The ZEro Trust DECision Making (ZEDEC) Method: Selecting Relevant Zero Trust Concepts to Mitigate High-Priority Risks

Master's thesis
Business Informatics

Supervisors
dr. K. Labunets
dr. S. Jansen

External supervisor
E. de Vries

Author
B. van Dijen
6019633

August, 2023

Abstract

Zero trust is a security principle that allows organisations to be more resilient to cyber threats than the traditional perimeter-based security solutions, by ensuring that users and devices are not trusted by default. In order to minimize security risks, 72 % of the companies were planning to implement zero trust capabilities into their security solutions in 2020. However, zero trust architectures have not yet succeeded in replacing traditional security solutions, because organisations have trouble overseeing how they should handle the migration process — the design of a zero trust architecture differs per organisation, depending on their needs. Therefore, this research proposes a ZEro trust DECision making (ZEDEC) method that helps these organisations decide which zero trust concepts they should integrate into their zero trust architecture. We follow the design science method to construct ZEDEC. In the first part of the research, we conduct a structured literature review and expert interviews to identify relevant method fragments. We identify multiple zero trust concepts that organisations should consider to integrate into their zero trust architecture and discover how organisations are currently migrating towards a zero trust architecture. We also discover that in the migration process towards a zero trust architecture, organisations mainly consider zero trust mitigations that address cyber security risks as factors to decide which zero trust concepts they want to include in their zero trust architecture. Therefore, we introduce a decision matrix that includes a mapping between the zero trust mitigations and their related zero trust concepts. Ultimately, we propose the ZEDEC method with the main activities: (1) create a vision on needed changes, (2) identify context, (3) perform risk assessment and (4) identify zero trust concepts. Through an evaluation with a security expert of a theoretical scenario, we confirm that the method is correct and useful.

Keywords— Zero trust architecture, risk management, decision making, zero trust migration, ZEDEC method, MITRE ATT&CK framework

Contents

Chapter 1

Introduction

With the growing impact of IT within companies, as well as the growing amount of employees that work from home, organisations are forced to restructure their security strategies. The increasing amount of data traffic and variety of endpoint locations requires strict policies in which IT systems are protected from inside and outside the network perimeter – defending the organisation's network from outsiders is not enough [16]. Zero trust is a security approach that meets those strict policies, stating that one should never trust an entity in the network, but verify its identity before granting access to resources [12, 44]. Furthermore, it requires an organisation to grant users the least amount of access they need and continuously monitor all logs and data traffic [4, 47, 61].

The zero trust approach helps minimize security risks, so organisations that harshly rely on their IT systems or data security are currently implementing zero trust principles or taking steps to do so. This trend was confirmed in 2020, when 72% of companies were planning to implement zero trust capabilities [14, 18]. Big tech-organisations also come with zero trust implementations, like Google with the *BeyondCorp* technology [63]. Although typical zero trust concepts like Multi-Factor Authentication (MFA) are in line with the zero trust tenets and have a positive impact on the security of the IT landscape of the organisations, the full potential of the zero trust approach is not yet extracted with just the implementation of a few concepts. To maximize the effect of the zero trust approach in a security landscape, more thorough architectural changes are required, with the creation of a *Zero Trust Architecture* (ZTA) as a result. Rose et al. [44] define a ZTA as *an enterprise's cybersecurity plan that utilizes zero trust concepts and encompasses component relationships, workflow planning, and access policies.*

When a ZTA is configured the right way, it can protect enterprise resources in a more effective way than traditional solutions. In the first place, strict authentication policies better protect the data against attacks and breaches. Second, the segmentation of the enterprise resources prevents lateral movement within the network. Third, the strict policies and network segmentation allow a ZTA to have fine-grained access models. Furthermore, the continuous logging and monitoring of the network allow to detect and interrupt suspicious behavior and irregularities. Finally, a ZTA may be better protected against DDoS attacks than traditional solutions [12].

1.1 Problem and Gap

ZTA's are proven to be effective and organisations are willing to implement zero trust capabilities [12]. Numerous zero trust solutions are available in current literature that can be used in ZTA's [44]. However, a comprehensive overview of these zero trust concepts is lacking, as well as guidance on how to select which zero trust concepts one wants to include in its IT landscape [25]. Especially research about experiences of organisations that migrated towards a ZTA is lacking in literature [12]. Because of this, *organisations are experiencing trouble with selecting the zero trust concepts that fit the needs of their IT landscape.* The subsequent uncertainty about the complex changes in technology, costs, workflow and employee training withhold organisations to make the step towards a ZTA [41]. This blockage could increase the risks to their threat landcapes.

1.2 Contribution

This research aims to support organisations with their migration from a perimeter-based security solution towards a zero trust architecture, by creating the ZEro trust DECision making (ZEDEC) method that helps them decide which zero trust concepts fit their needs. This method is based on data collected by a structured literature review (SLR) and data gathered from interviewing zero trust experts. We make contributions to scientific and practical research fields.

To construct the ZEDEC method, we first create an overview of currently known zero trust concepts that one could consider integrating into a ZTA. This helps organisations to get a comprehensive overview of the possible concepts they could consider to integrate. From a scientific viewpoint, this overview can be used for future research on (a subset of) zero trust concepts and to identify areas in which more developments are needed. Secondly, we discover how organisations are currently migrating towards a ZTA according to the SLR and by conducting expert interviews to complement these findings. Organisations can follow these steps to guide them in their own decision-making. The scientific contribution is that we create an overview of the current known processes, as well as gather data about the experiences of zero trust experts with this, filling the gaps mentioned in Chapter 1.1. The third building block for the ZEDEC method is that we identify what factors should be considered that impact the decision to integrate a zero trust concept or not. The results on the decision-making factors indicate that they are similar to mitigations that help reduce risks. Therefore, we introduce a decision matrix that maps zero trust mitigations from the MITRE ATT&CK framework [55] to the related zero trust concepts. The practical contribution is that organisations can use this matrix to decide which zero trust concepts they should integrate to mitigate the risks that are relevant to them. The scientific contribution is that we gather new data from expert experiences and we discover new applications for the MITRE ATT&CK framework, which could be used for future research as well.

The ZEDEC method could be used as an inspiration to improve or create other methods in the field of security, or potentially even other research areas. The method helps organisations to be resilient against cyber threats, and the consecutive advantage for society is that the actors in the IT landscapes of these organisations are provided with appropriate security measures on their data and services.

1.3 Outline

The outline of this research is as follows. Chapter 2 provides the basics of a ZTA and explains what the advantages are over the traditional perimeter-based approach. Chapter 3 describes the research approach and the methods used. In Chapter 4 we present the results from the SLR and Chapter 5 provides the findings from the expert interviews. In Chapter 6 we construct the decision matrix and assemble the method fragments that we identified in the research to create ZEDEC. In Chapter 7 we validate to which extent this method works by demonstrating that the method works in a fictional scenario and an evaluation session with a security expert. Finally, in Chapter 8 we present and discuss the conclusions, and provide opportunities for future research.

Chapter 2

Background

Currently, many organisations rely on a perimeter-based security architecture. This chapter explains what the shortcomings in this approach are and how a ZTA handles these problems. Furthermore, it illustrates what a ZTA looks like and how it works.

2.1 Shortcomings of Traditional Solutions

In the traditional perimeter-based approach all subjects (entities that request data from enterprise resources) are considered trusted once they enter the internal network of an organisation. Despite restrictions to enter the internal network, intruders have many entry points available to potentially break in. This is because nowadays organisations enable external access to internal resources, including IoT-devices, services from external parties and devices from employees working from home [41]. The protection of the internal network is as strong as the weakest protected entry point, meaning that just one vulnerable asset with access to the internal network is enough for intruders to get in [16, 51]. Cybercrime techniques like phishing and outdated software exploitation can effectively gain intruders access to at least one of the assets, allowing them into the internal network. Once inside, they are free to move within the network and access all enterprise resources where the compromised subject is allowed through lateral movement, because they are considered trusted entities [16, 42]. Depending on the access rights of the subject, the intruders can start exploiting the targeted resources or utilize the available data and resources to gain access to other network layers to reach their goal. Therefore, the trust that a subject gets within the internal network is a threat on itself.

2.2 Basics on Zero Trust Architecture

A zero trust architecture eliminates the threat of trusting subjects that are present in the internal network. By default, no user, network or interface is trusted [30]. In order to achieve this, the traditional internal network perimeter is replaced by a centralised portal that handles all connection requests to enterprise resources. No enterprise resource can be reached unless access is explicitly granted. Rose et al. [44] describe that the main components of a ZTA are the *Policy Decision Point* (PDP) and the *Policy Enforcement Point* (PEP). An abstract visualisation of a ZTA with these components is shown in Figure 1. It includes the following components:

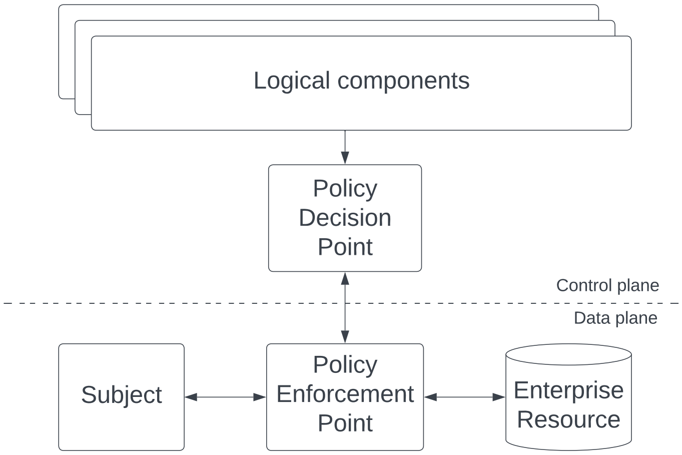

Figure 1: Zero trust architecture main components

- A **subject** is any entity that wants to access an enterprise resource. This could be a user, but also an enterprise resource on itself. For example, when an on-premise server wants to send data to an on-premise printer, the server is considered a subject. A subject communicates with a Policy Enforcement Point (PEP) to get access to an enterprise resource.

- The **Policy Enforcement Point (PEP)** is responsible for enabling, monitoring and terminating connections between subjects and enterprise resources on the data plane. It is configured by the Policy Decision Point (PDP), which is responsible for the decision on granting access to the subject. The PEP is either one element (as shown in Figure 1) or is separated into two components: one on the client side (e.g., an agent on the client's device) and one on the enterprise side (e.g., a gateway that controls access to resources). Multiple PEPs may exist within the ZTA.

- The **enterprise resources** consist of all data sources and computing devices belonging to the organisation. An enterprise resource is typically accessed by a PEP or other resources that might be behind the same PEP.

- The **Policy Decision Point (PDP)** operates at the control plane and is responsible for making the decision on granting a subject access to an enterprise resource. For every request, it executes a *Trust Algorithm* (TA) that gives a decision as output. As input, it considers data from logical components. After the TA generates a decision, the PDP logs it and configures the relevant PEPs to establish or terminate a connection.

- **Logical components** provide the PDP with information that is relevant for the decision-making. Examples of logical components are *identity management systems*, *SIEM systems* and *continuous diagnostics and mitigation (CDM) systems*. Within at least one of the logical components policies are stored. Policies formulate criteria that a subject should satisfy in order to be granted access. They have to be strictly followed by the TA before access is granted to a subject. Based on the policies and the other input variables, the PDP makes a decision. A total overview of the logical components is described in Chapter 4.

In addition to the component, Figure 1 shows a distinction between the *data* and *control plane*. The data plane is responsible for the transmission of data between enterprise resources and subjects, while the control plane is isolated from this process and judges on the allowance of data transmissions. The only way that the control plane can interact with the data plane, is by direct communication between a PEP and a PDP.

The main components form the basis of a ZTA. Multiple variations of the model exist, depending on the wishes of the organisation and the users of their systems. In addition to the abstract visualisation of a ZTA, Rose et al. [44] formulate extra rules that have to be followed to reach the best security level of the ZTA. They are extracted into the following list:

- All communication has to be secured, regardless of network location. Communication between enterprise resources needs to be secured as well.
- Access to enterprise resources is granted per session. Every time a subject wants to access another resource, access has to be reconsidered.
- Access is only granted after a strict authorization process. Policies, MFA and behavioral and environmental attributes should be part of this.
- Subjects should be allowed access to the least amount of resources necessary (according to the principle of *least privilege*). In case an intruder gets access to an account, only limited harm can be done.
- As much as possible data and access should be collected about the assets and networks. The more data a PDP can use, the more reliable the decision is.
- All assets and network information should be monitored. Irregularities can be detected and considered during the access decision process.

As the rules show, the main goal of a ZTA is to create an IT landscape that is as strictly secured as possible. Although the extent to which these rules are integrated into the IT landscape of an organisation differs on the needs of the organisation, this situation can be considered as the ideal solution. The method we propose in this research must identify which concepts fit the needs of the organisation best, in order to decide which concepts need to be integrated into their ZTA.

Chapter 3

Research Approach

This chapter describes the objective of the research, the research questions and the method of how the objective of the research is achieved. Finally, we state what the threats to the validity are to the research approach we use.

3.1 Objective

As we discussed in Chapter 1, organisations want to move towards a zero trust architecture, but have trouble with the migration towards it. This research seeks to support these organisations. Therefore we formulate our main objective based on Wieringa's template [65] as follows:

Main objective: *This thesis aims to improve* the migration process for organisations that move from a perimeter based security architecture to a zero trust architecture *by* proposing a systematic method *that* guides them in deciding which zero trust concepts to include in their zero trust architecture *in order to* be resilient against cyber threats.

3.2 Research Questions

We translate the main objective in the following research question.

RQ: *How can organisations decide which security concepts to integrate in their zero trust architecture when migrating from a perimeter-based security architecture?*

We formulate multiple sub-questions that build on each other in order to answer the main research question. We describe them below.

SQ1: *What are zero trust concepts that organisations should consider to integrate in a zero trust architecture?*
We create an overview of all concepts that one should consider to integrate in a zero trust architecture.

SQ2: *How do organisations currently decide which zero trust concepts to integrate in a zero trust architecture?*
We explore how organisations currently migrate from a perimeter-based security approach to a zero trust architecture.

SQ3: *Which factors do organisations consider before deciding which concepts to integrate in a zero trust architecture?*
We discover what elements are taken into consideration during the decision-making process. They should contribute to the selection of relevant zero trust concepts.

SQ4: *How are the organisational decision-making factors (of SQ3) related to the existing zero trust concepts (of SQ1)?*
We create an overview of concepts that might be affected by the factors. The overview is a matrix that indicates which factors should be considered for the different concepts.

SQ5: *What is a method to decide which concepts to integrate in a zero trust architecture?*
We design the ZEro trust DECision making (ZEDEC) method that helps organisations decide which concepts to integrate in their zero trust architecture.

SQ6: *To what extent does the ZEDEC method work?*
This is the validation step to discover to what extent the ZEDEC method works and whether it is considered useful for the end-users.

3.3 Research Method

This research is based on the design science method, as proposed by Wieringa [65]. The method is used to design artifacts (such as methods) in a certain context, as is the case in this study. We follow the steps of the engineering cycle, which describes how we should design our decision-making method (see Figure 2). It states that we have to start with the problem investigation, then design the method, validate the method and finally implement and evaluate it. The treatment validation and evaluation steps might trigger a refinement of the treatment and thereby a new iteration of the engineering cycle. We handle each step as follows:

1. **Problem investigation.** In this step we explore phenomena that are relevant in the context of migrating towards a ZTA. SQ1-4 states the information we want to discover

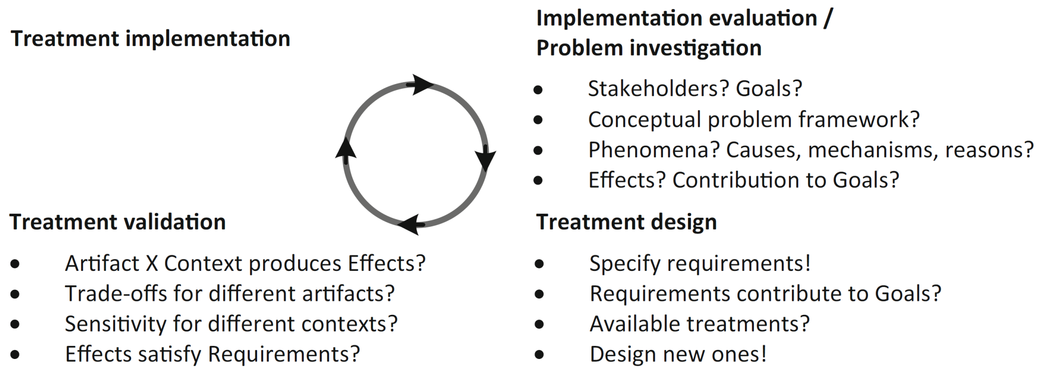
Treatment implementation
Implementation evaluation / Problem investigation
Stakeholders? Goals?
Conceptual problem framework?
Phenomena? Causes, mechanisms, reasons?
Effects? Contribution to Goals?
Treatment validation
Artifact X Context produces Effects?
Trade-offs for different artifacts?
Sensitivity for different contexts?
Effects satisfy Requirements?
Treatment design
Specify requirements!
Requirements contribute to Goals?
Available treatments?
Design new ones!

The question marks indicate knowledge questions and the exclamation marks indicate design problems

Figure 2: The engineering cycle as described by Wieringa [65]

before designing ZEDEC. On these questions we perform a structured literature review (SLR) to gain knowledge about the topics and then do back- and forwards snowballing on the found papers. Finally, we conduct additional expert interviews to complement the findings from the SLR on SQ2-4.

2. **Treatment design.** We use method engineering to design our initial treatment by assembling the method fragments identified in the *problem investigation*. This step is related to SQ5.

3. **Treatment validation.** We validate the ZEDEC method by interviewing both experts and potential users of the method to see if they find the method relevant, correct and complete. According to their feedback, changes to the initial design might be made. Furthermore, we demonstrate how ZEDEC works in a fictional scenario, to confirm how it works in theory. This step is covered by SQ6.

4. **Treatment implementation and evaluation.** Now the ZEDEC method is ready to be executed in a real-world environment and evaluated. Due to the long time span of the ZTA migrations, we do not have enough time to test it in a real-world environment. Such a case study might be interesting for future research.

Figure 3 shows how the sub-questions ultimately lead to the creation of our treatment. Table 1 shows an overview of the methods used to answer the sub-questions. The remaining part of this section is dedicated to describing these methods in detail.

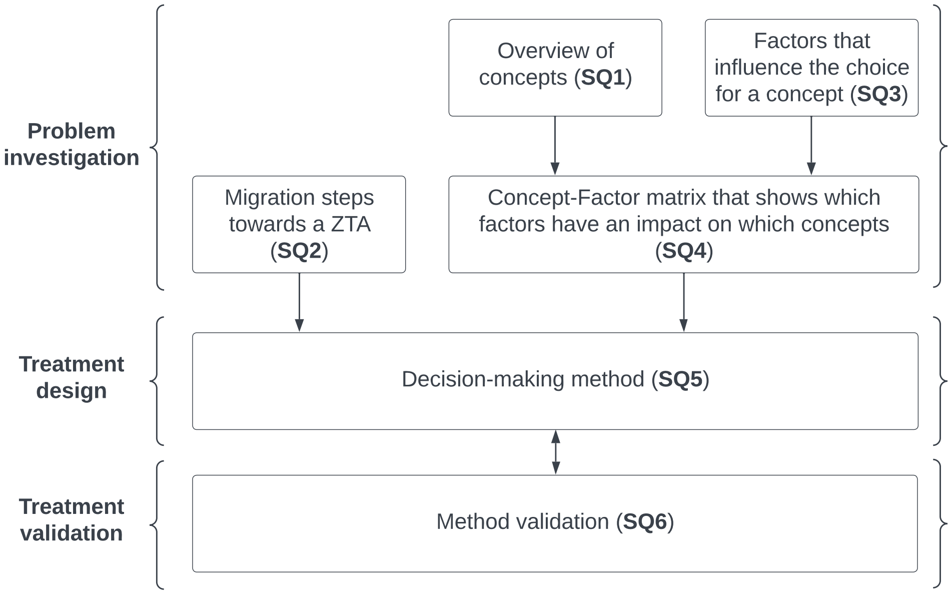
Problem investigation
Overview of concepts (SQ1)
Factors that influence the choice for a concept (SQ3)
Migration steps towards a ZTA (SQ2)
Concept-Factor matrix that shows which factors have an impact on which concepts (SQ4)
Treatment design
Decision-making method (SQ5)
Treatment validation
Method validation (SQ6)

Figure 3: Research approach per engineering cycle step

Table 1: An overview of the research methods used per engineering cycle step

	Problem investigation				Treatment design	Treatment validation
	SQ1	**SQ2**	**SQ3**	**SQ4**	**SQ5**	**SQ6**
Structured literature review	X	X	X	X		
Back- and forwards snowballing	X	X	X	X		
Interviews		X	X	X		X
Method engineering					X	

3.3.1 Structured Literature Review and Snowballing

The SLR is the first part of the *problem investigation* step, with the goal to explore all known aspects regarding zero-trust concepts (SQ1), ZTA migration (SQ2), factors to consider (SQ3) and the relation between the factors and concepts (SQ4). We focused on academic literature, since we wanted to ensure that the concepts are scientifically established and substantiated. Based on the work of Gusenbauer and Haddaway [23], we selected the search engines that are (1) relevant in the field of zero trust and (2) considered a principal (essential) search system. The search engines are: *Scopus, Web of Science, ACM Digital Library, IEEE Xplore* and *Wiley Inter Science Journal Finder*. Since the academic research performed on zero trust is limited, we formulated the search terms used to perform the SLR broadly: *"Zero-Trust" OR "Zero Trust" OR BeyondCorp*. The advantage of this broad term is that the chances of missing information are low and the literature review can be performed parallel on all four sub-questions. The *BeyondCorp* search term refers to Google's ZTA implementation. In order to capture information about these concepts as well, the search term was included.

Following the SLR guidelines proposed by Kitchenham et al. [31], we defined the inclusion and exclusion criteria to filter the title, keywords, abstract and conclusion of the identified papers. Exclusion criteria applied to the search are:

1. **Publication years before 2010.** The first time the zero-trust term was used in the context of a ZTA was in 2010 [12], so papers before that year are not about the right topic;

2. **Main focus is on AI and Blockchain.** Although these concepts are promising, most are not yet widely adopted by organisations;

3. **Article types other than journal articles and conference papers**;

4. **CiteScore rankvalue is below Q2 or unknown (for journal articles).** We calculated the CiteScore rankings by dividing the journal's citations in the previous 4 years by the number of publications in the previous 4 years. The higher the score, the more reliable the journal is. For this research, we included the top half of journals (quartiles 1 and 2) based on the CiteScore.

