they could not contribute fully to a direct comparison between adaptive and traditional recovery methods. They were still considered relevant because they provided theoretical or contextual insights into cyber resilience, automated incident response, self-healing systems, SDN-based adaptation, or the limitations and risks of cyber defence. These publications were still identified through the same Boolean search and passed the screening criteria, even if their main contribution was not a direct comparison of AI-supported and traditional recovery methods. Because of the narrow publication period, this review does not provide a proper overview of earlier traditional recovery methods. Traditional approaches are instead considered only when they are discussed or used for comparison in the recent publications identified through the SLR.

Publications were excluded if they were not peer-reviewed, published outside of the time period mentioned above, or not written in English. Studies were also excluded if they did not specifically focus on, or have a clear connection to topics such as network recovery, response, resilience or adaptive network defence. This included publications concerned with general cybersecurity awareness or training, blockchain security, encryption methods, or threat detection without relevant recovery methods.

4.1.5 Quality Assessment

As the final part of the screening process, the remaining 22 publications were examined using a quality assessment checklist with five questions that was developed with guidance and example scales by Carrera-Rivera et al. (2022). The checklist was defined to evaluate how clear the research was to its study purpose, the description of the research method, supporting evidence and conclusions, as well as how any limitations were acknowledged. Table 2 shows the questions used to evaluate each publication.

Table 2

Quality Assessment Evaluation Questions

	Quality Assessment Evaluation Questions
1.	Is the study purpose or research objective clearly stated?
2.	Is the research method or approach adequately described?
3.	Is the evaluation or supporting evidence appropriate for the study's type and claims?
4.	Are the conclusions supported by the results or evidence presented?
5.	Does the publication acknowledge relevant limitations or threats to validity?

Each assessment question was scored using a point system which could result in a maximum of five points. The minimum score set for inclusion was three points. A score of Yes (1) was given when the criterion was clearly met, Partially (0.5) when it was only partly met or important information was missing, and No/unclear (0) when the question criteria were not met or could not be determined. The evidence was assessed based on the type of contribution each publication made. Out of the 22 assessed publications, 21 met the minimum score for inclusion and one was excluded. The full assessment is available in Appendix D.

4.1.6 Data Extraction Form

The recommendations of Carrera-Rivera et al. (2022) specify that a data extraction form should be used to record information needed in order to answer the research question. The recorded fields included the publication names, authors and publication years, findings, relevancy to the research question, codes, subthemes and final themes. Each finding was recorded separately from others and linked to its original location in each publication to better maintain traceability during the process.

4.1.7 Study Selection

The database searches resulted in 71 records from ScienceDirect, 14 from IEEE Xplore and 43 from Web of Science, giving a total of 128 records. After the results from the three databases were combined and 19 duplicates were removed, 109 unique publications remained for title and abstract screening.

In order to identify and remove the duplicates, the database results were combined and compared through their publication titles and authors. These records were checked against the inclusion and exclusion criteria, first by their respective titles as well as a brief look at their abstracts where needed. Publications that were excluded in this section were because they were too focused on general cyber security, education, regulation, encryption or threat detection without a clear connection to recovery, response, cyber resilience or adaptive network defence.

Publications that were deemed to be about adaptive network defence, cyber resilience, automated incident response, self-healing systems, reinforcement learning, LLM supported response techniques, software defined network-based adaptivity or other related methods relevant to this study and for future evaluation were kept for closer review. This review of both titles and abstracts excluded 75 records and therefore reduced the total number of relevant publications from 109 to 34.

The full texts of the remaining 34 publications were then retrieved so that they could be reviewed more carefully by reading through their abstracts, methods, results and conclusions. One potentially relevant publication could not be accessed, reducing the number to 33. Since the analysis required the full text of each included publication, the inaccessible publication could not be included in the final selection.

Out of the remaining 33 publications, 11 were excluded because they did not have a close enough connection to network recovery or adaptive techniques, were not relevant to the recovery and response process after a cyber incident or did not provide enough technical or conceptual relevance to the research question. This left 22 publications for the final quality assessment stage of screening. However, one publication did not meet the minimum quality threshold as explained in Section 4.1.5, resulting in a final selection of 21 publications.

The complete selection process is shown in a PRISMA diagram in Figure 1, and the list of included publications is presented in Appendix A.

Figure 1

PRISMA 2020 flow diagram of the literature selection process (adapted from Page et al., 2021)

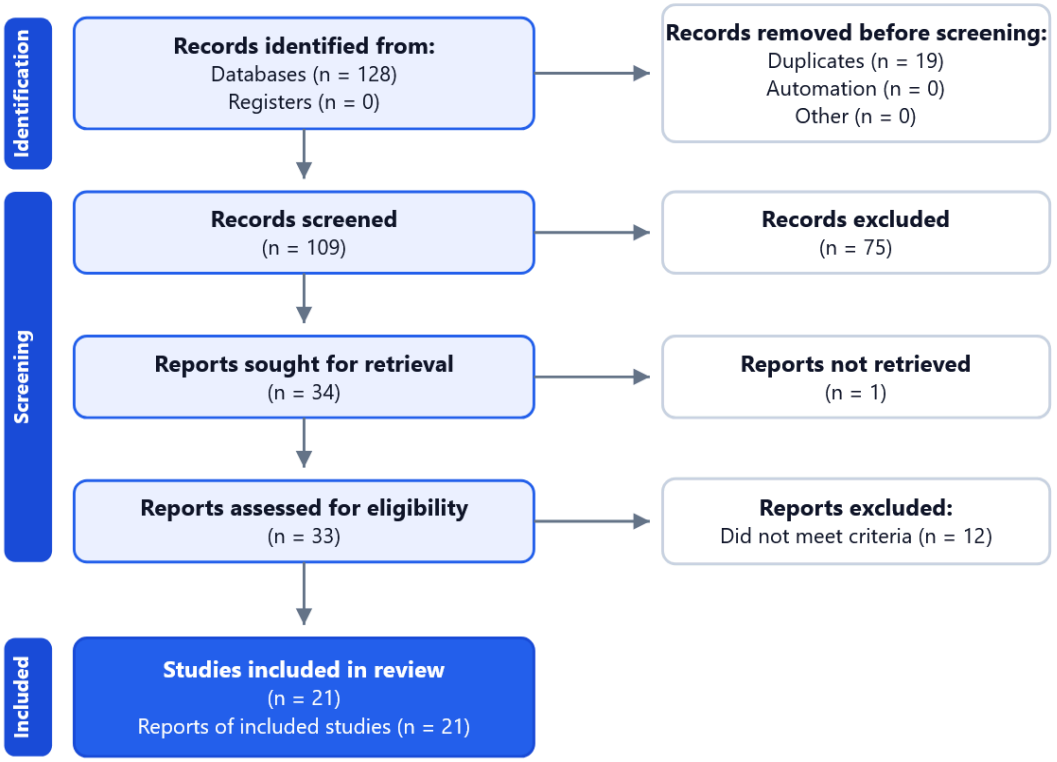

4.2 Thematic Analysis

Thematic analysis was used to analyse the data that was extracted from the 21 publications listed above. This is a method for identifying and interpreting recurring patterns within qualitative data. Naeem et al. (2023) suggest a systematic process of six steps that connects quotations and keywords to codes, themes, interpretations, and a conceptual model. The six steps are described below.

1. *Transcription, Familiarisation With the Data, and Selection of Quotations*: The first step involves becoming familiar with the data and selecting information that is relevant to the research objectives. As this study only analysed publications rather than interviews, transcription was not required. Relevant passages and findings concerning adaptive recovery techniques, reported outcomes, evaluation methods, comparisons, and limitations were chosen at this stage.

2. *Selection of Keywords:* In the second step, important or recurring words and expressions are identified within the selected data. In this study, keyword selection was integrated with the coding stage. Important terms and expressions within each finding became an initial code but were not kept at this level.

3. *Coding:* In this step, short words or phrases, which are known as codes, are assigned to relevant parts of the data. Codes summarise the central meaning of a finding and transform the extracted data into something more manageable for analysis according to Naeem et al. (2023). Examples of codes used in this study include “automated restoration”, “human approval required”, “traditional baseline comparison” and “faster automated response”.

4. *Theme Development:* The fourth step compares and groups related codes in order to identify broader patterns within the data. Codes that describe similar functions, techniques, benefits, or limitations were then grouped into subthemes. These subthemes were then combined to create main themes that represent the most common patterns identified across the 21 publications.

5. *Conceptualisation Through Interpretation of Keywords, Codes, and Themes:* The fifth step involves interpreting the connections between keywords, codes, and themes to get a deeper understanding of the findings. In this study, themes were interpreted with the research question in mind in order to explain which adaptive recovery techniques were proposed, how they were evaluated, and how they compared with traditional recovery methods.

6. *Development of a Conceptual Model:* The final step in the process is based upon combining all the identified concepts into a conceptual model that represents the findings. This step was not applied because the purpose of this review was to identify and compare patterns within existing research rather than develop a new model. The resulting themes and subthemes were used to answer the research question.

This thematic coding process by Naeem et al. (2023) was meant for qualitative research and has been adapted in this study to suit a systematic literature review. It is important to note that the research question guided which findings were deemed relevant, while the extracted findings were organised into codes, subthemes, and main themes through the steps described above. The full analysis resulted in four themes and 20 subthemes. Section 5.2 presents the findings according to the four main themes produced through the thematic analysis. Within these themes, the coded findings were summarised by comparing the adaptive techniques, evaluation methods, reported outcomes and limitations across the included publications. Appendix B presents the theme and subtheme mapping, while Appendix C presents a sample from the thematic codebook.

In order to assess how the identified adaptive techniques may contribute to cyber resilience, findings related to resilience were identified during data extraction and thematic analysis. The findings were then grouped into subthemes related to resilience, allowing the publications to be compared according to the type of contribution they reported rather than only by whether they used the term “cyber resilience”. The resulting resilience-related subthemes and the publications contributing to each of them are presented in Appendix B.

