UiA University of Agder

Designing the Extended Zero Trust Maturity Model

A Holistic Approach to Assessing and Improving an Organization’s Maturity Within the Technology, Processes and People Domains of Information Security

Jarand Nikolai Jansen
Simen Tokerud

SUPERVISOR
Marko Ilmari Niemimaa

University of Agder, 2022
Faculty of Engineering and Science
Department of Information Systems

ABSTRACT

Zero Trust is an approach to security where implicit trust is removed, forcing applications, workloads, servers and users to verify themselves every time a request is made. Furthermore, Zero Trust means assuming anything can be compromised, and designing networks, identities and systems with this in mind and following the principle of least privilege. This approach to information security has been coined as the solution to the weaknesses of traditional perimeter-based information security models, and adoption is starting to increase. However, the principles of Zero Trust are only applied within the technical domain to aspects such as networks, data and identities in past research. This indicates a knowledge gap, as the principles of Zero Trust could be applied to organizational domains such as people and processes to further strengthen information security, resulting in a holistic approach. To fill this gap, we employed design science research to develop a holistic maturity model for Zero Trust maturity based on these principles: The EZTMM. We performed two systematic literature reviews on Zero Trust and Maturity Model theory respectively and collaborated closely with experts and practitioners on the operational, tactical and strategic levels of six different organizations. The resulting maturity model was anchored in prior Zero Trust and maturity model literature, as well as practitioner and expert experiences and knowledge. The EZTMM was evaluated by our respondent organizations through two rounds of interviews before being used by one respondent organization to perform a maturity assessment of their own organization as a part of our case study evaluation. Each interview round resulted in ample feedback and learning, while the case study allowed us to evaluate and improve on the model in a real-world setting. Our contribution is twofold: A fully functional, holistic Zero Trust maturity model with an accompanying maturity assessment spreadsheet (the artifact), and our reflections and suggestions regarding further development of the EZTMM and research on the holistic application of Zero Trust principles for improved information security.

ACKNOWLEDGEMENTS

We would like to give a special thanks to our supervisor, Marko Ilmari Niemimaa for the detailed and valuable feedback and guidance provided to us throughout the process of writing this thesis. We would also like to thank all our respondents and the organizations that contributed to this model – You know who you are, and that this model would not exist without you. Lastly, we would like to thank our friends and families for being understanding and supportive through the trying process of developing the EZTMM in addition to full-time work and other commitments.

CONTENTS

ABSTRACT.....I

ACKNOWLEDGEMENTS.....II

CONTENTS.....III

FIGURES.....V

TABLES.....VI

GLOSSARY AND DEFINITIONS.....VII

1 INTRODUCTION.....1

1.1 Research gap and research questions.....2

2 KNOWLEDGE BASE.....4

2.1 Maturity Models.....4

2.2 Existing Information Security Maturity Models.....5

2.2.1 Cybersecurity Capability Maturity Model.....5

2.2.2 ZeTuMM.....6

2.2.3 CTI-SOC2M2.....7

2.2.4 Zero Trust Maturity Model.....7

2.3 Zero Trust.....8

2.3.1 Network segmentation.....9

2.3.2 Access Control & Least Privilege.....10

2.3.3 Dynamic Access Policies.....10

2.3.4 Monitoring and logging.....11

2.3.5 Multi-factor Authentication.....11

2.3.6 Device Verification.....11

2.3.7 Threat protection.....12

2.3.8 Encryption.....12

3 DESIGN SCIENCE RESEARCH.....13

3.1 Literature Review.....13

3.2 Design Process.....19

3.3 Data Collection.....26

3.3.1 The Respondents.....26

3.3.2 Interviews.....26

3.3.3 Case study evaluation.....27

3.4 Data Analysis 27
3.4.1 Interview and written feedback 28
3.4.2 Inclusion and Exclusion Criteria 28
3.4.3 Case study evaluation 29
4 DESIGNING THE EXTENDED ZERO TRUST MATURITY MODEL ..30
4.1 Iteration 0: Literature Review and existing model comparison 31
4.2 Iteration 1: Gathering and implementation of respondent feedback on the initial EZTMM draft: Restructuring the model 32
4.3 Iteration 2: Gathering and implementation of respondent feedback on the second EZTMM draft: Polishing the model 33
4.4 Iteration 2.5: Expert feedback on Identity 35
4.5 Iteration 3: Case study evaluation...35
5 DISCUSSION 39
5.1 Reflections on the Design of EZTMM 39
5.2 The Components of Zero Trust...39
5.3 Existing maturity models and the existing knowledge gap 40
5.4 The EZTMM 40
5.5 Acceptance Evaluation of the EZTMM...41
5.6 Practical implications 43
5.7 Future Research...44
5.7.1 Zero Trust Principles for organizational aspects.....................44
5.7.2 Expansion of the EZTMM...44
5.7.3 Compatibility with existing standards.....................................45
5.7.4 Weighted focus areas 45
5.7.5 Case studies on the application of Zero Trust principles to all organizational processes 45
5.8 Limitations...45
6 CONCLUSION 47
7 REFERENCES..49

FIGURES

Figure 1: PeP and PDP Diagram 11
Figure 2: EZTMM design research framework.. 20
Figure 3: EZTMM Iteration Overview... 30
Figure 4: The EZTMM Evaluation Sheet Control Questions.......................... 36
Figure 5: EZTMM Evaluation Sheet original Maturity Dashboard 37
Figure 6: EZTMM evaluation sheet updated Maturity Dashboard.................. 38

TABLES

Table 1: Maturity Model Comparison 8
Table 2: Zero Trust Concept Search Strings 14
Table 3: Zero Trust Maturity Model Search Strings 14
Table 4: Inclusion and Exclusion Criteria for Literature Review Articles...... 15
Table 5: Zero Trust Concept Matrix.. 16
Table 6: Inclusion and Exclusion Criteria for Maturity Model Theory........... 17
Table 7: Articles on Maturity Model Theory 18
Table 8: Maturity model Theory Search Strings 18
Table 9: Maturity Model Theory Articles 19
Table 10: Acceptance Criteria for the EZTMM 24
Table 11: Overview of Respondents by Role.. 26
Table 12: Interview Overview 27
Table 13: Inclusion and Exclusion Criteria for interview feedback 28

GLOSSARY AND DEFINITIONS

EZTMM: The Extended Zero Trust Maturity Model. A holistic approach to assessing and improving Zero Trust maturity in organizations developed during this thesis.

Holistic: An approach that aims to consider all aspects of a topic. In this thesis: An approach that considers both the technical and the organizational aspects of information security and Zero Trust, meaning that the technology, processes and people domains are all considered.

Implicit Trust: A type of trust that is implied, for example by the request originating from within the organization's internal network: If the organization's internal network is considered safe, it is then implied that any request originating from this network also is safe and does not require further authentication.

Network Segmentation: The act of segmenting the organization's network into smaller pieces. Often categorized into macro and micro segmentation. Macro segmentation involves segmenting the network into larger zones accommodating different security needs, or separating different environments such as production, test and development. Micro segmentation involves segmenting the network into much smaller pieces based on either application workloads, services or individual servers.

Perimeter Defense: An approach to defending the organization's infrastructure by implementing a defensive perimeter that protects against external attacks.

Zero Trust: An approach to information security that tries to eliminate all implicit trust.

Zero Trust Component: A component is a part of a greater whole. In terms of Zero Trust components, this thesis argues that any technical measure or organizational policy, process or procedure that incorporates Zero Trust principle is a component of that organization's Zero Trust adoption.

Zero Trust Principles: The core principles upon which a Zero Trust architecture or strategy is built. There are many different versions of these principles. This thesis uses the Microsoft Zero Trust principles of verify explicitly, use least privileged access and assume breach.

1 INTRODUCTION

Digitalization has increased massively in the last decades, with advancements in internet-and communication technologies being a major contributing factor. As organizations have moved their business, activities, and interactions to cyberspace, the need for cybersecurity has increased rapidly (Li & Liu, 2021). For years, cybersecurity professionals have built their defenses based on the mantra "trust but verify" (Warren, 2021). This approach involves trusting users and endpoints within the organizational network after initially authenticating their identity. However, trends such as ransomware attacks and credential theft show that security incidents are often caused by exploitation of this trust (Verizon, 2022).

The traditional way of building computer network defenses, relying on perimeters and trusting everything on the inside, is a practice ready for evaluation. Considering the inside of the perimeter secure implies less effort is needed to protect against attackers or malicious insiders with access to the internal network. No network segmentation and little access control enables easier lateral movement and bigger potential damages (Ferretti, Magnanini, Andreolini, & Colajanni, 2021). There are also big shifts in technology trends, such as cloud computing and remote work, making the bounds of the perimeter harder to define (Teerakanok, Uehara, & Inomata, 2021).

Zero Trust has been coined as the solution to the problem (Buck, Olenberger , Schweizer, Fabiane, & Torsten, 2021) and has transitioned from being mystical and exciting to being a model many companies aspire to adopt. The Zero Trust model traditionally suggests assuming all networks, endpoints, identities, and solutions are compromised, treating both internal and external requests equally. Trust is no longer implicit; it is earned through rigorous verification (Teerakanok, Uehara, & Inomata, 2021).

The fundamental principles of Zero Trust can be traced back to the origins of the internet. One example is the change introduced to RFC 1122 in 1989 (Internet Engineering Task Force, 1989): "In general, it is best to assume that the network is filled with malevolent entities that will send in packets designed to have the worst possible effect." When John Kindervag introduced the term Zero Trust (Kindervag, 2010) the focus was to eliminate the idea of trusted and untrusted networks and see everything as untrusted.

He introduced three foundational concepts: Ensure that all resources are accessed securely regardless of location, adopt a least privilege strategy and strictly enforce access control, and inspect and log all traffic. These principles might be interpreted as network centric. Microsoft (Microsoft, 2019) has defined three similar principles to describe Zero Trust:

1. **Verify explicitly:** Organizations that verify explicitly use all data and information available to reduce uncertainty and implicit trust.
2. **Use least privileged access:** Using least privilege is always providing the least number of permissions necessary.
3. **Assume breach:** Assuming breach is when you already consider your digital environment compromised.

While it is possible to build a security strategy based on Zero Trust, the model traditionally only addresses the weakness of implicit trust in a communications network. Because of this, some security researchers have called this approach to security architecture fundamentally flawed (SABSA, 2022). Multiple maturity models haven been proposed with purpose to help organizations assess and improve their Zero Trust capabilities. Modderkolk (Modderkolk, 2018) presents a maturity model for assessing overall information security maturity with Zero Trust as the focal point, developed as part of academic research. However, the model does not try to apply the ideas of Zero Trust to other domains within information security. This makes the model hard to distinguish from other modern maturity models meant to assess information security. Other relevant models mainly focus on technology and were either developed by government organizations or for commercial purposes. Building an overall information security strategy based on such models is challenging as they do not address organizational aspects such as processes and people. This is supported by Zahoor et al. (Zahoor, Mahmood, & Javed, 2015) which highlights the importance of access control, defining security policies, security awareness and training, and argues that a holistic approach is necessary in information security management. As May & Dhillon observed in 2010, information security has become a multidimensional discipline where both social and technical considerations must be considered in a coherent manner (May & Dhillon, 2010).

1.1 Research gap and research questions

Zero Trust and information security maturity models have been focused on the technology domain, with little to no effort being made to apply the same principles to organizational domains such as the people and processes. Our research sets out to fill this gap by designing a maturity model to measure Zero Trust maturity in both technological and organizational aspects. Our intension is to extend the Zero Trust term by applying the foundational principles suggested by Microsoft to organizational domains within information security. Combining these organizational aspects with the technology aspects of traditional Zero Trust maturity models enables a holistic assessment and improvement of Zero Trust capabilities. The research presented in this thesis will cover how the Extended Zero Trust Maturity Model (EZTMM) was designed, along with the finished product.

Our research sets out to address the identified gap with the following research question:

RQ: How can a model be designed for organizations to assess and improve their technical and organizational Zero Trust maturity?

To answer this question, we have identified three supporting research questions that must also be answered:
SRQ1: What are the components of Zero Trust?
SRQ2: Which Zero Trust maturity models exist?
SRQ3: Which organizational processes can benefit from the Zero Trust principles?

Supporting research questions number one and two serve as the foundation of our knowledge base and will be answered through literature reviews. Answering the third supporting research question is crucial for developing a Zero Trust maturity model that can be used to assess organizational aspects. This question will be answered with the help of subject matter experts through several iterations of feedback.

Design science research, which focuses on creating unique problem-solving artifacts, was chosen as the research approach. The maturity model serves as our unique artifact, and it was designed following guidelines proposed by Hevner et al. (Hevner, Ram, March, & Park, 2004). However, seven requirements specifically defined for designing maturity models in design science research served as the main research criteria (Poeppelbuss & Roeglinger, 2011). Literature studies covering the theory of maturity models along with a comparison of existing ones, and Zero Trust principles and components were performed. These formed the knowledge base used to create the initial draft of the EZTMM. Several iterations were developed based on feedback from numerous subject matter experts. A case study evaluation was conducted to see how the model would perform in a real-life setting and if the model complied with predefined acceptance criteria. Results of our research includes a comprehensive maturity model to assess and improve overall Zero Trust maturity in an organization, along with a tool that organizations can use for self-evaluation.

2 KNOWLEDGE BASE

The following chapter contains a literature review serving as the foundation and knowledge base for this master thesis and the development of a maturity model. Research performed in this master thesis revolves around designing a maturity model based on Zero Trust and its' principles. Dividing the literature study into three parts therefore seemed natural – theory on maturity models, a comparison of existing models and Zero Trust.

A description of the research method for the literature reviews can be found in chapter 3.1.1 – Literature review.

2.1 Maturity Models

Maturity models are used to assess organizations as-is situation of current capabilities in a certain domain. Further use includes prioritizing future improvements and measuring progress (Poeppelbuss & Roeglinger, 2011). Different levels are used to describe a path from the initial state to improved maturity of capabilities (Becker, Knackstedt, & Pöppelbuß, 2009). According to Gottschalk (Gottschalk, 2009)" some models suggest that organizations progress through stages while others argue that there may be multiple paths through the stages".

Software Engineering Institute (SEI) published the "Capability Maturity Model for Software" in 1993 (Paulk, Curtis, Chrissis, & Weber, 1993). The model has served as a blueprint for a large amount of other maturity models developed since then (Poeppelbuss, Niehaves, Simons, & Becker, 2011). This has led to maturity models being subject of criticism due to their perceived redundancy. Other points of criticism have been the lack of documentation of design process and principles used during development (Becker, Knackstedt, & Pöppelbuß, 2009).

De Bruin et al. (Bruin, Freeze, Kulkarni, & Rosemann, 2005) suggest three different application-specific purposes for maturity models: descriptive, prescriptive, and comparative. Models that are descriptive are purely used to describe the as-is, and do not make any suggestions to improve maturity. Prescriptive models are used by organizations to improve their capabilities, while comparative models can be used to compare practices across industries and industry standards. A maturity model can have one or multiple purposes.

2.2 Existing Information Security Maturity Models

Among the points of criticism of maturity models are the lack of documentation of the design process. Due to this reason, a comparison of existing maturity models was performed to investigate their characteristics and development process. The requirements for designing maturity models introduced by Becker et al. (Becker, Knackstedt, & Pöppelbuß, 2009) were used as inspiration for comparing each models' design process. The model application-specific purpose (descriptive, prescriptive, or comparative) was included as well as model structure and content. The latter two were included to gain an overall better understanding of models developed for the information security domain. This comparison could identify if the proposed research was redundant or confirm an existing need. It also served as a valuable knowledge base for developing the initial iteration of the maturity model.

Maturity models of two kinds are included in the comparison. The first group comprises models documented and developed by academic research. These models contribute to the knowledge base by providing great insight into leveraged design processes and decisions made during development. The second group consists of models created by government agencies. They do not offer the same detailed documentation of the development process. However, they do provide valuable insight into security controls that are considered best practices, especially related to Zero Trust.

2.2.1 Cybersecurity Capability Maturity Model

The U.S Department of Energy developed the C2M2 maturity model in collaboration with the Department of Homeland Security and subject-matter experts from the electricity subsector (U.S Department of Energy, 2021). The model saw its first release in 2012, while version 2.0 was published in July 2021. C2M2 aims to help organizations build better cybersecurity programs, benchmark capabilities, and prioritize future actions and investments. Although the Department of Energy was a significant contributor to creating the model, the intended audience is described as any organization regardless of sector.

The C2M2 Version 2.0 has been released in multiple versions and is being improved continuously, therefore one could argue that the developers are using an iterative procedure. It is stated that more than 60 industry experts gave their feedback on the model after version 2.0 was developed through a series of working sessions. In addition to expert feedback the model is said to be built upon existing cybersecurity resources. The model was developed for giving descriptive guidance. However, it could be argued that the provided control questions could be used to create roadmaps for improvement. In that case, the model also serves a prescriptive purpose.

C2M2 has four different stages of maturity, ranging from zero to three. The first stage of maturity MIL0 describes a level where no practices are performed whereas MIL3 is the most advanced state of maturity. The model includes control questions to evaluate maturity in 10 different domains within information security. Maturity within domains is independent of each other.

2.2.2 ZeTuMM

ZeTuMM is a model for assessing Zero Trust maturity (Modderkolk, 2018). The author describes it as a model designed according to Zero Trust principles. Its' main goal is assessing organizations cybersecurity capabilities. The chosen design process for ZeTuMM is very well documented because it was developed as part of a master thesis.

As part of the design process, Modderkolk performs a comparison of six different maturity models related to information security. The comparison along with a literature study serves as the knowledge base for the initial drafts of the model. A case study is used to both further develop and verify the models' usefulness.

The model ranks maturity on three different levels. 53 capabilities are hosted within 15 domains (referred to as focus areas). 428 control questions are listed and used to determine maturity. Controls suggested within the model have been derived from established frameworks such as CIS Controls and the NIST Special Publication 800-53. The purpose is mainly assessing current capabilities, and the model is therefore of the descriptive kind.

While the model is said to center around Zero Trust and its principles, little effort seems to have been put into applying the Zero trust principles to each domain. Zero Trust is only present for areas related to technology, while organizational areas are based on best practices identified in other frameworks and models. Distinguishing this model from other maturity models related to information security is therefore challenging.

2.2.3 CTI-SOC2M2

Instead of providing a maturity model to assess overall cybersecurity in an organization the CTI-SOC2M2 model focuses on one very specific domain. It aims to improve the integration between cyber threat intelligence (CTI) and security operation centers (SOC). It does so by mapping CTI data in different formats to services often provided by security operation centers (Vielberth, Schlette, & Pernul, 2021).

A comparison of existing maturity models is performed as an initial activity in the development phase. The comparison briefly describes several models related to information security. Additionally, a literature study is performed and together these activities are used to create develop the model. Once developed, the model was evaluated using a mixed approach with a quantitative user study, and qualitative evaluation based on expert interviews. There is no mention of application purpose in the documentation, but the description of use indicates both descriptive and prescriptive purposes.

Six different levels are used to rank the ability to integrate each CTI format into different SOC services. An overall maturity level is then evaluated based on capabilities within each service. The overall maturity level is divided into four stages – Initial, core, extended and visionary. Highest level of maturity is achieved when capability level four is reached for all services.

2.2.4 Zero Trust Maturity Model

The Zero Trust Maturity Model was developed by The Cybersecurity and Infrastructure Security Agency (CISA) and suggests a path for organizations to transition to Zero Trust. The model was developed as a response to the Executive Order 14028 “Improving the Nation’s Cybersecurity” which embraces Zero Trust as the desired model for security (CISA, 2021).

Unfortunately, there is no formal documentation of how the model was designed and developed. Five domains are used to cover the areas to be assessed. The included domains are Identity, Device, Network/Environment, Application Workload, Data. Capabilities related to the categories “Visibility and analytics”, “Automation and orchestration", and "Governance" are suggested for each of the five domains. Maturity is measured in three stages – traditional, advanced, and optimal.

While the model aims to provide a holistic approach for adopting Zero Trust it is still very technology focused. Even though it extends beyond network centric applications, it lacks the adaptation of Zero Trust principles on processes and other organizational aspects. A summary of our maturity model comparison can be found in the table below (Table 1):

Requirement	Cybersecurity Capability Maturity Model (C2M2)	CTI-SOC2M2	ZeTuMM	Zero Trust Maturity Model
Design process	Initial draft developed by industry advisory group. Second draft based on expert feedback	Maturity model comparison and literature study	Maturity model comparison, literature study, and case study validation	N/A
Specific-Application purpose	Mainly descriptive	Descriptive	Descriptive	Descriptive and prescriptive
Levels and domains	4 levels of maturity and 10 domains	4 levels of maturity. 6 domains (SOC services)	3 levels of maturity and 15 domains	3 levels of maturity and 5 domains
Content	General information security capabilities	Specific domain within information security (CTI and SOC services)	General information security capabilities with Zero Trust as focal point	Information security capabilities related to Zero trust. Technology focused

Table 1: Maturity Model Comparison

2.3 Zero Trust

The literature reviewed suggests that creating perimeters and trusting everything and everyone on the inside is a practice ready for evaluation. This generosity of trust is described as a vulnerability that should be eliminated (Campbell, 2020). Ferretti et al. (Ferretti, Magnanini, Andreolini, & Colajanni, 2021) describes how focusing on security at the perimeter may result in a lack of network segmentation and access control on the internal network. If an attacker or malicious insider gains access to the network, this weakness can be used as an advantage as it enables lateral movement and more impactful attacks. Buck et al. argues that organizations using the perimeter approach are only as secure as their least secure device or application. Also, with increased remote work force and organizations moving their workloads to cloud service providers, defining a clear perimeter becomes a very challenging task (Buck, Olenberger , Schweizer, Fabiane, & Torsten, 2021).

Zero Trust is a security model that suggests security professionals stop granting trust implicitly, and instead provide access to resources based on continuous evaluation and verification. The mantra "never trust, always verify" (Buck, Olenberger , Schweizer, Fabiane, & Torsten, 2021) is often used to describe this way of thinking.

Kindervag introduces three fundamental principles of the Zero Trust model (Kindervag, 2010). The principles are removing all trust from your networks, introducing strict access control, and inspecting and logging all traffic. While Kindervag's proposed principles may be interpreted as network-centric, Microsoft introduces similar, more generally applicable principles. They suggest "Verify explicitly, use least privileged access, and assume breach" being the core of Zero Trust (Microsoft, 2019). Following is a description of these principles, and other components often mentioned as part of a Zero Trust architecture.

2.3.1 Network segmentation

One of the main guiding principles of Zero Trust is to always assume breach (Microsoft, 2019). Since all traffic must be assumed to be threat traffic until authorized, inspected and secured (Kindervag, 2010), each individual asset should have a protective perimeter around it. Having a protective perimeter around each individual resource is often referred to as micro-segmentation (Tyler & Viana, 2021). The ideal scenario would be to have fully distributed ingress/egress cloud micro-perimeters and deeper micro-segmentation (Microsoft, 2019), but this may not always be possible. Partial micro-segmentation is still superior to the traditional flat and wide-open internal network (CISA, 2021).

The proper use of micro-segmentation can contribute to ensuring that all lateral traffic within the network is authorized, inspected, and secured. However, micro-segmentation is not a standalone measure: To have a fully functional Zero Trust Architecture, other concepts such as identity governance and access policies are required (Buck, Olenberger , Schweizer, Fabiane, & Torsten, 2021).

2.3.2 Access Control & Least Privilege

The essence of access control is making resources available or restricted to the requester through the process of authentication and authorization. Kindervag introduced access control as a fundamental concept of Zero Trust, which our review also confirms (Kindervag, 2010).

The principle of least privilege is frequently mentioned in relation to access control. Buck et al. describes the principle as "access only granted to those resources required to perform functions" (Buck, Olenberger , Schweizer, Fabiane, & Torsten, 2021). Following the principle of least privilege means that just enough rights and permissions are given to users for them to be able to complete their tasks. One way of implementing the principle of least privilege is through access control lists (ACLs) to ensure devices in the network are only allowed to communicate with required resources (Tyler & Viana, 2021).

2.3.3 Dynamic Access Policies

Dynamic access policies allow dynamically granting access to a resource based on the context of the request. Access is accepted or denied based on different factors such as identity and credentials, device health, geographic location, previous access information, defined access policies, and more (Buck, Olenberger , Schweizer, Fabiane, & Torsten, 2021) (Teerakanok, Uehara, & Inomata, 2021).

Two components are central in this dynamic method of granting access: a policy decision point (PDP) and a policy enforcement point (PEP). When a subject requests access to an enterprise resource, the request goes to the PEP, serving as a gate between them. The policy enforcement point forwards the request to the PDP which collects contextual information and based on that, decides whether to accept or deny the request. The decision is sent back to the PEP which then establishes or terminates the connection between the subject and resource (Teerakanok, Uehara, & Inomata, 2021). The figure below (figure 1) describes how an untrusted request becomes trusted after going through the policy enforcement point and policy decision point (referred to as the Zero Trust engine in the figure).

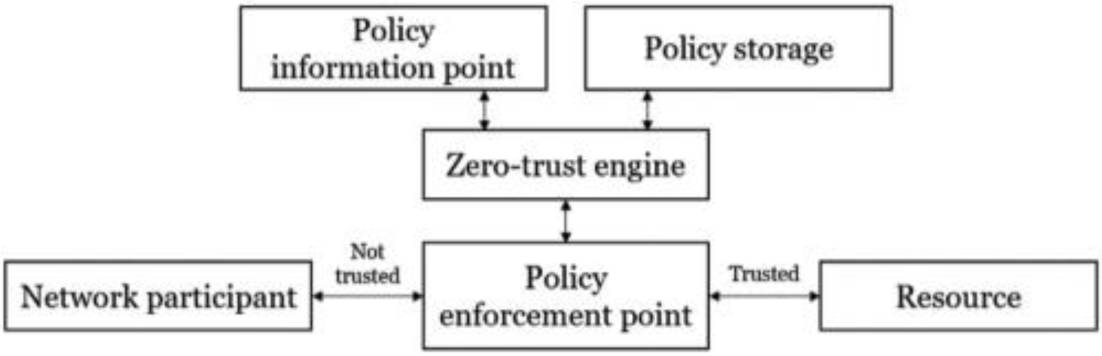

Figure 1: PeP and PDP Diagram

Retrieved from (Buck et al., 2021)

2.3.4 Monitoring and logging

Logging and inspection were introduced as the third fundamental concept of the Zero Trust (Kindervag, 2010). The enterprise should collect as much information as possible on the state of devices, network traffic and access requests (Rose, Borchert, Mitchell, & Connelly, 2020). This information can be used to improve the access policy decision (see 3.3.3), but also to detect both external and internal security threats. Sending logs to a security information event management solution will help detect suspicious behavior and respond to it at a quicker pace (Kindervag, 2010).

2.3.5 Multi-factor Authentication

Multi-factor authentication (MFA) is commonly used to improve the authentication process and ensure strict access control. MFA is the concept of using multiple factors to prove identity. The most common factor is username and password (D'Silva & Ambawade, 2021) also known as something you know, but other factors can be something you have, like a device or something you are, like your fingerprint.

