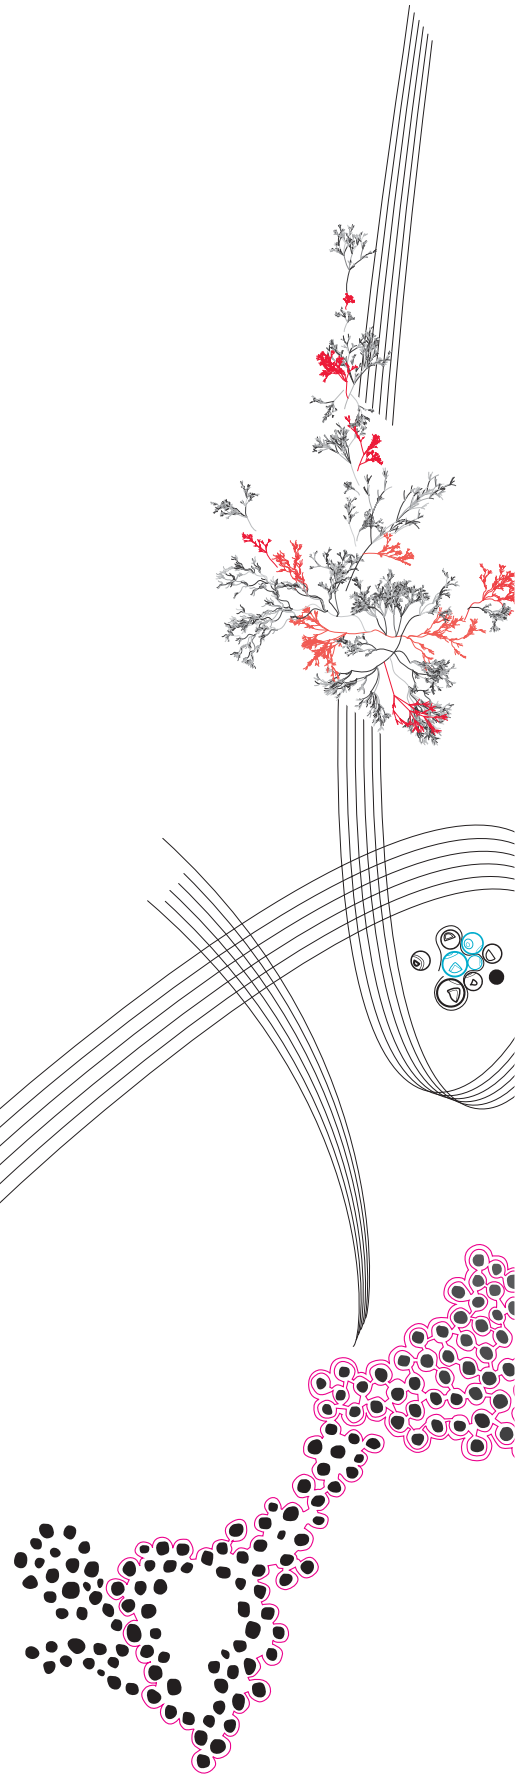

MSc Business Information Technology
Master Thesis

Utilizing a Digital Model to Support Business Continuity Management Processes and Enhance Cyber Resilience

Remo van den Berg

Examination Committee:
Dr. Ing. M. El-hajj,
Dr. A. Abhishta,
T.M. Itäpelto,
Dr. Ir. M.J. van Sinderen

External Supervisor:
M.G.M. Dekker

October 4, 2024

Department of Computer Science
Faculty of Electrical Engineering,
Mathematics and Computer Science,
University of Twente

UNIVERSITY OF TWENTE.

NORTHWAVE
CYBER SECURITY

Abstract

Cyber attacks are considered a big threat to the business continuity of organizations. Preparing and testing systems and people for these situations to take place is challenging due to their disruptive effect on business operations. This makes the preparation for realistic cyber resilience scenarios difficult to achieve. This study shows the possibility of using a Digital Twin-like solution to support and improve system capabilities for business continuity management processes and how it can enhance the cyber resilience of organizations. More specifically, a literature review was conducted on the topic from which the insights initiated a design study to develop a conceptual digital model architecture which can improve and integrate with three business continuity management processes: disaster recovery, disaster recovery test and business impact analysis. We call this conceptual solution the Resilience Digital Model Architecture (RDMA), which we have described using Enterprise Architecture models. By implementing the RDMA, organizations can get more insights into their IT interdependencies and reduce disaster recovery time in a disrupting cyber-attack scenario.

Keywords: Digital Twin, Digital Model, Business Continuity Management, Cyber Resilience, Cyber Security, Rapid Review, Literature Review

Acknowledgements

In writing this section, it has become evident to me that my years of student life have come to pass. A period of personal development, endless fun experiences and academic growth which I will always remember. As a last stretch of this period, it has been a great pleasure and knowledgeable experience to write my thesis about such an interesting topic. This would not have been possible without the guidance and support of my supervisors, family and friends who have helped me along the way. I would like to give a special thanks to these people.

To start, I would like to thank Northwave for allowing me to experience the world of cyber security first-hand from the so-called 'vakidioten' in the field. The knowledge that I have gained, especially in the Cyber Resilience team, has contributed to my professional growth and the thesis project. The professional, enthusiastic and family-like atmosphere at Northwave has made a great impression on me and has been instrumental in shaping my thesis period. A special thanks to Milou Dekker and Pascal Renckens for giving me the opportunity to start the project with the Cyber Resilience team. It was a great pleasure to learn from all members of the team and to experience what being "resilient" actually means. Thanks to Milou and Jesse for supervising my thesis and making me feel at home at Northwave. Your insights have helped me to make sense of the field and contributed greatly to the end results of my research.

Secondly, I would like to thank my supervisors at the University of Twente. To Mohammed El-hajj: when we first met, you already showed your excitement towards the project and believed in the opportunities of the research. You have kept this attitude to this day and helped me keep a positive mindset during the challenges along the way. It was my great pleasure to work with you. Thanks to Taru Itäpelto for being my daily supervisor and helping me with the questions and concerns I had throughout the project. You assessed my decisions critically which has kept me on top of my game. Last, thanks to Marten van Sinderen for being critical and having an outside view on the thesis. You were a vital addition to the correctness of the models and the thesis structure. Each one of you showed such a kind attitude, and made a great deal of time investment into the project, for which I am deeply grateful.

Third, my time as a student couldn't have been more fun without my friends, making this such a memorable period of my life. And finally, a big thank you to my family, who have always been there for me during my student life and supported me every step along the way. You laid the foundation for critical thinking and a positive attitude, even recently through the difficult times as a family.

Enjoy the read!

Remo van den Berg

Utrecht

September, 2024

Contents

Abbreviations

AI Artificial Intelligence.

ART Advanced Red Teaming.

BC Business Continuity.

BCM Business Continuity Management.

BCP Business Continuity Plan.

BIA Business Impact Analysis.

CCI Cyber Critical Infrastructure.

CERT Computer Emergency Response Team.

CM Crisis Management.

CPS Cyber-Physical System.

CR Cyber Resilience.

CTI Cyber Threat Intelligence.

DM Digital Model.

DS Digital Shadow.

DT Digital Twin.

EA Enterprise Architecture.

ICS Industrial Control Systems.

IIoT Industrial Internet of Things.

IoT Internet of Things.

IR Incident Response.

ISO International Organization for Standardization.

IT Information Technology.

NIST National Intitute of Standards and Technology.

OT Operational Technology.

RR Rapid Review.

SOAR Security Orchestration, Automation and Response.

SSI Semi-Structured Interview.

Chapter 1

Introduction

With the rise of Industry 4.0, new technologies enable increased automation, efficiency, and productivity in manufacturing operations [70]. The interconnectedness of cyber-physical systems continues to create new opportunities for businesses and industries but also increases the threat landscape of organizations leading to increased vulnerabilities to cyber-attacks [70]. As organizations embrace transformative technologies and digital advancements, they must also become more cyber-resilient and prepare for incidents to occur. One of the prime technologies Industry 4.0 offers is the Digital Twin (DT). Consequently, research into this field has expanded greatly in recent years. In this study, we want to explore how this technology can enhance cyber resilience in the context of business continuity management.

Preparation for and recovery from cyber incidents have always been core topics within cybersecurity. There is no assurance that organizations will not suffer major cyber attacks, as the field of information systems is constantly evolving and vulnerabilities and exploitations are continuously being discovered. By preparing for cyber incidents, organizations can minimize the impact of these incidents and recover more effectively. Business Continuity Management (BCM) governs the steps that need to be taken to ensure resilience within an organization and especially focuses on minimizing the impact of an adverse situation by preparing for adverse situations to happen. BCM assures this by developing a Business Continuity Plan (BCP) as a preparation method that states procedures, roles, responsibilities, recovery methods, and tasks that must be taken after an incident or discontinuity. Within cyberspace, BCM enhances cyber resilience (CR) to deal with cyber attacks. As the number of malware attacks increases across various fields [64], BCM is starting to receive more attention in the context of cybersecurity. The new European Union guidelines for countermeasures against cyber threats, the NIS 2 directive [20], also emphasizes the importance of business continuity and incident handling as minimal countermeasures against potential threats [20]. Organizations are constantly looking for ways to improve their CR with new technologies.

One such technology that has been at the centre of attention is the Digital Twin (DT). This is one of the prime technologies Industry 4.0 offers. DTs and comparable solutions offer real-time insights into the behaviour and performance of physical objects or processes and have been popularized in recent years. DTs are seen as virtual representations of real-world counterparts, which can be an object, system, or process [11]. This technology reaches across a broad spectrum of industries, including manufacturing, healthcare, aerospace, automotive, networking, and education. As businesses implement this technology, they are not only confronted with the associated cyber threats due to the increased threat landscape but can also explore the potential of leveraging the DT to enhance their cybersecurity [34].

To create a distinction between types of DTs, Kritzinger et al. [42] classifies three types. A fully operational DT is defined by a bi-directional automatic data flow between the system of interest and the digital virtualization environment [39]. A Digital Shadow (DS) is only connected through a one-way automatic information flow from the system of interest to the digital virtualization environment [42]. Lastly, a Digital Model DM does not have any automatic data flow between the systems, and can also be seen as a simulation environment [42]. This research will explore the capabilities of the DT, DS and DM in the context of BCM and CR.

Northwave Cyber Security (hereafter: Northwave) is a company that aims to aid medium to large-size organizations with intelligent cybersecurity services to protect them against malicious actors, including CR services. The CR services guide companies to evaluate their cyber risks, implement plans in case of an incident and provide various exercises based on the recovery plans to make employees more skilled in dealing with adverse situations. The team offers a wide variety of exercises and tests from plan walkthroughs to semi-live exercises. Northwave also maintains a Computer Emergency Response Team (CERT) that provides round-the-clock support for European customers who experience a cyber crisis.

In this study, we initially aimed to understand how a DT, DM, and specifically a DS can be integrated with BCM processes to enhance cyber resilience. This motivated us to conduct a literature review on the key contributions that have been made in literature on this topic. Note that the review will focus on the DS specifically, as a DT poses security vulnerabilities itself [34] due to the direct ability to adjust the real-world system through the DT. This is something that Northwave finds important to consider. As we want to mitigate these concerns, the DS can be considered safer as it doesn't need a direct ability to adjust the real-world system. However, as these classifications are currently not adopted by all researchers as later will be explained in chapter 4.1, all DT-related technologies will also be considered to enhance the contribution of this research. This can be seen as the first phase of the study. In the analysis of the results of this review, some research gaps were found. These gaps were then used for the second phase of this study. It also became clear from the literature review results that the DM has the best potential to improve BCM processes. In the second phase, one research gap was chosen to investigate in more depth. The chosen gap indicated that there were no studies found in the literature review that showed the integration of one DM with multiple BCM processes. While taking stakeholder goals into account and considering the time constraints of the researchers, we utilized the identified research gap to initiate another research involving the integration of a DM with three specific BCM processes instead of BCM as a whole. These three processes are the business impact analysis (BIA), disaster recovery and disaster recovery test. In each of the processes, a problem investigation was executed with the help of Northwave and their knowledge and experience with real-world cases.

1.1 Problem statement

Now that we have discussed the research context, we can describe the problem in more depth. Considering the first phase of the study, we want to investigate the state-of-the-art on how a DT-related technology can enhance CR in the BCM. Related work has been done to explore this work in a more general approach. General opportunities and challenges of the DT for cybersecurity have been identified by Holmes et al. [34]. Allison et al. [1] takes a deeper dive into Incident Response (IR) and has proposed an integration of the DT capabilities into the National Institute for Standards and Technology (NIST) incident response

life-cycle and framework for cyber-physical systems. This covers a great amount of work delving into technical implications that can be considered as a sub-discipline of BCM but does not cover BCM as a whole. As such, the framework does not include risk assessments, business impact analysis, or defining alternatives to critical functions as preparation methods for discontinuities. Another study by Cali et al. [13] examines the benefits of DT for smart cities and energy systems, including benefits towards cybersecurity and CR. Lastly, Faleiro et al. [24] explores the literature on the use and concerns of DT for cybersecurity, and proposes challenges and future directions toward the topic. These articles give a great understanding of the possibilities and challenges that DT technology presents. However, with the rapid development of the technological field and the importance of BCM, the literature lacks a holistic view of the possibilities of utilizing the DT or DS to enhance CR through BCM as a whole. A theoretical framework on the scope of this study and the relation between these concepts can be found in section 2.5. Considering the research gap, the first phase of this study aims to understand the current literature on the possibilities of using a DS to enhance CR in the context of BCM with a Rapid Review (RR) and additive Semi-Structured Interviews (SSI).

Now transitioning to the second phase of the thesis, we want to explore the opportunity to improve BCM processes with a single DM design. In this design science study, we chose three specific processes as mentioned before. These three specific BCM processes were chosen due to the insights from Northwave on the possibility of improving certain aspects with a DM. To explore the problems behind these improvement possibilities, we investigated the problem context in more detail. Meaning that we wanted to understand the process in more detail and what the exact problem is. We will explain each of the problems briefly.

1.1.1 Disaster Recovery

The purpose of disaster recovery in the context of cybersecurity is to restore critical IT infrastructure and operations following a major disruption or disaster to ensure business continuity. The primary challenge that organizations face in this process is the time required to set up an isolated environment for the application recovery process. This environment is required to eradicate the threat actor from a specific application while being able to run applications that have already been through this process in the production environment. Reducing the amount of time needed to set up this environment is critical, especially in a ransomware scenario where every hour counts. Organizations struggle with the implementation of repressive measures to speed up the recovery process. Additionally, organizations often lack a clear understanding of which IT systems are needed to enable critical business functions, leading to delays in recovery as this can take time to figure out.

1.1.2 Disaster Recovery Test

In a disaster recovery test, the main goal is to ensure that the recovery of critical business operations aligns with the organizational objectives of this process after a disruptive event. The test can be considered a success or failure, based on the specific objective of the test. In either case, the evaluation can point out improvement steps that can be taken in the future to further improve the capability to deal with disruptions. In a cyber security context, a separate realistic testing environment needs to be created based on script requirements to leave the production environment unaffected. This process is time-consuming and resource-intensive, often resulting in tests that are not fully realistic. The limited resources available mean that only small parts of procedures are tested, which diminishes the effectiveness of

the tests. A more accurate and realistic test could improve the capabilities to reduce the recovery time in a real-world scenario.

1.1.3 IT Inter-Dependencies

The main objective of the business impact analysis is to understand the dependencies of critical business functions and to specify organizational business continuity requirements for these functions. The challenge within the business impact analysis centres around the complexity of identifying dependencies within the IT infrastructure. As many IT systems can be dependent on each other, it is a complex process to accurately assess them which creates potential knowledge gaps in disaster recovery procedures and risk assessments. Organizations often lack a comprehensive database of their IT infrastructure dependencies, making it difficult to understand vulnerabilities and recovery requirements. This lack of understanding can delay recovery times and makes evaluation of risks less accurate.

To try and improve these problem contexts, our objective is to design a conceptual digital model architecture with enterprise architecture models.

1.2 Research Goals

To solve the problems that we just defined, we first need to define the goals of the research. In this section, we take into regard both phases of the research. To understand what has been done in literature, we need to investigate the state-of-the-art in this domain. Next, as mentioned in section 1.1, we will investigate the possibility of designing a generalized digital model to enhance various BCM processes. Specifically, this model will be evaluated for its use as a recovery measure, a business continuity testing environment, and a tool to enable more accurate BIA by assessing IT interdependencies. The following research goals will be addressed:

1. Identify key-trends of digital shadows to bolster business continuity management in the context of cyber security enhancing cyber resilience.
2. Identify and model the key processes involved in cyber recovery, business continuity plan testing and business impact analysis.
3. Identify system requirements for a digital model that facilitates and enhances cyber recovery, business continuity plan testing and business impact analysis.
4. Develop a conceptual digital model architecture that integrates with business continuity management processes.
5. Design a framework for implementing the digital model into disaster recovery, disaster recovery testing and business impact analysis processes.

1.3 Research Questions

From the introduction, which provides motivation and justification for the research, a research question can be composed. The main question that we are aiming to answer is the following.

MQ. *How can a digital model support and improve business continuity management processes to enhance cyber resilience?*

Answering the main research question will require us to answer multiple sub-questions. First, we need to understand what the literature discussed about the use of DT/DS/DM in the context of BCM and research the state-of-the-art on the subject. Studying all three types will broaden our knowledge of the subject. This research question is exploratory in nature and will provide insights into the usefulness of each technology. Initially, the digital shadow was the main aim of the study, as explained in section 1.1 of this chapter. However, a transition from DS to DM was made after the results of this first question were analyzed.

SQ 1. *What are the key trends on how a Digital Shadow can enhance Cyber Resilience in the context of Business Continuity Management for Organizations?*

Next, we are looking to investigate the three specific problems defined in section 1.1 in more detail.

SQ 2. *Why would organizations want to integrate a digital model into business continuity management processes?*
SQ 3. *What are the key processes and components in a disaster recovery, disaster recovery test and business impact analysis?*

Lastly, the following research questions will address the design of the digital model for the problem context and propose an integration into the problem context.

SQ 4. *What are the requirements and key components of a digital model to enhance business continuity management processes?*
SQ 5. *How can a digital model improve and integrate with disaster recovery, disaster recovery test and business impact analysis processes?*

1.4 Structure

The report is structured as follows. Chapter 2 provides a detailed background, covering essential concepts of Cyber Resilience, Digital Twin, Business Continuity Management and Enterprise Architecture. Chapter 3 describes the methodology employed in the research. Chapter 4 presents the literature review findings, highlighting general findings, key trends and research gaps. Chapter 5 dives into the problem investigation, discussing each part of the problem statement in more depth. Chapter 6 introduces the proposed solution, detailing the requirements, design choices, functional models, and architectural models to explain the solution and how it improves the problem context. Chapter 7 validates the proposed solution through expert opinions. Next, chapter 8 offers a discussion on study-related topics and study limitations. Lastly, chapter 9 concludes the research by answering the research questions and giving suggestions for future work.

Chapter 2

Background

Before we can collect, understand, and analyze the articles that we will be studying, we need to understand the context and define the theories and phenomena that are being researched. This chapter will dive into three topics of the Digital Twin, Cyber Resilience, and Business Continuity Management, and end with a conclusion on how these concepts interact.

2.1 Digital Twin, Digital Shadow and Digital Model

The technology artefact we want to understand in this study is the DT. With the rise of Industry 4.0, the concept of a real-time virtual representation of a physical object, system, or process has been popularized. A DT was first described as a virtual representation of a physical object or system, which is used to simulate, predict, and optimize its performance in real-time [39]. The technology can be an enabler for organizations to gain insights, optimize performance, and make informed decisions [39]. This concept was introduced by Michael Grieves and John Vickers of NASA in 2003, and it has since been widely adopted in various industries, including manufacturing, healthcare, and transportation [39].

To get a comprehensive understanding of the DT concept, a systematic literature review was conducted by Jones et al. [39], which identifies 13 characteristics that many papers have adopted. These include *Physical Entity/Twin; Virtual Entity/Twin; Physical Environment; Virtual Environment; State; Realisation; Metrology; Twinning; Twinning Rate; Physical-to-Virtual Connection/Twinning; Virtual-to-Physical Connection/Twinning; Physical Processes; and Virtual Processes*. The rate at which data flows from the physical to the virtual system or the other way around is called the twinning rate, and the fidelity of a DT depicts how realistic and trustworthy a DT is to its real-world counterpart. [39]

In recent years, researchers have tried to decompose the DT to distinguish different types and specify definitions. For example, some papers describe a DT as lacking a connection between virtual and physical components, while other authors do specify this connection [39]. The study by Kritzinger et al. [42] attempted to create a distinction between these specifications and was later recognized among researchers. To break down the DT, the concepts of a Digital Model (DM) and a Digital Shadow (DS) emerged as subsets of the DT based on different levels of integration [42]. Many researchers and practitioners are not using the subsets as their research topic since the classifications have not been fully adopted and presumably because the term 'Digital Twin' is known more publicly. With this in mind, Kritzinger et al. [42] try to classify the DT as follows:

A DM is a digital representation of a physical object that does not have any automated data exchange with the physical object. Changes in the physical object do not directly

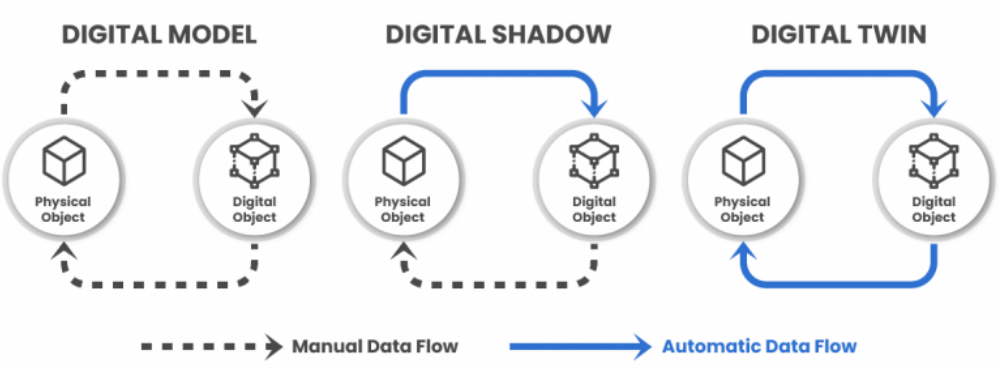

FIGURE 2.1: DT categorisation by Kritzinger et al. [42] visualized by Andrade et al. [2]

affect the DM, and vice versa. The data exchange is done manually. [42]

A DS goes a step further than a DM. It involves an automated one-way data flow from the physical object to the digital object. Changes in the physical object lead to changes in the digital object, but not vice versa. As disconnecting the virtual-to-physical connection can be easily done, the DT can often be adapted to create a DS. [42]

A DT is the most integrated form of these digital representations. It involves fully automated and integrated data flows in both directions between the physical object and the digital object as shown in figure 2.1. Changes in the physical object directly lead to changes in the digital object, and vice versa. The digital object, as shown in 2.1, can also act as a controlling instance of the physical object. Using the term DT for both the overarching concept and the automatic bi-directional form can be very confusing. [42]

Gu et al. suggest using the definition Digital Manager instead of DT for this sub-classification to distinguish these concepts [31]. For this research, the explanation by Kritzinger et al. [42] will be used as it specifies types of DT better.

In the context of cybersecurity, a frequently used reference within literature written by Dietz and Pernul [17] describes how the DT can be used for Industrial Control Systems (ICS) security. Within this study, they tried to identify different modes of security operations a DT can have. They did not include the suggested classification types of Kritzinger et al. [42] to differentiate between types of DT. However, the descriptions imply that the modes correspond to some classifications. They modeled these modes of operations in figure 2.2 and identified three:

The *historical data analytics and optimization* mode uses both real-time and stored historical data for various analytical purposes and optimizations. By using techniques like machine learning and artificial intelligence, it can detect anomalies, predict maintenance needs, and forecast system health. It also serves as a tool for security purposes, detecting potential security threats using historical data analytics. [17]

Simulation mode provides unique security opportunities as it is based on a model of a real-world asset [17]. They allow for repeated tests, compress time intervals, and reveal a system's behaviour under a range of conditions [17]. These simulations can be used to detect potential vulnerabilities and misconfigurations and test new components in a virtual environment [17]. Using the categorization by Kritzinger et al. [42], we can identify that the simulation mode corresponds to a DM as the description does not show an automatic data flow between the twins.

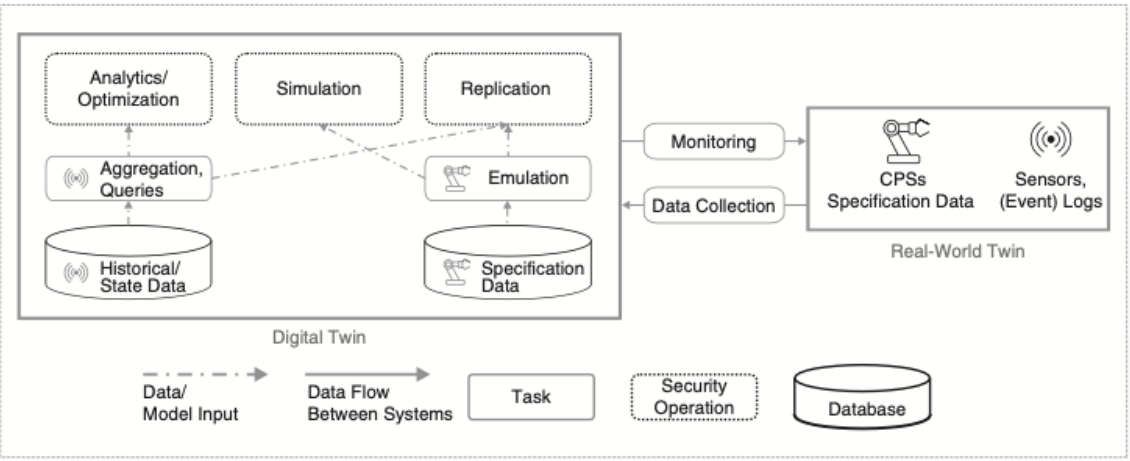

FIGURE 2.2: Digital Twin Security Operation Modes by Dietz and Pernul [17]

Replication mode uses specification data and real-world information to emulate the behaviour and state of the real-world twin [17]. It allows for the detection of threats and attacks [17]. Identifying divergences between the state of the DT's replication mode and the real-world state can reveal possible attacks or system failures [17]. We can classify this mode of operation as a DS or DM, because it can automatically retrieve and use real-time data, but does not send any information into the system automatically. However, this is not always true, as it can also be characterized as a review of historical behaviour, which would classify it as a DM rather than a DS.

2.2 Cyber Resilience

In this study, we aim to examine the application of a DT in enhancing an organization's CR. To identify areas for improvement, it is essential to gain a comprehensive understanding of CR and explore strategies for enhancement.

Resilience is a broad concept used in various fields. In general, resilience refers to the ability of a system, whether it's an individual, a community, or an organization, to withstand disruptions and return to a stable state after a disruption [9]. It's closely related to the capability of an element to adapt to turbulence and discontinuities [9]. However, the precise definition of resilience varies depending on the scope and context of its application [10].

Now transitioning towards the digital world, we apply the same way of thinking for CR. CR, also known as cyber resiliency, is a specific application of the concept of resilience in the field of cybersecurity [54]. CR is defined by the NIST as:

> *"The ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises on systems that use or are enabled by cyber resources"* [54].

This definition needs to be understood and adapted properly, for which frameworks can offer guidance. The concept has gained attention in the last decade, leading to the development of various CR frameworks [60]. NIST Special Publication 800-160 [55] defines the goals and objectives of CR for guidelines in cyber resilience engineering. These were adapted by The MITRE Corporation in their Cyber Resiliency Engineering Framework

[10], which illustrates the goals at the top and objectives on the bottom in figure 2.3. The objectives serve to accomplish the goals. They also described various techniques that can be employed to fulfil the objectives within their framework, though these details are beyond the scope of this research.

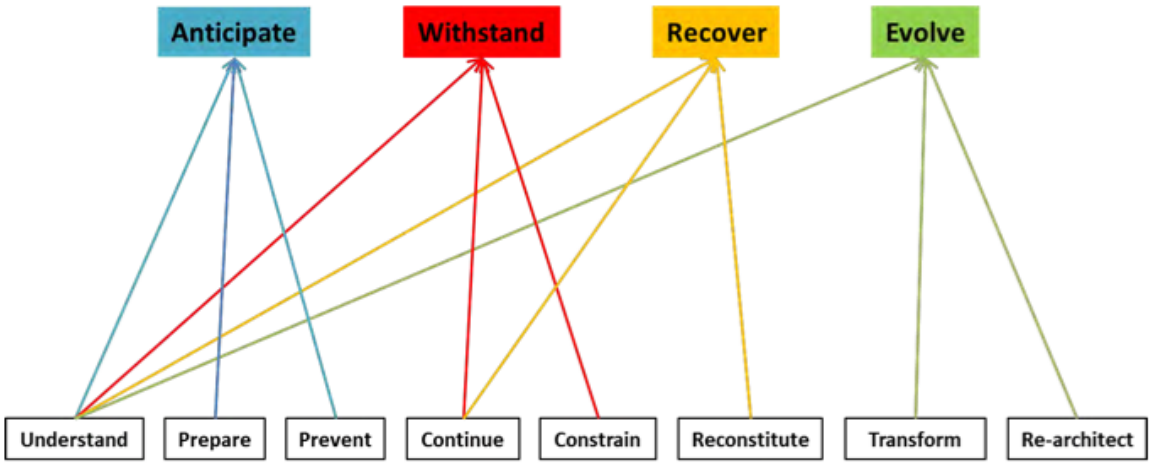

FIGURE 2.3: CR goals (top) and objectives (bottom) [10]

The goals are explained as follows:

Anticipate: This involves maintaining a state of preparedness to prevent compromises to mission/business functions resulting from adversary attacks. The key objectives are to predict, prevent, and prepare for attacks. This requires organizational capabilities to obtain and analyze threat intelligence, monitor the mission environment for adversary activity, and prevent attacks from being executed. Prevention tactics include basic security hygiene, system adjustments to reduce the attack surface, and preparing alternate cyber courses of action conducted by cyber defenders. [10]

Withstand: This goal involves continuing essential business functions despite a successful attack by an adversary. The goal is to *fight through* an attack or maintain functionality amidst adversary actions. This requires maintaining minimal essential capabilities even in a degraded or alternative mode and working to contain the threat and defeat adversary activity. [10]

Recover: This goal is about restoring mission/business functions to the maximum extent possible after a successful attack. The objectives are to determine damages, restore capabilities, and determine the degree of confidence that can be accorded to the restored capabilities. Damage determination involves forensic analysis and coordination with external organizations if necessary. Restoration can involve rolling back to a known acceptable state or recreating capabilities. [10]

Evolve: The final goal involves changing missions/business functions and/or the supporting cyber capabilities to minimize adverse impacts from actual or predicted adversary attacks. This involves transforming existing processes and behaviour and re-architecting systems in response to changes in the threat environment, the system environment, and

the technology environment. [10]

While this framework provides a thorough explanation of the goals and objectives of CR, Northwave adapts a different approach. Even though they agree with the CR definition mentioned above by NIST [55], the Cyber Resilience Team focuses on the ability of organizations to endure and recover from threats effectively. In contrast with the framework by NIST [55], Northwave splits CR into two sub-fields: the capability to withstand a cyber threat and the capability to endure and recover from a cyber incident or crisis (stated as "weerbaarheid" and "veerkracht" in Dutch) [53]. While Northwave as a company provides clients with the capability to withstand difficulties and adverse situations [53], the CR team focuses on enduring and recovering from a cyber incident. In their white paper about CR, they define *enduring and recovering* as the:

> *"ability to face and cope with adversity, adapt to change, recover, learn, and grow from cybersecurity incidents and crises"*

The major difference is that the framework by the MITRE Corporation [10] puts the goals in terms of phases in CR (where for example the preparation for recovery is put into the anticipate phase), while Northwave separates CR into the two subjects of withstanding and endurance and recovery [53].

Tracking back to the core definition of resilience, two main business goals can be recognized. The ability to resist or not be affected by an adverse situation, and the ability to reduce the impact of an adverse situation when it does occur. This logical reasoning resonates better with the separation of the concept by Northwave, rather than the four goals of NIST [55] and the MITRE corporation [10].

2.3 Business Continuity Management

To scope down how CR can be improved, this study will focus on the context of BCM. The use of the term BCM has had a long history of different definitions and implications [32]. The first task one should answer to understand BCM is the implication of Business Continuity (BC), which can be broadly addressed in ISO 22301 [36] as the

> *"capability of an organization to continue the delivery of products and services within acceptable time frames at predefined capacity during a disruption"* [36].

The 2010 study authored by Tammineedi [67], which is often referred to in literature, outlines a structured methodology for designing a BCM system offering a consistent approach for organizations. This paper provides a deeper understanding of the British standard: BS 25999-2 [12], which for a long time was the main standard that was used in the BCM field [32]. Later on, the International Organization for Standardisation (ISO): ISO 22301 [36] replaced the original BS 25999-2 [12], and built upon the fundamentals of that the British Standard Institution started [38]. In the ISO 22301 [36], BCM is defined as the

> *"the holistic management process that identifies potential threats to an organization and the impact those threats, if realized, can cause on business operations, and provides a framework for building organizational resilience with the capability of an effective response that safeguards the interests of key interested*

parties, reputation, brand and value-creating activities" [36].

This definition is very broad and Galaitsi et al. [26] suggest it "seeks to protect against, reduce the likelihood of the occurrence of, prepare for, respond to, and recover from disruptions when they arise". The inclusion or exclusion of certain phases and areas is disputed amongst researchers [26], although this definition or very similar definitions are accepted and used widely among scholars [47]. In the Good Practice Guidelines (GPG) Edition 7.0 about BCM by the Business Continuity Institute (BCI) [68], they exclude the reduction of the likelihood of the occurrence of an incident in their description of BCM.

From the perspective of one of the key stakeholders, Northwave, another definition is given. They define BC as the "ability to survive and restart, or to continue providing services/products (at a minimum acceptable level) in the face of a serious sudden event". Northwave illustrates the definitions in figure 2.4, where it also shows the difference and correlation between IR, Crisis Management (CM), and BC. While these phases are separated in this figure, they should be considered highly correlated.

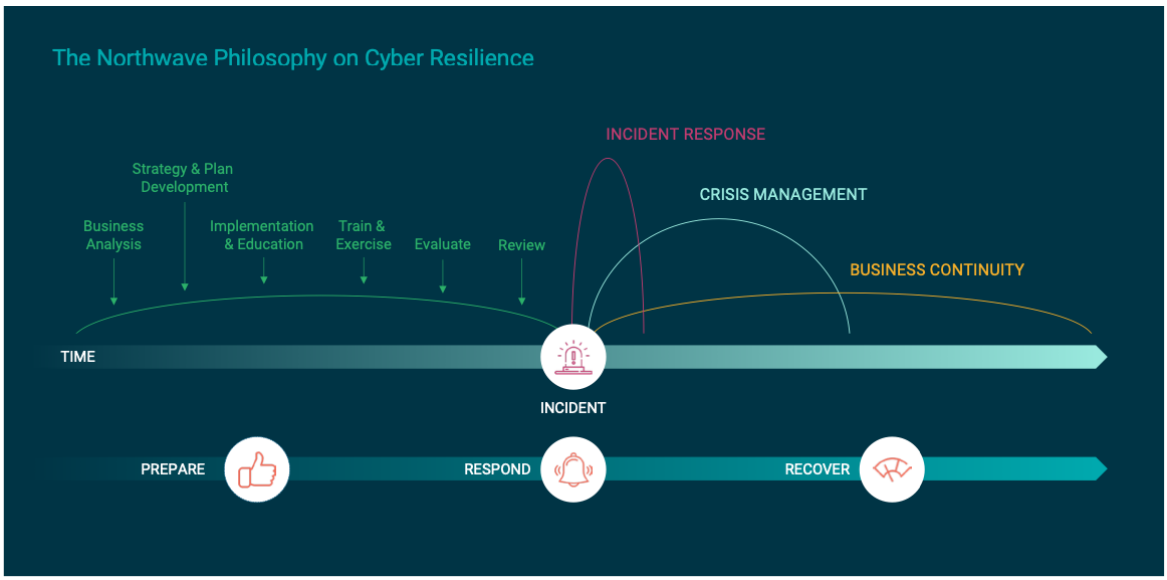

FIGURE 2.4: Three streams during cyber incidents [53]

According to Northwave, BCM is defined as follows:

> *"The organization's ability to have a proper understanding of the impact incidents/crises will have on key processes and to continue the delivery of products or services at an acceptable, predetermined level after a disruptive incident"*

This definition also implies that BCM excludes the reduction of the likeliness of cyber incidents, and is mainly aimed towards the preparation for and recovery from cyber incidents. BCM ensures awareness and resiliency when an incident takes place, providing personnel and systems to be prepared for a speedy recovery of the business practices.

Within BCM literature, various frameworks have been developed that provide procedural guidance for creating plans that prepare for, respond to, manage, and recover a business from any disruptions [25]. They are designed to be adaptable, usable, and beneficial based on an organization's existing asset and process landscape [26]. It is also suggested that researchers mainly focus on the preparation for incidents, rather than the response to them

[47]. However, few frameworks define the components and activities that take place in BCM. The ISO 22301 [36] standard offers a framework that describes the key considerations of BCM that need to be taken care of, but it does not provide a holistic overview of the activities. The paper by Russo et al. (2022) [56] tries to tackle this problem by developing a framework to describe the BCM components that are relevant to be able to develop a BCP. Figure 2.5 presents a model including the components and how they interact [56]. Russo et al. (2024) [57] additionally executed a systematic literature review to develop a comprehensive framework for the multidisciplinary evaluation of organizational maturity on Business Continuity Program Management to further develop and validate their initial model. In the systematic literature review, Russo et al. (2024) [57] present an updated model, including all the main components that were extracted in the review, which is shown in figure 2.6. This model will be used as a basis which will be used to describe where technologies can enable the components.