5. **CORE rank below B or unknown (for conference papers).** The CORE value rates conferences, of which rank B means that the conference is *good to very good*, while C scores only meet *basic standards*. Therefore, we only include conferences with score B, A or A*.

6. **Languages other than English;**

7. **Papers that are not available for download.**

We formulated the inclusion criteria per filtering step. In the first filtering step, the inclusion criterion was that the title or keywords refer to a ZTA. For papers where this was the case and exclusion criteria were not applicable, filtering step 2 is applied. Here, we formulated two inclusion criteria. When at least one criterion was met, it was included in the next filtering step. The inclusion criteria were that (1) the abstract and conclusion should include one or more concepts that are used in a ZTA and (2) one or more implementation or migration steps towards a ZTA. We did not find papers that explicitly have factors as the main topic, so we did not formulate explicit inclusion criteria for SQ3 and SQ4. Then these papers were full-text reviewed to find information that helps answer the sub-questions. Finally, we performed back- and forwards snowballing on the used papers, in order to find more relevant papers to answer the sub-questions. We performed the same filtering steps as used in the initially retrieved papers in the snowballing process. Figure 4 illustrates how we conducted the SLR and Figure 5 shows the distribution of the papers found over time per year.

3.3.2 Expert Interviews

The SLR showed that only limited research is done on practical cases of zero trust. Therefore we did additional interviews with zero trust experts to complement the findings from the literature and verify whether they are in line with the practical experiences of the experts. In order to answer SQ2, we wanted to know how the experts migrate towards a ZTA. Furthermore, we wanted to know what factors they consider before they decide which concepts

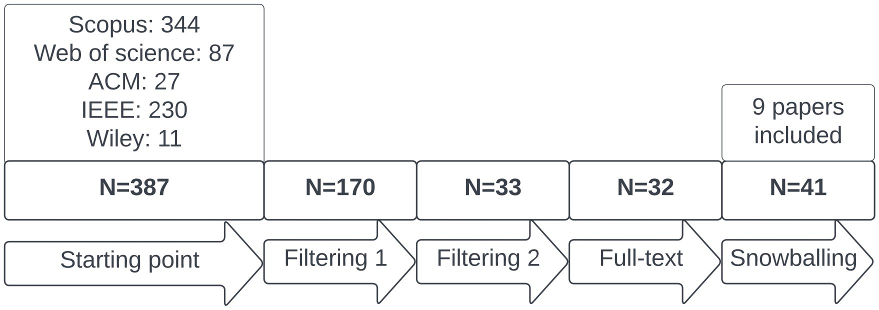

Starting point	Hits per search engine (N without duplicates)
Filtering 1	Inclusion term: title + keywords relate to a ZTA
Filtering 2	Inclusion terms: abstract + conclusion refer to (1) concepts or (2) implementation/migration steps
Full-text	Full-text review to answer SQ1-SQ4
Snowballing	Adding relevant papers by back- and forwards snowballing

Figure 4: SLR approach

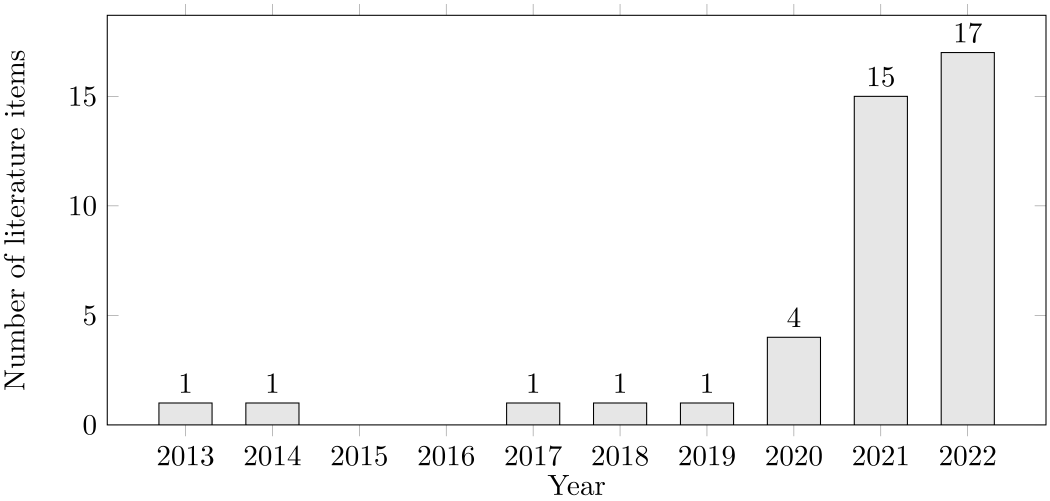

Figure 5: Distribution of the literature found in the SLR over time

they want to integrate (to answer SQ3) and we mapped the concepts that are impacted by the factors to answer SQ4. We used a semi-structured interview approach.

Expert Selection

We constructed the sample of experts based on multiple non-probability sampling methods, because we only wanted to select participants that have experience with a zero trust migration. We used *convenience sampling*, since the partners were selected from SURF to participate in the interviews. We also used *purposive sampling*, because we selected experts from organisations with characteristics that were as varying as possible. We aimed to interview six experts as recommended by Guest et al. [22], who state that most codes (elements for a meta-analysis) are discovered within this number of interviews. In Chapter 5.1 we provide an overview of the experts we selected for the interview sessions.

Consent from Participants

Before we started the interviews with the participants, we sent them an information letter to inform them of their rights before, during and after the interview. The participants had to fill in a form, in which they agreed with these terms. The information letter and consent form that was shared with the participants is provided in Appendix A.

Interview Protocol

The interview protocol consists of three main parts. In the first part, we gathered a general overview of the experiences of the expert in the field of zero trust. We created the user profiles based on this data. In the second part, we explored the main steps that the expert takes during the migration towards a ZTA. We used a shared diagram builder program [1] to create a workflow diagram together with the participant. The shared workspace helped us to have meaningful discussions about the steps that the participant described and increases the likelihood that our interpretation of the answers of the participant is correct, because the participant is able to indicate mistakes in the created workflow. The results from the interviews were used to identify how organisations decide which zero trust concepts they integrate in their ZTA (SQ2). Finally, we asked the experts what factors play a role in choosing in favor of the different concepts identified in SQ1. We asked them to mention the concepts *out of the blue* — without mentioning possible factors that could have an impact on their decision from literature — to make sure that we did not bias their answers. The list of concepts answers SQ3. The list is written down on the vertical axis of a matrix that has the concepts on the horizontal axis. We ask the participant to mark all concepts where the factor is applicable. This data is used to answer SQ4. The full interview protocol, including examples of the workflow and matrix presented to the participant, is shown in Appendix B.

Before we did our interviews with the experts, we did two pilot interviews with a security expert from SURF and a PhD student from Utrecht University. The goal of this was to detect flaws in the interview protocol and increase the quality of the interview guide. Based on the interviews we made some changes to improve the interview protocol. We increased the duration of the interview from maximal 60 minutes to maximal 90 minutes for example, because we found out that we needed more time. Furthermore, we improved the descriptions

[1] Hosted in the platform of `http://www.whimsical.com`.

of the elements that we used to identify the migration steps and the explanation for the factors that the participants should identify.

We used Microsoft Teams as our platform to conduct and record the interviews. The recordings were stored until they were transcribed, for a maximum of six months. We conducted the interviews in Dutch, since this is the native language of both the interviewer and the participants.

Analysis of the results

After we conducted the interviews, we transcribed them and coded them in the Atlas.ti platform. To ensure the reliability of the results, we coded the documents with a second independent coder, who is well-grounded with ZTA's. In Chapter 5.2 we describe in detail how we did the analysis and in the rest of Chapter 5 we present the results we gathered.

3.3.3 Method Engineering

Here we start the second step of the engineering cycle (*treatment design*), by building our initial ZEDEC method. The goal of the method is defined in the objective of the research (see Chapter 3.1). Then we designed it based on the method engineering approach by Weerd and Brinkkemper [64]. They state that in order to create a method, one should first analyze the project situation, as we already did in the problem investigation step of our research. Then, we select relevant method fragments that fit the needs of our own treatment. These fragments consist of (1) method fragments from articles about ZTA concepts, (2) method fragments from articles about a ZTA migration and (3) method fragments from the experiences of experts. Finally, we assemble the fragments by creating a process deliverable diagram (PDD). A PDD models an activity flow and the needed concepts to execute the activities. It allows multiple activity types, for example, concurrent activities (using AND-gates) and conditional activities (using OR-gates). Figure 6 shows the main components of a PDD.

The activities in our PDD are the migration steps discovered in SQ2 (the result of the SLR, snowballing and expert interviews). In addition, we improve the decision-making process by introducing the factors that have an impact on the concepts (SQ4) to the ZEDEC method. This way the method contributes to the main objective of the research (see Chapter 3.1).

3.3.4 Validation Interviews and Application

For the *treatment validation*, we demonstrated how the ZEDEC method could be applied to a fictional scenario, to prove whether the ZEDEC method works in theory. After this, we did an evaluation session with a security expert from SURF, following the guidelines of Wieringa [65] to validate the method using an expert opinion. During the session, we evaluated whether the expert deemed the method useful and if it would fulfill the needs of the organisations. Furthermore, we asked whether the ZEDEC method needed some adaptations to increase its correctness. After the session, we adapted the method where needed.

3.4 Threats to Validity

For the identification of the threats to the validity to the research method, we use guidelines stated by Runeson and Höst [45]. They describe four categories in which the threats to the

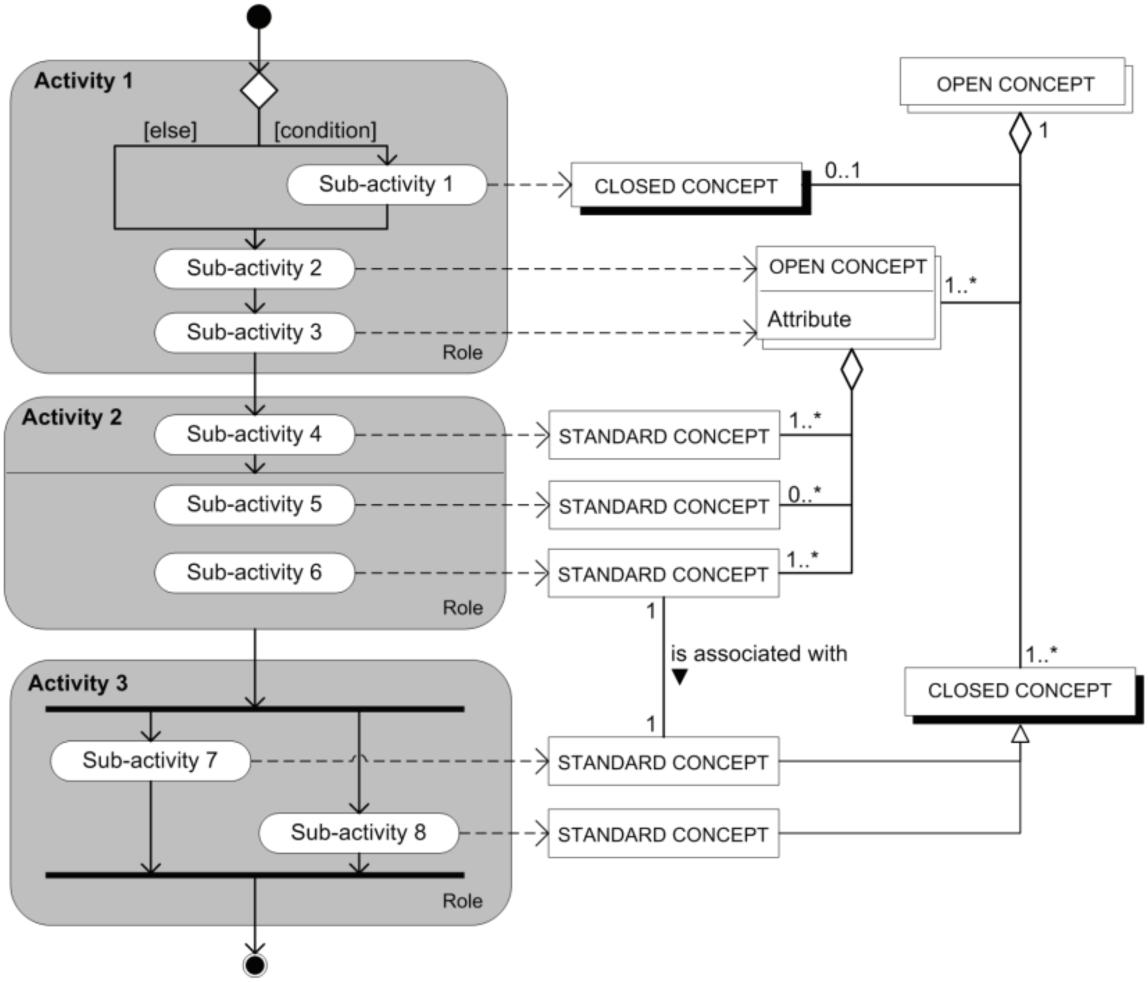

Figure 6: The basic components of a PDD as explained by Weerd and Brinkkemper [64]

validity can be identified. In this section, we describe the threats applicable to our research approach and mention how we mitigate them. Apart from the work of Runeson and Höst, we used the work of Wohlin et al. [66] to find additional threats to the validity of this research.

- The **construct validity** reflects to which extent the measures that the researchers have in mind to study, are indeed studied in the research. A threat in this research could be that participants have a different perception of the meaning of a ZTA than the researchers. Possible reasons for this are the novelty of the topic or the wide variety of concepts that are covered in ZTA's. This threat is minimized by the flexible capabilities of semi-structured interviews, which allows the interviewer to identify possible misinterpretations. If needed, extra explanations on the definitions of a ZTA in this research could be mentioned during the interviews, to make sure that the participants indeed have the same notion of ZTA's.

- An **internal validity** threat occurs when an extra factor might affect the investigated factor. In this case, the role of the participant in the organisation might affect the answers it gives to the questions in the interview. To mitigate this problem, we interview participants with different roles from different organisations, to capture the vision from different points of view. Another threat to the internal validity is that participants might react differently as time passes. They could for example get bored or tired during the interview, since they could take up to 90 minutes in total. We mitigate this threat by including tools like the shared diagram builder program to make the protocol more varied, instead of only asking questions. A final threat to the internal validity is that we use tools like Microsoft Teams and whimsical.com. When bugs or disruptions occur in these tools, the interview results could be impacted by them. Although we do not have full control over this threat, we minimize the risk by making sure that our own internet connection is as stable as possible.

- Threats to the **external validity** concern the extent to which the findings can be generalized and the extent to which they are of interest to people from a different case. A threat in this research is that the participants in this research are partners from SURF, and therefore they might have a certain bias towards the sector that SURF operates in. Partners of SURF are mainly *publicly financed educational institutions in the Netherlands* and relevant *governmental ministries and advisory parties.* For this research, we pick organisations from both sectors and within the sectors we also vary the organisations as much as possible. For example, in the sector of *publicly financed institutions in the Netherlands* we selected both *higher professional education institutions* and *secondary vocational education institutions.*

- Threats to the **reliability** indicate that different findings might be gathered out of the research once it is repeated by another researcher. We aimed to define our approaches and protocols extensively and therefore this threat should be minimal. A threat to the reliability is that we might not discover all relevant codes at the expert interviews, since we do not aim for full saturation. However, because the study of Guest et al. [22] states that most codes are discovered within six interviews, this threat is minimized.

Chapter 4

Structured Literature Review Results

In this chapter we describe our findings from the SLR. We first describe our findings on the identified zero trust concepts that one should consider to integrate into a ZTA (SQ1). Then, we describe our findings on the current decision-making method for zero trust (SQ2), the factors they consider in this process (SQ3) and the mapping from these factors to the concepts (SQ4).

4.1 Zero Trust Concepts

This section describes concepts one should consider to integrate in a ZTA, thereby answering SQ1. We conducted the SLR described in Chapter 3.3 to collect studies about zero trust concepts. In this chapter, we mention all concepts related to a ZTA in these studies. We categorize the concepts based on the main components mentioned in Chapter 2, starting with discovering the *logical components* and thereafter the concepts that exist for the *PDP*. Finally, we cover the *subject, PEP and enterprise resources* in the *data plane* section. We describe these three elements in one section because they are tightly bound to the design choices one makes in the ZTA.

4.1.1 Logical Components

The logical components serve as input for the PDP. Based on the data they gather, the TA ultimately makes a decision on whether a subject gets access to a resource or not. From a comprehensive literature review [56] and a report of the NIST [44], we found the main types of logical components that a ZTA consists of. Other sources confirmed the relevance of those types and provided additional information. We mention the main types below to give an overview of the logical components one can integrate to feed the TA with the required information. In practice, multiple components might be combined in a federated system. An overview of all concepts is shown in Figure 7.

Identity Management

Identity management systems are responsible for creating, storing and managing user accounts for the organisation. The systems also store information that belongs to the accounts, like

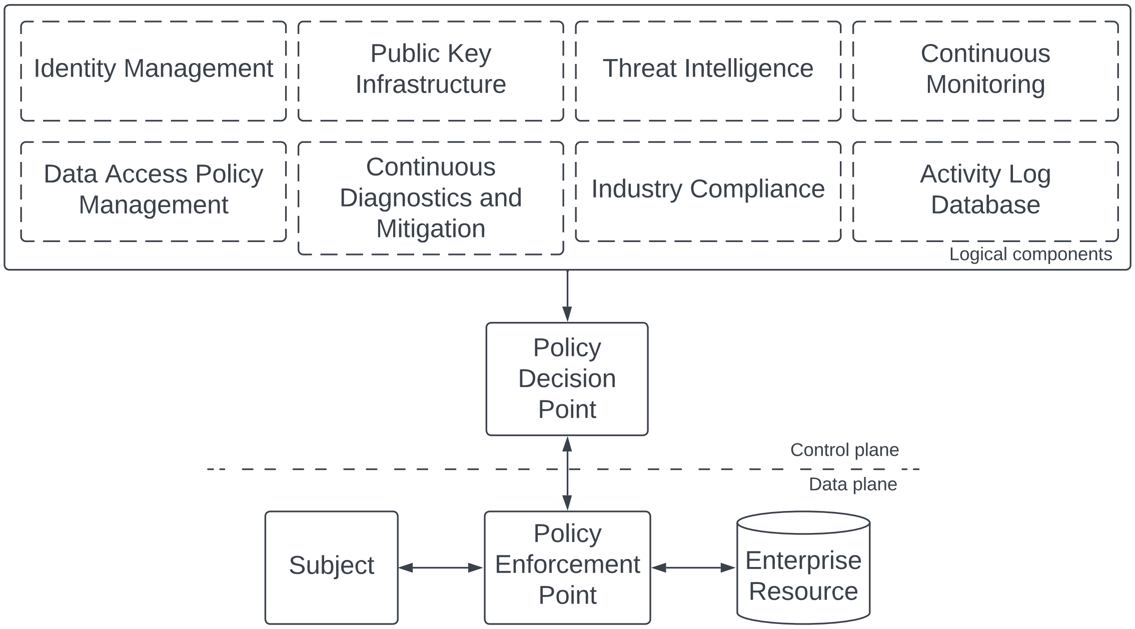

Figure 7: Logical components in a ZTA

subject attributes and enterprise information (e.g., role, access attributes and assigned assets) [44]. An identity management system goes hand in hand with an Access Control Mechanism (ACM), which deploys multiple methods to define and enforce access control policies [56]. The identity management system should fit the structure of the ACM in such a way that the policy rules have access to the required information.

We found many ACM schemes in the papers of the SLR. Many of them lacked the fine-grained functionalities needed for a ZTA (e.g., Discretionary Access Control (DAC) and Mandatory Access Control (MAC) [67]) or encountered other limitations. Therefore we only mention schemes that have proven to be effective in ZTA's. Xiao et al. [67] state that in principle Role Based Access Management (RBAC) or Attribute Based Access Management (ABAC) schemes should be used in a ZTA. Syed et al. [56] state that ABAC and Usage Control (UCON) are the main schemes that fit with the fine-grained and context-aware functionalities that a ZTA requires. We describe the schemes here, including the SADAC scheme that builds upon ABAC:

- **Role Based Access Control (RBAC)** is a scheme that associates subjects with certain roles. A role defines which resources can be accessed by the subjects that are assigned to the role. A subject might be assigned to multiple roles and vice versa. The main advantages of this scheme are the simplicity by which minimum privileges can be set and that it is widely adopted among organisations. Disadvantages of the RBAC scheme are that (1) in complex situations there might be too many roles activated to manage (also called a *role explosion*), (2) it does not support the use of contextual information about the subjects and (3) it is not fine-grained enough to fulfill the needs of the *least privilege* zero trust tenet [25, 56].

- **Attribute Based Access Control (ABAC)** is a policy-based scheme that uses attributes of the subjects as input. The attribute-based nature allows the scheme to apply fine-grained policies that go beyond role access rights. Roles can still be used as attributes for the subjects, meaning that RBAC can be combined with ABAC. Other access schemes might also be combined with ABAC in the same way. The fine-grained access possibilities result in a more complex policy handler [25, 56].
- **Security Attribute-based Dynamic Access Control (SADAC)** is an extension on the ABAC scheme [21]. It introduces three main novel features on top of ABAC, namely: (1) additional attributes regarding configuration and operation, (2) a supervision procedure to evaluate the security profile of a subject over time based on behavior and (3) a procedure that informs the subject how to adapt the security profile to a satisfactory threshold. The design of SADAC is based on zero trust principles.
- **Usage Control (UCON)** is an alternative to ABAC that is also context-aware and allows for fine-grained access [56]. In contrary to ABAC, this model allows the attributes of subjects to be updated during sessions, instead of only at the moment of authentication. This scheme requires an attribute manager on each node/asset, which may cause problems at constrained devices (like IoT devices).

The above-mentioned schemes do not necessarily rule out other access control systems — combinations might be used.

Data Access Policy Management

The policy database contains the attributes, rules and policies about accessing enterprise resources collected in all logical components. They form the starting point for granting access to resources. The policies should be based on the needs and risks of the organisation and should restrict the subject to only having access to the resources it needs [44, 58]. The policies are either manually or automatically managed. Automatic management techniques are preferred because of the high effort of manual management, and are typically based on artificial intelligence and blockchain technologies [19]. The eXtensible Access Control Markup Language (XACML) supports the storage of policies. It defines how policies could be formulated based on the XML language and is proven to be compatible with ABAC-like access schemes [52].

Lee et al. [33] present a policy management framework called FURZE (Fuzzy Risk Framework for ZTA). They facilitate a fuzzy risk evaluation that takes adaptation to dynamically changing contexts into account. The framework uses XACML and a Risk Adaptable Access (RAdAC) access scheme. Ferretti et al. [19] mention policy management frameworks based on blockchain technologies.