4.3 Validity and Reliability

The validity and reliability of this study depend on how consistently and transparently the review process was conducted. The chosen criteria were considered to be appropriate given the current fast-moving landscape of adaptive network techniques due to the increasing popularity of AI. The final publications were also chosen for their relevance to the research question, helping to ensure that the studies would contribute meaningfully to this thesis.

Reliability is supported by having documented the search query used to generate the list of initial studies before screening and excluding publications. The databases used, number of search results generated and the selection process were documented in full. The same extraction fields were used for each publication before the findings were compared and thematically coded. Reliability was further strengthened by checking the original Boolean syntax query through a few shorter searches, using the term “self-healing” as it was at the very end of the query.

The shorter searches only confirmed that the databases processed the query as it was entered. They did not confirm that the chosen terms covered all of the terminology used across the field. The use of exact phrases and the requirement for publications to match all three groups could therefore have excluded relevant publications that used different wording, such as in the British spelling for “defence”, referred to recovery or restoration using other terms, or

described adaptive recovery without referring to AI or any kind of machine learning in general. Limiting the search to IEEE Xplore, ScienceDirect and Web of Science may also have excluded relevant publications indexed in other academic databases.

The focus on recent adaptive and AI-supported recovery techniques meant that the search did not cover earlier research on traditional recovery methods to the same extent. Sources used in Section 2 to explain concepts and provide background information were identified separately from the systematic search. These sources were not part of the 21 selected publications, were not included in the thematic coding and were not used as evidence when answering the research question. Comparisons with traditional recovery methods in the results are therefore limited to cases where the selected publications described these methods or used them as baselines.

The publications in this review also used different techniques, environments and evaluation methods. This creates a validity threat to the conclusions as the reported outcomes were not always directly comparable. The findings were therefore analysed thematically and interpreted in relation to the settings in which they were evaluated, rather than combined into one overall measure of effectiveness. As the review included 21 publications, the results should be seen as an overview of the recent research identified throughout the chosen search strategy, rather than a complete representation of the entire field.

Another limitation is that it is difficult to avoid some degree of personal perspective or subjectivity when grouping together studies based on what is relevant to the research and when interpreting their findings. To try to eliminate this as much as possible, the analysis followed the adapted thematic process described in Section 4, based on the approach by Naeem et al. (2023). The same extraction fields were used for every publication to document and find relevant excerpts. Findings with similar meanings were then coded and grouped into subthemes. Publications could contribute to more than one theme or subtheme where their reported findings were relevant, as seen in Section 5. The selection criteria during the review process have been applied as consistently as possible to reduce the risk of a misleading or selective choice of publications. The included studies have also been handled with the goal of reflecting their contents in a fair way when reporting their findings, including both the benefits and limitations of adaptive recovery techniques. Overall, this study examines how recent publications evaluate adaptive recovery in relation to traditional methods, rather than providing a full comparison of both research areas.

5 Results

This section presents the findings from the 21 publications that remained after the quality assessment. The publications were analysed according to the recovery techniques they presented, their reported findings and the ways in which they contributed to the research question. Publications could contribute to more than one theme when they examined several techniques or reported findings relating to both technical approaches and their limitations. Section 5.1 provides an overview of the selected publications. Section 5.2 presents the findings according to the four main themes produced through the thematic analysis, and Section 5.3 summarises the results as a whole.

5.1 Overview of Selected Studies

The final selection of studies consists of 21 publications. These publications were analysed in order to identify which adaptive recovery techniques were presented in relevant research, how they were evaluated and how they compare to more traditional recovery methods. The chosen publications did not all contribute in the same ways. Some publications only presented more direct techniques for adaptive recovery or incident response, whereas others contributed to wider perspectives in cyber resilience, critical infrastructure, human oversight or limitations in autonomous cyber defence.

The chosen publications show that current research surrounding adaptive network recovery techniques is focused upon multiple recurring technical areas. A big part of the literature is about reinforcement learning and autonomous cyber defence agents, where response measures are chosen dynamically instead of through the use of static playbooks or predefined rules (Ren et al., 2025; Jaber et al., 2025; Iturbe et al., 2026; Loevenich et al., 2025). Another recurring approach is the use of LLMs within incident response, where they are used in order to support planning of actions or measurements to be taken, playbook generation and reasoning of workflows within security operations (Paduraru et al., 2025; Baral et al., 2025; Srinivas et al., 2025).

A smaller but still relevant group of publications focuses on self-healing and resilience oriented architectures, especially within cyber physical systems and critical infrastructure where service continuity and less downtime are important factors (Alijoyo et al., 2024; Alhidaifi et al., 2026; Senarak, 2025). Other publications are about adapting on a network level through the use of SDN, MTD, address hopping and other deceptive mechanisms that are designed to minimise predictability and limit attackers possibilities of succeeding (Xi et al., 2025; d'Ambrosio et al., 2025).

In the chosen literature, adaptive recovery is generally presented as an answer to limitations in traditional recovery methods. Many publications describe traditional methods as dependent on static rules, predefined playbooks, manual intervention or fixed network configurations. Adaptive methods are instead presented as more dynamic and context-aware as they can change their response methods depending on the state of the system, the attack behaviour, network conditions or other previously learned information. The literature also

shows that adaptive recovery is not unproblematic. Some recurring challenges include the cost of training LLMs, achieving realism in simulations, transferability to other environments along with complexity in implementation and a continued need of human oversight and monitoring.

5.2 Thematic Findings

The thematic analysis produced four main themes called “Adaptive Response and Recovery Mechanisms”, “Reported Resilience Functions and Outcomes”, “Evaluation Methods and Supporting Evidence”, and “Operational Limitations and Safeguards”. Each theme is presented through the five subthemes produced during the thematic analysis. Appendix B shows the complete mapping between the themes, subthemes, and publications.

5.2.1 Adaptive Response and Recovery Mechanisms

This theme covers the different technical ways in which adaptive response and recovery were presented in the selected publications. The findings include automated restoration, learning-based response, LLM and agent support, network reconfiguration and approaches that consider risk and the importance of different system functions.

Automated Restoration and Self-Healing

One technical approach in the chosen literature is self-healing and resilience oriented architectures. In comparison to reinforcement learning methods and LLM-driven methods, self-healing systems are more focused on how recovery is built into the system. The goal is for the system to be able to detect intrusions or disturbances, respond to them and restore functionality with less of a need for manual intervention.

Alijoyo et al. (2024) is the publication that most directly focuses on self-healing, with a study concerning cyber-physical systems. These are especially relevant as such systems are often used within critical infrastructure where downtime can lead to serious consequences. The authors of the study present adaptive self-healing as a way that could reduce downtime associated with security incidents. They describe how traditional security methods within cyber-physical systems have often been based on static rule-based approaches and that recovery processes tend to be manual and time-consuming. Their proposed solution is a self-healing system based on Long Short-Term Memory networks, where the system is intended to identify, respond to and recover from cyber-physical attacks autonomously. Detection, fast response and automation of the recovery phase are central aims of this study (Alijoyo et al., 2024).

Explicit self-healing was directly examined only by Alijoyo et al. (2024), while Boshnakov et al. (2025), Mejia-Ruiz et al. (2025) and Zhang et al. (2025b) contributed other findings connected to automated restoration and self-healing. The self-healing aspect is very relevant for adaptive network recovery, as it would mean that recovery becomes less of a separate measure during an incident and more of an integrated part of the system.

Learning Based Response Optimisation

One of the clearest technical approaches in the chosen literature is the reinforcement learning based recovery and autonomous response method. These studies often treat recovery and incident response as problems of decision-making, as a system or agent needs to choose the appropriate defence mechanisms based on the condition of the network, sophistication and development of an attack and previously observed information. This is different from traditional recovery methods where response measures are usually based upon predefined rulesets, static threshold values or manual decisions.

Ren et al. (2025) present ARCS, a framework that uses deep reinforcement learning to optimise automated incident response strategies. The framework combines a representation of security events along with hierarchical decision-making in order to link attack patterns to suitable defence measures.

Loevenich et al. (2025) also regards autonomous cyber defence and response, but through the use of a hybrid AI-architecture where deep reinforcement learning is combined with LLMs, rule-based systems and cyber security knowledge graphs. Their autonomous cyber defence agent can perform measures such as monitoring, analysis, decoy-distribution, removal of services and recovery of compromised systems. The study shows that autonomous response is not only about identifying threats, but also about being able to act within the network and the ability to recover affected services.

Other publications in this subtheme used learning based or cognitive approaches for recovery and response selection (Du et al., 2025; Iturbe et al., 2026; Jafar et al., 2025; Zhang et al., 2025a). These studies show how learning based methods can support adaptive network recovery through more dynamic decision making and less dependence on fixed response rules.

LLM and Agent Response Support

Another technical approach in the chosen literature is LLM-driven incident response and the generation of response plans. In these studies, LLMs are mainly used to help interpret incidents, figure out best measures and structure the response plan. This is relevant to adaptive network recovery as multiple parts of incident response may still tend to be manual, time-consuming and dependent on the assessment of security researchers.

CyberPlaybookLLM, as presented by Paduraru et al. (2025), is a system that uses LLMs to generate playbooks from structured incident data. The study relies on the theory that the latter parts of incident response, such as response planning, generation of playbooks and implementation are often still very manual, time-consuming tasks and that they are prone to inconsistencies. CyberPlaybookLLM uses a JSON-based intermediate representation to transform information from an incident into executable mitigation procedures. The study also shows how post-processing and prompt optimisation can lead to better overall consistency in playbook generation.

Baral et al. (2025) focus on LLM-driven incident response in their study through a framework based on Reasoning and Acting, or ReAct agents. In this case, the language model is not just used to generate text, but to analyse security alerts, reason around possible measures and to carry through responses via an integrated cyber security tool. The

framework also used monitoring tools along with a messaging system in order to maintain fast threat detection. The study shows that the ReAct agent can generate strategies to implement measures and autonomously carry out these strategies, for example through the use of firewalls. Their main value is the fact that they can reason about alerts and select actions dynamically, otherwise known as the reason-action loop.