2.3.6 Device Verification

To provide secure access to resources in the network, the security posture of the requesting devices must be verified and taken into consideration in the decision process. If the enterprise controls the device, it can make sure it has the latest security patches and anti-malware installed (Teerakanok, Uehara, & Inomata, 2021). Because of this, some companies only allow corporate owned or managed devices to access internal resources.

2.3.7 Threat protection

Threat intelligence is part of the supplementary information provided to the policy engine to aid in the access decision process (Teerakanok, Uehara, & Inomata, 2021). Information on malware, IP addresses and domains related to malicious activity are examples of information that could be used to alter the trust algorithm in the Policy decision point (Rose, Borchert, Mitchell, & Connelly, 2020).

2.3.8 Encryption

Assuming breach and removing all implicit trust zones means one should act as attackers are watching all network flow and communication on the network. This means communication should be treated as if it was leaving the enterprise network and going out on the internet. Encryption is therefore important to protect the confidentiality of the data (Rose, Borchert, Mitchell, & Connelly, 2020).

3 DESIGN SCIENCE RESEARCH

In this chapter we will describe our chosen research approach: The Design Science Research method used to develop the EZTMM. We will also describe the Systematic Literature Review process used to create the Knowledge Base used as a foundation upon which the EZTMM was developed.

3.1 Literature Review

Our systematic literature reviews were conducted following the search and filtering strategy as described by Kitchenham & Charters (Kitchenham & Charters, 2007). We chose this strategy because it was known to us from prior systematic literature reviews and is well suited for finding all relevant literature on a given topic. The method also allows for thorough documentation of the search. Furthermore, we chose a concept-centric approach as described by Webster & Watson (Webster & Watson, 2002) to break down Zero Trust into concepts which would serve as a starting point for our focus areas. This concept-centric approach also proved valuable when assessing which concepts were included in the various Zero Trust maturity models we analyzed during the literature review.

For the literature review on Zero Trust components conducted Autumn 2021, we selected Scopus, Web of Science and AIS eLibrary as our literature search databases based on recommendations from our thesis supervisor and our previous experience using these databases. We then developed search strings through several iterations, attempting to filter out articles of little relevance to our scope while including the relevant ones. Separate search strings for Zero Trust maturity models were created for each database with no results. This prompted a further literature review into maturity model theory at the start of Iteration 0. The final search strings for both Zero Trust Concepts and maturity models (not to be confused with the maturity model theory search strings in table 8) for each database can be seen in the two tables below (Tables 2 and 3):

Database	Search String	Number of results
Web of Science	TI=(Zero Trust) AND AB=(Zero Trust) AND DOP=(2020-01-01/2021-11-01)	22
Scopus	TITLE-ABS-KEY (zero AND trust) AND PUBYEAR > 2019 AND (LIMIT-TO (SUBJAREA , "COMP")) AND (LIMIT-TO (EXACTKEYWORD , "Network Security") OR LIMIT-TO (EXACTKEYWORD , "Zero Trust") OR LIMIT-TO (EXACTKEYWORD , "Access Control") OR LIMIT-TO (EXACTKEYWORD , "Network Architecture") OR LIMIT-TO (EXACTKEYWORD , "Continuous Authentications") OR LIMIT-TO (EXACTKEYWORD , "Dynamic Access Control") OR LIMIT-TO (EXACTKEYWORD , "Trust Modeling") OR LIMIT-TO (EXACTKEYWORD , "Micro-segmentation") OR LIMIT-TO (EXACTKEYWORD , "Zero Trust Security") OR LIMIT-TO (EXACTKEYWORD , "Zero Trust") OR LIMIT-TO (EXACTKEYWORD , "Access Control Mechanism") OR LIMIT-TO (EXACTKEYWORD , "Authentication Mechanisms") OR LIMIT-TO (EXACTKEYWORD , "Authorization") OR LIMIT-TO (EXACTKEYWORD , "End-to-end Security") OR LIMIT-TO (EXACTKEYWORD , "Identity Authentication") OR LIMIT-TO (EXACTKEYWORD , "Identity Management") OR LIMIT-TO (EXACTKEYWORD , "Security Architecture") OR LIMIT-TO (EXACTKEYWORD , "Security Model") OR LIMIT-TO (EXACTKEYWORD , "Trust Frameworks") OR LIMIT-TO (EXACTKEYWORD , "Trust Management") OR LIMIT-TO (EXACTKEYWORD , "Zero Trust Architecture") OR LIMIT-TO (EXACTKEYWORD , "Architecture") OR LIMIT-TO (EXACTKEYWORD , "Boundary Protection") OR LIMIT-TO (EXACTKEYWORD , "Computer Architecture") OR LIMIT-TO (EXACTKEYWORD , "Computer System Firewalls") OR LIMIT-TO (EXACTKEYWORD , "Continuous Authentication"))	75
AIS eLibrary	title:(Zero Trust) OR abstract:(Zero Trust)	1

Table 2: Zero Trust Concept Search Strings

Database	Search String	Results
Web of Science	ALL=("Zero Trust") AND ALL=("maturity model")	0
Scopus	(TITLE-ABS-KEY ("Zero Trust") AND TITLE-ABS-KEY ("maturity model"))	0
AIS eLibrary	"Zero Trust" AND "maturity model"	0

Table 3: Zero Trust Maturity Model Search Strings

All available resulting articles from the above search strings were downloaded. The resulting articles then underwent a four-phase filtering process:

1. Duplicates were removed.
2. Inclusion and exclusion criteria based on our research questions and the scope of the literature review were then used to filter out the less relevant articles.
3. The remaining articles were red thoroughly and a decision was made on whether to include them or not based on their relevance for our scope.
4. Recommendations from our supervisor were added along with relevant articles from a backwards search.

Our inclusion and exclusion criteria can be found in the table below (Table 4):

Inclusion	Exclusion
Must be found through the defined search criteria (including backwards search) in the defined databases or provided/approved by our supervisor	No authors
Must be peer reviewed	Languages other than Norwegian and English
Must be either conference proceedings, journal articles or books	Older than 2009 for RQ1
Must be freely available	Older than 2020 for RQ2 and RQ3
Zero Trust must be the main topic of the article	Paid articles
	Unavailable articles

Table 4: Inclusion and Exclusion Criteria for Literature Review Articles

A total of 54 articles were found to not meet our inclusion and exclusion criteria. The 18 articles that remained after the filtering process were then used to create our systematic literature review on Zero Trust: We analyzed each article, creating a concept matrix with an overview of which Zero Trust concepts each article discussed. We then used this concept matrix to identify the core concepts of Zero Trust and described each concept in detail based on the knowledge from the articles. A complete list of the included articles and the concepts they cover can be seen in the concept matrix below (Table 5):

Article	Category	Zero Trust Concepts								Maturity Models	Implementation	
	Concept	Network segmentation	Access Control & least privilege	Dynamic access policies	Monitoring and logging	MFA	Device verification	Threat protection	Encryption		Case Studies	Recommendations
Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero-trust		X	X	X	X	X	X		X			
Survivable zero trust for cloud computing environments			X	X		X	X		X			
Federating trust: Network orchestration for cross-boundary zero trust				X			X		X			
Zero trust: Never trust, always verify			X	X			X					
Building A Zero Trust Architecture Using Kubernetes		X	X	X	X	X	X		X			
Zero-Trust Principles for Legacy Components: 12 Rules for Legacy Devices: An Antidote to Chaos		X	X		X		X	X	X			X
Zero Trust in the Context of the Utility Industry		X		X	X							
Survey on Zero-Trust Network Security			X	X	X		X	X				
Beyond Zero Trust: Trust Is a Vulnerability			X	X							X	X
The zero trust supply chain: Managing supply chain risk in the absence of trust			X	X	X							
Migrating to Zero Trust Architecture: Reviews and Challenges			X	X	X	X	X	X	X		X	X
Trust No One? A Framework for Assisting Healthcare Organisations in Transitioning to a Zero-Trust Network Architecture		X	X		X	X			X			X
Zero Trust Architecture		X	X	X	X	X	X	X	X			
Build Security Into Your Network's DNA: The Zero Trust Network Architecture		X	X	X	X				X			
No More Chewy Centers Introducing The Zero Trust Model Of Information Security			X	X	X				X			X
CISA Zero Trust Maturity Model_Draft		X	X	X	X	X	X	X	X	X		X
Microsoft Zero Trust Maturity Model		X	X	X	X	X	X	X	X	X		X
Palo Alto Networks Zero Trust Maturity Model										X		

Table 5: Zero Trust Concept Matrix

To identify relevant literature on maturity model theory, we used the following keywords in search strings on Scopus:

1. Maturity
2. Model
3. Design
4. Principles
5. Theory

The following criteria was used for articles related to maturity model theory (Table 6):

Inclusion	Exclusion
Title and abstract must appear relevant for our literature review on maturity models	No authors
Must be peer reviewed	Languages other than Norwegian and English
Must be either conference proceedings, journal articles or books	Paid articles
Must be freely available	Unavailable articles
Must be found through the defined search criteria (including backwards search) in the defined databases or provided/approved by our supervisor	Duplicates

Table 6: Inclusion and Exclusion Criteria for Maturity Model Theory

Articles were then chosen from search results based on relevance in headers and abstracts.

We removed duplicate articles and ones that did not match out inclusion criteria. After filtering, we ended up with six articles that were thoroughly read and reviewed. Additionally, the article “Design Science in Information Systems Research” by Hevner et al. was recommended to us by our supervisor. The table below shows the identified articles on maturity model theory (Table 7).

Title	Author
Maturity models in business process management	Röglinger, Maximilian; Pöppelbuß, Jens; Becker, Jörg
What makes a useful maturity model? A framework of general design principles for maturity models and its demonstration in business process management	Röglinger, Maximilian; Pöppelbuß, Jens;
Maturity levels for interoperability in digital government	Petter Gottschalk
A set theoretical approach to maturity models: Guidelines and demonstration	Lasrado, Lester Allan; Vatrapu, Ravi; Andersen, Kim Normann;
Developing maturity models for IT Management	Pöppelbuß, Jens; Jörg Becker; Ralf Knackstedt;
Maturity models in IS research	Niehaves, Björn; Pöppelbuß, Jens; Jörg Becker; Simons, Alexander
Design Science in Information Systems Research	Hevner, Alan R; March, Salvatore T; Park, Jinsoo;

Table 7: Articles on Maturity Model Theory

To identify existing maturity models for the comparison we used the following two search strings on Scopus (Table 8):

Search string	Number of results
(TITLE-ABS-KEY (cyber AND security) AND TITLE-ABS-KEY (maturity AND model) AND TITLE-ABS-KEY (design)) AND PUBYEAR > 2014	20
(TITLE-ABS-KEY (cybersecurity) AND TITLE-ABS-KEY (capability AND maturity AND model)) AND PUBYEAR > 2014	36

Table 8: Maturity model Theory Search Strings

Our goal was to identify articles describing maturity models used to assess information security capabilities. There were specifically two criteria for inclusion. The selected articles should preferably include scientific documentation of the design process. Second, the maturity models described had to be related to information security, and preferably Zero trust. As the comparison was partly performed to inspire the contents of our maturity model, we decided to limit results to articles no older than 2015. Due to the same reason we wanted to include one or more maturity models developed by special interest groups or organizations highly regarded within information security.

Based on these criteria we ended up with the following remaining articles (Table 9):

Title	Author
CTI-SOC2M2 – The quest for mature, intelligence-driven security operations and incident response capabilities: CTI-driven SOC capability maturity model	Schlette, Daniel; Vielberth, Manfred; Pernul, Günther
(CSM2-RA-R2-TI): Cyber Security Maturity Model for Risk Assessment Using Risk Towards a maturity model for health-care cloud security (M2HCS)	Lakshmi Prasanna B; Reddy, M. Saidi Akinsanya, Opeoluwa Ore ;Papadaki, Maria; Sun, Lingfen
Secure design and development cybersecurity capability maturity model (SD2-C2M2): Next-generation cyber resilience by design	Gourisetti, Sri Nikhil Gupta; Mix, Scott; Mylrea, Michael; Bonebrake, Christopher; Touhiduzzaman, Md
Incorporating Systems Thinking into a Cyber Resilience Maturity Model	Shaked, A., Tabansky, L., Reich, Y.
Evaluating and improving cybersecurity capabilities of the energy critical infrastructure	Curtis, P.D., Mehravari, N.

Table 9: Maturity Model Theory Articles

After reading each article we decided to include the models CTI-SOC2M2 and C2M2 (From the article “Evaluating and improving cybersecurity capabilities of the energy critical infrastructure”). Neither of the models were related to Zero Trust, so we turned to Google Scholar with the search term “Zero Trust Maturity Model”. The top result was the research paper “Zero Trust Maturity Matters: Modeling Cyber Security Focus Areas and Maturity Levels in the Zero Trust Principle” by M.G. Modderkolk. This is a master thesis describing the development of a Zero Trust maturity model making it a necessity in the comparison. We wanted to include at least two models related to Zero Trust and decided to include the Zero Trust Maturity Model published by CISA. This was identified during our first round of reviewing literature.

3.2 Design Process

Design science is the chosen research approach for this master thesis. It is an approach revolving around creating unique and problem-solving artifacts. IT artifacts are broadly defined as constructs (vocabulary and symbols”, models (abstractions and representations), methods (algorithms and practices), and instantiations (implemented and prototype systems) (March & Smith, 1995). These artifacts are supposed to improve IT practitioners' ability to better understand and implement information systems successfully (Hevner, Ram, March, & Park, 2004). Problem solving solutions are created through an iterative build and evaluation process. This thesis proposes the development of a maturity model as a unique artifact to help organizations understand their current security posture, and identity future actions to improve their capabilities.

Research by Hevner et al. (Hevner, Ram, March, & Park, 2004) proposes a framework for IS research, along with guidelines for performing design science. The framework consists of three main components: An environment, IS research (In figure 2, known as "Designing the EZTMM") and a knowledge base (Figure 2).

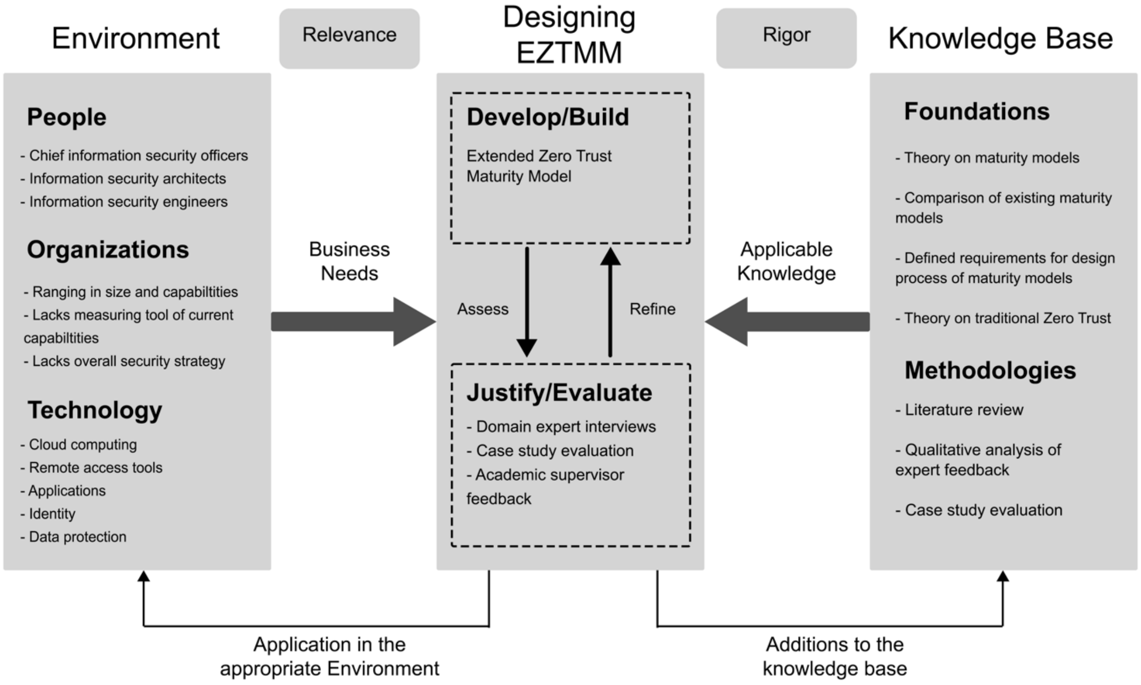
Environment
Relevance
Designing EZTMM
Rigor
Knowledge Base
People
- Chief information security officers
- Information security architects
- Information security engineers
Organizations
- Ranging in size and capabiltities
- Lacks measuring tool of current capabiltities
- Lacks overall security strategy
Technology
- Cloud computing
- Remote access tools
- Applications
- Identity
- Data protection
Business Needs
Develop/Build
Extended Zero Trust Maturity Model
Assess
Refine
Justify/Evaluate
- Domain expert interviews
- Case study evaluation
- Academic supervisor feedback
Applicable Knowledge
Foundations
- Theory on maturity models
- Comparison of existing maturity models
- Defined requirements for design process of maturity models
- Theory on traditional Zero Trust
Methodologies
- Literature review
- Qualitative analysis of expert feedback
- Case study evaluation
Application in the appropriate Environment
Additions to the knowledge base

Figure 2: EZTMM design research framework

Retrieved from Hevner et al. (2004)

The environment consists of people, organizations, and technology. This is the problem space where business needs are defined. Business needs are defined by goals, problems, tasks, and opportunities and how people in the organization consider each of them. Business needs are then evaluated from an organizational perspective taking strategy, culture, and processes into consideration. They are then aligned with the infrastructure, applications, communications architecture, and development capabilities that already exist in the organization. Creating an artifact based on an articulated business need assures relevance of the research.

In our research we considered the environment to be organizations working to secure their business by measuring capabilities and improving information security. The maturity of their capabilities may range from just starting out to being highly experienced in building defensible architectures. The roles in the environment may be technical personnel such as security engineers and architects, but also managerial positions such as chief information security officers and similar. What the stakeholders have in common is shared interest and an active role in protecting the business from security incidents. Increased remote work and new technology trends such as cloud computing are among the stakeholders' challenges. These are both strategic directions shared by many organizations. These technology trends along with an advanced threat landscape create a need for an improved overall information security strategy. Our proposed solution is using Zero Trust as the strategy and offer a way to measure relevant capabilities and create action plans. During our literature review we could not identify any maturity models that would assess and improve an organizations' overall information security by using Zero Trust as the core strategy. Technologies like remote access tools, cloud computing, identity management and data protection are very relevant in our environment. The proposed model is therefore of extra relevance for organizations involved with such technologies.

The IS Research component (Designing the EZTMM in figure 2) of the framework consists of two activities, Develop/Build and Justify/Evaluate. This aligns with suggestions by Mark and Smith (March & Smith, 1995) which argue that research activities in design science are split into two parts: build and evaluate. "Build" is the activity of creating a unique artifact while "Evaluate" consists of developing criteria and then measuring against these too see how well the artifact performs. When an artifact has been developed, it must be decided if it works, and it must be evaluated scientifically. Mathematical evaluation may be appropriate in some types of research, while empirical and qualitative methods may be suitable in others.

The artifact developed as part of this research is a maturity model. The goal for designing the model was to help organizations address the business need for a new and improved information security strategy. The evaluation of the artifact was based on a qualitative approach. Multiple iterations of semi-structured and unstructured interviews with subject matter experts served as the main form of evaluation. These interviews were used to gather data on the artifact's usability, relevance, and overall quality. Feedback from the interviews was reviewed, considered, and implemented in the next iteration of the build phase. Co-creating the model with highly experienced domain experts ensured relevance in our research and greater level of correctness in the model. A case study was used as another form of evaluation towards the end of the design process. In this phase we aimed to put the artifact into the context it was intended for to see how well it performed. This phase was conducted in cooperation with an organization that used the maturity model to assess their current extended Zero Trust capabilities. This revealed the model's efficiency, usability, completeness, and fit in the organization.

The knowledge base heavily impacted how the artifact was built and evaluated. It consisted of knowledge on existing artifacts and prior research. For this study, knowledge on existing maturity models and how they were developed was of importance. Making the content of the model usable and efficient also required knowledge on Zero Trust. These topics were addressed in a literature study, serving as the foundation for our knowledge base. Each phase of building and evaluating added new knowledge to the base, which again sparked a need for changes to the artifact.

Hevner et al. (Hevner, Ram, March, & Park, 2004) also propose seven guidelines to be used by researchers to enable more effective design science research. While leveraging the framework and adhering to these guidelines, we did not consider them the main criteria of our research. Instead, we focused on following a set of requirements defined in the research paper "Developing Maturity Models for IT Management – A Procedure Model and its Application" (Becker, Knackstedt, & Pöppelbuß, 2009). The seven guidelines presented by Hevner et al. served as the main foundation for the requirements. However, the guidelines have been interpreted and modified to better suit the development of maturity models and address points of criticism (Becker, Knackstedt, & Pöppelbuß, 2009).

The first requirement is that a comparison of existing maturity models must be performed. Knowing what maturity models already exist will uncover if there is a need for the proposed model. There is a chance that models like the proposal already exist, making the research redundant. However, the comparison may be used to draw inspiration from, or discover existing models that can be improved. To satisfy the requirement, a literature study on maturity model theory was conducted as part of this research, including a comparison of existing maturity models. The comparison included models which goal is to assess and improve organizations information security capabilities. Inclusion of existing models related to Zero Trust was highly desirable. For models to be included in the comparison, documentation of their development would preferably be obtainable. However, we decided that this requirement could be omitted if considered necessary. Models were compared on their content, structure, and development methodology. The literature study also served as foundation for creating the first iteration of the model.

The second requirement is having an iterative procedure. We complied with this requirement by developing the model in numerous iterations. The time frame for developing the artifact was limited to the duration of a master-thesis, meaning the number of iterations was also limited. However, a total of five iterations were completed by the end of our research.

During each iteration, the EZTMM was evaluated by the researchers along with security professionals. After completion of each iteration, it had to be decided if the model was accepted or if another iteration was required. This evaluation is the third requirement in the design process. For each iteration, the respondents (subject matter experts and practitioners) were asked if the model was satisfactory in terms of overall quality, usability, and efficiency. We also asked for additional feedback and if they saw any need for further changes. When consensus of satisfactory quality was reached among the respondents, and no further changes were imposed, the second phase of evaluation could start. This phase involved the performance of a case study where the model was tested in a respondent organization. The participants carried out the assessment on their own, while we observed and took notes on how the model was used. These observations, as well as feedback from the respondent performing the assessment, provided data used to improve the model in one last iteration. After these modifications, the model was compared against the acceptance criteria listed in the table below. It is worth noting that these requirements evolved throughout the research, in accordance with the growing knowledge base. The requirements in the table are the total requirements, including both the respondent acceptance requirements after each iteration and the other requirements tested during the case study.

Acceptance Criteria	Rationale
Applicable in a broad range of organizations	For a maturity model to be adopted and considered useful by organizations, it needs to be applicable. Therefore, a broadly applicable model is a prerequisite for broad adoption.
Can be used to self-assess current Zero Trust Maturity	The main purpose of our maturity model is that it can be used for assessing organizations' Zero Trust maturity level.
Can be used to improve overall information security capabilities	An important use for maturity models is to assess the organization's maturity and based on this assessment identify improvements, low-hanging fruits and plan further implementation. The model therefore needs to be able to facilitate this.
Applies Zero Trust principles to both technical and organizational domains within information security	The main knowledge gap identified in our systematic literature review was that existing models do not take organizational domains sufficiently into account. This model was intended as a possible solution.
Security controls suggested in the model are placed at appropriate levels and described sufficiently and correctly	The model's usefulness depends on it being correct. Controls placed at the wrong level lowers the usability and credibility of the model.
Respondents being presented the model show desire to leverage the model in their own organization	Respondents showing interest in using the model gives an indication of relevance and potential value provided by the research.
Respondents have few or no additional suggestions for changes when being presented the latest draft of the model.	Receiving few suggestions for changes indicate that consensus is reached among the respondents, and that the model has reached an acceptable level of correctness.

Table 10: Acceptance Criteria for the EZTMM

The fourth requirement is having a multi-methodological procedure. This requirement is fulfilled through the use of literature reviews, semi-structured and unstructured interviews, and a case study.

The fifth and sixth requirement is having a problem definition and demonstrating its relevance. The problem definition was summarized as the following: "Developing a maturity model to assess current capabilities and prioritize future actions to adopt the Zero Trust principles and improve an organizations' information security in both technical and organizational areas". It is important to note that this problem definition describes the prospective application of the maturity model and is not related to the research questions addressed in this thesis. Problem relevance is described earlier in this chapter but was also addressed as part of the semi-structured interviews. Respondents were asked to describe their perception of Zero trust relevance, and if they saw any use for the proposed model.

A targeted presentation of the maturity model is the next requirement. It is stated that the presentation must meet the needs of the users and the conditions of the model's application. For this reason, we decided to present the model in two parts. The main part, a document describing the model in its entirety. This document includes an introduction, all the different focus areas to be assessed and control questions. Organizations may use this document to guide the assessment process but also to plan future actions to improve their capabilities. The document was meant to fulfill the model's prescriptive application purpose. In addition, we developed an excel spreadsheet for organizations to self-assess their capabilities, fulfilling the model's descriptive application purpose. The spreadsheet includes a collection of control questions and automatically calculates a maturity score for each focus area as well as an overall maturity score when filled out.

The last requirement is scientifically documenting every part of the development process. We complied with this by documenting how every step described in this chapter was carried out. The documentation includes which parties were involved in the different parts, which decisions were made and why. Furthermore, we have documented the changes made to the model itself for each iteration (attached in appendices E and G).

3.3 Data Collection

The data collection and analysis for our maturity model was conducted during January-May 2022 and divided into four iterations. In this chapter we will detail who our respondents were, which data collection methods were used and how the data was analyzed.

3.3.1 The Respondents

We interviewed nine respondents from six different organizations. The respondents were all security professionals, ranging from technical specialists to security managers and Chief Information Security Officers. All our respondents had at least five years of experience working within the field of information security, with some having as much as 30 years' experience, which is significant given that the shift in focus towards information security is a recent phenomenon. The respondents were recruited from our own professional networks and recommendations from our thesis supervisor. An overview of our respondents by role can be found below (Table 11):

Role	Number of respondents
CISO	3
Manager	4
Technical Specialist	2

Table 11: Overview of Respondents by Role

3.3.2 Interviews

We conducted two types of interviews and one workshop with our respondents: One-hour semi-structured interviews based on the included interview guide in appendix A were used for the initial feedback on EZTMM. 30-minute unstructured interviews where the respondents could highlight topics and feedback within their field of expertise were used for returning respondents providing additional feedback to the model. In addition to this, a one-hour workshop where we presented Zero Trust for a group of ten security professionals and gathered feedback afterwards was arranged as a part of the initial feedback gathering in iteration 1. As depicted by table 12 below, we conducted a total of one workshop and 10 interviews, where one was a group interview:

Interview type	Total Number of interviews	Group interviews
One-hour semi-structured	6	1
30 minute unstructured	4	0
Workshop	1	1

Table 12: Interview Overview

3.3.3 Case study evaluation

In addition to the interviews, we conducted a short case study evaluation, where we asked one of our respondent organizations to use the EZTMM to conduct a Zero Trust Maturity Evaluation. We instructed the respondent organization to conduct the maturity assessment using the accompanying maturity assessment spreadsheet and refer to the model itself for further information if needed. We worked closely with the organization doing the evaluation, answering questions, collecting feedback, discussing maturity levels and observing how our model was used.