Public Key Infrastructure

Since no trusted internal perimeter exists, all traffic is considered vulnerable to attacks. Therefore all communication has to be secured, regardless of network location — ubiquitous encryption is required to guarantee data secrecy [6]. The Public Key Infrastructure (PKI) system is responsible for generating and logging certificates to all subjects, resources, services and applications within an organisation, verifying their trustworthiness. The enterprise PKI might build upon X.509 certificates, but this is not necessary [44].

Data must be encrypted at rest, in transit and during processing. AES (encryption), SHA-256 (hashing) and RSA/Elliptic Curve (signing) are effective techniques on devices with adequate computational capacities [56]. For smaller devices that have less computational power, those techniques might not be supported (e.g., IoT devices). These require lightweight authentication techniques, which are computationally less demanding. Although this comes at the cost of the security being compromised, the NIST states that the techniques are desired in constrained devices [59]. They also mention multiple lightweight authentication techniques based on hashing techniques and Syed et al. [56] mention a list of encryption-based lightweight authentication techniques.

Organisations could consider using *homomorphic encyption* when working with privacy-sensitive data. This technique allows data to be processed while staying encrypted. For example, a homomorphic-encrypted number can be doubled (or modified any other way) without the processing party knowing the initial number nor the outcome during the process. Apart from advantages in privacy-sensitive situations, homomorphic encryption is also resilient to quantum computing which might become relevant in the future [56].

Activity Log Database

The zero trust tenets state that an organisation should collect as much as possible data about the assets and networks. Collecting log data contributes to this tenet. Valuable logs typically provide real-time feedback on the security posture of the network. Examples are: (1) network traffic, (2) resource access actions and (3) database usage logs. An organisation might want to add sensors to collect data from network segments with low network visibility. Data from the activity logs might be used by the TA or processed by continuous monitoring systems [36, 44].

Continuous Monitoring

Continuous monitoring systems typically analyse activity logs and behavioral attributes, and process them to gain insights on the entire network [15, 37]. The system generates warnings when it detects potential security issues [6] and can contribute to policy refinement processes in the deployment phase of the ZTA [56]. It might use threat intelligence feeds to help detect possible weak spots in the architecture. Continuous monitoring techniques are divided into two categories [48]:

- **Intrusion Detection Systems (IDS)** analyse network activities and detect potential attacks on enterprise assets. A Security Information and Event Management (SIEM) system falls in this category [44, 58, 67]. SIEMs automatically generate warnings when potential malicious behavior is detected and can be adapted to various platforms [56].

- **Intrusion Prevention Systems (IPS)** are an addition on IDS's. Apart from detecting intrusions, the systems also react to the alerts and solve them if necessary. Typically, Security Operation Center (SOC) teams consider alerts from IDS systems and mitigate the problems as soon as possible. Security Orchestration, Automation and Response (SOAR) systems could assist SOC teams with the mitigation activities, by orchestrating standardised response procedures and detecting false negative alarms using machine learning [14, 56].

An IDS is essential to keep an overview of the network traffic between the sources. An IPS is an addition to an IDS which has more features, mainly focused on responding to detected problems. In a ZTA, the continuous monitoring concept should cover as much as possible assets of the enterprise.

Continuous Diagnostics and Mitigation

Organisations should collect as much as possible data about the assets and networks. Continuous Diagnostics and Mitigation (CDM) systems help identify assets on the organisation's network and manage them using Hardware Asset Management (HWAM) programs. Asset management programs are essential to building the capacities needed to move towards a ZTA, because these programs help to detect potential malicious assets in the network [2, 58, 60]. Another feature of a CDM system is that it can apply updates on configuration and software components of the organisation's assets when necessary. The system is also responsible for creating policies on non-enterprise-owned devices that want access to enterprise resources, as well as sharing information about device statuses with the TA [44].

Threat Intelligence

Threat intelligence provides an organisation with the most recent information regarding cyber threats and releases instructions on how to solve these threats [3, 44, 48]. This includes discovered malware, issues in software and reported attacks. Multiple sources should be used to get a comprehensive overview from both inside and outside the organisation. Syed et al. [56] mention multiple threat intelligence concepts that collect feeds from different sources and automate responses to the threats.

Industry Compliance

An industry compliance system ensures that the organisation stays compliant with any regulations it might fall under. This could be laws from the government, as well as guidelines and restrictions in specific industries. Some compliance systems also distribute this information to relevant employees. The system prevents the organisation from getting fines or potential loss in revenue [3, 34, 44].

4.1.2 Concepts Within PDP

The PDP is responsible for deciding whether a subject gets access to a requested resource and thereafter instructing the relevant PEP's on the actions they have to perform on the access request. This subsection focuses on the state-of-the-art concepts that exist for the PDP.

Trust Algorithm

Every time a subject makes an access request, the PDP executes the TA to decide on the follow-up actions. Therefore, the accuracy of the decisions are of high importance for an effective deployment of a ZTA. On the one hand, the TA should be strict enough to not allow unauthorized access, but on the other hand, should be loose enough to disallow legitimate users access to resources they need. Multiple variations on the structure of a TA exist, which are divided by the following characteristics [44]:

- **Criteria- versus score-based.** Criteria-based TA's have a list of requirements that have to be met in order to generate an access decision. Depending on the organisation's preferences all criteria have to be met, or a predefined percentage. The score-based variant is less binary than a criteria-based TA and allows for a weighted approach to the policies. More important policies might have a higher impact on the decision than less important ones. When the total score is above a certain threshold, the subject is allowed access [58].

- **Singular versus contextual.** The main difference between these two characteristics is that a contextual TA takes user history and behavior into account, while singular TA's handle requests only per session, without considering the history and behavioral attributes. Although contextual TA's have a higher chance of detecting attacks, singular TA's allow faster evaluations due to their low complexity. Furthermore, contextual algorithms require extensive support from the logical components, which might not always be available [15, 21, 37].

- **Binary- versus tier-based decisions.** The decision that a TA generates could be binary (yes or no), but can also be more gradual. For example, a tier-based algorithm could decide to require a (partial) re-authentication from the subject when not enough confidence is gained or certain criteria are not met. Multiple tiers might exist [63, 67].

Trust algorithms are typically based on machine learning technologies [56, 58]. He et al. [25] compare established evaluation methods for trust algorithms. They state that popular trust evaluation methods are based on *the subjective logic model, information entropy model, weighted average model, Bayesian model, fuzzy theory, game theory* and *machine learning.*

Authentication techniques

Authentication techniques verify the identity of subjects that make access requests. Once verified, more certainty is gained on the legitimacy of the subject's identity. A key technology in a ZTA is Multi-Factor Authentication (MFA), meaning that more than one authentication step needs to be performed before finishing the authentication process. The more steps, the higher the certainty that the subject's identity is legitimate [6, 49, 50, 56]. The accuracy of the TA is dependent on the authentication steps that an organisation requires from their subjects; when the organisation decides to only authenticate based upon a password, the gained trust is not as accurate as when a combination of different techniques is used. Therefore the techniques should be selected carefully. This subsection summarizes the authentication techniques found in the SLR that organisations could consider to integrate. Syed et al. [56] mention the following types (specific authentication techniques are summarized in Table 2):

- **Conventional authentication** is the traditional way of authenticating. Typically this consists of something you *know* (e.g., a password), something you *are* (e.g., biometric data like fingerprint or iris scans) or something you *have* (e.g., a key or token owned by only the subject) [50]. Although these techniques are widely adopted, it has repeatedly been proven that those techniques are vulnerable to subversion.

- **Context aware authentication** takes attributes of the situation of the subject into account (e.g., location and time). An advantage compared to conventional techniques is that the subject does not need to actively participate in these authentication steps.

- **Continuous authentication** is considered a key authentication type in a ZTA [44]. In contrary to the previous types, it does not only provide entry-point authentication (at the beginning of a session), but it authenticates the user throughout the whole session. Continuous authentication techniques are typically based on subject behavior, like patterns in typing or gestures on a touchscreen.
- **Device authentication** focuses on device-to-device authentication, like verifying the identity of IoT devices. Due to the different characteristics of devices compared to humans, different authentication techniques are required. Some techniques from other categories might be applicable to devices as well.

Table 2 shows authentication techniques per authentication type found in the papers from the SLR. The PUFDCA technique by Alshomrani and Li [9] uses a combination of PUF-techniques and continuous authentication techniques to create a device fingerprint.

PDP Design

The centralized nature of a PDP makes the element a potential single point of failure. The result of a PDP failure could lead to resources being unavailable, data leakages and data modifications. To limit the risk of a PDP failure, organisations should consider various security techniques to protect the element from attacks (like DDoS attacks) [44, 56]. In addition, they could distribute the trust over multiple PDPs to secure the availability of the resources when a single PDP gets compromised. In the DistriTrust ZTA architecture as proposed by Sengupta and Lakshminarayanan [49], this is achieved using the notion of threshold signatures. It splits the key into key-shares and distributes them among the available PDP's, to make sure that the full key is not available to one single PDP.

4.1.3 Concepts Within Data Plane

This subsection focuses on the concepts that apply to the subject, PEP and resources. Most concepts relate to architectural approaches that enable the PDP with different capabilities. We point out (1) what deployment models exist and what their capabilities are, (2) what types of (micro)segmentation exist and (3) what resource locations exist, such as cloud solutions.

Segmentation

Organisations should seek to isolate enterprise resources as much as possible from each other, because it prevents lateral movement within a network and allows fine-grained access controls. Ideally, all resources should be allocated to their own segment and are only accessible via their own gateway. Different models on how to deploy segmentation are explained in Chapter 4.1.3, but here we first state what techniques exist to establish segmentation between resources. A combination of these techniques might be used [56]:

- **Native micro segmentation** is based on the hard- and software characteristics of a resource. Infrastructural characteristics might limit access from other systems for example (e.g., there is no connection between two resources) and operating systems might allow a resource to be isolated as well.

Table 2: Authentication techniques

Authentication type	**Approach**	**Description and source**
Conventional	Password/PIN	A string of characters that must be recalled by a subject to gain access [6, 8, 12, 49, 50].
	Biometric data	The subject must provide human characteristics before access, like scanning a fingerprint or iris [49, 50].
	Key/token	An artifact owned only by the subject, providing required data to access a resource [6, 50].
Context aware	Geolocation	An unusual or far off access request location might reduce trust [6, 15, 24, 48, 49, 56].
	Time of access request	An usual date and time for an access request might reduce trust [12, 48, 56, 58].
	Device update status	Devices without the latest updates (e.g., OS, firmware and software) might be less trustworthy [8, 15, 48, 49, 56].
	Screen-lock time	Devices with a long time before it goes into sleep-mode might be less trustworthy [8].
	Subject penalty history	Subjects with suspicious behavior in the past might reduce trust [48].
	Phone placement	The position of a phone (e.g., hand, table or pocket) might influence the established trust [56].
Continuous	Typing behavior	Sudden changes in the way the subject types (e.g., speed) might reduce trust [12, 29, 50, 56].
	Touchscreen behavior	Sudden changes in the gestures the subject makes might on a touchscreen reduce trust [29, 50, 56].
	Gait behavior	Changes in the way that a subject walks (e.g., average speed or steps per kilometer) might reduce trust [29, 56].
	Physiological signals	Changes in physiological signals like EEG (brain signals) and EMG (muscle usage) might reduce trust [29, 56].
	Radio signals	Sudden changes on radio-frequency hardware might reduce trust [56].
	Resource usage history	When a subject wants to access resources that it usually does not want to access, trust might be reduced [12, 24, 48, 58, 67].
	Camera	Changes in the appearance or environment of the subject might reduce trust[29].
Device	PUF-based	PUF (physical unclonable function) is a unique digital fingerprint for IoT device that identifies subjects and increases trust [9, 15, 56].
	Cryptography-based	Cryptographic keys are exchanged between devices to increase trust [56, 57].
	Device fingerprinting	Data about the devices are collected (e.g., battery percentage) and changes might reduce trust [9, 25, 56, 57].

- The **third party model** uses hardware components to establish micro segmentation. Frequently used devices are intelligent switches and next-generation firewalls (NGFW) [44].
- The **overlay model** makes use of software-defined perimeters. A central component on the control is responsible for the generation of VLAN's for every subject, allowing them only access to allowed resources. This model establishes greater visibility than the other solutions [62].

Deployment Models

Here we explain how access to the different segments can be managed on the data plane. We discuss multiple deployment models as stated by Rose et al. [44], which vary in the structure of the PEP and the way how the subject and resource segments are connected to it. An organisation should choose a deployment model that fulfills its specific needs and capabilities. The different segmentation techniques are applicable in these models. Figure 8 shows an abstract visualization of the deployment models.

- **Device agent/gateway-based deployment**. This model requires all subjects to have an agent installed on its organisation-owned asset (e.g., a laptop). The PDP only allows for (secured) communication with an agent, so assets that do not have it installed are unable to access resources. The agents enable advanced tracking possibilities on the subject's device, allowing for fine-grained access policies in the PDP [4]. Once access is granted by the PDP, the agent connects to a gateway that is located in front of an enterprise resource. In this approach, every resource has its own gateway to accomplish maximal segmentation.
- **Enclave-based deployment**. This is a variation on the *device agent/gateway-based deployment* model and allows a gateway to be placed in front of multiple resources (an enclave of resources). This allows for less fine-grained access control, since the subject is more likely to access other resources behind the gateway than intended. This model is typically suitable for legacy systems or on-premise data centers that can not have individual gateways for their resources.
- **Resource portal-based deployment**. In this model, all assets are allowed to communicate with a gateway portal. This means that the asset does not require to be managed by (agents of) the organisation, allowing any device to access resources and restricting the possibilities to collect advanced data about the devices. The gateway portal might be placed in front of individual resources or enclaves of resources.
- **Device application sandboxing**. This model illustrates how to handle sandboxed environments in a ZTA. It allows one asset to run multiple isolated applications at the same time by using virtual machines, containers or other implementations. Ideally, each application is separated by its own sandbox, and each sandbox has its own gateway. Depending on the desires of the organisation, any of the previously mentioned deployment models could be used as PEP for the deployment. Application sandboxing deployment is typically useful when it is not possible to scan the whole asset for vulnerabilities and the organisation still wants to protect applications from the host asset. A disadvantage

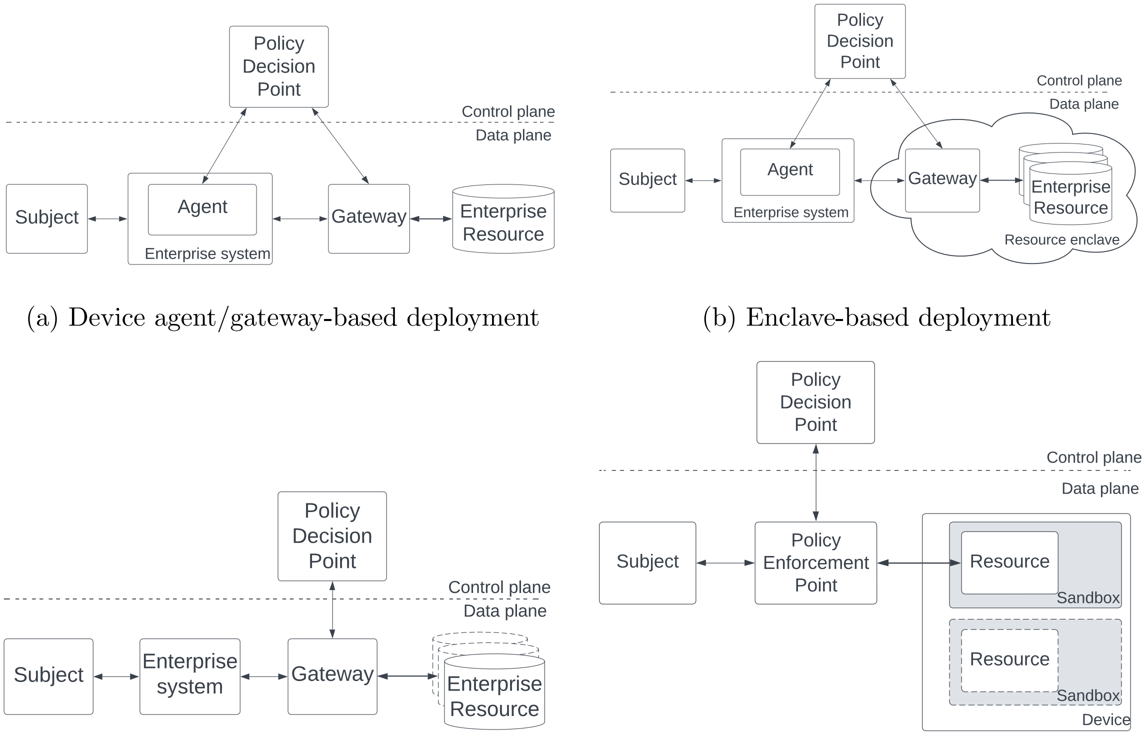

(a) Device agent/gateway-based deployment

(b) Enclave-based deployment

(c) Resource portal-based deployment

(d) Device application sandboxing

Figure 8: Deployment models on the data plane [44]

of this method is that all individual sandboxes need to be maintained, which could be a high-effort task [68].

Another element that might be present on the data plane, but is not yet explicitly mentioned in the deployment models, is a Single Sign-On (SSO) handler. This would allow the subject to stay authenticated on multiple systems at once, without having to re-authenticate every time it wants to access another resource. SSO would significantly increase usability for the subject, while security is retained when implemented correctly. The use of a tier-based TA would contribute to making a balanced trade-off between maximizing usability advantages and retaining security certainties by enforcing (partial) re-authentication [19]. In contrary to the NIST, Google's *BeyondCorp* ZTA model does explicitly mention a SSO component [63].

Resource Location

Enterprise resources on the data plane might be located on-site or in the cloud. Traditionally, servers are situated on-premise, but updating the hardware and software of the servers is a demanding task. Therefore current trends show that organisations are moving towards cloud solutions [13]. Although the maintainability of cloud servers is less demanding, other challenges arise. Sarkar et al. [48] state that most challenges focus on security issues regarding network visibility, segmentation and deployment models. The visibility challenges are similar to ZTA problems and segmentation should be handled as mentioned in Chapter 4.1.3. For the deployment of the cloud infrastructure they distinguish two main types: the *private-* and *public cloud infrastructure*. The private cloud infrastructure is only accessible by a limited amount of users (similar to an agent-based deployment model), while the public cloud infrastructure allows every subject to access the network (similar to an agentless deployment model). Furthermore, a hybrid approach exists between these two. During the migration towards a ZTA, the nature of the cloud environment should fit the structure of the ZTA in these aspects. Ferretti et al. [19] propose a survivable zero trust architecture for cloud computing that follows the zero trust tenets.

4.2 Decision Making Method

In this section, we state our findings from the SLR regarding the decision-making methods that organisations currently use in order to decide which zero trust concepts they want to integrate in their ZTA. Thereby we answer SQ2 from the viewpoint of the literature.

The SLR shows that the migration process towards a ZTA consists of four main steps. The first step is *identify the protection surface*. An organisation should get an as clear as possible picture of all actors, assets and workflows that are relevant to the context where the migration takes place. The second step is *perform a risk assessment and prioritisation*. In this step, the organisation identifies what the risks in the identified context are and what risks need the most priority to be mitigated. The third activity is *identify zero trust solutions*. The organisations selects the solutions that it wants to use to mitigate the risks with the highest priority. The final step is *deploy and review*. Here, the ZTA solutions are deployed and monitored. The final activity is out of scope for this research, because it happens after the decision-making on which zero trust concept the organisation selects and therefore does not contribute to the main research question. These steps are performed as a never-ending cyclic process, because when the migration of a concept is finished, another zero trust concept might be integrated. In the following subsections, we describe in detail how the activities can be carried out. An overview of the steps is illustrated in Figure 9.

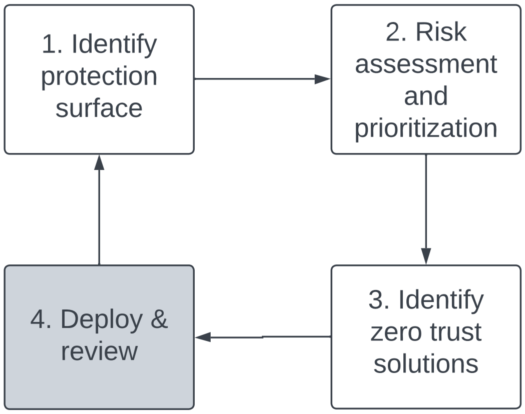
1. Identify protection surface
2. Risk assessment and prioritization
3. Identify zero trust solutions
4. Deploy & review

Grey-colored boxes are out of scope to answer the main research question

Figure 9: Method to migrate towards a ZTA as found in the SLR

4.2.1 Identify Protection Surface

In the first step, an organisation should identify relevant elements in the context where the migration takes place. In the SLR, multiple detection fields were mentioned of which organisation should get an overview. These are:

- **Actors**. The organisation should identify all actors that make use of the system. These actors can either be humans, or non-person entities [2, 12, 17, 44, 58].
- **Enterprise assets**. An organisation should get a clear overview of all assets in the system, and the connections between them. This includes resources like databases, servers, network hardware, etc. Both enterprise-owned and non-enterprise-owned assets should be identified [2, 12, 17, 44, 58].
- **Crown jewels**. Organisations should also get an overview of the **crown jewels**, so that they know which resources require priority for protection measures. High-value assets typically have (1) value to the owner, (2) value to an attacker or (3) steps that lead to other assets [32, 58].
- **Key processes**. An organisation should identify the key processes and transaction flows that belong to it. This helps the organisation to understand how actors use the assets and refine workflows that need improvement [13, 14, 17, 44].

4.2.2 Risk Assessment and Prioritization

The risk assessment and prioritization activity has the purpose to identify risks to critical business processes and to prioritize which of those risks need mitigation first [2, 13, 17, 44, 58]. The ISO/IEC27005 standard on risk management provides a methodology on how to do this [53]. This standard states that after the context has been established, an organisation should first identify the risks that threaten its IT landscape and establish a list of risks. This step is called *risk identification*. Then, a *risk analysis* has to be performed, where the likelihood and the consequences of the risks are assessed. The combination of the likelihood and consequence results in an impact. The higher the likelihood and consequences are, the higher the impact. Organisations are free to determine how the reflection upon the combination of the likelihood and consequences leads to a certain level of impact. Finally, based on the impact of the risks, an organisation should determine whether the risk is acceptable or needs mitigation and a prioritization of the risks needs to be made in the *risk evaluation* step.