The subtheme also includes LLM and agent support presented by Boshnakov et al. (2025), Castro et al. (2025) and Srinivas et al. (2025). Together, these publications show how LLM-driven methods and agents can help contribute to adaptive network recovery by creating more structured response plans with less manual input and supporting more context aware incident response.

Network Reconfiguration and Deception

Another technical approach in the chosen literature is adaptive mechanisms on a network-level. This approach focuses on the ability to change the network environment in order to reduce an attacker's possibilities of observing, predicting or using a system. The adaptability is not just centered around choosing a measure of response, but through dynamically changing the network addresses, flow of traffic or exposed services.

Xi et al. (2025) present an SDN-based model used for address hopping in IoT networks. The study is based on how more traditional anonymising methods are often based on static identity models, making them more vulnerable to inference attacks through long-term observation. In order to make this less predictable, Xi et al. (2025) make use of a MTD model and SDN in order to periodically change network addresses and routes. This means that device identifiers, communication routers and different parts of the network topology can change over time, making it harder for an attack to use previously collected or observed information.

d'Ambrosio et al. (2025) present another type of network-based adaptation through the use of SMASH, a framework that combines SDN, MTD and honeypots. The study is based on how traditional security tools such as firewalls and intrusion prevention systems are still important, but that they have clearer limitations when cyber attacks become more complex. SMASH proposes a solution to this as it uses dynamic redirection of traffic and deception-based techniques to drive malicious traffic away from the production network and into an isolated threat intelligence network environment. This means that an attacker can be isolated from sensitive systems at the same time as the organisation and relevant cyber security researchers have the opportunity to observe the actions of an attacker in a controlled environment.

The framework also supports automatic distribution, configuration and monitoring of honeypots. Overall, these publications suggest how adaptive network recovery is not just about AI-based decision making, but also about making networks more dynamic and less predictable. Alijoyo et al. (2024), Mejia-Ruiz et al. (2025) and Zhang et al. (2025b) also contributed findings to this subtheme.

Risk and Mission-Aware Management

The final subtheme within adaptive response and recovery mechanisms concerns methods that take account of risk, resilience, mission requirements or the operational importance of the system. This subtheme included Alhidaifi et al. (2026), Burbano et al. (2025), Hausken et al. (2025), Jaber et al. (2025), Mejia-Ruiz et al. (2025) and Senarak (2025).

Burbano et al. (2025) help to explain in their study why autonomous cyber defence agents are relevant for recovery, describing how traditional incident response often leads to delays as human analysts need to check for indications of cyber attacks and then follow up with incident response plans. They also bring to light how human led recovery is often a slow process. In contrast, autonomous cyber defence agents can automatically block or place malicious traffic in quarantine, isolate any affected endpoints, reset login credentials and restart or even reinstall devices. This connects autonomous response directly to cyber resilience as the goal is to reduce the amount of time it takes between detection and the subsequent removal of an attacker from a network (Burbano et al., 2025).

Senarak (2025) used expert opinion in order to identify the most important response and recovery tasks. The other publications in this subtheme focused more on risk, resilience, infrastructure and mission-aware response rather than focusing on one specific recovery action.

5.2.2 Reported Resilience Functions and Outcomes

This theme focuses on the ways in which the selected publications reported that adaptive techniques may contribute to cyber resilience. The findings mainly concerned containment, adaptability, recovery coordination, response speed and the ability to restore or maintain services.

Containment and Threat Suppression

Baral et al. (2025) show that the ReAct agent can generate strategies to implement measures and autonomously carry out these strategies, for example through the use of firewalls. This was evaluated against attacks like port-scanning, botnet intrusions and secure shell brute force attacks.

Iturbe et al. (2026) investigated learned actions but then described the start of the recovery process as future work. Their contribution is therefore more closely connected to containment and response than demonstrated restoration. Jaber et al. (2025) also reported response outcomes, although their training process was not clearly described to support a comparison with traditional methods.

The results from Xi et al. (2025) show that their dynamic model reduced the exposure of device identity and scan success rates. Xi et al. (2025) reported a 27.3% reduction in IP hit rate and a 60% reduction in communication delay under high attack intensity compared with their static baseline. This makes the publication relevant for adaptive network recovery as it shows how network configurations can be used actively in order to preserve communications and cut off an attacker's possibilities during an ongoing attack.

d'Ambrosio et al. (2025) reported that SMASH reduced the number of real services discovered by approximately 77% in one evaluated setup and 85% in another. They also report that SMASH was able to mitigate multiple stages of the attack chain by mimicking production assets behind a dedicated Threat Intelligence Network. This allows for both low and high-interaction honeypots to gather attacker intelligence. Jafar et al. (2025) also contributed findings connected to containment and threat suppression.

Decision Quality and Adaptability

Ren et al. (2025) gave the clearest comparison to a traditional rule-based method. In their offline evaluation, ARCS reported a response effectiveness of 89.2% compared to 72.3% for the traditional rule-based baseline. The average response time decreased from 45.6 in their traditional rule-based baseline to 23.5 seconds. Their reported false-positive rates were also lower at 3.9%, compared to 8.9% in the traditional method. The framework is also described as being especially strong in situations where it had to manage more complex attacks with multiple steps and it was able to efficiently balance incident resolution time, system stability and defence effectiveness through its reward mechanism.

In the evaluation by Loevenich et al. (2025), the autonomous agent performed better than baseline implementations, and the deep reinforcement learning based models could both mitigate and recover compromised systems. Loevenich et al. (2025) also made use of an LLM as a generation based chatbot, able to function as a user interface between the system and human cyber security experts. This means that LLMs are not only used directly for autonomous response, but also as a way to support interpretation of events and to make the systems more user-friendly for human operators.

Paduraru et al. (2025) evaluated the response playbooks that were generated and reported a validity of 96.8% and an execution-path completeness of 92.5%. These results mainly show how well the playbooks were structured rather than how well they would actually perform during a real incident. Iturbe et al. (2026) also contributed findings connected to decision quality and adaptability.

Recovery Coordination and Verification

The experiment conducted by Du et al. (2025) reported recovering exploited hosts in varying amounts of steps depending on the type of agent used. Teams working with cognitive agents were able to complete the task with an average of 1.45 steps, heuristic agents with an average of 2.15 steps, and random agents managed with an average of 3.92 steps.

Senarak (2025) used expert opinion in order to identify important response and recovery tasks. Together, these publications show that recovery coordination can involve both human-agent recovery actions and organising important recovery tasks.

Response Speed and Workload Reduction

Castro et al. (2025) compared two adaptive techniques in their CAGE 4 simulation. This is a simulated test environment in which their reinforcement learning agents took action around 104 times faster than the LLM agents. Baral et al. (2025) reported that their ReAct agent was able to apply mitigation actions within seconds, but there was no direct comparison or traditional response baseline provided.

Compared to traditional methods where playbooks are often created and updated by hand, an LLM-based system can lessen the load and lead to a more flexible and faster response. The framework presented by d'Ambrosio et al. (2025) also supports automatic distribution, configuration and monitoring of honeypots, which can reduce work that would otherwise be carried out manually. Paduraru et al. (2025) and Zhang et al. (2025a) also contributed findings to this subtheme.

Service Restoration and Continuity

Most of the outcomes based on recovery concerned the amount of time or the number of steps taken in order to restore assets that had been compromised. A study on malware propagation by Jafar et al. (2025) combined reinforcement learning and a model-predictive-control approach that reported restoring all 60 of their simulated devices in as little as 6 seconds. The study's comparative results gave restoration times that ranged from 14 to 38 seconds for other evaluated approaches. Zhang et al. (2025a) reported an average recovery time of 6.2 seconds, compared with 12.7 seconds for a deep reinforcement learning-based intrusion detection system, which was used as an adaptive baseline. Snort, a traditional intrusion detection system that uses predefined rules and attack signatures, did not provide automatic recovery (Zhang et al., 2025a).

Boshnakov et al. (2025) investigated response time, recovery time and service availability in their automated cloud-recovery workflow, although their reported results were compared with separate figures within the industry rather than traditional recovery methods tested in the same environment. Alhidaifi et al. (2026) used recovery and resilience scores from their model to compare strategies in simulated infrastructure. These values were produced by a probabilistic resilience model and did not represent actual recovery in a real operational network.

Du et al. (2025) and Zhang et al. (2025b) also contributed findings to service restoration and continuity. Overall, the publications show how adaptive techniques are intended to support faster restoration, less dependence on manual response and better continuity during and after an incident.

5.2.3 Evaluation Methods and Supporting Evidence

This theme focuses on how the adaptive techniques were evaluated and what kind of evidence was used to support the reported findings. The publications used different baselines, simulations, testbeds, offline datasets, expert input and review-based evidence, which affected how directly the results could be compared.

Method and Baseline Comparison Evidence

Ren et al. (2025) clearly compared their adaptive method against a traditional rule-based approach. This shows that ARCS supported faster and more effective responses than the traditional rule-based method within the evaluated dataset (Ren et al., 2025). Loevenich et al. (2025) compared the performance of their autonomous agent with baseline implementations, while Paduraru et al. (2025) compared the quality of generated playbooks rather than recovery performance.

In the comparison by Castro et al. (2025), the training time and LLM prompt-engineering were excluded, so only the time needed for the agents to select and carry out actions during the simulation was compared. This comparison did not include a traditional recovery method and is therefore more useful for showing the differences between adaptive approaches than for directly answering how adaptive recovery compares with traditional methods.

The comparison reported by Xi et al. (2025) was made with Honey-MTD, which is another adaptive approach rather than a traditional recovery method. The study also included a static baseline, but its evaluation mainly concerned anonymisation intended to reduce exposure before an attack succeeds rather than recovery. Iturbe et al. (2026) also contributed method and baseline comparison evidence.

Out of the 21 publications in this review, the three clearest comparisons with traditional or static methods for a response or recovery related outcome were Jafar et al. (2025), Ren et al. (2025) and Zhang et al. (2025a). Many of the other publications instead compared different adaptive methods or evaluated other parts of response and recovery.