The main purpose of this case study evaluation was twofold:

1. Gather feedback on the maturity assessment spreadsheet that accompanies the model, as well as the control questions for each maturity level
2. Verify that the model meets our meta requirements

3.4 Data Analysis

In this chapter we will describe how the data gathered from both the interviews and the case study evaluation was analyzed, and which inclusion and exclusion criteria were used to determine whether the feedback was implemented or not.

3.4.1 Interview and written feedback

Our respondents provided us with two types of feedback during and after the interviews: Written feedback and oral feedback. The written feedback was converted into comments that were then added to our "live" version of EZTMM. The interviews were transcribed, with feedback being extracted from these transcriptions and input as comments into our "live" document along with the written feedback. The feedback we received was sorted into three categories:

1. Structure Improvements
2. Corrections and Consistency
3. Focus Area Improvements/Suggestions

A complete list of all the feedback we received on the model sorted by category and iteration can be found in appendix D.

3.4.2 Inclusion and Exclusion Criteria

When all the feedback from our interviews had been imported into the "Live" document through comments, we went through each comment and decided on whether to implement them or not based on our scope and the following exclusion and inclusion criteria (Table 13):

Inclusion	Exclusion
Improves clarity of the text or removes ambiguity	Duplicate feedback
Corrects mistakes	Feedback already implemented elsewhere
Adds context or additional insight to the discussed focus area or measure	Outside of Scope
For larger proposed changes or new sections: Adheres to scope and is an important or impactful addition to the EZTMM	The proposed section or change is challenging or impossible to relate to Zero Trust principles

Table 13: Inclusion and Exclusion Criteria for interview feedback

Once a comment was determined to pass the inclusion and exclusion criteria, one of the authors was assigned to implement the comment and the other to review and verify the implementation. A color-coded version of the EZTMM with all changes by iteration can be found in appendices E (original file download) and G.

3.4.3 Case study evaluation

The observations made during the case study evaluation were converted to a similar format as the feedback from the interviews: Any questions asked by the respondent organization and their answers were input as comments, so that the model could be clarified to better address the question if needed. Any difficulties that the respondent organization experienced using the model were noted and converted to improvement suggestions. The discussions regarding maturity levels were also recorded and converted into comments. Lastly, all these comments were evaluated using the same inclusion and exclusion criteria detailed in the previous chapter before being implemented.

Furthermore, we included a short interview with the practitioner doing the maturity assessment using the EZTMM where we attempted to validate that the model met our defined acceptance criteria. The complete case study notes including the interview can be found in appendix B.

Lastly, the improved model was evaluated against our acceptance criteria defined in chapter 3.1.2. The evaluation against our acceptance criteria can be found in the discussion. The final version of the EZTMM can be found in appendices E (original file download) and F.

4 DESIGNING THE EXTENDED ZERO TRUST MATURITY MODEL

In this section, we will describe the process used when designing the Extended Zero Trust Maturity Model. The design process was performed in four iterations: The literature review and existing model comparison, the initial interviews, the second interviews and the case study validation. An overview of the iterations can be found below (Figure 3), and each iteration is described in detail in the following chapters.

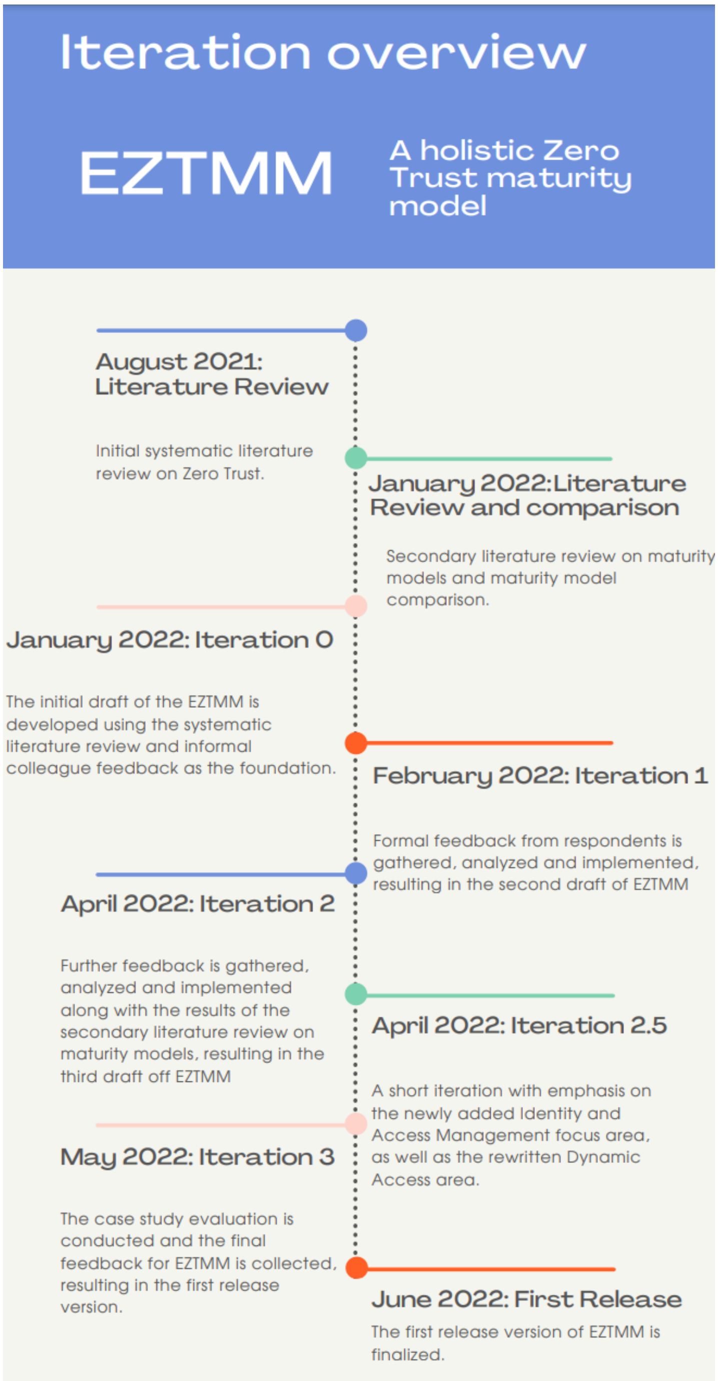

Iteration overview

EZTMM

A holistic Zero Trust maturity model

August 2021: Literature Review

Initial systematic literature review on Zero Trust.

January 2022:Literature Review and comparison

Secondary literature review on maturity models and maturity model comparison.

January 2022: Iteration 0

The initial draft of the EZTMM is developed using the systematic literature review and informal colleague feedback as the foundation.

February 2022: Iteration 1

Formal feedback from respondents is gathered, analyzed and implemented, resulting in the second draft of EZTMM

April 2022: Iteration 2

Further feedback is gathered, analyzed and implemented along with the results of the secondary literature review on maturity models, resulting in the third draft off EZTMM

April 2022: Iteration 2.5

A short iteration with emphasis on the newly added Identity and Access Management focus area, as well as the rewritten Dynamic Access area.

May 2022: Iteration 3

The case study evaluation is conducted and the final feedback for EZTMM is collected, resulting in the first release version.

June 2022: First Release

The first release version of EZTMM is finalized.

Figure 3: EZTMM Iteration Overview

4.1 Iteration 0: Literature Review and existing model comparison

The foundation of our Extended Zero Trust Maturity Model was a systematic literature review conducted in Autumn 2021. In this literature review, we identified the core concepts of a Zero Trust architecture and compared existing Zero Trust maturity models. During this work, we found that Zero Trust principles are mainly applied to the technology domain in existing maturity models. We wanted to create a more holistic maturity model, considering an organization's processes and people, as well as its technology.

Due to the lack of academic articles on maturity models in the systematic literature review conducted in Autumn 2021, we also conducted another systematic literature review into maturity model theory following the process described in chapter 3.1.1. The second literature review was completed in January 2022.

We then based our initial draft on the core concepts of a Zero Trust architecture identified in our literature review and the information found in our maturity model theory literature review. The core Zero Trust principles were applied to organizational aspects (from the end of iteration 1 known as the processes and people domains) of information security that were not traditionally associated with Zero Trust. This was our first attempt at creating a more holistic Zero Trust maturity model.

The result of this work was compared against prominent security frameworks and models such as ISO 27001, NSMs grunnprinsipper for IKT-sikkerhet and CIS Controls. The goal of this comparison was to identify gaps in our own maturity model and get an understanding of existing security controls. Based on this understanding, we defined three maturity levels for each focus area, along with control questions the organization can use to assess their current maturity level. The maturity levels were inspired by the levels defined by CISA in their Zero Trust Maturity Model. We then informally asked our closest colleagues for feedback and suggestions related to the focus areas they were experts and practitioners in. This became our initial draft of the Extended Zero Trust Maturity Model.

The initial draft of the EZTMM included a total 72 control questions within 10 focus areas. Iteration 0 ended on the 13th of February 2022.

4.2 Iteration 1: Gathering and implementation of respondent feedback on the initial EZTMM draft: Restructuring the model

The initial draft was then sent out for review to our respondent organizations. A total of eight respondents from five different organizations were included in the first feedback round. All the respondents were security professionals, ranging from technical specialists and engineers to managers and CISOs. The goal of this iteration was to improve all aspects of our model, including restructuring the model, adding new focus areas and rearranging maturity level requirements and control questions.

After a two-week review period, we received written feedback on our initial draft and arranged one-hour interviews with all the respondents following a structured interview guide to obtain more feedback. We also arranged a presentation and workshop in one of our respondent organizations with over ten security professionals. During the workshop we presented the topic of Zero Trust while introducing our own maturity model and gathering feedback.

The feedback on the initial draft could generally be classified in three categories:

1. Structure Improvements
2. Corrections and Consistency
3. Focus Area Improvements and Suggestions

All the comments and feedback gathered from each iteration can be found in appendix C.

The first type of feedback was new focus area suggestions:

> Backup and restore is a crown jewel in most organizations, this should be emphasized further in the model. Perhaps backup and restore should have its own focus area? Immutable backups are a good example of Zero Trust principles being applied to backups.

In this case we decided that an increased emphasis on backups in the Data Governance Focus Area would be sufficient. We received similar comments on out-of-band-communications:

> How about collaboration tools? If an attacker has compromised your AD, how do you communicate? The attacker is likely to have access to your collaboration tools such as email, teams/slack/confluence. Perhaps establishing out-of-band collaboration solutions.

Similarly to the first suggestion, this topic was included and emphasized in our existing Incident Management Focus Area.

Corrections were also common, especially among the technical specialists. These were generally minor corrections such as “Add also VRFs. VRFs play a key role at route-segmentation level.”, most of which were implemented. There were also comments related to the consistency of the writing, all of which were implemented:

> Think about the style in which the maturity levels are defined. Now, at least in many requirements, the text describes what an organization "should" do to be in specific maturity. At least some maturity models would describe what an organization "is" doing when they are at specific level. Think for instance risk management: maturity 1) risk assessments are conducted on ad-hoc basis; 2) risk assessments are planned, systematic and periodic; 3) risk assessment results are used to guide organizations future activities

Lastly, we received a lot of comments regarding the structure of the model. Many respondents commented that they struggled to see the connections between our focus areas: “When I read these focus areas, I see a good mix without seeing the connection between them”. Based on this feedback, we restructured our model by adding three domains: Technology, Processes and People. These domains allowed us to easily categorize each focus area and highlight the importance of a holistic approach to Zero Trust.

During iteration 1 we made a lot of big changes: We restructured the model and introduced the three domains. We also implemented several corrections, improved the consistency of the writing, and put more emphasis on important topics such as backups, using a risk-based approach and out-of-band communications. The result of implementing the feedback from our respondents was the second draft of the Extended Zero Trust Maturity Model.

The second draft of the EZTMM included a total of 85 control questions within 10 focus areas and three domains. 13 new control questions were added, and three questions were revised based on respondent feedback. Iteration 1 ended the 31st of March 2022.

4.3 Iteration 2: Gathering and implementation of respondent feedback on the second EZTMM draft: Polishing the model

The second draft of our maturity model was then sent out to all the respondents from the previous round of feedback as well as three new respondents from three new organizations. Any changes made from the initial draft were marked with yellow, saving the respondents who participated in the first round of feedback time and allowing us to decrease the review window to one week. The goal of this iteration was to improve and polish most aspects of the model, making sure everything was clear and correct without making large changes.

After receiving more written feedback at the end of the review window, we scheduled 30-minute follow-up interviews based on the written feedback from each respondent. The suggestions collected from the written feedback and follow-up interviews were inserted as comments and evaluated using the same process as described in iteration 1.

There were three recurring themes in the iteration 2 feedback:

1. Insufficient coverage of identities
2. Negative experiences with phishing simulation tools
3. Insufficient focus on the Zero Trust principles in the focus areas under the people domain

The lack of focus on identity in our framework was pointed out by multiple respondents: “I miss a bit more about IAM here? Maybe particularly identity governance and tying authentication to other tools such as EDR and MDM?” (This comment was translated from Norwegian) Due to this feedback, we decided to add a new focus area dedicated to the governance of identities: Identity and Access Management. Furthermore, we rewrote the Dynamic Access section under technology to emphasize more clearly how identities are used to determine access. We were not planning on introducing such big changes so late in the model but felt that this feedback was very accurate and could not be ignored.

Several of the respondents this time around also had bad experiences with phishing simulation tools: “I have strong opinions about the “usefulness” of these tools…reach out if you want to hear why I mean they are indeed useless if not even damaging the security culture of an organization.” Based on our respondents’ negative experiences with phishing, we decided to do further research on the topic. As a result, phishing simulators were removed as a maturity level requirement in our model based on recent findings that suggest such tools can have negative impacts (Lain, Kostiainen, & Capkun, 2021).

Lastly, the focus areas in the people domain were rewritten, increasing the emphasis on Zero Trust principles. This was done to address comments such as: “Where is zero trust here? What you describe makes sense also in terms of maturity level, but where is the “assuming compromise” or other “ideas” of zero trust?” The result of this process is the third draft of the Extended Zero Trust Maturity model.

The third draft of the EZTMM included a total of 104 control questions within 11 focus areas and three domains. 19 new control questions were added and 13 were revised based on respondent feedback. Iteration 2 ended on the 24th of April 2022.

4.4 Iteration 2.5: Expert feedback on Identity

To ensure the quality of the newly added focus area, we performed an additional short iteration, where we asked two of our subject matter experts for written feedback on the new focus area and the rewritten dynamic access area in particular.

The feedback was overall positive, validating that we had captured the essence of identity in Zero Trust. In addition to receiving formal acceptance for the newly added and rewritten focus areas, we received some minor suggested additions along with some corrections and comments on the consistency of the writing: “Mention something about assigning rights based on the role extracted from the HR system” (This comment was translated from Norwegian) and “Be consistent…either short version first and then explanation…or the opposite…as you have it for CSF”. These comments were addressed and implemented prior to conducting the case study evaluation. Iteration 2.5 ended on the 30th of April 2022.

4.5 Iteration 3: Case study evaluation

In the third iteration, we conducted a case study review of our maturity model by asking one of our respondent organizations to perform a maturity review using our model. We were tightly involved in the maturity review process, gathering data on how the model was used as well as direct feedback from the people using it to assess their organization’s maturity. After the evaluation had been performed, we interviewed the person who conducted the maturity assessment about the use of the spreadsheet and recorded the answers. The full case study notes including the post-evaluation interview are attached in appendix B.

The evaluation was done using our EZTMM Evaluation Sheet. This spreadsheet was developed as an evaluation tool based on the EZTMM and includes every control question introduced in the model in a color-coded format based on which domain the question belongs to. Below is an example screenshot including all control questions for one focus area from each domain (Figure 4):

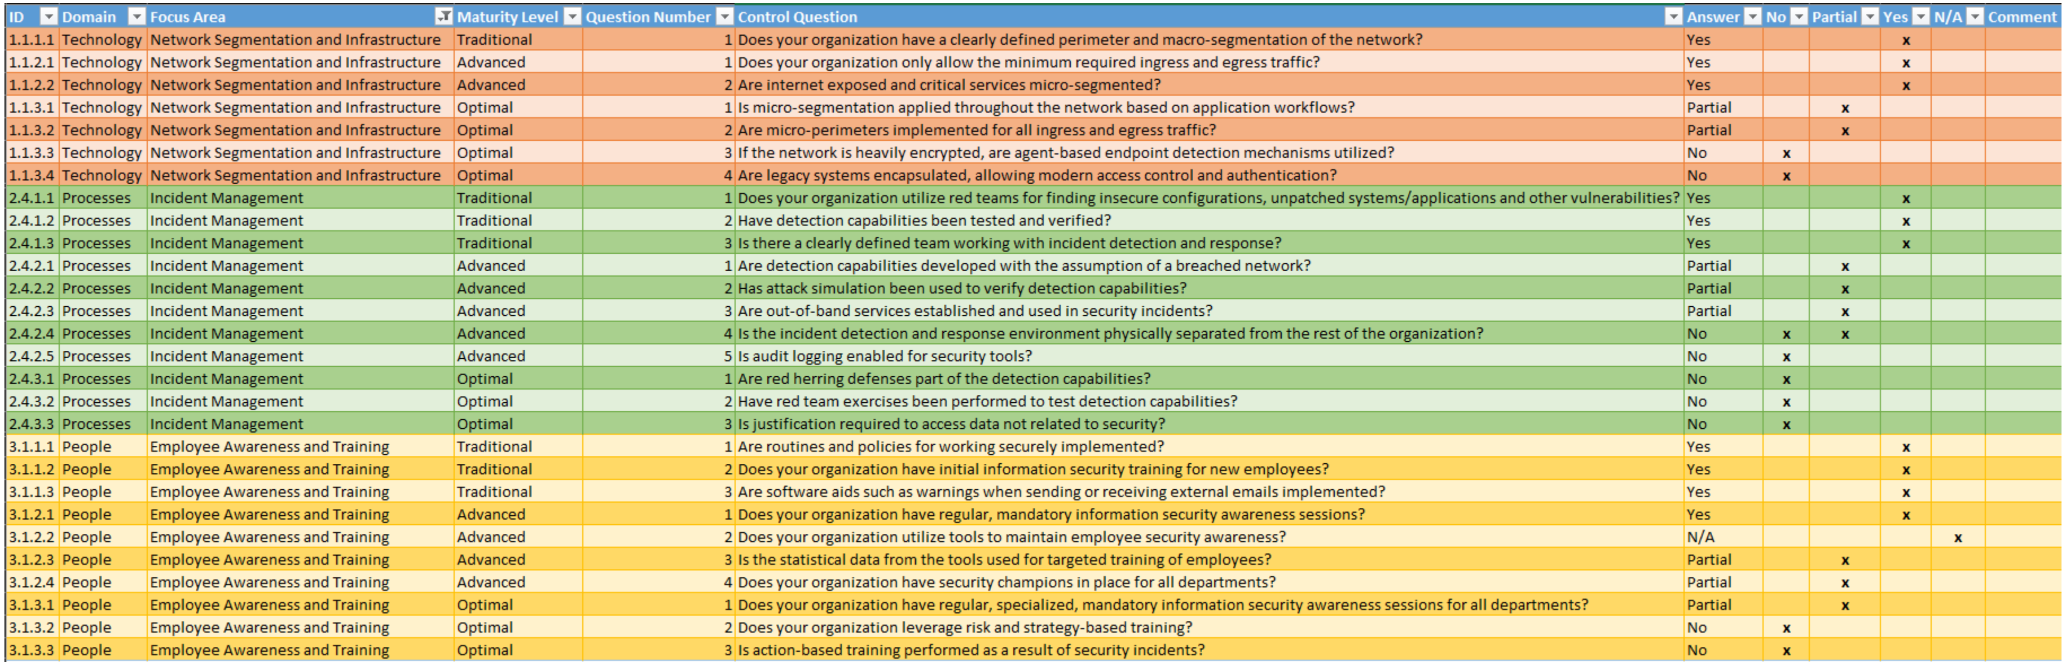

ID	Domain	Focus Area	Maturity Level	Question Number	Control Question	Answer	No	Partial	Yes	N/A	Comment
1.1.1.1	Technology	Network Segmentation and Infrastructure	Traditional	1	Does your organization have a clearly defined perimeter and macro-segmentation of the network?	Yes			x		
1.1.2.1	Technology	Network Segmentation and Infrastructure	Advanced	1	Does your organization only allow the minimum required ingress and egress traffic?	Yes			x		
1.1.2.2	Technology	Network Segmentation and Infrastructure	Advanced	2	Are internet exposed and critical services micro-segmented?	Yes			x		
1.1.3.1	Technology	Network Segmentation and Infrastructure	Optimal	1	Is micro-segmentation applied throughout the network based on application workflows?	Partial		x			
1.1.3.2	Technology	Network Segmentation and Infrastructure	Optimal	2	Are micro-perimeters implemented for all ingress and egress traffic?	Partial		x			
1.1.3.3	Technology	Network Segmentation and Infrastructure	Optimal	3	If the network is heavily encrypted, are agent-based endpoint detection mechanisms utilized?	No	x				
1.1.3.4	Technology	Network Segmentation and Infrastructure	Optimal	4	Are legacy systems encapsulated, allowing modern access control and authentication?	No	x				
2.4.1.1	Processes	Incident Management	Traditional	1	Does your organization utilize red teams for finding insecure configurations, unpatched systems/applications and other vulnerabilities?	Yes			x		
2.4.1.2	Processes	Incident Management	Traditional	2	Have detection capabilities been tested and verified?	Yes			x		
2.4.1.3	Processes	Incident Management	Traditional	3	Is there a clearly defined team working with incident detection and response?	Yes			x		
2.4.2.1	Processes	Incident Management	Advanced	1	Are detection capabilities developed with the assumption of a breached network?	Partial		x			
2.4.2.2	Processes	Incident Management	Advanced	2	Has attack simulation been used to verify detection capabilities?	Partial		x			
2.4.2.3	Processes	Incident Management	Advanced	3	Are out-of-band services established and used in security incidents?	Partial		x			
2.4.2.4	Processes	Incident Management	Advanced	4	Is the incident detection and response environment physically separated from the rest of the organization?	No	x	x			
2.4.2.5	Processes	Incident Management	Advanced	5	Is audit logging enabled for security tools?	No	x				
2.4.3.1	Processes	Incident Management	Optimal	1	Are red herring defenses part of the detection capabilities?	No	x				
2.4.3.2	Processes	Incident Management	Optimal	2	Have red team exercises been performed to test detection capabilities?	No	x				
2.4.3.3	Processes	Incident Management	Optimal	3	Is justification required to access data not related to security?	No	x				
3.1.1.1	People	Employee Awareness and Training	Traditional	1	Are routines and policies for working securely implemented?	Yes			x		
3.1.1.2	People	Employee Awareness and Training	Traditional	2	Does your organization have initial information security training for new employees?	Yes			x		
3.1.1.3	People	Employee Awareness and Training	Traditional	3	Are software aids such as warnings when sending or receiving external emails implemented?	Yes			x		
3.1.2.1	People	Employee Awareness and Training	Advanced	1	Does your organization have regular, mandatory information security awareness sessions?	Yes			x		
3.1.2.2	People	Employee Awareness and Training	Advanced	2	Does your organization utilize tools to maintain employee security awareness?	N/A				x	
3.1.2.3	People	Employee Awareness and Training	Advanced	3	Is the statistical data from the tools used for targeted training of employees?	Partial		x			
3.1.2.4	People	Employee Awareness and Training	Advanced	4	Does your organization have security champions in place for all departments?	Partial		x			
3.1.3.1	People	Employee Awareness and Training	Optimal	1	Does your organization have regular, specialized, mandatory information security awareness sessions for all departments?	Partial		x			
3.1.3.2	People	Employee Awareness and Training	Optimal	2	Does your organization leverage risk and strategy-based training?	No	x				
3.1.3.3	People	Employee Awareness and Training	Optimal	3	Is action-based training performed as a result of security incidents?	No	x				

Figure 4: The EZTMM Evaluation Sheet Control Questions

Most of the feedback was related to specific questions in our model, such as “Is it possible to split this into multiple questions? We use asset management tools, but do not have a fully implemented CMDB. Related: What is the difference between a configuration manager and an administrator?)” (This comment was translated from Norwegian) These comments were addressed, resulting in a revision of several control questions.

We also received requests for clarification: “This point is poorly explained, could perhaps have been elaborated on in the document.” (This comment was translated from Norwegian) These requests were also addressed, leading to further additions to the descriptions of some maturity levels in the EZTMM.