We identified two risk assessment frameworks that were used in the context of zero trust, that assist the *risk assessment and prioritization* activity. The first network is the MITRE ATT&CK framework, as stated by Køien [32]. This framework is a curated knowledge base for cyber adversary behaviors for industrial control systems. It provides an overview of known threats and maps these to mitigations that can control the threats [5, 55]. Another identified risk assessment framework is the NIST Risk Management Framework (SP800-37) [27, 44]. This framework prescribes similar steps to assess risks as the ISO/IEC27005, and also provides a list of controls that can be used to mitigate threats in SP 800-53 [28]. Both the MITRE and NIST framework provide a list of controls or mitigations, while only the MITRE ATT&CK framework also provides a mapping to the threats that can be reduced by the mitigations.

4.2.3 Identify Zero Trust Solutions

This is the final step, where the decision making takes place. In terms of a risk assessment, the organisation should here select a zero trust treatment or mitigation that eliminates a risk, or downgrades the risk to an acceptable level. While selecting the solution, the organisation can already start identifying which policies might be applicable on the implemented solutions [13, 17, 44].

4.3 Factors and Connection to Concepts

In the SLR we searched for factors that organisations consider when deciding which zero trust concepts they want to integrate (SQ3). Furthermore, we wanted to find a mapping from the factors that showed which concepts are influenced by which factors (SQ4). In the papers that we selected in this SLR, we did not find information or data about these topics.

4.4 Chapter Takeaways

In the first section of this chapter, we answered SQ1 by describing the main concepts of a ZTA, divided over three categories: (1) the logical components, (2) the PDP and (3) data plane. The first category is a list of logical components that provide data to the PDP. In the second category, multiple concepts are described on how to handle this data. The last category about the data plane covers the flow of how a subject gets access to a resource. Multiple segmentation strategies and deployment model variations exist to enable this process in different ways. An overview of the discovered concepts is shown in Table 3. We use the concepts later on to create the matrix that connects the factors that have an impact on deciding to select a concept for the migration or not.

After that, we described how the decision-making on which zero trust concepts an organisation wants to integrate goes according to the literature (for SQ2). We identified three main steps: (1) identify protection surface, (2) risk assessment and prioritization and (3) identify zero trust solutions. We can use these activities in the eventual method we create in SQ5. We did not find any factors related to SQ3 and SQ4 in the literature.

Table 3: Overview of detected ZTA concepts

Legend: white = logical component; lightgrey = concept within PDP; grey = concept in data plane

Concept	Description
Identity management	Identity management systems are responsible for creating, storing and managing user accounts. The access scheme should support the storage of data that other elements in the ZTA need.
Data Access Policy Management	The policy database contains the attributes, rules and policies about accessing enterprise resources collected in all logical components. The policies should be formulated as strictly as possible to allow each user least privilege access rights.
Public key infrastructure	The Public Key Infrastructure (PKI) system is responsible for generating and logging certificates to all subjects, resources, services and applications within an organisation. It is responsible for encrypting all data at rest and in transit.
Activity logs	Logs in which relevant data about the network, subjects and resources are stored for processing purposes.
Continuous monitoring	Continuous monitoring systems typically analyse activity logs and behavioral attributes, and process them to gain insights on the entire network. The system generates warnings when it detects potential security issues and is typically divided in *intrustion detection systems (IDS)* and *intrusion prevention systems (IPS)*.
Continuous diagnostics and mitigation	Continuous Diagnostics and Mitigation (CDM) systems help identify all assets on the organisation's network and manage them using Hardware Asset Management (HWAM) programs.
Threat intelligence	Threat intelligence provides an organisation with the most recent information regarding cyber threats and releases instructions on how to solve these threats.
Industry compliance	An industry compliance system ensures that the organisation stays compliant to any regulations it might fall under. Examples are laws and industry restrictions.
Trust algorithm	Every time a subject makes an access request, the PDP executes the TA to decide whether access should be granted or not. The TA should be configured strict enough to prevent access to adversaries, and meanwhile loose enough to allow legitimate users to get access to the resources they need.
Authentication Techniques	Authentication techniques verify the identity of subjects that make access requests. The more authentication techniques used, the higher the certainty that a subject's identity is legitimate.
PDP design	The PDP needs to be designed properly to mitigate the risks of becoming a single point of failure.
Segmentation	Organisations should seek to isolate enterprise resources as much as possible from each other. We discussed multiple types of segmentation.
Deployment models	The deployment model enables data traffic between subject and resource. The deployment model must fit the needs of the organisation.
Resource location	On-premise and cloud solutions are widely used. On-premise servers are more flexible but hard to maintain and cloud solutions have challenges with network visibility, segmentation and deployment.

Chapter 5

Interview Results

In this chapter, we describe the experts that we selected for the interviews and explain how we analysed them. Furthermore, we illustrate the findings, in which we aim to find out how organisations decide which zero trust concepts to migrate in a ZTA (SQ2), what factors they consider about ZTA concepts before they make the decision (SQ3) and to map these factors to the concepts on which they are relevant (SQ4). The interviews are conducted as described in chapter 3.3.

5.1 Expert Selection

We selected the zero trust experts by using convenience sampling and purposive sampling, as described in Chapter 3.3.2. We selected six participants that have experience with zero trust and were willing to participate in the research. We anonymised their identities and created a general profile about the participants. We identified their identities with a letter (A-F) and categorized the organisation types. The used categories are: *governmental ministries and advisory parties* (M&A) and publicly financed educational institutions from SURF's community. The latter category is divided into *higher professional education institutions* (HPEI) and *secondary vocational education institutions* (SVEI). We identified their working experience in their current role, their experience with ZTA's and their experience in a security-related job in general. The profiles of the participants are shown in Table 4. On average, the participants have 3.5 years of experience with zero trust. All participants are male and all work at different organisations except for participants D and E. Although their role is the same, they focus on different areas of zero trust. Participant E focuses more on the technical solutions of the ZTA, while participant D focuses more on governance.

5.2 Interview Analysis

After we conducted the interviews, we first transcribed them and then coded them in the ATLAS.ti platform to detect categories related to migration steps (SQ2), factors (SQ3) and the relation between factors and concepts (SQ4). We also coded the diagrams obtained during the interviews in addition to the transcripts. To ensure reliability of the coding results, we checked the coding with a second independent coder, who is well-grounded with ZTA's. We did this in multiple iterations. In the first iteration, we coded one transcript and created a

Table 4: Participants' demographics and background

Exp. = experience (in years); Org. = organisation; SVEI = secondary vocational education institution; HPEI=higher professional education institution; M&A = governmental ministry or advisory party.

Participant	Org. type	Role	Role exp.	ZTA exp.	Security exp.
A	SVEI	Security officer	2-5	2-5	10-20
B	HPEI	Information security officer	0-2	2-5	10-20
C	M&A	Enterprise architect	2-5	2-5	10-20
D	HPEI	Solution architect	5-10	2-5	5-10
E	HPEI	Solution architect	5-10	2-5	10-20
F	M&A	Enterprise architect	2-5	2-5	2-5

codebook, which we sent to the second coder together with the transcript file (without preset quotations). He coded the transcript file with the codebook and afterwards we compared the agreement between the coded files manually. We had a discussion about the differences between our codes and discovered that most differences between our coding were caused by (1) misinterpretation of the codes or (2) because we selected different sentences for our quotations. Therefore, we improved the descriptions of the codes and made clearer arrangements on how what sentences we selected for our quotations.

In our second iteration, we used two different transcripts and repeated the process from the first iteration. When comparing the coded documents from both parties, we again identified problems with (1) misinterpretations from the new codes and (2) the selection of different sentences for the same codes. For the final iteration, we made the new codes clearer and decided to provide the second coder with a transcript document that has pre-selected quotations. The task of the second coder was then to assign the codes from the codebook to the pre-defined quotes. Consequently, the second problem would thereby be eliminated.

In the final iteration, we included all remaining transcripts and all workflow models. A calculation of the Holsti index resulted in an intercoder agreement of 57.4% [26]. To increase the reliability of the coding process, we discussed the remaining differences in our coding. We improved the description of the codes in case of misinterpretations and came changed the assigned codes in such a way that we both agreed.

5.3 Findings on the Current Decision-Making Method

In the first part of the interview, we asked the participants what steps they performed in order to decide which zero trust concepts they integrated. We coded the interviews and detected four main activities that the participants carry out in order to make this decision. Furthermore, we detected what the inputs and outputs are, as well as the actors that perform the activities. We illustrate the activities and inputs in Figure 10.

5.3.1 Activities and Inputs

We describe the activities and inputs as shown in Figure 10 as follows:

- **Create vision on needed changes (5/6)**. Five participants mentioned that a vision is needed on the direction in which the migration phase should head. They initiated this as the first step of the process: *"Well, for us [the first activity] is [...] to define*

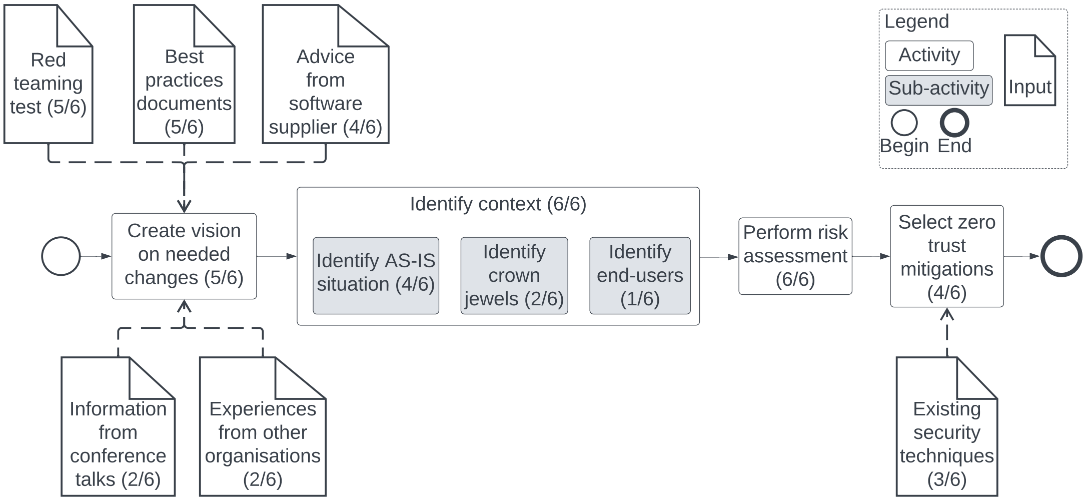
Red teaming test (5/6)
Best practices documents (5/6)
Advice from software supplier (4/6)
Legend
Activity
Sub-activity
Input
Begin
End
Create vision on needed changes (5/6)
Identify context (6/6)
Identify AS-IS situation (4/6)
Identify crown jewels (2/6)
Identify end-users (1/6)
Perform risk assessment (6/6)
Select zero trust mitigations (4/6)
Information from conference talks (2/6)
Experiences from other organisations (2/6)
Existing security techniques (3/6)

Figure 10: Zero trust decision making process as identified in the interviews

some more policy frameworks on that, right? [...] So the first question is often [...] what do I have to comply with?" [C]. The instantiation of this vision differentiates across the participants. For example, *participant D* calls the vision a *"policy document" [D]*, while *participant B* calls it a *"vision strategy" [B]*. Other participants did not mention how they call the document, or point out that the vision is never formally documented.

We identified multiple elements that contribute to the establishment of the vision. The first input is a *red teaming test (5/6)*, where the report about the shortcomings in the current IT landscape indicates what changes need to be made: *"Well, what has also been a very big driving force for us here is that we deliberately chose to have a red teaming action done by a third party" [E].* Another input is *best practice documents (5/6)*. Participants mentioned that they follow instructions from documents from advisory parties like the NCSC or NIST to create a vision on the needed changes. Likewise, they use *advice from software suppliers (4/6)*, *experiences from other organisations (2/6)* and *information from conference talks (2/6)* as an input.

- **Identify context (6/6)**. All participants mentioned that a context of the to be improved system area needs to be identified, or that they used documentation on the context that already existed. We identified three sub-activities in the interview that contribute to this activity. Four participants mentioned that they *identify the AS-IS situation (4/6)*, referring to a process in which they identified the currently available systems and other resources. Two other participants said that they *identify the crown jewels (2/6)* so that they can prioritize improvements regarding these system areas. The final sub-activity is that organisations want to *identify the end users (1/6).*

  > *"You have to know what the crown jewels of the organization are. Well that analysis, that's where it starts" [F].*

- **Perform risk assessment (5/6)**. Five participants state that they perform a risk assessment in the process. In order to do so, *participant C* mentions that he tries to comply with the ISO norms. Furthermore, he states that he uses the MITRE ATT&CK framework to identify the risks in his IT landscape: *"Yes we are actually looking at the MITRE ATT&CK framework. So basically a kind of the market standard that looks at: what are the risks on the cloud? What are the risks on Microsoft 365? What are the risks on so many of those things?" [C].*

- **Select zero trust mitigations (4/6)**. Four participants stated that they selected mitigations to migrate into their ZTA. Although the other two participants mentioned that they implemented zero trust solutions in a later stage, they did not explicitly mention that they selected mitigations and were therefore not included in the *select mitigation* code. An identified input for selecting the mitigations are the *existing security techniques (3/6)*: *"Look, we like to use things that are already there" [F].*

  > *"But from the architecture flow, we are indeed looking from: What risks do we identify, and what control measures? What control measures should you take on that in terms of architecture?" [C].*

5.3.2 Actors

In the interviews, we identified a high variety of actors that play a role in the discovered activities. Table 5 shows which actors were identified and how many participants mentioned them. Furthermore, it shows at which activities the actors were involved. Some participants associated one actor with multiple activities, which is why some actors are linked to more activities than the total amount of times they were mentioned by distinct participants. It also occurred that participants assigned multiple actors to an activity, which could explain the high variety of actors per activity. Another reason could be that organisations use different names for similar roles.

5.3.3 Outputs

In the interviews, only one type of output was mentioned. Four participants pointed out that they wrote or updated *documentation (4/6)* about the activities. For example documentation about their vision strategy, policy document or the identified context.

> *"When you create policy, it produces a document, of course." [C]*

5.4 Findings on the Factors

In the second part of the interview, we asked the participants what factors they considered about each concept when they were deciding to integrate it. We allowed them to mention factors that had both a positive and a negative effect on the decision. When we analysed the results from the interviews, we noticed that most mentioned factors were similar to risk mitigations. For example, participant F mentioned that he considers *"privileged access management" [F]* as a factor, or participant C considers *"multi-factor authentication" [C]* as a factor. This corresponds with the findings from the decision-making process, because the participants indicated that the last activity of the decision-making process is to *select mitigations*. Therefore mitigations in the context of zero trust are useful factors to consider in this process.

Table 5: Mentioned actors in the interviews

The total column represents the amount of distinct participants that mentioned the role

		Times linked to activity			
Role	Total	Create vision	Identify context	Risk analysis	Select mitigations
Administrator	(2/6)		1		2
Advisor	(3/6)	1	2	1	2
Architecture team	(1/6)			1	
Chief information officer	(2/6)	1	1		2
Chief information security officer	(1/6)	1			
ICT manager	(1/6)				1
Infrastructure team	(1/6)	1			1
Management	(1/6)		1		
Privacy team	(1/6)	1			
Security analyst	(1/6)				1
Security architect	(3/6)	1	1	2	1
Specialist from external party	(1/6)			1	
Stakeholders	(1/6)		1		
Supplier	(1/6)			1	
System owner	(1/6)	1	1	1	1

> *"[...] customized access, so being able to give the right authorizations to the right people at the right time" [A].*

We used the MITRE ATT&CK Framework as our starting point to identify the mitigations. We analysed all enterprise mitigations available in the framework (as stated by MITRE [40]), and made a selection of all mitigations that match with the properties of the identified zero trust concepts from the SLR (SQ1). A subset of 32 mitigations were identified. We added this subset to our codebook and detected which mitigations were mentioned by the participants. We searched through the entire interview to find these mitigations, not limiting ourselves to the final part of the interview where we asked what factors they considered. We did this because the participants also mentioned mitigations they used when asking about milestones, the decision-making process or at other stages of the interview. This way we were able to capture all mentioned mitigations, including the ones that the participants did not think of at the final part of the interview. The results of the coding analysis is shown in Table 6. The results show that 22 out of the 32 mitigations were mentioned by at least one of the participants.

We chose to use the mitigations from the MITRE ATT&CK Framework because in the first place it is a widely adopted framework in the field of risk management, which increases the likelihood that organisations are able to adopt our approach in their business processes. Second of all, the terminology that MITRE uses for the mitigations is widely adopted in the context of cybersecurity, which helps organisations to understand the meanings of the mitigations correctly. Finally, the MITRE ATT&CK Framework includes a mapping from known risks in the field of cybersecurity towards relevant mitigations that reduce the risk. Considering that the second last step mentioned in the decision-making process is *perform risk analysis*, the mapping from MITRE allows organisations to identify which mitigations they can use to mitigate the risks that they identified in this step [55].

Table 6: Mitigations related to zero trust as identified in the interviews

#	**Mitigation**	**Description according to the MITRE ATT&CK framework [40]**	**Occurence in interviews**
1	Limit Access to Resource Over Network	Prevent access to file shares, remote access to systems, unnecessary services. Mechanisms to limit access may include use of network concentrators, RDP gateways, etc.	6
2	Network Segmentation	Architect sections of the network to isolate critical systems, functions, or resources. Use physical and logical segmentation to prevent access to potentially sensitive systems and information. Use a DMZ to contain any internet-facing services that should not be exposed from the internal network. Configure separate virtual private cloud (VPC) instances to isolate critical cloud systems.	6
3	Update Software	Perform regular software updates to mitigate exploitation risk.	6
4	Active Directory Configuration	Configure Active Directory to prevent use of certain techniques; use SID Filtering, etc.	5
5	Network Intrusion Prevention	Use intrusion detection signatures to block traffic at network boundaries.	5
6	Restrict File and Directory Permissions	Restrict access by setting directory and file permissions that are not specific to users or privileged accounts.	5
7	Account Use Policies	Configure features related to account use like login attempt lockouts, specific login times, etc.	4
8	Filter Network Traffic	Use network appliances to filter ingress or egress traffic and perform protocol-based filtering. Configure software on endpoints to filter network traffic.	4
9	Multi-factor Authentication	Use two or more pieces of evidence to authenticate to a system; such as username and password in addition to a token from a physical smart card or token generator.	4
10	Password Policies	Set and enforce secure password policies for accounts.	4
11	Privileged Account Management	Manage the creation, modification, use, and permissions associated to privileged accounts, including SYSTEM and root.	4
12	Remote Data Storage	Use remote security log and sensitive file storage where access can be controlled better to prevent exposure of intrusion detection log data or sensitive information.	4
13	User Account Management	Manage the creation, modification, use, and permissions associated to user accounts	4
14	Encrypt Sensitive Information	Protect sensitive information with strong encryption.	3
15	Application Isolation and Sandboxing	Restrict execution of code to a virtual environment on or in transit to an endpoint system.	2
16	Credential Access Protection	Use capabilities to prevent successful credential access by adversaries; including blocking forms of credential dumping.	2
17	Threat Intelligence Program	A threat intelligence program helps an organization generate their own threat intelligence information and track trends to inform defensive priorities to mitigate risk.	2
18	Vulnerability Scanning	Vulnerability scanning is used to find potentially exploitable software vulnerabilities to remediate them.	2
19	Antivirus/ Antimalware	Use signatures or heuristics to detect malicious software.	1
20	Disable or Remove Feature or Program	Remove or deny access to unnecessary and potentially vulnerable software to prevent abuse by adversaries.	1
21	Execution Prevention	Block execution of code on a system through application control, and/or script blocking.	1
22	Exploit Protection	Use capabilities to detect and block conditions that may lead to or be indicative of a software exploit occurring	1

Table 6: Mitigations related to zero trust as identified in the interviews

#	Mitigation	Description according to the MITRE ATT&CK framework [40]	Occurence in interviews
23	Behavior Prevention on Endpoint	Use capabilities to prevent suspicious behavior patterns from occurring on endpoint systems. This could include suspicious process, file, API call, etc. behavior.	0
24	Code Signing	Enforce binary and application integrity with digital signature verification to prevent untrusted code from executing.	0
25	Data Loss Prevention	Use a data loss prevention (DLP) strategy to categorize sensitive data, identify data formats indicative of personal identifiable information (PII), and restrict exfiltration of sensitive data.	0
26	Environment Variable Permissions	Prevent modification of environment variables by unauthorized users and groups.	0
27	Limit Software Installation	Block users or groups from installing unapproved software	0
28	Restrict Library Loading	Prevent abuse of library loading mechanisms in the operating system and software to load untrusted code by configuring appropriate library loading mechanisms and investigating potential vulnerable software.	0
29	Restrict Registry Permissions	Restrict the ability to modify certain hives or keys in the Windows Registry.	0
30	Software Configuration	Implement configuration changes to software (other than the operating system) to mitigate security risks associated to how the software operates.	0
31	SSL/TLS Inspection	Break and inspect SSL/TLS sessions to look at encrypted web traffic for adversary activity.	0
32	User Account Control	Configure Windows User Account Control to mitigate risk of adversaries obtaining elevated process access.	0

5.5 Findings on the Factors Mapped to Concepts

In the interview, we asked the participants what factors they considered during the decision-making. This resulted in the factors as identified in Table 6. After that, we asked the participants to make a mapping for all factors to the zero trust concepts they are applicable on. They could choose all concepts we identified in the SLR, as stated in Table 3. Together with each participant, we created a matrix as shown in Figure 17 in Appendix B, where we map the factors to the relevant zero trust concepts. In this section, we analyse the transcripts and the results from the matrices created at the interview sessions to answer SQ4.

We checked all mitigations that the participants mentioned and added an "X" to all concepts that the mitigations were mapped to. We present the results of this analysis in Figure 11, where each "X" indicates that at least one participant mentioned mapped the mitigation to the concept. In addition to the obtained mapping, we added selected clusters of mitigations with similar relevant concepts. They give an insight into potential areas of interest for organisations to work on and help us classify the mitigations that were not mentioned by the participants (zero occurrences in the interviews). We identified the following classifications: blue — identity and access management; and yellow — network monitoring and device management.

The results show that all concepts are mapped to at least one mitigation, except for *Industry compliance* and *PDP design.* The main reason for the *industry compliance* was that organisations were not familiar with the software: *"Yes, yes, we really have very little experience with that" [E].*

The policy decision point, on the other hand, was familiar to the participants. Multiple participants mentioned that they wanted the PDP to be designed as centralized as possible, but did not mention a mitigation in this context. Therefore, this concept is not mapped as well: *"Yeah yeah, It's anyway we're trying to do that centrally.* [...] *Ultimately it has to land with us in our own in, in our current infra, where that fits. Yeah so I'm not inclined to say that we require really something specifically different" [E].*

We did not include the mitigations in the matrix that were mentioned by zero participants, because no participants included these factors in the matrix during the interview. Since these mitigations were not mentioned by the participants, we were not able to detect the relevant concepts based on the interviews. In chapter 5.5 we describe what concepts these mitigations should be connected to based on literature and our own expertise that we build throughout the research, and motivate why this would be the case. There we add these to the matrix and highlight concept clusters.