Emulation, Testbed and Prototype Evidence

Baral et al. (2025) evaluated their framework against attacks like port-scanning, botnet intrusions and secure shell brute force attacks in simulated environments. Boshnakov et al. (2025) investigated an automated cloud-recovery workflow, while Iturbe et al. (2026) investigated learned response actions.

Xi et al. (2025) evaluated an SDN-based address hopping model for IoT networks. d'Ambrosio et al. (2025) evaluated SMASH using an isolated threat intelligence network environment in which the actions of an attacker could be observed under controlled conditions. Zhang et al. (2025a) also contributed evidence from an emulated or testbed environment.

These evaluations show how the methods worked in their tested environments, but they do not show how well the same techniques would work in real production networks.

Evidence Limitations and Readiness for Real-World Use

Some publications researched components that could support recovery without measuring recovery itself. Alijoyo et al. (2024) describe in their study how self-healing mechanisms can be a way to reduce downtime in relation to security incidents and report results between 99% and 100% across accuracy and precision for the Long Short-Term Memory model. These values concerned identifying abnormal system behaviour, while no restoration time, downtime, service availability outcome or corrective action was reported. The findings

therefore show the model's ability to classify anomalies rather than demonstrate complete self-healing recovery.

The playbooks studied by Paduraru et al. (2025) were not tested in a live security orchestration, automation and response (SOAR) environment and still required human oversight and approval before they could be used. Iturbe et al. (2026) investigated learned actions but described the start of the recovery process as future work. Jaber et al. (2025) also reported response outcomes, although their training process was not clearly described to support a comparison with traditional methods.

The results reported by Ren et al. (2025) are promising and show that adaptive response can be effective in the dataset that was used, but they should not be seen as evidence of how well the method would perform in a real incident response environment. The findings from Xi et al. (2025) were mainly related to preventive address hopping and maintaining availability rather than recovery after a compromised service. The values produced by Alhidaifi et al. (2026) came from a probabilistic resilience model and did not represent actual recovery in a real operational network. Boshnakov et al. (2025), Senarak (2025) and Srinivas et al. (2025) also contributed findings concerning evidence limitations and readiness for real-world use.

Evidence for complete self-healing recovery was more limited than evidence for reinforcement learning and LLM-driven response because fewer publications directly focused on self-healing network recovery. The available evidence is therefore not sufficient to conclude that self-healing systems improve recovery compared with traditional methods.

Human, Expert and Review Evidence

The experiment by Du et al. (2025) was reportedly conducted in a small simulated network with participants that were not experts in the area, but it still gives a comparison of recovery performance in this setting. Senarak (2025) used expert opinion in order to identify the most important response and recovery tasks.

Hausken et al. (2025) and Srinivas et al. (2025) contributed review-based findings rather than new recovery experiments. These publications provide a wider view of recovery and resilience, but not the same type of evidence as direct technical evaluations.

Numerical Simulation and Offline Evaluation

The findings by Jafar et al. (2025) are based on a malware propagation model rather than a real network, which limits how directly they can be applied to an operational network. Ren et al. (2025) evaluated ARCS using an offline dataset rather than a real incident response environment.

Loevenich et al. (2025) mainly evaluated their autonomous defence system through accumulated reward in simulated attack scenarios. The study also evaluated the chatbot that was included for cyber security experts, but this focused on explanation quality rather than actual recovery performance. Paduraru et al. (2025) evaluated generated response playbooks, while Alhidaifi et al. (2026) used a probabilistic model to compare strategies in simulated infrastructure.

5.2.4 Operational Limitations and Safeguards

This theme focuses on a recurring pattern in the reviewed publications. Adaptive recovery is often presented as faster and more flexible than traditional methods, but the publications also describe several risks and limitations that mean the techniques may not yet be ready for widespread use or be free from risk.

Adversarial and Automation Risks

Burbano et al. (2025) raise an important point being the risks of autonomous cyber defence agents. They show that the agents themselves can be targets for attacks, especially if the attacker can manipulate the information that the agent uses in order to make decisions or interfere with the measures that the agent tries to implement. This is explained through the use of the term meta-attacker, where the attacker does not just try to attack the network, but also its defence systems. This is important as it shows that an autonomous recovery method in this case is not only working as a protection, but that it can also become another new vector for attacks.

Generated response plans may need structural validation in order to be useful in practice (Paduraru et al., 2025). An LLM-based response can reduce manual work, but it does not necessarily replace human judgement. Network deception can also introduce automation risks. Honeypots need to be isolated properly in order to make sure that an attacker cannot use them as a pivot point into a production network. Deeper network monitoring may create privacy risks, while false positives may cause legitimate users to be redirected to honeypots (d'Ambrosio et al., 2025). Boshnakov et al. (2025) and Castro et al. (2025) also contributed findings to this subtheme.

Data, Computing Needs, Scalability and Integration

Some recurring challenges include the cost of training LLMs and the complexity of implementing adaptive techniques. d'Ambrosio et al. (2025) explain that honeypots and deception-based techniques can be difficult to distribute and manage in bigger network environments. Manual configuration of honeypots is described as being time and resource consuming.

The publications also identified wider requirements concerning data, computing, integration and scalability (Hausken et al., 2025; Iturbe et al., 2026; Jafar et al., 2025; Mejia-Ruiz et al., 2025; Ren et al., 2025; Srinivas et al., 2025; Zhang et al., 2025b). These requirements affect whether adaptive recovery techniques can be implemented and maintained outside the environments in which they were evaluated.

Applicability Across Different Settings and Scenarios

Many of the evaluated methods have only been tested in simulated or controlled environments. This means that their actual effectiveness in a real production network could still be affected by factors such as the cost of implementation, training and differences between operational environments.

Baral et al. (2025) stress that ReAct agents are in an early stage for cyber security and that their use still requires further validation. The study also discusses risks connected to

generative AI, such as prompt injection, jailbreaks, bias, high computational costs and the possibility for the language model to be misused by attackers. Loevenich et al. (2025) stress that the quality of autonomous cyber defence depends on how realistic the attack environment is during training and testing.

Self-healing methods may also depend on the specific environment, data and system architecture used during evaluation. These systems could therefore be difficult to transfer into other network environments without further testing (Alijoyo et al., 2024). This is important, as incorrect automatic measures could potentially have serious consequences in a production environment, affecting both digital and physical processes.

The applicability of the findings is also affected by the limited attack scenarios and settings covered by the publications. Alhidaifi et al. (2026), Castro et al. (2025), Du et al. (2025), Hausken et al. (2025), Jaber et al. (2025), Jafar et al. (2025), Mejia-Ruiz et al. (2025), Ren et al. (2025), Senarak (2025) and Zhang et al. (2025a) also contributed findings concerning applicability across different settings and scenarios. The complete mapping is shown in Appendix B.

Human Oversight, Trust and Governance

The publications show that LLM-driven response should not be seen as being completely independent or risk free. Generated measures and response plans should still be validated and reviewed by human analysts, especially considering how potentially bad response plans in a network could harm or wrongfully affect the availability or security of a network. Human oversight may be especially important where sensitive data is handled.

Human oversight is not only about approving measures, but also about being able to understand and control a system's decision processes. The techniques therefore need validation and a clear responsibility structure before they are used for adaptive recovery. Boshnakov et al. (2025), Du et al. (2025), Iturbe et al. (2026), Paduraru et al. (2025), Ren et al. (2025), Srinivas et al. (2025) and Zhang et al. (2025a) contributed findings to this subtheme.

Security and Performance Balance

Xi et al. (2025) report that network-based methods must be balanced with security, anonymity, performance and availability. The same study also examined CPU load and the balance between hopping frequency, exposure and delay. Even if their SDN-based address hopping model aims to keep a high level of data availability while strengthening security, this shows that adaptive network mechanisms must be assessed based on several factors at the same time.

Loevenich et al. (2025) also report how adaptive cyber defence systems can have issues with explainability. They discuss how deep reinforcement learning (DRL)-based autonomous cyber defence models can serve as black boxes, making them harder to interpret and trust in real environments. In order to better handle them, their architecture includes a chatbot, meant to help cyber security experts understand events, defence decisions and various other defensive techniques. Mejia-Ruiz et al. (2025) also contributed findings concerning the balance between security and operational performance.

This final theme shows how adaptive network recovery cannot only be assessed by speed or technical effectiveness. The techniques need to be understood in relation to the risks, human control, cost of implementation and operational reliability. Compared to traditional recovery methods, adaptive methods can give more flexibility and faster responses, but that can also introduce new vulnerabilities and management requirements.

5.3 Summary of Key Findings

One of the clearest patterns was that most of the recent literature focused more on adaptive decision-making and automated response than on full self-healing recovery. The mechanisms were represented through automated restoration and self-healing, learning based response optimisation, LLM and agent response support, network reconfiguration and deception, and risk and mission-aware management. Since publications could contribute to more than one subtheme, the findings overlap and should not be treated as separate publication groups.

The majority of the publications seemed to agree that adaptive techniques can be used to support more flexible responses that are better adapted to a situation. They also agree that this could decrease the amount of dependency needed on manual human intervention should an event occur, and help to improve response time or restoration speed within their particular environments of evaluation. However, the effectiveness of the tests was defined differently between the studies, and they can therefore not be compared as if they reflected the same type of recovery outcome. For example, reinforcement learning studies commonly measured rewards, response effectiveness, response time or restoration outcomes. LLM studies more often measured reasoning quality, playbook structure or workload reduction, while network-level studies measured communication delay, availability or resource use.

The evidence from the publications did not show that one adaptive technique was better than the others. Castro et al. (2025) reported higher rewards and faster action for reinforcement learning agents than for LLM agents in their simulation. For comparison, the publications focused on LLMs mainly reported benefits connected to incident reasoning, playbook generation and support for human operators. The approaches may therefore have different strengths. Comparisons with traditional recovery were also uncommon, as only three publications used a traditional or static baseline for a recovery or response related outcome (Jafar et al., 2025; Ren et al., 2025; Zhang et al., 2025a). These findings therefore provide examples where adaptive methods achieved stronger reported outcomes, but not enough evidence to say that adaptive recovery is generally better across networks.