Lastly, we received some comments on the maturity assessment spreadsheet itself. One such comment was that the dashboard could have been more detailed. “Would possibly have been useful to have some additional details in the dashboard for when the sheet has been filled out. Instead of only getting the maturity level achieved, we could get how much we are missing, if we were close or not or something similar.” (This comment was translated from Norwegian) We addressed this comment by increasing visibility into the scoring on each maturity level for each focus area. Before and after screenshots can be seen below (Figures 5 and 6):

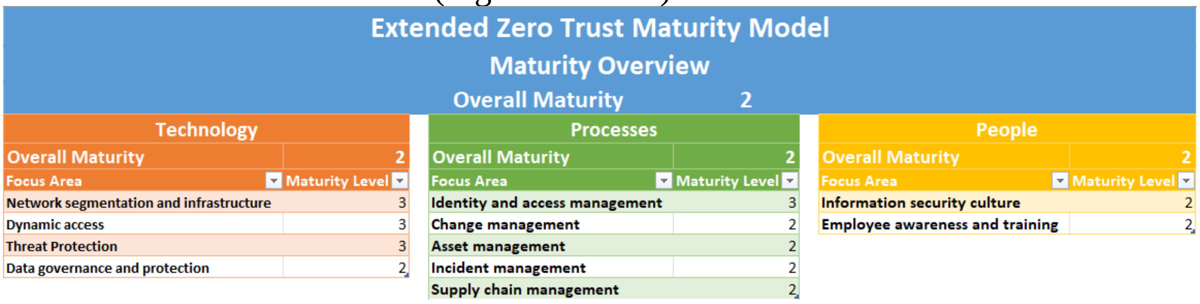

Extended Zero Trust Maturity Model

Maturity Overview

Overall Maturity 2

Technology

Overall Maturity	2
Focus Area	Maturity Level
Network segmentation and infrastructure	3
Dynamic access	3
Threat Protection	3
Data governance and protection	2

Processes

Overall Maturity	2
Focus Area	Maturity Level
Identity and access management	3
Change management	2
Asset management	2
Incident management	2
Supply chain management	2

People

Overall Maturity	2
Focus Area	Maturity Level
Information security culture	2
Employee awareness and training	2

Figure 5: EZTMM Evaluation Sheet original Maturity Dashboard

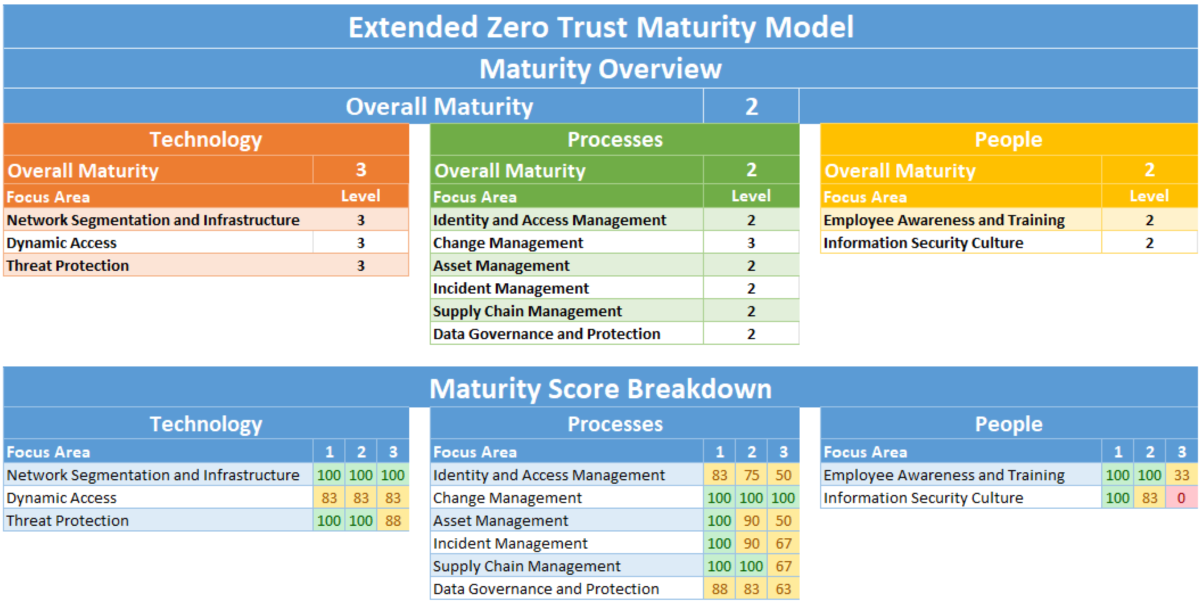

Extended Zero Trust Maturity Model

Maturity Overview

Overall Maturity	2

Technology	
Overall Maturity	3
Focus Area	Level
Network Segmentation and Infrastructure	3
Dynamic Access	3
Threat Protection	3

Processes	
Overall Maturity	2
Focus Area	Level
Identity and Access Management	2
Change Management	3
Asset Management	2
Incident Management	2
Supply Chain Management	2
Data Governance and Protection	2

People	
Overall Maturity	2
Focus Area	Level
Employee Awareness and Training	2
Information Security Culture	2

Maturity Score Breakdown

Technology			
Focus Area	1	2	3
Network Segmentation and Infrastructure	100	100	100
Dynamic Access	83	83	83
Threat Protection	100	100	88

Processes			
Focus Area	1	2	3
Identity and Access Management	83	75	50
Change Management	100	100	100
Asset Management	100	90	50
Incident Management	100	90	67
Supply Chain Management	100	100	67
Data Governance and Protection	88	83	63

People			
Focus Area	1	2	3
Employee Awareness and Training	100	100	33
Information Security Culture	100	83	0

Figure 6: EZTMM evaluation sheet updated Maturity Dashboard

Several corrections were also made to the spreadsheet, such as the Data Governance and Protection initially being incorrectly placed in the Technology domain in the original model, and the maturity level calculations of every subsequent focus area being incorrectly calculated as a result.

The only improvement suggestion we were unable to implement due to a lack of time was weighted focus areas:

> It is a good tool for providing a snapshot of where an organization is and the competency that is there. However, I think it was difficult to determine in which end to start and how to prioritize improvements. I would have liked to see a form of weighting for the various focus areas, and maybe some recommendations on what to prioritize based on where the organization has a low score, and which focus areas are most important. This way we would get a better impression of what is important for us. (This comment was translated from Norwegian)

Implementing this suggestion would require a lot of additional research into how to correctly weight the focus areas relative to each other. We decided that this would not be possible given the time constraints for the thesis and therefore mention it in our chapter on future research instead.

Implementing the feedback from the case study evaluation resulted in a total of 107 control questions within 11 focus areas and three domains. Three new control questions were added and eight were revised based on respondent feedback. The result of these changes was V1.0 of the EZTMM, our first release version and the final version to be developed during the writing of this thesis. Iteration 3 ended on the 29th of May 2022.

5 DISCUSSION

In this chapter we present our reflection on the design of the EZTMM, how we address an existing knowledge gap, and what we now consider the components of Zero Trust. We discuss the importance of a holistic approach to cybersecurity and describes how we successfully applied the Zero Trust principles to organizational domains. The model is evaluated against pre-defined acceptance criteria, before we discuss practical implications, future research, and the limitations of our study.

5.1 Reflections on the Design of EZTMM

The results show how design principles and requirements from design science research were used to develop a model for assessing and improving organizations' technical and organizational Zero Trust maturity. Design of the model was carried out adhering to the seven requirements proposed by Becker et al (Becker, Knackstedt, & Pöppelbuß, 2009). A knowledge base was founded by a comprehensive literature review covering theory on maturity models, comparison of existing maturity models and Zero Trust. An initial iteration of the model was created based on the knowledge base and feedback from subject matter experts was used to improve the model in several iterations. A case study evaluation was then performed to see if the model would comply with predefined acceptance criteria. Small adjustments were also made to the model based on observations and feedback during the case study. The result of our research includes the Extended Zero Trust Maturity Model, a self-evaluation tool and thorough documentation of the entire design process.

5.2 The Components of Zero Trust

One of the research questions we set out to answer was "what are the components of Zero Trust?" (SRQ1). Through our research we have discovered that Zero Trust may be seen as a mindset founded in the idea that trust is not granted implicitly, but rather gained through rigorous verification. There is no complete list of what components a security architecture based on Zero Trust should contain. However, our literature review showed that there are technologies supporting this mindset, making them relevant as components. Examples of such technologies are policy enforcement engines, tools for network segmentation and multi-factor authentication. Our research has shown, in addition to this, that the core principles of Zero Trust can be applied to other domains within information security. Our view is therefore that every change made to technology, processes or policies based on the Zero Trust principles can be interpreted as a component of an organizations' Zero Trust strategy.

5.3 Existing maturity models and the existing knowledge gap

In our comparison of existing maturity models, we discovered a gap in models created for assessing Zero Trust maturity where none focused on applying Zero Trust to organizational aspects. Although, research by Moddelkolk (Modderkolk, 2018) has similarities with the research performed in our thesis. Modderkolk's research revolves around designing the ZeTuMM model, which can be used to assess overall information security and is based on Zero Trust. However, the model only addresses Zero Trust in a technological aspect, and uses best practices, not related to Zero Trust, for other domains within information security. In addition, the model is purely descriptive, meaning it is purposed to assess the current situation as-is, and not for improving capabilities. ZeTuMM was the only model we identified for assessing Zero Trust that had detailed documentation of the models' design process, probably because the other models were not developed as part of academic research.

5.4 The EZTMM

While Zero Trust traditionally is limited to technology, the EZTMM extends the term to include the human side that is people and processes. This extension is important because it makes Zero Trust possible to use as a holistic approach to information security. Using a security model which only focuses on technology must be seen as a weakness in the current threat landscape, where humans are often considered the weakest link (Zahoor, Mahmood, & Javed, 2015). Research by Zahoor et al. argues that a holistic approach must be leveraged to address the incidents caused by human error. This is further confirmed by data presented in Verizon's Data Breach Investigations report stating that 82% of the breaches they observed were caused by the human element, including social attacks, errors, and misuse (Verizon, 2022). The respondents participating in our study were very positive to the inclusion of the people and processes domains, and one of the participants stated that while many organizations have the technical controls in place, what they really struggle with are the processes.

During the early stages of development, we realized there was limited information available on applying zero trust principles to aspects beyond technology. Reviewing best practices suggested for organizational domains in industry accepted security frameworks helped us identify controls relevant to Zero Trust. Numerous brainstorming sessions were used to apply the Zero Trust principles to different organizational areas within information security. The relevance of the identified areas and suggested controls were then validated with expert feedback. In total, we were able to identify and apply the Zero Trust principles to the following eight organizational areas:

1. Identity and Access management
2. Change Management
3. Asset Management
4. Incident Management
5. Supply Chain Management
6. Data Governance and Protection
7. Employee Awareness and Training
8. Information Security Culture

Every respondent was asked about the relevance of each included focus area, and if they had any suggestions for areas to be added. There were a couple of suggestions for additional focus areas, some were implemented while others were not. Read more about each focus area and why they were excluded or changed during development in Appendix D: Focus Area Changes and suggestions.

5.5 Acceptance Evaluation of the EZTMM

Evaluation of the model is among the requirements for design science proposed by Becker et al. (Becker, Knackstedt, & Pöppelbuß, 2009). Acceptance criteria was therefore developed during the design phase to have something to compare and measure the final iteration against. A total of seven requirements were defined, along with a rationale for inclusion (see table 10 for the full list).

Being applicable in a broad range of organizations was the first defined criteria for acceptance. The model was designed to be used by all types of organizations, differing in size and industries. Respondents were brought in from different organizations to include various perspectives and ideas. Only one organization was used in the case study validating the model. However, the organization is characterized by having a very complex infrastructure, being comprised by multiple business units, and leveraging many different technologies. Being able to successfully use the model in such an organization may be used as an argument for its applicability, as it is likely that it would suffice for less complex organizations as well.

Organizations being able to self-assess their current Zero Trust Maturity using our model was a goal from the beginning and was also defined as an acceptance criterion. We therefore developed a self-assessment tool in addition to the EZTMM. The solution was developed in excel, where control questions can be answered, and the maturity score will be calculated automatically. The self-assessment tool was tested as part of the case study validation, and while it was considered successful, several improvements were suggested. Some of these improvements were implemented, while others are suggested for future research.

Another relevant criterion was being able to use the model to improve overall information security. The EZTMM gives detailed background information on each domain, and descriptions of every suggested control. Having this information should enable organizations in prioritizing future steps to improve their capabilities in various areas within information security. However, the case study revealed that the respondent felt overwhelmed by the results and found it challenging to figure out which area to start improving. A possible solution was proposed by the respondent and is discussed under "Weighted focus areas" in the future research section.

The models' success depends on its correctness. Respondents were asked if the suggested security controls were described correctly and placed at the appropriate level. Every suggestion for change was considered, and many was implemented. All respondents were asked to review the model and suggest changes for structure, focus areas, description and level of security controls, and consistency and correctness. In the last iteration of the model, we received very few suggestions for changes. We concluded that consensus had been reached among the respondents and that additional iterations would provide little extra value.

The last acceptance criterion we defined was "respondents being presented the model show desire to leverage the model in their own organization". Respondents were already showing interest in using the model after being presented the initial draft. However, it was stated that further development would be required for it to be considered. After finalizing the last iteration, multiple respondents (including one CISO) expressed actual intent of using the model for assessing and improving their organizations information security. Founding principles of design science research proposed by Hevner et al. (Hevner, Ram, March, & Park, 2004) highlights the importance of demonstrating the problem relevance. There was a shared agreement among our respondents that it existed a need for a model as the we proposed. Our research has also caught interest in the cybersecurity community in Norway, and we have been asked to present our findings at multiple events, as well as in a podcast.

Our opinion is that comparing the results against the predefined acceptance criteria makes a good argument for accepting the model. Reviewing how the model was designed show that the process described prior to development was followed in detail. We believe that choosing design science as the research approach and sticking to it throughout the study were important factors in creating an acceptable holistic model for assessing and improving Zero Trust maturity. However, even though we consider the model accepted, it does not mean there is no room for improvements.

5.6 Practical implications

This thesis has practical implications for any organization looking to improve their information security posture, whether they are just starting their Zero Trust implementation or are well underway. The benefits of the model are threefold:

1. Improved identification of low-hanging fruits and problem areas
2. A basis for planning and evaluating improvements to the organization’s information security posture and tracking progress
3. Improved reporting on the organization’s security posture

Organizations can use the EZTMM (All the original files can be found in appendix E) to evaluate their own maturity and identify problem areas: The included evaluation spreadsheet can automatically calculate an organization’s overall maturity, along with maturity levels for each domain or focus area. Furthermore, the dynamic maturity dashboard in the evaluation spreadsheet will provide a scoring for each focus area, allowing the organization to easily identify problem areas and low-hanging fruits.

This assessment can then serve as a basis for planning and evaluating improvements and mitigations, improving the organization’s overall security posture. As the improvements are implemented, the control question answers can be updated, allowing the organization to track their progress towards the next maturity level.

The maturity dashboard is also ideal for creating reports on the organization's security posture: The dashboard is designed with multiple layers of abstraction, from the highest level (overall maturity) all the way down to scoring on the maturity levels for each focus area. This allows the report to be easily interpreted by anyone from upper management to technical specialists. For even more details, the readers of the report can refer to the control question list. These abstraction layers make the dashboard suitable for high level security posture reports or arguing for higher security budgets with upper management, as well as giving details on exactly which measures are needed to improve the organization's security posture.

5.7 Future Research

In this chapter we will discuss potential directions for future research on the topic of holistic Zero Trust maturity models and the EZTMM.

5.7.1 Zero Trust Principles for organizational aspects

Over the course of writing this thesis, it has become increasingly apparent that there is a lack of research on applying Zero Trust principles to organizational aspects of information security. While the application of Zero Trust principles to data, networks and identities is important, we have found it equally important to apply the principles to the organization's processes and people domains. These domains are according to many of our respondents the domains where organizations struggle in terms of information security, and we think Zero Trust is a part of the solution to this. We would therefore like to see more research on this topic.

5.7.2 Expansion of the EZTMM

The EZTMM could be developed to become even more comprehensive with further additions to both the processes and people domains. Developments could also be made in the technology domain to make the model more capable of deeply assessing the technological Zero Trust maturity of an organization. As seen in appendix D, we received several suggestions for new focus areas throughout developing this model, some of which were included and some not. Further research on the inclusion of these and other potential focus areas is warranted.

5.7.3 Compatibility with existing standards

The EZTMM could also be further improved by aligning the control questions with already existing standards such as the CIS controls, ISO 27001 and NIST CSF, allowing for easier adoption in organizations that have built their cybersecurity practices on one or more of these standards. This could be done by mapping the EZTMM control questions to existing controls in either framework and adapting the EZTMM where necessary. Such a mapping could allow for an even more thorough assessment of certain focus areas by leveraging other security frameworks, or even boost adoption of the EZTMM due to its compatibility with existing cybersecurity initiatives in organizations.

5.7.4 Weighted focus areas

The EZTMM could be further developed to have weighted focus areas or control questions, giving the organization a better idea of which areas are the most important. The model could then provide improvement suggestions based on each organization's scoring in these weighted focus areas, allowing it to be become even more prescriptive. This weighted approach would require more research into how each focus area should be weighted.

5.7.5 Case studies on the application of Zero Trust principles to all organizational processes

We would also like to see more case studies on attempting to apply Zero Trust principles to every security-related process in an organization, possibly through the implementation of EZTMM as the organization's maturity assessment framework. Seeing the benefits of a holistic approach to Zero Trust implementation could contribute to raising adoption rates and provide valuable insights into the problems and advantages of such a holistic approach to Zero Trust.

5.8 Limitations

The empirical data gathered during the work on this thesis was collected from Norwegian organizations, or the Norwegian branch of an international organization. Organizations from other countries may be further along or behind in their Zero Trust implementations and may therefore have different ideas of which measures are appropriate at the various maturity levels.

In addition to the data not being generalizable across countries, the data is also not generalizable across sectors. We have attempted to include respondents from different sectors such as consulting, finance, process industry and education, but the qualitative approach taken along with the limited number of respondents make generalization impossible.

Another limitation related to the respondents is that they were all recruited through our own professional networks or through supervisor recommendations. We found that when contacting organizations without a known reference, the organizations were significantly less cooperative. This could be because information security is a sensitive topic for most organizations and providing such sensitive information to unknown students being deemed too great of a risk.

The organization participating in the case study has been a respondent since iteration 1, and therefore knows the model well. We have had several discussions with them regarding the model's structure and the reasoning behind each focus area. This could lead to a different result than if the organization conducting the maturity assessment saw the model for the first time.

The EZTMM itself, however, must be able to be used in various countries and organizations. We have therefore attempted to mitigate the generalization issues as much as possible by basing our work on international maturity models and standards. However, due to time constraints we were not able to develop a full mapping between the EZTMM control questions and other maturity models. This limitation could potentially hurt adoption of our maturity model in organizations that have already based their Zero Trust initiatives on existing standards and frameworks.

6 CONCLUSION

In this thesis we have shown how a model can be designed for organizations to assess and improve their technical and organizational Zero Trust maturity. Using design science research, we developed the EZTMM, a fully functional holistic maturity model with an accompanying maturity assessment tool. This model is an important contribution to both the maturity model literature, where a higher emphasis on organizational domains is needed, and to organizations wanting to improve their security posture using a much-needed holistic approach, rather than the technical approach suggested by other maturity models. This model meets the meta requirements defined in table 10 and is designed to address the limitations of current maturity models by applying Zero Trust principles to both organizational and technical domains **(RQ)**.

Furthermore, we have found that Zero Trust components can be both technical solutions such as multi-factor authentication, policy enforcement engines and network segmentation, but they can also be any other manifestation of the core principles of Zero Trust, whether this be a technical solution or a change in the organization's processes, policies, or procedures **(SRQ1)**:

1. Verify explicitly
2. Use least privileged access
3. Assume breach

During our literature reviews on Zero Trust and maturity model theory we identified several maturity models, four of which we compared: The C2M2 by the U.S Department of Energy and the Department of Homeland Security, the CTI-SOC2M2 by Vielberth, Schlette & Pernul, the ZeTuMM by Modderkolk and the Zero Trust Maturity Model by CISA. The two latter models were Zero Trust-focused while the former two were general information security maturity models. The fundamental principles of Zero Trust can be traced back to the origins of the internet. We found that these principles are incorporated to some degree in the two former models in a more general sense **(SRQ2)**.

Working closely with practitioners and experts from our respondent organizations, we found that Zero Trust principles could be applied to the following organizational processes, resulting in them becoming the focus areas in our processes and people domains:

1. Identity and Access Management
2. Change Management
3. Asset Management
4. Incident Management
5. Supply Chain Management
6. Data Governance and Protection
7. Employee Awareness and Training
8. Information Security Culture

More processes such as Risk Management, Hardening and Configuration Management and Patch Management were also considered. We found that the latter two processes can indeed benefit from Zero Trust principles but decided to incorporate them into existing focus areas **(SRQ3)**.

7 REFERENCES

Køien, G. (2021). Zero-Trust Principles for Legacy Components 12 Rules for Legacy Devices: An Antidote to Chaos. *Wireless Personal Communications*.

Becker, J., Knackstedt, R., & Pöppelbuß, J. (2009). Developing Maturity Models for IT Management. *Business & Information Systems Engineering*, 213–222.

Bruin, T. d., Freeze, R. D., Kulkarni, U., & Rosemann, M. (2005). Understanding the Main Phases of Developing a Maturity Assessment Model. *Association for Information Systems* .

Buck, C., Olenberger , C., Schweizer, A., Fabiane, V., & Torsten, E. (2021). Never trust, always verify: A multivocal literature review on current knowledge and research gaps of Zero Trust. *Computers & Security*.

Campbell, M. (2020). Beyond Zero Trust: Trust Is a Vulnerability. *IT INNOVATION*.

CISA. (2021). *Defending Against Software Supply Chain Attacks* . CISA.

CISA. (2021). *Zero Trust Maturity Model.* Cybersecurity and Infrastructure Security Agency.

Collier, Z., & Sarkis, J. (2021). The zero trust supply chain: Managing supply chain risk in the absence of trust. *International Journal of Production Research*.

D'Silva, D., & Ambawade, D. (2021). Building A Zero Trust Architecture Using Kubernetes. *International Conference for Convergence in Technology.* Pune.

Durán, J., & Jeferson, M. (2021). *Software Supply Chain Attacks, a Threat to Global Cybersecurity: SolarWinds' Case Study.* IETA.

Ferretti, L., Magnanini, F., Andreolini, M., & Colajanni, M. (2021). Survivable zero trust for cloud computing environments. *Computers & Security*.

Gartner. (2021, April 1). *Gartner Forecasts Global Devices Installed Base to Reach 6.2 Billion Units in 2021*. Retrieved from Press Release Newsroom: https://www.gartner.com/en/newsroom/press-releases/2021-04-01-gartner-forecasts-global-devices-installed-base-to-reach-6-2-billion-units-in-2021

Gottschalk, P. (2009). Maturity levels for interoperability in digital government. *Government Information Quarterly*, 75-81.

Hevner, A. R., Ram, S., March, S. T., & Park, J. (2004). Design Science in Information Systems Research. *Management Information Systems Quarterly*.

Internet Engineering Task Force. (1989, October). *Requirements for Internet Hosts -- Communication Layers - RFC 1122*. Retrieved from https://datatracker.ietf.org/doc/html/rfc1122

Kindervag, J. (2010). Build Security Into Your Network's DNA: The Zero Trust Network Architecture. *Security & Risk Professionals*.

Kindervag, J. (2010). No More Chewy Centers: Introducing The Zero Trust Model Of Information Security. *Security & Risk Professionals*.

Kindervag, J. (2010). No More Chewy Centers: Introducing The Zero Trust Model Of Information Security. *For Security & Risk Professionals*.

Kitchenham, B., & Charters, S. (2007). *Guidelines for performing Systematic Literature Reviews in Software Engineering*. Keele University & Durham University.

Kumar, N., & LaRoy, N. (2021). Zero Trust in the Context of the Utility Industry. *Springer*.

Lain, D., Kostiainen, K., & Capkun, S. (2021). *Phishing in Organizations: Findings from a Large-Scale and Long-Term Study*. ETH Zurich, Switzerland: Department of Computer Science.

Li, Y., & Liu, Q. (2021). A comprehensive review study of cyber-attacks and cyber security; Emerging trends and recent developments. *Energy Reports*, 8176-8186.

March, S., & Smith, G. (1995). Design and Natural Science Research on Information Technology. *15*(10.1016/0167-9236(94)00041-2).

May, J., & Dhillon, G. (2010). A Holistic Approach for Enriching Information Security Analysis and Security Policy Formation. *18th European Conference on Information Systems.*

Mehravari, P. D. (2015). Evaluating and Improving Cybersecurity Capabilities of the Energy Critical Infrastructure. Waltham, MA, USA: IEEE International Symposium on Technologies for Homeland Security.

Microsoft. (2019). *Zero Trust Maturity Model.* Microsoft.

Modderkolk, M. (2018). ZERO TRUST MATURITY MATTERS MODELING CYBER SECURITY FOCUS AREAS AND MATURITY LEVELS IN THE ZERO TRUST PRINCIPLE. *LegalIT.* Utrecht: Universiteit Utrecht.

Olson, K., & Keller, E. (2021). Federating Trust: Network Orchestration for Cross-Boundary Zero Trust. *SIGCOMM.*

Paulk, M. C., Curtis, B., Chrissis, M. B., & Weber, C. (1993). *Capability Maturity Model for Software (Version 1.1).* Software Engineering Institute.

Poeppelbuss, J., & Roeglinger, M. (2011). What makes a useful maturity model? A framework of general design principles for maturity models and its demonstration in business process management. *European Conference on Information Systems (ECIS).* Helsiniki, Finland.

Poeppelbuss, J., Niehaves, B., Simons, A., & Becker, J. (2011). Maturity Models in Information Systems Research: Literature Search and Analysis. *Communications of the Association for Information Systems*.

Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero Trust Architecture. *NIST Special Publication 800-207*.

SABSA. (2022). *The Attributer's Blog – Zero Trusted*. Retrieved from SABSA Enterprise Security Architecture: https://sabsa.org/the-attributers-blog-zero-trusted/

SABSA. (2022, 04 20). *The Attributer's Blog – Zero Trusted*. Retrieved from SABSA Enterprise Security Architecture: https://sabsa.org/the-attributers-blog-zero-trusted/

Symantec. (2019). *Internet Security Threat Report.* Symantec.

Teerakanok, S., Uehara, T., & Inomata, A. (2021). Migrating to Zero Trust Architecture: Reviews and Challenges. *Security and Communication Networks*.

Tyler, D., & Viana, T. (2021). Trust No One? A Framework for Assisting Healthcare Organisations in Transitioning to a Zero-Trust Network Architecture. *Applied Sciences*.

U.S Department of Energy. (2021, July). *Cybersecurity Capability Maturity Model (C2M2).* Retrieved from Energy.gov: https://www.energy.gov/sites/default/files/2021-07/C2M2%20Version%202.0%20July%202021_508.pdf

Van Niekerk, J., & Von Solms, R. (2010). Information security culture: A management perspective. *Computers & Security*.

Verizon. (2021). *2021 Data Breach Investigations Report.* Verizon.

Verizon. (2022). *2022 Verizon Data Breach Investigations*.

Vielberth, M., Schlette, D., & Pernul, G. (2021). CTI-SOC2M2 – The quest for mature, intelligence-driven security operations and incident response capabilities. *Computers & Security*.

Warren, R. L. (2021, December 03). *Zero Trust Security: Moving From 'Trust but Verify' to 'Never Trust, Always Verify'*. Retrieved from New Jersey Law Journal: https://www.law.com/njlawjournal/2021/12/03/zero-trust-security-moving-from-trust-but-verify-to-never-trust-always-verify/?slreturn=20220431123157

Webster, J., & Watson, R. (2002). ANALYZING THE PAST TO PREPARE FOR THE FUTURE: WRITING A LITERATURE REVIEW. *MIS Quarterly*.

What is IT asset management (ITAM)? (2022). Retrieved from Atlassian: https://www.atlassian.com/itsm/it-asset-management

Wylde, A. (2021). Zero trust: Never trust, always verify. *International Conference on Cyber Situational Awareness and Data Analytics* . Cardiff: Cardiff University.

Xiangshuai, Y., & Wang, H. (2020). Survey on Zero-Trust Network Security. *Springer*.

Zahoor, A., Mahmood, H., & Javed, A. (2015). Information security management needs more holistic approach: A literature review. *International Journal of Information Management*.