5.6 Chapter Takeaways

For SQ2 we identified as main steps: (1) create vision on needed changes, (2) identify context, (3) perform risk analysis and (4) select mitigations. Relevant inputs to establish a vision on the needed changes are *red teaming tests*, *best practice documents*, *advice from software suppliers*, *information from conference talks* and *experiences from other organisations. Currently existing techniques* are an input to select mitigations. A high variety of actors were identified and the main outputs for the activities are documents.

In SQ3 we identified factors that participants consider for the decision-making of which zero trust concepts they want to integrate into their ZTA. We found that the factors were similar to mitigations, and therefore we identified relevant zero trust mitigations from the MITRE ATT&CK framework. We identified a subset of 32 mitigations that are relevant in the context of zero trust, from which 22 are mentioned at least by one of the participants. Table 6 shows all mitigations and their occurrence in the interviews.

In SQ4 we made a mapping that shows which zero trust concepts are linked to the mitigations (see Figure 11). We made two clusters; (1) identity and access management and (2) network monitoring and device management. The fragments identified in this chapter can be used to construct the final decision-making method.

Figure 11: Mitigations (from SQ3) mapped to concepts (from SQ1), based on interview results

Blue cluster : identity & access management; yellow cluster : network monitoring & device management

Mitigation	Identity management	Data access policy management	Trust Algorithm	Authentication techniques	Public key infrastructure	Activity logs	Continuous monitoring	Continuous diagnostics & mitigation	Threat Intelligence	Industry compliance	Policy decision point design	Segmentation	Deployment models	Resource location
Account use policies	X	X	X	X		X								
Active directory configuration	X	X	X	X										
Limit access to resources over network	X	X	X	X										
Multi-factor authentication	X	X												
Password policies	X	X												
Privileged account management	X	X		X										
Restrict file and directory permissions	X	X	X	X										
User account management	X													
Credential access protection					X									
Encrypt sensitive information					X									
Antivirus/antimalware								X						
Disable or remove feature or program								X						
Execution prevention								X						
Exploit protection							X	X						
Filter network traffic			X			X	X	X	X				X	
Network intrusion prevention							X	X						
Update software								X						
Vulnerability scanning							X	X						
Threat intelligence program									X					
Application isolation and sandboxing												X		
Network segmentation												X		
Remote data storage													X	X

Chapter 6

ZEDEC Method

In this chapter, we construct the ZEro trust DECision (ZEDEC) method, based on the findings in the SLR and the interviews. The main goal of the method is to guide organisations in deciding which zero trust concepts to include in their ZTA to be optimally resilient against cyber threats, as stated in Chapter 3.1. We first finalize the decision matrix in which the factors are mapped to the zero trust concepts. Although most mitigations were already mapped when analysing the interview results, some need to be added based on logical mappings based on literature, supported by our expertise gained during the research. Then we assemble the identified method fragments obtained from the literature and the interviews — including the decision-making matrix — and model them in a PDD.

6.1 Decision Matrix

In Chapter 5.4 we identified the mitigations that were relevant to zero trust in the MITRE ATT&CK framework. We created a subset of 32 mitigations from the framework that had shared properties with the zero trust concepts that we identified in the SLR. In Chapter 5.5 we identified a mapping from the mitigations to the zero trust concepts for all 22 mitigations that were mentioned by the participants in the interview. Still, there are 10 mitigations that were not mentioned by the participants, but have a logical connection to the zero trust concepts based on literature. Therefore, these mitigations also need to be considered in the ZEDEC method. In this section we add them to the already identified matrix as illustrated in Figure 11, to finalize the decision matrix (see the updated matrix presented in Figure 12).

In order to identify the mappings from the non-mentioned mitigations, we used the literature found in the SLR to identify to which zero trust concept a mitigation should be mapped. In Table 7 we describe for each mitigation the relation to the relevant zero trust concepts and explain why this is the case. The newly added mitigations were also included in the clusters we made, because they are in line with the clusters we already identified. To the “identity and access management” cluster (blue), we added the mitigations “environment variable permissions,” “restrict library loading,” “restrict registry permissions” and “user account control.” We also added “behavior prevention on endpoint,” “data loss prevention,” “limit software installation,” “software configuration” and “SSL/TLS inspection” to the “network monitoring and device management” cluster (yellow).

Table 7: Mapping for mitigations that were added to the decision matrix based on literature

Mitigation	**Related concepts (bold text) and reason**
Behavior Prevention on Endpoint	This mitigation refers to endpoint management activities, which are applied in the concept of **continuous diagnostics and mitigation** [2, 58, 60]. Other mitigations related to device management were also mapped to CDM systems by the participants in the interviews.
Code Signing	Code signing uses digital signatures techniques, which are part of the **public key infrastructure** concept [56].
Data Loss Prevention	Data loss protection protects data from adversaries, by detecting breaches on both the endpoint and the network [35]. **Continuous monitoring** covers the network protection part [15, 37], while a **CDM system** covers the endpoint protection part[2, 58, 60].
Environment Variable Permissions	This mitigation is about permission management for certain user(groups)s. The user(group)s are typically managed by **identity management** components [44], while the permission policies are stored in the **data access policy management** component [58].
Limit Software Installation	This mitigation refers to device management activities, which are applied in the concept of **continuous diagnostics and mitigation** [2, 58, 60]. Other mitigations related to device management were also mapped to CDM systems by the participants in the interviews.
Restrict Library Loading	This mitigation is about permission management for certain user(groups)s. The user(group)s are typically managed by **identity management** components [44], while the permission policies are stored in the **data access policy management** component [58].
Restrict Registry Permissions	This mitigation is about permission management for certain user(groups)s. The user(group)s are typically managed by **identity management** components [44], while the permission policies are stored in the **data access policy management** component [58].
Software Configuration	This mitigation refers to device management activities, which are applied in the concept of **continuous diagnostics and mitigation** [2, 58, 60]. Other mitigations related to device management were also mapped to CDM systems by the participants in the interviews.
SSL/TLS Inspection	Packet inspection over the network is typically connected to intrusion detection systems and intrusion prevention systems [48]. In the SLR we found that these elements belong to the **continuous monitoring** concept.
User Account Control	User account control is linked to controlling user accounts, which is managed in the **identity management** component [44].

Figure 12: Decision matrix to select relevant zero trust concepts

X: mapping originates from interviews; O: mapping originates from literature; Blue cluster : identity & access management; yellow cluster : network monitoring & device management.

Mitigation	Identity management	Data access policy management	Trust Algorithm	Authentication techniques	Public key infrastructure	Activity logs	Continuous monitoring	Cont. diagnostics & mitigation	Threat Intelligence	Industry compliance	Policy decision point design	Segmentation	Deployment models	Resource location
Account use policies	X	X	X	X		X								
Active directory configuration	X	X	X	X										
Environment variable permissions	O	O												
Limit access to resources over network	X	X	X	X										
Multi-factor authentication	X	X												
Password policies	X	X												
Privileged account management	X	X		X										
Restrict file and directory permissions	X	X	X	X										
Restrict library loading	O	O												
Restrict registry permissions	O	O												
User account control	O													
User account management	X													
Code signing					O									
Credential access protection					X									
Encrypt sensitive information					X									
Antivirus/antimalware								X						
Behavior prevention on endpoint								O						
Data loss prevention							O	O						
Disable or remove feature or program								X						
Execution prevention								X						
Exploit protection							X	X						
Filter network traffic			X			X	X	X	X				X	
Limit software installation								O						
Network intrusion prevention							X	X						
Software configuration								O						
SSL/TLS inspection							O							
Update software								X						
Vulnerability scanning							X	X						
Threat intelligence program									X					
Application isolation and sandboxing												X		
Network segmentation												X		
Remote data storage													X	X

6.2 Method Construction

In this section, we construct the ZEDEC method, by assembling the method fragments that we collected in this research. They originate from the findings in the SLR, interviews and the decision matrix that we created in Chapter 6.1. Based on these fragments, we identified four main activities, which we discuss separately in this section. The main activities are: (1) create a vision on needed changes, (2) identify context, (3) perform risk assessment and (4) identify zero trust concepts. The total ZEDEC method is illustrated in Figure 13. We also included a description of the activities (see Table 8) and a description of the concepts (see Table 9).

6.2.1 Create Vision on Needed Changes

The first activity of the ZEDEC method is "create a vision on needed changes". This step was mentioned by five out of six participants in the interview as starting point of the migration process. In this activity, an organisation starts to realise that it needs to add zero trust capabilities to the current IT landscape and they start to identify in what areas of the network these changes are required. The vision is a starting point that provides the organisation with a general guideline for the direction that the migration process is heading to. This includes properties like a scope on the context in which the migration takes place and might also include objectives that the organisation wants to achieve (apart from the overall objective to integrate zero trust components into the IT landscape). The way that the vision is created and instantiated can differ per company, therefore the activity is modeled as a blackbox activity. Instantiations of the activities that were mentioned during the interviews were a *policy documents*, *vision strategy* or they mentioned that the vision is a *general idea that is not formally established or documented*.

In the interviews, the participants also mentioned several inputs that helped establish the vision. These are: (1) red teaming tests, (2) best practice documents, (3) advice from software suppliers, (4) information from conference talks and (5) experiences from other organisations.

6.2.2 Identify Context

Once the vision is created the security team of the organisation should start identifying the context. The vision might include a scope on the area in which the context needs to be identified, but this does not necessarily need to be the case — context identification can also be done across the entire IT landscape. The context identification activities include fragments that we discovered in the interviews and the SLR. In the SLR we discovered that *enterprise assets* (including the *data flows* between them), *crown jewels*, *key processes* and *actors* need to be discovered in this step. In the interviews, we found that in this phase the *crown jewels* and *end-users* need to be discovered. In ZEDEC we consider identifying the *end-users* part of the activity "identify actors," because the end-users are a subset of the actors. Furthermore, the participants mentioned that *AS-IS models* were discovered, referring to the current systems and resources in the network. We modeled the elements that need to be identified as sub-activities to the *identify context* activity in an unordered way, since it does not matter which element is identified first.

The literature does not provide a clear subset of actors that should perform this activity of the ZEDEC method and the interviews also showed a high variety of potential actors.

Therefore, the broad term *security team* should encompass all actors that might contribute to this activity. Possible actors according to the interviews in this activity are: administrators, advisors, chief information officers, management, security architects, stakeholders and system owners.

6.2.3 Risk Assessment

This step is mentioned in both the literature and the interviews. Based on the SLR results, in ZEDEC we follow the risk management steps of the ISO/IEC27005 [53] and ISO31000 standard [54]. This framework states that the main activities in risk assessment, after context identification, are that organisations should perform a (1) risk identification, (2) risk analysis and (3) risk evaluation. The ISO norm provides a general overview of how these steps can be performed and leaves room for frameworks to be integrated into the method. Because our decision matrix is based on mitigations from the MITRE ATT&CK framework, we incorporated steps from the "assessment and engineering" process, as stated by Applebaum [10]. This method provides guidance on how to select relevant tactics and techniques that are a threat to the IT landscape and provides a mapping to the mitigations that address these threats. In ZEDEC we already incorporate the MITRE ATT&CK framework at the "risk identification" phase, so that the selection of mitigations in the decision matrix can be done based on the mapping that MITRE provides. In this subsection we explain the activities that the organisations have to perform to perform a risk assessment by using the MITRE ATT&CK framework.

Risk Identification

In this step, the goal is to find, recognize and describe risks. A risk is typically built up from a combination of threats, vulnerabilities and assets [53]. To reach the goal of identifying the risk, organisations need to establish what threats exist and what vulnerabilities are present on the assets. The organisation already identified the assets in the *context identification* step, so we need to identify the threats and vulnerabilities to the current assets in order to establish the risks that are present.

Following the assessment and engineering method, the security team should identify the **threats** by using the MITRE ATT&CK framework [10]. In the first step, they identify the tactics that are applicable within the scope of the assessment. Tactics represent the highest level of abstraction in the MITRE ATT&CK model and are the tactical goals that an adversary may has during an operation [55]. All tactics present a list of techniques that describe actions that adversaries take to achieve their tactical goals. These techniques are formulated as threats. In the second step, the security team selects threats that are applicable to the scope of the risk assessment. Irrelevant threats might be techniques on assets that are not present in the system or techniques to which mitigations are already in place. In the third step, the security team should identify **vulnerabilities** in the current assets to the threats. When a vulnerability is detected, it should be noted together with the asset that it is applicable to.

The organisation now has an overview of the threats, vulnerabilities and assets. The combination of these elements leads to one or more risks. An example of a risk is as follows: given the threat "drive-by compromise" and the vulnerability "insecure browser plugins" on the asset of "laptops from users with high privilege access to resources", the following risk could be identified: “An adversary uses a drive-by compromise technique on insecure browser plugins at laptops of employees with high privilege access.”

Risk Analysis

In this step, the security team assesses the risk level of each risk. The higher the risk level, the more priority it might get to be mitigated. They establish the risk level based on the likelihood that the risk occurs and the consequence of the risk in case it would occur. Once the likelihood and consequence levels are established, the level of risk can be determined. A pre-defined table or formula should prescribe how the level of risk can be extracted from the values of the likelihood and consequence [53, 54].

Risk Evaluation

In the risk evaluation phase, the security team decides whether the risk requires follow-up steps or not. Examples of follow-up steps that the security team could choose are: *accept risk*, *require further investigation* or *mitigate risk*. In the end, an ordered list is provided of the risks that have the highest priority to be mitigated. The risk level assessment in the “risk analysis” step helps identify this order [53, 54].

6.2.4 Identify Zero Trust Concepts

After the risk assessment, ISO/IEC27005 states that a risk treatment needs to be selected [53]. Furthermore, we found literature in the SLR that states that this should happen [13, 17, 44]. Also in the interviews, the participants mentioned that in this stage they select zero trust mitigations. Therefore the step “identify zero trust concepts” is included in ZEDEC. Depending on the available (human) resources, the organisation might want to select one or more risks to focus on. To select relevant mitigations, they use the MITRE ATT&CK framework. For the risk that the organisation wants to mitigate, it should select the threat that the risk is linked to in the MITRE ATT&CK framework. For each threat, the framework shows a list of mitigations, out of which the identified zero trust mitigations in the decision matrix are a part. The security team can now select a zero trust mitigation that is not yet (fully) present in the current IT landscape and helps mitigate the risk.

Once the security team selected the mitigation, they look into the decision-making matrix (as depicted in Figure 12) to identify which zero trust concepts need to be integrated, in order to meet the zero trust principles. It might occur that a zero trust concept is not integrated at all. For example, when the mitigation *filter network traffic* is identified, but no continuous monitoring component is yet in place. In that case, this concept needs to be added to the current network. It might also occur that a continuous monitoring component is already there, but is not yet monitoring the area in which the risk is identified. In that case, the component needs to be enhanced.

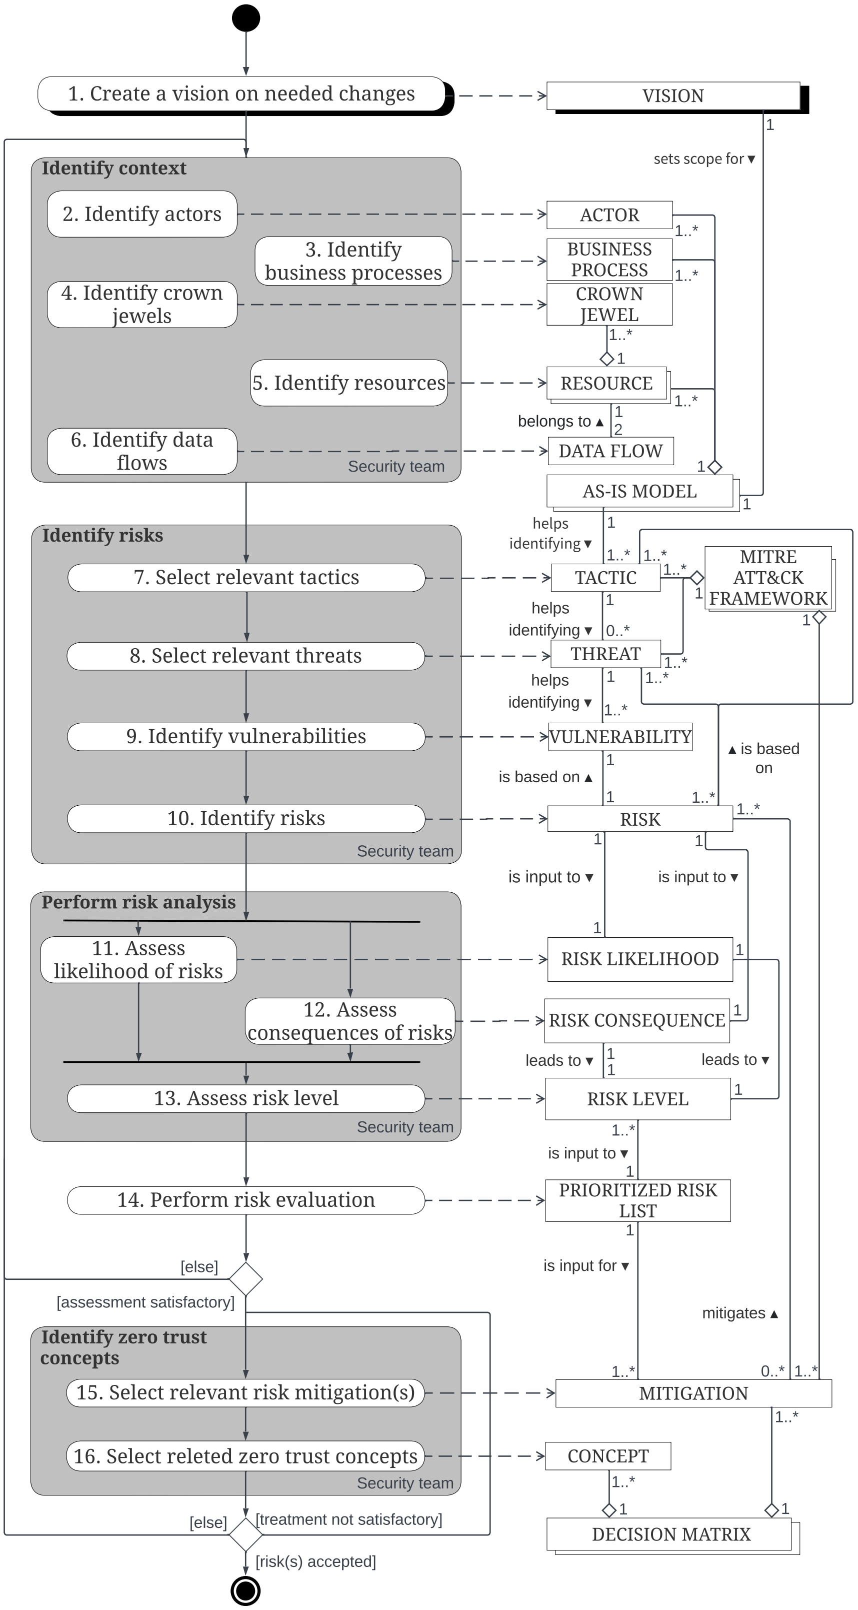

Figure 13: The draft ZEro trust DECision making (ZEDEC) method

Table 8: Activities in ZEDEC method

#	Activity	Description
1	Create a vision on needed changes	An organisation comes up with an idea where improvements in the IT landscape are required. The created vision forms a direction and scope on the area where the migration takes place. This VISION determines the scope in which the context is identified in the next step.
2	Identify actors	The security team gathers a complete overview of all ACTORS that have access to the relevant area of the IT landscape.
3	Identify business processes	The security team gathers a complete overview of all BUSINESS PROCESSES that are carried out in the relevant area of the IT landscape.
4	Identify crown jewels	The security team gathers a complete overview of all CROWN JEWELS that are present in the relevant area of the IT landscape.
5	Identify resources	The security team gathers a complete overview of all RESOURCES that are present in the relevant area of the IT landscape.
6	Identify data flows	The security team gathers a complete overview of all DATA FLOWS between the RESOURCES in the IT landscape.
7	Select relevant tactics	The security team identifies all relevant TACTICs from the MITRE ATT&CK FRAMEWORK that are applicable on the VISION.
8	Select relevant threats	The MITRE ATT&CK FRAMEWORK contains a list of relevant THREATs to each TACTIC. In this activity, the security team identifies all THREATS that are relevant given the identified context.
9	Identify vulnerabilities	For each THREAT, identify VULNERABILITIES that are present in the current system. Knowing what VULNERABILITIES exist in your current assets is essential to get a complete overview of the RISKs in the network.
10	Identify risks	Based on the identified THREATs and VULNERABILITIES, identify which RISKs are present in the network. The RISKs should also state which assets are at risk.
11	Assess likelihood	For all RISKs, assess the likelihood of them to happen.
12	Assess consequences	For all RISKs, assess how extensive the consequences would be in case it occurs.
13	Assess risk level	The total impact of each RISK should be determined based on the combination of the LIKELIHOOD and CONSEQUENCE.
14	Perform risk evaluation	Based on the IMPACT, the security team determines what the follow-up actions are on the RISKs; this could be to accept a risk, to require further investigation on the risk or to mitigate a risk.
15	Select relevant risk mitigations	For all RISKs that need mitigation, select relevant MITIGATIONS. During this process, an organisation should consider the current security measures that already exist. MITIGATIONS that already exist in the network do not need to be selected.
16	Select related zero trust concepts	Out of the decision matrix, the organisation should select which zero trust CONCEPTS are mapped to the MITIGATIONS. Now they know which zero trust CONCEPTS they need to migrate into their architecture.