Another recurring pattern was that much of the evidence focused on only one part of the recovery process. Many of the publications investigated detection or response planning without actually demonstrating restoration of services. Most of the technical evaluations were also conducted using simulations, offline datasets, emulated networks, prototypes or small controlled experiments. This may mean that complete recovery in production networks is still underrepresented.

Finally, every included publication contributed at least one limitation. The most common concerns included transferability to other environments, limited attack scenarios, data and computational requirements, integration and scalability, security and performance compromises, risks of manipulation and the continued need for human oversight or manual intervention. Overall, the results show that adaptive techniques have potential uses for response, recovery support and network reconfiguration. However, the strongest available evidence currently comes from specific controlled settings such as in the included publications rather than complete recovery in production networks. These findings provide the basis for the following discussion about the significance, benefits and limitations of adaptive network recovery.

6 Discussion

This section discusses the main findings presented in Section 5, how they relate to the research question, previous research along with ethical and societal aspects.

The theme, “Adaptive Response and Recovery Mechanisms”, showed that the selected publications present several ways in which adaptive techniques may support network response and recovery, although the strength of the evidence differs between the techniques and settings during evaluation. Reinforcement learning and autonomous agents mainly supported the selection and execution of actions, while LLMs were more focused on reasoning and response planning. The few self-healing approaches focused mainly on maintaining or restoring system functionality, while network-level techniques changed the network environment in order to reduce exposure or maintain communication. This shows that adaptation can take place at several different stages of response and recovery rather than through one single technique.

Earlier research also shows that adaptation can work in different ways. This can include systems that are able to manage parts of their own behaviour, as well as networks that can be changed more dynamically through software (Kephart & Chess, 2003; Kreutz et al., 2015; Cho et al., 2020). This supports the idea that adaptive recovery is not based around one single technique, but can involve different types of response, reconfiguration and recovery.

However, the findings within the “Evaluation Methods and Supporting Evidence” theme show that the publications included in this review do not provide enough comparable evidence to be able to determine whether adaptive recovery is more effective than traditional recovery. The selected publications show that traditional methods were typically described as static, manual or dependent on predefined playbooks, while adaptive techniques were presented as more flexible and better at adjusting to changing situations within their evaluated settings. Only three publications actually compared an adaptive method with a traditional or static baseline for a recovery or response related outcome (Jafar et al., 2025; Ren et al., 2025; Zhang et al., 2025a). The baselines used in these comparisons also differed in their functions, and one did not make use of automatic recovery. Because of this, the comparisons were not always between methods with the same capabilities. At best, the reported advantages are results achieved under specific experimental conditions, rather than evidence that adaptive techniques would perform better than well-known and established recovery methods in operational networks.

The three direct comparisons with traditional or static methods show why the results need to be interpreted carefully. Ren et al. (2025) reported better response effectiveness and shorter response times than a rule-based baseline, but the evaluation was based on an offline dataset and only one traditional comparison method. Jafar et al. (2025) reported improved results for their combined reinforcement learning and Model Predictive Control approach within a numerical malware propagation model. Zhang et al. (2025a) reported faster recovery than a DRL-based system, while Snort did not provide automatic recovery. Therefore, the comparison with Snort mainly shows a difference in functionality. Together, these studies show improved results in certain settings, but they do not prove that adaptive techniques generally perform better than current recovery methods.

As described in Section 3.3, the search focused on recent adaptive and AI-supported approaches and did not investigate traditional recovery research to the same extent. Descriptions of traditional methods in the reviewed material mainly came from publications proposing an adaptive alternative. These publications may have naturally placed greater emphasis on the limitations of the methods they aimed to improve. Another limitation of this study is that it cannot determine whether the traditional methods used for comparison were the most effective options currently available. Therefore, this study looks at how recent adaptive recovery research describes and evaluates these traditional methods, but it cannot provide a complete comparison of the two research areas or confirm that adaptive recovery is generally better than the other.

This theme also showed that the settings used during evaluation limited how widely the results could be applied. Most of the technical evaluations were based on simulations, offline datasets, cyber ranges, emulated networks or controlled prototypes. These environments make it possible to test disruptive actions in a safer way, but they cannot fully represent factors such as older systems, approval processes within an organisation, or specific dependencies between different services. The repeated concerns about transferability in the selected publications support this interpretation (Iturbe et al., 2026; Loevenich et al., 2025; Ren et al., 2025). The reported effectiveness should therefore mainly be seen as applying to the environments in which the techniques were tested, rather than to operational networks in general.

Similar concerns have also been raised in earlier research on machine learning for network security. Sommer and Paxson (2010) explain that results from controlled test environments may not always transfer well into real operational networks, as real networks can be much more complicated and unpredictable. This supports the idea that the results in this review should mainly be understood in relation to the environments where the techniques were actually tested.

The theme, “Reported Resilience Functions and Outcomes”, showed that the chosen literature appears more developed around adaptive response and containment than around full recovery. Many of the publications evaluated detection, selection of actions, isolation, response planning or reconfiguration without showing that the affected service was fully restored (Alijoyo et al., 2024; Iturbe et al., 2026; Paduraru et al., 2025; Xi et al., 2025). This matters because a fast response is not necessarily the same as a successful recovery. Recovery also concerns whether essential functionality remains available, whether services are returned to an acceptable operational state and whether the recovery actions may have caused further disruption. This may therefore make the available evidence appear more complete than it really is.

The findings can also be interpreted in relation to cyber resilience that was introduced in Section 2.1.1. Cyber resilience involves being able to anticipate, withstand, recover from and adapt to adverse conditions rather than only detecting or responding quickly to an attack (Ross et al., 2021). However, many of the chosen publications evaluated only one part of this process. Alijoyo et al. (2024) measured anomaly classification without showing an executed recovery action, and Paduraru et al. (2025) evaluated generated playbooks without actually testing them in an active recovery environment. Iturbe et al. (2026) described the beginning of a recovery process as future work, while Xi et al. (2025) focused on preventative network

reconfiguration. These operations may contribute to cyber resilience, but they should not be interpreted as evidence that complete recovery was achieved. Altogether, these examples support the finding that the literature is currently more focused on response and containment rather than complete service restoration.

As such, the different techniques may be better understood as helpful parts of the recovery process rather than as real alternatives. Reinforcement learning may support fast action selection, LLMs may support incident interpretation and planning, network-level mechanisms may preserve communication or isolate threats, and self-healing mechanisms may support restoration (Jafar et al., 2025; Loevenich et al., 2025; Paduraru et al., 2025; Xi et al., 2025). Castro et al. (2025) reported advantages for reinforcement learning defenders over LLM defenders in action speed and reward, while the LLM-focused publications reported different benefits that were related to reasoning, response planning and support for human operators (Baral et al., 2025; Loevenich et al., 2025; Paduraru et al., 2025). Hybrid approaches also combined multiple operations within the same system (Jafar et al., 2025; Loevenich et al., 2025). These differences make it difficult to determine whether one technique performs better overall than the others. Instead, the findings suggest that practical cyber resilience may depend on combining adaptive techniques with existing recovery methods, rather than expecting one method to perform all parts of the recovery process.

This can also be compared with established guidance for cyber incident recovery, where recovery still includes planned restoration, rebuilding, backups and human decision-making (Bartock et al., 2016; Nelson et al., 2025). This supports the idea that adaptive techniques may be more realistic as additions to existing recovery processes rather than as complete replacements for them.

The theme, “Operational Limitations and Safeguards”, supports a gradual and supervised use of adaptive recovery. Low risk and reversible tasks, such as gathering incident information, recommending actions or generating response plans, may be more suitable for automation in its current state. On the opposite end, actions that isolate essential systems, reset credentials, fully rebuild compromised systems or change critical network routes should require more thorough verification and human approval before being executed (Boshnakov et al., 2025; Du et al., 2025; Paduraru et al., 2025). This approach could allow organisations to benefit from faster information gathering and response planning while maintaining human control over decisions that could otherwise have serious consequences. Adaptive techniques may therefore currently be most useful when they support established recovery processes rather than replacing them entirely.

This can also be related to earlier research on human interaction with automation. Parasuraman et al. (2000) describe automation as having different possible levels rather than being completely manual or completely autonomous. This supports the idea that the amount of human involvement in adaptive recovery should depend on the type of action being carried out and the possible consequences if something goes wrong.

6.1 Ethical and Societal Aspects

The results show that adaptive recovery could provide some societal benefits if it was used in systems that organisations and the public rely on. Faster containment and recovery could reduce downtime, limit data loss and help maintain access to important digital services. This could be especially valuable in cloud infrastructure, ports for handling cargo, energy networks and other environments where disruptions can affect many organisations or users (Boshnakov et al., 2025; Mejia-Ruiz et al., 2025; Senarak, 2025; Zhang et al., 2025b). However, it must be mentioned again that since most of the evidence came from controlled settings, these benefits should still be seen as possible outcomes rather than confirmed effects in society. The findings were also reported cautiously in order to avoid overstating the effectiveness of adaptive or AI-supported recovery techniques beyond the evidence provided by the selected publications.

One of the main ethical risks is the consequences of incorrect automated actions. An adaptive system may respond faster than a human, but an incorrect decision could isolate users, interrupt an important service or restore a system using a bad configuration. Boshnakov et al. (2025) warn that automated recovery could act on an incorrect classification or restore on an insecure baseline, and d'Ambrosio et al. (2025) identify a risk of users being redirected into honeypots. Burbano et al. (2025) also show that attackers may be able to manipulate the observations or actions of autonomous defence agents. If these techniques were to be used in operational or production networks, the recovery system itself could therefore become another target for attack and an additional point of failure. This also raises questions about responsibility, as it may be unclear who is accountable when an automated recovery decision causes harm. Human oversight and clear limits on which decisions may be automated remain very important, particularly when these types of services or sensitive data could be affected (Du et al., 2025; Srinivas et al., 2025).