APPENDIX A: INTERVIEW GUIDE

Extended Zero Trust Maturity Model – Intervjuguide

Introduksjonsspørsmål

0.1 Kan du fortelle oss litt om din bakgrunn innen informasjonssystemer?

0.2 Hvor mye erfaring har du med informasjonssikkerhet?

Generelt

1.1 Er Zero Trust et begrep du/dere er kjent med?

1.2 Har du/dere brukt modenhetsmodeller tidligere?

1.3 Viser modellen at Zero Trust prinsipper er overførbare til andre områder innen informasjonssikkerhet?

1.4 Hva tenker du/dere om denne modellens relevans og nytteverdi?

Modellens innhold

2.1 Hva synes dere om modellens dekning av områder innenfor informasjonssikkerhet?

2.2 Var det noen av områdene som ble beskrevet dere tenker er mindre relevante fra et Zero trust perspektiv?

2.3 Var det fokusområder dere følte vi manglet, som dere tenker kunne passet inn i modellen?

2.4 Er de ulike tiltakene beskrevet godt nok, eller skulle hvert tiltak vært beskrevet grundigere?

2.5 Er de ulike tiltakene plassert på riktig nivå av modenhet?

2.6 Var det noen tiltak dere savnet, eller noen dere mente ikke hørte hjemme?

2.7 Kommer det klart nok frem hvordan en oppnår de ulike gradene av modenhet?

Modellens struktur

3.1 Hva tenker dere om bruken av tre nivåer for ulik modenhetsgrad?

Bruk av modellen

4.1 Kunne det være aktuelt for deres egen del å gjøre en selvevaluering av deres modenhet?

Avsluttende spørsmål

5.1 Er det noe mer du ønsker å tilføye? For eksempel temaer du mener er viktig som ikke har blitt tatt opp til nå.

5.2 Vet du om noen andre sikkerhetseksperter som du vil anbefale oss å ta kontakt med?

5.3 Vi ønsker å få tilbakemeldinger på endringene vi utfører etter disse intervjuene. Er det greit for deg at vi sender en oppdatert utgave av modellen med tydeliggjorde endringer, for en tilbakemelding?

APPENDIX B: CASE STUDY NOTES

EZTMM Case Study Maturity Evaluation Notes

The case study was performed by a security architect with broad knowledge of the subject organization. One of the thesis authors was present during the entire evaluation to take notes and answer questions when necessary.

Several notes were taken regarding different control questions:

1.1.3.3 – Denne er oppfylt, til tross for at advanced ikke er oppfylt. Skal det være mulig å ha tiltak på de øvre nivåene når de laveste ikke er oppfylt?
1.2.3.3 – Litt vanskelig formulering – Ville vurdert å legge inn application access – Applikasjonen trenger ikke nødvendigvis å være designet for dette, fordi man kan bruke andre verktøy for å få til dette.
1.3.3.3 – Dette punktet er litt dårlig forklart, kunne kanskje vært utdypet ytterligere i dokumentet.
1.3.3.4 – På mobiltelefoner også? All assets, hva mener man med det, er det 100% av alle assets? Hvis 95% av alle alle assets, så kan vi sette yes, men hvis det er 100% må vi sette Partial.
2.2.1.2 – Kanskje reformulere til employees?
2.2.2.1 – Dette er et stort spørsmål, fordi her antyder man at alle som bestiller en change kjenner til Zero Trust principles. Og så må de som utfører changen vite at denne changen ikke følger zero trust prinsippene, og deretter be sikkerhet om å gjøre en review.
Forslag til future research: Ettersom dette området i stor grad er skrevet om fra tidligere utkast burde vi kanskje hatt en review av dette området av noen som er veldig modne på change management.
2.2.3.2 – If a more lenient – Denne setningen er veldig tungt formulert. Vanskelig for respondent å forstå. Også, hva betyr shortly after?
2.3.2.1 – Mulig å dele dette opp? Vi bruker asset management tools, men vi har ikke en fullstendig implementert CMDB.
Relatert: Hva er forskjellen på en configuration manager og en administrator?
2.3.3.4 – Var dette litt lite konkret for fokusområdet?

2.4.3.2 – Duplikat spørsmål?
Data governance er et ganske vanskelig område å svare ut ettersom vi har store mengder data.
Supply chain – remove uneccessary dependencies – Dette er en jo tidlig steg I Zero trust og er ikke noe som gjøres tradisjonelt. Dette tiltaket må nok endres. Tradisjonelt beskriver kanskje ikke så godt hva som faktisk er tradisjonelt under Supply chain? Må revurderes
2.5.1.4 (2.5.1.3 after change) – Vil noen noensinne svare nei på denne?
Vi har glemt å legge til Kontrollspørsmål på backup.
3.2.1.1 – Hva er all tasks? Man vil aldri få Yes på denne. Man definerer heller ikke policier for en oppgave, man definerer prosedyrer.
3.2.2.1 – Svare partially på regularly?

EZTMM Case Study Interview Notes

After the maturity assessment, a semi-structured interview was conducted with the purpose of validating the model against the acceptance criteria.

Er EZTMM tilstrekkelig som støttedokument for å kunne svare ut kontrollspørsmålene?

- Ja, jeg synes det. Det var en del ganger jeg måtte gå tilbake til hoveddokumentet for å undersøke hva som mentes med de ulike kontrollspørsmålene, og som regel fikk jeg svar på det lette etter. Det var imidlertid et par ganger jeg måtte forhøre meg med andre, eller gjøre et raskt søk på internett.

Føler du at kontrollspørsmålene følger en naturlig progresjon i modenhet?

- Dere har gjort en god jobb med å utforme kontrollspørsmålene. Det var et par ting vi la merke til underveis, men stort sett var progresjonen bra.

Var det mulig for deg å svare ut kontrollspørsmålene på egenhånd, eller måtte du hente inn hjelp på visse områder?

- Svarte ut på egenhånd, men det var et par områder der jeg måtte gjette litt, hvor det kanskje hadde vært mer hensiktsmessig å hente inn hjelp fra andre avdelinger. Jeg er en del av en teknisk avdeling, og noen av områdene kanskje spesielt innenfor people var litt vanskelig å svare ut alene.
-

Føler du at modellen effektivt kan brukes til å måle Zero Trust modenhet?

- Ja. Føler først og fremst at jeg har lært mye om Zero Trust, nyttig å lese gjennom dokumentet og deretter å svare ut spørsmålene.Organisasjoner som føler de har oppnådd Zero trust kan nok kjøre gjennom denne modellen her og få en wake-up call.

Hva tenker du om målingen av Zero Trust modenhet på de organisatoriske områdene (processes og people?)

- Jeg synes det gir mening. Det hadde ikke gitt mening å ikke skulle ha med de organisatoriske områdene dersom man ønsker å få et helthetlig inntrykk. Synes det er kult at dere har klart å vise at Zero trust er mer enn bare bits and bytes.

Vil modellen kunne brukes til å forbedre den totale informasjonssikkerheten i en organisasjon?

- Det er et godt verktøy for å gi et nåbilde av hvor en organisasjon står og kompetansen som finnes. Synes dog at det var vanskelig å vite hvilken ende man skulle begynne fra når man skal begynne å prioritere forbedringer. Det jeg hadde likt å sett var en slags vekting på de ulike områdene, og kanskje at det ble gitt anbefalinger på hva man burde prioritere å forbedre basert på hvor man scorer dårlig og hvilke områder som er viktig. På den måten hadde man fått et innstrykk av hva som var viktigst for oss.

Hva følte du var den største utfordringen i evalueringen?

- Det er et veldig omfattende tema, og man må ha bred kunnskap om organisasjonen man evaluerer. Man må også kjenne godt til hvilken kunnskap organisasjonen besitter innenfor de ulike områdene. Derfor var det noe utfrodrende å fylle ut denne alene, og ved flere anledninger ble det noe gjetting

- En annen utfordring vi la fort merke til var når man fylte ut modellen var spørsmål som f.eks "Benytter alle applikasjoner MFA". Det er svært vanskelig å svare ja på dette. Det kan hende man har 200 applikasjoner der 199 av dem bruker MFA. Skal man da svare "Partially" på dette spørsmålet, eller "Yes"?

Mindre utfordringer:

Det som passet oss dårlig var det supply chain management. Vi jobber jo lite med software utvikling osv, men vi har mye supply chains. Derfor var dette begrepet litt forvirrende.

Generelle utfordringer:

Hadde muligens vært nyttig å ha noe mer detaljer i dashboardet når man er ferdig å fylle ut. I stedet for at bare får listet ut tallet på nivået man er på, kunne man fått oppgitt hvor mye man manglet, om man var nære eller ikke eller lignende.

Kommentarer på utregning:

Det ser ut til å være noen feil med formlene for utregning.

Vi har fått 0 i score på incident management, når vi har svart yes på alle på spørsmålene i traditional.

Samme gjelder IAM. 2 yes og 1 partial på tradisjonell gir oss modenhetsnivå 0.

Samme gjelder Employee and Awareness.

APPENDIX C: INTERVIEW COMMENTS PER ITERATION

Iteration 1

Structure Improvements	Corrections and Consistency	Focus Area Improvements/Suggestions
Defining maturity levels (what does traditional, advanced, optimal mean in general). Are there some more well-known maturity levels that you could use and refer to. This also makes it easier when describing the development process, i.e., you can state that they are based on xxx (preferably research paper) and modified to fit this context (e.g., to include the "traditional" to show what the traditional non-zerotrust way is).	Think about the style in which the maturity levels are defined. Now, at least in many requirements, the text describes what an organization "should" do to be in specific maturity. At least some maturity models would describe what an organization "is" doing when they are at specific level. Think for instance risk management: maturity 1) risk assessments are conducted on ad-hoc basis; 2) risk assessments are planned, systematic and periodic; 3) risk assessment result are used to guide organizations future activities.	Several "sub-requirements" in one requirement: think about how this would influence the use of the maturity model, e.g., what would it mean in terms of results to have good internal software development processes but less-than-optimal software supply chain/procurement processes (requirement 10).
Organizing the requirements: the requirements could be categorized to orgnaizaional/technical requirements	As dicussed, the changes to previous comment might change this, but also think what is this a maturity model of. Is it a maturity model to evaluate organization's implementation of zerotrust, a maturity model to evaluate organization's maturity to implement zerotrust, or a maturity model to evaluate organizational/process zerotrust maturity (or perhaps combination of these or none of these)	Bør det vurderes å være eksplisitt på hvordan dette bør fungere for skybasert infrastruktur? (kan gjenbrukes tankegangen her, og stort sett lettere å implementere i sky fremfor on-prem)
Når jeg leser disse områdene så ser jeg en god blanding, uten at jeg ser en rød tråd.	Application or workload segmentation should also be in focus. best practice today is to implement micro segmentation.	Man kan også gjøre kultur-assessments. Og da bruker man et behavioral framework som definerer hvordan man ønsker at folk

Structure Improvements	Corrections and Consistency	Focus Area Improvements/Suggestions
Mer konkret: Dere har tatt nettverk, og da er det ren arkitektur. Men dette er så å si den eneste dere har med om arkitektur. Det har ikke noe med om prosessen f.eks. Modenhet på hvordan man skal gjøre et design av en arkitektur. Og dere har tatt nettverk, uten å ha noe med om applikasjon. Og applikasjon er veldig viktig. Den integrasjonen mot andre applikasjoner, den delen hvor man utveksler informasjon mellom applikasjoner og tredjeparter, den delen mener jeg er «key». I praksis kan man bygge opp en arkitektur der alt er tredjeparter fordi man stoler ikke på noe. Og da blir det lettere å applisere den i en ny hverdag. Da kan man enkelt ta inn nye cloud applications og sette de inn i våre zero trust architecture.	I would add workload segmentation as another focus point.	tenker. Så det ligner på et sikkerhetsrammeverk som NIST CSF, der mennesker er i fokus. Du kan ha en 30-day sikkerhetschallenge, du kan ha et sikkerhets newsletter, coaching. Så det er eksempler på en deler av en actionplan man definerer. Man definerer hvordan man ønsker at de tar til seg informasjon, hvordan ønsker man at de handler basert på informasjonen, hvordan ønsker du at de jobber. Hvordan ønsker du at de kommuniserer. Det er overfladiske spørsmål som hjelper deg å bestemme hvordan man ønsker at de ansatte skal oppføre seg.
Dere har tatt litt av prosesser. Dere har med noen viktige: - asset management - change management (Jeg betrakter det mer som IT) - network access (Betrakter dette som access management) Her har dere tre nøkkelprosesser. + Sikkerhetsprosesser som er mer relatert til secops i threat protection og incident detection.	Better say, supply chain management. Risk factor in real scenarios is 3rd party supply chain which is NOT limited to software. This is a very overlooked topic in real cases. Very good point you brought up!	Insentiver slår Kultur any day

Structure Improvements	Corrections and Consistency	Focus Area Improvements/Suggestions
Da vil jeg også si: Sett de sammen og se det i sammenheng. Som i «prosessdelen»		
Could it be an idea to divide the measures into areas such as identities, infrastructure, software etc.? Maybe develop some sort of tagging system for each measure?	Add also VRFs. VRFs play a key role at route-segmentation level.	Muligens ikke riktig sted å legge det, men '(external) attack surface management' er en videreføring av asset management. Det var via manglende kontroll på dette at vi ble truffet to ganger av ransomware i fjor, f.eks.
Dere valgte 3 når det vanligvis er 5. De fem som vanligvis velges er basert på prosesser. Vet ikke om det er bevisst eller ei, men jeg synes deres tilnærming er bedre enn de internasjonale standarder.	Due to the increased utilization of remote-working services, the perimeter is becoming increasingly difficult to define and control.	Ville lagt til 'resolving gaps'. 'Assets' bør også deles i nivå, f.eks. data => komponent => system => verdikjede.
Da blir spørsmålet, hva er scope i et zero trust environment? Skal det være hele organisasjonen, hele infrastrukturen. Prøv å få med dere en beskrivelse av hva dere ser for dere som scope, av hva man skal vurdere modenhet.	Inaccurate, it is not just for servers, I would use computing resources or workloads. It contains Kubernetes and containers plus other types of computing resources. This is exactly the point that most implementations fail to deliver.	Good asset management could be seen as a prerequisite of Zero Trust.
Det andre jeg hadde – Hvor fant dere deres områder? Kunne ikke finne det beskrevet hvordan disse områdene ble valgt ut.	This is a vague term to use. you should define what you really mean by microperimeters with some examples. This is not a standard term in audit and compliance world. I'd define it in prior sections before using it here.	I think you have included all the most essential focus areas already. Initially some of them seemed to not quite fit in (Security awareness and culture), but when explained during the interview they make very much sense to include.
Modenhetsnivåer? Keep it simple, og med det mener jeg 3.	Not necessarily, I'd change it to: assumes all applications are accessed via untrusted network such as Internet	Å ta inn risk management blir jo å ta inn en helt ny dimensjon. Og dere har jo inkludert en god del tekniske detaljer her. Greit å få med i oppgaven hvorfor man ikke har tatt med risk management f.eks. Vise at det er tatt et aktivt valg rundt.

Structure Improvements	Corrections and Consistency	Focus Area Improvements/Suggestions
Data protection og governance – Litt forvirrende å kombinere disse to med mindre du mener data governance.	Actually three: Authentication, authorization, and accounting	Det var ikke en seksjon for risikohåndtering, og strategi som er en av de første tingene man bør definere. Fordi alt annet man gjør etter det vil være basert på det du har definert som dine største risikoer. Det vil si at man vil kunne akseptere å være på et lavt nivå på noen områder, men ha ønske om å være på høyere nivåer andre steder.
Synes det går helt fint med tre nivåer. Det gjør det lettere å vurdere hvor man er. Føler ordene som er brukt for å beskrivene nivåene er veldig bra.	Identity and access management process. Identity is the key part in Zero Trust architecture	Security champions – En security champion er en definer sikkerhetsansvarlig for hver avdeling. Så selvom de jobber i finansavdelingen eller markedsavdelingen så er de personer som har fått ekstra opplæring i sikkerhet og har ekstra fokus på dette.
Viktig å huske at det er en stor kobling mellom alle områdene. Det vil si at hvis du gjør noe med nettverkssystemene dine, så skjer det endringer i asset management og change management er involvert. Så enten så skjer dette manuelt, eller så skjer det automatisk.	Vague statement.	At man har aksjonsbasert trening basert på employee behavior. Så hvis du har gjort noe som ledet til en sikkerhetshendelse så får du trening basert på det du har gjort feil.
Det er lett å glemme at alle disse områdene henger sammen, så det er viktig å huske på det større bildet. Så hvis en bedrift er helt optimal, så snakker alle disse tingene sammen automatisk. Tradisjonelt så snakker de ikke sammen og man gjør alle endringer manuelt.	… from outside and over internet.	Kanskje si noe om behavior analytics – som er veldig optimalt i dag.
	Policy enforcement engine	Og så i 5.2 så har dere et punkt som sier Hva ansatt må gjøre og hvordan som er egentlig veldig bra. Jeg hadde bare skrevet det ut litt mer, der jeg hadde skrevet at en policy er hvorfor, en prosess er hva de skal

Structure Improvements	Corrections and Consistency	Focus Area Improvements/Suggestions
		gjøre og hvem og hvor de skal gjøre det, og en procedure er hvordan de gjør det, steg for steg. Det er de tre forskjellige nivå.
	Logged and reviewed	AD fungerer ofte som en form for backend for PEPs; hva skjer om hele AD ryker? I ransomwarehendelser og rettede angrep, ryker normalt AD => hvilke konsekvenser får dette for helheten?
	InfoSec har en lang historikk med 'løsninger' som skal gjøre ting sikkert (AV->IDS->IPS, etc), men som allikevel kommer til kort på grunn av at trusselbildet tilpasser seg. Kanskje lurt med en litt mer kritisk tilnærming?	Hadde vurdert å legge på en break glass prosess. Så selv om du er en ingeniør så har du ikke mulighet til å gjøre endringer på din vanlige konto, men må logge inn på en separat konto.
	'The Network' gir ikke nødvendigvis mening i en verden der perimeteret er relativt porøst – vi må anta at alt kan kompromitteres – endepunkter, identiteter, løsninger.	Asset management – hadde vurdert å skrive noe om asset classification og ranking basert på risikoer. Det er vanlig at bedrifter har sine assets, med hardware og applikasjoner i en database. Når endringer blir gjort i nettverket, eller en applikasjon blir oppdatert så oppdaterer verktøyet denne databasen. Et steg videre vil være å ha rangert assets etter kritikalitet og definere hva er mine crown jewells, hva har jeg ikke råd til å miste et sekund, eller en time.
	Det er en lang historikk før Kindervag; Open Groups Jericho Forum ikke minst, men spor tilbake til iallefall 70-tallet.	
	Dere bør kanskje iallefall nevne Forresters Zero Trust eXtended (ZTX) fra 2018 slik at det ikke forveksles?	
	Vet ikke helt om jeg kjøper denne påstanden; selv for 20+ år siden var det aktiv soneinndeling av også interne nettverk.	
	Som betyr? DPI eller lignende => drakamp mellom kryptert trafikk (som vi	

Structure Improvements	Corrections and Consistency	Focus Area Improvements/Suggestions
	ønsker) og mulighet/enkelhet på inspeksjon ut over metadata (src, dst, port, etc.)	
	Ransomwareangrep er generelt ikke spesielt sofistikerte, de bare gjenbruker teknikker som fungerer alt for godt i mange organisasjoner. Normal flyt er noe lateral bevegelse east/west, men så rask tilgang til domain admin; derfra bli det top-down-deployment, f.eks. via GPOer.	
	Skeptisk til dette begrepet; av hvem og hvordan må nesten besvares.	
	Hvem har ansvar for å modellere/beskrive disse?	
	Ofte vil legacy-løsninger ha et assortert utvalg avhengigheter til andre systemer (DNS, DHCP, NTP, oppdateringsservere etc.), og litt usikker på hvor mye verdi en ‘zero trust proxy’ vil kunne tilføre i realiteten.	
	Dette må man nok veldig langt tilbake i tid for å se; finnes fortsatt, men AD er kjernen i de aller fleste organisasjoner.	
	Dette impliserer at det vil være en menneskelig funksjon å overvåke => skalerer ikke. Om det er maskinelt, hvordan skiller det seg fra eksisterende løsninger?	
	Målet bør være at applikasjonene er tilgjengelig for bruker, uavhengig av de er; en direkte eksponering av applikasjoner er ikke et mål (alt som er eksponert må antas å kunne kompromitteres). Klassisk VPN er døende, men mer moderne løsninger	

Structure Improvements	Corrections and Consistency	Focus Area Improvements/Suggestions
	gjør dette mye mer transparent for brukeren, og man bruker i praksis bare Internet som transitt (e.g. WireGuard, TailScale, etc.)	
	Would it be possible to elaborate more on concrete tools in the measures? Be more specific, but not so specific that the information is outdated in a year.	
	Employee awareness and training – Veldig bra skrevet. Jeg følte at det var litt tradisjonelt i forhold til approachen. Det er litt vanskelig å skjønne hvis en ikke har vært ute for å implementere disse tingene. Tenker nok at trening og simulering av phishing er tradisjonelt, og ikke advanced eller optimal. Det som er mye brukt i dag er risiko og strategi-basert trening. Så hvis du definerer at dine største risikoer er malware f.eks så fokuserer man trening på innhold relatert til dette.	
	Det er viktig at man har en policy, en prosess og at man går gjennom policyen, men det er ikke nødvendigvis en del av kulturen. Et dokument i seg selv gjør ikke nødvendigvis noe for kulturen. Det statuerer kun hvordan man gjør ting og hva man forventer. Det som endrer en kultur er f.eks et awareness program hvor målet er å endre atferd. I forhold til hvordan vi definerer kultur så er det mindset, values og behaviours. En policy er en del av behaviour, fordi den sier hva man bør gjøre, men hvis man ikke har et program for skape awareness rundt policyen så vil ikke folk lese den.	

Iteration 2

Structure	Corrections and consistency	Focus Areas
The biggest problem with any model is the "SCOPE"…take this…what if I have network at Traditional and one at Advanced in my organisation. At which level is my organization? To continue. Ref above comment…how many is "some"…If I have 2 services segmented in my in network am I at advanced already? You got the "criticality" aspect, but then I would expect that all identified critical services are segmented.	In which state of projects do you think this model is most applicable? Should it be used before starting implementation or during the implementation of Zero Trust?	Look to the CIS Maturity Benchmark and "sjekkliste for informasjonssikkerhet og personvern" for inspiration when creating the spreadsheet for maturity assessment.
…these are all yes/no question…it that intended? You then have a check list more then "control questions" May be semantic…as long as it is what you want.	I would almost argue the past decade in this case, although the last years have made it even more clear.	Should identity be listed as a domain?
You make a good effort here…but try to focus on definition more the "judgment" and suggesting how to adhere to a defined maturity level.	…an infrastructure build with a secure perimeter …	Onboarding og opplæringsprosesser? Eller tenker dere at disse faller inn under people? Configuration management og hardening ligger kanskje inn under en av disse?
…you are expressing a judgement…focus on the consequences… e.g. lists are often inaccurate and requires…	What is Zero Trust and what was coined as a solution…Term? Concept? Idea? Model?	Hva med andre teknologier, sånn som identitet og tilgangsstyring som virker å være veldig sentral innenfor zero trust? At alt skal ligge bak en autentisering, rollebasert tilgangsstyring, least priviledge, ++
…a report giving suggestions? ;)	Difficult to read, missing a comma?	Savner litt mer om IAM her? Kanskje spesielt identity governance og å knytte

Structure	Corrections and consistency	Focus Areas
		autentisering opp mot andre verktøy som EDR og MDM?
So is PM part or not of this maturity level?	Generally, I like to write acronymous extentdet the first time in a document A la Cyber Security Framwork (CSF) You decide if this is relevant for you/this document	One point we are missing is still accounting. Should this be included here? In newer trust models, accounting is a crucial part. Read a bit about AAA (Authentication, Authorization and Accounting) and see if you would like to include this.
You should be describing the maturity level not advocating for it.	…here you put yourself in trouble…"partial" is already 1% of 100% and so is "partial" 99% of 100%. Try to avoid "some", "somewhat", "partial", not completely and so on. Here you might use for example "substantial", "existing", "consistent"	Could possibly include accounting here
However, to what?	Some degree of micro-segmentation? So is it possible to measure micro-segmentation? How?	Litt mer om identity governance?
May…is this part of the definition or not?	Does firewalls with inspection mechanisms (assuming that is what scrutiny of traffic using policy enforcement points refer to means)? Since all traffic is terminated on a host, can hosted based detection mechanisms and mutual authentication for services rather be emphasized?	Should we emphasize the need for telemetry and detection engineering in order to have trust in a zero trust network?
Where is the focus areas int the "description"? you may make them more explicit	Basic? …less adjectives make "definition" easier	Overall, I think this section is too 'network focused', the network vendors are doing a lot to control the narrative on zero-trust but it is important to consider their relevance in a zero-trust world if you look at it more conceptually

Structure	Corrections and consistency	Focus Areas
You describe a situation here? Where is the definition of the "level" and where are the focus areas	This can also be a vulnerability, if an organization with little to no security has exposed their services publicly without thinking about defensive mechanisms	I would suggest this to eb less network focus, more focused on detecting threats on based on telemetry. Machine learning may also be overrated in this scenario but can contribute.
…right…but then…so what? Where is the definition? How are we supposed to look for attackers in the network already?	Change to "Assume that users access all applications via untrusted networks, such as the internet"	This is the most challenging part where a lot of organizations struggle. Generally all businesses are very big on processes, but they do not have full control. All organizations have infrastructure (firewalls, proxies etc.), but they may be poorly configured. Look to ISO 27001 and 27002 for compliance-related topics and "soft security". Some control questions can be taken from there.
While important how is this (only?) relevant for supply chain Management	Is this a condition? This might be true also for traditional?	Should a risk assessment process be included?
Very personal opinion. The only thing that flies here is the "inventory, classification and access". Encryption is a "control" that may be part of all levels Data recovery may be a chapter in itself but need more "thoughts".	Which device compliant to what?	What about the configuration management process?
Where is zero trust here? What you describe makes sense also in terms of maturity level, but where is the "assuming compromise" or other "ideas" of zero trust? Only in Optimal you get some "hint" of zero trust, but yet you do not focus on how "zero trust" is applied to awareness and training.	I could have a few hours speech "demonstrating" threat landscape is actually quite static and possibly being stable for the last 2000 years…or so ;) What is ever changing is technology and techniques…and then I agree that reactive protection is no longer adequate.	I would have expected more here… like how the change management process can be designed so that it can go fast and still preserve the zero trust in the infrastructure. Without architects need to be involved to evaluate all change…. E.g. Taking some inspiration from you

Structure	Corrections and consistency	Focus Areas
		Different level of changes defined in advance that make possible to respect principles of Zero Trust. If change do not fit in the predefined the Architects are involved Optimal. Change process is automated and changes are automatically scanned for issues. Just a suggestion
Ref previous comment. This chapter make sense, but where is the "zero trust"?	This should be the case already in a well-designed change process without taking into account zero trust	Since the advanced section mentions user accounts, should this section mention service principals and how integrity of changes is maintained?
	So how is this relevant for this level?	By enforcing any change to be done 'as code' you can enforce reviews to preserve the integrity and deployments can strictly be limited to reviewed code, I would consider this for the Optimal level.
	I don't think multi-factor should be considered sufficient in most cases due to the long lifetime of a session token and potential for abuse on an asset management system prior to detection and response. Privileged Access Workstations (PAWs)/Secure Admin Workstations (SAWs) vastly reduce the risk and is considered a best practice	Also consider most of the software bought today is indeed "SaaS"…so how to secure those connections?
	I've personally seen the need for a war room become heavily scrutinized during the last 2 years of pandemic and now incident responders adapting to permanent remote. Workloads are also becoming virtualized and the need for performing physical forensics is slowly being	Diskutere litt mer rundt phishing øvelser? Kan i mitt hode ofte skade like mye som det forbedrer noe

Structure	Corrections and consistency	Focus Areas
	removed. I think handling of information and artifacts during an incident is important but may not necessarily be as physical as it used to be.	
	Threat emulation	I have strong opinions about the "usufullness" of these tools…reach out if you want to here why I mean the are indeed useless if not even damaging the security culture of an organization.
	Background info	Tilpasset opplæring? At feks økonomi trenger annen type opplæring enn HR
	…outgoing from where? I thought the perimeter was dead! ;)	
	Transmission=Transit? Or how are they diffent?	
	…could start a long digression about "sentitive"… Let's reduce to "all information of value to the organization".	
	Do you take a cut for the advertisement? ;)	

Iteration 2.5

Structure	Corrections and consistency	Focus Areas
	Be consistent…either short version first and then explanation…or the opposite…as you have it for CSF	….a bit of a punishment…I think better something like…the reason why the attack was successful is assessed and if needed specific training is developed…or something in the general "improvement" direction
	Formatting…keep together ;)	Partner identiteter, Servicekontoer er nevnt i liten grad.
	Security health or just "security"?...is connected to security health checks? You may want to check the definition…if any	Nevne noe i advanced om at IAM løsningen brukes til å federere og autentisere identiteter for applikasjoner både on-premise og i cloud løsninger?
	As is written here it seems this is not the case for other level…reading further you mean actually something else I believe	Nevne noe om forskjellen på personlige kontoer og ikke personlige tilganger.
	Logically not a second point…	Nevne noe at det gis rettigheter basert på rollen som hentes fra HR systemet.
	Implies a way to measure it	
	Very specific…usually "regularly" is used….	