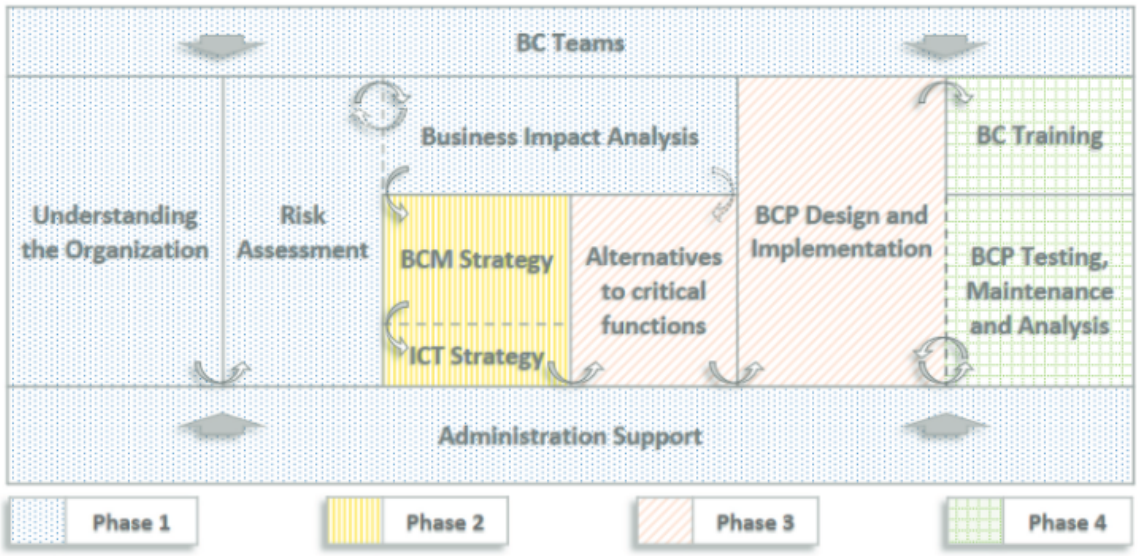

FIGURE 2.5: BCM Components in the proposed BCM methodology by Russo et al. (2022) [56]

The initial phase (that can be found in figure 2.5) involves gaining administrative support and forming BC teams while developing an understanding of the organization's structure, operations, and risks. The second phase identifies BC strategies to enable suitable responses for each business process or activity. The third phase focuses on developing and implementing a BCM response, analyzing alternatives to critical functions, and defining a robust response to manage incidents. The final phase exercises, maintains, reviews, and audits BCM and BCP solutions. This enables the organization to assess the implementation of its strategies and plans and identify areas for improvement. [56] [57]

These phases are based on the framework by Hiles [3], which describes these phases in a model for the BCM life-cycle, shown in figure 2.7. One may notice that the life-cycle includes a step from the fourth phase to the first phase, which is not included in the framework by Russo et al. (2024) [57] in figure 2.6. The BCP-focused perspective of the framework could explain this issue. It emphasizes the development of one BCP, which halts before the creation of a new BCP development cycle. However, the life-cycle by Hiles [3] does not clarify whether significant organizational change should prompt the development of a new BCP, or if it should trigger a reassessment of the existing BCP, incorporating

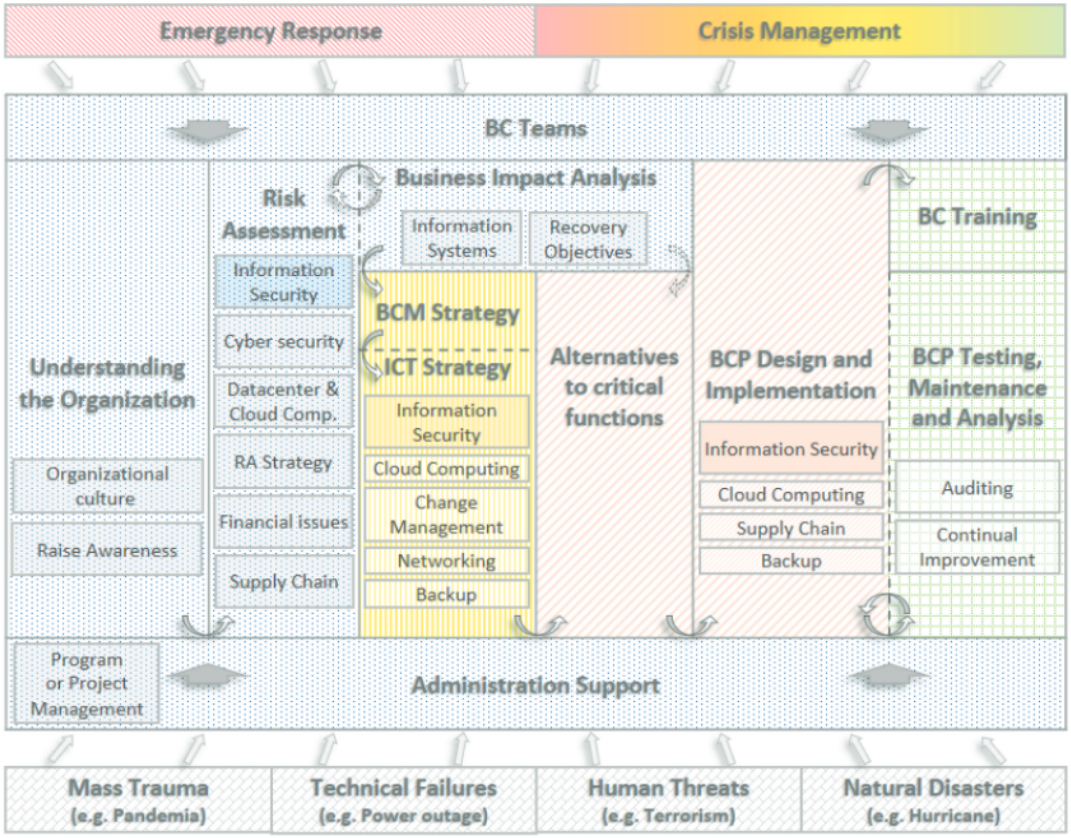

FIGURE 2.6: BCM Components by Russo et al. (2024) [57]

insights from the organizational change. Some might suggest that the framework by Russo et al. (2024) [57] should include a process flow from phase four components back to phase one components. This will not be validated or implemented within the scope of this study.

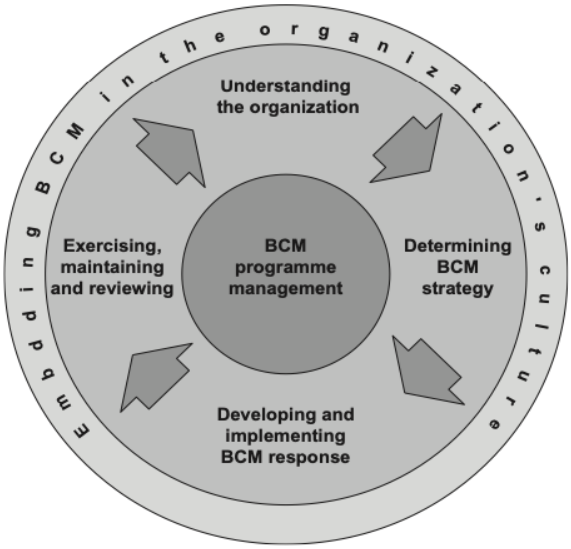

FIGURE 2.7: BCM lifecycle by Hiles [3]

Russo et al. [56][57] explained the BCM components in the following way:

Phase 1

Understanding the organization: This involves a detailed analysis of the operational environment, services, products, and dependencies. By investigating the organization's resources, processes, and interdependencies, an understanding of how the organization functions can be understood. This also shows where potential vulnerabilities may lie, which can be explored in the risk assessment. [56][57]

Risk Assessment: In the systematic literature review by Russo et al. (2024) [57], the risk assessment component is found most often in literature. This process involves identifying potential risks that could interrupt the organization's operations. analyzing these risks in terms of their potential impact, and assessing their likelihood is the most used approach. This component is narrowly related to the business impact analysis component, as they feed each other to gain more insight. [56][57]

Northwave noted that they always conduct a risk analysis and a business impact analysis together, never only one of the two.

Business Impact Analysis: Business impact analysis is a systematic process to identify and prioritize critical business functions. It also determines the potential impact if those functions become unavailable due to a disaster or disruption, providing insights into what resources need to be maintained and within what time frame. [56][57]

Risk assessment and business impact analysis have a clear relationship as stated by Torabi et al. [69]. Their results are jointly used to develop business continuity plans to manage identified risks. In other words, the outputs of business impact analysis, such as critical business functions (activities essential to business operations), dependencies (among systems, people, data, access in buildings), risk appetite (how much risk are you willing to accept), Minimum Business Continuity Objective (how much do you aim to produce during disruption), and Maximum Tolerable Period of Disruption (how long can a disruption take before irrecoverable losses), are used together with the results of risk assessment to prepare the most suitable response plans. These can then be used as guidelines to understand the impact of risks. [69]

Noticeably, Northwave defines different BC requirements which are based on Good Practices Guidelines 7.0 by the Business Continuity Institute [68] than Torabi et al. [69]. Although very similar, two other requirements were added. As depicted in figure 2.8, there are four requirements defined which will be used throughout this research.

- Maximum Tolerable Period of Disruption (MTPD) is the maximum duration of a disruption to production or service after which the organisation considers the damage to be unacceptable.
- Maximum Tolerable Data Loss (MTDL) is the maximum acceptable time that data can be lost before the business is impacted.
- Recovery Point Objective (RPO) is the amount of data which, in a time period, is acceptable to lose after recovery from an incident. This should be shorter than MTDL to ensure it is not reached.
- Recovery Time Objective (RTO) is the time period after a disaster in which production, services or activity must be resumed, or resources must be restored. This should be shorter than MTPD to ensure it is not reached.

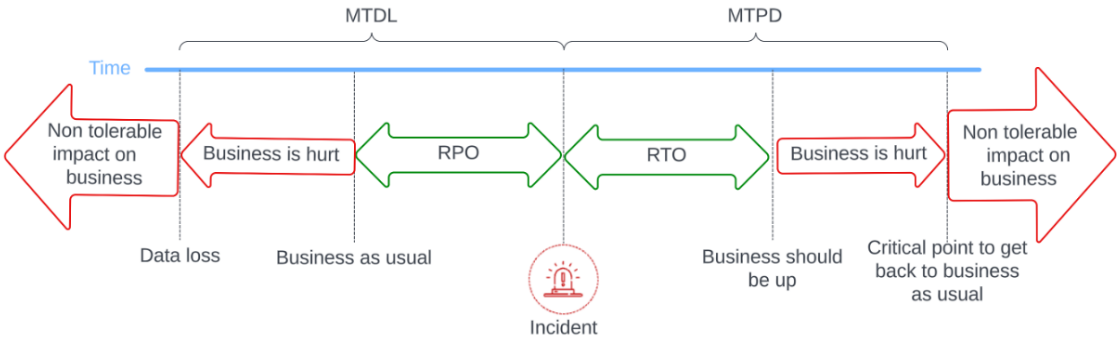

FIGURE 2.8: Business Continuity Requirements by Northwave

Phase 2

BCM Strategy: The BCM strategy is defined in many ways. Mostly, it includes defining a comprehensive set of procedures designed to protect critical business resources. This strategy is informed by the risk assessment and business impact analysis, and it outlines which kind of procedures should be set up to enable an organization to respond to and recover from disruptions. A strategy can outline what kind of plans need to be made like the Crisis Management Plan, Disaster Recovery Plan, or BCP. The activity also constructs a high-level recovery and business continuity strategy that guides what kind of IT strategy and alternative to critical function will be used to minimize the impact of a cyber incident. [56][57]

Although not explicitly described in either article by Russo et al. [56][57], we would also expect this strategy to provide guidelines towards long-term improvements for alternatives to critical functions, how often the BCP should be updated, and an exercising strategy.

ICT Strategy: Given the importance of IT in organizations, the IT Strategy is an integral part of the BCM plan. It ensures that technological development and digitization are incorporated into the organization's BCM strategy. It focuses particularly on maintaining the availability and integrity of critical ICT systems and data. [56][57]

Phase 3

Alternatives to critical functions: This involves identifying and implementing alternative solutions to ensure the continuity of critical business processes in the event of a disruption. This could include designing optional secondary processes, selecting substitute technologies, or outsourcing certain functions. [56][57]

BCP Design and Implementation: This stage involves drafting a detailed specification of the BCP. The BCP contains all the steps and procedures the organization needs to follow to recover from a disruption. Scenarios of unexpected situations, whether incidents, disasters, or crises, should be described in the BCP. With these scenarios, the initiation, contingency actions of teams, and use of technological systems should be predefined. These unexpected situations should only be handled when the assessed cost of the risk is higher than the potential cost it will take to handle this situation. Furthermore, a BCP should be simple, clear, unambiguous, and comprehensive. [56][57]

Northwave mentions that the implementation should not be part of this component, as they argue a BCP is only actually implemented when it is tested and exercised. They define the implementation as a finalized status of the BCP, where the plan should be designed, exercised, and adapted.

Phase 4

BC Training: This component involves developing and implementing awareness and skill-building programs to ensure that all staff members are prepared to execute the BCP when required. This component focuses more on the capabilities of a staff member to execute procedures that are required in a BCP, rather than exercising with a BCP or testing whether certain pre-defined goals can be acquired (like a certain recovery time in a scenario). [56][57]

Northwave also defines resilience training as the "training of certain skills and behaviour".

BCP Testing, Maintenance, and Analysis: This is an ongoing process where the BCP is regularly tested, updated, and improved to ensure its continued effectiveness and relevance to the organization's evolving needs and circumstances. This process is crucial for identifying gaps in the plan making necessary adjustments and raising awareness amongst the BC teams. Successful tests will make BC teams confident in their strategy. The description by Russo et al. also uses the terms exercising, testing, and simulating as ways to conduct this component. [56][57]

Northwave describes exercising as "the simulation of a situation to let people act within the scope of the exercise and learn to orient, work together or experiment" and testing as "the testing of a workflow and/or procedure to evaluate on which components the participants are sufficiently exercised or insufficiently exercised". Examples within the context of Northwave are tabletops, simulations, semi-live exercises, and disaster recovery tests.

2.4 Enterprise Architecture

Enterprise Architecture is a strategic discipline that investigates and models business processes, information systems, and enterprise goals. This involves the principles, methods, and models that guide the design and implementation of enterprise systems. These enable organizations to respond effectively to changes in business environments and technological advancements. Primary processes are the visualization, analysis, and documentation of enterprise structures and processes that facilitate business improvement processes. [62]

To create understandable and insightful representations of these structures and processes, various institutions created EA modelling languages. EA modelling languages are formal notations used to represent enterprise architectures, which can assist with decision-making [62]. These languages provide the tools to map complex systems and their interdependencies.

Enterprise architecture frameworks provide the methodologies and tools necessary to create and maintain these blueprints. These frameworks offer a structured approach to documenting and analyzing an enterprise's architecture, ensuring consistency, clarity, and comprehensiveness. They facilitate communication among stakeholders, support decision-making processes, and promote best practices in architecture management.

One of the most prominent EA modelling languages is ArchiMate. ArchiMate is an

open and independent modeling language for enterprise architectures, developed and maintained by The Open Group. It offers support for a multi-layered approach that addresses different parts of an enterprise. It is designed to provide a clear and comprehensive way to describe, analyze, and visualize the relationships among business domains, based on the TOGAF framework as shown in figure 2.10. In their approach, three main layers and three main aspects are involved. This is the core framework, where ArchiMate includes the following core layers according to the specification [5]:

Business Layer: This layer represents the business processes, organizational units, business functions, and business objects that make up the operational landscape of an organization. It captures how the business operates and how different elements within the business interact with each other.

Application Layer: This layer focuses on the software applications used within the organization. It maps out the various applications and how they support different business processes, and it identifies the interactions between different software applications.

Technology Layer: This layer is concerned with the technology infrastructure of the organization. It details the hardware, communication networks, and software platforms (like databases and operating systems) that underpin the application layer and facilitate the business layer.

Across these layers, the framework defines three core aspects that describe the type of a certain structure according to the specification [5]:

Active Structure Aspect: This describes the structural concept of an organization, including the business actors or application components.

Behaviour Aspect: This captures the behaviour of an organization, such as business processes or application functions.

Passive Structure Aspect: This represents the objects (business, data, or technology) that the behaviour acts upon.

As an extension to the core framework, ArchiMate added a strategy layer, implementation & migration layer, and motivation aspect in their full framework, which is depicted in figure 2.9. In this research, we will mainly be using the complete framework to explain the processes behind the problem contexts and proposed solution architecture.

The ArchiMate framework has many components that explain a certain behaviour. Interaction between these components can also have different meanings, which are explained by various types of arrows indicating a certain relationship. A detailed explanation of each component and relationship can be found in their specification [5]. For the purposes of this report, a summary and overview of these components and relationships can be found in appendix E.

2.5 Conclusion

Concluding, BCM and CR are interrelated concepts, where the DT might offer enhancements that play a critical role in the era of Industry 4.0. CR, the capacity to withstand,

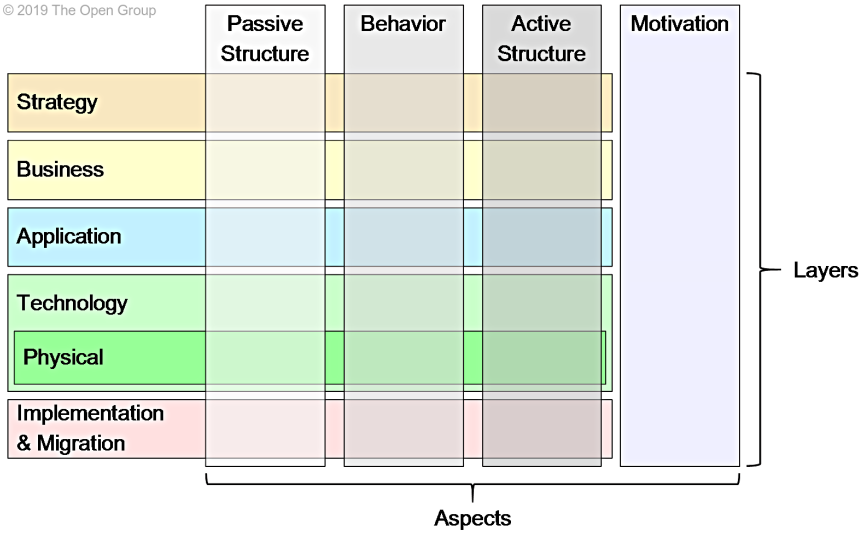

FIGURE 2.9: ArchiMate Full Framework [5]

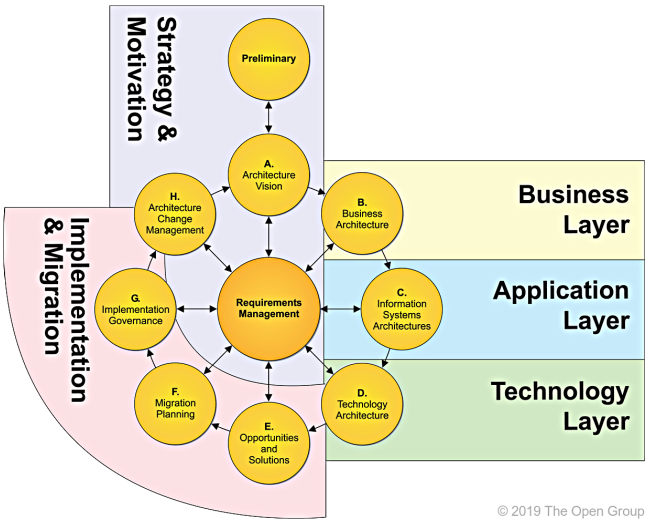

FIGURE 2.10: Correspondence between ArchiMate Language and the TOGAF ADM [5]

endure, recover from, and adapt to cyber threats or incidents can be strengthened using BCM. BCM focuses on managing the minimization of the impact of business interruptions using various activities. It does not necessarily focus on prevention or reduction of the likelihood of incidents but rather focuses on the results in the case of a disruption. To conceptualize these interactions, the theoretical framework shown in figure 2.11 was created. Understanding these concepts and their interactions is important for the decision-making process for effective strategies to deal with potential cyber threats. In this study, we will be focusing on minimizing the impact of a CI. In other words, we want to evaluate how a DT can enhance BCM activities in a cyber context to enhance CR.

In the theoretical framework, we use the explanation of Northwave on the concept of CR as mentioned in section 2.2 to explain how this relates to BCM. Although CR can be explained as a capability to show a certain behaviour, and BCM is aimed towards a process that manages certain abilities, their goals can be aligned. BCM governs the processes that need to be taken that enable the ability to endure and overcome discontinuities like cyber

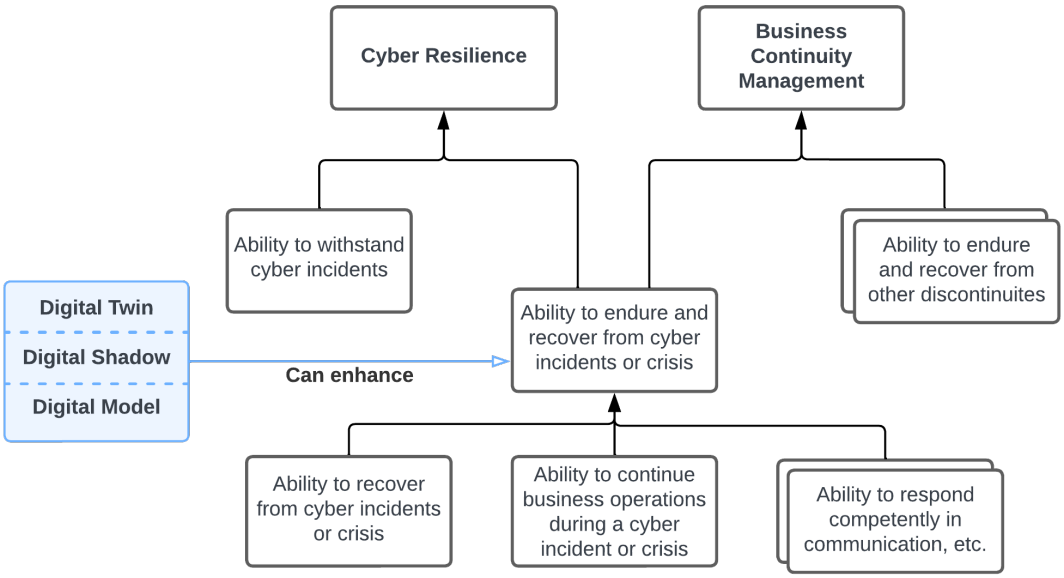

FIGURE 2.11: Theoretical framework of the study context relating BCM, CR, and the DT technologies

incidents or for example fire hazards. Differently, CR is focused solely on cyber threats with goals to withstand, endure and overcome these threats. The overlap is in the ability to endure & recover from a cyber incident or cyber crisis. This ability can then be divided into many sub-abilities as shown in the framework. Finally, as this the main area of interest for this research, we are studying the possibility of using the DT, DS, and DM to enhance the ability to endure and recover from cyber incidents.

Chapter 3

Methodology

Having established the background context of the research topic, we now turn our focus to the research design and methodologies of this study.

3.1 Research Design

As stated in the introduction, the objective of this research is to explore the possibility of using a DM to enhance BCM processes and, consequently, improve CR. To begin this project, we conducted a literature review to explore the key contributions in existing research. This review offered a comprehensive understanding of the research context and highlighted gaps that served as the problem context for the second part of our study. Together, these two parts contribute to addressing the main research question MQ.

In summary, this research is structured in two main parts:

Literature Review: The first phase involves a descriptive study to assess the current state-of-the-art in the field. This review identified research gaps, which led to the next phase of the study.

Design Study: The second phase involves developing a conceptual digital model aimed at improving three specific BCM processes.

The methodologies employed in both phases of this research will be detailed in the following sections.

3.2 Literature Review

In the literature study, a primarily descriptive methodology was adopted. Specifically, the research was conducted in the form of a Rapid Review according to the methodology authored by Cartaxo et al. [14]. The research aims to explore what literature exists about how DTs, DSs and DMs can enhance BCM processes. Instead of a systematic literature review, the study opts to perform a rapid review due to time constraints and the specific context within Northwave. Rapid reviews provide quick evidence to address practical problems efficiently, which aligns well with the objectives and constraints of the researcher. As the research subject is quite new, grey literature was also involved in the rapid review.

3.2.1 Rapid Review

In software engineering, rapid reviews (RRs) serve as quick secondary sources of evidence for practitioners. In contrast with comprehensive systematic literature reviews, RRs are designed to address practical issues promptly and often arise out of immediate needs [14]. In this study, an RR was selected over a systematic literature review because the specific focus of the study is on the practical problems facing a cybersecurity company. The RR methodology was chosen with time limitations in mind so that prompt insights can be provided to effectively address the identified problem. Unlike systematic reviews, which require a careful and time-consuming approach, RRs have a streamlined process that allows for rapidly collecting relevant information and informing decisions [14]. In this regard, RR methodology suits well Cartaxo et al.'s [14] objectives and constraints.

Parsifal, a web-based software application, was used to facilitate the structured execution of the RR. Parsifal is tailored to support researchers engaged in Software Engineering Systematic Literature Reviews [48]. Furthermore, Mendeley served as a library for organizing papers and references, and Notion was used as a comprehensive project management toolkit.

Research Objectives and Research Question

The PICOC (Population, Intervention, Comparison, Outcome, Context) method has been applied to define the research question and search criteria for this research described by Wohlin et al. [74]. As this study specifically examines the DS technique, there is no direct comparison with other intervention methods. Therefore, the comparison field is not applicable in this context. This research aims to evaluate whether organizations can achieve CR by implementing the DS intervention within the context of BCM. It is important to note that this research is conducted within the broader context of BCM, which emphasizes the importance of maintaining uninterrupted operations and mitigating potential disruptions caused by cyber threats. Therefore, the target population for this study consists of '*Organizations*' operating in various industries. The primary focus of this research will be on evaluating the possibilities of the intervention method known as '*Digital Shadow*' in enhancing '*Cyber Resilience*' within a company in the context of '*Business Continuity Management and Cybersecurity*'.

Based on the PICOC methodology, the research question has been formulated as follows:

> *SQ.1: How could a Digital Shadow enhance Cyber Resilience in the context of Business Continuity Management for Organizations?*

Search Strategy

The search strategy aims to cover a broad amount of literature speedily [14]. This means that the search query and database selection will need to be optimized to extract a narrow amount of results while still covering most of the important papers [14]. The risk that some relevant papers will be excluded, will be taken as a consideration and limitation. Next to this, Northwave also recommended searching for grey literature on this subject. Within the cybersecurity industry, companies and institutes can provide valuable and new insights, where research can lag behind or primarily show positive results due to publication bias [74]. The research question covers multidisciplinary fields, and therefore Scopus, and Web Of Science were selected as academic databases. Next to this, Perplexity has been selected as the search engine to search for grey literature.

Perplexity AI is an artificial intelligence (AI) tool that serves as both a chatbot and an intelligent search engine [50]. It employs advanced natural language processing and machine learning techniques to comprehend and address user queries conversationally [50]. Its purpose is to deliver accurate and informative responses by searching the internet for real-time, relevant information [50]. However, it is important to acknowledge that while Perplexity AI can be a valuable tool in the literature search process, it is essential to validate any article citations generated by the AI using reliable sources and to critically review the generated text.

Generative AI tools continue to be developed, and continuous reviews of newly developed AI tools are hard to keep up with. Perplexity AI is one of the newest tools that is not comprehensively reviewed in academic research. The objectives of the tool promise to be very interesting for a grey literature search strategy. Especially in the context of a RR given its timely nature. In this paper, Perplexity AI will be tested as a grey literature review search method and results will be reviewed.

To develop the query, the PICOC methodology was utilized. With the PICOC methodology, the target population, intervention method, and enhanced outcome are used [74]. Synonyms for each of the definitions were added to the search query to enhance results. As mentioned in chapter 2.1, the concept of DS is not used widely yet. This is why the term DT was used in the query, and the results of the DT and DM will be evaluated as well. Various variations of queries have been tested to evaluate the amounts of results and the type of results. After some iterations, a few topics were discovered that were not relevant to the goal of this research. These topics were added to the exclusion criteria, including mobile communication, virtual reality, and virtual environment. These topics are conducted towards their research fields, with mobile communication focusing on communication technologies, virtual reality focusing on incident prevention education, and virtual environment on the technical field of DT development. The following search queries were used within the different libraries.

> **Scopus:** TITLE-ABS-KEY*(("Organization" OR "Company" OR "Enterprise" OR "Entity" OR "Industry") AND ("Digital Twin" OR "Digital Shadow") AND ("Cyber Resilience" OR "Business Continuity" OR "Cyber Security" OR "Cybersecurity" OR "Incident Response") AND NOT ("Mobile communication" OR "Virtual Reality" OR "Virtual Environment"))*
>
> **Web of Science:** TOPIC:*((Organization OR Company OR Enterprise OR Entity OR Industry) AND ("Digital Twin" OR "Digital Shadow") AND ("Cyber Resilience" OR "business continuity" OR "cyber security" OR cybersecurity OR "Incident Response") AND NOT ("Mobile communication" OR "Virtual Reality" OR "Virtual Environment"))*
>
> **Perplexity AI:** *"How could a Digital Twin enhance Cyber Resilience in the context of Business Continuity Management for Organizations?"*

The search queries that were used for Scopus and Web of Science are identical, although their syntax is different due to each library utilizing a different querying language. On the contrary, as Perplexity is a chatbot search engine, we needed to feed it questions to receive proper results. After testing some queries, we discovered that the chatbot can confuse synonyms and produce confusing results. So, to keep the query straightforward, while still including every part of the PICOC methodology, the above-mentioned question was used.

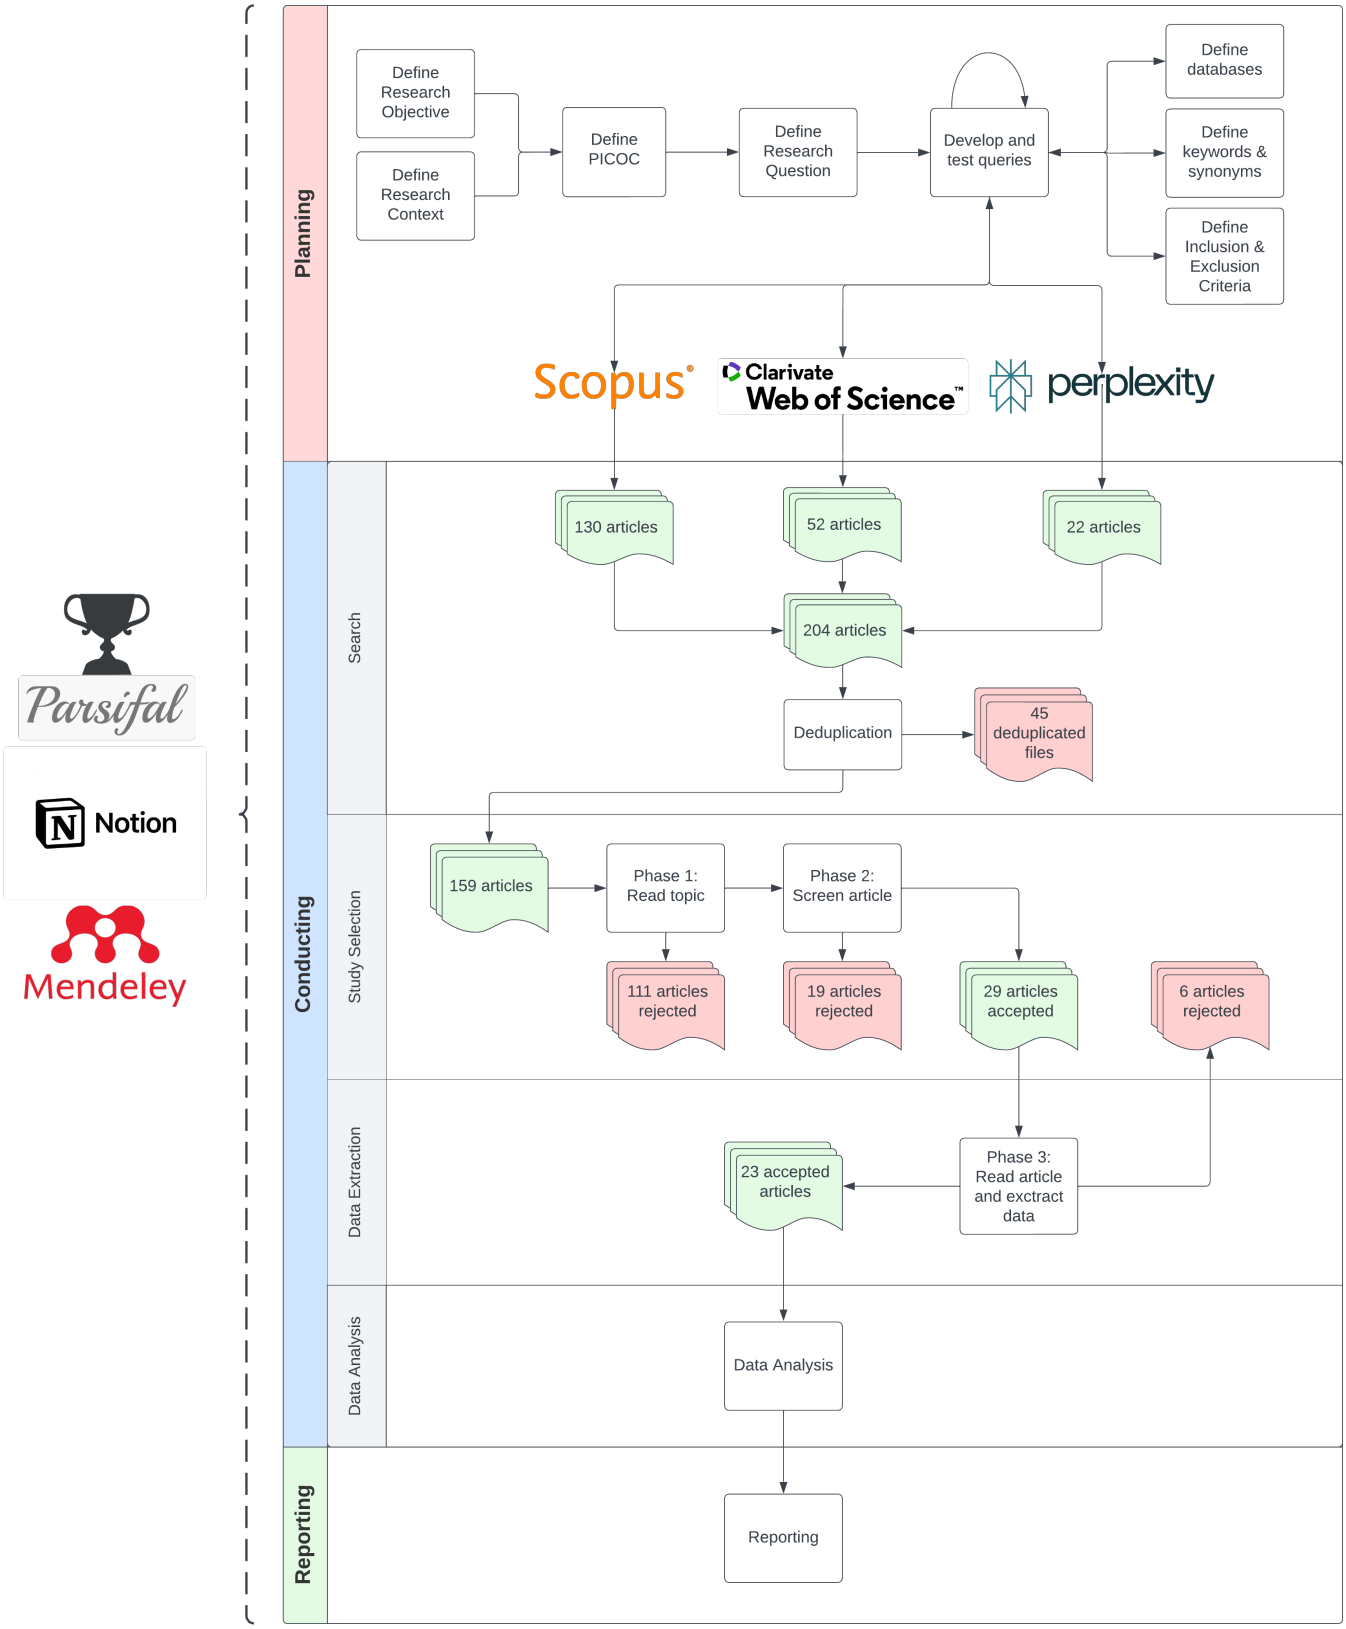

FIGURE 3.1: Rapid Review Methodology process

As a final note, the term "Digital Twin" was used instead of "Digital Shadow", as the DT is the most widely used term to describe the technology.

Inclusion and exclusion criteria were also defined to remove irrelevant papers as mentioned in the methodology [14]. First of all, the papers need to be written in English to be able to be read and the papers need to be published in the last six years (from 2017 onward) to remove outdated material. The concept of the DT is very new, and due to the continuous advancements in technology, the last six years should include the most relevant papers. Papers that fall into the categories of retracted material, editorial material, or are inaccessible to the University of Twente will be excluded. Next to this, exclusion criteria exclude papers that are not in the context of this research and papers that do not use the DT technique or do not enhance CR. The inclusion and exclusion criteria are listed in table 3.1.

Inclusion Criteria	**Exclusion Criteria**
English language studies	Article not in the context of BCM
Studies published in the year 2017 onward	Article not in the context of Cybersecurity
	Article does not propose a DT solution
	Article that are not accessible with the University of Twente library
	Editorial material
	Retracted publications

TABLE 3.1: Inclusion and exclusion criteria used in rapid review

Selection Procedure

The resulting papers from the query that was defined in the previous phase need to be examined to identify if they are relevant and contributing to the research question, or if they can be excluded from the research [14]. Based on the defined inclusion and exclusion criteria concerning the research goals, the relevance of a paper is evaluated and will be included or excluded accordingly.

The selection procedure consists of three substeps to reduce the number of papers on which data extraction will be conducted. Cartaxo [14] mentions that many papers might be lost if the first round only includes reading the title. For this reason, we opted to read the title, abstract, and keywords in the first substep to increase the accuracy of the review. In the second substep, the papers were screened to evaluate their relevance in more detail. In the last round, the paper was read whole, where the data was also extracted to increase time efficiency.

The query resulted in 130 articles in Scopus and 52 articles in Web of Science. Next to this, 22 results were considered from Perplexity AI. After putting them in a list, 45 duplicates were found resulting in 159 articles. These articles were assessed if they were relevant to the research question in two rounds. The first round included reading the abstract, title, and keywords and the second round included a screening of the article. After the first round (reading the title, keywords, and abstract), 111 articles were rejected and 48 were selected. After the second round (screening the article) 19 were rejected and 29 articles were selected to be read more thoroughly to extract data. Lastly, during the next step, data extraction, every article was read completely and 6 more articles were rejected as a consequence.

Data Extraction

Following the selection phase, the data extraction phase involved reviewing and analyzing the selected papers to identify key findings and insights [14]. Data was extracted from the literature, including each study's problem, objectives, solution, challenges, limitations, recommendations, and contributions to various components of BCM.

Next to this, as mentioned in subsection 2.1, not every researcher has adopted the classification by Kritzinger et al. [42]. To ensure we have a correct understanding of the article and its contribution, we want to identify if the actual solution is a DM, DS, or DT. This was a challenging task. As the classification differentiation focuses on the difference between automatic and manual twinning, we needed to define where we extracted this data from. While reading the articles, often conceptual models of systems did not show whether a twinning connection was updated manually or automatically. This is why we classified the types as follows. The classification was adopted *when both virtual-to-physical and physical-to-virtual connections were stated to be automatic or manual in text.* If this was not the case, the article was labelled as *"Not explicitly defined in description by the author(s)"*. Sometimes, the classification method is mentioned in the background knowledge subsection of an article, however not adopted throughout the rest of the article. In these cases, the text should still clearly state the connection types to be classified accordingly.

The next challenge was that some articles did propose a solution to the problem context with a certain type of DT but only proposed it as an idea without a conceptual framework to validate the approach. To differentiate these contributions, we identified whether an article created a conceptual model or implementation of the solution architecture. The articles that did not do this, were often literature reviews that propose a solution without thorough research on the solution at hand.

3.2.2 Semi-Structured Interviews

Since the RR may not have captured all significant papers in the field, we employed Semi-Structured Interviews (SSI) with experts on the topic as an additional methodology to ensure that every relevant paper with high importance is included in this research. Figure 3.2 can be found in how the SSIs relate to the RR, showing that the SSIs were conducted for the RR. These SSIs were conducted following the proposed framework by Kallio et al. [40]. The SSI methodology is a versatile and flexible data collection method that can be used in various contexts [40]. Due to the amount of literature that was found in the RR, and the exploratory nature of this study, the SSI methodology is a sufficient way to extract data from the experts. The goal of this methodology is to find important articles that were not captured by the RR methodology on the topic of this research. this goal doesn't require a rigorous, detailed approach, making SSIs well-suited for the task. [40].

Retrieving and utilizing the previous knowledge

To be able to retrieve appropriate information from the SSIs, background knowledge about the subject is needed to be able to set up adequate questions [40]. Two types of previous knowledge were required before starting this methodology. First of all, a general background context on the processes and technologies of this research topic should be acquired [40]. This was done before this study and described in chapter 2. The second knowledge that should be acquired is the articles that were selected in the RR study [40]. This is why this methodology was only started after both of these knowledge requirements could be answered. This also ensured that there was sufficient knowledge on the topic to develop the SSI guide.