If an autonomous system were to cause an outage or remove access for legitimate users, it may be unclear whether the responsibility belongs to the developer, the organisation or the person who approved the system. This becomes especially important when reinforcement learning decisions are difficult to explain or when an LLM generates an incorrect response plan. This is why several publications included human approval, explanation tools or manual review before more serious actions were carried out (Du et al., 2025; Loevenich et al., 2025; Paduraru et al., 2025; Zhang et al., 2025a). Human oversight should be seen as an important safeguard rather than only something that limits automation. Real-world implementations would also need clear responsibility, records of these automated decisions and a way for authorised staff to stop or reverse harmful actions if and when they should occur.

Privacy and unequal access are also important societal concerns. Adaptive recovery systems may rely on detailed security logs, network traffic and information about user behaviour. If this information is not properly protected while being collected or processed, sensitive organisational or personal data could be exposed. This is a serious privacy concern that may harm affected individuals and reduce trust in the organisations using these systems. Systems based on deception may also have privacy concerns if legitimate users are incorrectly redirected into monitored environments (d'Ambrosio et al., 2025; Srinivas et al., 2025). The results also showed that some adaptive techniques need large amounts of data and

computational resources. This may make these systems easier for large organisations to implement than for smaller organisations with fewer resources (Ren et al., 2025; Srinivas et al., 2025). Unequal access to these techniques could therefore create differences in the level of cyber resilience available to organisations and, in turn, the users who depend on their services.

The results can mainly be linked to two United Nations Sustainable Development Goals for a more sustainable, fair and secure world by 2030. The first is Goal 9, which concerns innovation and resilient infrastructure. Adaptive recovery may support this goal by helping digital infrastructure maintain and restore functionality during cyber incidents. The findings of this study can also be connected to Goal 16, which concerns reliable institutions and public services. This is relevant as adaptive recovery could be used within public institutions or critical public services, where automated decisions need to remain transparent and controlled (United Nations, n.d.-a, n.d.-b). However, this study does not show that adaptive recovery directly achieves either of these goals. Implementing services with unsafe automation that could have unequal access or unclear responsibility could end up working against these goals.

Overall, the ethical and societal value of implementing systems that use adaptive recovery depends on how the techniques are applied and regulated. The results support gradual and supervised automation, where systems are tested in realistic environments and where serious actions still need further testing and validation. Clear responsibility, human oversight and control, protection of privacy, and the ability to reverse inaccurate actions are necessary before adaptive techniques and recovery can be fully trusted in important production networks (Boshnakov et al., 2025; Du et al., 2025; Srinivas et al., 2025).

7 Conclusion

The aim of this study was to identify adaptive network recovery techniques presented in recent research and examine how they may support cyber resilience. Comparisons with traditional or static recovery methods were also examined where reported. This systematic literature review included 21 publications and identified four main themes through the thematic analysis. These themes were further divided into 20 subthemes.

The chosen publications reported several ways in which adaptive techniques may support cyber resilience. This included picking response actions according to changing conditions, generating response plans, reconfiguring networks, and helping restore system functionality. However, the findings were more focused on response, containment, and making decisions rather than complete system recovery. Metrics such as detection accuracy, performance score, playbook validity, and reduced network exposure may measure useful functions, but they do not by themselves show that a disrupted service was actually restored. Although eight publications contributed to subthemes related to automated restoration, service continuity and self-healing, only one focused specifically on a self-healing recovery approach (Alijoyo et al., 2024).

Only three publications actually compared an adaptive technique with a traditional or static baseline for a recovery or response related outcome. Some of these baselines were quite simple or did not provide the same recovery functions as the adaptive system. The review also could not determine whether they represented the strongest traditional methods that are currently available and in use. The results therefore show that some adaptive systems achieved better reported outcomes than their chosen baselines under certain controlled conditions, but they do not show that adaptive recovery is generally more effective than traditional recovery.

These conclusions are limited by the small and recent set of publications and by a search strategy that focused mainly on adaptive and AI-supported methods. Most of the technical evidence came from simulations, offline datasets, cyber ranges and prototypes rather than real production networks. It is also possible that successful or newer approaches were more likely to be published than unsuccessful or unclear results, and this review cannot know how many negative results were never published. As this is the case, the findings should be understood as a structured summary of what the 21 selected publications reported, rather than as a general success rate for adaptive recovery across the wider research field.

Overall, the findings provide a careful answer to the research question. Recent research has presented multiple adaptive techniques that can support different parts of network response and recovery. In the few studies that directly compared adaptive and static methods, the adaptive techniques reported better results for certain outcomes. However, only three publications made this type of direct comparison, and the studies used different evaluation methods in simulated or controlled environments. The current evidence therefore suggests that adaptive techniques can improve certain parts of network response and recovery in the settings where they were evaluated, but it does not show that they generally perform better than traditional methods in real production networks. Adaptive recovery should therefore be seen as a developing area that can support and complement existing recovery methods,

especially if the techniques have been properly tested, compared against suitable traditional baselines, applied with good security in mind and used with enough human oversight and with the possibility of manual intervention.

7.1 Future Work

Future work surrounding this topic or area of research could be promising, especially as more empirical studies are published in the coming years. Future research should preferably aim to improve the comparability of the evaluation results. A study focused on a common set of evaluation measures across different adaptive recovery techniques could be beneficial to the field as they would provide a stronger view of cyber resilience rather than single measures that cannot be applied consistently across different techniques. These measures could include response time, restoration time, availability of a service during an incident, failure and false-positive rates, resource use and the cost of implementation.

More direct comparisons with strong traditional methods are also needed in this field as adaptive methods should be tested under the same attacks, network conditions and operational requirements as traditional methods. Where relevant, studies should compare against current rule-based, manually supervised and hybrid approaches, while also explaining why each baseline was chosen. Researchers on this topic should also test each part of hybrid systems separately in order to determine whether an improvement actually comes from the adaptive techniques or from another part of the architecture.

Self-healing was one of the areas intended to be examined in this study and was also included in the search terms. However, as mentioned previously, it was underrepresented in the selected literature. Future research should therefore give more attention to the complete recovery of networks and systems to make up for this. Any evaluations should follow and document an incident from detection all the way to restoration and then return to acceptable service conditions. Researchers should also step away from small simulations and offline datasets towards larger and more realistic simulations and controlled trials. Testing the methods across different types of networks, while also reporting unsuccessful actions and unexpected results, would help provide stronger evidence about how well they scale and integrate with existing systems.

Another important area for future work is how humans and other organisational factors work alongside autonomous techniques. These new studies should examine which actions can be automated safely where typically human oversight and approval is required and whether they help operators to make better decisions or if they create increased workload and dependence. Adaptive recovery techniques should also be tested against malicious activities such as manipulation, prompt injection and harmful inputs to see how the system reacts. Future work could also evaluate logs, backups and rollback procedures as part of the recovery system, as these could provide additional support when automated actions fail or cause unexpected outcomes.

Finally, a future literature review could also include a larger publication period with additional databases. Such a study could also have a separate systematic search including traditional recovery methods. On a smaller scale, this study could be continued by selecting one or two of the adaptive techniques identified in this review and comparing them with a traditional or rule-based recovery method in the same controlled environment. Using the same attack scenarios and evaluation measures for both approaches could provide a more direct comparison and help to address one of the limitations of this study. A wider review would also help address the fact that the publications used in this study mainly focused on adaptive approaches and techniques, while also giving a more balanced basis for comparison. If future studies start using more common outcome measures and similar baselines, it may also become possible to make stronger quantitative comparisons across the literature.

References

Ahmed, R., & Boutaba, R. (2011). A survey of distributed search techniques in large scale distributed systems. *IEEE Communications Surveys & Tutorials, 13*(2), 150–167. https://doi.org/10.1109/SURV.2011.040410.00097

Alhidaifi, S. M., Asghar, M. R., & Ansari, I. S. (2026). Cyber resilience quantification: A probabilistic estimation model for IT infrastructure. *Reliability Engineering & System Safety, 265*, 111473. https://doi.org/10.1016/j.ress.2025.111473

Alijoyo, F. A., Kaur, C., Anjum, A., Vuyyuru, V. A., & Bala, B. K. (2024). Enhancing cyber-physical systems resilience: Adaptive self-healing security using long short-term memory networks. *2024 International Conference on Advances in Computing, Communication and Applied Informatics* (ACCAI). https://doi.org/10.1109/ACCAI61061.2024.10602467

Arulkumaran, K., Deisenroth, M. P., Brundage, M., & Bharath, A. A. (2017). Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine, 34*(6), 26–38. https://doi.org/10.1109/MSP.2017.2743240

Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile* (NIST AI 600-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.600-1

Baral, S., Saha, S., & Haque, A. (2025). Autonomous cyber incident response using reasoning and action. *2025 International Wireless Communications and Mobile Computing* (IWCMC). https://doi.org/10.1109/IWCMC65282.2025.11059476

Bartock, M., Cichonski, J., Souppaya, M., Smith, M., Witte, G., & Scarfone, K. (2016). *Guide for cybersecurity event recovery* (NIST Special Publication 800-184). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-184

Boshnakov, V., Wrona, K., & Baser, C. E. (2025). Next-generation AI-enabled incident responses for military cloud networks. *2025 IEEE Military Communications Conference* (MILCOM), 1620-1626. https://doi.org/10.1109/MILCOM64451.2025.11310692

Burbano, L., Sasahara, H., & Cardenas, A. A. (2025). Steerability of autonomous cyber-defense agents by meta-attackers. *2025 IEEE Conference on Artificial Intelligence* (CAI). https://doi.org/10.1109/CAI64502.2025.00194

Carrera-Rivera, A., Ochoa, W., Larrinaga, F., & Lasa, G. (2022). How-to conduct a systematic literature review: A quick guide for computer science research. *MethodsX, 9*, 101895. https://doi.org/10.1016/j.mex.2022.101895