APPENDIX D: FOCUS AREA CHANGES AND SUGGESTIONS

Focus Area	Status	Comment
Risk Management	Removed during iteration 0.	We found that risk management is an area where the Zero Trust principles would make little difference: A risk manager is already assuming breach and making plans and mitigations for the worst-case scenarios. We would instead argue that a risk-based approach should be taken when approaching any of the other focus areas. Discussions had with respondents in Iteration 1 supported this reasoning.
Hardening and Configuration Management	Not implemented. Considered during iteration 2	This suggestion was considered, and we found it to be better to implement aspects of hardening and configuration management in the asset management and supply chain management areas.
Patch Management	Not implemented. Considered during iteration 2	Patch management was mentioned as a possible focus area in iteration 2. We had already touched upon patch management in the asset management focus area. We decided to expand upon this and explain the cases where Zero Trust principles could improve patch management as a part of the asset management focus area. The respondent feedback on this change was positive.
Security Architecture	Not implemented. Considered during iteration 2	One of the respondents suggested including “Designing Security Architecture” as a focus area under Processes. We did research on the topic, tried to identify possible controls, and discussed how the Zero Trust principles could be applied. We ended up not including it as a focus area, as we did not manage to gather controls and content with enough emphasis on Zero Trust.
Network Segmentation and Infrastructure	Implemented since iteration 0	Network segmentation was identified early on during our initial literature review as a core component of Zero Trust and can greatly benefit from the Zero Trust principles.
Dynamic Access	Implemented since iteration 0	Dynamic access was yet another component of Zero Trust identified during out literature review.
Threat Protection	Implemented since iteration 0	Threat protection was also identified during the literature review on Zero Trust and can greatly benefit from Zero Trust principles.
Identity and Access Management	Implemented since iteration 2	Multiple respondents brought in for feedback on the second draft commented that identity was not covered well enough throughout the model. The initial plan was covering identity in the Dynamic access focus area, having sub areas for dynamic access and identity governance. However, we realized from the comments that the

Focus Area	Status	Comment
		comprehensiveness of identity governance required a separate focus area and therefore created this.
Change Management	Implemented since iteration 0	Change management was brought up during the informal discussions with colleagues during iteration 0. During these discussions, several approaches to implementing Zero Trust principles in change management processes were discussed, and so it became a focus area in our model.
Asset Management	Implemented since iteration 0	Asset management was discussed during our informal conversations with colleagues similarly to change management with the conclusion that asset management processes would benefit greatly from Zero Trust principles as well.
Incident Management	Implemented since iteration 0	Incident management is yet another focus area that was initially included after informal conversations with our colleagues.
Supply Chain Management	Implemented since iteration 0	“Third-party services” added as sub focus area after feedback saying the model lacked focus on cloud services, as most software today is offered through “software-as-a-service” solutions.
Data Governance and Protection	Implemented since iteration 0	Data governance and protection is a key process in every organization, and we found this to be an area where Zero Trust principles were highly relevant and could make a big difference. This was confirmed by respondent feedback during iteration 1.

APPENDIX E: TABLE WITH ORIGINAL EZTMM AND MATURITY EVALUATION SHEET FILES

Note: The file objects only function in the .docx version of this thesis. Use the download links if viewing the PDF-version. All files can alternatively be provided upon request by the authors.

Filename	Description	Version	Date	Comment	File	Download Link
EZTMM V1 First Release.docx	The EZTMM	V1.0	29.05.2022	The holistic maturity model developed (the artifact) during our research.	EZTMM V1 First Release.docx	Control click the arrow below to download ↓
EZTMM Evaluation Sheet V1 First Release.xlsx	The accompanying evaluation spreadsheet for EZTMM.	V1.0	29.05.2022	The maturity assessment spreadsheet accompanying the EZTMM. Open in Microsoft Excel for the best possible compatibility.	EZTMM Evaluation Sheet V1 First Relea:	Control click the arrow below to download ↓
EZTMM V1 First Release Color-Coded.docx	A color-coded version of the EZTMM with changed per iteration highlighted	V1.0	29.05.2022	Change color codes are: No color: I0 Yellow: I1 Green: I2 Red: I2.5 Blue: I3	EZTMM V1 First Release Color-Co	Control click the arrow below to download ↓

Appendix F: EZTMM

EXTENDED ZERO TRUST MATURITY MODEL

Jarand Jansen and Simen Tokerud

UNIVERSITETET I AGDER

1 INTRODUCTION

The past decade has shown that an infrastructure with a secure perimeter protecting a less secure core is ineffective against today's threats. Zero Trust has been coined as the solution to the problem and has transitioned from being mystical and exciting to being the model most companies aspire to adopt. The Zero Trust model traditionally suggests assuming that all networks, endpoints, identities and solutions are compromised, treating both internal and external requests equally. Trust is no longer implicit; it is earned through rigorous verification. While it is possible that Zero Trust can contribute to improved security, the model only addresses the weakness of implicit trust in a network. Some security researchers have called this network-centric approach to security architecture fundamentally flawed (SABSA, 2022). This is because the strong focus on network makes the model hard to use as an overall strategy. In this model we suggest applying the foundational principles of Zero trust to other areas of information security.

The EZTMM (Extended Zero Trust Maturity Model) is designed with two scenarios in mind: Organizations that would like to start their implementation of Zero Trust and organizations that have already started their implementation and need a way to evaluate their progress. By performing regular maturity assessments using our model, the organization's management can get a good picture of the organization's current and past Zero Trust maturity. If the model is used at the start of the Zero Trust implementation process, the organization can also define goals based on the maturity levels defined in the model and perform a maturity assessment to validate that the goals have been reached.

1.1 Zero Trust Principles

The fundamental principles of Zero Trust can be traced back to the origins of the internet. One example is the change introduced to RFC 1122 in 1989: "In general, it is best to assume that the network is filled with malevolent entities that will send in packets designed to have the worst possible effect." When John Kindervag introduced the term Zero Trust in 2010 the focus was to eliminate the idea of trusted and untrusted networks and see everything as untrusted. He introduced three foundational concepts: ensure that all resources are accessed securely regardless of location, adopt a least privilege strategy, strictly enforce access control and inspect and log all traffic. While these concepts describe practices likely to improve the majority of organizations' security architecture, we considered them too network-centric and hard to adopt in other areas. Microsoft has defined three similar principles to describe Zero Trust:

- **Verify explicitly:** Organizations that verify explicitly use all data and information available to reduce uncertainty and implicit trust.
- **Use least privileged access:** Using least privilege is always providing the least number of permissions necessary.
- **Assume breach:** Assuming breach is when you already consider your digital environment compromised.

1.2 The Extended Zero Trust Maturity Model

In this maturity model, we apply Microsoft's key principles of Zero Trust to organizational aspects of information security and combine those organizational aspects with the networking aspects of traditional Zero Trust maturity models. This allows organizations to perform a more comprehensive assessment of their overall Zero Trust maturity. We call this the Extended Zero Trust Maturity Model (Not to be confused with Forresters Zero Trust eXtended (ZTX) from 2018). The model is divided into three main domains: Technology, Processes and People. The domains are divided into focus areas with three maturity levels. Each focus area has accompanying control questions, allowing organizations to assess their own maturity.

1.3 Scoping and domains

The goal of the Extended Zero Trust Maturity Model is to take a more holistic approach to information security utilizing Zero Trust principles across the entire organization. When taking a holistic approach, the focus areas can be categorized within three different domains: Technology, Processes and People. The technology domain explains the technical solutions and components of a modern Zero Trust architecture. The process domain dives into many organizational processes and explains how to apply Zero Trust principles to these processes. The people domain sheds light on how the people of an organization can use Zero Trust principles to improve the organization's overall security posture. Below is a full list of focus areas within each domain.

Technology

- Network Segmentation and Infrastructure
- Dynamic Access
- Threat Protection

Processes

- Identity and Access Management
- Change Management
- Asset Management
- Incident Management
- Supply Chain Management
- Data Governance and Protection

People

- Employee Awareness and Training
- Information Security Culture

1.4 Identifying the focus areas

A literature study performed as part of the initial research laid the groundwork for identifying the focus areas. The research aimed to identify key components of Zero trust where the technical focus areas were discovered. Another noteworthy discovery was the lack of organizational aspects mentioned in both existing research and Zero trust maturity models. Thereafter multiple brainstorming sessions were held which resulted in a list of areas related to information security. Part of these sessions also included reviewing internationally recognized security frameworks for inspiration such as CIS(Center for Internet Security) Controls and NIST (National Institute of Standards and Technology) CSF (CyberSecurity Framework) among others. The Zero Trust principles were then tried on each area and its relevance was decided. The relevance of each focus area was then confirmed in several interviews with security professionals.

1.5 Defining the levels of maturity

The structure of our maturity model is heavily inspired by other maturity models such as the C2M2 (Cybersecurity Capability Maturity Model) model (Mehravari, 2015) and CISA's (Cybersecurity and Infrastructure Security Agency) Zero Trust Maturity Model, as well as feedback from our many respondents. As a result, we have chosen to define three levels of maturity: Traditional, advanced and optimal.

Traditional is meant to describe the traditional information security practices prior to implementing Zero Trust. The traditional level is often characterized by manual configurations and static security policies. Networks are often only segmented on the macro level with widespread implicit trust on internal networks. The policy enforcement is often proprietary and inflexible. Incident response and mitigation is done manually. The organizations have clearly defined policies, processes and procedures, but review intervals are limited, and no active awareness campaigns are utilized.

Advanced depicts an organization that has started the implementation of Zero Trust principles across the technology, processes and people domains. Centralized visibility and policy enforcement is implemented, incident response is partially automated through some pre-defined mitigations and the principle of least privilege is becoming more prominently adhered to. Some micro-segmentation of assets based on criticality is implemented, while egress and ingress is traffic is reduced to a minimum. Employees have a good grasp of the organization's security policies, processes and procedures, and these are regularly reviewed and updated.

Optimal describes the current ideal situation. The optimal state is ever-changing, and an organization will never truly be done implementing the Zero Trust principles in the most effective manner. In this maturity level, configuration and attribute assignment is fully automated. Access to resources is granted dynamically considering numerous factors including devices' security posture, threat intelligence, previous logging behavior, authentication, and authorization. Least privilege is dynamically enforced through open standards for interoperability across focus areas. Centralized visibility complete with extensive logging allows for point-in-time-recollection of state. The organization's information security culture is in focus and anchored in the organization with high awareness amongst employees. Employees are highly capable of scrutinizing information requests and mindful of where and when they discuss sensitive information.

CONTENTS

1 INTRODUCTION 2

1.1 Zero Trust Principles 2

1.2 The Extended Zero Trust Maturity Model 3

1.3 Scoping and domains 3

1.4 Identifying the focus areas 4

1.5 Defining the levels of maturity 4

2 TECHNOLOGY 9

2.1 Network Segmentation and Infrastructure 9

2.2 Dynamic Access 13

2.3 Threat Protection 17

3 PROCESSES 19

3.1 Identity and Access Management 19

3.2 Change Management 23

3.3 Asset Management 26

3.4 Incident Management 30

3.5 Supply Chain Management 34

3.6 Data Governance and Protection 39

4 PEOPLE 42

4.1 Employee Awareness and Training 42

4.2 Information Security Culture 45

5 REFERENCES 49

FIGURES

Figure 1: Micro-segmentation 10
Figure 2: Legacy System encapsulation... 12
Figure 3: Using a PEP for determining access... 15
Figure 4: Example of a Supply Chain Attack 35
Figure 5: Information Security Culture Illustration 46

TABLES

Table 1: Network Segmentation and Infrastructure Maturity Levels 9
Table 2: Network Segmentation and Infrastructure Control Questions 13
Table 3: Network Access Maturity Levels 13
Table 4: Zero Trust Network Access Control Questions 16
Table 5: Threat Protection Maturity Levels 17
Table 6: Threat Protection Control Questions 18
Table 7: Identity and Access Management Maturity Levels 19
Table 8: Identity and Access Management Control Questions 22
Table 9: Change Management Maturity Levels 23
Table 10: Change Management Control Questions 25
Table 11: Asset Management Maturity Levels 26
Table 12: Asset Management Control Questions 29
Table 13: Incident Detection Maturity Levels 30
Table 14: Incident Detection 33
Table 15: Software Supply Chain Management Maturity Levels 34
Table 16: Software Supply Chain Management Control Questions 38
Table 17: Data Protection and Governance Maturity Levels 39
Table 18: Data Protection and Governance Control Questions 41
Table 19: Employee Awareness and Training Maturity Levels 42
Table 20: Employee Awareness and Training Control Questions 45
Table 21: Information Security Culture Maturity Levels 45
Table 22: Information Security Culture Control Questions 48

2 TECHNOLOGY

In the technology domain we will explore technical solutions and components in a modern Zero Trust architecture. The technology domain is at the core of any Zero Trust architecture.

2.1 Network Segmentation and Infrastructure

Focus area	Traditional	Advanced	Optimal
Network Segmentation and infrastructure	• Perimeter-based security • Legacy applications have no added security	• Limited egress and ingress traffic • The most critical internal services are micro-segmented	• Micro-segmentation based on application workflows • Full usage of micro-perimeters for ingress and egress traffic • Encapsulated legacy systems

Table 14: Network Segmentation and Infrastructure Maturity Levels

2.1.1 Background

Traditionally, network segmentation and infrastructure were configured around a perimeter-based model: Internal traffic is generally trustworthy and only egress or ingress traffic is decrypted if necessary and inspected as it passes through the perimeter. Further segmentation of the internal network happened on the macro-level, often through different VLANS (Virtual Local Area Networks) and VRFs (Virtual Routing and Forwarding) with varying attributes.

The perimeter-based model has proven ineffective at preventing threat actors already on the inside of perimeter from moving laterally (East-West movement). This is particularly noticeable in today's ransomware and supply chain attacks, where an attacker can spread their payload to large parts of the internal network with relative ease and with little need for employing sophisticated methods. This is usually achieved through lateral movement, followed by a top-down deployment after obtaining domain admin access. Furthermore, due to the increased utilization of remote-working services, the perimeter is becoming increasingly difficult to define and control. One solution to these problems is micro-segmentation, where all traffic regardless of its origin and destination will be heavily scrutinized by a Policy Enforcement Point described in the chapter on Dynamic Network and Application Access.

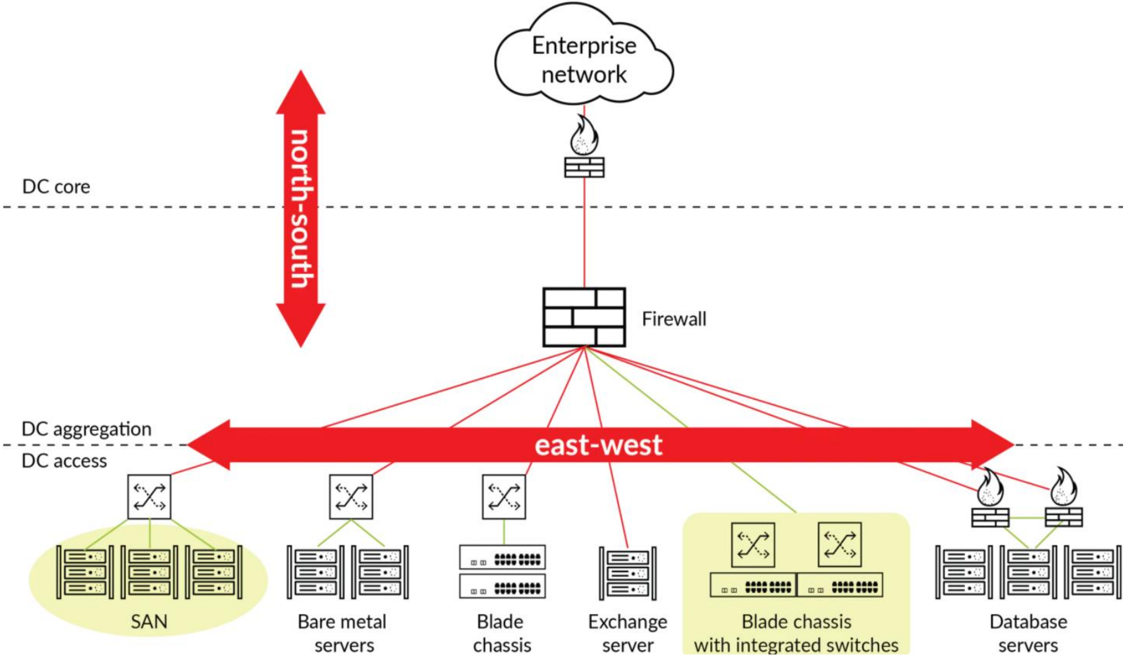

Figure 7: Micro-segmentation

Retrieved from https://www.paloaltonetworks.com/cyberpedia/what-is-microsegmentation

2.1.2 Traditional

Network segmentation is largely done at the perimeter, with little to no micro-segmentation. Internal networks may be macro-segmented utilizing different VLANS, VRFs or additional hardware to suit the organization's needs. An example of such macro-segmentation is the separation of development, test and production environments. Legacy applications have no added security and trust internal traffic implicitly.

2.1.3 Advanced

The organization’s most critical assets are micro-segmented, with particular emphasis on internet-exposed services. The organization’s less critical assets remain macro-segmented. The micro-segmentation of critical assets is handled by physical firewalls that will decrypt and analyze traffic if necessary.

2.1.4 Optimal

The entire network is micro-segmented utilizing virtual or physical firewalls for each computing resource or workload. The micro-segmentation is done based on application workflows, with full scrutiny of both internal and external traffic using Policy Enforcement Points. The application workflows are described through traffic analysis done by policy designers.

In cases where the network traffic is heavily encrypted, the organization has transitioned from traffic inspection on the firewall to using agent-based endpoint detection mechanisms. Since all encrypted traffic is terminated on the recipient, the bottleneck caused by firewalls having to decrypt all traffic that passes through them is removed.

Legacy systems are encapsulated, allowing modern access control and strong authentication to be used when accessing the devices. Encapsulation tailored for each legacy system depending on its dependencies on DNS, DHCP, NTP etc. Below is an example of legacy system encapsulation using forward and reverse proxies deployed on the device running the apps.

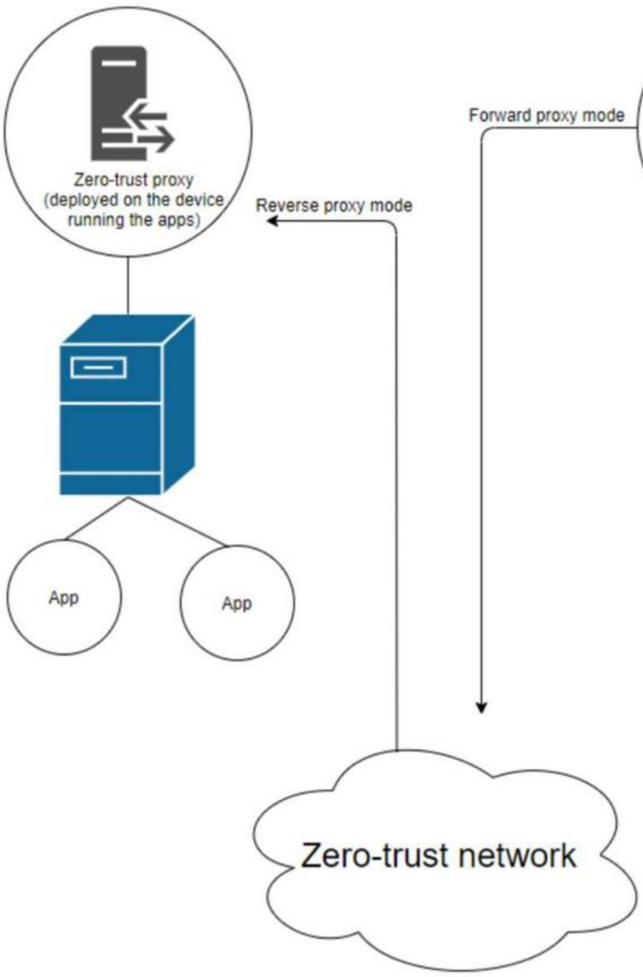
Zero-trust proxy (deployed on the device running the apps)
Reverse proxy mode
Forward proxy mode
Zero-trust proxy (deployed on the device running the apps)
App
App
App
App
Zero-trust network

Figure 8: Legacy System encapsulation

Retrieved from Trust No One? A Framework for Assisting Healthcare Organizations in Transitioning to a Zero-Trust Network Architecture (Tyler and Viana 2021)

2.1.5 Control Questions

In the event of multiple networks with differing security measures, the control questions work best when applied to each network individually. Depending on each network's purpose and need for security, the organization can choose to either focus on improving the security of the lowest scoring network until they are all equally secure or consider different maturity levels as acceptable for different networks.

Level	Question
Traditional	1. Does your organization have a clearly defined perimeter and macro-segmentation of the network?
Advanced	1. Does your organization only allow the minimum required ingress and egress traffic? 2. Are internet exposed and critical services micro-segmented?
Optimal	1. Is micro-segmentation applied throughout the network based on application workflows? 2. Are micro-perimeters implemented for all ingress and egress traffic? 3. If the network is heavily encrypted, are agent-based endpoint detection mechanisms utilized? 4. Are legacy systems encapsulated, allowing modern access control and authentication?

Table 15: Network Segmentation and Infrastructure Control Questions

2.2 Dynamic Access

Focus area	Traditional	Advanced	Optimal
Dynamic access	• Access decisions are not centralized • Mainly authenticating identities with passwords • Most on-premises applications are accessible through VPN	• Centralized policy engine used to make access decisions • Multi-factor Authentication • Device compliance with specified security policy	• Access decisions are continuously reviewed and verified • Password-less authentication • Assume that users access all applications via untrusted networks, such as the Internet

Table 16: Network Access Maturity Levels

2.2.1 Background

NIST defines Zero Trust in the following way: “Zero Trust (ZT) provides a collection of concepts and ideas designed to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services in the face of a network viewed as compromised." It is clear from this definition that making good access decisions and removing uncertainty from this process is the main goal of Zero Trust. This was also the big selling point when Kindervag first introduced the term in his article "No More Chewy Centers: Introducing the Zero Trust Model of Information Security".

If we dive into the tools for removing uncertainty, the major components are authentication, authorization, and accounting. Authentication is the process of providing proof of the claimed identity and the most common form is when the claimer provides a username with an associated password. The main purpose of authorization is defining required permissions to access resources and then enforcing this for access requests. In Zero Trust, authorization is based upon the principle of least privilege, and access is granted on a need-to-know basis, which means permissions and access to data are limited to only what is required for someone or something to perform their operations. Accounting is storing information on each user's consumption of resources, and what actions they perform during access. It is the combination of authentication and authorization put into a system supported by other context-enriching tools that enable dynamic access decisions for every request. Access can be accepted or denied based on different factors such as identity and credentials, device compliance, geographic location, previous access information, and more. However, being able to fully leverage this technology requires organizations to have strong control of identity of users and applications both on-premises and in their cloud solutions.

2.2.2 Traditional

Traditional organizations often do not have any central system for authorization, but instead, hand this responsibility over to the individual applications. The applications perform a check against some static values and the decision for access is made only once. Identities are stored and managed in identity providers on-premises. Identities are authenticated using weak authentication methods, usually username and password. Some external endpoints like Virtual Private Networks (VPNs) might require additional factors when authenticating, however, very few applications are exposed on the internet. This means remote users must connect to the on-premises network through a VPN to access most applications.

2.2.3 Advanced

A centralized policy engine is the heart of the process of granting access. The system should take multiple signals into consideration when making an access decision. Most important is authentication, authorization, and device security compliance. However, access policy compliance is only enforced on the first access, and not verified continuously. Access to sensitive data is only granted to managed devices to maintain sufficient security on the connecting devices. Security health checks may consider patch level of device and existence of an endpoint detection and response tool. The figure below shows the concept of access requests going through a policy engine, becoming trusted by evaluation and verification.

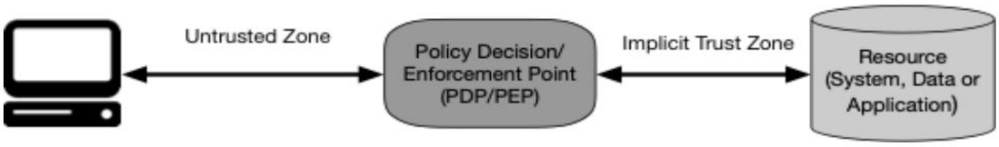
Untrusted Zone
Policy Decision/ Enforcement Point (PDP/PEP)
Implicit Trust Zone
Resource (System, Data or Application)

Figure 9: Using a PEP for determining access

Retrieved from NIST Special Publication 800-207: Zero Trust Architecture (NIST 2020)

Multi-factor authentication is enforced for all services, but Single Sign-On (SSO) can be enabled to improve workflow. Privileged roles should be limited, and just-in-time access should be utilized to follow the principle of least privilege. The organization federates user and application identities with cloud and on-premises solutions. Most of the applications on-premises are internet-facing but some are still only accessible through VPN.

2.2.4 Optimal

Advanced policy engines consider additional context-based signals in the decision-making process, including data on previous behavior collected for accounting purposes and threat information. Instead of only granting access at first request, continuous evaluation and verification are performed. This is also true for device compliance and security health, meaning devices are checked for compliance in real-time for every session. If there is a change of context during the connection access may be revoked. Users prove their identities with multi-factor password-less authentication methods. It is assumed that users access all applications and services from untrusted networks, such as the Internet. Emphasis is therefore placed on keeping applications up to date and behind defensive mechanisms such as web application firewalls and application proxies.