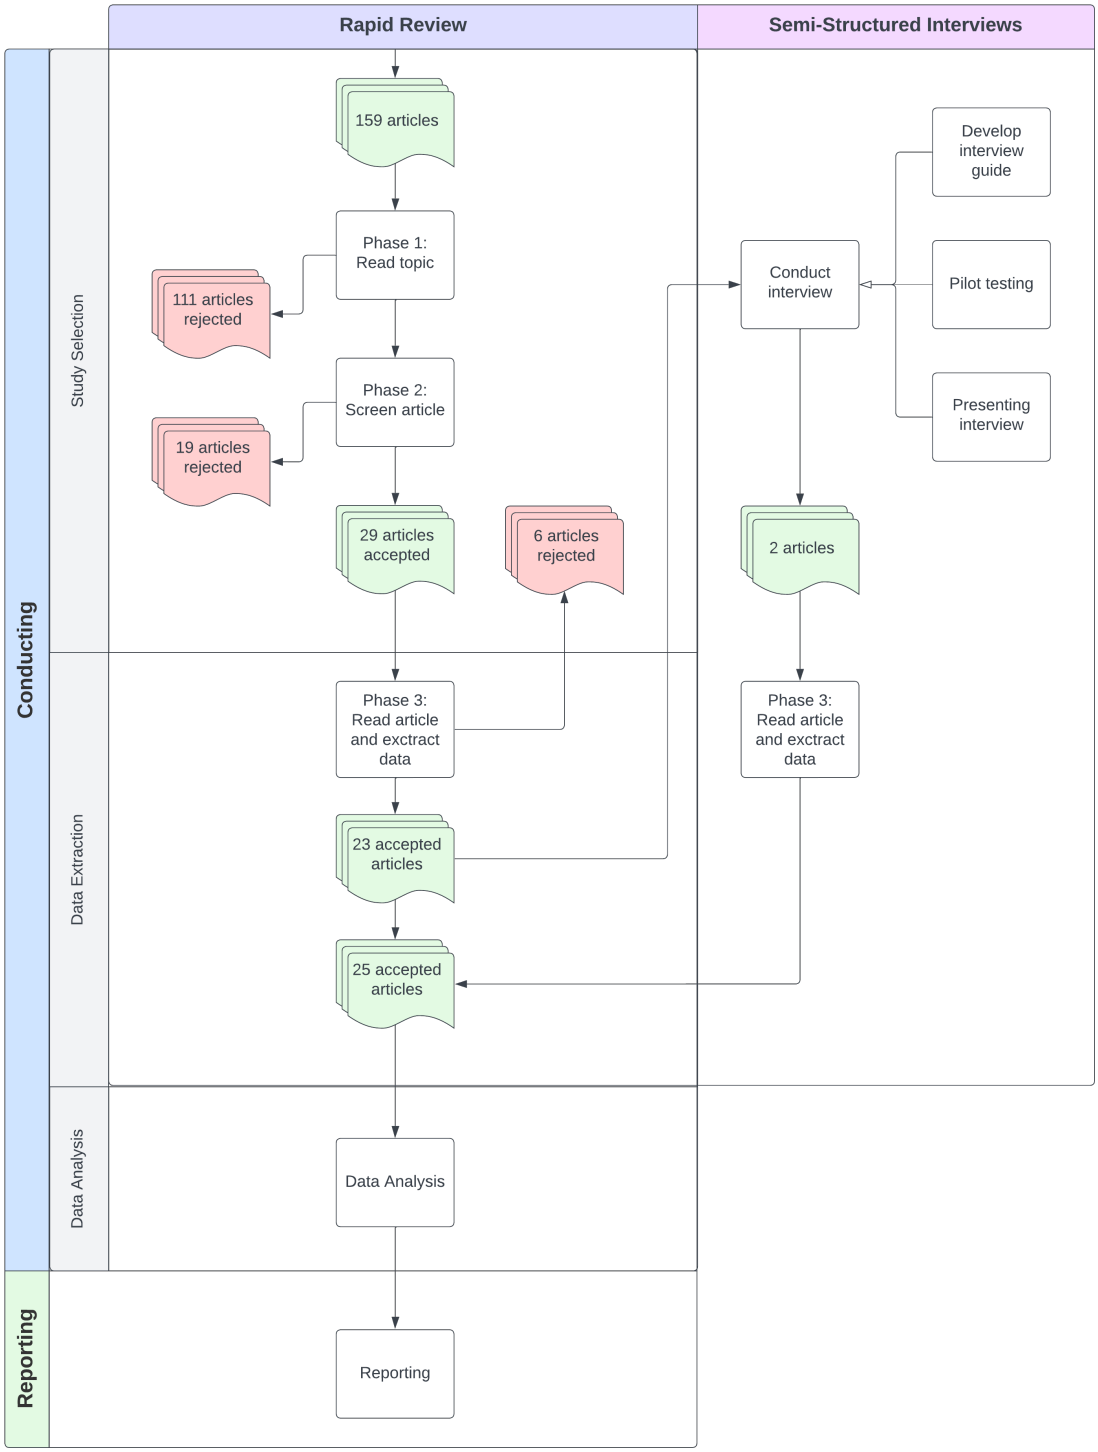

FIGURE 3.2: Literature collection process

Formulating of the preliminary interview guide

The development of the SSI guide was quite straightforward, as the goal of the SSI is clear and simple. One challenge that needed to be addressed, was that the interviewee would need time to read through the articles that were already discovered in the RR and would

need time to search for articles that were not included in the RR. This is why the list of references for the RR-selected articles was sent to the experts beforehand.

Pilot testing

After the development of the interview guide, the guide by Kallio et al. [40] states that the guide needs to be tested before it can presented. The pilot testing of the interview guide was done through the expert testing technique. This was executed as a discussion with a colleague at Northwave. After the discussion, questions were adapted with the feedback accordingly.

Presenting the complete interview guide

To extract data, three experts were interviewed in this methodology. Two of the participants were selected due to their academic experience with the topic at the University of Twente, and one person as a business practitioner was selected at Northwave.

Data collection

First, the articles needed to be analyzed to decide whether they were relevant for this study. If this would be the case, the data extraction method would be started in the same manner as the RR. In total, two articles were extracted and both were accepted for data extraction after analysis.

3.3 Designing a Digital Model Solution Architecture

The second part of this research paper addresses a design problem. Given this context, we need to choose an appropriate research design methodology to add structure and scientific foundation to the research. There are several design science methodologies for information systems. Specifically, the Design Science Research Methodology by Peffers et al. [49] and the design science methodology by Wieringa [73] are popular amongst researchers within the field of information systems. Where Peffers et al. [49] provides a six-step iterative process with a general structure focused on developing and implementing a solution, Wieringa [73] proposes a more thorough four-step guideline with knowledge questions as a basis. As this research is still quite exploratory, the methodology by Wieringa [73] was chosen as it better suits design studies that try to answer exploratory problems. The design science methodology by Wieringa [72] provides three main steps: *problem investigation*, *treatment design* and *treatment validation*. It can be extended into the engineering cycle defined by Wieringa [73] with the inclusion of *treatment implementation* and *implementation evaluation*. The engineering cycle does not apply to this research, as the implementation phase is out of scope.

The first step in the cycle is to understand the problem context. In addition to the extensive data that was acquired from the literature review, we will be gathering additional data in our research from semi-structured interviews with experts at Northwave, informal conversations and studying cases through observations of past projects. Do note that these are different interviews than those used in the literature review (mentioned in sub-section 3.2.2). The interviews and informal conversations will provide a more in-depth understanding of the problem context, allowing us to gain practical insights into the processes and components regarding the topic of this study. The case studies will help us to evaluate and understand how the processes are being applied and what results are created from the

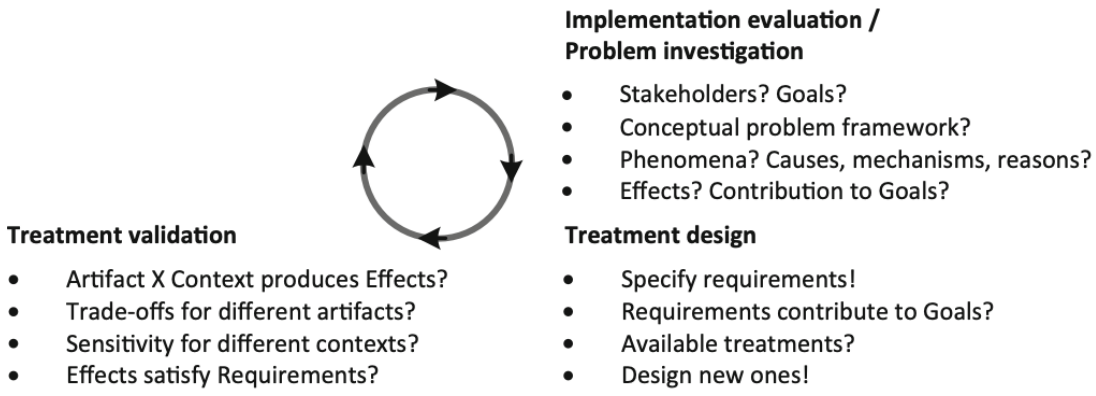

FIGURE 3.3: Design Science Research Methodology Cycle by Wieringa [73]

processes and components within BCM. In practice, the case studies with document analysis, semi-structured interviews and informal conversations will lay a basic understanding of the problem context that will be investigated.

This step will give us a comprehensive understanding of the problem context, which will allow us to design the solution architecture based on requirements. These requirements were acquired from the results of the problem investigation, in combination with the literature review. In the process of modelling the proposed solution, we will be using Archimate, which has been explained in the chapter 2.

Creating an implementation prototype of the solution will be out of the scope of this research due to time constraints, and thus we will conduct the validation process with expert opinions. This means we will be gathering insights from individuals with specialized knowledge and experience relevant to the subject [73]. This process begins with identifying and selecting experts who possess the requisite expertise, often determined by their professional background, academic knowledge, or industry experience. Once selected, these experts are engaged through structured methods such as interviews, surveys or focus groups. This methodology helps ensure that the findings are robust, credible, and grounded in practical and theoretical knowledge. The validation will be further explained in section 3.3.2.

3.3.1 Case studies

In this study, we aim to describe and explain the problem contexts thoroughly to gain a holistic understanding of the challenges. Due to the explanatory nature of this study, we will use the case study methodology developed by Yin [76]. Yin's [76] case study methodology offers a structured framework to gain a comprehensive understanding of complex problems. In the framework, researchers can take a case of a certain context of interest to identify and explain phenomena. His approach uses *triangulation*, taking multiple sources of evidence which adds to the validity of the methodology. We will adopt this framework to get a better understanding of the problem contexts. Considering his approach, we will also consider multiple sources of evidence in every problem context. The cases that were chosen for each of the problem contexts are unrelated. This means we have taken a single case approach for each respective problem context with a single unit of analysis. A single-case approach can be representative of more cases and thus provide interesting results. However, it should be noted that Yin sees it as a "trap" [76], to try and select a representative case for generalization. He argues there never is a case that represents all cases,

but these representative cases can be used to add to theory. Although a multiple-case approach provides the study with even more credibility, this was out of the scope of this research project. From the results of the case studies, generalization will be conducted on the collected data to develop enterprise architecture models. [76]

Case Selection

Selecting an appropriate case is a critical step in ensuring the validity of the findings in this study [76]. For each business process, a single case has been chosen that is intended to represent typical conditions and practices found across typical cases. The cases were selected based on the expert recommendation, who identified the chosen case as an exemplifying case context and results.

Northwave has conducted each of the BCM processes multiple times. Do note that this was always done from a cybersecurity perspective. Considering this study is also conducted in the context of cybersecurity, this is representative in this context. Due to the sensitivity of the information included in the BCM processes, the cases and results will be described with strict confidentiality. Details related to the client have been altered or omitted to protect the privacy of all involved.

Data Collection

To comprehensively explore each BCM process for each case, a multi-method approach will be employed to induce triangulation. This approach ensures a rich understanding of the phenomenon by triangulating data from different sources. The three primary data collection sources used in this study are documentation, SSIs, and informal conversations. How the data will be analysed is described in sub-sub-section 3.3.1.

1. Documentation: By collecting documentation about the case, we will get a better understanding of the problem context by identifying key components and processes. This data source provides a foundational understanding and contextual background, which will contribute to the semi-structured interviews where a deeper understanding will be gathered.

For both the disaster recovery test and business impact analysis, documentation on the case has been reviewed thoroughly. Both had documents on the process itself and the results of the process for a certain organization. Examples are project plans, interview documentation, and project reports. Only for the disaster recovery, we were unable to gain access to the case due to confidentiality.

2. Semi-Structured Interviews: SSIs are a core component of the data collection process, providing in-depth insights into the experiences, perceptions, and practices. Similarly to the methodology used in section 3, the SSIs that were conducted for the case studies followed the proposed framework authored by Kallio et al. [40]. As mentioned previously in section 4.1, the SSI methodology is a versatile and flexible data collection method that can be used in various contexts [40]. However, contrary to the implementation of the methodology in the literature review, the goal of this study context is to extract data from experts in the field of BCM and disaster recovery on the processes and components within these business functions. As these interviews can hardly be fully structured beforehand due to the exploratory nature of the interview, the interviews will take place in the semi-structured form [40]. For each of the problem contexts, we want to extract the same information and use the same interview guide. In the following paragraphs, we will

explain how the methodology was implemented.

Retrieving and utilizing the previous knowledge: In this methodology, background knowledge was acquired in multiple ways. The same contextual knowledge was acquired for the literature review methodology explained in sub-sub-section 3.2.2. In addition, knowledge from the whole literature review, case documents, and informal conversations was acquired and utilized before entering the interviews.

Formulating of the preliminary interview guide: For all three problem contexts, the same interview guide was developed and used. It includes sections for introducing the interviewee, extracting data related to the problem context, and gathering data on specific requirements. The interview guide can be found in appendix C.
Pilot testing: Pilot testing of the data collection methods was conducted through informal conversations, which provided feedback on the clarity and relevance of the questions. As mentioned above, the explanatory nature of the SSIs allowed for straightforward questioning to extract relevant data. With the option to ask additive questions during the interview, it ensures a complete understanding of the components and processes can be acquired.

Presenting the complete interview guide: To extract data, two experts were interviewed in a total of three interviews for each of the cases. In this research, two experts were selected to conduct the validation. One of which is an expert with multiple years of experience in the field of CR. The expert has executed various types of projects related to CR and BCM, including business impact analysis and disaster recovery tests. The second expert has multiple years of experience as a member of the Northwave CERT assisting organizations that are dealing with cyber attacks. This expert also has experience with supervising disaster recovery tests, next to his experience with many disaster recovery procedures at clients of Northwave.

3. Informal Conversation: Informal conversations serve as a supplementary data collection method that offers additional context and insights that might not emerge through formal methods [76] [66]. Especially gaining background knowledge on the subject is often captured from this type of conversation. Next to this, it was used to ask for clarification on certain elements that arose from the analysis of the other two data collection methods. These conversations are typically spontaneous and occur in day-to-day settings, or during other knowledge-sharing meetings.

Data Analysis Techniques

The data analysis process in this study is designed to explore and integrate information from multiple sources, including documents, semi-structured interviews, and informal conversations as mentioned in section 3.3.1. The analysis is tailored to these data collection methods to develop conceptual models that accurately reflect the constructs and processes of the problem contexts. After analysing the data from each of the data collection methods, the findings were integrated to develop grounded conceptual models. Triangulation was employed to validate the findings by cross-verifying information from different sources, enhancing the credibility of the conclusions [76].

1. Document Analysis The initial phase of the data analysis involved the examination of relevant documents related to the problem context cases. The purpose of this

analysis was to establish a foundational understanding prior to the interviews. We employed content analysis as our primary analytical technique, focusing on identifying key themes, patterns, and constructs within the documents [35]. The content analysis mostly adopted a conventional approach, although insights from the literature review were used as a basis for understanding certain concepts [35]. By analyzing these elements, we developed an understanding of the problem context. These findings were noted down, and since many concepts required complementary explanations, the data was not directly transferred to the conceptual models. Instead, it was combined with insights from other data collection techniques mentioned at the beginning of this section to induce triangulation [76].

2. Semi-Structured Interviews Following the document analysis, a semi-structured interview was conducted for each problem context to get insight from an expert on the problem context. In these interviews, the documents were used as a guideline for asking clear questions about the process [40]. With this preliminary insight, more specific processes could be identified and understood. It was also asked what names certain components should have if the names of processes were unclear. These were noted down either on paper or directly into conceptual models. Next to this, the interviews were recorded and transcribed to conduct further analysis of the content [40].

Given that only one interview was conducted per problem context, the analysis focused on combining the findings from the document analysis. Pattern matching was employed to compare the insights from the interviews with the concepts identified in the documents [76]. This ensures that the insights align and that there are no discrepancies between the two data collection methods.

Explanation building was also part of this phase, where the data from the interviews was used to construct explanations for the phenomena under investigation [76]. These explanations were iteratively refined as the analysis progressed, to remove any uncertainties.

3. Informal Conversation Informal conversations complemented the formal data sources by providing additional context and clarification [76] [66]. These insights were not always formally noted down fully, but for some conversations, recordings were made for analysis purposes. The data gathered was integrated with the data from the documents and interviews. Mostly, clarifications were directly implemented into the models, as the results from the other two data collection methods sometimes led to wrong interpretations. The informal conversations often provided added depth to the understanding of specific issues and improved the models iteratively [66].

3.3.2 Validation

For the validation of the solution design, we employed a qualitative approach based on expert opinions. Three experts were consulted varying in fields like CR and IR. Each expert was provided with the reported solution detailing the requirements, design choices, complete reference architecture, challenges and considerations. They were asked to carefully review this document, focusing on whether the proposed solution complies with the stated requirements and noting the effects that they consider to be produced if the solution design is implemented. During their review, the experts were encouraged to note any corrections, observations, or additional ideas that emerged. This method allowed for a thorough examination of the reference architecture, leveraging their expertise to ensure its validity and applicability. The results will be collected and described in chapter 7.

Expert Opinions

In this research, the two experts that were also consulted in the SSIs during the case studies mentioned in sub-sub-section 3.3.1, will also be consulted to execute the validation procedure. As both experts have much experience in the research context, they can give realistic validation feedback for the proposed solution design.

3.4 Conclusion

In this chapter, we have described the overall research methodology and the research methods that will be used to answer our research questions. Table 3.2 provides a mapping of each chapter and their corresponding research questions, design science research methodology phase, and research methods. This structured approach not only aids in tracking the progress of the research but also demonstrates the logical flow and coherence of the study.

Chapter	DSRM Phase	Research Methods	Research Questions
1. Introduction	-	-	-
2. Background	-	-	-
3. Methodology	-	-	-
4. Literature Review	-	Rapid Review: Semi-Structured Interviews	SQ1, SQ2
5. Problem Investigation	Problem Investigation	Case Studies: Document Analysis, Semi-Structured Interviews, Informal Conversations	SQ2, SQ3
6. Proposed Solution	Treatment Design	-	SQ4, SQ5
7. Validation	Treatment Validation	Expert Interviews	SQ4, SQ5
8. Discussion	-	-	MQ1
9. Conclusion	-	-	MQ1

TABLE 3.2: Methodology structure aligned with chapters and research questions

Chapter 4

Literature Review

4.1 Rapid Review Results

In this chapter, we will discuss the general findings, article contributions, challenges & limitations and recommendations that were extracted from the RR. In the general findings, we will discuss some notable and sometimes unexpected notions that can be said about the body of literature in subsection 4.1.1. The article-specific contributions will be presented as key trends in subsection 4.1.2. To give this subsection a clear structure and differentiate between contribution areas, they were listed according to the BCM framework by Russo et al. (2024) [57]. Afterwards, challenges & limitations and recommendations were noted as collective results from multiple articles in subsection 4.1.3 and 4.1.4 respectively.

4.1.1 General Findings

In recent years, there has been a growing number of literature studies exploring the role of the DT in the cybersecurity field. Figure 4.1 demonstrates that the amount of articles that were found in the RR per year has been significantly increasing in recent years. This is probably due to the increase of attention towards cybersecurity in cyber-physical systems, as this is one of the main applications where DTs with a cybersecurity goal are used.

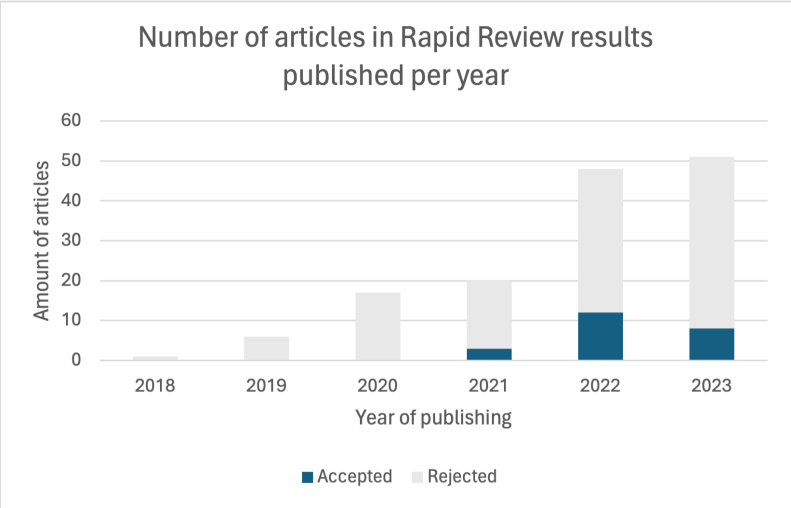

FIGURE 4.1: Number of articles published per year in RR results

Out of the 159 papers that were analyzed and evaluated in the queries, 136 were

rejected based on the exclusion criteria. The majority of these papers were rejected because they did not contribute to BCM and instead focused on incident prevention or incident detection systems which, while important, did not align with the objectives of this study. Furthermore, some papers were deemed irrelevant due to the study not being conducted in the context of cybersecurity altogether, as they did not address the specific challenges and concerns within this domain.

Among the papers that were selected, many presented novel implementation frameworks for a DT, DS or DM to enhance CR. However, given the complex nature and broad definition of these technologies, frameworks adjust to the context and goals of a study, resulting in high variation. This includes the context in which BCM is used, such as incident response, and the industries where it might be applied, like the Internet of Things (IoT) for healthcare. Some of the frameworks try to reach high fidelity, while others prioritise usability.

One significant finding is that almost all articles used the term DT throughout their work. In reality, these could more accurately be classified as a DS or DM as explained in subsection 3.2.1. This is clearly illustrated in figure 4.2, where the classification of DT types by Kritzinger et al. [42] reveals that many articles actually developed or designed a DM. Furthermore, many articles did not specify how the data flow was executed in general (automatic or manual). This indicates that the terminology used within the current literature, as discussed in subsection 2.1, remains unclear. It highlights the importance of clearly defining and differentiating between these terms to ensure accurate communication. A second notion that could be made is that some articles are only theorizing about the capabilities of a DT, without showing a specific framework or implementation of their proposed solution. As this is not necessarily an issue for these articles, since they are often theorizing conceptually, we would like to differentiate these results to get a clearer understanding of the implications. As we can find in figure 4.2, the articles that did not propose a DT framework or implementation in their work did also mostly not describe the type of data flow connection.

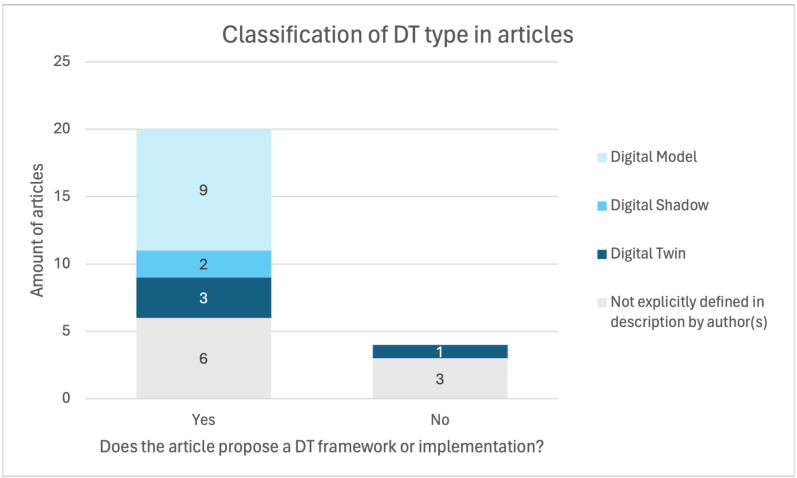

FIGURE 4.2: Classification of DT type in articles

Another finding is that 17 out of 25 articles include a notion of Operational Technology (OT) or Cyber-Physical Systems (CPS), rather than just Information Technology (IT) systems. These articles have mentioned the explanation of one or more OT devices in their proposed DT, DS, or DM solution. This is logical due to the nature of the DT technology and its focus on the virtualization of physical objects. Moreover, it can also be explained

by the focus of many studies on critical infrastructure, where OT is often used to operate machines. Nevertheless, it also highlights a potential research gap in understanding the role of these technologies in enhancing the CR of IT systems, which form the backbone of most modern organizations.

Another notable observation from the RR is the lack of literature exploring the potential of DTs for system recovery. Given the inherent capabilities of DTs to simulate and test various scenarios without risking real-world systems, one might expect that their usage in system recovery or as an alternative to critical functions would be a point of attention in research. However, this topic appears to be underrepresented in literature as there are no studies that suggest contributions to the 'alternatives to critical functions' component in the BCM framework, and only few suggest how it can enhance recovery methods in the IR component. This represents a significant gap in our understanding and utilization of DTs and suggests opportunities for further exploration in future studies.

Despite the increasing interest in integrating DTs into BCM components, there have been no studies performed on the holistic integration of the DT into BCM as a whole. This absence of research leaves a gap for research area-wide understanding of how DTs can enhance BCM practices, particularly in areas such as risk assessment, business impact analysis, BCM strategy and BC training & testing. Furthermore, such a conceptual framework could develop guidelines for the successful incorporation of DTs into BCM, giving practitioners a reference procedure.

Three articles by Holmes et al. [34], Faleiro et al. [24] and Management Events [44] provided literature reviews that give an overview of research done in the context of this research and suggest potential benefits and challenges of the DT technology. Holmes et al. identify challenges and benefits of DTs in cybersecurity that are found in literature, and suggest future research based upon these findings. Many cybersecurity challenges were identified that would be created by the DT, with main concern themes: *Availability, Integrity, Confidentiality, Data ownership and IP leakage, and Safety.* But they also observed many security challenges that could be solved through a DT, including themes: *Improved patch management, Improved risk management, Active Cyber Defence, Advanced Training and Incident Response Capability, Anomaly Detection, Virtual Commissioning, Autonomy, and Predictive Analytics.* Faleiro et al. [24] look into the uses of the DT for enhancing CR. They suggest four main use cases: *Intrusion Detection Systems, Simulation Testing and Training, Privacy and Legal Compliance, and Security for the Factory of the Future.* What should be noted is that the authors do include measures that can be associated with security mitigation or prevention measures, which are not included in this study. Also, the study by Faleiro et al. [24] was published in 2022, but conducted in 2020. It resulted in 13 articles, which is considerably less than the amount included in this study. The paper does include benefits, which are also results in this study, confirming the capabilities of a DT in the context of cyber resilience. Lastly, the article by Management Events states uses for the DT enhancing CR including: *Risk Assessment, and Security Testing and Validation.* These articles are not added to the overview shown in key trends, as they refer to certain articles themselves that should have been captured in this methodology as well. The ideas for future studies are noted in subsection 4.1.4.

4.1.2 Key trends

To be able to structurally group the key contributions of the reviewed articles, the framework by Russo et al. [57] was used. The key contributions were mapped with the specified components of the BCM framework and can be found in figure 4.3.

Noteworthy, in the analyzed literature there were no studies that presented contribu-

tions to the utilization of DT, DS or DM in understanding the organization, alternatives to critical functions, and BCP design and implementation. This gap in current research indicates a need for further exploration into these topics. On the contrary, many articles study made contributions to BCP testing, maintenance and analysis.

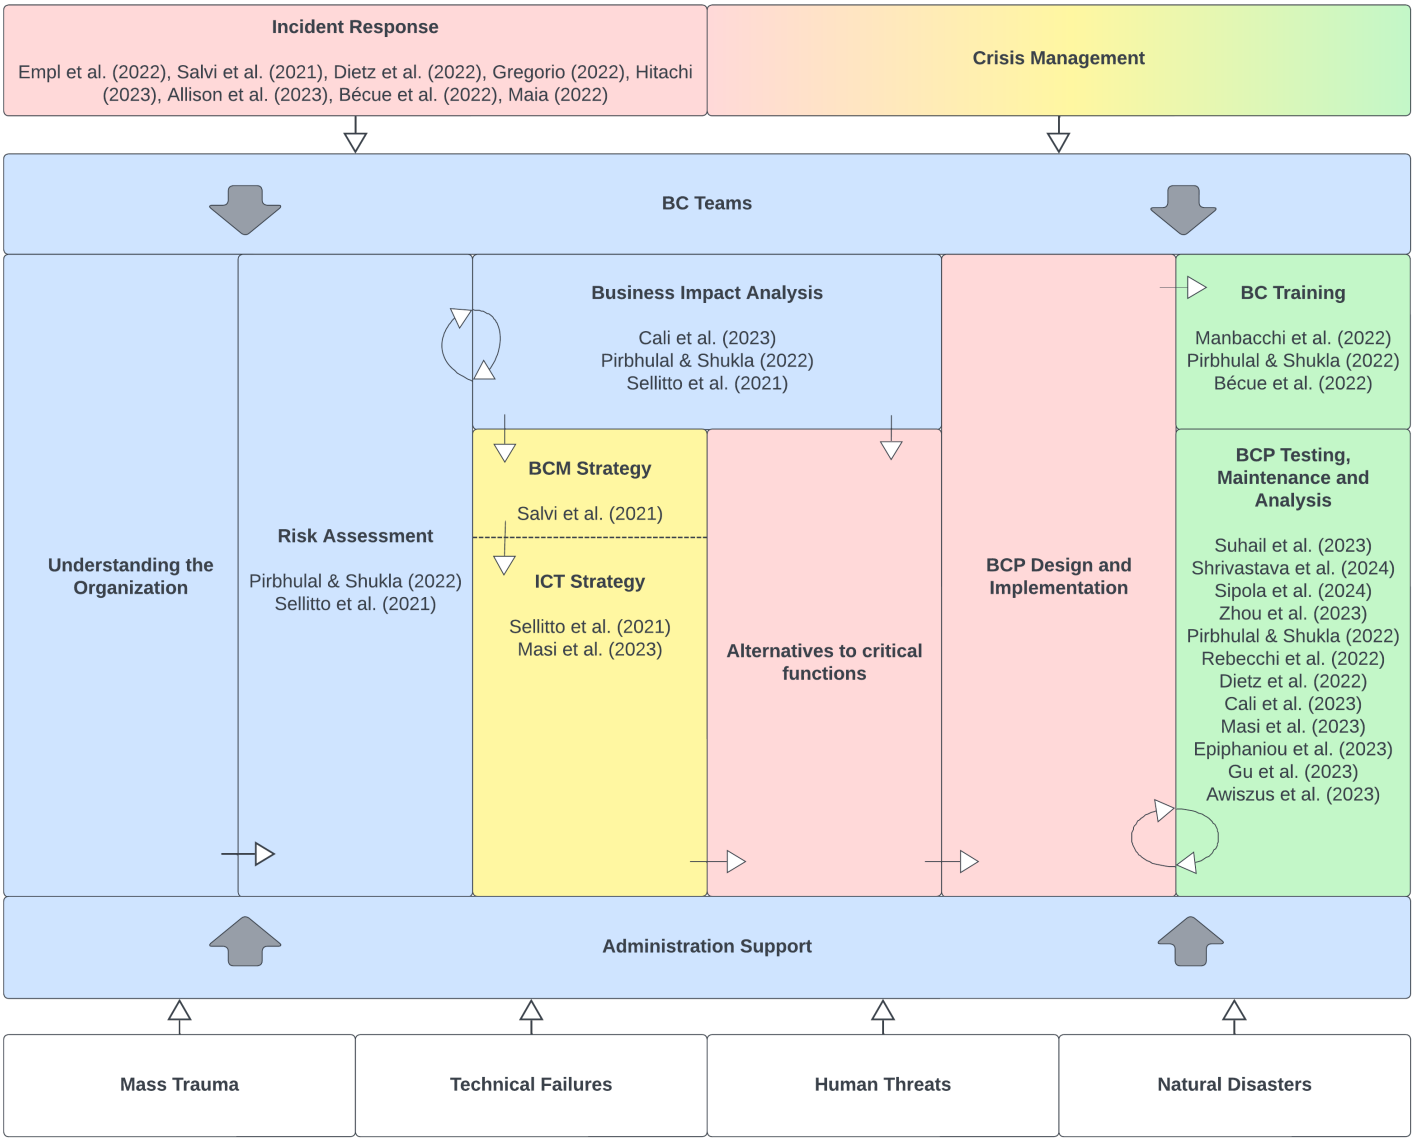

FIGURE 4.3: RR and SSIs papers distributed in BCM framework by Russo et al (2024) [57]

Risk Assessment

In the study authored by Pirbhulal and Shukla [51], they develop a novel DT framework in IoT-based healthcare applications. In this paper, the researchers propose to use the DT for risk simulation to enable risk assessment. By using "the input data and system model, a threat scenario is designed and integrated with the DT world" [51]. They learn from this step to identify potential vulnerabilities, risks and threats which can enhance

risk assessment processes as they argue. [51]

As also explained in sub-subsection 4.1.2, Sellitto et al.'s [59] article offers a methodology to model a security-oriented DT of critical infrastructure. The DT, which is based on an Enterprise Architecture EA blueprint, can be employed to simulate cyber threats and assess countermeasures' effectiveness. This method can be used to assess risks iteratively by incorporating simulation results back into the EA model to assess improvements to system security. This approach enables in-depth risk assessments without interfering with the operational systems, thus enhancing the risk identification and mitigation strategy process which can be used within the business impact analysis. [59]

Business Impact Analysis

Cali et al. [13] looked at the possible benefits of DT for the future of energy systems and smart cities. They argue that proactive cybersecurity studies can be done on an energy system using the DT of the operating grid model. These granular feeder DT models can be leveraged to build NESCOR failure scenarios. These failure scenarios, compiled by the US National Electric Sector Cybersecurity Organization Resource (NESCOR), are an example of a comprehensive security assessment that includes business impact analysis [37]. These scenarios can then be used to forecast failures, perform predictive studies, conduct tabletop exercises and other applicable applications. [13]

In the paper authored by Pirbhulal and Shukla [51], the authors do not only propose to use their DT framework for risk assessment explained in sub-subsection 4.1.2, but also propose to use it for Business Impact Analysis. They use the same input data and system model to create a threat scenario and assess the impact it could have on the business and its recovery time. They argue that their "DT module then analyses the current performance of the system and suggests improvements accordingly which can be applied to the physical world" [51].

As mentioned in the previous sub-subsection, the paper by Sellitto et al. [59] demonstrated how to analyze security risks and mitigation strategies using their EA modelling method. Their solution "is an Enterprise Architecture model of the system specifically targeted at providing a sound base for simulations in order to devise proper countermeasures without any outage of the physical infrastructure." [59] Like mentioned, their study also highlighted the opportunity to evaluate various mitigation strategies. Although not explicitly mentioned, this method could be integrated into a business impact analysis to understand the value of each mitigation strategy. [59]

BCM Strategy

Cyber critical infrastructure (CCI) is of high risk to cyber-attacks and should aim to be able to respond to attacks quickly. In the paper by Salvi et al. [58], the researchers aim to make CCI more resilient by bridging incident prevention and IR in a conceptual framework that addresses three organizational layers. The conceptual framework includes a DT which provides hardening or cyber-physical systems, risk scenarios for policy requirements and training for organizational requirements. This framework can contribute to the BCM strategy by addressing the organizational workflows of different elements in BCM. [58]

ICT Strategy

Sellitto et al. [59] describe that Enterprise Architecture (EA) modelling languages have little in common with the need to be able to model cybersecurity DTs. In their paper,

they "propose a methodology to derive a digital twin of a critical infrastructure, aimed at performing simulations for cyber security and visual threat modelling, starting from an architectural blueprint of the system" [59]. With the newly created Cyber Security EA View, it is possible to model the cyber security DT and execute cybersecurity analysis using cyber-attack simulations with the MAL (Meta Attack Language). In a second paper by the same authors, they extend this methodology to critical infrastructure and decouple the process from design/construction to be able to include legacy systems [46].

Incident Response

The most prominent paper, which was not extracted from the RR, was the study by Allison et al. (2023) [1]. It proposes an integration of the DT technology into the IR life cycle for cyber-physical systems. The researchers do this based on the NIST IR Life Cycle [15]. They use the three types of DT security modes for cybersecurity, mentioned in subsection 2.1, to identify how the technology can empower the steps in the IR life-cycle. They recognized that this integration can aid practitioners throughout the IR workflow, and contribute to this field by analyzing IR workflows and finding appropriate DTs for the specific workflow items. [1]

In the paper by Empl et al. [19], the researchers state that IoT assets are often poorly secured and IoT devices often become part of a botnet. Organizations often use Security Orchestration, Automation and Response (SOAR) platforms to handle events and deal with incidents. In this paper they are investigating the use of SOAR for IoT, and how this can be enabled by DTs due to its promising features and IR capabilities that they found in the literature. While considering different options to secure IoT, they found that DTs "provide a feasible, lightweight solution abstracting heterogeneous assets". They propose a novel SOAR4IoT framework that uses a DT-based middleware to secure IoT assets. [19]

As explained in the paragraph about the paper of Dietz et al. [18] in sub-subsection 4.1.2, it is explained that the Cyber Threat Intelligence (CTI) can be shared among organizations with the DT security simulations. They argue that this is also applicable to IR case scenarios. [18]

In terms of grey literature, Gregorio identifies that attack vectors are always changing in critical infrastructure. Organizations have to move beyond preventive measures and "use model-based system engineering (MBSE) techniques to create a digital twin for the SoS" [30]. They argue that by using a "proactive risk management approach, system architects can enable higher operational availability so that even when cyber attacks occur, the most important operations within critical infrastructures can keep running". [30]

The article by Hitachi also suggests to use of a DT to test OT countermeasures [33]. As traditional security countermeasures can cause unexpected system downtime, they try to solve this problem with a so-called 'Security Digital Twin'. With thin DT, they can assess the side effects of security countermeasures before they are used in practice. This can be done preemptively or during an incident. They propose a technical implementation of this model in their article. [33]

The article about this activity by Maia et al. [43], demonstrates they developed the SMS-DT, a platform that simulates and monitors industrial conditions within a DT-based architecture to enhance cyber situational awareness in future-oriented factories. This platform captures, analyzes, and correlates a vast range of events tracked by sensors and systems across the factory. It includes components that simulate key aspects of industries, like network analysis, energy optimization, and worker behaviour. Tested in a textile industry environment, the platform proved to intervene in giving a comprehensive view of an incident, demonstrating its potential to enhance security operations to enhance IR. [43]

Salvi et al.'s [58] article introduces a model for enhancing CR in CCI using DTs. Their contribution is focused on showing how prevention and response strategies of cyber incidents correlate on a multi-layer model, in which the DT is utilized as a security measure. The model shows the utilization of fine-grained data analytics for the early detection of threats and a system for causal analysis to identify root causes. Thus, facilitating rapid cyber threat mitigation, both as an incident response and preventive measure. The use of DTs creates improved situational awareness and promotes better coordination across teams and organizations. This approach not only reduces response time and impact of a cyber incident but also facilitates continuous learning, bridging the gap between prevention and response. [58]

In the article authored by Becue et al. [7], they discuss the security operations of a DT in the context of aerospace manufacturing and OT systems. The researchers propose an implementation to use a DT that "If maintained in operation, a DT can support decision making for how to react to unexpected attacks", thus enhancing the incident response capabilities. The simulation capacity of the DT is used to facilitate the design of a secure OT environment where threats can be identified and mitigated accordingly. [7]

BC Training

Manbachi et al. looked into the possibility to train people to work with cyber-physical environments from a remote location [45]. During the COVID-19 pandemic, there was an increasing amount of demand for a possibility to educate people to familiarize themselves with cyber-physical systems remotely. This prompted the researchers to look at the general education of employees using a DT in an online environment. To enable this, they developed a real-time platform called the Virtualized Experiential Learning Platform (VELP), which provides a hands-on virtualization tool to interact with physical counterparts. Here researchers can also conduct studies, and employees can be educated in industries like smart grids. [45]

One of the applications of the DT identified by Pirbhulal and Shukla [51] is to use the technology for advanced training in IoT-based healthcare applications. In the paper, the researchers propose a novel framework to build a DT which can be used as a cyber range in the context of IoT-based healthcare applications. The concept of a Cyber Range and how it relates to a digital twin will be discussed in section 8.5. This could be used for both the training and testing component of the BCM framework, which is left undefined by the author. [51]

Also in the context of aerospace manufacturing, the cyber-range based DT was studied by Bécue et al. [7] also proposed use for a training environment. They mention that it can be used for training employees, and as an instrument for decision-making in IR scenarios using DT and cyber ranges as mentioned in the previous sub-subsection. They propose a novel design for a DT of cyber-physical systems in the aerospace environment. An important contribution is their realistic attack scenario they presented. The technique they present involves a highly heterogeneous OT environment. [7]

BCP Testing, Maintenance and Analysis

In the paper by Suhail et al. [65], the researchers try to expand on existing DT gamifying techniques to develop a security-aware learning environment that can be used for automated security testing and enable explainable DT assessment for security analysts. They propose ENIGMA, (sEcuriNg dIgital twins through GaMification Approach). ENIGMA "leverages DT as an offensive security solution by launching overt attacks on DT instance(s)

in a simulated, interactive, and controlled environment, hence called gamification" [65]. With ENIGMA, the researchers present a novel approach to enable cyber exercises to test the security of a DT. They argue that "such DT assessments must be performed in an isolated environment without negatively affecting the operation of DT replication mode, where it is critical to register changes between the states of the twin and its physical counterpart to avoid invalid interpretations". Hence, the DT in their framework that represents a cyber-physical system is used to instantiate gaming scenarios environments. This assures there is no risk to the critical infrastructure to be affected. Another key contribution in this paper is the AI-enabled component in the attack/defence scenario of the gamification approach. [65]

The paper by Epiphaniou et al. (2023) [21] identified that is hard to understand the cyber effects that cyber attacks can have on cyber-physical systems and the interactions between the physical, electronic and cyber domains. Within their research, they try to tackle this challenge by examining an AI-enabled DT to identify threat source characterisation. Next to this, they identify "existing gaps when integrating standard security description references for attack analysis (e.g. Cyber standards) with simulation standards" [21]. As a result, they present a review of cyber modelling and simulation standards for cyber attack analysis and a list of recommendations for cyber resilience assessment with a DT. This study can contribute to the testing of a BCP and analyse the resilience level of an organization. [21]

Shrivastava et al. identified many security preparedness exercises revolved around IT, leaving OT less attended, while cyber-physical attacks have expanded [61]. While there is a lot of information on the subject, there is a gap in structure for cyber-physical exercises (CPX). As a solution, the researchers provide us with a CPX playbook that enables "an ideal battleground for attackers (red teams), defenders (blue teams), evaluators (white teams), forensics (yellow team), and infrastructure builders (green teams)" [61]. In this paper, a DT was built of a SWaT testbed to enable the CPX. [61]

Another paper by Sipola et al. [63] identified that cyber exercises are a great way to enhance the knowledge and skills of staff in the food supply chain sector. To create a technical infrastructure to enable such an exercise, referred to as a Cyber Range or Cyber Arena, they used the requirements from Karjalainen and Kokkonen [41]. One of the requirements states that the arena should be as realistic as possible. They propose that a DT can be used to enable high realism of the cyber arena, however, they aimed to use a Digital Model as this was more appropriate for their research while still delivering a good amount of realism. They argue that "such real-time approaches are not quite feasible for cyber security situations, not least because of the often sensitive nature of the domain data". In the paper, they presented a novel approach to develop a DM to create a cyber arena with the use of standard web technologies following a web service architecture. [63]

In the paper by Zhou et al. [77], the researchers want to deal with the two main challenges of virtual-real connected Industrial IoT (IIoT) cyber ranges. The first problem is that industrial devices on the cyber range are expensive, which means that every instance of such a cyber range will have repeatable high costs to create (including costs for experts to set up the environment). The second challenge is that network attacks can cause damage to the industrial site, and in such cases, scenario restoration costs a lot of time. To deal with these challenges, they use a DT-based approach to create the cyber range. They argue the two main benefits of using a DT-based approach for IIoT cyber ranges are "On the one hand, the digital twin is needed to substitute the real devices to react to the attacks from cyber ranges to avoid unnecessary damages because of their financial value. On the other hand, to deploy cyber ranges freely without the limitations of the devices or locations, a

digital twin-based cyber range is a promising approach ideally. Therefore, a framework for a digital twin-based IIoT cyber range is proposed." [77].