Table 9: Concepts in ZEDEC method

Concept	**Description**
VISION	The VISION is a starting point that provides the organisation with a general guideline for the direction that the migration process is heading to. It helps set the scope for the context in which the migration takes place and might also include objectives that the organisation wants to achieve.
AS-IS MODEL	This model describes all ACTORS, BUSINESS PROCESSES, RESOURCES, CROWN JEWELS and DATA FLOWS that are relevant to the scope of the method.
ACTOR	ACTORS are actors that make in any way use of the IT system on which the ZEDEC method is applied.
BUSINESS PROCESS	BUSINESS PROCESSES illustrate how the system is used. This helps to assess the importance of the different assets in the IT landscape.
RESOURCE	RESOURCES are the present devices, applications and devices in the IT landscape.
CROWN JEWEL	CROWN JEWELS are the RESOURCES that are the most valuable to the organisation.
DATA FLOW	When data between two resources is transmitted, this is done via a DATA FLOW.
MITRE ATT&CK FRAMEWORK	MITRE ATT&CK is a globally-accessible knowledge base of adversary TACTICS and techniques/THREATS based on real-world observations. Furthermore, the framework provides a list of MITIGATIONS to each THREAT [39].
TACTIC	Tactics represent the highest level of abstraction in the MITRE ATT&CK model and are the tactical goals that an advisory may has during an operation [10].
THREAT	All TACTICS contain a list of techniques that describe actions that adversaries take to achieve they tactical goals. In the ZEDEC method, we use these techniques to identify the THREATS [10].
VULNERABILITY	A VULNERABILITY is a weakness or opportunity in the IT landscape that adversaries can exploit to gain access to a system [7].
RISK	A RISK is a combination of THREATS, VULNERABILITIES and RESOURCES. One can establish the LIKELIHOOD and CONSEQUENCE of a RISK [53].
RISK LIKELIHOOD	The RISK LIKELIHOOD is an assessment that considers the likelihood that a risk occurs and the likelihood that adverse impacts are achieved. The higher the likelihood, the higher the RISK LEVEL of a RISK might be [53].
RISK CONSEQUENCE	The RISK CONSEQUENCE is the assessment of how big the effect of the risk occurring would be on the organisation. The higher the consequence, the higher the RISK LEVEL might be [53].
RISK LEVEL	The RISK LEVEL is the overall impact of the risk. It considers both the LIKELIHOOD and CONSEQUENCE of the RISK. The higher the RISK LEVEL, the more priority the RISK might have in the PRIORITIZED RISK LIST [53].
PRIORITIZED RISK LIST	This list includes all RISKS and orders them from most important RISKS to be mitigated, to least important RISKS to be mitigated. For the most important RISKS, mitigations are selected in the *select risk mitigations* phase [54].
DECISION MATRIX	The DECISION MATRIX maps the possible zero trust MITIGATIONS to the zero trust CONCEPTS that they are relevant to. The matrix is shown in Figure 12.
MITIGATION	A zero trust MITIGATION represents security concepts and classes of technologies that can be used to prevent a RISK from occurring, in the context of zero trust [40]. MITIGATIONS are part of the DECISION MATRIX.
CONCEPT	A CONCEPT is a zero trust concept that we discovered in the SLR. The concepts are part of the DECISION MATRIX.

6.3 Chapter Takeaways

In this chapter, we first finished the *decision matrix*. This matrix maps the zero trust mitigations to relevant zero trust concepts (see Table 6.1). Based on this matrix and other fragments discovered in the SLR and interviews, we built the ZEro trust DECision-making method (ZEDEC). This method consists of four main activities: (1) create a vision on needed changes, (2) identify context, (3) perform risk assessment and (4) identify zero trust concepts. The eventual resulting method is shown in Figure 13. In the following chapter, we evaluate this method to establish to which extent it works.

Chapter 7

Method Evaluation

In this chapter, we evaluate the ZEDEC method to answer SQ6. We do an evaluation session with a security expert from SURF and improve the method based on the findings. After that, we demonstrate how the ZEDEC method could be applied in a fictional scenario, to demonstrate that the method works in theory.

7.1 Evaluation Session

One way to evaluate the ZEDEC method, is to use expert opinions to validate its correctness and usefulness [65]. Therefore, we did an evaluation session with a security expert from SURF for the method assessment. He has been working in different security-related roles for more than 20 years and has expertise in zero trust and risk management. The session lasted an hour in which we showed the ZEDEC method and asked him to give feedback on its various components. The main goal was to identify whether the current method is correct and useful to organisations that want to migrate towards a ZTA. Later in this section, we explain in detail what the objectives of the evaluation are and what the impression of the expert is regarding them.

The evaluation session consisted of four phases. First, we showed the ZEDEC method (Figure 13) and the decision matrix (Figure 12) together and briefly explained how they worked. This way we allowed the expert to get familiar with the main components of the method and to talk about his general impression on these elements. In the second phase, we dived deeper into the ZEDEC method and we discussed every step in detail. We illustrated the activity flow of the ZEDEC method by explaining an example scenario of how this method should work (see Chapter 7.2). This way the expert was able to give in-depth feedback about the activities and to mention strengths and weaknesses of them. In the third phase, we dived deeper into the decision-making matrix and again had an in-depth discussion on the weaknesses and strengths. Finally, we asked the expert his opinion on the usefulness of the method and whether he thinks that the partner organisations from SURF would use this. Furthermore, we allowed him to give additional comments that we did not talk about in the other part of the session.

The main objectives of the session were to assess the correctness and usefulness of the ZEDEC method. For the *correctness* our main objective was to capture whether the activities, concepts and activity flows of the PDD were correct. We also dedicated a moment to evaluate the correctness of the decision matrix. The first three phases of the session contributed to achieving these objectives. For the *usefulness* we wanted to establish whether the ZEDEC method is useful to organisations that want to migrate towards a ZTA and we also wanted to figure this out for the decision matrix. Finally, we asked the expert whether he found that the method achieves its main objective; guiding organisations in deciding which zero trust concepts to include in their ZTA, in order to be optimally resilient against cyber threats. The usefulness was discussed in the fourth phase of the session.

7.1.1 Evaluation Session Results

In this section, we describe the findings from the evaluation session. We first describe our findings on the **correctness** of the ZEDEC method. The expert recognized the risk assessment from the ISO steps and agreed that the activities regarding *perform risk analysis* and *perform risk evaluation* are indeed in line with his experiences in how an organisation performs a risk assessment. Furthermore, he agreed that some kind of scope had to be created before the risk assessment process, as is the case in the *create a vision on needed changes* activity. He is familiar with the existence of the MITRE ATT&CK framework, but has not used it in practice. However, once we showed him how we applied the framework in the *identify risk* activity, he confirmed that this was a correct way to identify risks. When focusing on the *identify zero trust solution* activity, he mentioned that he understood how a mitigation follows up on the risk and how the mitigation maps to a zero trust concept in the decision matrix. He did however mention that it was not completely clear to him how he should implement the identified zero trust concept in a way that it contributes to the main goal of migrating towards a ZTA. He gave us an example that the zero trust concept *identity management* was not necessarily connected to zero trust, but also to other security solutions. After a discussion, we came to the conclusion that the initial ZEDEC method does not imply that the zero trust concepts need to be in line with the zero trust rules as identified in Chapter 2 and the identified capabilities of the concepts in the SLR (see chapter 4). Therefore, we agreed that an extra activity needs to be added to the ZEDEC method that explicitly states that the organisation must consider how the zero trust concept complies with the zero trust rules and capabilities.

When looking at the activity flows, the expert mentioned that they were logical and correct. He made one remark; the vision on the needed changes could change at any time and the current data flows do not support this option. Therefore, he mentioned that the arrow back to the *identify context* activity should be placed before *create a vision on needed changes*. He did not have any remarks on the correctness of the concepts on the deliverable-side of the PDD. We also evaluated the correctness of the decision matrix. The expert found the mitigations relevant in the context of zero trust and did not have any remarks on the relationships between the mitigations to the zero trust concepts.

We also discussed the **usefulness** of the method. He mentioned that the method contributes to the main goal to decide which zero trust concepts an organisation should integrate. He also pointed out that the ISO standard is widely adopted and that this would help organisations to easily adopt this method. Finally, he stated that the decision matrix is a good addition that helps organisation identify zero trust concepts specifically.

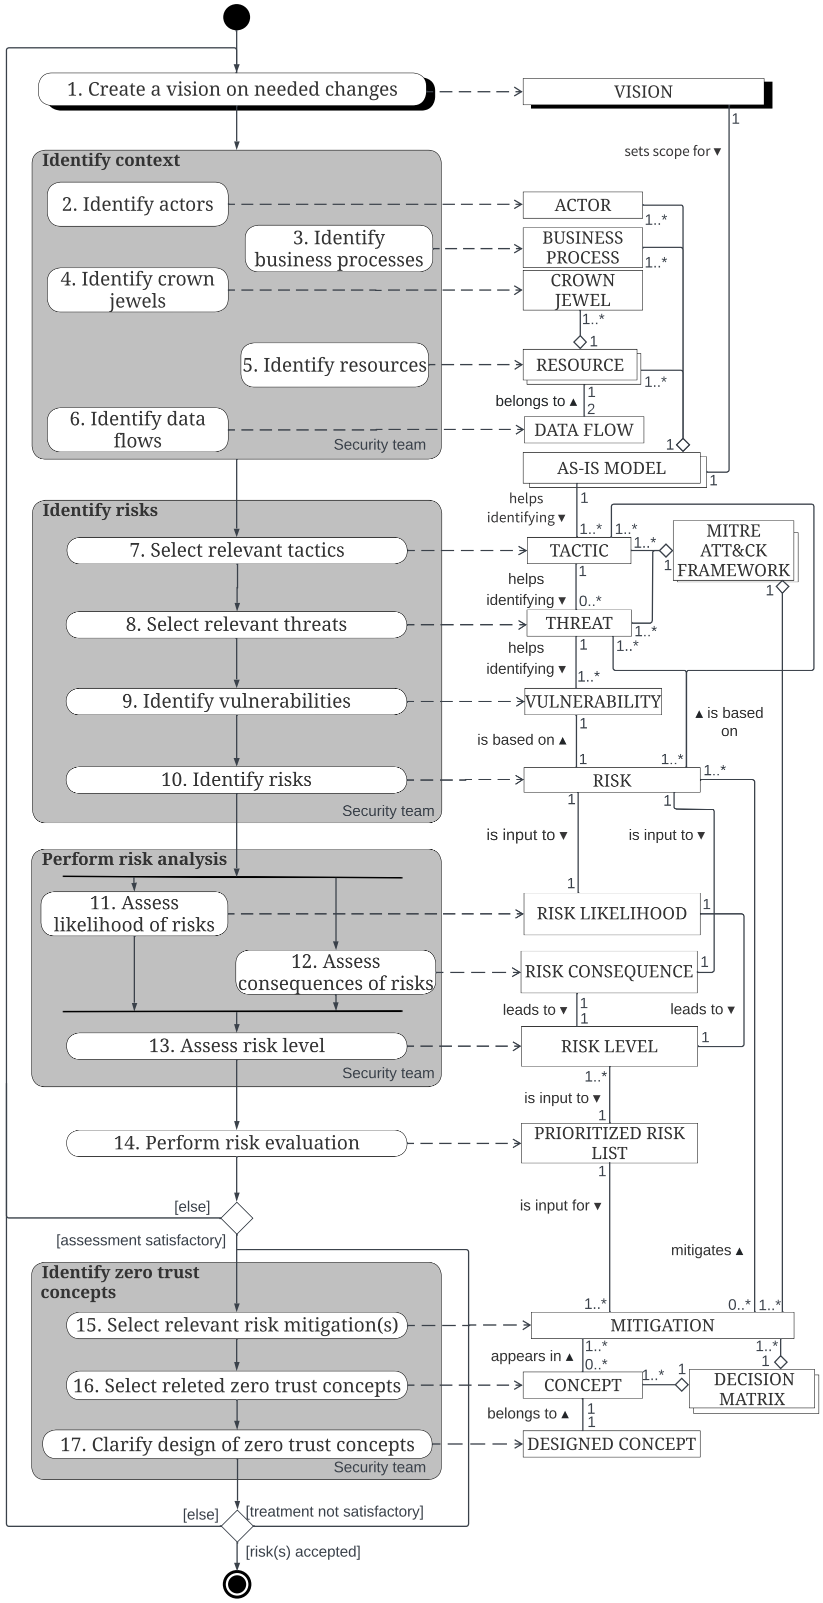

Figure 14: The refined ZEro trust DECision making (ZEDEC) method

7.1.2 Improvements to ZEDEC Method

Based on the results of the evaluation session, we made some changes to the initial ZEDEC method. First, we added the activity *Clarify design of zero trust concept*. Here, the organisation establishes how the zero trust concept is designed and it should make sure that the design follows the zero trust rules and identified capabilities from the SLR. Another change to the initial ZEDEC method is that we moved the data flow that initially pointed to *identify context*, back to *create a vision on needed changes*. This improvement is in line with the methods that we currently found in the SLR, because similar data flows point back to the beginning of the method [53, 44]. Therefore, this data flow could point to the *identify context* activity and all other activities before it, as is the case for create vision on needed changes. Furthermore, no experts mentioned contradicting data flows in the first expert interviews. We show the updated version of the method in Figure 14.

7.2 The ZEDEC Method Application

In this section, we illustrate how the method should be used, based on a fictional scenario. The scenario is about a higher professional educational institution (HPEI) in the Netherlands. This institution started in 2014 and has about 8000 students enrolled. The application starts with the first step of the method: *create vision on needed changes*. The chief information security officer (CISO) of the institution attended a conference about zero trust and is convinced that a zero trust architecture is an ideal solution to improve the security of the current IT landscape. Furthermore, he read the report of the NCSC that states that one of the four main national security risks in the Netherlands is *unauthorized access to information* [1]. He realizes that the institution is also vulnerable to this risk and decides to discuss with the rest of the security team the idea to integrate some elements of a ZTA into the current perimeter-based security architecture. The team decides to use the ZEDEC method to identify which zero trust concepts are relevant to add to the current security architecture. Since the institution is pretty small, they decide to identify the context for the entire system and aim to prevent unauthorized access to information in the system with information about the students. For this migration, data protection for other actors than students is of less importance.

Now that the vision is established, the security starts to *identify the context*. First, they identify the end-users in the system, which are: (1) teachers and professors, (2) students, (3) the management, (4) the security team and (5) other employees at the institution. Furthermore, they identify the resources of the current system. They create a map of all databases, infrastructure hardware, cloud providers, applications and the data flows between the elements. Crown jewels that the security team identified are: (1) the identity management databases, (2) the grade processing system OSIRIS and (3) the online exam software Remindo, in which students make their exams. Relevant business processes that the security team identifies are: (1) filling in grades to OSIRIS, (2) the processing of enrollments for study programs and (3) the de-enrollment of graduated students.

With the identified context for the migration, the security team starts to identify the risks in the current architecture to achieve the objective of their vision, namely preventing adversaries to get access to (sensitive) student data. They first identify the most relevant tactics from the MITRE ATT&CK framework [40]. They decide to focus on the (1) initial access, (2) credential access, (3) privilege escalation and (4) collection tactics, because they are closely related to their objective. From the tactics, they identify threats that are applicable to the current IT landscape. Furthermore, they identify the vulnerabilities related to the threats that are present in the current architecture. The vulnerabilities are inspired on the procedural examples that MITRE provides for all tactics. Table 10 gives an overview of the identified threats and vulnerabilities for the *initial access* tactic. For illustrative purposes, we focus only on the initial access tactic.

Table 10: Identified threats and vulnerabilities in the *initial access* tactic in the MITRE ATT&CK framework [40]

Threat	**Vulnerability**
1. Drive-by compromise	Outdated web browsers at laptops of employees.
1. Drive-by compromise	Insecure browser plugins at laptops of employees.
1. Drive-by compromise	Malicious advertisements (malvertisements) opened on employee devices.
2. Exploit public facing application	Remote File Inclusion to osiris server.
3. External remote services	Lack of vendor security controls at eduVPN.
3. External remote services	Misconfigured permissions and privileges in active directory manager.
4. Phishing	Malicious attachments and links per email to users with access to sensitive data or applications.
5. Replication through removable media	Lack of endpoint monitoring at devices with access to sensitive data or applications.
5. Replication through removable media	Plug-in unknown media on devices with access to sensitive data or applications.
6. Supply chain compromise	Lack of supplier vetting at coding environments like Spider and Visual Studio.
7. Valid accounts	Weak credentials for users with high privilege access.
7. Valid accounts	Credential reuse by users with high privilege access.

The ISO standards state that typical elements that should be considered while identifying risks are *threats*, *vulnerabilities* and *related assets* [53, 54]. Therefore, in ZEDEC risks are composed of a combination of these elements. The security team identifies the following risks, based on the discovered combinations of threats, vulnerabilities and assets in Table 10:

- R1.1: An advisory uses a “drive-by compromise” technique on outdated web browsers at laptops of employees with high privilege access.
- R1.2: An advisory uses a “drive-by compromise” technique on insecure browser plugins at laptops of employees with high privilege access.
- R1.3: An advisory uses a “drive-by compromise” technique by gaining access via malicious advertisements on laptops of employees with high privilege access.
- R2: An advisory uses the “exploit the public facing application” technique by including remote files into the OSIRIS application.
- R3.1: An advisory leverages “external remote services” by exploiting a lack of security controls at eduVPN.

- R3.2: An adversary leverages "external remote services" by misconfigured permissions in the active directory manager.
- R4: An advisory uses the "phishing" technique to send malicious attachments and links to users with high privilege.
- R5.1: An advisory uses the "replication through removable media" technique by adding malicious files to removable media devices like USB sticks, to target users with high privilege.
- R5.2: An adversary uses the "replication through removable media" technique by plugging in removable media with malicious hardware into devices from users with high privilege access.
- R6: An adversary uses the "supply chain compromise" technique by utilizing vulnerabilities at the Spyder coding environment or the Visual Studio coding environment.
- R7.1: An adversary uses the "valid account" technique by trying weak passwords at accounts with privileged access.
- R7.2: An adversary uses the "valid account" technique by trying previously leaked passwords from accounts with privileged access.

Now that the security team identified the risks, they can assess them. To do so, they first establish the likelihood and consequence of the risk. Based on these two inputs, the risk level is determined. The security team considers the likelihood and consequence of equal importance, and therefore the risk level is determined by the middle value between the likelihood and consequence. For the risk evaluation, the security team establishes whether a risk is accepted or if they want to implement a mitigation. To guide the decision-making in this step, they create a table that states for each combination of likelihood and consequence what the follow-up step is. This table is shown in Table 11. The results of the risk assessment are shown in Table 12. The risk assessment and outcome are based on the procedures described by Fikri et al. [20], which demonstrates a case study on a risk assessment using ISO27005.

Table 11: Risk acceptance criteria

	Consequence				
Likelihood	Very Low	Low	Medium	High	Very High
Very Low	Accept	Accept	Accept	Accept	Mitigate
Low	Accept	Accept	Accept	Mitigate	Mitigate
Medium	Accept	Accept	Mitigate	Mitigate	Mitigate
High	Accept	Mitigate	Mitigate	Mitigate	Mitigate
Very High	Accept	Mitigate	Mitigate	Mitigate	Mitigate

Table 12: Risk assessment for R1-R7

Threat	Risk	Likelihood	Consequence	Risk level
Drive-by compromise	R1.1	Low	Medium	Low
	R1.2	Medium	Medium	Medium
	R1.3	Very Low	Medium	Low
Exploit public facing applications	R2	Medium	Medium	Medium
External remote services	R3.1	Very Low	High	Low
	R3.2	Low	Medium	Low-Medium
Phishing	R4	Very High	Medium	High
Replication through removable media	R5.1	Low	Medium	Low
	R5.2	Very Low	Medium	Low
Supply chain compromise	R6	Low	Low	Low
Valid Accounts	R7.1	High	Medium	Medium-High
	R7.2	High	Medium	Medium-High

The outcome of the evaluation is that R4 has the highest priority in mitigation, since the risk level of this mitigation is the highest. Then come R7.1 and R7.2. The selected risks with the lowest priority are R1.2 and R2.

For the final activity, *identify zero trust concepts*, the security team uses the MITRE ATT&CK framework to identify mitigations to the risks [40]. For each threat that corresponds with the selected risk, the security team looks at the list of mitigations the framework provides. They select mitigations that are present in the decision matrix (see Figure 12). They identified four mitigations that help mitigate all four risks: (1) Exploit protection, (2) Update software, (3) Privileged account management and (4) Network intrusion protection. The connected zero trust concepts to these elements are: *identity management*, *data access policy management*, *authentication techniques*, *continuous monitoring* and *continuous diagnostics and mitigation.* These are the concepts that the institution should include in its zero trust architecture in order to be resilient against these cyber threats. As a final step, the security team clarifies the design for these concepts as follows:

- **Data access policy management:** in the current situation the institution saves the policies on multiple places in a static way. They decide to integrate a policy manager that centrally saves the policies for all users, as is the case in a ZTA (see chapter 4). This way, it is the easiest to manage the policies and change them in such a way that accounts have the least needed privilege.

- **Identity management:** the organisation should include an identity manager that is able to provide all needed attributes to the policy manager. The institution currently already has a sophisticated RBAC system on which does not need extra integration activities. Some changes might be needed to connect the user (groups) to facilitate the policies in the data access policy management component.

- **Authentication techniques:** the policies that allow least privileged access management should also include appropriate authentication techniques. Resources that are of high value to the organisation should require authentication from the more sophisticated techniques. The institution decides to integrate two extra context aware authentication techniques based on *geolocation* and *time of access request*, based on the techniques mentioned in Table 2.

- **Continuous monitoring:** The institution integrates a Security Operational Center (SOC) system, which is an intrusion prevention system that keeps track of data traffic and identifies malicious patterns. When malicious patterns are detected, blocking activities take place to prevent the risk of data being breached. This concept is in line with the findings on continuous monitoring as described in chapter 4.
- **Continuous diagnostics and mitigation:** The institution enhances the current management strategies for devices in their IT landscape. They require employees to have managed devices, in order to be able to enforce updates and configuration settings that solve vulnerability issues on the devices. This way they enhance their Hardware Asset Management system. Furthermore, the device management techniques allow the SOC system to detect even more malicious behavior in the network.

7.3 Chapter Takeaways

In this section, we demonstrated how the ZEDEC method works in theory, by applying it to a fictional scenario. We also did an evaluation session with a security expert from SURF in order to find out whether he found the method's correctness and usefulness. The expert was positive about the correctness and usefulness of the method. One expert is not enough to yet draw conclusions on the effectiveness of the method yet, for this more research is needed in the future.

Chapter 8

Discussion

In this chapter, we present our findings on the (sub)questions of the research. Furthermore, we discuss how they should be interpreted and what the implications and limitations of this research are. Finally, we point out directions in which future research could take place in the scope of this research.

8.1 Findings

In this research, we aimed to improve the migration process for organisation that move from a perimeter-based security approach to a zero trust architecture, by proposing a systematic method that guides them in deciding which zero trust concepts to include in order to be resilient against cyber threats. To achieve this objective we answered the main research question: *how can organisations decide which security concepts to integrate in their zero trust architecture when migrating from a perimeter-based security architecture?* In order to answer this question, we answered six sub-questions, of which we describe our findings in this section.