2.2.5 Control Questions

Level	Question
Traditional	1. Is authorization required for application access? 2. Is at least one authentication factor used to authenticate users? 3. Is it possible for remote users to access internal services securely, for example via a VPN?
Advanced	1. Is a centralized policy enforcement engine used to make access decisions? 2. Is MFA used for identity authentication? 3. Is the health and antivirus status of the requesting device considered in the access decision?
Optimal	1. Are access decisions being logged and reviewed continuously in real-time? 2. Have usernames and passwords been replaced with password-less authentication methods? 3. Are directly available to users regardless of logical and physical location in a secure manner?

Table 17: Zero Trust Network Access Control Questions

2.3 Threat Protection

Focus area	Traditional	Advanced	Optimal
Threat Protection	- Static traffic filtering - Known threats	- Basic analytics identify new and unknown threats - Use of End Point Detection and Response agents for critical assets	- Machine learning used for threat identification - Dynamic traffic filtering based on context - Use of End Point Detection and Response agents for all assets

Table 18: Threat Protection Maturity Levels

2.3.1 Background

Threat protection has traditionally been about protecting the network from threats, often using databases containing known threat signatures or static traffic filtering configured on the firewalls by network technicians. As the threat landscape becomes ever more changing and complex, reactive threat protection is no longer adequate. This has given rise to threat analytics and the use of machine learning to identify new threats on the fly and dynamic traffic filtering based on context.

2.3.2 Traditional

Threat protection is largely based on static traffic filtering and known threats. The approach is purely reactive, and any newly discovered threats will need to be added to the database of known threats. The traffic filtering is done manually through firewall configuration.

2.3.3 Advanced

Basic analytics is deployed to proactively discover new threats. The combination of proactive threat analytics and reactive static traffic filtering and known threat databases provides significantly better coverage than a purely reactive approach. The dark web is regularly checked for any signs of data leaks that may relate to the organization.

Threat protection is further enhanced by deploying End Point Detection and Response (EDR) agents for critical assets.

2.3.4 *Optimal*

Telemetry is heavily utilized in the threat detection process. The emphasis on telemetry allows organizations to detect potential threats as they establish connections to potential Command and Control (C&C) servers or try to download malicious code. Furthermore, usage data and other contextual data such as the time of access, the location the request originates from or the account that is used to make the request is analyzed to make a decision on whether to let the traffic through or not.

All compute resources that support it now have EDR agents with the exception of the employees' personal mobile phones (if they are provided by the organization).
Machine learning is leveraged to have significantly more accurate proactive threat detection. The machine learning algorithms continuously improve their detection rate by analyzing data, learning threat signatures and predicting other similar threats. Traffic filtering now considers context-based signals such as application workflows.

2.3.5 *Control Questions*

Level	Question
Traditional	1. Does your organization perform traffic filtering? 2. Does your organization utilize a database of known threats for threat protection?
Advanced	1. Does your organization deploy analytics to proactively discover new threats? 2. Does your critical assets have EDR agents installed?
Optimal	1. Does your organization utilize telemetry for threat detection? 2. Does your organization utilize machine learning to improve threat analytics? 3. Is the traffic filtering based on contextual data such as application workflows, telemetry and usage patterns? 4. Do all your assets except the employees' personal mobile phones have EDR agents installed?

Table 19: Threat Protection Control Questions

3 PROCESSES

The processes domain explores the organizational processes where Zero Trust principles can be applied to enhance information security.

3.1 Identity and Access Management

Focus area	Traditional	Advanced	Optimal
Identity and access management	• Central directory of identities • Manual process for granting access and permissions • Separate accounts for administrative tasks	• Permissions granted following principle of least privilege • Time-limited roles and permissions • Manual access reviews	• Automated access reviews performed periodically • Account activity is monitored and inactive accounts are deactivated • Penetration tests to harden the Identity and Access Management solution • Just-in-time activation of privileged roles

Table 20: Identity and Access Management Maturity Levels

3.1.1 Background

Confidently granting access dynamically to resources in the network requires good Identity and Access Management (IAM). This includes having full insight into all existing identities in the organization, be it employees, applications, or machines. Additionally, each of these identities has its own rights and permissions that must be managed. An identity and access management solution is often leveraged to manage this process. The IAM solution is preferably fed data from an HR (Human Resources) system, serving as a source of truth. The reason for using the HR system is its' detailed information on all employees, their full name, department, role, start and stop date, and so on. Having information on roles for all users enables easier granting of permissions and privileges in the IAM solution.

After compromising legitimate user accounts, attackers will start examining what permissions they have and look for servers and systems accessible to them. The principle of least privilege is central in the prevention of this and can severely reduce the impact of such compromises. Following this principle, identities should only have access to what is required at minimum to perform their purposed operations. A database administrator should have access to work on databases, and employees in HR have the right to administrate the HR systems. These permissions should however not be mixed. A challenge often occurring is when individuals stay at an organization for many years and take on different roles. In these cases, it can happen that permissions are kept from previous roles, and they end up having very privileged accounts. To prevent this, organizations can establish frequent access and permission reviews where employees must justify the further need for their current permissions and roles.

A possible threat is former employees who are not offboarded properly and thus still have access to the organization's systems and resources. The employee may use this access to cause damage to assets or steal information. A proper offboarding process can be used to mitigate this threat. It is also recommended to track accounts for inactivity to detect employees who are no longer with the organization or service accounts no longer in use.

Assuming accounts will be compromised is also a good reason for having separate accounts for performing administrative tasks. Accounts that are used for business purposes such as e-mail and browsing the internet have a higher likelihood of compromise and should therefore not be given privileged roles. However, using separate accounts for admin tasks is no guarantee against compromise, and privileged roles should therefore not be given permanently. Instead, roles should be available for activation by using multi-factor authentication.

3.1.2 Traditional

The organization keeps a record of all users, applications, and machines in a central directory service. There is a manual process for granting access and permissions to identities. Instead of limiting permissions to only what is required, extensive permissions are granted to guarantee sufficiency. Users perform all tasks using the same account. Lacking off-boarding processes may enable previous employees the ability to access internal systems after leaving the organization. Employees changing roles keep their permissions, eventually ending up with highly privileged accounts. Separate accounts are used for administrative tasks.

3.1.3 Advanced

The IAM solution is used for federating identity and authorization for applications, making it central in the process of granting dynamic access to applications and systems. Permissions are granted to identities through a formal and automated process where approval from one or more parties is required. Project-related roles are time-limited, meaning they will expire when the project ends. Business justification must also be provided when requesting new permissions. There is a strong focus on only providing identities with the minimum required permissions for performing purposed operations. Permissions are automatically granted to employees based on their role, which is information collected from the HR systems. Automated processes are established for employees changing positions, joining, or leaving the organization.
Manual access reviews are performed regularly with the objective to remove permissions that are no longer necessary based on the employee’s role and tasks.

3.1.4 Optimal

Automated access reviews are performed periodically where the employee or owner of a service account must justify why their roles and permissions are still needed. Accounts are continuously monitored for inactivity, and dormant accounts are deactivated.

Privileged roles are not standing, meaning they must be activated when needed, which may require additional authentication. This process if often referred to as just-in-time activation. If the role is required for a project or task, eligibility should be set for a specified period.

Both automated and manual penetration tests are performed regularly to investigate the available paths for attackers after successfully compromising an account. Automated tests are great for finding weaknesses that can easily be identified by scanning tools, while manual tests are performed by domain experts to identify weaknesses with higher complexity. The results are used to harden and improve the security of the IAM solution.

3.1.5 *Control Questions*

Level	Question
Traditional	1. Is a central directory service used to manage identities? 2. Is there an existing process for granting identities permissions and roles? 3. Are separate accounts used for administrative tasks?
Advanced	1. Are permissions granted following the principle of least privilege? 2. Does the process for granting permissions require business justification and approval from manager and/or system owner? 3. Are permissions time-limited? 4. Are manual access reviews performed regularly?
Optimal	1. Are access reviews automated and run regularly? 2. Are automated and manual penetration tests performed and used to harden the IAM solution? 3. Are just-in-time activation leveraged for permissions and roles?

Table 21: Identity and Access Management Control Questions

3.2 Change Management

Focus area	Traditional	Advanced	Optimal
Change Management	• Formally defined change management process for regular and urgent changes • The people allowed to request changes are clearly defined	• Criteria for compliance with Zero Trust principles and security requirements are defined for changes, non-compliant changes reviewed by security architects • Functional accounts are used to make changes	• Changes are handled using Configuration or Infrastructure as Code pipelines • The strict pipeline review process replaces the architect reviews in previous maturity levels.

Table 22: Change Management Maturity Levels

3.2.1 Background

Managing changes in the organization is about controlling the changes and making sure that they go through the correct approval processes. Changes can be initiated both internally (An employee suggests an improvement to the existing system architecture or a new component) and externally (A customer requests a change to infrastructure that your organization operates for them).

3.2.2 Traditional

A formal change management process is defined, along with a set of people who are allowed to initiate and request changes. Both the internal employees able to make changes and any customer representatives allowed to request changes are clearly defined. A special process is defined for handling urgent changes, bypassing or escalating some of the testing and quality checks of the change.

3.2.3 Advanced

Criteria for compliance with Zero Trust principles and the organization's security requirements are defined. Whenever a change fails one of these requirements, the change must be examined and verified by security architects to ensure that it is adequately secure before deployment. If a change request fails to adhere to the Zero Trust principles, or will otherwise result in reduced security, the request is denied or put on hold pending a workshop to improve the request and make it conform to the security requirements.
Furthermore, the number of people both internally and externally who can issue a change request is as low as possible, following the principle of least privilege.
When making changes, the administrators use functional administrator accounts specifically designed for making that type of change, not their personal accounts. An administrator's personal account has no elevated privileges.

3.2.4 Optimal

Changes are now handled using either Configuration as Code (CaC) or Infrastructure as Code (IaC) depending on the change. This means that any changes to the infrastructure or configuration goes through a pipeline that is defined with code. This pipeline is coded such that the proposed changes will have to adhere to Zero Trust principles and the organization's security requirements to be implemented. Only the code defining the pipeline is reviewed by the architects, freeing up resources and making it possible to implement changes faster and more securely.

CaC and IaC automate the security review process of each change, allowing it to be used even for emergency changes. However, based on the organization's needs, a more lenient review pipeline for emergency changes can be coded to ensure an even faster implementation of the change. If this is the case, the security of the change is further improved to meet the stricter requirements in the main pipeline as soon as possible after implementation.

3.2.5 Control Questions

Level	Question
Traditional	1. Does your organization have a defined change management process for both regular and urgent changes? 2. Are the employees allowed to make changes and the customer representatives (if applicable) allowed to request changes clearly defined?
Advanced	1. Are security criteria for changes clearly defined and based on Zero Trust principles, as well as the organization’s security needs? 2. Are internal and external change requests reviewed and verified by security architects when not conforming to the requirements in question 1? 3. Are the people allowed to request changes both inside and outside the organization limited according to the principle of least privilege? 4. Are functional administrator accounts used?
Optimal	1. Is IaC and CaC utilized for automated security reviews of changes? 2. If a less strict IaC or CaC pipeline for emergency changes is implemented, are these changes required to adhere to the stricter requirements of the main pipeline as soon as possible after implementation?

Table 23: Change Management Control Questions

3.3 Asset Management

Focus area	Traditional	Advanced	Optimal
Asset Management	• Manually updated asset inventory	• Automated asset inventory • Asset classification based on criticality • Separate patching regimes for critical assets • Limited access to Asset Management tools	• Asset Management tools are only accessible using privileged access workstations or secure administrator workstations • Changes to assets are verified automatically using a pipeline • Asset classification based on abstraction levels • Use of red teams

Table 24: Asset Management Maturity Levels

3.3.1 Background

IT asset management is the process of ensuring an organization's assets are accounted for, deployed, maintained, upgraded, and disposed of when the time comes (What is IT asset management (ITAM)?, 2022). Without having a complete overview of an organization's assets, it is impossible to define protect surfaces and perform the necessary network segmentation a Zero Trust architecture requires. However, many aspects of Zero Trust such as MFA can still be implemented without a complete server and software inventory. Implementing multiple improvements in parallel using a piecemeal approach is important for implementing a Zero Trust architecture in a reasonable timeframe.

Attack surface management can be seen as an extension of asset management where the organization approaches security from the attacker's perspective. Organizations generally employ red teams for this purpose. The red team utilizes various attack surface management tools to quickly find and close potential attack vectors that a real threat actor could exploit. Since the red team is employed by the organization and has easier access to internal information and a good overview of the infrastructure, they have a noticeable advantage over an external threat actor.

3.3.2 Traditional

In small to medium environments, asset lists or inventories are often maintained and created manually. One example of this is to use Microsoft Excel or similar spreadsheet tools to create a structured list containing information such as IP addresses, server names, FQDNs (Fully Qualified Domain Name) and operating system. This manual document is used as input for patch management. Strict and rigid documentation processes are implemented to avoid the asset inventory being out of sync with the real deployment. An asset list must be completely accurate for company to know exactly which assets need to be protected. Any inaccuracies in the asset list can lead to vulnerabilities not being patched properly and zombie servers.

3.3.3 Advanced

Asset management is done using an automated asset management tool in all environments. All newly created servers and decommissioned servers are automatically updated via the asset management tool. The risk or employees making mistakes or not following the strict documentation processes of the traditional level is removed. The inventory generated by the asset management tool is used as input for patch management, which are essential to keeping the organization's systems updated, reducing exposure to known and unknown exploits. Gathering data through agents on each server or client and storing them in a Configuration Management Database (CMDB) is a common way of doing this. Any access to the asset management system requires MFA.

Critical assets are tagged in the asset management tool to allow employees to immediately identify them. The top ranked critical assets can be defined as the organization's crown jewels. Any downtime on the crown jewels is highly detrimental to the organization and should be avoided at all costs. These critical assets are enrolled in a separate patching regime, with pilot testing being done prior to the rollout, removing the implicit trust in the software developers supplying the patches.

Following the principle of least privilege, only configuration managers and other employees who need to have access are granted access. Some employees may only need partial access to a specific system or group of systems, and others may only need read permissions for reporting purposes. The level of access each employee has to the asset management tool is reviewed regularly based on the organization's security requirements. Any permissions that are not strictly necessary are removed, and any employees changing roles within the organization or being offboarded will have their permissions reviewed.

3.3.4 Optimal

Any access to the asset management system is done through Privileged Access Workstations / Secure Admin Workstations (PAWs/SAWs). The number of people who can access an organization's asset management system is limited because of the high value such information will provide for a potential attacker.

The organization always assumes breach for all changes made through the asset management system. Therefore, the changes are all processed in a defined pipeline as described in chapter 3.2.4. Any changes to this pipeline must be reviewed and approved before they take effect. The pipeline ensures the integrity of any changes made.

In an effort to perform gap analysis and resolve any gaps, assets are further classified in abstraction levels such as data, component, system, value chain.

Furthermore, the organization employs red teams and to spot vulnerabilities and insecure configurations, unpatched systems/applications and other vulnerabilities from the attacker's perspective. These red teams work closely together with the asset management teams to remedy any potential attack vectors that are discovered.

3.3.5 Control Questions

Level	Question
Traditional	1. Does your organization have at least a manually updated asset inventory that is fully accurate? 2. Is the asset inventory used as input for patch management?
Advanced	1. Does your organization utilize asset management tools and to automatically keep the asset inventory updated and minimize the risk of human errors? 2. Has your organization implemented a CMDB to keep track of all assets? 3. Does your organization classify and rank assets based on criticality or other criteria? 4. Are separate patching regimes and piloting employed to secure the most critical assets? 5. Does your organization implement MFA for configuration managers? 6. Does your organization limit the number of people with access to the asset management system according to the principle of least privilege?
Optimal	1. Is access to the organization's asset management tool only granted when the request originates from a PAW/SAW? 2. Does your organization handle asset management using IaC and CaC? 3. Does your organization classify assets in abstraction levels to facilitate gap analyses? 4. Does your organization utilize red teams for finding insecure configurations, unpatched systems/applications and other vulnerabilities?

Table 25: Asset Management Control Questions

3.4 Incident Management

Sub-Focus area	Traditional	Advanced	Optimal
Security incident detection	• Establish logging and monitoring	• Assume breach detection capabilities.	• Red herring defenses.
Verification of detection	• Manual testing of detection capabilities	• Attack simulation	• Red and purple team exercises
Governance and information control	• Incident management team with clearly defined roles and access	• Restricted physical areas • Audit logging in security tools • Out-of-band communication during incidents	• Justification for accessing data seemingly not related to incident detection.

Table 26: Incident Detection Maturity Levels

3.4.1 Background

The ability to detect security incidents is an important part of any digital defense against cyberthreats. Failing to detect a security breach could make the impact of the incident a lot worse. When Kindervag (Kindervag, No More Chewy Centers: Introducing The Zero Trust Model Of Information Security, 2010) first introduced the term Zero Trust, one of the foundational concepts was to inspect and log all traffic. According to Verizon's 2021 Data Breach Investigations Report (Verizon, 2021) 20% of breaches analyzed were not detected before months had gone by. The zero-trust assumption of compromise requires us to think differently about how we develop our detection mechanisms. Many organizations spend large on security products and services and trust them to detect security incidents in their own environment. Unfortunately, there is no one-size-fits-all solution for security detection which is why trust in detection capabilities should not be implicit but gained with testing and confirmation. Also, security products often have access to large amounts of systems and sensitive data. This is a gold mine for attackers and malicious insiders, and access should be protected.

3.4.2 Traditional

Organizations have established security monitoring both for network traffic and endpoints of different kinds. Organizations collect and forward logs to a central location where security analysis can be performed. Security tools for endpoint detection and response are leveraged.

Detection capabilities are tested regularly. Manual testing is performed for some basic verification where adversarial behavior is simulated.

Security tools for detection and response have access to vast amounts of log data and end-client systems. Strict access control must be enforced, which requires a defined team of incident managers. Only members of this team will have access to the security portals but the principles of least privilege and need-to-know still apply. So even though security tools are very powerful and provide broad access, the security team only have access to the data and information that is required to do their job.

3.4.3 Advanced

Applying the assume breach mindset to the detection development is a game-changer. An organization that assumes that clients or network have been compromised is required to shift their focus away from trying to detect if someone is trying to get in. Instead, they are focusing on detecting behavior deeper into the attack chain. Typical behavior to look for is attackers performing internal reconnaissance, moving laterally, dumping credentials on endpoints, communicating with command-and-control servers, or already completing their objectives. This could be preparing and exfiltrating data or encrypting files as part of a ransomware attack.

Simulated attacks in a lab environment can be used to further increase confidence and reduce implicit trust in the detection capabilities. There are several open-source solutions that deploy virtualized environments combined with attack simulation test frameworks.

Communication between team members is key amid a security incident but secure communication is hard to achieve in an environment assumed to be compromised. Staying one step ahead of the adversary is hard if they can tap into the investigation. Severe incidents like a ransomware attack may also take down services used for communication, crippling teams' cooperation capabilities. Out-of-band services for communication and cooperation are therefore established and leveraged when deemed necessary by the incident response team.

Incident management teams work with sensitive information and often have a need for a visual representation of data. During incidents or handling of sensitive information, the incident detection and response team relocates to a physical area where access is controlled. Audit logs are stored to keep control of who accesses the physical area, who accesses which logs, and actions are performed in the security tools and services.

3.4.4 Optimal

The assumption is breached networks, and it is used as an advantage against the attackers. Red herring defenses can be used to distract attackers from their actual objectives, but also to make them step into a trap and set off the alarms. Honeypots are decoy systems in the network that appear legit, but their entire purpose is to lure attackers. Honeypots are deployed and whoever interacts with them is detected. The same principle can be used for sensitive files or files that appear sensitive before you open them. A benefit of using files is that they can also be used to expose malicious insiders.

Reaching the final stage of maturity requires regularly performing red team exercises to test detection capabilities. Red teams will simulate real threat actor activity and the ability to detect a red team can be used as an indicator of your capability’s efficiency. It is important from a Zero Trust perspective that the red team tests the organization’s ability to detect activity that implies compromise.

To gain better control of who accesses what data, analysts must provide justification on why access to data is required. This should only be applied to data that is seemingly not related to cybersecurity.

3.4.5 Control questions

Level	Question
Traditional	1. Are systems established to monitor and detect cybersecurity incidents? 2. Have detection capabilities been tested and verified? 3. Is there a clearly defined team working with incident detection and response?
Advanced	1. Are detection capabilities developed with the assumption of a breached network? 2. Has attack simulation been used to verify detection capabilities? 3. Are out-of-band services established and used in security incidents? 4. Is the incident detection and response environment physically separated from the rest of the organization? 5. Is audit logging enabled for security tools?
Optimal	1. Are red herring defenses part of the detection capabilities? 2. Have red team exercises been performed to test detection capabilities? 3. Is justification required to access data not related to security?

Table 27: Incident Detection

3.5 Supply Chain Management

Focus area	Traditional	Advanced	Optimal
Software development and dependencies	• Software dependency inventory	• Separate software environments • Vulnerability scanning of software dependencies. • Remove unnecessary dependencies	• Review software dependencies • Separate application by services.
Vendor purchased software	• Identify software in use • Download software from vendor using HTTPS.	• Only use third-party certified software • Security requirements for software vendors. • “Software bill of materials”	• Isolate and monitor all software
Third-party services	• Inventory of third-party services	• Security requirements for service providers • Assess security of service provider	• Continuous assessment of service providers • Cloud access security broker

Table 28: Software Supply Chain Management Maturity Levels

3.5.1 Background

Buying and developing software and services or outsourcing a part of the IT environment to a third party are both common practices in modern IT strategy. While making any of these decisions you are making a choice to trust one or more parties. If you are developing your own software, it is likely that you are using libraries or dependencies developed by others. When buying software, you trust a third party by running their code on your systems. In both cases, you must trust that their intentions are pure and that their set of security standards matches yours. Making use of third parties means increased risk because your attack surface is growing. If the company offering you software or services gets compromised, it could potentially mean you getting compromised. An attacker might alter the source code of legitimate software to gain unauthorized access and detecting this is usually a lot harder than detecting regular malware. Many software development companies maintain a great level of security, but that does not mean they cannot be compromised.

The research paper “Software Supply Chain Attacks, a Threat to Global Cybersecurity” (Durán & Jeferson, 2021) suggests the reuse of code being the main problem in software supply chain attacks. The author provides the following explanation: “... from 85% to 97% of the code currently used in the software development industry comes from the reuse of open-source code frameworks, repositories of third-party software and APIs, creating potential vulnerabilities in the development cycle of a software product”. Very often developers tend to import code written by others to perform simple tasks they could have written on their own or import large blocks of code while only using a small part of its functionality. Recent attacks and high amounts of code reuse show that this risk is real. The following section shows how applying the Zero Trust principles and mindset when working with cyber security supply chains can reduce this risk significantly.

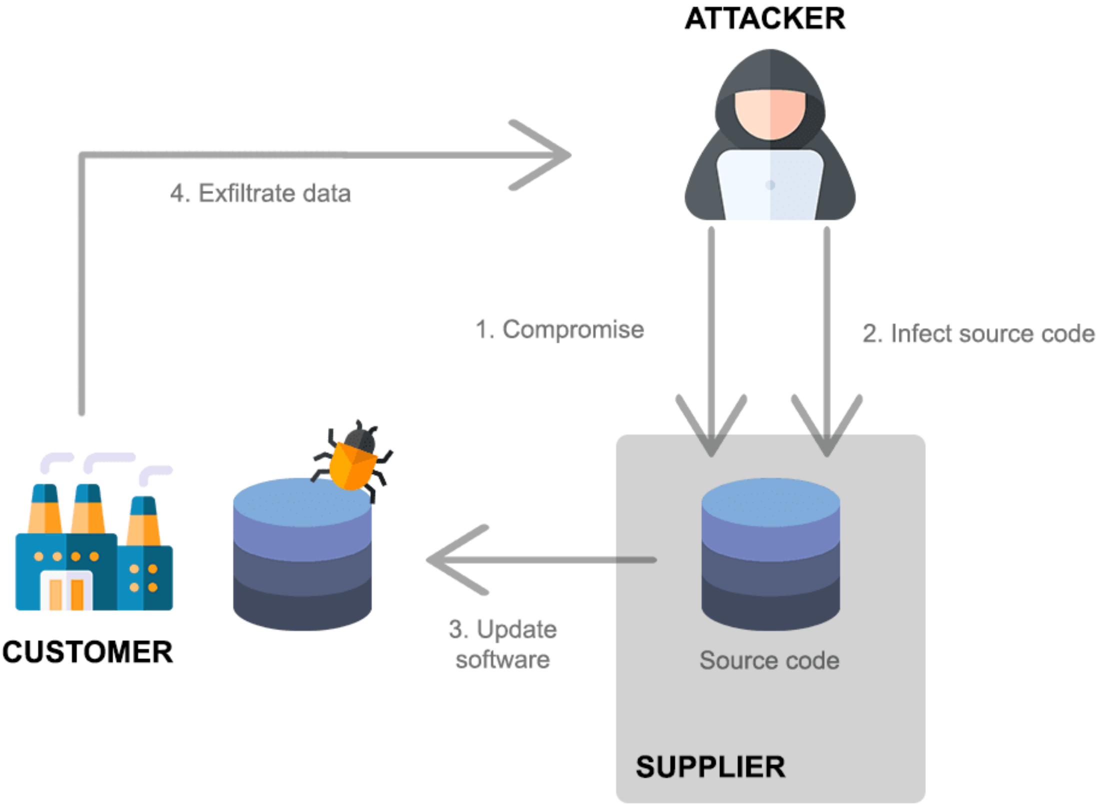

Figure 10: Example of a Supply Chain Attack

Recreated based on Threat Landscape for Supply Chain Attacks (ENISA, 2021)

3.5.2 Traditional

An inventory of software dependencies is established to identify existing vulnerabilities and reduce the risk of trust in the software supply chain. This is often done by storing source code on a central location preferably in a source code management service like GitHub or Bitbucket.

The risk of supply chain attacks is also present in software and services bought from external vendors. Organizations keep track of service providers and software in use, as a first step to mitigate this risk. It is also made sure that software is downloaded for official sources using HTTPS. Manual comparison of hashes is used to verify the integrity of downloaded software.

3.5.3 Advanced

Changes in dependencies can have a major impact on your systems. If a developer decides to delete a software package that you depend on, or one of your dependencies depend on, it could break the application. To mitigate this risk, organizations have established separate environments for development, testing and production. A set of different types of tests are ran in the testing environment revealing potential security and stability issues. It is great to combine these environments with a software composition analysis tool that can scan imported packages for known vulnerabilities and often includes different tests to reveal breaking changes to the application.

Using libraries and dependencies developed by others introduce trust to another party. This can in itself be problematic, but the real issue occurs when dependencies include dependencies of their own. This causes the supply chain to grow, as well as the number of trusted parties. To reduce this risk, organizations advocate the practice of removing unnecessary dependencies in code to their developers.

Buying software and third-party services from external vendors involve a certain amount of trust. With a Zero Trust mindset, we want to transform that trust from being given implicitly to something that is gained. Organizations may use two different practices to achieve this. The first practice is a control suggested in a security framework published by the Norwegian National Security Authority. It suggests that organizations “aim to only use software evaluated and certified by a third party. An example of such a certification regime is Common Criteria.” The second practice is defining a set of security requirements that external vendors must comply with when acquiring new software or services. CISA (CISA, 2021) suggests including some of the following requirements: Description of a software development lifecycle, vulnerability program, patch management capabilities and details on management of supplier lists. Reviewing providers' standardized assessment reports such as Service Organization Control 2 (SOC2) or using custom security questionnaires are also performed when acquiring new services. Purchased software should also include a “Software bill of materials” which is similar to a nutritional list and describes all the software components that make up the software.