One more paper researched the possibility of using a cyber range with a DT-based approach. In the paper by Rebecchi et al. [52], the researchers have identified that there is a large amount of 5G cyber security threats that need to be addressed. They want to be able to conduct exercises in a realistic and safe 5G environment, be able to track the performance of employees using this exercise and be able to develop exercises with AI components. As a solution to these challenges and goals, they propose an architectural approach of a so-called SPIDER cyber range, taking "full advantage of advanced network orchestration, log-processing data pipeline, cyber risk assessment frameworks, and applies advanced machine learning techniques in support of its hands-on learning objectives." [52]. With this approach, red teams and blue teams can emulate attack and defence scenarios to train for real scenarios.

The research by Dietz et al. [18] identifies a gap in the utilization of digital twins for enhancing cybersecurity through CTI. Organizations are hesitant to share CTI due to potential negative repercussions and trust issues. Additionally, Industrial Control Systems (ICS) are targeted more by cyberattacks and are often equipped with outdated security measures, making them vulnerable. By integrating DT security simulations with standardized structured CTI, organizations can use collaborative threat intelligence for preventing, mitigating, and remediating security incidents. In technical cyber exercises, new intelligence can be gathered and this technique will help organizations collaborate to use this CTI. [18]

As described in the subsubsection 4.1.2, the initial paper by Selitto et al. [59] described an EA DT modelling approach. In their paper, they also proposed a vulnerability assessment and security test using a threat-oriented DT and security reasoning. With this analysis, they demonstrated a use case and identified three attack scenarios, which they could mitigate by changing the design method. With this context, in a second paper by the same authors, Masi et al. [46] try to tackle the issue that simulation and assessment of BCP of critical infrastructure is usually restricted to the design/construction phases of a system. This is often due to the inability to halt a critical system to perform the business continuity test. They mainly extended the methodology towards critical infrastructure, and decoupled it from the architectural development phase so it can also be used for legacy systems. This contribution in their second article enhances the testing capability of BCPs. [46]

Gu et al. (2023) [31] notes that there is no comprehensive and precise way to evaluate cyber-physical resilience. In their paper, they present the cyber-physical resiliency metric (CPRM), which can be used to evaluate resilience based on the problem description. They specifically use it in their study for microgrids, which are seen as critical infrastructure. They demonstrate how DT-based simulation can be used to measure this resiliency. This can be related to BCP testing and analysis to measure the amount of resilience of cyber-physical systems. [31]

Another paper by Awiszus et al. (2023) [6] also looks at cyber resilience assessment for complex cyber systems. Their concern is coming from a legislator and regulator perspective that finds it hard to evaluate cyber systems. In their paper, they show how networked systems can be modelled to evaluate cyber infection spread from a DT-based paradigm. Using network science and contingency theory, they introduce the artificial cyber lab as "an experimental framework to analyze the impact of both security-related and topology-based interventions" [6]. This can be used to analyze a BCP and its resilience.

In the work by Pirbhulal and Shukla [51], they discuss the option of testing and sim-

ulation for a DT in a healthcare system. They explain how threat scenarios are designed based on input data and system models, and then tested in a simulated DT environment. The testing includes traffic generation to identify bugs and vulnerabilities, as well as attack generation to consider potential threats. The results can be used as an impact assessment, consisting of identification, evaluation, and prediction stages. The DT module analyses the security performance of the system and suggests improvements which can be applied to the physical world. [51]

The article authored by Cali et al. [13], mentioned the potential benefit of using the DT to test NESCOR scenarios and evaluate the results. They argue that "Such a study will lead to forecasting failures, ..., workforce advancements through tabletop exercises, etc." In this way, a tabletop exercise scenario can be developed through previously earned intelligence. [13]

Conclusion

Concluding this extensive subsection, it highlights the innovative methods introduced by different researchers to leverage DTs, DSs, and DMs in the respective contexts of each study. For instance, using DTs for risk simulation in IoT-based healthcare applications [51], or the integration of DT technology into the IR life-cycle for CPS [1], shows the multitude of solutions it can offer. As already mentioned at the start of this subsection, these results show that some components received more attention in literature than others. It suggests that the technology can be more useful as a training environment using the simulation mode of operation (by Dietz and Pernul [17]) to test and train teams. On the other hand, the results show a low amount of interest in the use of the technology for the BCM components of understanding the organization, alternatives to critical functions, BCP design and implementation, and crisis management. In the case of understanding the organization and crisis management, this is quite logical due to the management-level contexts of these components. However, the other two components which did not receive much attention would be expected to have some degree of research. Especially the idea that one might be able to use or restore a system from an uncompromised DT, DS, or DM, as an alternative to a critical function is a research gap.

Also, interestingly, an integration of the DT into the design and implementation was not found. However, as this is still a new research area, and no holistic study towards the use of DT, DS, or DM was done towards BCM, this does not come as a big surprise. As can be seen in figure 4.3, many papers contribute to BCP testing, maintenance, and analysis components. These contributions for varying types of exercising and testing procedures use a DT, DS, or DM in various ways. The requirements of the technology for the different exercises lack a complete overview of requirements and guidelines. A conceptual framework on the integration of DTs to BCP testing, maintenance, and analysis component is lacking, which could be used by practitioners to understand why and how a certain type of DT could be used for a specific type of BC exercise or test.

4.1.3 Challenges & Limitations

Next to the contributions to various BCM components, many authors also mentioned challenges and limitations in their work. One significant challenge faced by researchers is ensuring the cybersecurity of DTs. Given that DTs provide a connection from physical systems to a digital environment, they present new vulnerabilities that cyber attackers could potentially exploit [24] [13] [19] [34]. Ensuring the security of DTs is a complex

task. Users of a DT need to continuously update their security protocols to respond to the rapidly evolving landscape of cyber threats.

Another challenge in developing DTs is the high development costs [13]. The creation of a DT takes a lot of resources because it is made for a specific system. The complexity of such a system can vary, which needs to be explored and understood beforehand. In general, there still needs to be an approach developed that makes DT development feasible and economical [34].

Developing a standard process for creating a DT presents challenges due to the inherent complexity of systems [13] [63]. Cali et al. [13] identifies the lack of "common technical specifications and protocols to ensure interoperability and compatibility among different systems and processes". Without this, DTs might not fully optimize efficiency and performance [13]. The main challenge that Sipola et al. [63] pointed out in their paper about DT-enabled cyber exercises, was that the system was designed specifically for one cyber range, the RGCE. The solutions employed in this range could be unique to it, making adaptability to other environments potentially challenging. Hence, they recommend studying if the cyber exercise could be generalized. Similarly, Suhail et al. [65] also recommends generalizing their DT based on the context to reduce the complexity of the task.

One of the major challenges in the development of DTs lies in accurately capturing the actions and processes of systems. [13] [46] [18] Systems are often subject to a range of parameters and conditions, which can cause changes in response to different events. For some systems, these complexities can be difficult to model, which could cause potential inaccuracies in the DT [13] [65]. These inaccuracies might affect its ability to effectively represent and predict the behaviour of the system it represents.

4.1.4 Recommendations

Finally, many authors also made recommendations for new studies. The article by Manbachi et al. [45] focuses on the Virtualized Experiential Learning Platform (VELP), to train people remotely. They recognized the opportunity to measure and assess the learning rate of employees through VELP to understand better how people interact with the platform and to give feedback to people using it, which they recommend as a future study. [45]

Empl et al. [19] recommend for future work to explore automated DT response strategies through their SOAR4IoT platform, which they are researching as suggested in their study. This can enhance their implementation to become more resilient. [19]

There is a call by Holmes et al. [34] in their literature study for more research work to implement complex systems through a DT approach both feasible and economical. Next to this, they recognize the DT to be an innovative study, which should be responsibly used given its ethical impacts. They suggest future work to develop an ethical framework for the DT. [34]

The study by Salvi et al. [58] also suggests the concept of "ontological reversal", not only mimics but can shape physical reality by the usage of a DT. Future research should delve into the impacts of using DTs for CR on organizations and society. [58]

The limitations of MAL and SecuriCAD are discussed by Sellitto et al. [59], pointing out that they cannot capture a detailed-level system design. As part of future work, there is an aim to define a way of grouping architectural assets to automatically determine the specifications of DT components to be reused for a more efficient modelling process, and to make the target modelling language parametric to avoid vendor lock-in. [59]

The recommendations for improving enterprise architecture of DTs by Masi et al. [46] include grouping architectural assets for reusable component specifications, making the modelling language parametric to avoid vendor lock-in, and employing Enterprise Cartography for better legacy system mapping.

To improve the evaluation of DTs, the study by Dietz et al. [17] suggests future research should define quality criteria for DTs and their usage in security so that benchmark studies can be realized. This will also include offering more advanced attack scenarios, merging different high-fidelity domain models into one digital twin simulation, and investigating aspects of recommender systems for systematic CTI generation. [17]

Allison et al. [1] note that they are researching automated DT response strategies, as they propose this as future work. This could potentially impact the design of playbooks through parallel execution and prioritization of information. [1]

Epiphaniou et al. (2023) [21] suggest to improve existing simulation standards as they do not clearly specify the data acquisition, measurement, and visualization process. Challenges exist in characterizing target assets, translating security requirements, threat information, and impact metrics in simulation methods. Further work is needed to standardize adversarial capabilities and objectives in simulations. [21]

In conclusion, the results have highlighted the general findings, key contributions, limitations, challenges and recommendations of the RR. In the next chapter, we will delve deeper into the key benefits of integrating the DS into BCM.

4.2 Integration of DT, DS, or DM in BCM

In this chapter, we will discuss the key benefits that DT, DS, and DM could bring when used within BCM. At the start of this study, we tried to focus specifically on the DS. However, as can be noted in subsection 4.1.1, only two articles clearly describe the use of DS by Maia et al. [43] and Epiphaniou et al. [21]. This is why we focus on all three classifications of the DT in this chapter.

Three findings are discussed below, which are extracted from the papers presented in the literature study. The specific benefits of each topic are mentioned here.

4.2.1 Realistic Risk Assessment and Business Impact Analysis

The DT, DM and potentially the DS can significantly enhance cybersecurity by simulating risks in a virtual environment that mirrors real IoT-based applications [51]. This allows for proactive studies of cyber-attack prevention and the development of strategies to enhance cybersecurity. DMs can also be used to build failure scenarios, such as those compiled by the National Electric Sector Cybersecurity Organization Resource (NESCOR), which include business impact analysis [13]. These simulations could potentially show system interdependencies, which are important to understand risks and recovery strategies. The technologies enable the risk assessment and business impact analysis to be more realistic, which can provide valuable insight which can be used in other BCM components.

The limitation of this finding is that the papers in which this is discussed lack a comprehensive study to demonstrate how this should be implemented. The authors only argue on the potential without demonstrating that the theory is correct.

4.2.2 Realistic BC Testing and Exercising Environments

In figure 4.3, it can be identified that many papers covered the BCP testing, maintenance and analysis component within BCM. This emphasis is not without reason, since many papers highlight the benefits of using a DT, DS or DM to enhance BC testing and exercising in particular. The types of tests or exercises vary, and the way the DT, DS, or DM are implemented depends on the goal of the training. Most of the papers focus on exercises for blue teams and red teams [65] [63] [52] [61], while some papers mention table-top exercises [13], general familiarization with a system [45], or no exact specific mention of a certain exercise [51] [77] [18] [7].

Although the type of exercise may differ, many of the papers mention that the realistic nature of a DT, DS and DM offers an immersive environment to train people and teams [63] [77] [52] [51] [7] [18] [58]. Recognizing that relatively many articles researched the functionality of the DT, DS, and DM for BCP testing, maintenance and analysis, it is interesting to see that many of the articles actually propose the use of DM as depicted in figure 4.4. Out of the articles that designed or implemented a DT, DS, or DM, six used a DM [52][6][63][77][18][46], one used a DS [21], two a DT [13][65], and two did explicitly define the type [61][51]. Lastly, one article did not propose a framework or implementation of a DT, DS, or DM, and did also not specify a type [31].

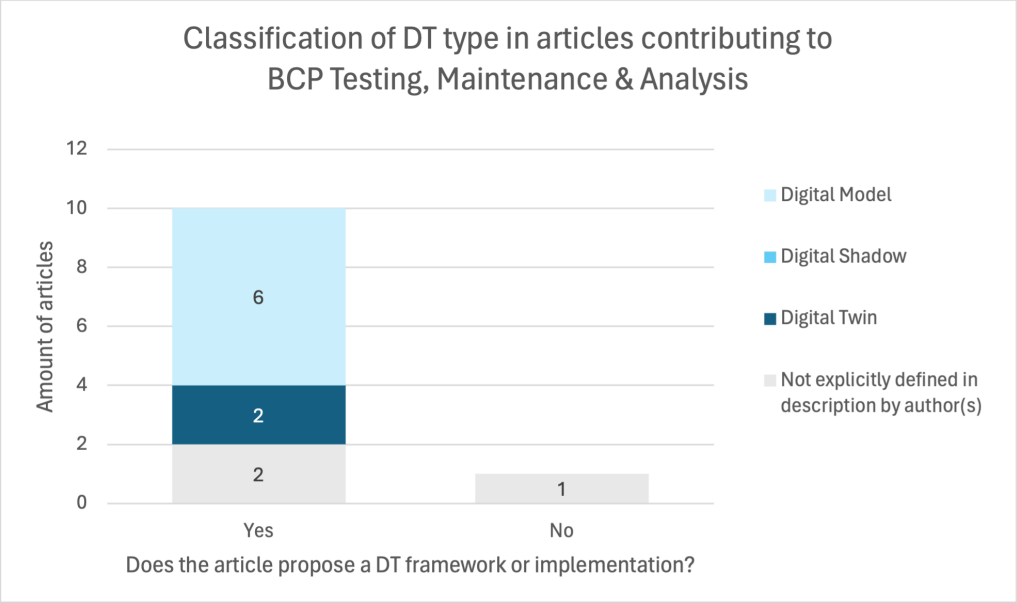

FIGURE 4.4: Classification of DT type in articles that contribute to BC Testing, Maintenance and Analysis component

In comparison with other cyber ranges, a DM-based cyber range can more realistically simulate IT environments and the effects a cyber attack can have. This simulation environment is a key benefit that other solutions do not offer. It should be seen as a kind of sandbox where teams can test their response capabilities without affecting the production environment of their IT infrastructure. However, it is also mentioned that the benefits need to outweigh the risks of implementing a DT, DS, or DM [65] [18], which are mentioned in chapter 4.3. So however it is implemented, these challenges should be considered and dealt with properly to release the real potential of the technology. Next to this, many papers state challenges that limit the fidelity of a system which is also mentioned in chapter 4.3.

An interesting benefit of using a DT for creating BC exercise and testing simulation

environments, as mentioned by Suhail et al. [65], is its ability to provide reproducible DM-based cyber ranges. Compared to the relatively costly development of testbeds, which are the primary alternative to DM-based cyber ranges, the DT offers a cost-effective alternative for simulating reconfigurable simulations of real systems.

4.2.3 Realistic and Testable Incident Response Strategy

Lastly, an essential advantage of the technologies is their ability to simulate and explore various system states and scenarios that would be impossible to conduct on real-world systems due to potential risks on system impact [1] [58] [18] [33][19]. This is particularly valuable in situations where testing in a real-world environment could entail substantial financial costs, or pose potential safety risks. Especially for a DM, researchers and practitioners can safely model and analyze a range of conditions and events, from routine operational situations to extreme events, to understand the implications of certain actions that can be taken in the case of an incident response.

Another benefit of DTs lies in their ability to simulate 'what-if' scenarios during system recovery [1] [33]. This functionality allows organizations to experiment with different recovery strategies in a simulated environment without risking disruptions in the real-world system. By doing so, potential issues can be identified and addressed before implementing the recovery strategy in the actual system. This capability enhances the understanding of system behaviours in cases of an incident and allows for better-informed decision-making during recovery.

Lastly, in general DT, DS and DM offer benefits for informed decision making during the IR life-cycle [1]. They allow for the simulation of various incident scenarios in a controlled, virtual environment that offers a comprehensive summary of the asset's security state [19].

4.3 Challenges, Limitations & Recommendations

In this chapter, the challenges and limitations that were encountered in this study will be discussed, and the general recommendations that resulted from this study.

4.3.1 Challenges

One of the challenges that occurred during the research is that a lot of papers use the term "Digital Twin" in a different way. Some explicitly mention that the DT is actually not connected to the real world and still use the term DT throughout their article, while some did not even explicitly show if their DT technology was connected to the real world. This made understanding the actual contribution of the articles quite hard.

In the RR, papers which claimed that the DT technology could be used for BCM components were included in the research. While most of the papers showed a novel approach to the topic, there were some papers whose contributions to BCM components were questionable. As mentioned in chapter 4.2, the papers that identified the use of DT for risk assessment primarily argued why a DT could be beneficial for the component without any validation to back up the claim.

4.3.2 Limitations

While developing the query to extract articles from the databases, it is a challenge to assess whether all relevant articles are included to extract all relevant information for this RR. Due to time constraints, some synonyms were not used to decrease the amount of articles

in the results. Although most of these synonyms did not seem relevant in this context, they might have excluded some relevant articles. Also, the two papers that were retrieved from the expert reviews, were not found by the specified queries. After looking into this matter, this was because no one synonym of "organization" was used in the article abstract, title or keywords. Given this information, in future literature studies, it might be useful to disband the "organization" term altogether, which will increase results significantly.

One of the exclusion criteria in this review was the context of BCM. Because of this exclusion, papers that discussed cyber risk mitigation measures using a DT were not considered. In the first selection phase, many papers that were categorized like this were excluded. However, it might be the case that some papers did not discuss BCM contributions in their abstracts but did mention them in the paper itself.

Another limitation is the use of Perplexity AI. This application has not been mentioned in research that explores the use of AI tools in literature reviews, which also was one of the motivations for testing the tool in this RR. Three out of the twenty-two results from the query in Perplexity were duplicates that were already found in the other queries, and three results were accepted to be used in the RR. That being said, a different approach towards the search of grey literature might have shown better results.

Although the BCM framework by Russo et al. (2024) shown in figure 2.6 is a great depiction of the components within BCM, some things are left unexplained. First of all, the paper by Russo et al. (2024) [57] lacks clearness how emergency response and crisis management relate to the rest of the model. The authors mention that these components are "two major research areas", but fail to explain why they were added as they are also excluded in the quantitative synthesis publications for the BCM components in table 1 [57].

One last limitation of this literature review is the lack of usage of specific BCM framework component terms in the search query of the RR. These terms could serve as alternatives to CR since they are part of BCM for which the fundamental aim of this field is to improve (cyber) resilience. Failing to include these terms could potentially exclude relevant studies that utilize the DT technology for the specific BCM component. In consequence, this may result in an incomplete review of the literature, thus limiting the comprehensiveness and accuracy of the findings. This might have increased results significantly, which could have caused problems for the researcher's personal schedule. This should be taken into consideration as well. Future literature reviews should consider including these terms to ensure a more precise exploration of the subject matter.

4.3.3 Recommendations

In this subsection, I will discuss the recommendations for further research within this research topic using research gaps identified from the papers or the results of this study.

An observation from the RR is the scarcity of studies specifically exploring the potential of DT, DS, or DM to enhance understanding of the organization, usage as an alternative to critical functions, and the design and implementation of BCPs. As the process of understanding an organization is aimed at creating a context to dive into the risk assessment and business impact analysis, the information gathered is often high-level without the need for a DT, DS, or DM. However, as a DT, DS, and DM provide a secondary environment that reflects the real system, this could provide organizations with an alternative to a critical function in case of an incident. This could also be seen as a recovery method, as the alternative to the critical function can keep systems running when used. Next to this, the integration of the DT, DS, or DM into a BCP could be explored to create guidelines on how to implement them. These suggested opportunities for future research can delve deeper

and uncover how DT, DS, or DM can be leveraged to further enhance BCM practices.

As can be identified from the results as well, there are some components in BCM where DT, DS and DM can be used. While existing studies have focused on specific aspects of DTs integration into BCM components, there remains an untapped potential for developing a comprehensive conceptual framework that encompasses the holistic integration of DTs, DSs, or DMs into BCM. Especially since the use cases proposed by many of the articles are aimed towards a DM, this is the primary technology that should be investigated. By addressing this gap in research, we can gain a deeper understanding of how DMs can enhance BCM practices across the board, offering valuable insights into optimizing resilience and response strategies.

Next to this, the results of this study show that DT, DS and especially the DM can be used for BCP testing, and exercising in many ways. Another exciting opportunity for future research is the development of an overview and guidelines towards using these technologies for the various ways to conduct BCP testing and exercises. This research gap can offer support to practitioners to understand which type of DT can be used for which exercise, answering the why and how.

Although some articles state that DT and DM can provide benefits to risk assessment, there is a gap in literature on the implementation and demonstration of this integration. A study that presents guidelines could provide significant benefits for practitioners interested in achieving the benefits stated by the articles in subsection 4.2.1.

One of the main concerns amongst scholars is the high cost of DT technology, due to its complexity of development and system-specific properties which makes it hard to generalize amongst more systems. To make this technology more economical, utilizing the option to use the system for more than one single use (even outside of cybersecurity) may lead to interesting opportunities. Noteworthy is that the properties and functionalities of the technology in unassociated technological domains might make the technology incompatible for interdisciplinary use. Future studies could look for similarities and develop guidelines for the interdisciplinary use of the DT to make the technology more economical.

Grasserli et al. [29] propose a methodology to implement a DT of industrial networks that can effectively and quickly assess, experiment and validate cybersecurity solutions. One use that the paper does not state explicitly, but might have potential, is the use of ad-hoc provisioning to be able to implement a cybersecurity solution rapidly in an IR scenario. They reduce the amount of time needed to do such provisioning from 3 minutes to 0.5 minutes [29]. This might be interesting for future research to look into.

4.4 Conclusion

This research aims to explore the current state-of-the-art on the potential of using a digital shadow to enhance cyber resilience through business continuity management, employing a rapid review and supplementary semi-structured interviews. Noticeably, the classification between a digital twin, digital shadow and digital model is rarely used within literature, which is why the digital twin and digital model are also included in this study. The results show that many studies tried to integrate the digital twin, digital shadow and digital model into various business continuity management activities. Especially the digital model was used often, due to its simulation capability. Moreover, many articles conducted their research in the context of an OT system, which was included in the proposed digital twin, digital shadow or digital model. The articles succeeded in showing the opportunity of using these technologies, in particular for *incident response and business continuity plan testing, maintenance and analysis*, and also moderately for *risk assessment, business im-*

pact analysis, business continuity management strategy, ICT strategy, business continuity training. Articles did however not cover the steps of *understanding the organization, alternatives to critical functions, business continuity plan design and implementation and crisis management.*

To answer the main research question, literature rarely reports on digital shadow implementations specifically. Only three articles explicitly used a digital shadow as the preferred technology. Two of which showed uses in incident response strategies for real-time monitoring and knowledge-based decision-making to enhance cyber resilience for organizations by Maia et al. [43] and Allison et al. [1]. The third article by Epiphaniou et al. [21] showed that a digital shadow could also be used for business continuity testing, maintenance and analysis through a cyber resilience assessment strategy. Noting that only three out of twenty-five articles explicitly mention that they use the digital shadow, it suggests a lack of use cases or limitations of this technology. On the contrary, neighbouring technologies (the digital twin and digital model) show many more use cases to enhance cyber resilience through business continuity management activities. The most likely reasoning is that a digital model provides a realistic simulation environment that can be utilized to evaluate realistic scenarios, in contrary to the digital shadow and digital twin which have to reflect the real-world situation continuously. For use-cases that would need such an environment, the power to adjust the real-world environment could be beneficial. Thus a digital twin would be preferred over a digital shadow. Noteworthy is that nine out of twenty-five articles did not explicitly mention the DT type that they actually used. Thus there may be more articles that have used the digital shadow, without it clearly being stated.

Although the number of studies and solutions towards the three technologies seem promising, many challenges were also identified. Most noticeably, we found three challenges that were faced most often. First, the high cost of developing and maintaining a digital twin, digital shadow or digital model is a challenge that was faced in many articles. Next to this, providing accurate behaviour of implementations, especially with high system complexity is a difficult task. Lastly, the cyber threats landscape that can be increased when implementing a digital twin, digital shadow and digital model should be taken into consideration. Specifically, the digital twin can extend the threat landscape specifically, due to its automatic data flow towards the real-world system. To make the technology an economical, safe and beneficial option is still a great concern amongst scholars, which may not be solved due to the complex nature of the technologies.

In light of these findings, further research is necessary to fully understand and leverage the potential of digital twins, digital shadows and digital models in enhancing business continuity management practices. A major research gap in the literature is the opportunity to use the digital twin, digital shadow and digital model as an alternative to a critical function in case of an incident. It would be especially interesting to investigate if a digital model, which can be changed without changing the real-world environment, could be used as an isolated environment to harden security and use it as an alternative to the production environment to enable business operations during an incident. A second use case for such an environment which could be investigated, would be to use it as an environment to test capabilities within business continuity plans that provide recovery strategies. For example, it is hard to get a grip on the time it will take to recover from a ransomware attack, this could provide teams with an environment to simulate such a scenario. Next to this, no studies delved into the practical implications, potential benefits, and challenges of implementing the different types of digital twins into business continuity management as a whole. A future study could research the opportunity to use one and the same digital twin, digital shadow or digital model for multiple business continuity management

processes, which can make the technology more economical. Furthermore, guidelines for the successful incorporation of the digital twin, digital shadow and digital model into the various ways to conduct business continuity plan testing and exercising could also provide benefits to practitioners.

Chapter 5

Problem Investigation

As we transition from the first phase to the second phase of the thesis, this chapter marks the beginning of the design science research. As thoroughly explained in chapter 1, one of the research gaps identified in the literature review serves as the foundation for this phase. Specifically, this research will address the gap concerning the lack of studies on a DM solution architecture that enhances multiple BCM processes, as discussed in section 4.4. Additionally, the potential for using DM as an alternative to critical functions will be explored. In this chapter, our goal is to gain a deeper understanding of BCM processes and investigate ways to improve them.

Due to the timeline of this project, we have scoped the research to investigate and improve three BCM processes. The problem contexts have been stated by Northwave, originating from their practical experience. These problems occur in the context of:

- Disaster recovery, which is part of the incident response component
- Disaster recovery test, which is part of the business continuity plan testing maintenance, and analysis component
- Business impact analysis, which is a component in itself

Every problem must be thoroughly investigated to gain a comprehensive understanding from both technical and process perspectives. During this investigation, three data-gathering methods were employed, as outlined in section 3.3.1. The first method involved a review of the literature, which is referenced in chapters 2 and 4, as well as within this chapter. The other data extraction methods were integrated into case studies, involving document analysis, informal conversations, and semi-structured interviews with experts to gain a deeper understanding of the general process. The results of this investigation are presented in four sections, corresponding to the motivational perspective of organizational resilience and the three problem descriptions. This chapter is descriptive, aiming to depict the current state through models that illustrate the situation as-is. Before proceeding, it is important to clarify some terminology that may be confusing. In the previous chapter, the term 'BCM component' was used to describe specific aspects of the BCM domain, following the framework by Russo et al. [57]. Moving forward in this design science research, we will be focusing on three BCM processes rather than three BCM components. As mentioned earlier, two of these processes are subsets of BCM components, while the business impact analysis is a component in itself. However, for simplicity, we will also refer to the business impact analysis as a BCM process. In each of these BCM processes, our goal is to identify and understand the challenges within them, which we refer to as 'problems', so that we

can later address them in the solution design during the design science research. These challenges exist within specific contexts, which we term 'problem contexts'. Each problem context corresponds to one of the three BCM processes, as each process provides its unique setting in which its respective challenges arise. While these contexts are considered distinct due to the individuality of the problems, they all fall within the broader BCM domain. Additionally, each BCM process is comprised of business processes that are detailed in the enterprise architecture models.

5.1 Organizational Resilience Motivation and Strategy

To understand how a DM can enhance the CR of organizations, it's essential to connect this with the underlying organizational motivation to prioritize resilience. By clarifying this perspective, we can better explain why organizations adopt specific strategies. Additionally, it's crucial to discuss the role of BC within this framework and its focus. This section will broaden the discussion from cybersecurity to the broader concept of organizational resilience. In essence, organizational resilience is risk-driven and involves implementing both preventive and repressive measures to minimize risk. BC, in particular, focuses on repressive measures, ensuring that a business can maintain high productivity during disruptions and swiftly restore normal operations.

5.1.1 Model Description

Based on the analysis of case studies and the literature, the organizational motivational and strategic perspectives have been synthesized and are presented in figure 5.1.

To gain a comprehensive understanding of the model, we should examine it systematically from top to bottom. At the top of the model, we represent an organization, which could be any entity. One key driver for organizations facing adverse events is the need for resilience [8]. 'Organizational Resilience' can be further dissected into 'Business Continuity' (BC), which focuses on maintaining business operations during disruptions and swiftly returning to a stable operational state, as detailed in section 2.3. This aspect is overseen by the 'BCM Officer', who is responsible for enhancing BC. Under the broader concept of 'Organizational Resilience', 'Risk Management' plays a critical role. This involves evaluating risks based on their likelihood and potential impact, and can be divided into 'Risk Likelihood Management' and 'Risk Impact Management'. Most organizations emphasize prevention, prioritizing 'Risk Likelihood Management'. However, since risks cannot always be completely avoided, organizations must also address 'Risk Impact Management'. Examining a layer further, risk assessments may reveal 'Unacceptable Risks' that exceed an organization's risk appetite. In such cases, the organization's objective is to 'Decrease Risks' that are deemed unacceptable. This involves two sub-goals: 'Decrease the Likelihood of Risks Occurring' and 'Decrease the Potential Impact of Risk Consequences', corresponding to the management of risk likelihood and impact. The primary strategy to influence these risks is to 'Apply Countermeasures', which can be classified as 'Repressive Countermeasures' or 'Preventive Countermeasures', as explained in section 2.2. Some countermeasures may serve both preventive and repressive functions, a concept that will be elaborated on in the following section. For simplicity, the model focuses on 'Decreasing Recovery Time,' a key aspect of BC. A practical example is the 'Implementation/Updating of a Disaster Recovery Plan,' which, when periodically tested, helps to 'Decrease Recovery Time.' This is a fundamental principle of BC, as discussed in Section 2.3, and represents a repressive measure used specifically during adverse situations.

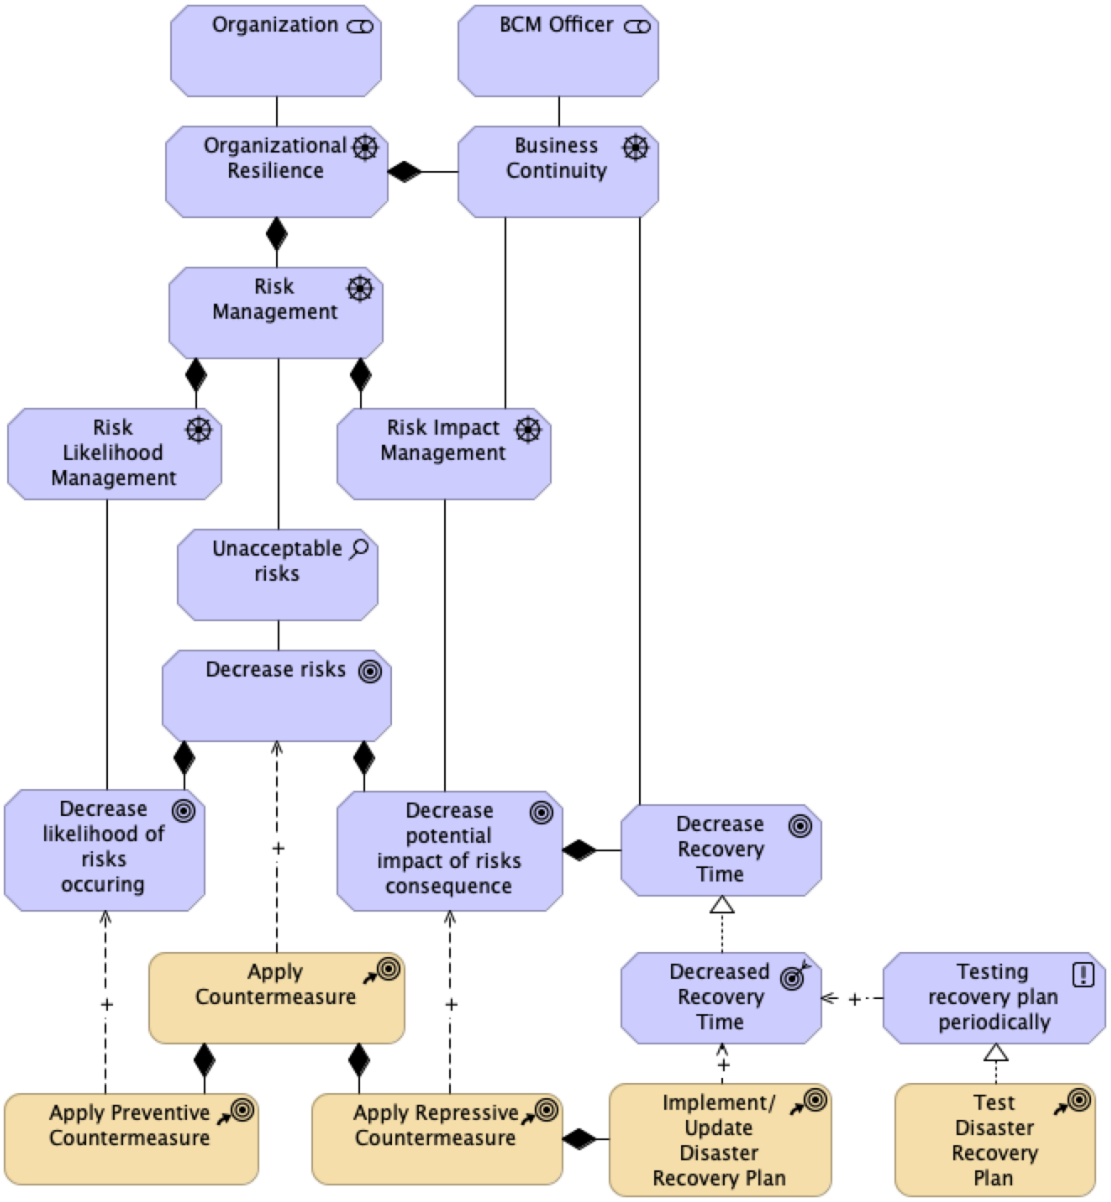

FIGURE 5.1: Organizational Resilience Motivation & Strategy Viewpoint

To gain a comprehensive understanding of the model, we should examine it systematically from top to bottom. At the top of the model, we represent an organization, which could be any entity. One key driver for organizations facing adverse events is the need for resilience [8]. 'Organizational Resilience' can be further dissected into 'Business Continuity' (BC), which focuses on maintaining business operations during disruptions and swiftly returning to a stable operational state, as detailed in section 2.3. This aspect is overseen by the 'BCM Officer', who is responsible for enhancing BC.

Under the broader concept of 'Organizational Resilience', 'Risk Management' plays a critical role. This involves evaluating risks based on their likelihood and potential impact, and can be divided into 'Risk Likelihood Management' and 'Risk Impact Management.' Most organizations emphasize prevention, prioritizing 'Risk Likelihood Management.' However, since risks cannot always be completely avoided, organizations must also

address 'Risk Impact Management'.

Examining a layer further, risk assessments may reveal 'Unacceptable Risks' that exceed an organization's risk appetite. In such cases, the organization's objective is to 'Decrease Risks' that are deemed unacceptable. This involves two sub-goals: 'Decrease the Likelihood of Risks Occurring' and 'Decrease the Potential Impact of Risk Consequences,' corresponding to the management of risk likelihood and impact. The primary strategy to influence these risks is to 'Apply Countermeasures,' which can be classified as 'Repressive Countermeasures' or 'Preventive Countermeasures,' as explained in section 2.2. Some countermeasures may serve both preventive and repressive functions, a concept that will be elaborated on in the following section. For simplicity, the model focuses on 'Decreasing Recovery Time', a key aspect of BC. A practical example is the 'Implementation/Updating of a Disaster Recovery Plan', which, when periodically tested, helps to 'Decrease Recovery Time'. This is a fundamental principle of BC, as discussed in section 2.3, and represents a repressive measure used specifically during adverse situations.

This can all seem a bit conceptual, so let's take an example to ensure the model is understood. By way of illustration, let's say organization X owns an office building and manages the risks that are tied to the office building. In this example, the organization needs the building to enable its core business operations, thus being a single point of failure. One risk could be that some old electronic components, which are used for business operations, are known to cause fire sparks. At other organizations, this spark has started fire hazards. Let's assume the organization does not have any countermeasures in place, and they identify that this risk is unacceptable to the organization. A goal would be to reduce this risk. To achieve this, a course of action could be to apply a preventive measure, to reduce the likelihood of the risk occuring which contributes to the decrease of the risk. Considering this example, the electronics might need replacement, which can be seen as a preventive measure. However, this might not diminish the likelihood of a fire hazard completely, and a repressive measure must also be applied to make the risk acceptable. For example, putting a fire distinguisher close to the electronic component can reduce the impact when a fire hazard just started. Furthermore, a countermeasure can be preventive and repressive, like a temperature sensor. If the electronic component gradually heats up towards a fire spark, it can be prevented by monitoring the temperature. If there suddenly is a fire spark which causes a fire hazard, monitoring can enable employees to act quickly, reducing the impact of the fire. Another way to reduce the impact is to reduce the recovery time. One course of action could be to put a plan in place of who should use the distinguisher, and how they should use it. Finally, a principle that influences this decreased recovery time, is to test the recovery procedure periodically, to ensure the plan still provides a sufficient recovery time and to train human capabilities.