SQ1: *What are zero trust concepts that organisations should consider to integrate in a zero trust architecture?*

We first identified which zero trust concepts an organisation should consider to integrate into its ZTA, to answer SQ1. Multiple zero trust concepts exist, of which eight concepts are logical components that the policy decision point takes as input for its decisions. These concepts are: (1) identity management, (2) data access policy management, (3) public key infrastructure, (4) activity logs, (5) continuous monitoring, (6) continuous diagnostics and mitigation, (7) threat intelligence and (8) industry compliance. Other zero trust concepts are located within the PDP. These are (1) the trust algorithm that ultimately decides whether access is granted to a subject, (2) authentication techniques that help the trust algorithm establish the trust level and (3) the design of the PDP component. Finally, three zero trust concepts exist on the data plane, namely (1) segmentation, (2) zero trust deployment models and (3) resource locations.

SQ2: *How do organisations currently decide which zero trust concepts to integrate in a zero trust architecture?*

Then we identified how organisations currently migrate towards a zero trust architecture to answer SQ2. We discovered steps from the SLR and interviews. The findings from the SLR show that in the process of identifying which zero trust concepts to integrate into a ZTA, the first activity is to identify the protection surface. Then a risk assessment and prioritization takes place. Based on the prioritized risks, the organisation identifies zero trust solutions to mitigate the risks. The interviews show that the first activity is to create a vision on the needed changes in the migration. Then the organisation identifies the context, performs a risk assessment and selects zero trust mitigations and concepts to mitigate the most important risks. Another finding from the interviews is that organisations tend to look at what zero trust concepts other organisations are selecting and do so similarly, rather than following a systematic method to select the concepts.

SQ3: *Which factors do organisations consider before deciding which concepts to integrate in a zero trust architecture?*

By answering SQ3, we determined what factors organisations consider when deciding which zero trust elements they want to integrate. The literature from the SLR does not provide any factors. The interviews, on the other hand, show that organisation consider factors that are similar to mitigations. The participants mentioned a total of 22 zero trust mitigations that also appear in the MITRE ATT&CK framework. Another 10 mitigations from the MITRE ATT&CK framework have a logical connection to zero trust concepts based on the literature found in the SLR.

SQ4: *How are the organisational decision-making factors (of SQ3) related to the existing zero trust concepts (of SQ1)?*

In SQ4 we identified how the mitigations were related to the zero trust concepts. A mapping of the relation between the mitigations and the zero trust components shows that all mitigations are part of one or more zero trust concepts. The result of the identified relations is the *decision matrix* that helps organisations identify which zero trust concepts they can integrate in their ZTA based on the mitigations they select. In the matrix we identified two clusters of zero trust concepts that are related to the same kind of mitigations. The first cluster is *identity and access management*, which encompasses the "identity management," "data access policy management," "trust algorithm" and "authentication techniques" concepts. The second cluster is *network monitoring and device management* which encompasses "continuous monitoring" and "continuous diagnostics and mitigation."

SQ5: *What is a method to decide which concepts to integrate in a zero trust architecture?*

The collection of method fragments discovered in the SLR and interviews, lead to the ZEro trust DECision making (ZEDEC) method, to answer SQ5. The main activities in the method are (1) create a vision on the needed changes, (2) identify context, (3) perform risk assessment and (4) identify zero trust concepts. The final step includes the decision matrix that helps detecting relevant zero trust concepts to integrate in the ZTA.

SQ6: *To what extent does the ZEDEC method work?*

In SQ6 we demonstrated that the method works in a fictional scenario. Furthermore, we did an evaluation session with a security expert, who confirmed the method's correctness and usefulness in achieving its objective. The ZEDEC method is a promising approach for assisting organisations in deciding which zero trust concepts to integrate in their ZTA.

8.2 Interpretation

In this section we compare the findings to previous studies. The discovered studies on the zero trust concepts in the SLR showed coherent data on elements that organisations should consider to integrate into their ZTA (SQ1). Therefore, the presented overview in this research captures all mentioned concepts in the papers found in the SLR, without major contradicting data. An interesting observation is that most papers found in the SLR build on the work of Rose et al. [44], which could explain the coherence among the studies.

The findings on how the current way that organisations are moving towards a ZTA (SQ2) correspond with the currently known methods in the literature, such as ISO 27005 and NIST SP800-207 [44, 53]. These studies state activities like *identify context*, *risk assessment and prioritization* and *identify zero trust solutions*, while the interviews showed similar activities such as *identify context*, *perform risk assessment* and *select zero trust mitigations*. The newly discovered activities and details from the interviews (as discussed in section 8.3) are not contradicted by the studies from the SLR. The discovered activity *create a vision on needed changes*, for example, fits well with the statement from Rose et al. [44] that the migration towards a ZTA requires thorough preparations. An unexpected finding in the interviews is that organisations do currently not follow a systematic method, but tend to look at what zero trust concepts other organisations integrate into their ZTA and do the same. A possible explanation for this is that before ZEDEC, no method for making these decisions existed.

We did not find previous studies on the factors that organisations consider during the decision-making process and their relation to the zero trust concepts (SQ3 and SQ4). Buck et al. [12] emphasize that this is the case in their SLR. An explanation for this could be that organisations have only recently started to integrate ZTA's into their IT landscape, and therefore organisations were not able to provide insights on these topics. The discovered factors derived from the MITRE ATT&CK framework have a relation with the *risk assessment and prioritization* step, as mentioned in multiple previous studies [2, 13, 17, 44, 58]. This suggests that mitigations are useful factors to consider in this process.

Overall, the elements that the ZEDEC method builds on are in line with findings from previous studies (SQ5 and SQ6). No studies were identified that contradict ZEDEC.

8.3 Implications

This section points out the implications of the research to science and practice. In this research, we provided a comprehensive overview of zero trust concepts that currently exist in the literature, which has not been done before [25]. The comprehensive overview, as summarised in Table 3, can be used as a starting point for future research that includes (a subset of) zero trust concepts. Furthermore, the overview can support researches that focus on other zero trust aspects, but require an understanding of zero trust concepts relevant to the field. From the practical point of view, these findings help organisations to get a clear overview of the possibilities of zero trust and concepts that they could consider, which helps them get an understanding of the possibilities that ZTA offers them.

The findings on the migration steps contribute to a detailed insight into the experiences of organisations and experts, by conducting expert interviews. Previous research already touches upon methods that could be used to integrate a ZTA, but research on the experiences of experts and organisations has not yet been done [12]. This research contributes to this field by identifying the new step *create vision on needed changes*, that precedes currently identified method activities [44, 53]. Furthermore, we discovered that the activities that experts mention are similar to the ones in the existing literature. Other additions to the current knowledge are that for each step we identified potential inputs that might affect them. Examples are *red teaming tests* and *experiences from other organisations.* Furthermore we identified that the outputs for the activities are typically documents. We also identified a list of actors that play a role in performing the steps. The newly gained data about the experiences of organisations with zero trust could be used for other research that takes the experiences from experts into consideration. Furthermore, the discovered steps could be used in researches for methods with similar properties and objectives.

We created the decision matrix, that maps zero trust mitigations from the MITRE ATT&CK framework to relevant zero trust concepts (see Figure 12). This matrix contributes to the field of experiences of organisations with ZTA's, since it is based on findings from the interviews. It shows what zero trust mitigations organisations consider when migrating towards a ZTA and what relations they see between the mitigations and zero trust concepts. Furthermore, it introduces a new application to the MITRE ATT&CK framework, which allows researches to improve the MITRE ATT&CK framework [55]. From the practical point of view, organisations can use this matrix as guidance to decide which zero trust concepts they should integrate into their ZTA.

The ZEDEC is a new method that includes newly discovered and other elements than existing approaches. As described in this section, these newly discovered elements can be used for future research, mainly in the field of experiences of users, applications in the real-world and the improvement of methods on decision-making in the field of cyber security. The method has a practical contribution to organisations that are moving towards a zero trust architecture. ZEDEC shows potential to guide them in deciding which zero trust concepts fit best in their ZTA, so that they can be resilient against cyber threats. The consecutive advantage for society is that the actors in the IT landscapes of these organisations are provided with appropriate security measures for their data and services.

8.4 Limitations

In this section, we state which limitations we encountered in the research. For the expert interviews in the *problem investigation* phase, we interviewed six security experts. We initially intended to do more expert interviews, in order to reach full saturation on the results. However, it was hard to find zero trust experts that wanted to participate in the research. A possible reason for this is that the topic is relatively new, and therefore experts on the topic are scarce. Eventually, we were able to interview six experts, which ensures that we found most important codes in the interviews [22]. The fact that findings on the current method of deciding which zero trust concepts to integrate in a ZTA were similar in the interviews and SLR, suggests that this was indeed the case.

In Chapter 3.4 we identified that a threat to the internal validity of this research is that zero trust experts with the same role or from the same organisations might bias the outcome of the interviews. To mitigate this threat we selected participants with different roles from different organisations. This way, we ensured the variety of our sample of experts. Participant D and E were from the same organisation, but focused on different areas of zero trust; governance and technical solutions respectively. This reduced the chances that the outcome of the interviews are biased, as well as the fact that we interviewed these experts separately — as we did with all experts. Although these mitigating factors were in place, the fact that the two experts came from the same organisation might still have an effect on the results.

Another limitation of this research is the evaluation of the ZEDEC method in the *treatment validation* phase. We intended to do a focus group with zero trust experts and persons that have less experience with zero trust, to evaluate to which extent the method works. We had this focus group planned with a set of participants, but we had to cancel it because the method was not ready in time. Because of time constraints in this research, we were unable to reschedule the focus group. We therefore decided to evaluate the ZEDEC method by demonstrating that the method works on a fictional scenario and doing an evaluation session with a security expert from SURF. Although expert opinions and simulations are potential methods to perform the treatment validation, a focus group with more participants would have had a stronger outcome [65].

8.5 Future Work

The insights gained in this research provide opportunities for future work. The first opportunity is to do additional validation sessions on the ZEDEC method. This way more data could be gathered to establish to which extent this method works and on how to improve the method. The objective of this would be to increase the usefulness of the method to organisations that want to decide which zero trust concepts they want to integrate into their ZTA. We recommend to organise a focus group, because this method suits the objective of exploring unknown territory [38]. We furthermore recommend to organise it with zero trust experts and persons that only have experience with a perimeter-based security architecture, in which participants are given a case on which they should apply the ZEDEC method. The fictional case as described in Chapter 7.2 could be used as a starting point for this. Ideally, the participants should work in groups, as is the case in reality. Observers should identify strengths and shortcomings on the method while the participants are working on the case, so that improvements can be made on the ZEDEC method afterwards [43].

Another future research opportunity is to try out the ZEDEC method in practice. This would contribute to the final step of the design cycle; *treatment implementation*. By implementing the ZEDEC method in a real environment, conclusions can be drawn as to which extent the ZEDEC method works in practice [65].

During the research, we also identified gaps in the current research on zero trust architectures. Although studies in recent years show an increase in researches performed on zero trust concepts, some concepts are still underexplored. A crucial concept in a ZTA that seems to be understudied is the *trust algorithm (TA)*. Current researches describe what properties exist for the TA, and state that AI solutions could be used to make access decisions [44]. Yet, there is no research focusing on how this algorithm could be implemented the most effectively. Given the fact that the TA plays a crucial role in a ZTA, it is of high importance for the ZEDEC method that organisations can integrate effective trust algorithms in their ZTA [44]. Therefore, we encourage researchers to study how the trust algorithm can be used the most effectively in a ZTA.

Another opportunity for future research is to discover the mental models about ZTA's for different user groups that work with it. A mental model is a representation to a person's knowledge about a topic [46]. Potential interesting user groups to investigate are parties that are involved in the development of ZTA's, as well as the users of ZTA's. A potential setup for this research could be inspired on the work of Binkhorst et al. [11], where multiple user groups are interviewed to establish and compare the mental models about ZTA's from different viewpoints.

Chapter 9

Conclusion

This research answered the research question: *how can organisations decide which security concepts to integrate in their zero trust architecture when migrating from a perimeter-based security architecture?* We used design science to create the ZEro trust DECision making (ZEDEC) method to answer it.

In the "problem investigation" phase, we conducted an SLR and interviews with six zero trust experts to establish the building blocks for ZEDEC. The first sub-question we answered was: *What are zero trust concepts that organisations should consider to integrate in a zero trust architecture?* Based on identified concepts in the SLR, we found that components that concern the policy decision point (PDP), which is responsible for making a decision on granting access to a subject, are: (1) the trust algorithm, (2) authentication techniques and (3) PDP design. For the logical components, that provide the PDP with input data, zero trust concepts are: (1) identity management, (2) data access policy management, (3) public key infrastructure, (4) activity logs, (5) continuous monitoring, (6) continuous diagnostics and mitigation, (7) threat intelligence and (8) industry compliance. On the data plane, we identify the concepts (1) segmentation, (2) deployment models and (3) resource location.

Then, we answered SQ2: *How do organisations currently decide which zero trust concepts to integrate in a zero trust architecture?* Based on the findings of the SLR and interviews we conclude that four main steps exist: (1) create vision on needed changes, (2) identify context, (3) perform risk analysis and (4) identify zero trust concepts.

The third sub-question is: *Which factors do organisations consider before deciding which concepts to integrate in a zero trust architecture?* We conclude based on the interviews that organisations consider zero trust mitigations during the decision-making. We introduced a list of 32 zero trust mitigations from the MITRE ATT&CK framework that the experts mentioned in the interviews.

In the fourth sub-question, *How are the organisational decision-making factors (of SQ3) related to the existing zero trust concepts (of SQ1)?*, we identified for each mitigation to which zero trust concepts they were related. We introduced the decision matrix that illustrates what zero trust concepts could be used to integrate the mitigation. All zero trust concepts are part of one or more zero mitigations. Two identified clusters of related mitigations in the matrix are: "identity and access management" and "network monitoring and device management."

In the "treatment design" phase we built the ZEDEC method based on the findings so far, answering SQ5: *What is a method to decide which concepts to integrate in a zero trust architecture?* The main steps of ZEDEC are: (1) create vision on needed changes, (2) identify context, (3) perform risk analysis and (4) identify zero trust concepts. The decision matrix is used in the final step as guidance to decide which zero trust concepts should be integrated.

In the "treatment validation" phase we andered the final sub-question: *To what extent does the ZEDEC method work?* Through an evaluation with a security expert of a theoretical scenario, we confirm that the method is correct and useful.

This research includes multiple scientific contributions. The ZEDEC method provides researchers with new insights on how to decide which zero trust elements to integrate into a ZTA, which they could use to improve and build other methods with similar properties and objectives. Other important new insights are (1) a complete overview of zero trust concepts that could be integrated into a ZTA, (2) an overview on zero trust mitigations that organisations consider to integrate into a ZTA and (3) a mapping that shows which zero trust concepts contribute to which mitigations. Furthermore, the findings in these contributions take experiences from organisations and experts into consideration, which has not yet been investigated in the field of zero trust integrations. The main practical implications are that ZEDEC could support organisations in the decision making on which zero trust concepts they want to integrate, so that they can be resilient against cyber threats. In order to utilize the potential of ZEDEC, we encourage researchers to continue improving it, by doing more evaluation sessions and eventually a case study on a real-world application of ZEDEC.

Bibliography

[1] National Cyber Security Centre (NCSC). *Cyber Security Assessment Netherlands*. Tech. rep. National Cyber Security Centre (NCSC), 2022. URL: `https://english.nctv.nl/binaries/nctv-en/documenten/publications/2022/07/04/cyber-security-assessment-netherlands-2022/Cyber+Security+Assessment+Netherlands+2022.pdf`.

[2] Z. Adahman, A.W. Malik, and Z. Anwar. "An analysis of zero-trust architecture and its cost-effectiveness for organizational security". In: *Computers and Security* 122 (2022). DOI: `10.1016/j.cose.2022.102911`.

[3] A. Alagappan, S.K. Venkatachary, and L.J.B. Andrews. "Augmenting Zero Trust Network Architecture to enhance security in virtual power plants". In: *Energy Reports* 8 (2022), pp. 1309–1320. DOI: `10.1016/j.egyr.2021.11.272`.

[4] Lampis Alevizos, Vinh Thong Ta, and Max Hashem Eiza. "Augmenting zero trust architecture to endpoints using blockchain: A state-of-the-art review". In: *Security and Privacy* 5.1 (2022), e191.

[5] Otis Alexander, Misha Belisle, and Jacob Steele. "MITRE ATT&CK for industrial control systems: Design and philosophy". In: *The MITRE Corporation: Bedford, MA, USA* 29 (2020).

[6] B. Ali, S. Hijjawi, L.H. Campbell, M.A. Gregory, and S. Li. "A Maturity Framework for Zero-Trust Security in Multiaccess Edge Computing". In: *Security and Communication Networks* 2022 (2022). DOI: `10.1155/2022/3178760`.

[7] Tejasvi Alladi, Vinay Chamola, Biplab Sikdar, and Kim-Kwang Raymond Choo. "Consumer IoT: Security vulnerability case studies and solutions". In: *IEEE Consumer Electronics Magazine* 9.2 (2020), pp. 17–25.

[8] A. AlQadheeb, S. Bhattacharyya, and S. Perl. "Enhancing cybersecurity by generating user-specific security policy through the formal modeling of user behavior". In: *Array* 14 (2022). DOI: `10.1016/j.array.2022.100146`.

[9] S. Alshomrani and S. Li. "PUFDCA: A Zero-Trust-Based IoT Device Continuous Authentication Protocol". In: *Wireless Communications and Mobile Computing* 2022 (2022). DOI: `10.1155/2022/6367579`.

[10] Andy Applebaum. *Getting Started with ATT&CK: Assessments and Engineering*. MITRE ATT&CK®. Aug. 1, 2019. URL: `https://medium.com/mitre-attack/getting-started-with-attack-assessment-cc0b01769cb4` (visited on 06/23/2023).

[11] Veroniek Binkhorst, Tobias Fiebig, Katharina Krombholz, Wolter Pieters, and Katsiaryna Labunets. "Security at the End of the Tunnel: The Anatomy of {VPN} Mental Models Among Experts and {Non-Experts} in a Corporate Context". In: *31st USENIX Security Symposium (USENIX Security 22)*. 2022, pp. 3433–3450.

[12] Christoph Buck, Christian Olenberger, André Schweizer, Fabiane Völter, and Torsten Eymann. "Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero-trust". In: *Computers & Security* 110 (2021), p. 102436.

[13] M. Bush and A. Mashatan. "From Zero to One Hundred". In: *Queue* 20.4 (2022), pp. 80–106. DOI: 10.1145/3561799.

[14] Mark Campbell. "Beyond zero trust: Trust is a vulnerability". In: *Computer* 53.10 (2020), pp. 110–113.

[15] B. Chen et al. "A Security Awareness and Protection System for 5G Smart Healthcare Based on Zero-Trust Architecture". In: *IEEE Internet of Things Journal* 8.13 (2021), pp. 10248–10263. DOI: 10.1109/JIOT.2020.3041042.

[16] Yang Chen, Hong-chao Hu, and Guo-zhen Cheng. "Design and implementation of a novel enterprise network defense system bymaneuveringmulti-dimensional network properties". In: *Frontiers of Information Technology & Electronic Engineering* 20.2 (2019), pp. 238–252.

[17] Z.A. Collier and J. Sarkis. "The zero trust supply chain: Managing supply chain risk in the absence of trust". In: *International Journal of Production Research* 59.11 (2021), pp. 3430–3445. DOI: 10.1080/00207543.2021.1884311.

[18] Bryan Embrey. "The top three factors driving zero trust adoption". In: *Computer Fraud & Security* 2020.9 (2020), pp. 13–15.

[19] L. Ferretti, F. Magnanini, M. Andreolini, and M. Colajanni. "Survivable zero trust for cloud computing environments". In: *Computers and Security* 110 (2021). DOI: 10.1016/j.cose.2021.102419.

[20] Muhamad Al Fikri, Fandi Aditya Putra, Yohan Suryanto, and Kalamullah Ramli. "Risk Assessment Using NIST SP 800-30 Revision 1 and ISO 27005 Combination Technique in Profit-Based Organization: Case Study of ZZZ Information System Application in ABC Agency". In: *Procedia Computer Science* 161 (2019), pp. 1206–1215. ISSN: 18770509.

[21] P. García-Teodoro, J. Camacho, G. Maciá-Fernández, J.A. Gómez-Hernández, and V.J. López-Marín. "A novel zero-trust network access control scheme based on the security profile of devices and users". In: *Computer Networks* 212 (2022). DOI: 10.1016/j.comnet.2022.109068.

[22] Greg Guest, Arwen Bunce, and Laura Johnson. "How many interviews are enough? An experiment with data saturation and variability". In: *Field methods* 18.1 (2006), pp. 59–82.

[23] Michael Gusenbauer and Neal R Haddaway. "Which academic search systems are suitable for systematic reviews or meta-analyses? Evaluating retrieval qualities of Google Scholar, PubMed, and 26 other resources". In: *Research synthesis methods* 11.2 (2020), pp. 181–217.

[24] K. Hatakeyama, D. Kotani, and Y. Okabe. "Zero Trust Federation: Sharing Context under User Control towards Zero Trust in Identity Federation". In: *IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events, PerCom Workshops*. Institute of Electrical and Electronics Engineers Inc., 2021. DOI: 10.1109/PerComWorkshops51409.2021.9431116.

[25] Yuanhang He, Daochao Huang, Lei Chen, Yi Ni, and Xiangjie Ma. "A Survey on Zero Trust Architecture: Challenges and Future Trends". In: *Wireless Communications and Mobile Computing* 2022 (2022).

[26] Ole R Holsti. "Content analysis for the social sciences and humanities". In: *Reading. MA: Addison-Wesley (content analysis)* (1969).

[27] Joint Task Force Transformation Initiative et al. *SP 800-37 Rev. 1. Guide for Applying the Risk Management Framework to Federal Information Systems: a Security Life Cycle Approach*. National Institute of Standards & Technology, 2010.

[28] Joint Task Force Interagency Working Group. *Security and Privacy Controls for Information Systems and Organizations*. Edition: Revision 5. National Institute of Standards and Technology, Sept. 23, 2020. DOI: 10.6028/NIST.SP.800-53r5.

[29] J. Junquera-Sánchez, C. Cilleruelo, L. De-Marcos, and J.-J. Martinez-Herráiz. "Access Control beyond Authentication". In: *Security and Communication Networks* 2021 (2021). DOI: 10.1155/2021/8146553.

[30] John Kindervag. "Build security into your network's dna: The zero trust network architecture". In: *Forrester Research Inc* 27 (2010).

[31] Barbara Kitchenham, O Pearl Brereton, David Budgen, Mark Turner, John Bailey, and Stephen Linkman. "Systematic literature reviews in software engineering–a systematic literature review". In: *Information and software technology* 51.1 (2009), pp. 7–15.