Castro, S. R., Campbell, R., Lau, N., Villalobos, O., Duan, J., & Cardenas, A. A. (2025). Large language models are autonomous cyber defenders. *2025 IEEE Conference on Artificial Intelligence* (CAI). https://doi.org/10.1109/CAI64502.2025.00195

Cho, J.-H., Sharma, D. P., Alavizadeh, H., Yoon, S., Ben-Asher, N., Moore, T. J., Kim, D. S., Lim, H., & Nelson, F. F. (2020). Toward proactive, adaptive defense: A survey on moving target defense. *IEEE Communications Surveys & Tutorials, 22*(1), 709–745. https://doi.org/10.1109/COMST.2019.2963791

d'Ambrosio, N., Lista, C., Perrone, G., & Romano, S. P. (2025). SMASH: An SDN-MTD framework for efficient honeypot deployment and insider threat mitigation. *Computer Networks, 269*, 111327. https://doi.org/10.1016/j.comnet.2025.111327

Du, Y., Prébot, B., Malloy, T., Fang, F., & Gonzalez, C. (2025). Experimental evaluation of cognitive agents for collaboration in human-autonomy cyber defense teams. *Computers in Human Behavior: Artificial Humans, 4*, 100148. https://doi.org/10.1016/j.chbah.2025.100148

Hasanov, I., Virtanen, S., Hakkala, A., & Isoaho, J. (2024). Application of Large Language Models in Cybersecurity: A Systematic Literature Review. *IEEE Access, 12*, 176751–176778. https://doi.org/10.1109/ACCESS.2024.3505983

Hausken, K., Welburn, J. W., & Zhuang, J. (2025). A review of game theory and risk and reliability analysis in infrastructures and networks. *Reliability Engineering & System Safety, 261*, 111123. https://doi.org/10.1016/j.ress.2025.111123

Iturbe, E., Rego, A., Llorente-Vazquez, O., Rios, E., Dalamagkas, C., Merkouris, D., & Toledo, N. (2026). Reinforcement learning in action: Powering intelligent intrusion responses to advanced cyber threats in realistic scenarios. *Expert Systems with Applications, 296*, 129168. https://doi.org/10.1016/j.eswa.2025.129168

Jaber, A., Endregard, M., Mancini, F., & Grov, G. (2025). Mission-aware cyber incident response generation using reinforcement learning. *2025 International Conference on Military Communication and Information Systems* (ICMCIS). https://doi.org/10.1109/ICMCIS64378.2025.11047892

Jafar, M. T., Yang, L.-X., & Li, G. (2025). An innovative practical roadmap for optimal control strategies in malware propagation through the integration of RL with MPC. *Computers & Security, 148*, 104186. https://doi.org/10.1016/j.cose.2024.104186

Kephart, J. O., & Chess, D. M. (2003). The vision of autonomic computing. *Computer, 36*(1), 41–50. https://doi.org/10.1109/MC.2003.1160055

Koo, R., & Toueg, S. (1987). Checkpointing and rollback-recovery for distributed systems. *IEEE Transactions on Software Engineering, SE-13*(1), 23–31. https://doi.org/10.1109/TSE.1987.232562

Kreutz, D., Ramos, F. M. V., Veríssimo, P. E., Rothenberg, C. E., Azodolmolky, S., & Uhlig, S. (2015). Software-defined networking: A comprehensive survey. *Proceedings of the IEEE, 103*(1), 14–76. https://doi.org/10.1109/JPROC.2014.2371999

Lee, E. A. (2008). Cyber physical systems: Design challenges. *2008 11th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing (ISORC)*, 363–369. https://doi.org/10.1109/ISORC.2008.25

Loevenich, J., Adler, E., Hürten, T., & Lopes, R. R. F. (2025). Design and evaluation of an autonomous cyber defence agent using DRL and an augmented LLM. *Computer Networks, 262*, 111162. https://doi.org/10.1016/j.comnet.2025.111162

Manadhata, P. K., & Wing, J. M. (2011). An attack surface metric. *IEEE Transactions on Software Engineering, 37*(3), 371–386. https://doi.org/10.1109/TSE.2010.60

Mejia-Ruiz, G. E., Marasini, G., Qu, Z., Kundu, S., & Pushpak, S. (2025). Cybersecurity challenges in power networks with distributed energy resources: A comprehensive survey. *Renewable and Sustainable Energy Reviews, 224*, 116100. https://doi.org/10.1016/j.rser.2025.116100

Naeem, M., Ozuem, W., Howell, K., & Ranfagni, S. (2023). A step-by-step process of thematic analysis to develop a conceptual model in qualitative research. *International Journal of Qualitative Methods, 22*, 1–18. https://doi.org/10.1177/16094069231205789

Nelson, A., Rekhi, S., Scarfone, K., & Souppaya, M. (2025). *Incident response recommendations and considerations for cybersecurity risk management: A CSF 2.0 community profile* (NIST Special Publication 800-61 Rev. 3). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-61r3

Paduraru, C., Dumitru, B., & Stefanescu, A. (2025). Automated generation of cybersecurity response playbooks via large language models. *Procedia Computer Science, 270*, 2987–2996. https://doi.org/10.1016/j.procs.2025.09.423

Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., … Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ, 372*, Article n71. https://doi.org/10.1136/bmj.n71

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans, 30(3)*, 286–297. https://doi.org/10.1109/3468.844354

Petersen, K., Vakkalanka, S., & Kuzniarz, L. (2015). Guidelines for conducting systematic mapping studies in software engineering: An update. *Information and Software Technology, 64*, 1–18. https://doi.org/10.1016/j.infsof.2015.03.007

Ren, S., Jin, J., Niu, G., & Liu, Y. (2025). ARCS: Adaptive reinforcement learning framework for automated cybersecurity incident response strategy optimization. *Applied Sciences, 15*, 951. https://doi.org/10.3390/app15020951

Rinaldi, S. M., Peerenboom, J. P., & Kelly, T. K. (2001). Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Systems Magazine, 21*(6), 11–25. https://doi.org/10.1109/37.969131

Ross, R. S., Pillitteri, V. Y., Graubart, R., Bodeau, D., & McQuaid, R. (2021). *Developing cyber-resilient systems: A systems security engineering approach* (NIST Special Publication 800-160, Vol. 2, Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-160v2r1

Runeson, P., & Höst, M. (2009). Guidelines for conducting and reporting case study research in software engineering. *Empirical Software Engineering, 14*(2), 131–164. https://doi.org/10.1007/s10664-008-9102-8

Schlette, D., Empl, P., Caselli, M., Schreck, T., & Pernul, G. (2024). Do you play it by the books? A study on incident response playbooks and influencing factors. *2024 IEEE Symposium on Security and Privacy (SP)*, 3625–3643. https://doi.org/10.1109/SP54263.2024.00060

Senarak, C. (2025). Leveraging advanced technologies and strategies for port cyber resilience: Strengthening incident response and recovery. *Transportation Research Interdisciplinary Perspectives, 34*, 101666. https://doi.org/10.1016/j.trip.2025.101666

Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. *2010 IEEE Symposium on Security and Privacy*, 305–316. https://doi.org/10.1109/SP.2010.25

Srinivas, S., Kirk, B., Zendejas, J., Espino, M., Boskovich, M., Bari, A., Dajani, K., & Alzahrani, N. (2025). AI-augmented SOC: A survey of LLMs and agents for security automation. *Journal of Cybersecurity and Privacy, 5*, 95. https://doi.org/10.3390/jcp5040095

United Nations. (n.d.-a). *Goal 9: Industry, innovation and infrastructure*. United Nations Sustainable Development. Retrieved July 29, 2026, from https://www.un.org/sustainabledevelopment/infrastructure-industrialization/

United Nations. (n.d.-b). *Goal 16: Peace, justice and strong institutions*. United Nations Sustainable Development. Retrieved July 29, 2026, from https://www.un.org/sustainabledevelopment/peace-justice/

Wohlin, C., Runeson, P., Höst, M., Ohlsson, M. C., Regnell, B., & Wesslén, A. (2012). *Experimentation in software engineering*. Springer. https://doi.org/10.1007/978-3-642-29044-2

Xi, Z., He, C., Wang, Y., & Zhang, B. (2025). A dynamic SDN-based address hopping model for IoT anonymization. *Computer Modeling in Engineering & Sciences, 144*. https://doi.org/10.32604/cmes.2025.066822

Zhang, J., Li, S., Huang, W., Jing, H., Zhang, Q., & Xia, X. (2025a). Design and computational modeling of an AI-based automated cybersecurity incident response system. *IEEE Access, 13*, 154383–154394. https://doi.org/10.1109/ACCESS.2025.3603975

Zhang, Z., Turnbull, B., Kermanshahi, S. K., Pota, H., Damiani, E., Yeun, C. Y., & Hu, J. (2025b). A survey on resilient microgrid system from cybersecurity perspective. *Applied Soft Computing, 175*, 113088. https://doi.org/10.1016/j.asoc.2025.113088

Appendix A List of Included Literature

Author/s	Publication Name
Alhidaifi et al. (2026)	Cyber Resilience Quantification: A Probabilistic Estimation Model for IT Infrastructure
Alijoyo et al. (2024)	Enhancing Cyber-Physical Systems Resilience: Adaptive Self-Healing Security Using Long Short-Term Memory Networks
Baral et al. (2025)	Autonomous Cyber Incident Response Using Reasoning and Action
Boshnakov et al. (2025)	Next-Generation AI-Enabled Incident Responses for Military Cloud Networks
Burbano et al. (2025)	Steerability of Autonomous Cyber-Defense Agents by Meta-Attackers
Castro et al. (2025)	Large Language Models Are Autonomous Cyber Defenders
d'Ambrosio et al. (2025)	SMASH: An SDN-MTD Framework for Efficient Honeypot Deployment and Insider Threat Mitigation
Du et al. (2025)	Experimental Evaluation of Cognitive Agents for Collaboration in Human-Autonomy Cyber Defense Teams
Hausken et al. (2025)	A Review of Game Theory and Risk and Reliability Analysis in Infrastructures and Networks
Iturbe et al. (2026)	Reinforcement Learning in Action: Powering Intelligent Intrusion Responses to Advanced Cyber Threats in Realistic Scenarios
Jaber et al. (2025)	Mission-Aware Cyber Incident Response Generation Using Reinforcement Learning
Jafar et al. (2025)	An Innovative Practical Roadmap for Optimal Control Strategies in Malware Propagation Through the Integration of RL With MPC
Loevenich et al. (2025)	Design and Evaluation of an Autonomous Cyber Defence Agent Using DRL and an Augmented LLM
Mejia-Ruiz et al. (2025)	Cybersecurity Challenges in Power Networks With Distributed Energy Resources: A Comprehensive Survey
Paduraru et al. (2025)	Automated Generation of Cybersecurity Response Playbooks via Large Language Models
Ren et al. (2025)	ARCS: Adaptive Reinforcement Learning Framework for Automated Cybersecurity Incident Response Strategy Optimization