5.1.2 Summary

The results of this investigation show how organisations deal with resilience, what course of action can be taken, and specifically how the fields of BC and disaster recovery relate to these concepts. This diagram will be used throughout this chapter to explain why the three problems are relevant to organisational resilience and how they can contribute the the field of BCM.

5.2 Disaster Recovery

The first use case that we are going to evaluate is the disaster recovery process of an organization in the context of cyber security. From the data extraction method, the problem has been identified that it is hard to provide organizations with repressive measures that aid disaster recovery capabilities, especially during a ransomware attack. As mentioned in the conclusion section 4.4 of the literature review, there has been no attempt by researchers to use a DM as an alternative production environment in the case of an IT disaster. To illustrate this opportunity in more detail, we will first explain how the process of disaster recovery works. Additionally, an IT infrastructure viewpoint has been provided to show the technical perspective of a recovery.

5.2.1 Goal

The purpose of disaster recovery in the context of cybersecurity is to restore critical IT infrastructure and operations following a major disruption or disaster to ensure BC. Its main focus is on providing availability of services, so critical business operations can be continued.

5.2.2 Model Description

Infrastructure Viewpoint

Explaining the results of the business process viewpoint of a disaster recovery procedure requires us to first explain what a typical IT infrastructure of an organization looks like. Organizational IT infrastructures can vary significantly based on the type and size of an organization. For instance, a small business with few dependencies on software programs can rely on basic setups with limited servers and software, while a large complex enterprise will rely on complex networks, large data centres, multiple locations and numerous applications. Despite these differences, many parts of IT infrastructure are universal, such as reliable networking, secure data storage, computational power, and identity management systems. These findings should be considered in the evaluation of the infrastructure viewpoint we will discuss. However, this case can be considered a common scenario. The infrastructure perspective of the problem context will be explained according to figure 5.2.

In this case, the IT infrastructure includes multiple physical locations. Let's first look at the network infrastructure. Each location is equipped with its own firewall and organizational location network through which communication flows. For simplicity reasons, the infrastructure details of the second location have been excluded in general. Communication that comes from the internet flows through the 'Firewall', which protects the internal network by monitoring and controlling incoming and outgoing network traffic based on predetermined security rules. It tries to keep the location infrastructure secure and isolated from potential threats.

The main location usually hosts a virtual machine (VM) cluster to run various applications that are used within business operations. We describe this cluster with the 'Virtualization' node. VMs are software-based emulations of physical computers that run operating systems and applications. Each VM operates on a 'Hypervisor', of which there are often multiple. These are specialized hardware or firmware solutions that facilitate computing power and software to create and manage VMs on an operating level. As such, the hypervisor does not necessarily manage what application runs on a VM but rather manages for example how much computing power and storage bandwidth it gets. Here

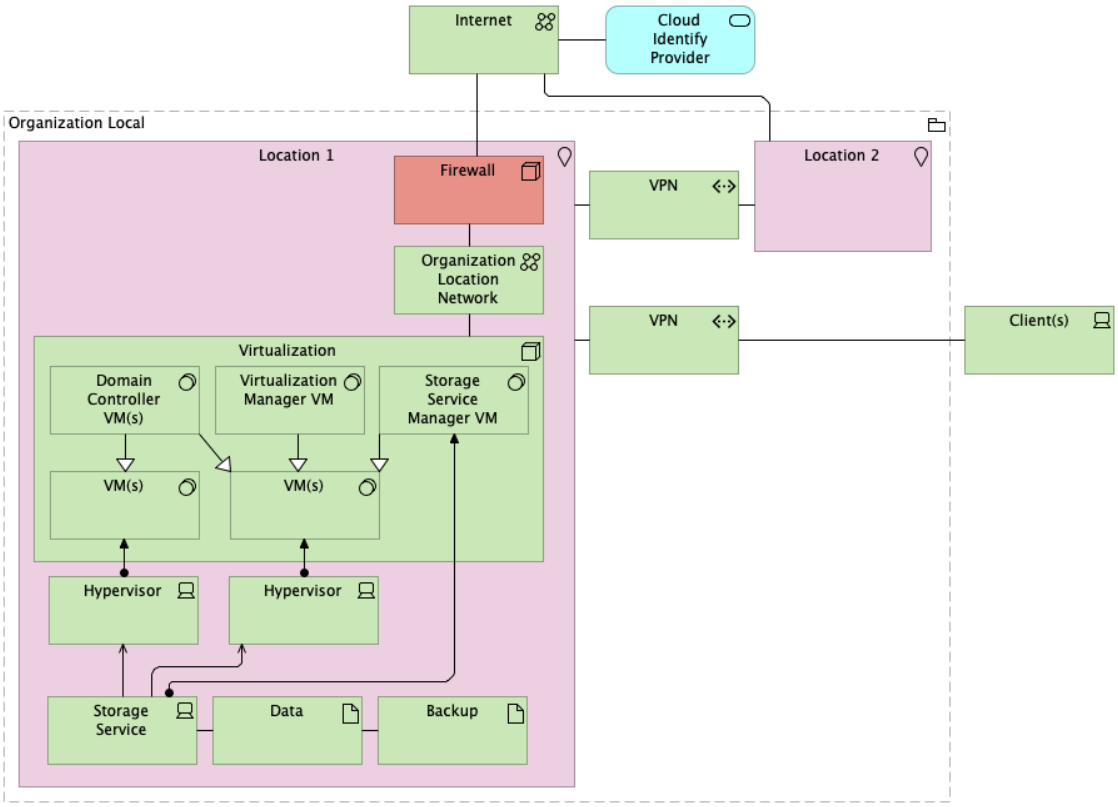

FIGURE 5.2: Disaster Recovery Infrastructure Viewpoint

we are depicting bare-metal hypervisors which include both hardware and software in one solution. To model this in our architecture viewpoint, the hypervisor assigns certain VMs and can manage the load on each of the VMs. These VMs can then run any application but are often dedicated to one system purpose (like a single application). We depict three essential VMs that are essential in any IT infrastructure of an organization which we will explain below.

To manage the hypervisors (including what VMs should run on them), organizations use a VM management application here depicted with the 'Virtualization Manager VM'. This application provides tools for managing the lifecycle of VMs and connects with the hypervisor. It enables provisioning, configuring, monitoring, and maintaining VMs, ensuring that VMs operate smoothly and scale according to the organization's needs. This setup maximizes hardware utilization, improves efficiency, and provides flexibility for deploying and managing various applications. As this application is running on a VM itself, this begs the question of whether it also manages itself as well. The answer is that this VM is the first that needs to be set up in the initial setup, after which it can be configured to be started first in a reboot.

Next to this, a domain controller (DC) manages the network security and user access. A domain controller is a server responsible for security authentication requests within a Windows Server domain. It serves as a management application of a directory and identity management service system that provides centralized management and security of user identities and network resources within an organization. It is a key component that handles the centralized management of network resources such as user accounts, computers, and printers. Often this is used in a hybrid solution combined with cloud identity providers,

like Entra ID by Microsoft. The domain controller is very powerful and it often plays a crucial role in ransomware attacks to execute the ransomware.

Lastly, next to the management of the virtualization, in this infrastructure, a 'Storage Service' is shared amongst the hypervisors. This way, a hypervisor can take over the activities almost instantly when the other hypervisor has some kind of failure. This storage also needs to be managed through a storage service manager, which run on a VM as well. The storage service stores data, making it accessible to authorized applications and users across the network. There are various ways of setting up a backup environment, however, often this is also managed by the database service.

Often, people require secure access to the internal organizational network from various locations, including remote offices, home offices, or while travelling. Virtual Private Networks (VPNs) are widely used to facilitate this secure access.

Business Process Viewpoint

From the data, the following business process model was constructed as shown in figure 5.3. In this case, we take a ransomware attack as the scenario.

In the context of cybersecurity, the process begins with a cyberattack on an organization, resulting in an 'IT disaster.' In such a scenario, the attacker gains access to the system through a vulnerability, such as an unpatched exploit, and obtains sufficient system privileges to either extract data or deploy ransomware to lock systems within the organization. Once the victim becomes aware of the attack, the organization must investigate how the threat actor gained access and initiate the recovery process according to the disaster recovery procedure.

In disaster recovery, the focus is solely on restoring the IT infrastructure. The first step is to contain the breach, which typically involves isolating the network by disconnecting VPNs, internet connections, and any other external links. Referring to the infrastructure depicted in Figure 5.3, this would involve severing the connection between the 'Internet' and the 'Firewall,' as well as disabling the VPN connections. Once containment is achieved, the 'Disaster Recovery' process can begin, led by the organization's 'Incident Response Team'. Firstly, two steps need to be taken in parallel which are the 'Pre Recovery Procedure' and the '(if unprepared) Create Application Priority List'. An 'Application Priority List' is required during the disaster recovery process so the IR team knows which VMs to recover first. This is based on the priority of applications which often corresponds to a list of applications required for critical business processes. Organizations occasionally created this list prior to the disaster, but often needs to be created during the incident. This list includes the 'VM(s)' that need to be recovered, which 'Business Process' is supported, the 'Location' of the recovery and all the 'Application Details' which are required for recovery (like the IP addresses of the applications). In parallel, the 'Pre Recovery Procedure' will be started.

This 'Pre Recovery Procedure' includes steps that are required to be taken before recovery of VMs can even start. The first step is to 'Prepare safe infrastructure'. This can include various steps depending on the IT infrastructure of an organization like reinstalling hypervisors and other hardware systems. Afterwards, the 'Disaster Recovery Secure Environment Setup' process can be started. In this process, the IR team needs to 'Set up isolated environment'. The isolated environment is needed to eradicate the threat actor from each VM while being able to run some applications on production in parallel. Understanding this step is essential to understand an efficient disaster recovery procedure.

Large organizations often have many VMs that they need to restore. However, not all of them are needed for the critical business functions. Organizations often want to

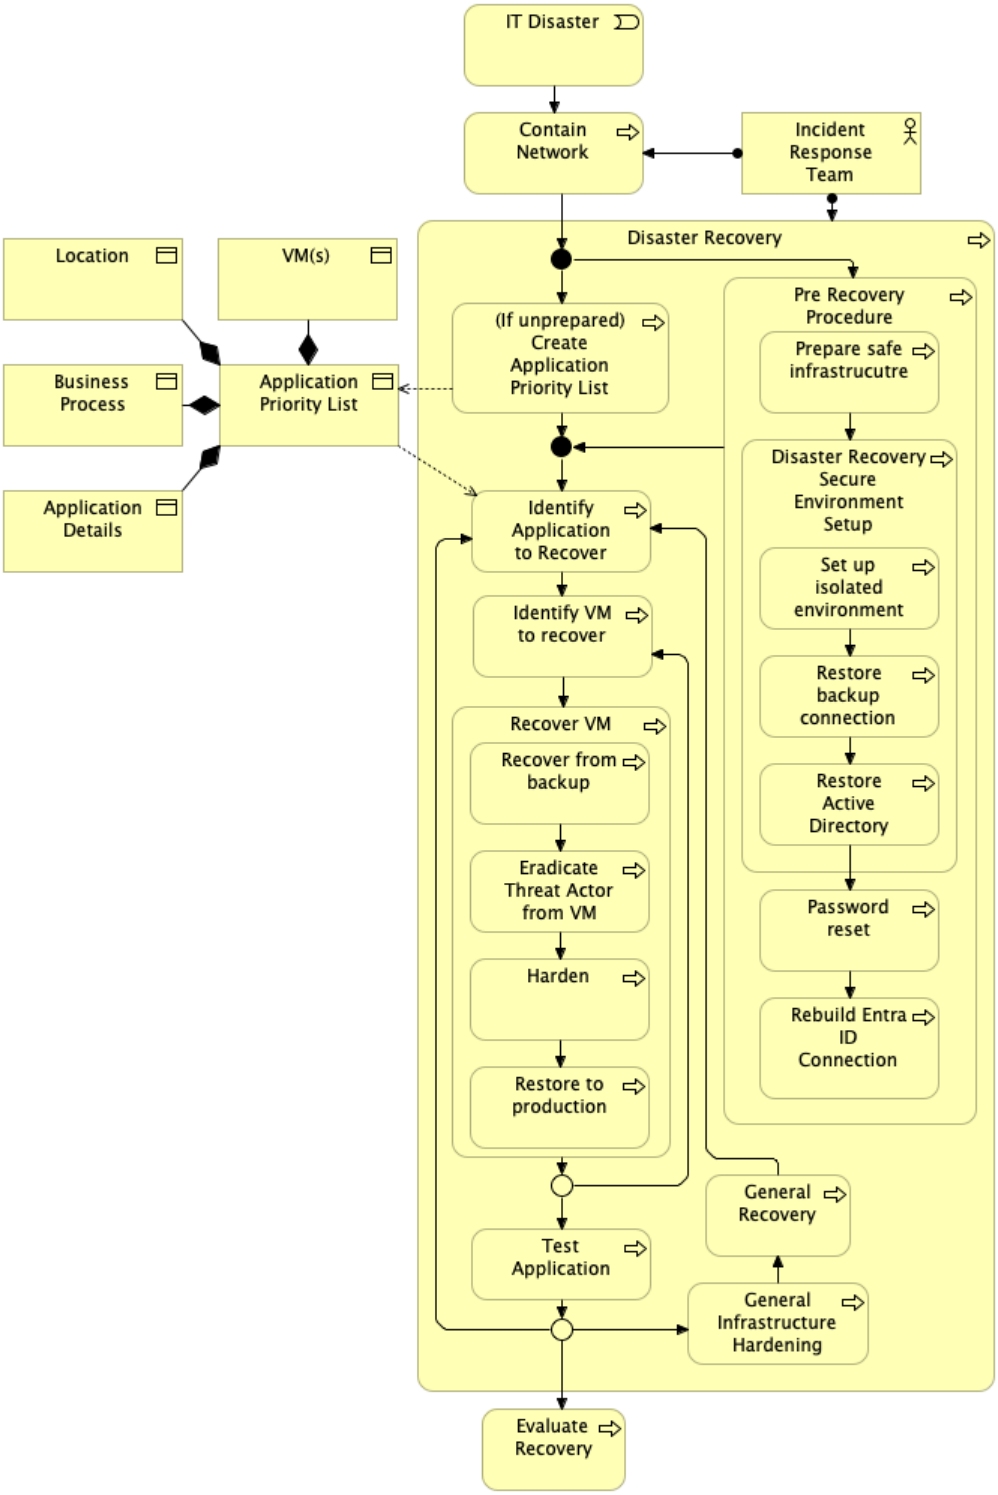

FIGURE 5.3: Disaster Recovery Business Process Viewpoint

restore their applications for critical business functions first, before all the other VMs are recovered. This requires these VMs to have internet access, as most applications rely on collaboration with other applications to function properly. If this were done on a single network environment, these VMs could function properly again, but removing the threat

actor from the other VMs would become dangerous, as the threat actor might gain access through the network again via a backdoor in one of the VMs that needs to be cleaned. This makes the separate isolated environment essential, as it provides a safe virtual space for threat eradication, while the cleaned VMs can go to production to provide the organization with critical business functions applications again.

After the isolated environment has been set up, 'Restore backup connection' needs to be established. How this is done depends heavily on the backup type, location and vendors. It can either be on-premise or in the cloud, which in the latter case can limit the recovery speed. Finally, 'Restore Active Directory' needs to be executed. As mentioned, the domain controller is essential in any organizational IT infrastructure, as it controls authentication, authorization, directory services and more. In a typical ransomware attack, the attacker has gained access to the accounts in the domain controller to execute its ransomware. Because of this, a 'Password reset' is required for all accounts to ensure the attacker cannot access the environment anymore. This is often quite a troubling process, as all users need to change their passwords after it has been reset for every user. While users change their passwords, the recovery environment needs to 'Rebuild Entra ID Connection', which will sync with the password updates so both are aligned again. This can sometimes take place later in the process, but this can already be done at this point.

After this process has been finished, applications can be recovered from the backups according to the priority list that was set up in parallel. From this list, the IR team will 'Identify Application to Recover'. For each application, there are a certain amount of VMs required, which will be restored one by one until each one is recovered. One by one the IR team will 'Identify VM to recover' to facilitate this process. With a large team, this can also be done in parallel. Then the 'Recover VM' process can start. In this process, the VM will be 'Recover from backup'. After restoring the backup, the other steps are essential to recover the production environment safely. In a cyber attack, a threat actor often installs a backdoor in some of the VMs. This consequently requires organizations to check each VM for such a backdoor and eradicate the access of the threat actor if one is found. Thus, the 'Erdicate Threat Actor from VM' will take place. Next to this, there might be VMs that run applications that have not received security updates. 'Harden' step will be applied to ensure these receive updates. Finally, the 'Restore to production' environment process will be executed. This is a quick process, as it only requires network configuration changes of the VM to be changed.

This process will be done for each VM (this is where the OR junction flows back to 'Identify VM to recover') required for an application until all are restored. Finally, the application needs to be tested to evaluate if it is functioning correctly. As the network configurations of an application are often based on the production environment, applications only work properly there. Only when they are restored to production, they can be tested. This will be done for every application (this is where the OR junction flows back to 'Identify Application Priority List') and its VMs according to the priority list until all critical business function applications are working properly.

Finally, after the critical business functions applications are restored, 'General Infrastructure Hardening' will be executed to prevent another attack. This also takes into consideration the investigation process that identifies where the threat actor gained access as mentioned at the start of this sub-section, which is often done parallel to the disaster recovery process. This can also include setting up security monitoring to prevent an attacker from escalating if a backdoor is missed. When this is done, the production environment can be put online to enable critical business functions. Afterwards, all the other VMs will be recovered in the isolated environment and returned to production. In figure 5.3, we

specifically address this case. If an organization only has critical business applications or does not want to recover all applications, this step will be omitted. Lastly, after every application has been restored, the disaster recovery process will be evaluated.

5.2.3 Problem Description

Looking back at the motivational viewpoint, any repressive measure aimed at supporting the disaster recovery process will reduce the impact of risks. Being able to recover from a disaster is defined by human capabilities and system capabilities. System capabilities, like an up-to-date backup strategy, are essential to restore systems. From the human perspective, a sufficient amount of knowledge about their IT architecture and technical know-how is needed to recover systems. Both capabilities can be enhanced by implementing repressive measures. The main challenge of this section focuses on the lack of system capabilities.

As shown in figure 5.3, many processes need to be taken to set up a separate isolated environment to enable the application recovery process. This pre-recovery procedure takes some time, depending on the system and human capabilities. In a ransomware scenario, every hour counts. This makes it crucial to implement repressive measures to ensure quicker recovery time. The main challenge is that few repressive measures can speed up the recovery process from a system capability perspective, especially for the pre-recovery procedure. This is because, for example in a ransomware scenario, the production environment will be encrypted. Some examples of repressive measures are quick, safe and reliable backups, which can decrease the time it takes to recover backup data. However, in the whole disaster recovery, there are many manual steps to interact with systems to recover. A second opportunity to increase recovery speed might be to use the isolated environment as the new production environment, instead of taking the last step in the application recovery process.

Next to this, during the process of recovering IT systems, organizations are often unprepared to have an understanding of which IT systems need to be restored to enable a critical business function. If an organization is thoroughly prepared, then this should be available. However, developing such a list and testing it is often very challenging. Understanding the dependencies of critical business functions is a responsibility of the business impact analysis, which will be explained in section 5.4. A takeaway should be that this process delays recovery by some period of time, depending on the size and type of organization. An estimate would be that it takes a couple of days for a large organization according to a Northwave employee.

5.3 Disaster Recovery Test

The second problem we will investigate is the disaster recovery test, which is highly related to the disaster recovery process. As mentioned in section 5.1, the disaster recovery test plays an essential role as an amplification of the effect of a disaster recovery plan as a countermeasure. Similar to the previous section, we will first explain the goal and process of the problem context, before we explain the problem itself.

5.3.1 Goal

The goal of a disaster recovery test is to ensure that the recovery of critical business operations aligns with organizational objectives after a disruptive event. The test evaluates if an organization can deal with a certain scenario according to its disaster recovery plan.

Scenarios are built to test procedures in the disaster recovery plan. After a test, the evaluation may be used as an enabler for a migration plan towards improved human capabilities, technical capabilities, or disaster recovery plan.

5.3.2 Model Description

The following model of the disaster recovery test business process was constructed as shown in 5.4.

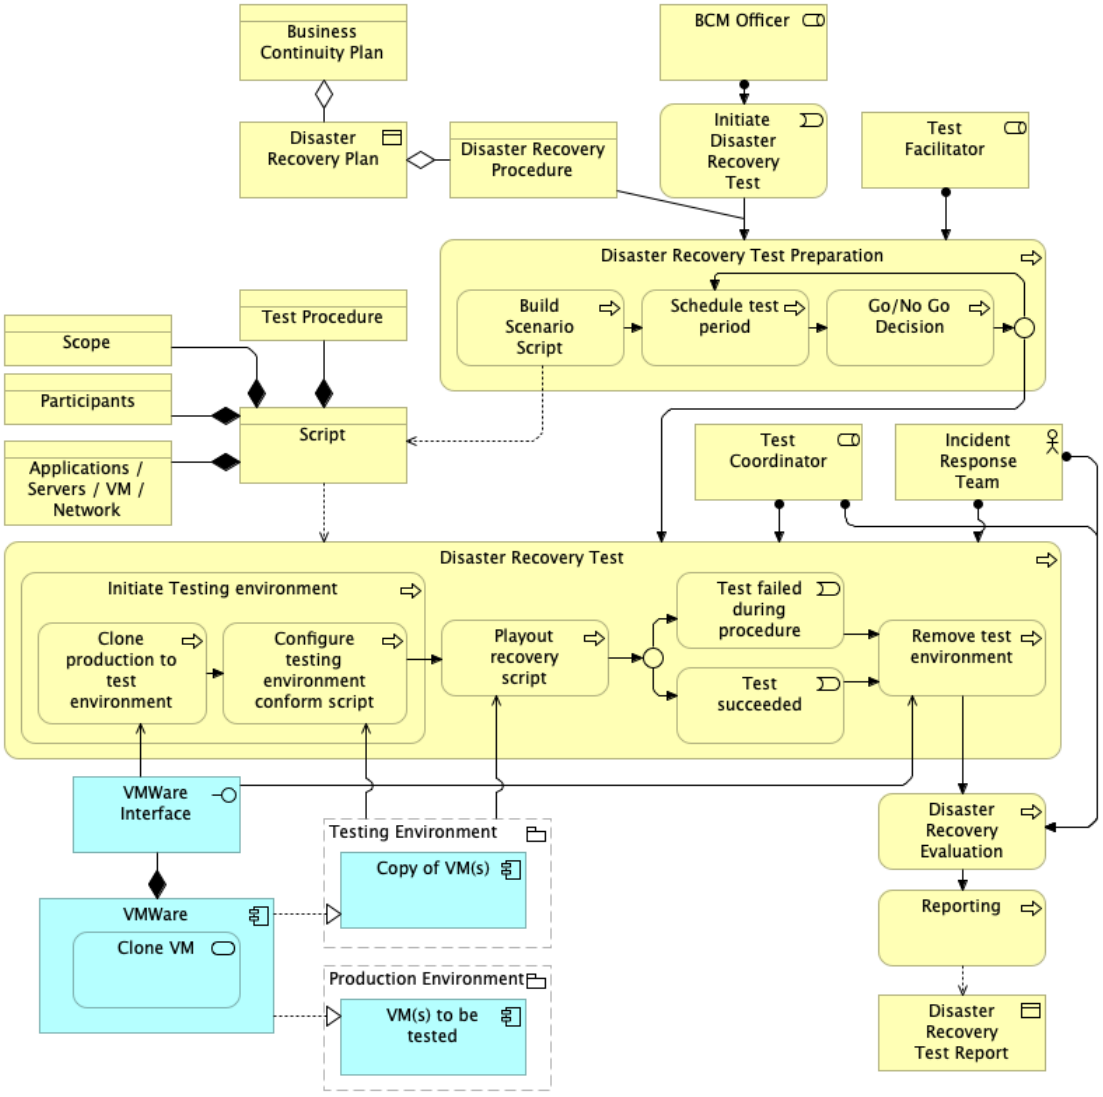

FIGURE 5.4: Disaster Recovery Test Business Process Viewpoint

The diagram will be explained from top to bottom. As an initial business event, the 'BCM Officer' will 'Initiates the Disaster Recovery Test' to assess whether the 'Disaster Recovery Procedure' can be effectively executed. This procedure is typically a component of the 'Disaster Recovery Plan', which itself is a part of the broader 'Business Continuity Plan'. Initiating the disaster recovery test triggers the 'Disaster Recovery Test Preparation' process, where all necessary elements for the test are organized. This includes creating a 'Script' that outlines the scenario to be tested, the 'Participants' involved, the 'Scope' of the test, and the specific 'Applications/Servers/VMs/Network' that will be utilized. This

preparation is primarily overseen by a 'Test Facilitator,' often the BCM Officer. After setting up the script, the test will be scheduled and there will be a 'Go /No Go Decision' process one day before the execution, depending on the conditions of the IT infrastructure. If there are any other problems around the IT infrastructure, the 'Incident Response Team' might have other obligations and the test will need to be rescheduled (flowing back from the OR junction to 'Schedule test period'. Otherwise, the 'Disaster Recovery Test' process itself will be started.

The test is conducted by the 'Incident Response Team', under the supervision of the 'Test Coordinator'. The script provides detailed instructions and delineates the steps to be executed throughout the test. The process starts by initiating the test environment. As stated earlier, the goal of the test is to test the recovery of certain IT applications, which requires various other systems to be built as well. For example, a Domain Controller as mentioned in figure 5.2 needs to be working and backups connections need to be established.

Next to this, the VMs which run an application often have certain network configurations that are essential to make them communicate with each other correctly. This means that the copies cannot be set up in the production environment, due to its potential disrupting effect towards existing VMs. Consequently, depending on the test, the VMs targeted for testing must be restored to a separate network environment that is a clone of the production environment. As mentioned in section 5.2, setting up such an environment is often part of the recovery strategy, making it an essential step which can be tested. However, it depends on the goals of the test and the test script whether or main this is the main focus. A key difference between the disaster recovery procedure and the DRT is that applications in the test environment cannot be deployed in production due to potential disruption. In a real recovery, applications are first restored in production before being tested. Therefore, the test environment's network configuration must replicate the production environment to accurately assess the application's functionality post-recovery. The feasibility of this replication varies depending on the IT infrastructure's capabilities.

The testing environment will be configured by the IR team so it provides sufficient components that can be recovered. The parts that need to be configured depend on the procedure that is being tested. For example, if the goal of the test is to evaluate if the IR team can set up a connection with the backup environment to restore an application from backups, then the backup connection should be removed, but VPN connections can still be active. When all the required systems are configured accordingly, the IR team can start the 'Playout recovery script' process. This procedure can either succeed or fail. When some part of the procedure is unable to be executed, this might cause the following part to be untestable. If so, the test can be stopped. In either case, the testing environment will need to be removed again by hand through the VMWare interface. When this has been finished, the 'Disaster Recovery Evaluation' will be started by the IR team and the test coordinator. This will be reported afterwards, upon which improvements and decisions can be made for the procedure. This will be reported in the 'Disaster Recovery Test Report'.

5.3.3 Problem Description

Looking back at figure 5.4, the main challenge is primarily related to the 'Initiate Testing Environment' process. Building a separate realistic testing environment based on script requirements can take much time. Especially as the setup needs to deal with many constraints on hardware resource usage, network settings that need to be configured, and invalid data manipulation to name some. The setups that are currently used are often not similar to real disaster situations, making the test less accurate. Next to this, due to the limited resources, only small parts of procedures are tested, which is less of a holistic

approach. Northwave would like to be able to test larger parts or the whole of the disaster recovery plan to have a more accurate and realistic test that can improve the capabilities and shorten the recovery time.

Connecting the problem to the motivation and strategy viewpoint in subsection 5.1, if the tests of the recovery plan can be done more realistically and cover a bigger part of the disaster recovery plan (so multiple procedures in one test), this could enhance the result of having a decreased recovery time.

5.4 Business Impact Analysis

Results from the problem investigation methodologies show that the business impact analysis is highly related to risk assessment and supports decision-making efforts on risk countermeasures. As a preparation step in the BCM field, it plays a valuable role in understanding organizational business processes and their dependencies. These results will be highlighted in this section.

5.4.1 Goal

One of the main goals of the business impact analysis is to understand the inter-dependencies within critical business operations and IT infrastructure. The results should note if dependencies are single points of failure in the case of systems or single points of knowledge in the case of humans or data. These dependencies can directly cause problems when the dependency fails to operate. Next, from this knowledge BC requirements can be stated. These are mainly related to recovery time and the data loss period. This will be explained later in this section.

As mentioned in section 5.2, many organisations experience a lack of grip on their IT infrastructure and especially its dependencies. When a response procedure needs to be executed, organizations need to know how they can rebuild their IT infrastructure to support their core business operations. Additionally, this is relevant for understanding the consequences of a discontinuity in one part of the IT infrastructure, and the effects on the rest of the business operations.

5.4.2 Model Description

To get a better understanding of the problem, the business process viewpoint was developed shown in figure 5.5.

The diagram will be explained from top to bottom, starting at the top-left 'Initiate Business Impact Analysis' event. A 'BCM Officer' is responsible for scheduling and starting the business impact analysis, which will be executed by at least the 'IT Manager' and 'Operational Manager' with guidance from the BCM officer. As most organisations have intricate and complex business operations, the BCM officer is also responsible for choosing a specific part of the business operations that they will analyse. Most often this is a 'Critical Business Function' below the initial event. For this research, we will use the term business function to describe a collective part of business operations which is built out of business processes that explain small parts of a business function. The first step in the business impact analysis of the critical business function we want to assess is to understand the business function and its processes. So the first step is to 'Analyse Critical Business Function' through a discussion. This will end up in a 'Business Function Diagram' which is part of the 'Critical Business Function', which is either described already beforehand or will be developed during the analysis at 'Develop Business Function Diagram'. To develop

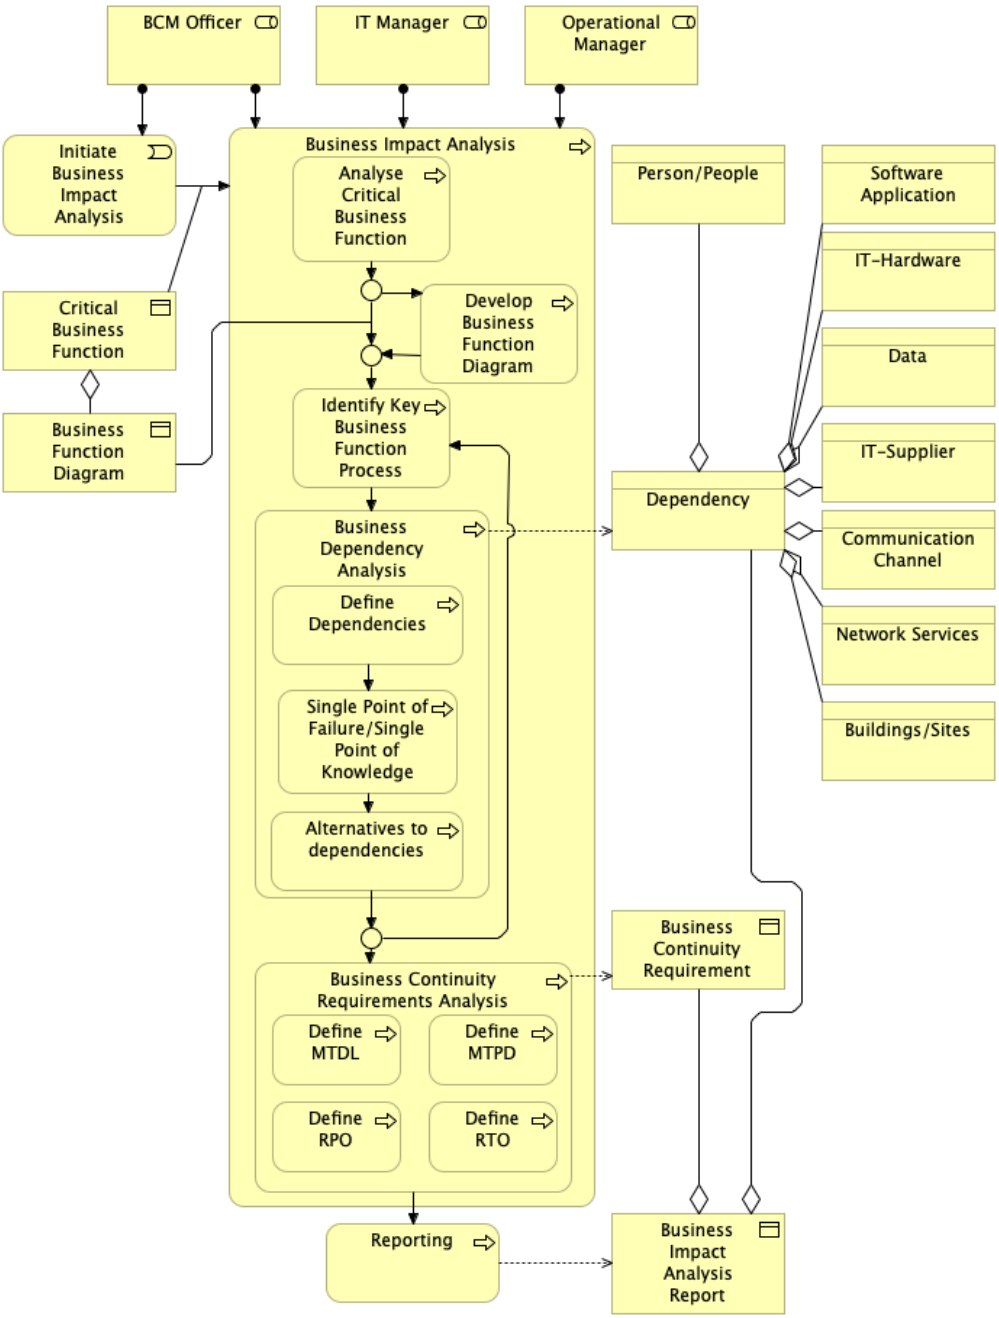

FIGURE 5.5: Business Impact Analysis Business Process Viewpoint

this diagram many steps need to be taken, which are irrelevant to the problem of this BCM process, so they are excluded from this model.

Continuing the business impact analysis, the members of the analysis need to 'Identify Key Business Function Process' of which a 'Business Dependency Analysis' will be conducted. This will be done for every key business function process. In this process, the 'IT Manager' and 'Operational Manager' will think of the dependencies they acknowledge for the business function process in 'Define Dependencies'. This 'Dependency' will be written down as shown to the right of the model, which can be of any type that is aggregated from 'Dependency'. For every dependency, it will be decided if it is a 'Single

Point of Failure/Single Point of Knowledge', depending on if the dependency is a system or datapoint respectively. A person can also hold specific knowledge which can be a single point of knowledge. After this process, 'Alternatives to dependencies' will be identified to determine if a dependency can be executed alternatively. Implicitly, this is only the case for dependencies that are not a single point of failure or knowledge. After all the business function processes have been evaluated, the 'Business Continuity Requirement Analysis' can be started.

In this final part, the recovery time objective (RTO), maximum tolerable period of disruption (MTPD), recovery point objective (RPO) and maximum tolerable data loss (MTDL) will be defined, which were explained in 2.3. While this process includes additional steps and the objectives depend on numerous variables, these details are not relevant to the problem context and will be excluded from this model. Finally, the analysis will be reported in the 'Business Impact Analysis Report', which includes the dependencies and BC requirements which were obtained during the analysis.

5.4.3 Problem Description

The main problem which was mentioned in interviews with Northwave employees is the complexity of identifying dependencies within the IT infrastructure. Not only does this make it harder to assess risks accurately, but it also brings about a potential knowledge gap in disaster recovery procedures. As mentioned in section 5.2, it can be of big value in a crisis situation to understand system dependencies to recover quickly. A scenario might arise where recovery of a certain IT system is needed, but some required servers, databases or networks lack from the disaster recovery procedure, causing a delay in recovery time. Organizations rarely have a database of their IT infrastructure dependencies at the ready. Mostly organizations with a high resilience maturity make these preparations. However, this only considers customers of Northwave, so the findings might not be transferable to every organization. To conclude, currently, it takes many resources to understand the interdependencies of IT systems in IT infrastructure, making it difficult to understand risks and write accurate disaster recovery plans.

5.5 Conclusion

In examining the three problem contexts of disaster recovery, disaster recovery test and business impact analysis, we have uncovered comprehensive insights into the core of each problem and the potential benefit of solving them. Collectively, these investigations not only provide a clearer understanding of the problems but can also be interpreted through the organizational motivation and strategy viewpoint as depicted in figure 5.1.

Firstly, our investigation into disaster recovery revealed that organizations often experience difficulties setting up an isolated recovery environment which delays the disaster recovery process. Also, the identified research gap in the literature review (section 4.4) to use the DM as a new production environment presumably does not offer many benefits. Moving a VM from an isolated and secure environment does not take much time. Next to this, it is even required to have at least two environments to recover critical business function applications first, before the non-critical applications.

Secondly, the analysis of the disaster recovery test revealed that it takes much time and resources to set up a realistic separate environment to conduct a disaster recovery test for specific procedures. Some disaster recovery tests might include this process, but some tests aim to evaluate some other specific disaster recovery procedure. Especially

because this step is repetitive for every test that does not include setting up a realistic separate environment, potential improvements can be made. Another problem is that applications often require specific network configurations that correspond to the production environment to work correctly. Being able to test if an application functions properly after a disaster recovery procedure requires a separate environment with the same configurations. This is also difficult to produce and often requires specific hardware or software.

Lastly, investigating the business impact analysis showed that organizations often have a lack of understanding of IT interdependencies. This makes it hard to understand the vulnerabilities of systems and which systems need to be recovered to enable certain business operations. This creates a lack of understanding of how long recovery times are and what consequences system disruptions have, making it challenging to evaluate risks effectively.

In this chapter we have assessed the problems independently from each other to get a thorough understanding of each problem and what happens if they are improved. However, when looking at these problems collectively, a connection between each problem can be recognized. The interaction between each problem was already briefly introduced in each section. As mentioned in section 5.2, improving the business impact analysis can not only improve the risk assessment accuracy but also shorten the disaster recovery process by providing a list of key IT interdependencies. This is a quite straightforward relationship between the two problems, where the results of the business impact analysis can be used within the disaster recovery as well. However, when relating the disaster recovery problem with the disaster recovery test problem, a different situation is created. The disaster recovery test problem addresses the lack of a realistic testing environment to test system and human capabilities. A DM might prove to facilitate such an environment wherein a cyber attack can be simulated to test if an organization conduct their disaster recovery procedures. However, if the proposed DM would also be used to improve the disaster recovery process, the disaster recovery test procedure would change as well. The separate test environment used in the disaster recovery test would become part of the disaster recovery test as it would also be used in a disaster recovery scenario. The creation of this environment still needs to be tested, next to all the other steps in the disaster recovery procedure. So this still requires practitioners to conduct the test of the disaster recovery procedure and even provides the opportunity to evaluate the recovery speed.

The proposed solution will investigate whether a DM can be used as a repressive measure, to improve risk assessments and shorten disaster recovery times. In the next chapter, we will discuss the requirements of the proposed solution, the conceptual model, and how it can be integrated into these BCM processes.

Chapter 6

Proposed Solution

In this chapter, we will design and introduce the proposed solution of a digital model architecture to improve the BCM problem context described in the previous chapter. By investigating the problem contexts identified in the previous chapter, we derive specific requirements needed to achieve the intended goals of improving the BCM processes. This chapter will provide a thorough explanation of the proposed solution, addressing its goal, requirements, design choices, functional models, architecture models, and theoretical implications.

The DM solution that will be proposed is not just a single software system or application; it is an integrated framework comprising multiple software systems to provide the required functionalities of a DM. Considering this, the solution can be better described as an architecture. This is why from this point forward we will describe this solution architecture as the Resilience Digital Model Architecture (RDMA). It serves as a conceptual solution umbrella, adaptable to various implementations using different automation of software and hardware systems. For instance, while this chapter might reference the Active Directory by Microsoft, the solution is conceptual and can integrate with alternative vendors as well. Or the multiplicity of methods to create an isolated network environment.