[32] G.M. Køien. "Zero-Trust Principles for Legacy Components: 12 Rules for Legacy Devices: An Antidote to Chaos". In: *Wireless Personal Communications* 121.2 (2021), pp. 1169–1186. DOI: 10.1007/s11277-021-09055-1.

[33] Brian Lee, Roman Vanickis, Franklin Rogelio, and Paul Jacob. "Situational awareness based risk-adapatable access control in enterprise networks". In: *arXiv preprint arXiv:1710.09696* (2017).

[34] S. Li, M. Iqbal, and N. Saxena. "Future Industry Internet of Things with Zero-trust Security". In: *Information Systems Frontiers* (2022). DOI: 10.1007/s10796-021-10199-5.

[35] Simon Liu and Rick Kuhn. "Data loss prevention". In: *IT professional* 12.2 (2010), pp. 10–13.

[36] Z. Liu, X. Li, and D. Mu. "Data-Driven Zero Trust Key Algorithm". In: *Wireless Communications and Mobile Computing* 2022 (2022). DOI: 10.1155/2022/8659428.

[37] S. Mandal, D.A. Khan, and S. Jain. "Cloud-Based Zero Trust Access Control Policy: An Approach to Support Work-From-Home Driven by COVID-19 Pandemic". In: *New Generation Computing* 39.3-4 (2021), pp. 599–622. DOI: 10.1007/s00354-021-00130-6.

[38] Lokanath Mishra. "Focus group discussion in qualitative research". In: *TechnoLearn: An International Journal of Educational Technology* 6.1 (2016), pp. 1–5.

[39] MITRE. *MITRE ATT&CK Framework*. July 2023. URL: https://attack.mitre.org/ (visited on 07/09/2023).

[40] MITRE. *MITRE ATT&CK Mitigations*. May 2023. URL: https://attack.mitre.org/mitigations/enterprise/ (visited on 06/30/2023).

[41] Abdallah Moubayed, Ahmed Refaey, and Abdallah Shami. “Software-defined perimeter (sdp): State of the art secure solution for modern networks”. In: *IEEE Network* 33.5 (2019), pp. 226–233.

[42] Jianli Pan and Zhicheng Yang. “Cybersecurity challenges and opportunities in the new” edge computing+ IoT” world”. In: *Proceedings of the ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization*. 2018, pp. 29–32.

[43] Fatemeh Rabiee. “Focus-group interview and data analysis”. In: *Proceedings of the nutrition society* 63.4 (2004), pp. 655–660.

[44] Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly. *Zero trust architecture*. Tech. rep. National Institute of Standards and Technology, 2020.

[45] Per Runeson and Martin Höst. “Guidelines for conducting and reporting case study research in software engineering”. In: *Empirical Software Engineering* 14.2 (Apr. 2009), pp. 131–164. ISSN: 1382-3256, 1573-7616. DOI: 10.1007/s10664-008-9102-8. URL: http://link.springer.com/10.1007/s10664-008-9102-8.

[46] Elizabeth B-N Sanders and Pieter Jan Stappers. *Convivial toolbox: Generative research for the front end of design*. Bis, 2012.

[47] Geoffrey Sanders, Timothy Morrow, Nataniel Richmond, and Carol Woody. *Integrating Zero Trust and DevSecOps*. Tech. rep. Carnegie-Mellon University Pittsburgh, 2021.

[48] S. Sarkar, G. Choudhary, S.K. Shandilya, A. Hussain, and H. Kim. “Security of Zero Trust Networks in Cloud Computing: A Comparative Review”. In: *Sustainability (Switzerland)* 14.18 (2022). DOI: 10.3390/su141811213.

[49] B. Sengupta and A. Lakshminarayanan. “DistriTrust: Distributed and low-latency access validation in zero-trust architecture”. In: *Journal of Information Security and Applications* 63 (2021). DOI: 10.1016/j.jisa.2021.103023.

[50] S.W. Shah, N.F. Syed, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss. “LCDA: Lightweight Continuous Device-to-Device Authentication for a Zero Trust Architecture (ZTA)”. In: *Computers and Security* 108 (2021). DOI: 10.1016/j.cose.2021.102351.

[51] Michal Shlapentokh-Rothman, Erik Hemberg, and Una-May O’Reilly. “Securing the software defined perimeter with evolutionary co-optimization”. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 2020, pp. 1528–1536.

[52] OASIS Standard. *Extensible access control markup language (xacml) version 3.0*. Jan. 2013. URL: http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.html.

[53] International Organization for Standardization (ISO). *Information security, cybersecurity and privacy protection — Guidance on managing information security risks (ISO/IEC 27005:2022)*. Tech. rep. International Organization for Standardization (ISO), Oct. 2022.

[54] International Organization for Standardization (ISO). *Risk management - Guidelines (ISO 31000:2018).* Tech. rep. International Organization for Standardization (ISO), Oct. 2018.

[55] Blake E Strom, Andy Applebaum, Doug P Miller, Kathryn C Nickels, Adam G Pennington, and Cody B Thomas. "MITRE ATT&CK: Design and philosophy". In: *Technical report.* The MITRE Corporation, 2018.

[56] N.F. Syed, S.W. Shah, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss. "Zero Trust Architecture (ZTA): A Comprehensive Survey". In: *IEEE Access* 10 (2022), pp. 57143–57179. DOI: 10.1109/ACCESS.2022.3174679.

[57] U. Tariq. "Zero-Tolerance Security Paradigm for Enterprise-Specific Industrial Internet of Things". In: *Electronics (Switzerland)* 11.23 (2022). DOI: 10.3390/electronics11233953.

[58] S. Teerakanok, T. Uehara, and A. Inomata. "Migrating to Zero Trust Architecture: Reviews and Challenges". In: *Security and Communication Networks* 2021 (2021). DOI: 10.1155/2021/9947347.

[59] Meltem Sönmez Turan, Kerry A McKay, Çagdas Çalik, Donghoon Chang, Lawrence Bassham, et al. "Status report on the first round of the NIST lightweight cryptography standardization process". In: *National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Rep.(NISTIR)* 108 (2019).

[60] D. Tyler and T. Viana. "Trust no one? A framework for assisting healthcare organisations in transitioning to a zero-trust network architecture". In: *Applied Sciences (Switzerland)* 11 (2021). DOI: 10.3390/app11167499.

[61] Steven Walker-Roberts and Mohammad Hammoudeh. "Artificial intelligence agents as mediators of trustless security systems and distributed computing applications". In: *Guide to Vulnerability Analysis for Computer Networks and Systems.* Springer, 2018, pp. 131–155.

[62] L. Wang, H. Ma, Z. Li, J. Pei, T. Hu, and J. Zhang. "A data plane security model of SR-BE/TE based on zero-trust architecture". In: *Scientific Reports* 12.1 (2022). DOI: 10.1038/s41598-022-24342-y.

[63] Rory Ward and Betsy Beyer. "Beyondcorp: A new approach to enterprise security". In: (2014).

[64] Inge van de Weerd and Sjaak Brinkkemper. "Meta-modeling for situational analysis and design methods". In: *Handbook of Research on Modern Systems Analysis and Design Technologies and Applications.* IGI Global, 2009, pp. 35–54.

[65] Roel J Wieringa. *Design science methodology for information systems and software engineering.* Springer, 2014.

[66] Claes Wohlin, Per Runeson, Martin Höst, Magnus C Ohlsson, Björn Regnell, and Anders Wesslén. *Experimentation in software engineering.* Springer Science & Business Media, 2012, pp. 104–112.

[67] S. Xiao, Y. Ye, N. Kanwal, T. Newe, and B. Lee. "SoK: Context and Risk Aware Access Control for Zero Trust Systems". In: *Security and Communication Networks* 2022 (2022). DOI: 10.1155/2022/7026779.

[68] J. Zhang, J. Zheng, Z. Zhang, T. Chen, K. Qiu, Q. Zhang, and Y. Li. “Hybrid isolation model for device application sandboxing deployment in Zero Trust architecture”. In: *International Journal of Intelligent Systems* 37.12 (2022), pp. 11167–11187. DOI: 10.1002/int.23037.

Appendix A

Information Letter and Consent Form

Information about the research

The interview you are asked to participate in is part of scientific research aiming to gain insights into the way organisations migrate towards a zero trust architecture. A zero trust architecture is defined by the NIST as: IT is defined as “an enterprise’s cybersecurity plan that utilizes zero trust concepts and encompasses component relationships, workflow planning, and access policies” (Rose et al., 2020). The research is conducted in collaboration with SURF. Although they are actively involved in the goals and dissemination of research, they are not involved in the carrying out of the study.

How will the study be carried out?

The interview will take 60 to 90 minutes, during which the researcher will ask questions in a semi-structured format. The interview will be recorded. After the recordings are transcribed, you will get the opportunity to remove any information from the text that should not be included in further analysis. Following the researchers’ analysis of these transcripts, you will be asked to evaluate and add to a summary of the results that are based on the interviews. You will not be reimbursed for your participation in this study.

What will we do with your data?

During this interview, data about your (and your organisation’s) experiences with zero trust will be collected. Although the objectives and design of this study do not require specific personally identifiable information, the data collected should be considered as such. The interview will be recorded before it is transcribed. Interview recordings will be retained for up to six months until transcribed. The non-pseudonymised transcripts will only be processed by UU researchers who are collaborating in the study, or who are responsible for assessing its implementation. After analysis, the transcripts will be further pseudonymised as described in the next section. There are no specific increased privacy risks related to the nature of the collected personal data or the processing that the data will undergo. The data is stored and processed exclusively in the EU and all third party applications used have an appropriate data processing agreement with Utrecht University.

Processed data will be retained for at least 10 years for the purposes of research integrity. Before this archival, all personal information that can reasonably be traced back to you or your organisation will have been removed or changed before the files are shared with other researchers or the results are made public. The researcher will keep a link that identifies you and your organisation with the information, but this link will be kept secure and only available to the researcher. Any information that can identify you will remain confidential. The information in this study will only be used in ways that do not reveal who you are. You and your organisation will not be named or identified in publications about this study or in documents shared with other researchers.

What are your rights?

Participation is voluntary. We are only allowed to collect your data for our study if you consent to this. If you decide not to participate, you do not have to take any further action. You do not need to sign anything. Nor are you required to explain why you do not want to participate. If you decide to participate, you can always change your mind and stop participating at any time, including during the study. You will even be able to withdraw your consent after you have participated. However, if you choose to do so, we will not be required to undo the processing of your data that has taken place up until that time. The research data we have obtained from you up until the time when you withdraw your consent will be erased.

Approval of this study

This study has been approved by the Science - Geo Ethics Review Board (SG ERB) at Utrecht University. If you have a complaint about the way this study is carried out, please send an email to the secretary of this Committee: etc-beta-geo@uu.nl. If you have any complaints or questions about the processing of personal data, please send an email to the Data Protection Officer of Utrecht University: privacy@uu.nl. The Data Protection Officer will also be able to assist you in exercising the rights you have under the GDPR. Please also be advised that you have the right to submit a complaint with the Dutch Data Protection Authority (https://www.autoriteitpersoonsgegevens.nl/en).

More information about this study?

In case you have additional questions, please contact Bjorn van Dijen(researcher and data controller for the study) b.vandijen@students.uu.nl or Kate Labunets (project supervisor for the study) at k.labunets@uu.nl.

References

Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero trust architecture (No. NIST Special Publication (SP) 800-207). National Institute of Standards and Technology.

Consent Form

1. I have read and understood the study information dated $currentdate, or it has been read to me. I have been able to ask questions about the study and my questions have been answered to my satisfaction: (yes/no)

2. I consent voluntarily to be a participant in this study and understand that I can refuse to answer questions and I can withdraw from the study at any time, without having to give a reason: (yes/no)

3. I understand that information I provide will be used for the report and publications in academic venues (like conferences or journals): (yes/no)

4. I understand that personal information collected about me that can identify me, such as my name or email address, will not be shared beyond the study team: (yes/no)

5. I additionally agree that my information can be quoted in research outputs (yes/no)

6. I give additional permission for the pseudonymised interview transcript that I provide to be archived in UU's YoushadorData so it can be used for future research and learning (yes/no)

7. My name:

Appendix B

Interview Protocol

In this guide, questions are marked as *italic text*. Non-italic texts indicate clarifications, instructions or explanations to the interviewee. The expected duration of the interview is 60 to 90 minutes.

Overview of interview setup

Before recording (1-5 min)
Introduction (5 min)
Experience with zero trust migration (5-10 min)
Migration steps (20-25 min)
Factors + concept matrix (15-20 min)
Interview end (1 min)

Before recording (1-5 min)

- Thank the interviewee for their participation and explain what we are going to do in the interview: 60-90min in total; first part: background questions; second and third part: more focussed questions.
- *Did you receive the informed consent form?*
- *Do you have questions about it?*
- *Did you fill in the informed consent form?*

Introduction (1-5 min)

- **Start recording**
- State that we need some background information (anonymous of course) to prove that the interviewee is indeed an expert on zero trust.
- About organization:
  - *How many employees work at the organisation?*
  - *What kind of service does the organisation offer?*
  - *Who are you offering this service to?*
- Personal questions:
  - *What is your current role in the organisation?*

  - *How long have you already been working at this function?*
  - *How long have you already been working with zero trust?*
  - *Do you have any other experiences with zero trust specifically?*
  - *Do you have any certificates or graduations that are relevant in the field of security?*

Experience with zero trust migration (5-10min)

- *What does your organisation do with zero trust?*
- *What is the main objective of doing a zero trust migration?*
- *When did you start the migration towards a ZTA?*
- *What are the main milestones you passed during the migration towards the zero trust architecture?* - The goal is to get a brief overview of the:
  - Timeline of the migration cycles that are made
  - The components that were included in the different migrations
- *What are elements you want to migrate in the future?*
- *Do you have recommendations for people that start with their first ZT migration?*

Migration steps (25-30 min)

For this step, we use a collaboration file in whimsical.com. The file allows the interviewer and the interviewee to understand the migration steps that the interviewee takes in a migration towards a ZTA. The interviewer asks the interviewee questions to detect migration steps and creates a workflow diagram while the interviewee answers. The interviewee sees those changes and is allowed to change the model of the interviewee at all times. For the sake of simplicity, AND/OR gates are not initially introduced, but are allowed if the interviewee wants to use them. An example of the intended workflow that is to be created, is depicted in Figure 15.

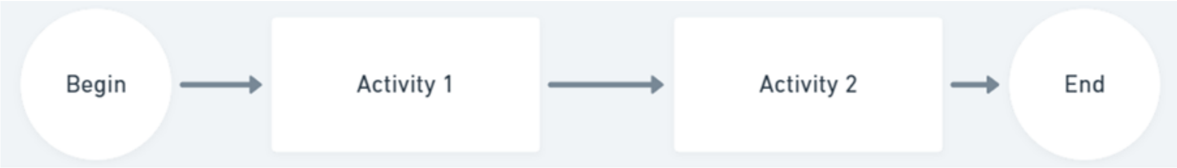

Figure 15: Intended workflow layout during the expert interviews

First, we identify a general overview of the starting point, end point and main activities using the following questions.

- State that the interviewee is asked to answer these questions in a general way, not too detailed.
- *What are prerequisites to start the migration?* - This might give an insight into the starting situation (begin) or helps identify extra activities
- *What are the main activities in migrating towards a zero trust activity in general?* - With this question we create an overview of the main activities.
- *When do you consider the migration as "finished," if you find it ever finished at all?* - This question helps us identify the end situation.

When this general overview is created, ask the interviewee to think about one or more specific migration cycles, and make more models like this. If applicable, the general model could be altered based on new insights gained in this process.

- *Can you describe a specific migration cycle you encountered, step-by-step? (continue until no more migration cycles exist or no more new activities are detected after starting a new cycle).*

Then we identify actors, inputs/tools and outputs for each activity. An example of the items that can be connected to the activities is illustrated in Figure 16. To discover these elements, ask the interviewee for all identified activities:

- *Who performs the activity? (actor)*
- *What inputs or tools do you need to execute the activity? (input/tool)*
- *What outputs are created after the activity? (output)*

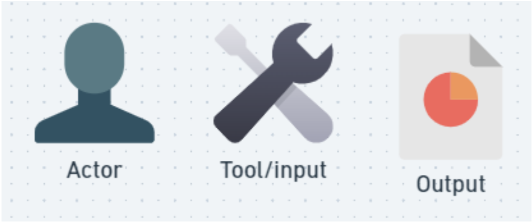

Figure 16: Icons used to identify actors, tools/inputs and outputs during the expert interviews

Factors + concepts matrix (25-30 min)

We want to identify what the factors are that the interviewee considers before he decides which zero trust concepts to choose. We want to know this for every single concept discovered in SQ1. In order to capture this information, we fill in the matrix as shown in Figure 17. The interviewees should identify the factors for each concept, and map the factors to other concepts that are applicable to the factor as well.

Explain the exercise by the following guidelines:

- Explain an example of the kind of factors that we mean in a different context. Use the following example: when you go outside, you consider factors like ”the sun is shining” and ”it is raining.” State that for all different concepts, this is the kind of factor what we are looking for.
- Explain to the interviewee how the matrix works. That we identify factors for all concepts mentioned in the matrix, and that the factors can be applicable on multiple concepts.

The procedure to capture the data in this step is as follows:

1. Explain the first concept in the list and asks the interviewee to reveal which factors have an impact on the decision on the inclusion of the concept in the final ZTA. *Ask the interviewee what the factors for the concept are.* Then, continue to the next concept until all concepts are explained.

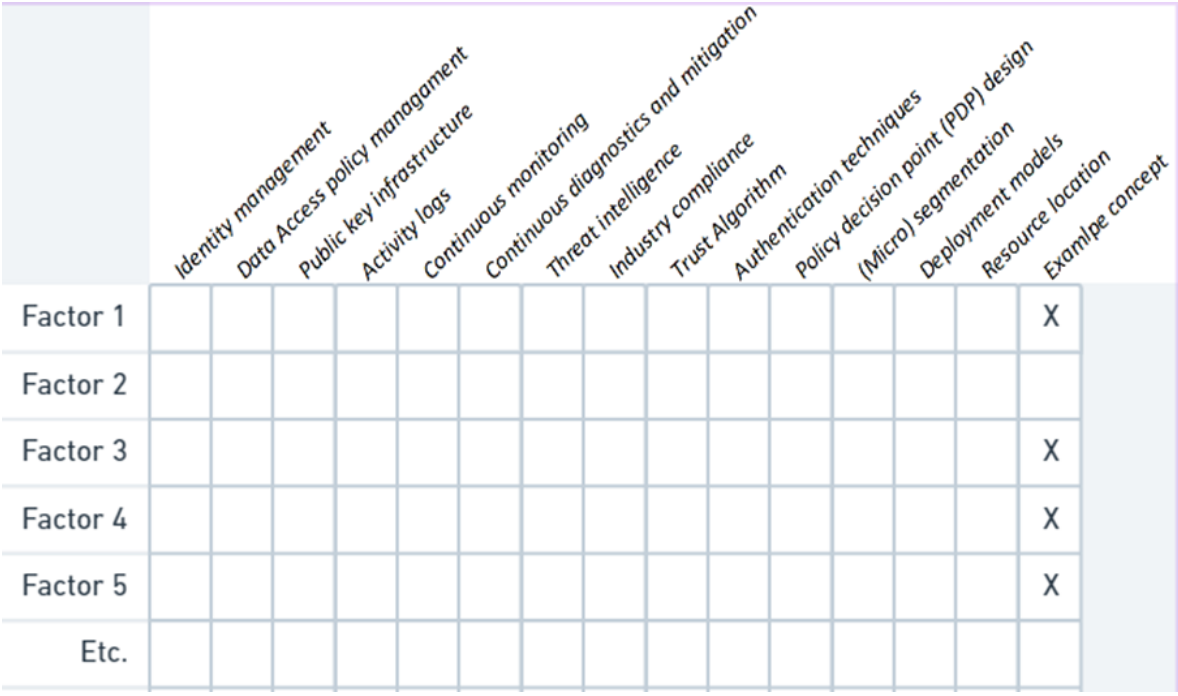

	Identity management	Data Access policy managament	Public key infrastructure	Activity logs	Continuous monitoring	Continuous diagnostics and mitigation	Threat intelligence	Industry compliance	Trust Algorithm	Authentication techniques	Policy decision point (PDP) design	(Micro) segmentation	Deployment models	Resource location	Examlpe concept
Factor 1															X
Factor 2															
Factor 3															X
Factor 4															X
Factor 5															X
Etc.															

Figure 17: Matrix used to capture factors that are applicable on the ZTA concepts

- We add an "X" to the factors that have an influence and add nothing to the factors that do not have an influence.
- We add a ">" when a participant mentions that a factor is relevant to all other concepts from that point onwards.
- We add all newly mentioned factors to the list.

2. When factors are identified for all concepts, *ask the interviewee to take another look at the identified factors and check if they might be applicable to other concepts.* If the time allows, go past every single factor to check this.
3. Some factors were found in the literature. Mention them to the interviewee and *ask the interviewee whether these are applicable on the concepts, and on which concepts.* The list of factors is mentioned below.

The factors that we found in the literature and want to validate with the interviewee are enumerated in this list. Please note that for the eventual method we decided to use a different kind of factors than the ones mentioned in this list.

- **Cost** - Organisations have a limited budget on their security expenses.
- **Interoperability** - Are devices and software compatible with other devices? Keep options open and prevent vendor lock-ins. Also considers data types.
- **Usability** - Limited frustrations from users, otherwise they start using other solutions than you want. Make sure they understand how it works and are willing to change to new solutions.
- **Processing latency** - When an access request is made, a complex TA might take a long time before it made a decision. Limit waiting times.

- **Contribution to risk mitigation** - How big is the chance that this concept is actually going to make the IT landscape more secure? The concept should have a sufficient contribution to risk mitigation to be worth integrating.
- **Applicability** - Is the concept always available for use? Or is it only applicable under certain conditions?
- **Organizational readiness** - Is there enough knowledge available within the organisation to integrate the concept?
- **Impact on stakeholders** - Does the integration of a concept have consequences or advantages for stakeholders and partners?
- **Laws and regulations compatibility** - The concepts need to be integrated in a lawful way.
- **Complexity** - An organisation might want to avoid the use of too many complex solutions.
- **Sustainability** - Limit the ecological footprint.
- **User privacy** - Make sure that the privacy of the users is guaranteed and no unnecessary data is processed (especially with microphone and camera using concepts).
- **Upcoming techniques** - Consider techniques that might decrease the effect of your solutions. For example quantum computing.
- **Resource limitations** - What limitations do your resources have? For example IoT devices and legacy systems.

Interview End (1 min)

- *Would you like to add anything else that was not discussed?*
- *Do you have any further questions?*
- Thank the interviewee for their participation.