Senarak (2025)	Leveraging Advanced Technologies and Strategies for Port Cyber Resilience: Strengthening Incident Response and Recovery
Srinivas et al. (2025)	AI-Augmented SOC: A Survey of LLMs and Agents for Security Automation
Xi et al. (2025)	A Dynamic SDN-Based Address Hopping Model for IoT Anonymization
Zhang et al. (2025a)	Design and Computational Modeling of an AI-Based Automated Cybersecurity Incident Response System
Zhang et al. (2025b)	A Survey on Resilient Microgrid System From Cybersecurity Perspective

Appendix B Theme and Subtheme Mapping

Theme	Subtheme	Contributing Publications
Adaptive response and recovery mechanisms	Automated restoration and self-healing	Alijoyo et al. (2024); Boshnakov et al. (2025); Mejia-Ruiz et al. (2025); Zhang et al. (2025b)
	Learning based response optimisation	Du et al. (2025); Iturbe et al. (2026); Jafar et al. (2025); Loevenich et al. (2025); Ren et al. (2025); Zhang et al. (2025a)
	LLM and agent response support	Baral et al. (2025); Boshnakov et al. (2025); Castro et al. (2025); Paduraru et al. (2025); Srinivas et al. (2025)
	Network reconfiguration and deception	Alijoyo et al. (2024); Mejia-Ruiz et al. (2025); Xi et al. (2025); Zhang et al. (2025b); d'Ambrosio et al. (2025)
	Risk and mission-aware management	Alhidaifi et al. (2026); Burbano et al. (2025); Hausken et al. (2025); Jaber et al. (2025); Mejia-Ruiz et al. (2025); Senarak (2025)
Reported resilience functions and outcomes	Containment and threat suppression	Baral et al. (2025); Iturbe et al. (2026); Jaber et al. (2025); Jafar et al. (2025); Xi et al. (2025); d'Ambrosio et al. (2025)
	Decision quality and adaptability	Iturbe et al. (2026); Loevenich et al. (2025); Paduraru et al. (2025); Ren et al. (2025)
	Recovery coordination and verification	Du et al. (2025); Senarak (2025)

Reported resilience functions and outcomes	Response speed and workload reduction	Baral et al. (2025); Castro et al. (2025); Paduraru et al. (2025); Zhang et al. (2025a); d'Ambrosio et al. (2025)
	Service restoration and continuity	Alhidaifi et al. (2026); Boshnakov et al. (2025); Du et al. (2025); Jafar et al. (2025); Zhang et al. (2025a); Zhang et al. (2025b)
Evaluation methods and supporting evidence	Method and baseline comparison evidence	Castro et al. (2025); Iturbe et al. (2026); Loevenich et al. (2025); Paduraru et al. (2025); Ren et al. (2025); Xi et al. (2025); Jafar et al. (2025); Zhang et al. (2025a)
	Emulation, testbed and prototype evidence	Baral et al. (2025); Boshnakov et al. (2025); Iturbe et al. (2026); Xi et al. (2025); Zhang et al. (2025a); d'Ambrosio et al. (2025)
	Evidence limitations and readiness for real-world use	Alhidaifi et al. (2026); Alijoyo et al. (2024); Boshnakov et al. (2025); Iturbe et al. (2026); Jaber et al. (2025); Paduraru et al. (2025); Ren et al. (2025); Senarak (2025); Srinivas et al. (2025); Xi et al. (2025)
	Human, expert and review evidence	Du et al. (2025); Hausken et al. (2025); Senarak (2025); Srinivas et al. (2025)
	Numerical simulation and offline evaluation	Alhidaifi et al. (2026); Jafar et al. (2025); Loevenich et al. (2025); Paduraru et al. (2025); Ren et al. (2025)

Operational limitations and safeguards	Adversarial and automation risks	Boshnakov et al. (2025); Burbano et al. (2025); Castro et al. (2025); Paduraru et al. (2025); d'Ambrosio et al. (2025)
	Data, computing needs, scalability and integration	Hausken et al. (2025); Iturbe et al. (2026); Jafar et al. (2025); Mejia-Ruiz et al. (2025); Ren et al. (2025); Srinivas et al. (2025); Zhang et al. (2025b); d'Ambrosio et al. (2025)
	Applicability across different settings and scenarios	Alhidaifi et al. (2026); Alijoyo et al. (2024); Baral et al. (2025); Castro et al. (2025); Du et al. (2025); Hausken et al. (2025); Jaber et al. (2025); Jafar et al. (2025); Loevenich et al. (2025); Mejia-Ruiz et al. (2025); Ren et al. (2025); Senarak (2025); Zhang et al. (2025a)
	Human oversight, trust and governance	Boshnakov et al. (2025); Du et al. (2025); Iturbe et al. (2026); Paduraru et al. (2025); Ren et al. (2025); Srinivas et al. (2025); Zhang et al. (2025a)
	Security and performance balance	Loevenich et al. (2025); Mejia-Ruiz et al. (2025); Xi et al. (2025)

Note. The table above contains the four themes and 20 subthemes that were produced through the thematic coding. These publications may contribute to more than one subtheme.

Appendix C Thematic Coding Sample

Authors	Text Extract	Codes	Subtheme	Theme
Ren et al. (2025)	Against the traditional rule-based baseline, reported response effectiveness increased from 72.3% to 89.2% and average response time fell from 45.6 to 23.5 seconds.	Traditional baseline comparison	Method and baseline comparison evidence	Evaluation methods and supporting evidence
Castro et al. (2025)	LLM agents hallucinated, misread observations, were sensitive to prompt/action wording and incurred high inference time in an environment designed around RL rewards.	Hallucination and output error risk	Adversarial and automation risks	Operational limitations and safeguards
Baral et al. (2025)	The agent executes a targeted UFW firewall rule through SSH to block a malicious source.	Automated threat containment	Containment and threat suppression	Reported resilience functions and outcomes
Boshnakov et al. (2025)	The implemented pipeline captures forensic snapshots, terminates compromised cloud instances, provisions clean replacements and reroutes traffic.	Automated restoration	Automated restoration and self-healing	Adaptive response and recovery mechanisms

Note. This table presents a sample of the thematic coding used during the analysis. It shows how the findings from the included publications were sorted into codes and then organised into subthemes and broader themes.

Appendix D Quality Assessment Results

Nr.	Publication Name	Q1	Q2	Q3	Q4	Q5	Total
1.	Cyber resilience quantification: A probabilistic estimation model for IT infrastructure	1	1	1	1	1	**5**
2.	Enhancing cyber-physical systems resilience: Adaptive self-healing security using long short-term memory networks	1	0.5	0.5	0.5	0.5	**3**
3.	Autonomous cyber incident response using reasoning and action	1	1	0.5	0.5	0.5	**3.5**
4.	Next-generation AI-enabled incident responses for military cloud networks	1	1	0.5	0.5	1	**4**
5.	Steerability of autonomous cyber-defense agents by meta-attackers	1	1	1	1	0.5	**4.5**
6.	Large language models are autonomous cyber defenders	1	1	0.5	0.5	1	**4**
7.	SMASH: An SDN-MTD framework for efficient honeypot deployment and insider threat mitigation	1	1	1	1	1	**5**
8.	Experimental evaluation of cognitive agents for collaboration in human-autonomy cyber defense teams	1	1	1	1	1	**5**
9.	Adaptive AI-driven cybersecurity framework for threat detection and response in networks	1	0.5	0	0	0.5	**2**
10.	A review of game theory and risk and reliability analysis in infrastructures and networks	1	0.5	0.5	1	0.5	**3.5**
11.	Reinforcement learning in action: Powering intelligent intrusion responses to advanced cyber threats in realistic scenarios	1	1	1	1	1	**5**
12.	Mission-aware cyber incident response generation using reinforcement learning	1	0.5	0.5	0.5	1	**3.5**
13.	An innovative practical roadmap for optimal control strategies in malware propagation through the integration of RL with MPC	1	1	1	1	1	**5**
14.	Design and evaluation of an autonomous cyber defence agent using DRL and an augmented LLM	1	1	1	1	1	**5**

15.	Cybersecurity challenges in power networks with distributed energy resources: A comprehensive survey	1	0.5	0.5	1	0.5	**3.5**
16.	Automated generation of cybersecurity response playbooks via large language models	1	1	1	1	1	**5**
17.	ARCS: Adaptive reinforcement learning framework for automated cybersecurity incident response strategy optimization	1	1	1	1	1	**5**
18.	Leveraging advanced technologies and strategies for port cyber resilience: Strengthening incident response and recovery	1	1	1	1	1	**5**
19.	AI-augmented SOC: A survey of LLMs and agents for security automation	1	1	1	1	1	**5**
20.	A dynamic SDN-based address hopping model for IoT anonymization	1	1	1	1	1	**5**
21.	Design and computational modeling of an AI-based automated cybersecurity incident response system	1	1	1	1	0.5	**4.5**
22.	A survey on resilient microgrid system from cybersecurity perspective	1	0.5	0.5	1	0.5	**3.5**