Throughout this chapter, we will explain how the RDMA works and what it aims to achieve, detailing its requirements, the functions it provides, and how it integrates with specific business processes to enhance them. We will also discuss the infrastructure design and potential challenges that may arise when implementing this solution design.

6.1 Requirements

To effectively develop a DM aimed at enhancing BCM processes, we need to establish a clear set of requirements. This section outlines the requirements that the solution must meet to improve the identified problem contexts. These requirements are derived from the insights gained during the literature review and problem investigation. To improve each of the problem contexts, every one of them will have its requirement list. These will all need to be considered for the solution architecture that needs to improve all three problem contexts. Combining these requirements will be the basis of the conceptual models to develop the RDMA.

Some requirements are not measurable, making it impossible to assess if they are satisfied. According to the methodology by Wieringa [73], these requirements need to be *operationalized*. This means that they require a measurement procedure that indicates the presence of the property. The metrics for which a requirement is measured are called *indicators* [73]. The specification of the amount it needs to achieve is called a *norm*. An

example would be a requirement that requires a system to process and output data quickly. In this requirement, 'quickly' needs to be operationalized. The indicator could be the number of seconds, and the norm could be 30 (seconds as defined by the indicator). In the requirements, we will establish some indicators for requirements that cannot be strictly measured. We will not suggest exact norms for every requirement as the amount of improvement can depend on a multitude of variables. For instance, one implementation might have a more complex infrastructure, with higher cost, but with a reduced risk of threats. The norms for such a case will differ from a lower-cost, higher-risk solution. However, in either situation, the requirements must at a minimum be improved by implementing the RDMA, compared to the situation without an RDMA. So, for example, the conceptual model (and thus, every implementation) must reduce the disaster recovery time by some amount as stated by requirement Req1.1, and when an implementation is designed for a specific context, the solution needs to be normed by for example to make the disaster recovery at least 20% quicker.

6.1.1 Disaster Recovery

Non-functional Requirements

Req1.1 **The solution must reduce the recovery time of a disaster recovery procedure.** This is the main aim of the solution. An indicator can be a literal amount of time or a percentage that the recovery time is reduced compared to the current situation. Norming this requirement is very situational to the complexity and design of an implementation. Depending on the type of implementation, the recovery time norm can be adjusted accordingly.

Req1.2 **The solution must be safe and unaccessible by a threat actor.** If a threat actor can access the DM manager, or a DM itself, the solution can be infected as well. This can cause the system to be compromised, extend the recovery time, or be used at a later point in time to conduct another attack. This is why an implementation of the solution architecture should be pen-tested by a certified independent security tester to comply with this requirement.

Functional Requirements

Req1.3 **The solution must have a management platform that allows users to interact with the system.** To manage the DM, there is a need for a management interface to configure systems settings so it can integrate with other systems and configure which VMs should be included in the DM that it manages. The functionalities that need to be provided by the management platform will be guided by the requirements below.

Req1.4 **The solution must be able to register accounts that have strict access to the solution architecture.** To manage who has access to what part of the DM and DM manager, identity and access management must be implemented. The management platform should include a module that gives admin users the option to manage which accounts have access to what.

Req1.5 **The solution must be able to simulate the production environment partly, or wholly, based on a configuration.** An IT infrastructure often includes many systems that are relevant to an organization, but not critical to business operations. In the recovery process, we want to run VMs that support

the critical business functions first, before recovering the rest of the IT infrastructure. To reduce the computing power needed to run a DM instance and shorten the process of restoring multiple VMs, being able to configure a DM to only restore business-critical VMs is a requirement for the architecture. This requirement can be disregarded if there is enough computing power to restore all VMs and if the time it takes to restore them is negligible.

Req1.6 **The solution must be able to change what part of the production environment it simulates.** In the IT disaster recovery process, VMs of applications need to be moved from the DM to production to restore them fully. This will change the DM, making it a dynamic simulation entity. An addition to the previous requirement, this requirement allows users to recover VMs from the backup to the DM during the disaster recovery, which can reduce the amount of VMs that are running at once in the DM.

Req1.7 **The solution must have a connection with the backup server to restore VMs.** When we want to restore and recover VMs, we need to have a connection with the backup system to be able to use it. Having this connection configured and established beforehand will decrease the time needed to restore the backups contributing to a shortened recovery.

Req1.8 **The solution must be able to set up a separate isolated network environment.** As one of the non-functional requirements is that the DM should be inaccessible by the threat actor, this functional requirement makes sure that this is possible. The solution architecture should be able to set up a separate network environment that is isolated so backdoors cannot be exploited during the eradication process.

6.1.2 Disaster Recovery Test

Non-functional Requirements

Req2.1 **The solution provides a realistic environment to test a disaster recovery procedure.** This requirement relates to the fidelity of the DM, which should be high to get a good understanding of the inner workings of the production environment.

Functional Requirements

Req2.1 **The solution must be able to register accounts that have strict access to the solution architecture.** To be able to give certain users access to the environment, this needs to be manageable.

Req2.2 **The solution must have a connection with the backup server to restore the system.** The same argument mentioned in the disaster recovery requirement can be given.

Req2.3 **The solution must have a management service that allows users to interact with configurations and the DM environment.** The same argument mentioned in the disaster recovery requirements can be given, however, more specifically the disaster recovery tests are often aimed to test a certain part of the production environment, making the configuration capability even more important.

Req2.4 **The solution must be able to simulate the production environment partly, or wholly, based on a configuration.** The same argument can be given as was done in the disaster recovery requirements, however, more specifically the disaster recovery tests are often aimed to test a certain part of the production environment, making the configuration capability even more important.

Req2.5 **The solution must configure the networking of each backup server such that it can still find other required servers in the DM.** One of the challenges with a disaster recovery test is that the VMs have specific network configurations to communicate with other VMs. These are essential to the functioning of an application, however, configured to be running in the production environment. Since these network settings are already being used in the production environment as it is still running, it can be difficult to use the same network addresses for the testing environment. If the application is not running correctly due to network constraints, it cannot be fully tested. The solution architecture should deal with this problem and be able to set up a clone of the network environment with the same IP range.

6.1.3 Business Impact Analysis

Non-functional Requirements

Req3.1 **The solution must provide a realistic environment to evaluate system interdependencies.** The same argument mentioned in the disaster recovery test requirement can be given.

Functional Requirements

Req3.2 **The solution must be able to register accounts that have strict access to the solution architecture.** The same argument can be given as was done in the disaster recovery test requirements.

Req3.3 **The solution must have a connection with the backup server to restore the system.** The same argument can be given as was done in the disaster recovery test requirements.

Req3.4 **The solution must configure the networking of each backup server such that it can still find other required servers in the DM.** The same argument can be given as was done in the disaster recovery test requirements.

Req3.5 **The solution must have a management service that allows users to interact with configurations and the DM environment.** The same argument can be given as was done in the disaster recovery test requirements.

Req3.6 **The solution must be able to simulate the production environment partly, or wholly, based on a configuration.** The same argument can be given as was done in the disaster recovery test requirements.

Req3.7 **The solution must configure the networking of each backup server such that it can still find other required servers in the DM.** Currently, it is challenging to run certain VMs of an application separately to evaluate system interdependencies due to specific network configurations of the VMs. These are essential to the functioning of an application, however, configured to be running

in the production environment. Since these network settings are already being used in the production environment as it is still running, it can be difficult to use the same network addresses for the testing environment. If the application is not running correctly due to network constraints, it cannot be fully tested. The solution architecture should deal with this problem and be able to set up a clone of the network environment with the same IP range. This allows organizations to test certain VMs for interdependencies.

6.2 Design Choices

Considering the problem contexts described in the previous two chapters, using the concept of a DM to improve the problem contexts can sound straightforward. The challenge is however that it needs to reduce the disaster recovery time for the disaster recovery procedure while keeping it safe and secure. This makes every step and operation required to create a DM a costly one. Due to this, some functional and architectural decisions have been made to optimize this process.

Many requirements state the need for a system in which the DMs and their configurations can be managed. Due to this, we suggest using a management system called the 'Digital Model Manager'. This management system will take care of the automation steps, which are essential to reduce the disaster recovery time during the disaster recovery process. As mentioned in the introduction section of this chapter 6, the combination of the DM Manager and the DMs itself is called the Resilience Digital Model Architecture (RDMA). In this section, we will explain what its main functions are.

A last consideration is the choice of hardware infrastructure on which the DM should run. Currently, mainly three hardware infrastructure options are available. The first is on-premises production hardware infrastructure that already exists and is available. The second is also on-premises, but separate from the production hardware infrastructure that already exists. Last is in the cloud, where computing power can be bought as a service. The pros and cons of each of these options will be explained in section 6.4.

6.3 Functional model

According to the requirements, the solution architecture needs to provide certain functionalities. To go a bit more in-depth, this section will cover the functionalities of the DM manager through an application function model. Next to this, we will also cover how the functionality model will integrate with the BCM processes which we are trying to improve. In these models, we have highlighted changes in the components. The components that have been added have been marked with a higher saturation compared to the default colour and the components wherein components have been removed were marked orange. Following Wieringa's methodology [73], this step involves proposing the treatment design and integrating it with the problem context to develop validation models. We hypothesize on the effects and evaluate it in validation chapter 7.

The application function model is based on the DT lifecycle management framework by Grasserli et al. [29]. This framework helps us to understand which phases a DT should have. We will explain the lifecycle management framework first before we explain how the application function model has been designed. This will help us to design the functionalities that the DM Manager needs to facilitate when managing the DMs.

6.3.1 Lifecycle management

The DT lifecycle management framework by Grasselli et al. [29] will be adopted in the RDMA. They propose a solution architecture with a management platform that orchestrates the creation of pre-configured DTs [29]. The schematic overview can be found in figure 6.1.

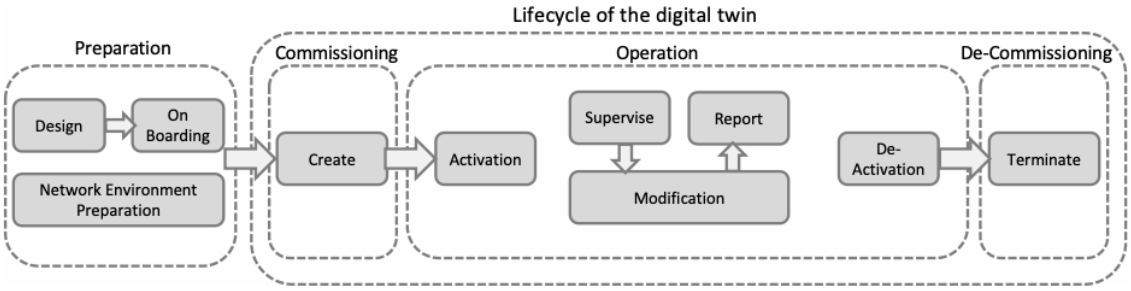

FIGURE 6.1: Schematic overview of DT lifecycle by Grasselli et al. [29]

The lifecycle management of a DT, inspired by 3GPP and ETSI standards for network slicing [22] [23], consists of four main phases: preparation, commissioning, operation, and de-commissioning. In the preparation phase, the DT's architecture is designed using templates called descriptors. After creation, these are onboarded onto the commissioning service with the necessary environment settings such as network configurations. The commissioning phase involves the automated instantiation of the DT based on these descriptors. During the operation phase, the DT is active, with its components monitored and reconfigured as necessary. Finally, the decommissioning phase entails the removal of the DT's components and releasing computational resources when the DT is no longer needed. [29] The life-cycle management has been integrated into the functional model that will be discussed in the next section.

6.3.2 Application Functions

In this section, we will discuss the functionalities of the DM Manager which will act as the main contribution to this solution design. Let's first explain the RDMA in a bit more detail before explaining the application function.

The RDMA functions primarily from a management platform designed to automate the creation and management of DMs. A DM in this context is a simulation of the necessary components required of a part of the production IT infrastructure, such as VMs and network configurations. So in essence, the RDMA is copying VMs from the production backup to a separate isolated environment where we can alter it to simulate certain scenarios without affecting the production environment. What part of the production IT infrastructure will be created depends on a given preset, which can be managed by the digital model manager. When needed, the DM is initiated using a preset configuration, which contains detailed information about what the DM should encompass, which allows a DM to be built in repetition if required. This automation allows organizations to simulate a test environment, thus facilitating rapid recovery and testing environments that support disaster recovery procedures. Later in this section, we will explain how the digital model manager can construct these DMs.

The DM manager is designed to streamline and manage DMs through three main services: a configuration service, a commissioning service, and a de-commissioning service.

Each of these services plays a crucial role in ensuring the efficient operation and lifecycle management of DMs within the solution architecture. The functions of the application are explained in an ArchiMate application function viewpoint depicted in figure 6.2.

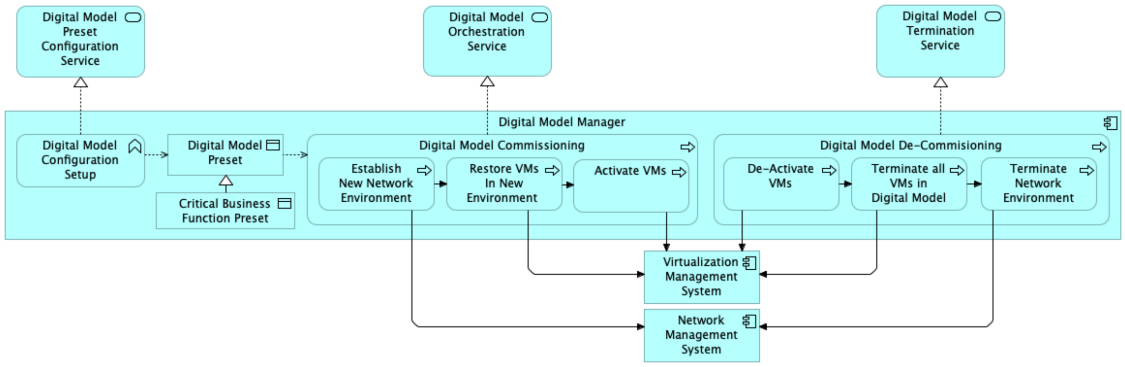

FIGURE 6.2: DM Manager Application Function Viewpoint

The 'Digital Model Preset Configuration Service' is responsible for setting up and maintaining DM presets. Each preset represents a specific part of the IT infrastructure, which are composed of VMs. One such preset that can be used for the disaster recovery process is the 'Critical Business Function Preset'. This service allows users to configure which VMs should be included in a preset, including the network configuration. Also, it should be able to configure the amount of computing power a DM can use. The preset can then be used to commission a DM according to the configuration. The configuration service should allow for easy adjustments and updates to the models' settings, enabling quick responses to changing requirements or improvements in technology. Additionally, it may include test mechanisms to ensure that any changes in the VMs do not introduce errors or inconsistencies in the DM.

The 'Digital Model Orchestration Service' facilitates the creation and integration of a DM into the existing infrastructure. As seen in figure 6.2, the service is effectively built out of three consecutive application processes using a DM preset that the user has selected. It connects directly with both the virtualization management system and the network management system, which allows the model to automate the process. By interfacing with these systems, the commissioning service ensures DMs are fully functional, and seamlessly integrated with other components of the infrastructure.

The de-commissioning service manages the orderly removal and deactivation of DMs from the IT infrastructure after their usage within the BCM processes. As with the commissioning service, it interfaces with the virtualization management system and the network management system to ensure that the decommissioning process can be executed. By doing so, the de-commissioning service helps prevent potential conflicts or resource wastage as they are not required for the BCM processes anymore.

6.3.3 Disaster Recovery Integration Model

In combining the RDMA into the disaster recovery problem context, some of the business processes have been changed or replaced by the DM manager. The proposed integration of the RDMA and disaster recovery mainly automates some steps of the process that will construct the isolated environment and run the VMs that need to be recovered. This integration is depicted in figure 6.3. The original business process model should be understood as explained in section 5.2, to understand the proposed improvements in the business pro-

cess. In this section, we will only cover the integration changes that have been made to the business process.

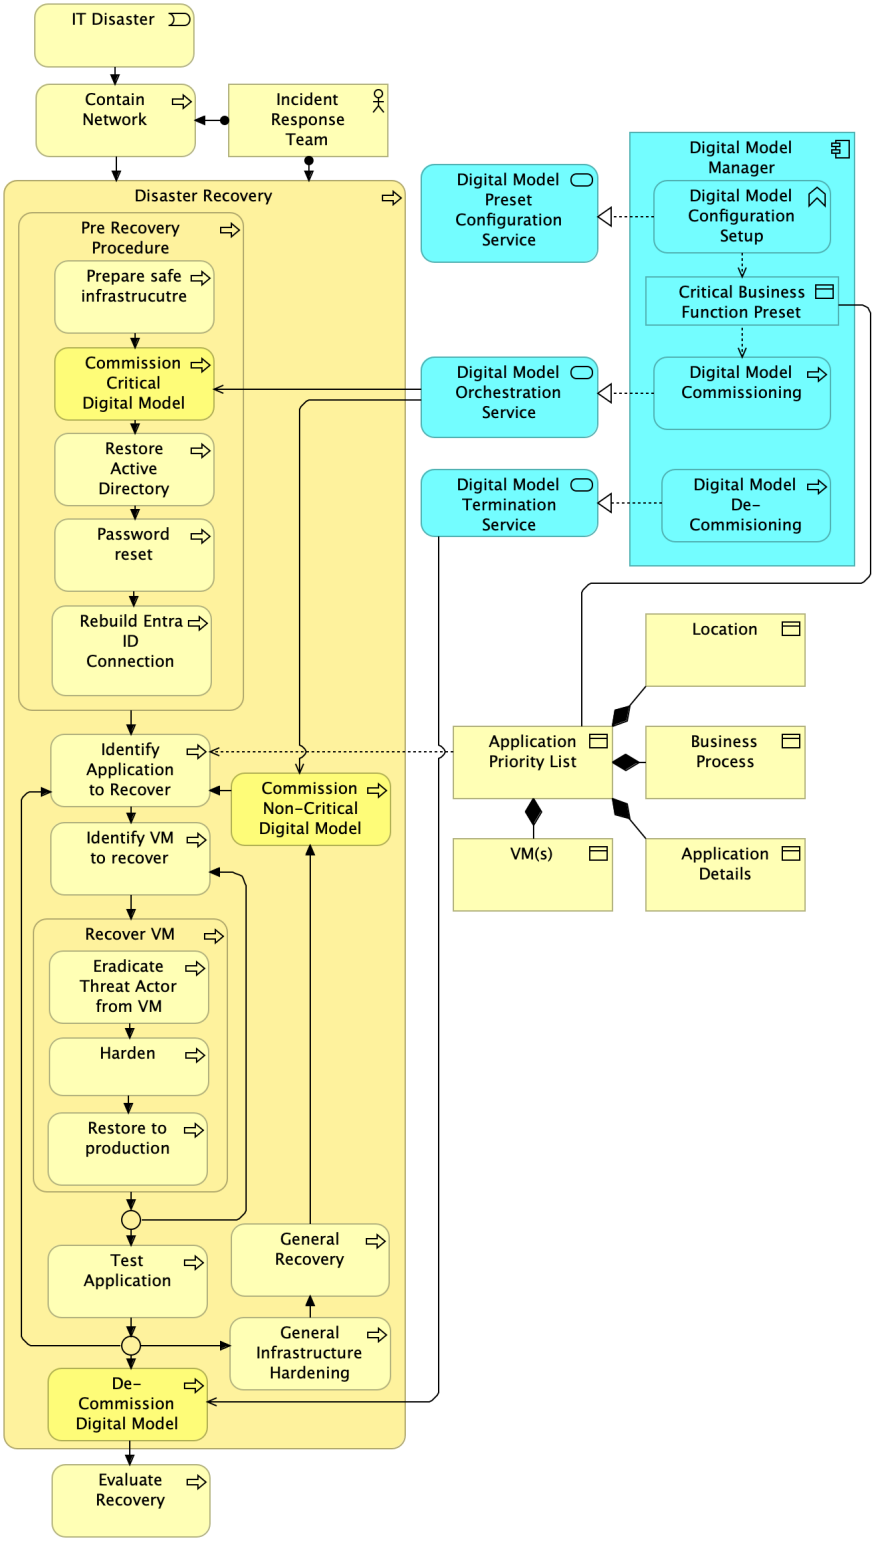

FIGURE 6.3: Proposed Solution Disaster Recovery Integration Viewpoint

First, we need to discuss what happens before and in the event of an IT disaster

to the RDMA. To ensure that we benefit from the automated processes to set up a DM, organizations need to invest time pre-emptively to create presets that can be commissioned in the disaster recovery process. One essential preset is the critical business function preset. This should include all the VMs that enable critical business functions. One may wonder why not all VMs are restored in a single DM. The counterargument is that this would cost a lot of computing resources and data transfer, which can cause a delay in the overall recovery process. To create this DM preset, organizations also need to figure out which applications are needed for critical business functions and thus create an application priority list. We propose to do this during the BIA, which also offers other benefits as explained in 6.3.5. Hence, the step to create an application priority list as was depicted in the original business process model in section 5.2 has already been conducted, so it will be excluded from the model.

In the pre-recovery procedure, we can find the most has changed. This is where the RDMA should make the biggest difference by automating some of the processes. As can be noted, the disaster recovery secure environment setup has been removed by the commissioning of a DM that includes all applications that support critical business functions. Still, some of the steps need to be executed manually. The active directory has to be restored, passwords need to be reset and a connection with Entra ID needs to be rebuilt.

The other processes have barely been changed. In the 'Recover VM' process, the 'Recover from backup' process has been removed as this is also something that the DM manager will provide. After the critical business function applications are running again and 'General Infrastructure Hardening' has been done next to 'General Recovery', a new DM needs to be commissioned for the non-critical business functions. This can be one or more DMs, depending on computing power configurations. Lastly, after every step has been executed in the disaster recovery, the DMs need to be de-commissioned.

In general, the integration of the RDMA into the disaster recovery process shortens recovery time by automating many of the critical steps involved. The RDMA allows for rapid instantiation of recovery environments, which facilitates the restoration and eradication of the threat actor of VMs. This automation reduces manual intervention and ensures a more efficient and reliable recovery process. This reliability is also maintained as the RDMA can also be used for disaster recovery tests, as explained in the next sub-section.

6.3.4 Disaster Recovery Test Integration Model

The integration of the RDMA into the disaster recovery test primarily changes the process of initiating the testing environment and removing the testing environment as shown in figure 6.4. This process will now be dictated by the DM manager. However, an extra step is needed to facilitate this in the preparation phase. Let's take a closer look.

In the 'Disaster Recovery Test Preparation' process, the DM now also needs to be prepared before the test will be started. A DM preset has to be tailored to the script and configured accordingly. As tests are often done periodically, DM presets can be reused, making the process redundant in these situations. The following change is in the 'Disaster Recovery Test' process itself.

In the disaster recovery test, the 'Initiate Testing environment' will be completely executed by the 'Digital Model Manager'. This is still part of the test itself, as it is important to know if the orchestration process still works properly. As the orchestration service has created the DM, the script can be played out in the same fashion as was done before the RDMA integration. Furthermore, the termination process will also be executed by the DM manager.

Generally, the process components haven't changed much, but the time and complexity

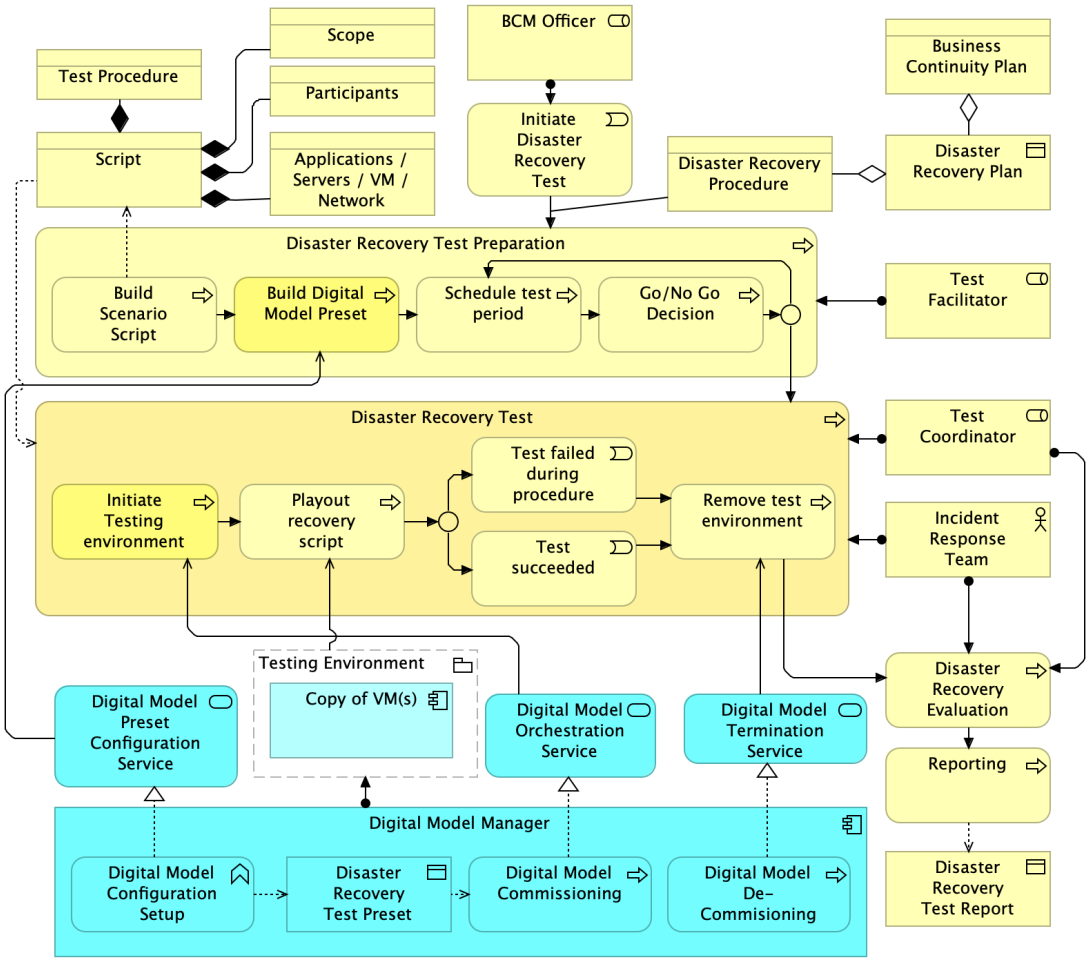

FIGURE 6.4: Proposed Solution Disaster Recovery Test Integration Viewpoint

it takes to initiate a testing environment and remove it again will be reduced largely by the implementation of the RDMA due to its automation properties. Incident responders are now not required to conduct the steps to create a separate environment with the specific settings that would have been needed before. This has all been taken care of by the 'Digital Model Manager', reducing the effort to conduct a realistic disaster recovery test significantly. Also, increasing the size of a disaster recovery test (how many applications are included for example), does not require as much manual work compared to the initial situation. Which VMs should be included in the test can be added within the 'Digital Model Preset Configuration Service', where only the names of the VMs are required.

6.3.5 Business Impact Analysis Integration Model

As mentioned in section 5.4 of the previous chapter, understanding application inter-dependencies is crucial for an accurate risk assessment, BIA, and improved disaster recovery procedures. To address this challenge, the RDMA has been designed to facilitate a newly defined 'Application Inter-dependency Analysis' as shown in figure 6.5. This process, which is aggregated from the 'Business Dependency Analysis,' can serve both the BIA and disaster recovery procedures independently. If the BIA is not executed, it can still function as a way to understand which applications are needed for critical business

processes to create a critical business function DM pre-set, which can then be used during a real disaster recovery scenario like mentioned in sub-section 6.3.3. Moreover, we include it as an essential part of the BIA. During this analysis, software application dependencies are identified, including the VMs essential for each application. A separate environment is required for this process, as explained in section 5.4.

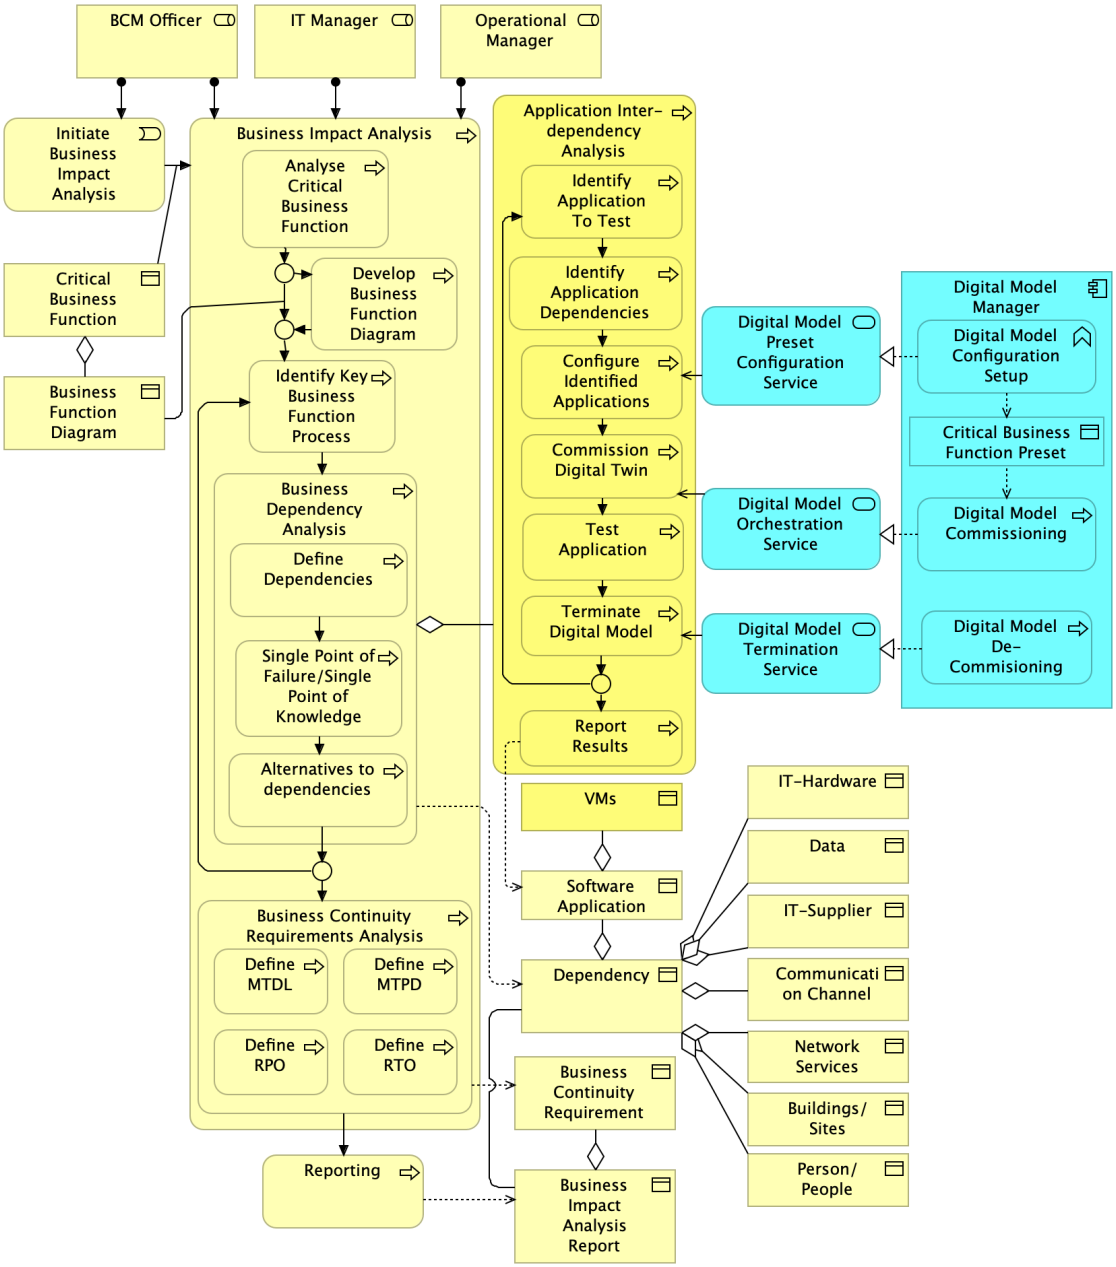

FIGURE 6.5: Proposed Solution Business Impact Analysis Integration Viewpoint

The application inter-dependency analysis involves several steps. First, the application needs to be identified to be tested. This is done in the same manner as initially, by a best-effort prediction of inter-dependencies according to the IT manager and operational manager. Next, the dependencies are determined on which it relies. These dependencies should be included in a DM preset, using the DM preset configuration service. This preset can then be commissioned through the DM orchestration service. Once the DM has been

started, the most critical step is testing the application to ensure it functions correctly. If so, this confirms the dependencies, although it may not rule out the presence of redundant VMs. To remove these redundant VMs, the test can be repeated with the removal of VMs that are presumed redundant. If the application functions properly, we know that these VMs are not required to recover critical business functions in disaster recovery. If the application fails, the tester must evaluate the issue and identify the necessary applications or VMs, which is depicted with a flow back to 'Identify Application To Test' in the model. Through iterative testing, a better understanding of the IT infrastructure will emerge, establishing a critical business function preset that supports the disaster recovery procedure. The final step is to 'Report Results'.

Furthermore, the BIA process does not require other changes. However, with the insights from the application inter-dependency analysis, practitioners can better understand the risks of their IT infrastructure, and reduce disaster recovery time with the critical business function preset in the DM manager. This can then be used directly during a disaster recovery process, which as explained in sub-section 6.3.3 will reduce disaster recovery time.

6.4 Architecture model

The IT architecture of the RDMA is essential to enable the functionalities defined in the requirements, for which three scenarios can be utilized. Implementing the RDMA to run large DMs can be a resource-intensive task. In the context of implementing the RDMA to improve the three problem contexts, the DMs should be able to run next to the production environment. As most organizations have computing capacity tailored to their IT infrastructure needs, it might become a problem to execute large tests or application inter-dependency analysis on production computing capacity. This brings us to the question of where the computing power should be established, and how to manage the security concerns according to requirement Req1.2.

To deal with this challenge, there are three primary options available: On existing hardware on-premises, on dedicated hardware on-premises, or cloud hardware. For each of the scenarios, an infrastructure model has been developed, from where we will state the pros and cons of the architecture.

6.4.1 On Existing Hardware Infrastructure

In figure 6.6, the implementation of the RDMA on existing hardware infrastructure has been depicted. On existing hardware, the RDMA can cause performance problems in the production environment. Next to this, there are security considerations to be made in using the same hardware as the production environment to host the DM manager.

Taking a closer look at the model in figure 6.6, we can see that there are some additions in the 'Virtualization' node. A 'DM Manager VM' has been added where the DM manager will run. It relates with the backup data block to show that it uses those backups to orchestrate the VMs, as well as with the firewall which it uses to create an 'Isolated DM Network Environment' for each of the DMs it creates. Lastly, a 'DM Environment' group has been added, which is used to explain that the DM runs in its own isolated DM network environment and can be commissioned more than once next to each other. As mentioned previously, the DMs will need their databases. In this IT architecture scenario, there needs to be a possibility to run these on top of existing database hardware. Considering this architecture option, we can identify some benefits and drawbacks.

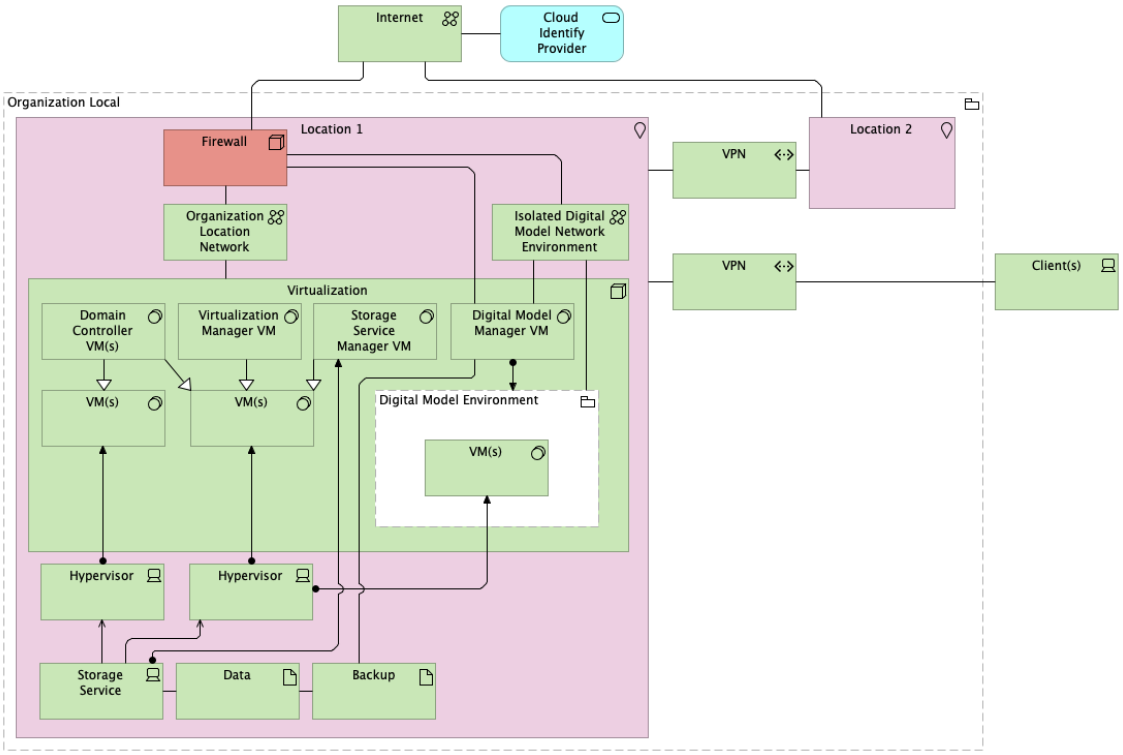

FIGURE 6.6: Architecture viewpoint of the proposed solution on existing hardware on-premises

Benefits

1. **Cost Efficiency.** Utilizing existing hardware eliminates the need for additional capital expenditure on new equipment, making it a cost-effective solution.

2. **Familiarity.** IT staff are already familiar with the existing infrastructure and do not have to manage additional resources. Next to the excluded maintenance costs, it reduces the possibility of potential errors during implementation and time to familiarize with the system.

Drawbacks

1. **Resource Constraints.** Existing hardware may not have the required capacity to handle the computational demands of large DMs next to the production environment, potentially leading to performance bottlenecks.

2. **Security Risks.** Integrating new applications into existing infrastructure can introduce vulnerabilities. Ensuring that security measures are up-to-date and robust enough to handle new threats is critical. It may also open opportunities for threat actors to confiscate the solution architecture and disable its functionality, which should be a key consideration when implementing this solution.

3. **Recovery Speed.** The goal of the RDMA is to reduce disaster recovery time. As we consider a ransomware scenario, we should take into account that the DM manager can be compromised during the attack if the solution runs in the production hardware and/or network. In such a case, the disaster recovery process will include a step to

recover the DM manager first, before it can commission a DM for disaster recovery purposes. This reduces the contribution of the solution to the disaster recovery problem context, defeating the purpose of the solution by some margin.

While integrating the RDMA into existing hardware infrastructure can be a cost-effective and efficient approach, it requires careful consideration of resource management, cybersecurity, and business continuity planning to ensure availability, integrity, and confidentiality.

6.4.2 On Dedicated Hardware Infrastructure

Integrating the DM into dedicated hardware infrastructure involves investing in and deploying specific hardware resources solely to run the RDMA. This approach can offer advantages in terms of recovery time and security but also comes with its own set of challenges.

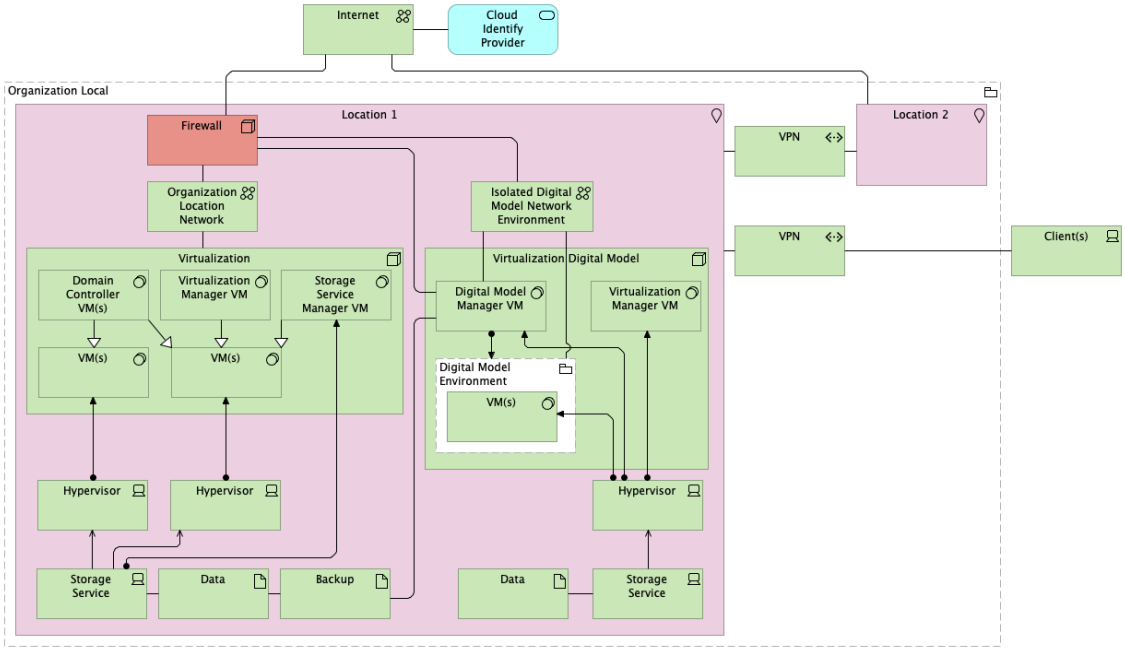

FIGURE 6.7: Architecture viewpoint of the proposed solution on dedicated hardware on-premises

Looking at the model in figure 6.7, in this architecture the DM environment separates itself from the production environment using its own 'Hypervisor', 'Virtualization Manager' and 'DM Manager VM' in the 'Virtualization DM'. In this virtualization, multiple DMs can be commissioned. The DM manager still connects with the firewall and backup system to create separate environments that facilitate the DMs. This option also brings some major benefits and drawbacks.

Benefits

1. **Enhanced Performance.** Dedicated hardware can specifically be designed to handle the computational demands of large DMs, ensuring optimal performance without impacting other critical business applications.

2. **Improved Security.** With dedicated hardware, the DM's environment can be isolated from the rest of the organization's IT infrastructure, reducing the attack surface and minimizing security risks. Security protocols can be tailored specifically to the needs of the DM.

3. **Scalability.** Dedicated hardware can be easily scaled up or down based on the requirements of the DM, providing flexibility in managing computational resources.

Drawbacks

1. **Higher Costs.** Investing in dedicated hardware requires significant investment, which can be a financial burden that will make the solution less attractive. Additionally, ongoing maintenance and upgrades also add to the annual costs.

2. **Complex Implementation.** Setting up and configuring dedicated hardware can be a complex and time-consuming process. It may require specialized knowledge and skills, increasing the dependency on IT staff and potentially delaying deployment.

3. **Resource Redundancy.** There is a risk of under-utilization of resources if the dedicated hardware is not fully leveraged. This can lead to resource redundancy, where the investment does not yield proportional benefits.

In conclusion, while integrating the DM into dedicated hardware infrastructure can offer superior performance and enhanced security, it requires a substantial financial commitment and careful planning to ensure that the resources are effectively utilized and the implementation is successfully executed.

6.4.3 On Cloud Infrastructure

Integrating the DM into cloud hardware infrastructure involves leveraging cloud computing resources provided by third-party vendors such as AWS, Azure, or Google Cloud to run the RDMA. This approach offers a flexible and scalable solution but comes with its own set of cybersecurity and business continuity considerations.

The model depicted in figure 6.8, shows that the DM is running off-premise in the cloud. The hardware infrastructure has been excluded as this heavily depends on the provider and has even more complex configurations. However, the DM manager VM will be running here. Some important benefits and drawbacks must be mentioned.

Benefits

1. **Scalability and Flexibility.** Cloud infrastructure can easily scale up or down based on the computational demands of the DMs. This allows organizations to manage resources efficiently and adapt to changing requirements without the need for high resource investments.

2. **Cost-Effectiveness.** Cloud computing follows a pay-as-you-go model, where organizations only pay for the resources they use. This can reduce costs associated with purchasing and maintaining dedicated hardware.

3. **Enhanced Security.** Leading cloud providers invest heavily in security measures and compliance, offering robust protection against cyber threats. They provide advanced security features such as encryption, identity and access management, and regular security updates.

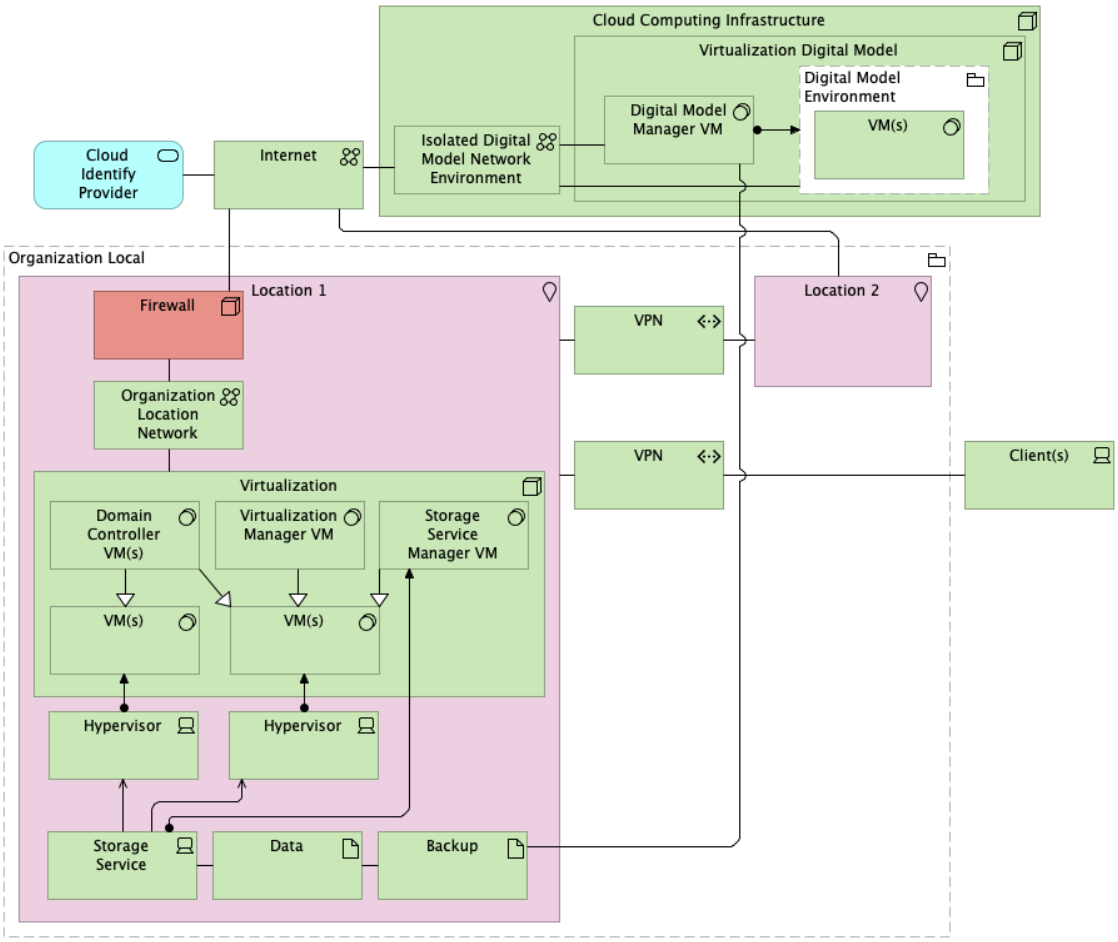

FIGURE 6.8: Architecture viewpoint of the proposed solution on cloud hardware

Drawbacks

1. **Dependency on Service Providers.** Relying on third-party cloud providers means that organizations are dependent on the provider's availability, performance, and security practices. Any disruption in the provider's service can impact the organization's operations.

2. **Data Privacy and Compliance.** Storing and processing data in the cloud raises concerns about data privacy and compliance with regulations such as GDPR or HIPAA. Organizations must ensure that their cloud provider meets all necessary legal and regulatory requirements.

3. **Recovery Speed.** The primary drawback of this option is that data transmission rates are dependent on data speeds from the organization to the cloud location. Especially because the solution aims to shorten the disaster recovery time, this option undermines that goal. Once again the impact of this drawback depends on the data speed and the amount of data that needs to be transferred.

6.4.4 Challenges

Multiple challenges have been recognized during the design of the RDMA on an architectural level.

Creating an implementation will be very costly due to the complex nature of the solution. Especially when it would be specialized to one specific organization, the cost of development might exceed the supposed benefits. It would be more economical if the implementation is provided as a service to multiple organizations. However, in doing so a new challenge is that organizations often have various IT infrastructures aligned with their needs.

Integrating with many different vendors will cause costs to rise in the development of the implementation. This causes organizations to have different vendors in terms of hardware and software for their backups, networking and virtualization. Each vendor typically has its proprietary technologies and standards, which may not always be compatible with each other. This lack of compatibility can lead to difficulties in ensuring seamless integration between the disparate systems, often requiring custom solutions or middleware to bridge the gaps. The need for specialized configurations and adjustments can increase the implementation time and require significant expertise, making the process labour-intensive and prone to errors. Moreover, the reliance on multiple vendors introduces complexities in terms of maintenance and support. The increased administrative burden of managing multiple systems also heightens the risk of security vulnerabilities, as maintaining a consistent and unified security posture across all systems becomes more challenging.

Ensuring the DM manager and the DMs themselves are isolated and safe from a threat actor is essential. In a disaster recovery scenario, the functionality of the DM manager can be perverted. This is why, after the development of the technology, the solution should tested thoroughly by a red-team to evaluate its security aspects.

6.5 Theoretical model

Looking at the RDMA from a theoretical viewpoint, we will explain how the RDMA should improve the problem context and integrate with the BCM processes defined by Russo et al. [57]. This will be done according to two EA models. Firstly, we will explain how the RDMA integrates into previously defined motivation & strategy viewpoint 6.9. After which we will show how the functions integrate with the BCM processes.

When implementing the RDMA, in theory, it should contribute in three ways towards improved organizational resilience. Two of these are aimed at achieving the outcome of having a "Decreased Recovery Time" and one positively influences the assessment of "Unacceptable Risks".

The RDMA can be used to automate processes during the disaster recovery of an IT disaster. This is a technical capability, which can be described as the "Disaster Recovery Automation Capability". It partly automates the disaster recovery process as explained in sub-section 6.3.3 using a preset that can be created with a BIA as explained in sub-section 6.3.5. This capability directly contributes to the outcome of a shortened recovery time. Secondly, implementing the RDMA facilitates the capability of conducting realistic disaster recovery tests, as explained in sub-section 6.3.4. This positively influences the principle of testing the recovery plan periodically, and consequently positively influences the decreased recovery time.

The third way positively influences organizational resilience differently. The IT interdependency analysis explained in sub-section 6.3.5, can give users improved insight into

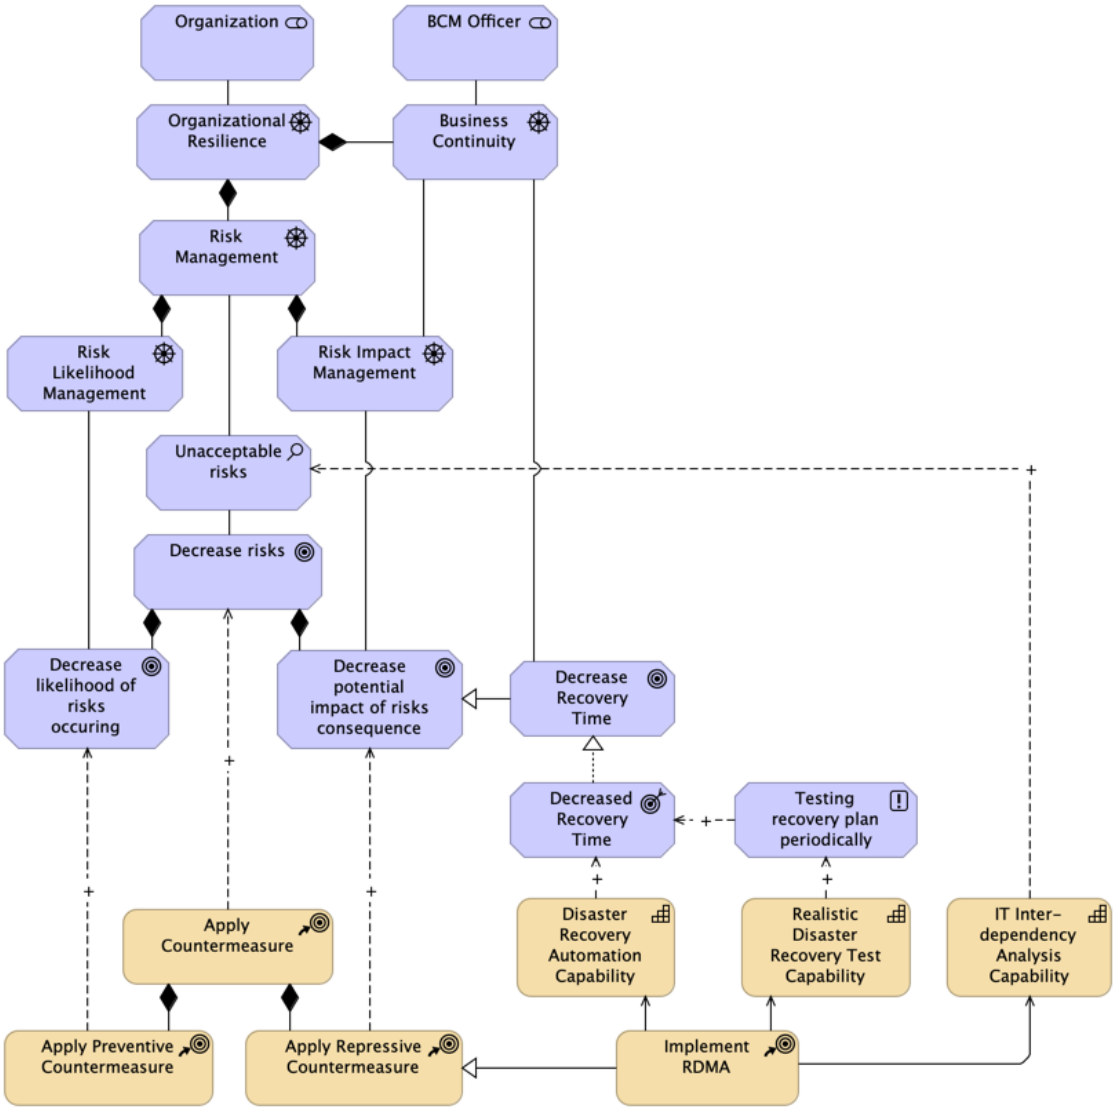

FIGURE 6.9: Proposed Solution Motivation & Strategy Viewpoint

their IT infrastructure and the possible risks attached. Having more accurate insights into risks can support practitioners to identify unacceptable risks and consequently apply appropriate countermeasures.

6.6 Conclusion

The RDMA that was proposed in this chapter aims to improve three BCM processes to improve organizational resilience during a cyber incident or cyber crisis. We explained what the requirements of such a system are, what functions it must have, how it should be integrated with the BCM processes, and what the IT infrastructure must look like to enable the DM architecture. Next to the benefit of having a shortened disaster recovery time, the DM environment can also provide organizations with a better insight into their IT infrastructure which can support them in assessing unacceptable risks.

By integrating the DT lifecycle management framework proposed by Grasselli et al. [29], the architecture provides a structured approach to the preparation, commissioning,

operation, and de-commissioning of the DMs. The DM manager provides users with a management system to provide configuration, commissioning, and de-commissioning services. As explained in the functional model depicted in figure 6.2, it should integrate with the network management system and the virtualization management system to automate the creation of a DM based on presets.

The RDMA can be used to improve disaster recovery, disaster recovery tests, and business impact analysis each in its respective way. By integrating the RDMA, the disaster recovery and disaster recovery test will supposedly shorten, thus reducing the time needed to execute the whole process. For the business impact analysis, an additional business function has been added, extending the analysis with an IT inter-dependency analysis. Results from this extension can be used to prepare for disaster recovery procedures and assess risks more accurately.

The document evaluates three primary IT infrastructure options for implementing the RDMA: existing on-premises hardware, dedicated on-premises hardware, and cloud hardware. Each option presents benefits and drawbacks. Existing hardware is cost-effective but may suffer from performance and security issues. Dedicated hardware ensures sufficient performance and improved security but comes with higher costs and complexity. Cloud infrastructure offers scalability and flexibility, yet raises concerns about dependency on service providers and data transfer speed.

Next to this, as explained in some of the sections in this chapter, some challenges and considerations need to be taken into account when implementing this solution. The implementation of the RDMA is inherently complex and costly, and integrating it into varied IT infrastructures across multiple organizations adds further challenges. A service-based approach could be more economical, but it necessitates addressing vendor compatibility issues and maintaining consistent security measures. At this hardware and software layer, vendors do not always use uniform protocols, standards or networking, which can increase the development costs of an implementation.

In the next chapter, the results of the validation cycle of the proposed solution will be described. Expert opinions will be used to evaluate if the validation models produce the proposed effects and assess if requirements are complied with.

Chapter 7

Validation

This chapter will present the results of the qualitative validation of the proposed solution in the previous chapter with expert opinions. The validation process evaluates our success in meeting the study objectives, verifies compliance with solution requirements, and assesses the practical utility and usability. As mentioned in the methodology chapter 3, it was out of scope for this study to develop a pilot implementation making quantitative validation difficult. However, considering the goals of this study, validation by expert opinions provides us with an appropriate level of understanding if the required effects are achieved with the proposed solution.

7.1 Expert Opinion Results

Expert one pointed out that DM might be required to run in the evaluation and reporting process of the BCM processes. Taking the BIA as an example, the proposed integration viewpoint as depicted in figure 6.5 has put the 'Remove test environment' business process as the last step of the disaster recovery test before the disaster recovery evaluation and reporting processes were started. The expert mentioned that the evaluation step might use the DM to evaluate what happened in certain steps and why they could or could not be executed. This can be of high significance when transferring this insight into improvement measures that can be taken in the future. The same note can be considered for the disaster recovery process.

Expert two mentioned that the benefit of familiarity with existing hardware infrastructure mentioned in sub-section 6.4.1 should be separated into two benefits. As the proposed infrastructure does not add any additional infrastructure, this doesn't come with any extra maintenance costs next to the benefit that IT staff is familiar with the existing infrastructure. This should be seen as an extra benefit for this infrastructure.

Similarly, expert two also stated that the drawback of higher cost with the dedicated hardware infrastructure in section 6.4.2 should consist of two drawbacks. First of all, high investments need to be made in extra hardware to facilitate the DMs, but there are also more maintenance costs due to it being a separate environment from production. IT staff need to understand how the infrastructure is set up and they will need more time to manage both production and the dedicated RDMA hardware.

Furthermore, expert two noted that the on-cloud infrastructure option mentioned in sub-section 6.4.3 will have another drawback. As cloud-infrastructure introduces a separate set of systems, tools, and processes that must be managed independently, it requires extra maintenance and knowledge. Especially since this needs to provide a special type of functionality, the cost of setting it up and maintaining it will be considerable.

Next to this, expert two mentioned that requirement Req1.7, which states that the solution must have a connection with the backup server to restore VMs, should state a more specific connection type. Currently, it does not specify how the connection should be set up. As such, if there is a bi-directional connection, the threat actor might be able to get from the backup server to the DM, thus causing some safety concerns. To make the connection safer, it should be done 'without being reachable from production, but in a "pull only" fashion'.

Last, a second note was made in the requirements list of the disaster recovery process. Requirement Req1.8 mentions that the solution must be able to set up a separate isolated network environment, without stating the need for the solution to have the same IP space as the production network environment. This was however stated in the related network requirements for the disaster recovery test and business impact analysis process (Req2.6 and Req3.7 respectively). Expert 2 mentioned that this property would also be nice to apply for the disaster recovery process to be able 'to have a lift and shift scenario'. Meaning that it would become less manual work to move a VM from the DM to production, thus speeding up the disaster recovery process. As the requirements are all considered for the solution architecture, these requirements should already be fulfilled by the RDMA.

7.2 Conclusion

Next to the suggestions and notes made by the experts, the validation process confirmed that the proposed solution effectively meets all established requirements and should achieve the expected effects as described in the proposed solution chapter. The feedback provides valuable insights and suggestions for improvement, which gives a clearer understanding of the solution without altering its core functionality. The experts did not identify any additional incompleteness, incorrect assumptions, or supposed effects during the validation, thereby validating that the proposed solution's effects are accurate.

Chapter 8

Discussion

After having validated the RDMA, this chapter will present a discussion on some research-related topics that are worth mentioning. These discussion topics are either of relevance to the study results or ideas that came up during informal conversations and reading materials online. Lastly, we will also mention the study limitations in this chapter.

8.1 Financial Considerations

In this research, we have proposed a solution which can improve and support BCM processes but will come at a significant investment cost as mentioned in 6.4. In a cybersecurity investment strategy, the potential benefits and costs of countermeasures should always be considered before implementing it [27]. Organizations should always aim for the highest return on investment for their cybersecurity strategy [27]. The costs for reducing the likelihood of risks and the potential consequences can exceed the expected cost of a cyber attack on an organization. This begs the question of how we can best invest in cyber/information security. One of the most adopted models is the Gorden-Loeb model [27].

The Gordon-Loeb model (GL model hereafter) is an economic model designed to determine the optimal amount an organization should invest in cybersecurity, which is depicted in figure 8.1. This model balances the cost of cybersecurity investments against the potential benefits in terms of reduced losses from cyber incidents. The fundamental principle behind the model is that organizations should allocate their cybersecurity budgets in a way that maximizes their return on investment (ROI). [27]

The Gordon-Loeb model consists of some key variables: the vulnerability level $v(0 \leq v \leq 1)$, the potential cost of a security breach L, expected loss before implementing security countermeasures vL, investment in cybersecurity z which will reduce v based on the productivity of the cybersecurity investment and finally $s(z, v)$, which is defined as the security breach probability function. This function indicates (as assumed by the researchers) that while cybersecurity investments reduce vulnerabilities, they do so with a decreasing rate of returns when increasing the investment. The model also assumes that even with increasing investments to a maximum, the probability of a breach can be reduced, but never to zero. The researchers argue that organizations should generally optimally invest less than or equal to about 37%, as at this point the marginal benefits of decreasing risks become less effective. [28]

As mentioned in the article by Gordon and Loeb [27], investments need to be compared towards their impact to reduce risk-related costs. In the context of the RDMA, organizations need to evaluate whether the investment into the architecture is a cost-effective investment of resources. Due to the highly estimated investment cost as mentioned in the

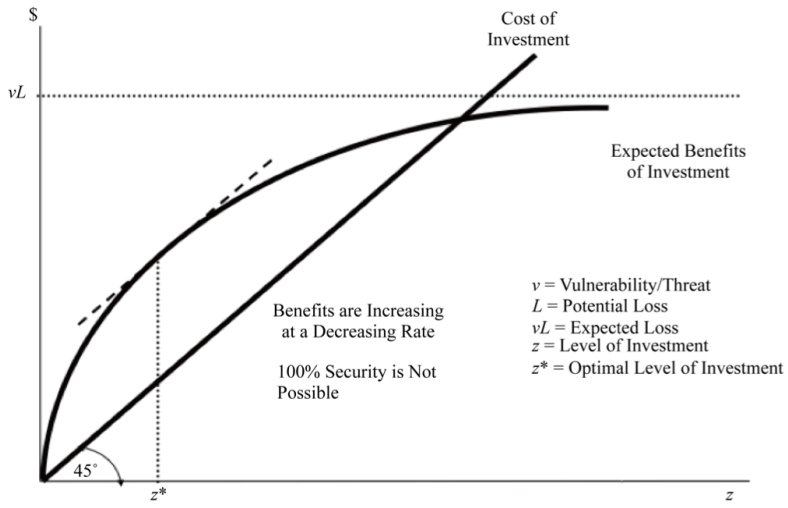

FIGURE 8.1: Gorden-Loeb Model [27], Benefits and costs of an investment in cyber/information security

proposed solution chapter 6, this solution is only relevant for organizations with very high expected losses. Furthermore, the solution should be compared to other solutions in terms of their security productivity. If it is an option that can be considered, organizations need to ensure it does not exceed the maximum investment threshold according to the GL model [27].

8.1.1 Organisational Maturity and Use of Solution

That brings us to the maturity of organizations that might be interested in implementing the solution. Especially in informal conversations, some colleagues of Northwave have mentioned that they expect companies should rather invest in other technologies that will enhance cyber security before implementing such a costly technology as the RDMA. Developing the RDMA will be a significant investment of time, money, and resources as mentioned in chapter 6.

Improving organisational cybersecurity maturity can be done in many ways. Most companies with low to medium maturity would have more motivation to invest in the three main areas of cybersecurity (according to Nothwave): business (compliance, policies, etc.), bytes (SOC, Red Teaming, etc.), and behaviour (phishing simulation, resilience training, etc.). These companies would rather allocate resources towards preventive measures, which focus on hardening their IT systems.

It was also mentioned in informal conversations that, most organizations, don't feel the motivation to be so well prepared against cyber attacks. Organizations prefer to focus on their core business operations and investing here to stay ahead of the competition. Direct investment into core business operations can improve profit, but for investments in cybersecurity, it is harder to grasp the benefit due to challenges in estimating the mitigated costs of a mitigated cyber attack. Organizations often only realize that cybersecurity investment is needed when a competitor or similar company has suffered a cyber attack. As mentioned at the start of this thesis, there is no such thing as being fully protected against cyber attacks.

Therefore, the RDMA is a technology that currently primarily offers opportunities to highly mature cybersecurity companies that understand cybersecurity risks well and

are willing to invest in new security efforts to increase their cyber resilience. For most organizations, this solution will probably exceed the optimal level of investment as depicted in the GL model in figure 8.1. In the future, hardware and software that are required to enable this solution might be offered at a lower price, allowing more organizations to adopt it as a cybersecurity enhancement.

8.2 VM2020

Another interesting discussion point is the company called VM2020 [71]. During the research, this company was found which offer a solution that could be explained to be an implementation of the RDMA.

The cornerstone of this solution is what they call the Thin Digital Twin technology, which creates safe, virtual copies of production systems for continuous vulnerability detection, remediation, and validation. These thin DTs replicate and simulate the real environment without the associated risks, allowing businesses to identify and address weaknesses in a controlled setting. Considering the classification by Kritzinger et al. [42], this solution architecture should in this context be called a DM.

Their DM manager is called CyberVR. As mentioned on its website, CyberVR distinguishes itself by providing a software automation and instrumentation platform that accelerates comprehensive and realistic cyber resilience validation. The platform leverages existing data protection solutions and virtualized workloads to create full-scale, instantly available DTs of production environments. These sandboxes, or snapshot clones, allow for extensive cybersecurity and DevOps testing without impacting live systems. By enabling interactive manipulation by multidisciplinary teams, CyberVR optimizes workflows, validates change management, and gathers critical evidence of cybersecurity strengths or vulnerabilities. This capability significantly reduces detection and remediation times, anticipates patching side effects, and provides essential forensic data for cyber insurance claims, thereby enhancing overall IT resilience and operational efficiency. [71]

VM2020 is a prime example of how the RDMA can be implemented and proves that there are already implementations out there providing the solution architecture. VM2020 integrates with specific vendors to facilitate the functionalities and can be seen as a service provider of the technology. Most importantly it shows the relevance of the technology and that VM2020 shows a business case can be made to help out organizations. During this research, VM2020 was used as an inspiration that supported the development of the RDMA.

8.3 Cyber Resilience Terminology

While conducting this study, there have been many situations where Cyber Resilience is used in a different context and is explained in various ways. As explained in section 2.2, where we go in-depth into the concept of Cyber Resilience, many institutions like NIST [55] and MITRE [10] dissect the concept into multiple goals and objectives. However, the approach to Cyber Resilience at Northwave is different by dissecting it into two areas: *withstand* and *endure & recover*. This was explained in section 2.2 based on a white paper by Northwave about Cyber Resilience [53]. Although these references explain Cyber Resilience thoroughly, the terminology used throughout all of the references can be confusing.

For example, the use of the term 'preventing' is one of the main issues, which is used by NIST [55] and MITRE [10] and within this research. Preventing implies that the

cause of an incident is being mitigated. But then the question arises, what is the cause exactly? Earlier in this study we took the example of an office building. Now let's say the office building is in an earthquake-prone area. Is the earthquake the cause of the adversity, or is the unstable foundation of the house the cause of the adversity? And in that same sense, does a preventive measure prevent the earthquake from initiating, or does it prevent the house from being affected in the case of an earthquake? We would argue that the earthquake is the root cause, and preventing it would mean that we take action to prevent the earthquake from taking place. Translating this to the cyber world, we use the terminology differently. Considering most literature, the term 'preventive measure' in the context of cyber resilience, is used for measures that reduce the risk of a threat actor creating adversity in an IT system.

We would suggest seeing Cyber Resilience and Business Continuity Management in the context that we should consider an adverse event if we talk about the topic. Withstanding means that organizations are not affected by adverse situations, which makes preventive measures aimed towards preventing being affected by adverse situations. Enduring and overcoming should be used when an adverse situation has any effect on the BC of an organization. Repressive measures can be used to reduce the impact and enhance the enduring and recovering capability of an adverse situation.

8.4 Disbanding the Context of Cybersecurity

This research was strictly conducted in the context of cybersecurity due to its relevance for Northwave. But what if we disband this context and take a more holistic look at the problem context for the whole of BCM? Could the solution architecture play a role here as well?

In section 5.1, we took the example of a fire hazard in the office of an organization to explain components within the field of organizational resilience. Let's re-use this example to explain how a DM can play the same role in a more general context. Just like in ransomware disaster recovery cases, where it is hard to create a realistic environment to simulate a disaster recovery scenario, the same problem exists for fire disasters. To create a realistic environment where an employee can test to distinguish such a fire, one would need a replica of the office that will be set on fire according to a scenario. As one can imagine, this would be practically impossible to execute.

Now what if one could create a DM of the office to simulate such a scenario using the right technologies, this might also bring the same benefits for some of the BCM processes. In terms of a disaster recovery test, using technologies like virtual reality or augmented reality could help employees test a realistic disaster recovery with the DM. In the case of the office building, using the DM in combination with technologies such as virtual reality could create a realistic environment in which employees could try and distinguish the fire, like in a real-world scenario. Although it might not provide the same opportunities as the disaster recovery process and we considered it in this thesis, it might offer more benefits in the other BC processes.

8.5 Cyber Range Similarities

In this thesis, we have mentioned the concept of a 'Cyber Range' multiple times. In the literature review section 4.1.2, multiple articles have mentioned its use. In section 4.2, the concept of a DM-based cyber range was proposed. To try and make sense of the

concepts and how they relate to each other, we propose a guideline on how researchers and practitioners should use both terms.

A Cyber Range is a virtual environment specifically designed for cybersecurity training, testing, and research [75]. It provides a platform where individuals can practice attacking or responding to cyber threats, test the resilience of systems against attacks, and develop new defensive strategies. Cyber ranges often include networks, devices, and applications that mimic a general real-world infrastructure.

Both DMs and cyber ranges serve the purpose of creating controlled, virtual environments for analysis and experimentation. However, while a DM may focus more broadly on simulating any aspect of a system or process, a cyber range is explicitly tailored for cybersecurity purposes. However, our suggested RDMA focuses on cyber security as well. This begs the question, what the exact difference is between the two concepts? In the current study context, the RDMA can be seen as a realistic cyber range. Where cyber ranges normally provide practitioners with a generic environment, our RDMA focuses on recreating the production environment partly or fully. This is why the RDMA could subsequently be called a DM-based cyber range. However, as the purpose of a cyber range is not so much about recovery but more aimed towards attacking or defending an IT environment, this can also be a bit of a confusing term. This is why the 'Resilience Digital Model Architecture' name was adopted.

To conclude, our solution could be considered a realistic organizational IT infrastructure-based cyber range, however, due to the difference in purpose, the terminology of cyber range was not adopted in this research.

8.6 Integration into ART Framework

In today's financial world, it is more important than ever for systems to be secure. The Advanced Red Teaming ART framework [16], developed by De Nederlandsche Bank (DNB), represents a recent development to improve cyber resilience in the financial sector. Designed to guide red teaming efforts to improve the security measures of financial institutions, the ART framework proposes a roadmap with components that should be included for a comprehensive red teaming assessment. Simulating realistic cyberattacks enables organizations to identify vulnerabilities and strengthen their defences against potential threats. Through a combination of advanced techniques and comprehensive assessments, ART is one of the most mature red teaming frameworks to test cyber resilience. In figure 8.2, an overview of the phases and modules is depicted.

As a red team never compromises the production environment to a significant extent, a blue team can never test a disaster recovery procedure according to the realistic attacking path of the red team. However, the reality is that a threat actor could actually compromise the system. Considering this, when a red team was successful in compromising high-level rights which can be exploited to initialize for example a ransomware attack, it might be worth extending the ART framework with a disaster recovery procedure.

This is where our solution might help. The RDMA can then serve as a realistic cyber range where the red team can execute malware according to the threat intelligence-driven ART framework in a safe and isolated environment. The disaster recovery procedure can be accurately tested during such a scenario to enhance cyber resilience. Considering the ART framework [16], this could be an extra module in the 'Purple teaming' phase called the 'Recovery Simulation'. The integration into the ART framework could be studied in future work.

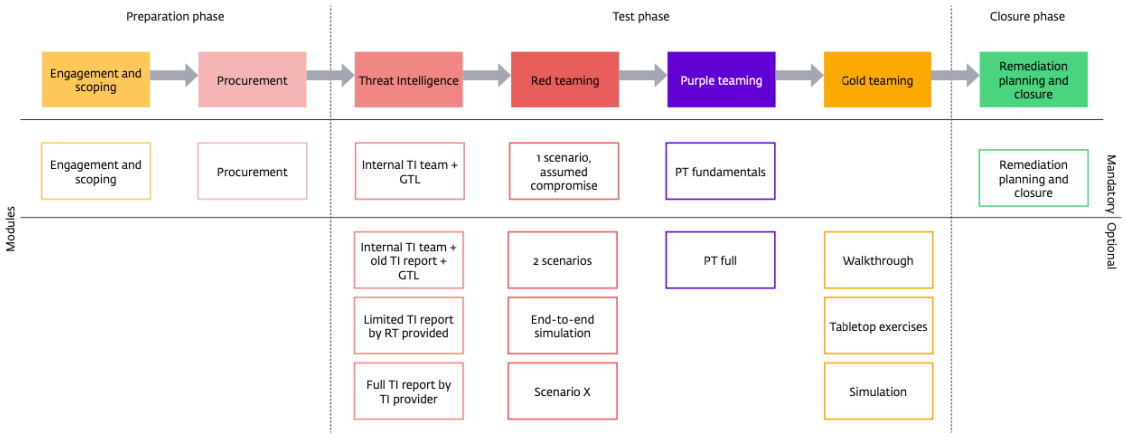

FIGURE 8.2: ART Phases and Modules [16]

8.7 Study Limitations

Transitioning from the general discussion, in this section we will discuss the study limitations.

8.7.1 Validation and Implementation

The first and foremost limitation of the proposed solution architecture is the lack of comprehensive validation. For the purposes of this thesis, the validation is sufficient, especially considering the available time of the researchers. The problem investigation data originates from three types of sources, the research incorporates expert opinions validation and an existing implementation (VM2020 [71]) of the proposed solution architecture was found online. All contributing to the validity of the results. However, there are some limitations to the validation. Only two experts were consulted for validation. Although they both are well experienced in the field, more expert opinions would lead to a more thorough validation. Also, as these experts are not very familiar with the EA modelling language Archimate [5], they might have missed some gaps in the models of the proposed solution.

Next to this, some other approaches could validate the proposed solution further. First, a conceptual demonstration of the solution architecture that improves a real-world case could provide more insights into the workings of the proposed solution. Secondly, a pilot implementation of the solution design could be developed to prove that the solution architecture can fulfil the requirements. This pilot implementation would significantly increase the credibility and validity of the solution architecture. Such a study would provide more insights into the challenges, feasibility, drawbacks and benefits of the RDMA. These validations are examples that ensure the robustness and reliability of the proposed solution in real-world scenarios.

8.7.2 Case Study Analysis

Within the case study, three data collection methods were conducted to induce triangulation to make the findings well-grounded. However, for each of the problem contexts we used a single-case approach and only interviewed a single expert per case due to time constraints around the researcher's schedule. This could have caused some amount of bias in our research, making the results less robust to other case contexts. Even still, the results are gathered from experts in the field with proven methods. However, it should be seen as

a limitation and it would be a suggestion for future studies to conduct more research into various cases interviewing multiple experts.

8.7.3 Hardware and Software Integrations

For the proposed solution to operate effectively, several hardware integrations must be established. These hardware components facilitate the functionalities that are needed for DMs to run properly, including the virtualization software, networking hardware/software (i.e. a firewall), and backup servers.

We have worked under the assumption that it is possible to integrate with each of the technologies. If these integrations cannot be achieved, the solution architecture will be required to provide these functionalities itself. For instance, if direct firewall integration is not available for a specific firewall vendor, the architecture might need to include a dedicated security module to manage network security. This makes the development of a solution implementation even more complex and expensive. However, considering this scenario, a firewall vendor might see the opportunity to extend their firewall to include the functionalities of a DM manager and integrate with backup and virtualization software.

Furthermore, we need to recognize that the selected hardware and software must possess specific functionalities that are not universally available. For example, not all firewalls can create isolated sub-networks in a network environment. Therefore, choosing the right hardware and software with the necessary capabilities is essential for the successful implementation of the proposed solution.

Chapter 9

Conclusion

This research has aimed to understand the possibility of using a digital model to support and improve business continuity management processes by means of a literature review and design science study. From the literature review, we identified that there has been no study into a single digital model architecture that can improve multiple business continuity management processes. As a result, we proposed a way to integrate and improve these processes with a digital model design. However, the proposed solution cannot just be described as a single digital model as it includes modules that create and manage digital models, but should rather be called an architecture. This is why we have called the proposed solution the Resilience Digital Model Architecture (RDMA).

In this chapter, we will conclude the research by answering the research questions, describing the contributions to theory and practice, and discussing suggestions for future work.

First, we will answer the questions to the sub-questions, before answering the main research question.

> **SQ 1.** *What are the key trends on how a Digital Shadow can enhance Cyber Resilience in the context of Business Continuity Management for Organizations?*

This sub-question was answered comprehensively in section 4.4, and we will repeat the most important part of what was said there.

The results show that many studies tried to integrate the digital twin, digital shadow and digital model into various business continuity management activities. Especially the digital model was used often, due to its simulation capability. Moreover, many articles conducted their research in the context of an OT system, which was included in the proposed digital twin, digital shadow or digital model. The articles succeeded in showing the opportunity of using these technologies, in particular for *incident response and business continuity plan testing, maintenance and analysis*, and also moderately for *risk assessment, business impact analysis, business continuity management strategy, ICT strategy, business continuity training.* Articles did however not cover the steps of *understanding the organization, alternatives to critical functions, business continuity plan design and implementation and crisis management.*

To answer the first sub-question, the literature rarely reports on digital shadow implementations specifically. Only three articles explicitly used a digital shadow as the preferred technology. Two of which showed uses in incident response strategies for real-time monitoring and knowledge-based decision-making to enhance cyber resilience for organizations

by Maia et al. [43] and Allison et al. [1]. The third article by Epiphaniou et al. [21] showed that a digital shadow could also be used for business continuity testing, maintenance and analysis through a cyber resilience assessment strategy. Noting that only three out of twenty-five articles explicitly mention that they use the digital shadow, it suggests a lack of use cases or limitations of this technology. On the contrary, neighbouring technologies (the digital twin and digital model) show many more use cases to enhance cyber resilience through business continuity management activities. Most likely, the reason is that a digital model provides a realistic simulation environment that can be utilized to evaluate scenarios which can be used for other use-cases than the digital shadow, and the digital twin provides users with more power to also change the physical world through the digital twin, while the digital shadow can not. Noteworthy is that nine out of twenty-five articles did not explicitly mention the DT type that they actually used. Thus there may be more articles that have used the digital shadow, without it clearly being stated.

> **SQ 2.** *Why would organizations want to integrate a digital model into business continuity management processes?*

Answering this sub-question is done best in two-fold. First by explaining which challenges occur during business continuity management processes and secondly by pointing out which functionalities of a digital model can help improve this process.

The business continuity management domain deals with discontinuities within organizations. Business continuity management officers will try to prepare for discontinuities by evaluating scenarios and creating procedures to deal with these situations. These will be put in a business continuity plan. However, as one might imagine, it is difficult to understand and deal with every single scenario of discontinuities in an organization. Moreover, testing the procedures built to deal with discontinuities will require an organization to either cause a discontinuity to happen or to simulate such a scenario. Taking the example of a fire hazard in an office building, simulating such a scenario is the only realistic option to test the procedure. In the context of cyber security, this is often the same case.

This is where a digital model can help. The digital model technology has two key characteristics: it is a realistic copy of the real-world system or object and it has a simulation capability, which means that we can use it to simulate certain scenarios which we do not want to create in our real-world system or object. In this research, we specifically investigated how the disaster recovery, disaster recovery test and business impact analysis processes can be improved.

> **SQ 3.** *What are the key processes and components in a disaster recovery, disaster recovery test and business impact analysis?*

In chapter 5, we identified the key processes and components in a disaster recovery, disaster recovery test and business impact analysis.

The disaster recovery process is quite complex and needs to be understood from a technical and process perspective. From a technical perspective, the main components that are important to understand are the data backup system, virtualization hardware and software, and the network infrastructure.

The disaster recovery test consists of two main phases. When the process is initiated, a certain disaster recovery procedure will be tested. With this, the first phase will be started which is the test preparation. In this phase, mainly the script will be developed which will be used in the next phase, which is the disaster recovery test itself. Here first

a separate network and VM environment will be created which can be used to execute disaster recovery procedure tasks. The script will be played out in which the incident response team can be either successful or unsuccessful. The recovery will be evaluated and reported from which improvement steps can be taken.

In the business impact analysis, there are also two main phases. Before the first process can start, to initiate the business impact analysis a critical business function needs to be chosen to be analysed. From this point onward, for each business process in this critical business function, the first phase will be started. This phase identifies the dependencies of this process in terms of people, software, data, suppliers, etcetera. After this phase, the second phase will identify the business continuity requirements to define recovery speed and data loss.

Understanding these processes and components lays the foundation for understanding the challenges and how they can be improved using a digital model.

> **SQ 4.** *What are the requirements and key components of a digital model to enhance business continuity management processes?*

From the question of understanding the key processes and components, which was answered in SQ 3., we will now answer the question of what requirements and key components a digital model needs to enhance these business continuity management processes. The requirements are laid out and explained in chapter 6, where the key components of the digital model are also described. For each of the problem contexts mentioned in the previous sub-question, requirements were set up that collectively needed to be achieved by our proposed solution. The most important requirements state that the digital model must reduce disaster recovery time Req1.1, unaccessible by the threat actor Req1.2, provide high fidelity of the production system Req2.1 Req3.1, and must provide a management platform where the digital models can be designed Req1.5 Req2.4 Req3.6 and managed Req1.3 Req2.3 Req3.5.

As stated in the requirements, the solution is not singularly about a digital model, but more particularly how to manage one or more digital models. This is why, as was mentioned at the start of this chapter, we should call our proposed solution an architecture: the Resilience Digital Model Architecture (RDMA). Key components are the Digital Model Manager, the connections with the network management system, the backup system and virtualization management system, and the hardware it runs on.

> **SQ 5.** *How can a digital model improve and integrate with disaster recovery, disaster recovery test and business impact analysis processes?*

A digital model can integrate and improve business continuity management processes by providing a secure environment where organizations can test and refine their disaster recovery and business continuity plans. This environment allows for realistic scenario testing without impacting the actual production environment, thereby identifying potential weaknesses and areas for improvement. Additionally, a digital model can help in understanding IT interdependencies, enabling more accurate risk assessments and better preparation for potential cyber incidents. By integrating a digital model into BCM processes, organizations can ensure faster recovery times, more effective incident response strategies, and a more resilient overall infrastructure.

For disaster recovery, a digital model can be used as an environment where threats can be eradicated from VMs that are being restored, while the production environment can

run the cleaned VMs. Instead of creating an isolated environment manually, the RDMA can be used to create a safe environment which also restores the VMs from the backup. This automated process can reduce the disaster recovery time if time has been put in pre-emptively.

In the context of disaster recovery testing, a digital model provides a separate and realistic environment to evaluate specific disaster recovery procedures. The realisticness helps in assessing whether an incident response team can execute disaster recovery procedures. Furthermore, by defining digital model pre-sets that can be re-used and automating the process of setting up this environment, a reduction of the time and resources required to execute a test can be realized.

For the business impact analysis, a digital model can provide a realistic environment where IT inter-dependencies can be tested by turning some applications on or off. This insight is crucial for identifying single points of failure and determining which systems are critical for maintaining business operations. By having a clear understanding of these interdependencies, organizations can make more accurate risk assessments and prioritize recovery efforts more effectively.

> **MQ.** *How can a digital model support and improve business continuity management processes to enhance cyber resilience?*

Reflecting on the main research question, this study has shown how a digital model architecture can conceptually improve three business continuity management processes by automating system processes and providing a realistic copy of the production environment. Using enterprise architecture models, the Resilience Digital Model Architecture has been explained from a functional, technical and theoretical perspective. By showing how the Resilience Digital Model Architecture can be integrated with three business continuity management processes, we show that these processes can be supported and improved to reduce disaster recovery times and a more accurate business impact analysis.

Although this can all sound very promising, some general challenges and limitations should be mentioned which are unrelated to specific sub-questions. First of all, one of the challenges we recognized in literature is that a digital twin, digital shadow or digital model solution in the context of business continuity management and cyber resilience, has high costs. We tried to deal with this challenge by making our proposed solution multi-functional. However, this study also identified that the solution architecture has significant costs. Not only does the software need to be developed, but it also requires high investment into hardware infrastructure setup and management (depending on the hardware configuration as mentioned in section 6.4 and the size of the software landscape).

These high costs might be mitigated by selling an implementation of the RDMA to multiple organizations, which then share the costs of development. However, another challenge arises. As most organizations make use of specific hardware infrastructure vendors that comply with their requirements, a shared implementation of the RDMA needs to be able to connect with each of them. Connecting with each vendor can raise the complexity, and thus, the costs. Moreover, the RDMA requires the hardware and software systems to have collaboration capabilities and system capabilities, which not all systems provide.

9.1 Contributions

In this academic work, we have contributed to theory in two ways.

First, the literature review provides researchers and practitioners with comprehensive insights into the state-of-the-art on how digital twins, digital shadows and digital models can enhance cyber resilience in the context of business continuity management. This review highlights new trends, identifies gaps, and suggests areas for future research. It serves as a valuable guide for researchers, practitioners, and policymakers interested in this field.

Second, our design science study fills a specific research gap in integrating digital model architectures into the business continuity management domain. We created a motivational and strategic model to show how cyber resilience and business continuity management align, contributing to the theoretical understanding of both fields. Furthermore, by examining three business continuity management processes, we found ways to make improvements and used these findings to propose a solution design. This proposed architecture helps practitioners understand how a digital model solution architecture can decrease disaster recovery speed. It also offers a balanced view of the benefits and potential challenges, aiding in informed decision-making and implementation.

Last, this research has contributed to practice by helping Northwave to understand the potential of using such to solution to improve business continuity. However, due to the fact that this solution design still faces many challenges as discussed in chapter 8, it currently is not a high priority for Northwave to conduct further investigations into this area. In specific, this solution is aimed towards a small target group, which we estimate to be too small for big advancements here as explained in section 8.1. Also, Northwave at this point in time does not advise organizations on the hardware that they should be using, for which this solution design does require specific functionalities (mainly being able to interact with them in a specific manner). Nonetheless, even considering all these notes, I do advise Northwave to keep track of the newest developments in this technological domain in order to identify opportunities when they arise. VM2020 is one provider of an implementation of the solution architecture and there might be more to come.

9.2 Suggestions for Future Work

Due to the exploratory nature of this study, many suggestions for future research can be made to build further on the theory. Next to the research gaps identified in the rapid review as mentioned in section 4.4, the design study also brought many ideas for future studies to light, which we will mention here.

First and foremost, we would suggest future studies to create a pilot implementation and test it. In doing so, insights can be gathered on the technical and practical feasibility of the solution and its performance. This will serve as another validation of the RDMA.

Next, future studies can look into the possibility of using the RDMA for other business continuity management processes next to the disaster recovery, disaster recovery test and business impact analysis. This can then be used to create a holistic framework to integrate the RDMA into business continuity management as a whole in the context of cyber resilience. This can offer practitioners even more benefits towards improving business continuity and cyber resilience. Additionally, if the integration also decreases recovery time further, this can raise the return on investment of the solution.

As mentioned in section 8.6 of the discussion chapter, the RDMA might also serve the ART framework in the form of a purple teaming environment to conduct cyber war games. Future studies could evaluate this possibility to add yet another use case to the collection of contributions to improve cyber resilience.

Bibliography

[1] David Allison, Paul Smith, and Kieran McLaughlin. Digital Twin-Enhanced Incident Response for Cyber-Physical Systems. In *ACM International Conference Proceeding Series*. Association for Computing Machinery, 8 2023. doi:10.1145/3600160.3600195.

[2] Andre Andrade. The 3 levels of the Digital Twin technology, 2022. URL: https://vidyatec.com/blog/the-3-levels-of-the-digital-twin-technology-2/.

[3] Andrew Hiles. The Definitive Handbook of Business Continuity Management. *The Definitive Handbook of Business Continuity Management*, 1 2012. URL: https://onlinelibrary.wiley.com/doi/book/10.1002/9781119205883, doi:10.1002/9781119205883.

[4] Archi. Resources – Archi. URL: https://www.archimatetool.com/resources/.

[5] ArchiMate. Introduction: ArchiMate® 3.2 Specification. URL: https://pubs.opengroup.org/architecture/archimate32-doc/.

[6] Kerstin Awiszus, Yannick Bell, Jan Lüttringhaus, Gregor Svindland, Alexander Voß, and Stefan Weber. Building Resilience in Cybersecurity – An Artificial Lab Approach. *Journal of Risk and Insurance*, 11 2022. URL: http://arxiv.org/abs/2211.04762.

[7] Adrien Bécue, Martin Praddaude, Eva Maia, Nicolas Hogrel, Isabel Praça, and Reda Yaich. Digital Twins for Enhanced Resilience: Aerospace Manufacturing Scenario. In *Lecture Notes in Business Information Processing*, volume 451, pages 107–118. Springer Science and Business Media Deutschland GmbH, 2022. doi:10.1007/978-3-031-07478-3{_}9.

[8] Rob Bemthuis, Maria Eugenia Iacob, and Paul Havinga. A Design of the Resilient Enterprise: A Reference Architecture for Emergent Behaviors Control. *Sensors 2020, Vol. 20, Page 6672*, 20(22):6672, 11 2020. URL: https://www.mdpi.com/1424-8220/20/22/6672/htmhttps://www.mdpi.com/1424-8220/20/22/6672, doi:10.3390/S20226672.

[9] Ran Bhamra, Samir Dani, and Kevin Burnard. Resilience: The concept, a literature review and future directions. *International Journal of Production Research*, 49(18):5375 – 5393, 9 2011. doi:10.1080/00207543.2011.563826.

[10] Deborah J Bodeau and Richard Graubart. Cyber Resiliency Engineering Framework. Technical report, MITRE Corporation, 2011.

[11] Stefan Boschert, Christoph Heinrich, and Roland Rosen. Next Generation Digital Twin. 2018. URL: https://www.researchgate.net/publication/325119950_Next_Generation_Digital_Twin.

[12] British Standards Institution. Business continuity management : Part 2., Specification. Technical report, British Standards Institution, 2007.

[13] Umit Cali, Berhane Darsene Dimd, Parisa Hajialigol, Amin Moazami, Sri Nikhil Gupta Gourisetti, Gabriele Lobaccaro, and Mohammadreza Aghaei. Digital Twins: Shaping the Future of Energy Systems and Smart Cities through Cybersecurity, Efficiency, and Sustainability. In *2023 International Conference on Future Energy Solutions, FES 2023*. Institute of Electrical and Electronics Engineers Inc., 2023. doi:10.1109/FES57669.2023.10182868.

[14] Bruno Cartaxo, Gustavo Pinto, and Sergio Soares. Rapid Reviews in Software Engineering. *Contemporary Empirical Methods in Software Engineering*, pages 357–384, 2020. URL: https://link.springer.com/chapter/10.1007/978-3-030-32489-6_13, doi:10.1007/978-3-030-32489-6{_}13/COVER.

[15] Paul Cichonski, Tom Millar, Tim Grance, and Karen Scarfone. Computer Security Incident Handling Guide Recommendations of the National Institute of Standards and Technology. Technical report, National Institute of Standards and Technology, 2012. URL: http://dx.doi.org/10.6028/NIST.SP.800-61r2, doi:10.6028/NIST.SP.800-61r2.

[16] DeNederlanscheBank. Advanced Red Teaming (ART) Framework for the financial sector. Technical report, DeNederlandscheBank, 2024. URL: https://www.dnb.nl/media/jxzbjyms/art-framework-april-2024-2.pdf.

[17] Marietheres Dietz and Gunther Pernul. Unleashing the Digital Twin's Potential for ICS Security. *IEEE Security and Privacy*, 18(4):20–27, 7 2020. doi:10.1109/MSEC.2019.2961650.

[18] Marietheres Dietz, Daniel Schlette, and Gunther Pernul. Harnessing Digital Twin Security Simulations for systematic Cyber Threat Intelligence. In *Proceedings - 2022 IEEE 46th Annual Computers, Software, and Applications Conference, COMPSAC 2022*, pages 789–797. Institute of Electrical and Electronics Engineers Inc., 2022. doi:10.1109/COMPSAC54236.2022.00129.

[19] Philip Empl, Daniel Schlette, Daniel Zupfer, and Günther Pernul. SOAR4IoT: Securing IoT Assets with Digital Twins. In *ACM International Conference Proceeding Series*. Association for Computing Machinery, 8 2022. doi:10.1145/3538969.3538975.

[20] ENISA. NIS 2 Directive, 2022. URL: https://eur-lex.europa.eu/eli/dir/2022/2555/oj.

[21] Gregory Epiphaniou, Mohammad Hammoudeh, Hu Yuan, Carsten Maple, and Uchenna Ani. Digital twins in cyber effects modelling of IoT/CPS points of low resilience. *Simulation Modelling Practice and Theory*, 125, 5 2023. doi:10.1016/j.simpat.2023.102744.

[22] ETSI. TS 128 530 - V15.2.0 - 5G; Management and orchestration; Concepts, use cases and requirements (3GPP TS 28.530 version 15.2.0 Release 15). Technical report, TSGS, 2019. URL: https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx.

[23] ETSI. TS 128 530 - V17.2.0 - 5G; Management and orchestration; Concepts, use cases and requirements (3GPP TS 28.530 version 17.2.0 Release 17). Technical report, TSGS, 2022. URL: https://portal.etsi.org/TB/ETSIDeliverableStatus.aspx.

[24] Rajiv Faleiro, Lei Pan, Shiva Raj Pokhrel, and Robin Doss. Digital Twin for Cybersecurity: Towards Enhancing Cyber Resilience. *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, LNICST*, 413 LNICST:57–76, 2022. URL: https://link.springer.com/chapter/10.1007/978-3-030-93479-8_4, doi:10.1007/978-3-030-93479-8{_}4/FIGURES/4.

[25] Silmie Vidiya Fani and Apol Pribadi Subriadi. Business Continuity Plan: Examining of Multi-Usable Framework. *Procedia Computer Science*, 161:275–282, 1 2019. doi:10.1016/J.PROCS.2019.11.124.

[26] S. E. Galaitsi, Elizaveta Pinigina, Jeffrey M. Keisler, Gianluca Pescaroli, Jesse M. Keenan, and Igor Linkov. Business Continuity Management, Operational Resilience, and Organizational Resilience: Commonalities, Distinctions, and Synthesis. *International Journal of Disaster Risk Science*, 14(5):713–721, 10 2023. doi:10.1007/S13753-023-00494-X.

[27] Lawrence A. Gordon and Martin P. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)*, 5(4):438–457, 11 2002. URL: https://dl.acm.org/doi/10.1145/581271.581274, doi:10.1145/581271.581274.

[28] Lawrence A. Gordon, Martin P. Loeb, and Lei Zhou. Investing in Cybersecurity: Insights from the Gordon-Loeb Model. *Journal of Information Security*, 07(02):49–59, 2016. doi:10.4236/JIS.2016.72004.

[29] Chiara Grasselli, Andrea Melis, Lorenzo Rinieri, Davide Berardi, Giacomo Gori, and Amir Al Sadi. An Industrial Network Digital Twin for enhanced security of Cyber-Physical Systems. In *2022 International Symposium on Networks, Computers and Communications, ISNCC 2022*. Institute of Electrical and Electronics Engineers Inc., 2022. doi:10.1109/ISNCC55209.2022.9851731.

[30] Tracy Gregorio. Digital Twins Key to Cyber Resilient Infrastructure, 12 2022. URL: https://www.govtech.com/opinion/digital-twins-key-to-cyber-resilient-infrastructure.

[31] Chan Gu, Chen Chen, and Wei Tang. Accurate and fast machine learning algorithm for systems outage prediction. *Solar Energy*, 251:286–294, 2 2023. doi:10.1016/j.solener.2023.01.014.

[32] Brahim Herbane. The evolution of business continuity management: A historical review of practices and drivers. *Business History*, 52(6):978–1002, 2010. URL: https://www.researchgate.net/publication/227608980_The_Evolution_of_Business_Continuity_Management_A_Historical_Review_of_Practices_and_Drivers, doi:10.1080/00076791.2010.511185.

[33] Hitachi. Development of Security Digital Twin Technology for Planning Security Countermeasures Ensuring Business Continuity : Research & Development : Hitachi, 2 2023. URL: https://www.hitachi.com/rd/news/topics/2023/2301_dt.html.

[34] David Holmes, Maria Papathanasaki, Leandros Maglaras, Mohamed Amine Ferrag, Surya Nepal, and Helge Janicke. Digital Twins and Cyber Security - solution or challenge? *6th South-East Europe Design Automation, Computer Engineering, Computer Networks and Social Media Conference, SEEDA-CECNSM 2021*, 2021. doi:10.1109/SEEDA-CECNSM53056.2021.9566277.

[35] Hsiu Fang Hsieh and Sarah E. Shannon. Three approaches to qualitative content analysis. *Qualitative Health Research*, 15(9):1277–1288, 11 2005. doi:10.1177/1049732305276687.

[36] International Organization for Standardization. ISO 22301:2019 - Security and resilience — Business continuity management systems — Requirements. URL: https://www.iso.org/standard/75106.html.

[37] Sumeet Jauhar, Binbin Chen, William G. Temple, Xinshu Dong, Zbigniew Kalbarczyk, William H. Sanders, and David M. Nicol. Model-Based Cybersecurity Assessment with NESCOR Smart Grid Failure Scenarios. *Proceedings - 2015 IEEE 21st Pacific Rim International Symposium on Dependable Computing, PRDC 2015*, pages 319–324, 1 2016. doi:10.1109/PRDC.2015.37.

[38] John Sharp. Moving from BS 25999-2 to ISO 22301,The new international standard for business continuity management systems. Technical report, The British Standards Institution, 2012.

[39] David Jones, Chris Snider, Aydin Nassehi, Jason Yon, and Ben Hicks. Characterising the Digital Twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology*, 29:36–52, 5 2020. doi:10.1016/J.CIRPJ.2020.02.002.

[40] Hanna Kallio, Anna Maija Pietilä, Martin Johnson, and Mari Kangasniemi. Systematic methodological review: developing a framework for a qualitative semi-structured interview guide. *Journal of Advanced Nursing*, 72(12):2954–2965, 12 2016. URL: https://onlinelibrary-wiley-com.ezproxy2.utwente.nl/doi/full/10.1111/jan.13031https://onlinelibrary-wiley-com.ezproxy2.utwente.nl/doi/abs/10.1111/jan.13031https://onlinelibrary-wiley-com.ezproxy2.utwente.nl/doi/10.1111/jan.13031, doi:10.1111/JAN.13031.

[41] Mika Karjalainen and Tero Kokkonen. Comprehensive Cyber Arena; the Next Generation Cyber Range. *Proceedings - 5th IEEE European Symposium on Security and Privacy Workshops, Euro S and PW 2020*, pages 11–16, 9 2020. URL: https://www.researchgate.net/publication/344886631_Comprehensive_Cyber_Arena_The_Next_Generation_Cyber_Range, doi:10.1109/EUROSPW51379.2020.00011.

[42] Werner Kritzinger, Matthias Karner, Georg Traar, Jan Henjes, and Wilfried Sihn. Digital Twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11):1016–1022, 1 2018. doi:10.1016/J.IFACOL.2018.08.474.

[43] Eva Maia, Sinan Wannous, Tiago Dias, Isabel Praça, and Ana Faria. Holistic Security and Safety for Factories of the Future. *Sensors*, 22(24), 12 2022. doi:10.3390/s22249915.

[44] Management Events. DIGITAL TWINS FOR CYBER SECURITY: STRENGTHENING CYBER RESILIENCE, 1 2021. URL: https://managementevents.com/news/digital-twins-for-cyber-security/.

[45] M. Manbachi and M. Hammami. Virtualized Experiential Learning Platform (VELP) for Smart Grids and Operational Technology Cybersecurity. In *Proceedings - 2022 IEEE 2nd International Conference on Intelligent Reality, ICIR 2022*, pages 54–57. Institute of Electrical and Electronics Engineers Inc., 2022. doi:10.1109/ICIR55739.2022.00027.

[46] Massimiliano Masi, Giovanni Paolo Sellitto, Helder Aranha, and Tanja Pavleska. Securing critical infrastructures with a cybersecurity digital twin. *Software and Systems Modeling*, 22(2):689–707, 4 2023. URL: https://link.springer.com/article/10.1007/s10270-022-01075-0, doi:10.1007/S10270-022-01075-0/TABLES/4.

[47] Marko Niemimaa. Interdisciplinary review of business continuity from an information systems perspective: Toward an integrative framework. *Communications of the Association for Information Systems*, 37:69–102, 2015. doi:10.17705/1CAIS.03704.

[48] Parsifal. About Parsifal. URL: https://parsif.al/about/.

[49] Ken Peffers, Tuure Tuunanen, Marcus A. Rothenberger, and Samir Chatterjee. A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3):45–77, 12 2007. doi:10.2753/MIS0742-1222240302.

[50] Perplexity. Perplexity Frequently Asked Questions. URL: https://www.perplexity.ai/hub/faq.

[51] Sandeep Pirbhulal, Habtamu Abie, and Ankur Shukla. Towards a Novel Framework for Reinforcing Cybersecurity using Digital Twins in IoT-based Healthcare Applications. In *IEEE Vehicular Technology Conference*, volume 2022-June. Institute of Electrical and Electronics Engineers Inc., 2022. doi:10.1109/VTC2022-Spring54318.2022.9860581.

[52] Filippo Rebecchi, Antonio Pastor, Alberto Mozo, Chiara Lombardo, Roberto Bruschi, Ilias Aliferis, Roberto Doriguzzi-Corin, Panagiotis Gouvas, Antonio Alvarez Romero, Anna Angelogianni, Ilias Politis, and Christos Xenakis. A Digital Twin for the 5G Era: the SPIDER Cyber Range. In *Proceedings - 2022 IEEE 23rd International Symposium on a World of Wireless, Mobile and Multimedia Networks, WoWMoM 2022*, pages 567–572. Institute of Electrical and Electronics Engineers Inc., 2022. doi:10.1109/WoWMoM54355.2022.00088.

[53] Pascal Renckens, Luca Ferentinos, and Alex Wielart. Never Waste a Good Crisis: The Six Success Factors of Cyber Resilience. Technical report, Northwave Cyber Security, 2023. URL: www.northwave-cybersecurity.com.

[54] Ron Ross, Victoria Pillitteri, Kelley Dempsey, Mark Riddle, and Gary Guissanie. NIST Special Publication 800-171 Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations. Technical report, National Institute of Standards and Technology, 2020. doi:10.6028/NIST.SP.800-171r2.

[55] Ron Ross, Victoria Pillitteri, Richard Graubart, Deborah Bodeau, and Rosalie Mcquaid. NIST Special Publication 800-160, Volume 2 Developing Cyber-Resilient Systems: A Systems Security Engineering Approach. Technical report, National Institute of Standards and Technology, 2015. doi:10.6028/NIST.SP.800-160v2r1.

[56] Nelson Russo and Leonilde Reis. Methodological approach to systematization of business continuity in organizations. *Research Anthology on Business Continuity and Navigating Times of Crisis*, 1:289–312, 1 2022. URL: https://www.researchgate.net/publication/358115955_Methodological_Approach_to_Systematization_of_Business_Continuity_in_Organizations, doi:10.4018/978-1-6684-4503-7.CH015.

[57] Nelson Russo, Leonilde Reis, Clara Silveira, and Henrique S. Mamede. Towards a Comprehensive Framework for the Multidisciplinary Evaluation of Organizational Maturity on Business Continuity Program Management: A Systematic Literature Review. *Information Security Journal: A Global Perspective*, 33(1):54–72, 1 2024. URL: https://www.tandfonline.com/doi/abs/10.1080/19393555.2023.2195577, doi:10.1080/19393555.2023.2195577.

[58] Andrea Salvi, Paolo Spagnoletti, and Nadia Saad Noori. Cyber-resilience of Critical Cyber Infrastructures: Integrating digital twins in the electric power ecosystem. *Computers and Security*, 112, 1 2022. doi:10.1016/j.cose.2021.102507.

[59] Giovanni Paolo Sellitto, Massimiliano Masi, Tanja Pavleska, and Helder Aranha. A Cyber Security Digital Twin for Critical Infrastructure Protection: The Intelligent Transport System Use Case. In *Lecture Notes in Business Information Processing*, volume 432 LNBIP, pages 230–244. Springer Science and Business Media Deutschland GmbH, 2021. doi:10.1007/978-3-030-91279-6{_}16.

[60] Daniel A. Sepúlveda Estay, Rishikesh Sahay, Michael B. Barfod, and Christian D. Jensen. A systematic review of cyber-resilience assessment frameworks. *Computers & Security*, 97:101996, 10 2020. doi:10.1016/J.COSE.2020.101996.

[61] Siddhant Shrivastava, Francisco Furtado, Mark Goh, and Aditya Mathur. The Design of Cyber-Physical Exercises (CPXS). In *International Conference on Cyber Conflict, CYCON*, volume 2022-May, pages 347–365. NATO CCD COE Publications, 2022. doi:10.23919/CyCon55549.2022.9811000.

[62] D Simon, K Fischbach, and D Schoder. An Exploration of Enterprise Architecture Research. *Communications of the Association for Information Systems*, 32(1):1–72, 2013. URL: https://doi.org/10.17705/, doi:10.17705/1CAIS.03201.

[63] Tuomo Sipola, Tero Kokkonen, Markku Puura, Kalle Eemeli Riuttanen, Kari Pitkäniemi, Elina Juutilainen, and Teemu Kontio. Digital Twin of Food Supply Chain for Cyber Exercises. *Applied Sciences (Switzerland)*, 13(12), 6 2023. doi:10.3390/app13127138.

[64] SonicWall. 2024 SonicWall Cyber Threat Report. Technical report, Inc, SonicWall, 2024. URL: https://www.sonicwall.com/medialibrary/en/white-paper/2024-cyber-threat-report.pdf.

[65] Sabah Suhail, Mubashar Iqbal, Rasheed Hussain, and Raja Jurdak. ENIGMA: An explainable digital twin security solution for cyber–physical systems. *Computers in Industry*, 151, 10 2023. doi:10.1016/j.compind.2023.103961.

[66] Jon Swain and Brendan King. Using Informal Conversations in Qualitative Research. *International Journal of Qualitative Methods*, 21:1–10, 3 2022. URL: https://us.sagepub.com/en-us/nam/open-access-at-sage, doi:10.1177/16094069221085056.

[67] Rama Lingeswara Tammineedi. Business continuity management: A standards-based approach. *Information Security Journal*, 19(1):36–50, 2010. doi:10.1080/19393550903551843.

[68] The Business Continuity Intitute. Good Practice Guidelines (GPG) Edition 7.0, 11 2023. URL: https://www.thebci.org/resource/good-practice-guidelines--gpg--edition-7-0.html.

[69] S. Ali Torabi, Ramin Giahi, and Navid Sahebjamnia. An enhanced risk assessment framework for business continuity management systems. *Safety Science*, 89:201–218, 11 2016. doi:10.1016/J.SSCI.2016.06.015.

[70] Saurabh Vaidya, Prashant Ambad, and Santosh Bhosle. Industry 4.0 – A Glimpse. *Procedia Manufacturing*, 20:233–238, 1 2018. doi:10.1016/J.PROMFG.2018.02.034.

[71] VM2020. VM2020 Solutions. URL: https://www.vm2020.com/.

[72] Gerit Wagner, Roman Lukyanenko, and Guy Paré. Artificial intelligence and the conduct of literature reviews. *Journal of Information Technology*, 37(2):209–226, 6 2022. URL: https://journals.sagepub.com/doi/full/10.1177/02683962211048201, doi:10.1177/02683962211048201/ASSET/IMAGES/LARGE/10.1177{_}02683962211048201-FIG1.JPEG.

[73] Roel J. Wieringa. Design science methodology: For information systems and software engineering. *Design Science Methodology: For Information Systems and Software Engineering*, pages 1–332, 1 2014. doi:10.1007/978-3-662-43839-8.

[74] Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Björn Regnell, and Anders Wesslén. Experimentation in software engineering. *Experimentation in Software Engineering*, 9783642290442:1–236, 7 2012. doi:10.1007/978-3-642-29044-2/COVER.

[75] Muhammad Mudassar Yamin, Basel Katt, and Vasileios Gkioulos. Cyber ranges and security testbeds: Scenarios, functions, tools and architecture. *Computers & Security*, 88:101636, 1 2020. doi:10.1016/J.COSE.2019.101636.

[76] Robert K. Yin. Case Study Research and Applications: Design and Methods. *SAGE publications*, (6th ed.), 2017.

[77] Haifeng Zhou, Mohan Li, Yanbin Sun, Lei Yun, and Zhihong Tian. Digital Twin-Based Cyber Range for Industrial Internet of Things. *IEEE Consumer Electronics Magazine*, 12(6):66–77, 11 2023. doi:10.1109/MCE.2022.3203202.

Appendix A

Semi-structured interview guide literature review

Semi-structured interview (interview guide)		
Preliminary	Study goal and context Selected articles from the rapid review	
Part I	Introduction interviewee	1. What is your profession and which role do you have within the organization that you work for? 2. What previous experience do you have with Digital Twins in a cybersecurity context?
Part II	Data extraction	3. How many articles do you recognise from the list of included articles that were discovered in the rapid review? 4. Are there any articles that you know are relevant to include in this study, which were not included in the rapid review? 4.a. If yes, which articles? 4.b. Why do you think they are relevant to the study?

Table A.1: Semi-structured interview guide

Appendix B

Semi-structured interview literature review results

Interviewee 1	
Name	Participant 2
Occupation	PhD student, researcher at the University of Twente
Previous experience on topic	PhD topic is on utilising DTs to improve cyber security of critical infrastructure
Recognized articles	13
Relevant articles	Allison et al., 'Digital Twin-Enhanced Incident Response for Cyber-Physical Systems' [1]
Why	Allison: Discusses IR and integrating DT as part of IR playbooks

Table B.1: Summarised results from interviewee 1

Interviewee 2	
Name	Participant 1
Occupation	Professional within computer science and cybersecurity. Lecturer in field of cybersecurity for the University of Twente.
Previous experience on topic	Engaged with DT since 2008, specifically within cybersecurity. Supervised many projects exploring the potential of the DT for cybersecurity purposes, especially for critical infrastructure.
Recognized articles	20+
Relevant articles	Allison et al., 'Digital Twin-Enhanced Incident Response for Cyber-Physical Systems' Epiphaniou et al. 'Digital twins in cyber effects modelling of IoT/CPS points of low resilience'
Why	Allison: Discusses IR and integrating DT as part of IR playbooks Epiphaniou: Using the DT for strategic cyber decision-making in enhancing resilience within the IoT/CPS system

Table B.2: Summarised results from interviewee 2

Interviewee 3	
Name	Participant 3
Occupation	Cyber Resilience Consultant at Northwave Cyber Security.
Previous experience on topic	No experience with DTs. Does have experience with BCM and CR.
Recognized articles	0
Relevant articles	None
Why	N/A

Table B.3: Summarised results from interviewee 3

Appendix C

Semi-structured interview guide case study

Semi-structured interview (interview guide)		
Preliminary	Study goal and context Case context	
Part 1	Introduction Interviewee	1. What is your profession and which role do you have within the organization that you work for? 2. What previous experience do you have with <problem context>?
Part 2	Data extraction problem context	3. What are the key processes, actors, components and software systems involved in <problem context>? 4. How do these processes, actors, components and software systems interact with each other? 5. What is the main driver and/or goal of <problem context>? 6. What are the main challenges in the context of business continuity management in <problem context>?
Part 3	Data extraction requirements	7. Considering the digital model technology, what requirements it need to comply with to improve <problem context>?

Appendix D

Semi-structured interview case study results

D.1 Disaster Recovery

The results discussed here are a summary of the transcription of the recorded sessions with the interviewee. For the general purpose of this research and due to the sensitive information that might be involved, only a summary is provided.

This interview was conducted in multiple parts due to time constraints. The interviewee is part of the Computer Emergency Response Team (CERT) within Northwave and has experience dealing with cyber incidents for multiple years. In the first session, the interviewee raised the importance of a wider understanding of how cyber attacks are executed and what actually happens before and during a cyber attack, to understand the process of a disaster recovery in depth. For this reason, the interviewee gave a presentation and in-depth overview of the processes and components involved in a cyber incident. The interviewee started with the process of a cyber attack. The so-called, in, through, and out phases were discussed here. To explain this process, the interviewee showed and explained a case using a drawing of the IT infrastructure.

After this explanation, the interviewee went in-depth into what organizations can do to recover from an attack. There are three main options: paying, recovering from backups or rebuilding the IT infrastructure. Considering the context of the study, we delved into the process of recovering from a backup. There are multiple phases in this process, and using the case that was used before, the pre-recovery steps were explained next to the challenges that organizations often face of not having a clear overview of their IT infrastructure. Afterwards, the typical eradication method was explained and the hardening steps to ensure a threat actor does not get access to the system again.

In a later session, some of the parts of the recovery procedure were discussed in more depth, and the requirements for the digital model technology were discussed. Here the interview also strayed into the general opportunities of such a solution, which was used as inspiration for the solution design.

D.2 Disaster Recovery Test

The results discussed here are a summary of the transcription of the recorded sessions with the interviewee. For the general purpose of this research and due to the sensitive information that might be involved, only a summary is provided.

The interviewee, an experienced professional in the field of cyber resilience, explained the process and components of a disaster recovery test. Preliminary knowledge was acquired during the document analysis, which was used as a foundation for the explanation of the process itself. To start, the interviewee emphasized the motivation and strategy behind conducting regular disaster recovery tests. They explained that these tests are essential for validating the effectiveness of an organization's disaster recovery plan and identifying potential weaknesses or gaps. In continuation, the interviewee outlined the key steps in a typical disaster recovery test process:

1. Planning and preparation: This involves defining the scope of the test, identifying key systems and procedures to be tested, and establishing clear objectives and success criteria. This also involves developing a script which lays out all these parts.

2. Test execution: The actual simulation of a disaster scenario and the implementation of recovery procedures. Here the interviewee also explained how the testing environment was usually set up. This is also where most of the challenges appeared to execute such a test. This may involve partial or full system recovery from backups, data restoration, and failover to secondary systems depending on the script.

3. Evaluation and documentation: After the test, a thorough analysis is conducted to assess the effectiveness of the recovery procedures, identify any issues or bottlenecks, and document lessons learned.

Lastly, considering the digital model technology, the requirements for such a system were discussed. Mostly the ease of use and the realisticness of a real disaster recovery were emphasized.

D.3 Business Impact Analysis

The results discussed here are a summary of the transcription of the recorded sessions with the interviewee. For the general purpose of this research and due to the sensitive information that might be involved, only a summary is provided.

The interviewee, an experienced professional in the field of cyber resilience and business continuity management, explained the process and components of a business impact analysis. In the same manner as the disaster recovery test process, documents were analyzed and used as a basis for the interview.

To start the process, the initiator of the business impact analysis chooses a critical business function in which they want to understand the dependencies and business continuity requirements. By analysing each process within the critical business function, a dependency analysis is executed where the involved parties discuss whether the process is dependent on an artefact.

After this is done for every process within the business function, the business continuity requirements are appointed. These define the objectives and maximum times of how long a discontinuity or information loss may take. Finally, all this information is gathered and reported.

Appendix E

ArchiMate 3.2 Specification Overview

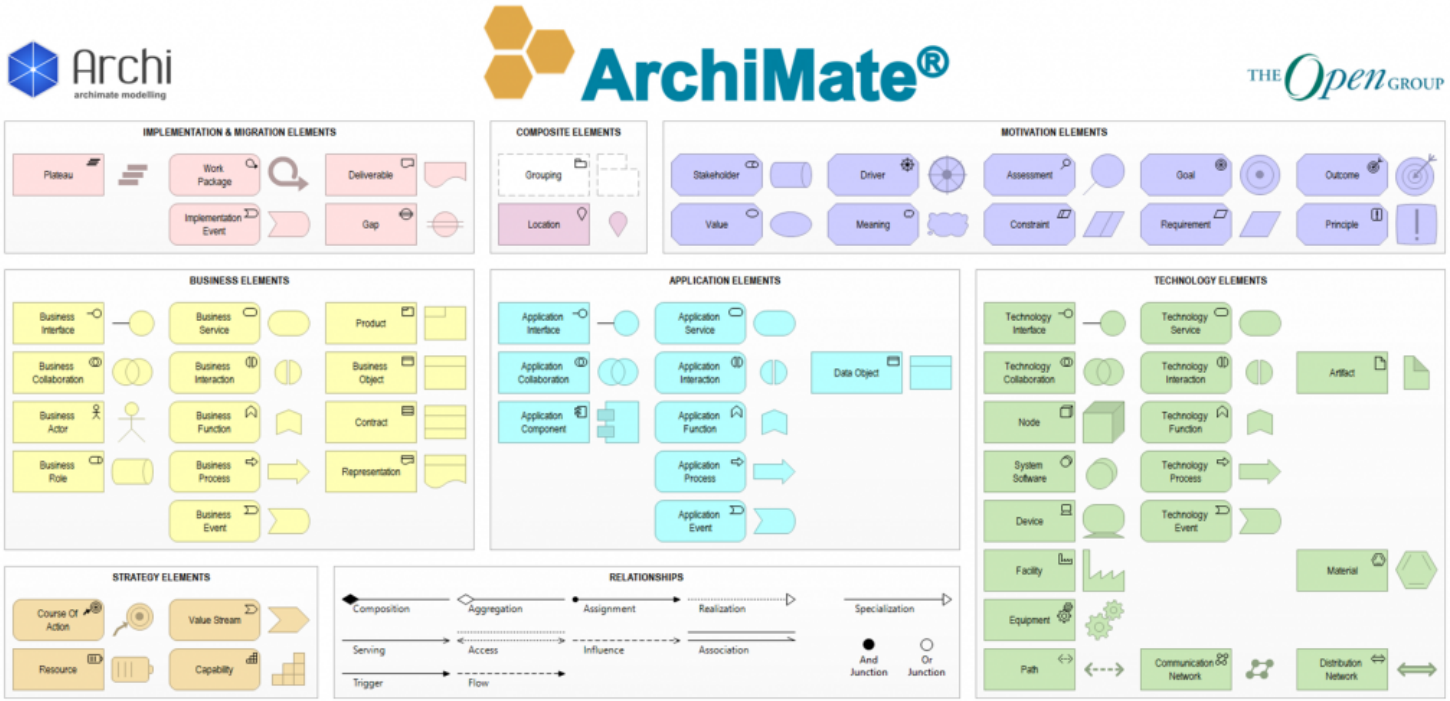

FIGURE E.1: ArchiMate 3.2 Specification Overview [4]