3.5.4 Optimal

Organizations at the highest level of maturity may review the actual source code of their dependencies. As this is a very costly operation it might only be possible to do for security critical functionality like access control and encryption. Taking the source code and including it into their own code enable testing with their own tools and allow manual reviewing. An article titled “Secure Your Software Supply Chain – Threats and Mitigations” published by Truesec refers to this mitigation as “Vendoring” and claims it can reduce the risk of malicious publishers and supply chain attacks. The same article suggests separating applications into different services. Dependencies are often used to solve tasks in one part of an application. Separating these different parts into services reduces the risk of the dependency being compromised or deleted by the developer.

Assuming compromise is one of the foundational principles of Zero Trust and it can also be applied for software supply chains. Organizations that assume purchased software is or will be compromised implement controls for both prevention and detection. The software is allowed to run as intended, but least privilege principles are applied to prevent non-legitimate connections. This includes preventing outgoing connections except for destinations required to receive updates from the vendor and other expected traffic. A baseline of how information is flowing between the software and other systems is established to detect anomalies.

Similarly, a fully mature organization assumes its third-party service providers are compromised. Cloud access security brokers sit between the users and the cloud services to detect unauthorized exposure of information and non-compliant behavior. Organizations also continuously assess their service providers' security compliance, read release notes, and monitor the dark web for related leaks.

3.5.5 Control questions

Level	Question
Traditional	1. Is source code located in a central source code management service? 2. Has there been established a software inventory? 3. Is software downloaded from official sources using HTTPS?
Advanced	1. Are there different software environments for development, testing, and production? 2. Are software dependencies being scanned for known vulnerabilities? 3. Is there a goal of only using software certified by third parties? 4. Are security requirements set for software vendors? 5. Have unnecessary dependencies been removed?
Optimal	1. Is the separation of services in developed applications used when possible? 2. Is software being monitored and isolated like it has been compromised? 3. Is a Cloud access security broker in use?

Table 29: Software Supply Chain Management Control Questions

3.6 Data Governance and Protection

Focus area	Traditional	Advanced	Optimal
Encryption	• End-user client encryption	• Encrypt data at rest	• Encrypt all data at rest and in transit
Data inventory, classification and access	• Data inventory and file classification system • Data access control lists	• Sensitive data categorized and protected • Dynamic access control	• All data is inventoried and access to sensitive data is continuously monitored
Data recovery	• Automated backups	• Isolated instances of recovery data • Restore capability tests	• Immutable backups • Multi-user authentication for modification

Table 30: Data Protection and Governance Maturity Levels

3.6.1 Background

Data is a valuable asset for any organization and protecting its confidentiality, integrity, and availability is of high importance. Encryption is an efficient tool to protect the confidentiality of data, especially in a network we assume to be compromised. However, being able to protect data sufficiently requires us to know what data exists in the organization, where it resides, and its level of sensitivity. To be able to answer these questions an organization should establish a data inventory keeping track of their data. Knowing what data you need to protect, the next step is to control who accesses it, and categorize it based on sensitivity. Having these tools and processes in place enables us to grant access to data and applications following the least-privilege principle and on a need-to-know basis.

Availability of systems and data is a critical part of information security. Adversaries may perform unwanted changes to applications and systems, or even destroy data as part of a ransomware attack. There is also the risk of employees making mistakes or performing sabotage with intent. Data recovery measures are recommended to reduce said risk.

3.6.2 Traditional

With remote work on the rise and a shift from office desktops to laptops (Gartner, 2021) it is safe to say that devices are leaving the physical perimeter of the enterprise more than ever. Organizations trust their employees to protect the security of their devices, but controls are implemented to reduce the necessary amount of trust. Encryption of disks on end-user devices are leveraged to mitigate the risk of them being lost or stolen. Devices containing sensitive data are prioritized.

Organizations have established a basic data inventory and file classification system. Data inventories describe where to find what data. Rule and keyword-based methods are used to discover sensitive data. A file classification system is used to categorize data by labels. “Public”, “Internal” and “Confidential” are labels often used by organizations. Labeling of data is performed manually at this level of maturity. Having these labels makes it easier to treat data the same way regardless of location. Access control lists are used to ensure only authorized users can access data.

Organizations have established regular automated backups, especially for assets considered sensitive. In case of an incident the organization can roll back to the last known good state.

3.6.3 Advanced

Encryption is taken one step further by encrypting all data at rest, also counting data on removable devices. Data at rest is defined as data not currently being used, or in a state of transit. Organizations at the advanced stage of maturity account for, categorize and protect all information of value to the organization. Access to data is governed by a policy enforcement engine considering the context of the request. Modern information protection tools can protect your data with encryption and authorization policies. The applied protection measures will follow the data, so it is protected regardless of location. Many security incidents are caused by human error where an e-mail is sent to the wrong recipient or uploaded to a cloud service outside of the organization’s control. Protection measures as such are applied to data categorized as sensitive to prevent the occurrence of these types of incidents.

Assuming compromise implies greater risk of both incoming ransomware attacks and tampering of systems’ configuration and data. Organizations store backup data in isolated environments such as cold storage, separate cloud solutions or completely separated sites. Recovering the backed-up data is tested in regular intervals to make sure that both the backup and restore process is properly functioning. Soft delete functionality is used to protect against unintentional deletes and human error. Retention for the backup data is then set for a given period, in which deleted backups can be restored.

3.6.4 Optimal

All data is encrypted both at rest and in transit. That way data is protected from attackers performing man-in-the-middle attacks or gaining physical access to equipment. Files and databases are encrypted and masked, especially those containing sensitive data. Fully mature organizations classify all data regardless of sensitivity using machine learning and automated processes. Data that is classified as sensitive is continuously monitored to detect unauthorized access.

Immutable backups are used to protect critical data from ransomware, accidental deletion, data corruption and more. For less critical data backups are protected with multi-user authentication, requiring more than one administrator to make changes to configuration or on the backup data directly.

3.6.5 Control questions

Level	Question
Traditional	1. Is disk-encryption enabled for end-user clients? 2. Has there been established a basic inventory of data? 3. Is manual file-classification being performed? 4. Is access to data controlled with static or dynamic lists? 5. Is there a process in place for automated backups?
Advanced	1. Is data at rest encrypted? 2. Is sensitive data categorized and protected? 3. Is access to data granted dynamically? 4. Is backup data stored in isolated environments, and restore functionality tested regularly?
Optimal	1. Are data both at rest and in transit protected with encryption? 2. Is all data accounted for, categorized and protected regardless of sensitivity? 3. Is access to sensitive data being monitored? 4. Are Immutable backups used to protect critical data?

Table 31: Data Protection and Governance Control Questions

4 PEOPLE

The People domain explores how an organization's security posture can be further improved by providing training and increasing the awareness of its employees. Not all security risks can be adequately mitigated by technical measures alone, making this an important, but often overlooked domain from a Zero Trust perspective.

4.1 Employee Awareness and Training

Focus area	Traditional	Advanced	Optimal
Employee Awareness and Training	• Routines, policies and controls • Software aids	• General awareness training of employees to evaluate internal and external requests for information based on ZT principles • Targeted training that focuses on underperforming departments or individuals • Appointed security champions	• Risk and strategy-based training, where topics focus are tailored to the organization's largest threats. • Specialized awareness training of employees and departments to evaluate internal and external requests for information based on ZT principles

Table 32: Employee Awareness and Training Maturity Levels

4.1.1 Background

Employee awareness and training is essential when attempting to remove implicit trust from an organization: Most employees interact with their colleagues and exchange information many times a day. With the current pandemic and prominence of working from home, a lot of these exchanges happen digitally instead of face-to-face. Furthermore, employees often need to communicate with people outside the organization itself. This type of communication has led to an increase in successful phishing attacks and several data breaches. As a response, many companies employ phishing training tools and awareness campaigns to alert their employees to potential attacks from the outside. This training is a good complement to existing technical measures designed to prevent security incidents.

4.1.2 Traditional

The organization has good routines and policies for working securely and handling internal and external communication. New employees receive basic security training as a part of the onboarding process, and that all employees receive training whenever the policies or routines change. Controls are implemented to ensure that employees follow these routines and policies. Software aids such as alerting the employee whenever they are about to send an email to an external actor, or when they receive an email from such an actor assist in making the employee more alert when dealing with external actors.

4.1.3 Advanced

Awareness is continuously maintained through regular security training sessions and workshops for all employees. Here the employees learn the importance of assuming breach and never trusting, always verifying when receiving requests for information. This training applies to both internal and external requests for information. They learn to carefully consider who the requester is, and why they would need the information in their requested role. If the request seems suspicious, they are instructed to confirm the request by contacting the requester face-to-face or via video call to confirm that their identity has not been stolen and used to extract information. This applies regardless how the request was made (email, instant message, etc.). The employees are taught to ask themselves questions such as:

- Where does the requester work, and does it make sense for them to request this information from me?
- Why are they requesting this information from me?
- How can I be sure that the requester’s identity has not been compromised and used to exfiltrate information

Organizations utilize tools to visualize the progress made from awareness campaigns and for tracking the success rate of confirmed phishing attack against employees and departments. This data is used for extra training and awareness sessions targeted at specific employees or departments.

Furthermore, security champions are appointed for each department. These champions receive additional security training and are instructed to keep an eye out for suspicious employee behavior and malicious insiders. They will also provide additional training and advice based on the employees' awareness campaign results.

4.1.4 Optimal

The organization realizes that phishing requests targeted at the HR or economy departments are different when compared to requests targeted at the IT department. Furthermore, the departments mentioned need different kinds of training. Therefore, the regular workshops from the advanced level are now specialized for each department: These workshops provide clever phishing examples that are tailored to each department, along with learnings from previous phishing attempts and other security incidents of relevance for that department.

Risk and strategy-based training is employed, meaning that the training given to employees is tailored to the largest threats the organization faces. Security champions are responsible for action-based training, meaning that all employees whose actions led to a security incident will receive specific training to prevent the incident from reoccurring.

The organization also assumes that their employees have been phished. There are no repercussions for employees who admit to being phished, and any successful phishing attempts against employees result in the reason why the attack was successful being assessed and specific training being developed if needed.

4.1.5 Control Questions

Level	Question
Traditional	1. Are routines and policies for working securely implemented? 2. Does your organization have initial information security training for new employees? 3. Are software aids such as warnings when sending or receiving external emails implemented?
Advanced	1. Does your organization have regular, mandatory information security awareness sessions? 2. Does your organization utilize tools to maintain employee security awareness? 3. Is the statistical data from the tools used for targeted training of employees? 4. Does your organization have security champions in place for all departments?
Optimal	1. Does your organization have regular, specialized, mandatory information security awareness sessions for all departments? 2. Does your organization leverage risk and strategy-based training? 3. Is action-based training performed as a result of security incidents?

Table 33: Employee Awareness and Training Control Questions

4.2 Information Security Culture

Focus area	Traditional	Advanced	Optimal
Information Security Culture	• Defined secure processes and policies	• Regular reviews of secure processes • External audits	• Shared security-first mindset • Effective controls • Targeted responses to non-compliance

Table 34: Information Security Culture Maturity Levels

4.2.1 Background

Technical solutions struggle with preventing security incidents that occur because of a conversation or phone call between colleagues in a public place. This is an argument for why Zero Trust principles are so important, also in organizational aspects such as information security culture. Furthermore, this heightened awareness among employees is yet another layer an attacker must breach before gaining access to an organization’s secrets. Information security culture can be viewed as a sub-culture of the corporate culture, and consists of four levels: Artifacts, Espoused Values, Shared Tacit Assumptions and Knowledge (Van Niekerk & Von Solms, 2010)

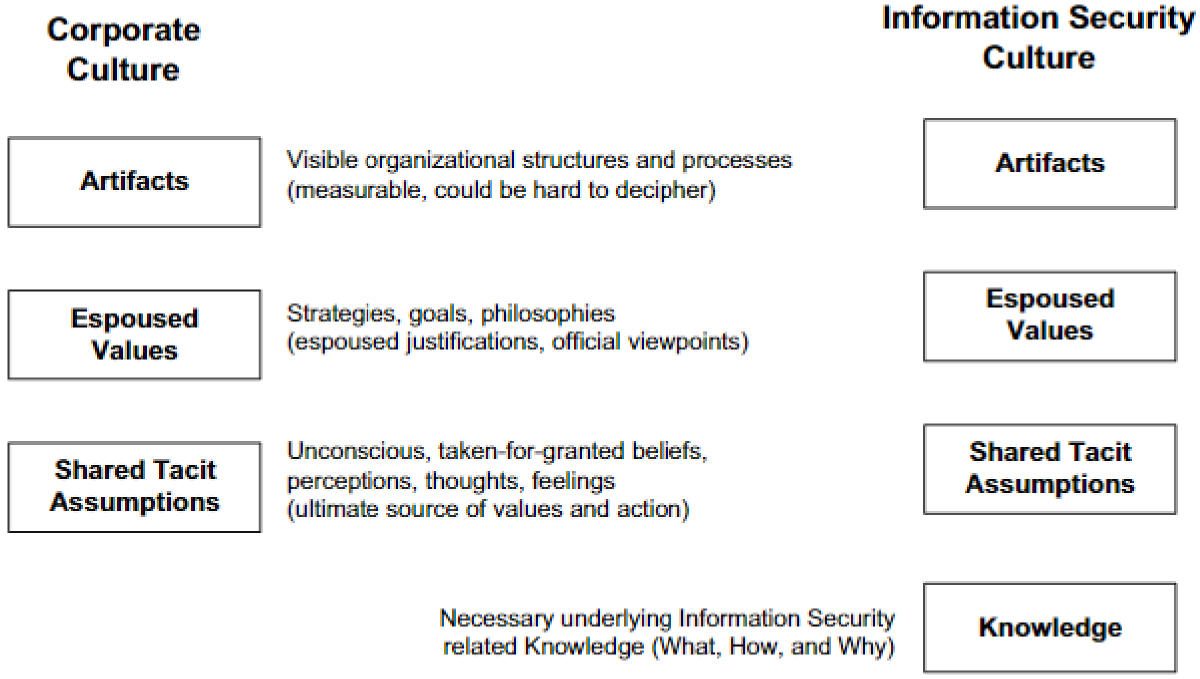
Corporate Culture
Information Security Culture
Artifacts
Visible organizational structures and processes (measurable, could be hard to decipher)
Artifacts
Espoused Values
Strategies, goals, philosophies (espoused justifications, official viewpoints)
Espoused Values
Shared Tacit Assumptions
Unconscious, taken-for-granted beliefs, perceptions, thoughts, feelings (ultimate source of values and action)
Shared Tacit Assumptions
Necessary underlying Information Security related Knowledge (What, How, and Why)
Knowledge

Figure 11: Information Security Culture Illustration

Retrieved from Information security culture: A management perspective (Van Niekerk & Von Solms, 2010)

4.2.2 Traditional

The organization has implemented policies, processes and procedures with at least some emphasis on information security. Input was gathered from internal security experts in this process. The policies detail the reasons **why** employees should adhere to the organization’s processes and procedures. The processes detail **what** the employees must do and **who** is responsible for doing it. Lastly, the procedures explain exactly **how** the task shall be carried out.

4.2.3 Advanced

Information security best practices are ever-changing, and the organization addresses this by doing regular reviews of its policies, processes and procedures. This updated knowledge is immediately shared with the employees, raising awareness around the organization's policies, processes and procedures. Knowledge sharing around information security contributes to building Shared Tacit Assumptions among the employees, assuring compliance with the organization's policies and processes. Following the principle of never trust, always verify, the organization's established Shared Tacit Assumptions are complemented with controls that prevent non-compliance with the organization's policies where possible.

Part of the review process described in the paragraph above involves limiting employee access according to the principle of least privilege: Employees should only have access to the data and tools that they require to perform their tasks. The reasoning behind the privilege removal must be communicated to the employees in such a way that they understand the necessity of following the principle of least privilege.

Consulting external security specialists when reviewing the secure processes or performing regular third-party audits can contribute to an even higher level of security and uncover weaknesses or faults in existing processes or controls.

4.2.4 Optimal

The majority of employees share the same security mindset and are updated on the latest security practices. They are very likely to comply with the secure policies, processes and procedures defined by the organization. Incentives are used to further increase compliance and create a sense of ownership to the organization's security posture both on an individual level and the department level. Furthermore, the incentives encourage even the employees that may not be completely aligned with the organization's security mindset to comply.

The organization assumes breach and knows that the possibility for malicious insiders exists. Therefore, the organization has implemented security controls designed to detect non-compliance with the policies, routines and processes. These controls make it harder for malicious insiders to exfiltrate information or cause damage to the organization, while also allowing for the detection of mistakes made by employees that result in non-compliance. The organization responds to these mistakes with targeted training as described in the chapter 4.1.3 and 4.1.4.

4.2.5 Control Questions

Level	Question
Traditional	1. Has your organization developed security policies and secure processes for all tasks considered important to the organization?
Advanced	1. Does your organization regularly review and improve the secure policies and processes? 2. Are external security experts consulted during the review process? 3. Are security audits carried out regularly?
Optimal	1. Are employees kept updated on the latest information security trends? 2. Are incentives leveraged? 3. Are controls implemented to detect and prevent non-compliance with the secure processes and policies? 4. Is non-compliance responded to?

Table 35: Information Security Culture Control Questions

5 REFERENCES

CISA. (2021). *Defending Against Software Supply Chain Attacks* . CISA.

Durán, J., & Jeferson, M. (2021). *Software Supply Chain Attacks, a Threat to Global Cybersecurity: SolarWinds' Case Study.* IETA.

Gartner. (2021, April 1). *Gartner Forecasts Global Devices Installed Base to Reach 6.2 Billion Units in 2021*. Retrieved from Press Release Newsroom: https://www.gartner.com/en/newsroom/press-releases/2021-04-01-gartner-forecasts-global-devices-installed-base-to-reach-6-2-billion-units-in-2021

Kindervag, J. (2010). No More Chewy Centers: Introducing The Zero Trust Model Of Information Security. *For Security & Risk Professionals*.

Mehravari, P. D. (2015). Evaluating and Improving Cybersecurity Capabilities of the Energy Critical Infrastructure. Waltham, MA, USA: IEEE International Symposium on Technologies for Homeland Security.

SABSA. (2022, 04 20). *The Attributer's Blog – Zero Trusted*. Retrieved from SABSA Enterprise Security Architecture: https://sabsa.org/the-attributers-blog-zero-trusted/

Symantec. (2019). *Internet Security Threat Report.* Symantec.

Van Niekerk, J., & Von Solms, R. (2010). Information security culture: A management perspective. *Computers & Security*.

Verizon. (2021). *2021 Data Breach Investigations Report.* Verizon.

What is IT asset management (ITAM)? (2022). Retrieved from Atlassian: https://www.atlassian.com/itsm/it-asset-management

CHANGE LOG

Version	Date	Change
0.5	13/02/2022	Initial Draft
0.6	31/03/2022	Second Draft: First round of comments from respondents addressed, changes and additions from V 0.5 highlighted in yellow (Color-coded version only).
0.7	24/04/2022	Third Draft: Second round of comments from respondents addressed, changes and additions from V0.6 highlighted in green (Color-coded version only).
0.8	30/04/2022	Review of identities with subject matter experts, changes and additions from V0.7 highlighted in red (Color-coded version only).
1.0	29/05/2022	First release: Case study feedback implemented, changes and additions from V0.8 highlighted in blue (Color-coded version only).

APPENDIX G: EZTMM WITH HIGHLIGHTED CHANGES PER ITERATION

EXTENDED ZERO TRUST MATURITY MODEL

Jarand Jansen and Simen Tokerud

UNIVERSITETET I AGDER

1 INTRODUCTION

The past decade has shown that an infrastructure with a secure perimeter protecting a less secure core is ineffective against today's threats. Zero Trust has been coined as the solution to the problem and has transitioned from being mystical and exciting to being the model most companies aspire to adopt. The Zero Trust model traditionally suggests assuming that all networks, endpoints, identities and solutions are compromised, treating both internal and external requests equally. Trust is no longer implicit; it is earned through rigorous verification. While it is possible that Zero Trust can contribute to improved security, the model only addresses the weakness of implicit trust in a network. Some security researchers have called this network-centric approach to security architecture fundamentally flawed (SABSA, 2022). This is because the strong focus on network makes the model hard to use as an overall strategy. In this model we suggest applying the foundational principles of Zero trust to other areas of information security.

The EZTMM (Extended Zero Trust Maturity Model) is designed with two scenarios in mind: Organizations that would like to start their implementation of Zero Trust and organizations that have already started their implementation and need a way to evaluate their progress. By performing regular maturity assessments using our model, the organization's management can get a good picture of the organization's current and past Zero Trust maturity. If the model is used at the start of the Zero Trust implementation process, the organization can also define goals based on the maturity levels defined in the model and perform a maturity assessment to validate that the goals have been reached.

1.1 Zero Trust Principles

The fundamental principles of Zero Trust can be traced back to the origins of the internet. One example is the change introduced to RFC 1122 in 1989: "In general, it is best to assume that the network is filled with malevolent entities that will send in packets designed to have the worst possible effect." When John Kindervag introduced the term Zero Trust in 2010 the focus was to eliminate the idea of trusted and untrusted networks and see everything as untrusted. He introduced three foundational concepts: ensure that all resources are accessed securely regardless of location, adopt a least privilege strategy, strictly enforce access control and inspect and log all traffic. While these concepts describe practices likely to improve the majority of organizations' security architecture, we considered them too network-centric and hard to adopt in other areas. Microsoft has defined three similar principles to describe Zero Trust:

- **Verify explicitly:** Organizations that verify explicitly use all data and information available to reduce uncertainty and implicit trust.
- **Use least privileged access:** Using least privilege is always providing the least number of permissions necessary.
- **Assume breach:** Assuming breach is when you already consider your digital environment compromised.

1.2 The Extended Zero Trust Maturity Model

In this maturity model, we apply Microsoft's key principles of Zero Trust to organizational aspects of information security and combine those organizational aspects with the networking aspects of traditional Zero Trust maturity models. This allows organizations to perform a more comprehensive assessment of their overall Zero Trust maturity. We call this the Extended Zero Trust Maturity Model (Not to be confused with Forresters Zero Trust eXtended (ZTX) from 2018). The model is divided into three main domains: Technology, Processes and People. The domains are divided into focus areas with three maturity levels. Each focus area has accompanying control questions, allowing organizations to assess their own maturity.

1.3 Scoping and domains

The goal of the Extended Zero Trust Maturity Model is to take a more holistic approach to information security utilizing Zero Trust principles across the entire organization. When taking a holistic approach, the focus areas can be categorized within three different domains: Technology, Processes and People. The technology domain explains the technical solutions and components of a modern Zero Trust architecture. The process domain dives into many organizational processes and explains how to apply Zero Trust principles to these processes. The people domain sheds light on how the people of an organization can use Zero Trust principles to improve the organization's overall security posture. Below is a full list of focus areas within each domain.

Technology

- Network Segmentation and Infrastructure
- Dynamic Access
- Threat Protection

Processes

- Identity and Access Management
- Change Management
- Asset Management
- Incident Management
- Supply Chain Management
- Data Governance and Protection

People

- Employee Awareness and Training
- Information Security Culture

1.4 Identifying the focus areas

A literature study performed as part of the initial research laid the groundwork for identifying the focus areas. The research aimed to identify key components of Zero trust where the technical focus areas were discovered. Another noteworthy discovery was the lack of organizational aspects mentioned in both existing research and Zero trust maturity models. Thereafter multiple brainstorming sessions were held which resulted in a list of areas related to information security. Part of these sessions also included reviewing internationally recognized security frameworks for inspiration such as CIS(Center for Internet Security) Controls and NIST (National Institute of Standards and Technology) CSF (CyberSecurity Framework) among others. The Zero Trust principles were then tried on each area and its relevance was decided. The relevance of each focus area was then confirmed in several interviews with security professionals.

1.5 Defining the levels of maturity

The structure of our maturity model is heavily inspired by other maturity models such as the C2M2 (Cybersecurity Capability Maturity Model) model (Mehravari, 2015) and CISA's (Cybersecurity and Infrastructure Security Agency) Zero Trust Maturity Model, as well as feedback from our many respondents. As a result, we have chosen to define three levels of maturity: Traditional, advanced and optimal.

Traditional is meant to describe the traditional information security practices prior to implementing Zero Trust. The traditional level is often characterized by manual configurations and static security policies. Networks are often only segmented on the macro level with widespread implicit trust on internal networks. The policy enforcement is often proprietary and inflexible. Incident response and mitigation is done manually. The organizations have clearly defined policies, processes and procedures, but review intervals are limited, and no active awareness campaigns are utilized.

Advanced depicts an organization that has started the implementation of Zero Trust principles across the technology, processes and people domains. Centralized visibility and policy enforcement is implemented, incident response is partially automated through some pre-defined mitigations and the principle of least privilege is becoming more prominently adhered to. Some micro-segmentation of assets based on criticality is implemented, while egress and ingress is traffic is reduced to a minimum. Employees have a good grasp of the organization's security policies, processes and procedures, and these are regularly reviewed and updated.

Optimal describes the current ideal situation. The optimal state is ever-changing, and an organization will never truly be done implementing the Zero Trust principles in the most effective manner. In this maturity level, configuration and attribute assignment is fully automated. Access to resources is granted dynamically considering numerous factors including devices' security posture, threat intelligence, previous logging behavior, authentication, and authorization. Least privilege is dynamically enforced through open standards for interoperability across focus areas. Centralized visibility complete with extensive logging allows for point-in-time-recollection of state. The organization's information security culture is in focus and anchored in the organization with high awareness amongst employees. Employees are highly capable of scrutinizing information requests and mindful of where and when they discuss sensitive information.

CONTENTS

1 INTRODUCTION 2

1.1 Zero Trust Principles 2
1.2 The Extended Zero Trust Maturity Model 3
1.3 Scoping and domains 3
1.4 Identifying the focus areas 4
1.5 Defining the levels of maturity 4

2 TECHNOLOGY 9

2.1 Network Segmentation and Infrastructure 9
2.2 Dynamic Access 13
2.3 Threat Protection 17

3 PROCESSES 19

3.1 Identity and Access Management 19
3.2 Change Management 23
3.3 Asset Management 26
3.4 Incident Management 30
3.5 Supply Chain Management 34
3.6 Data Governance and Protection 39

4 PEOPLE 42

4.1 Employee Awareness and Training 42
4.2 Information Security Culture 45

5 REFERENCES 49

FIGURES

Figure 1: Micro-segmentation 10
Figure 2: Legacy System encapsulation 12
Figure 3: Using a PEP for determining access 15
Figure 4: Example of a Supply Chain Attack 35
Figure 5: Information Security Culture Illustration 46

TABLES

Table 1: Network Segmentation and Infrastructure Maturity Levels 9

Table 2: Network Segmentation and Infrastructure Control Questions 13

Table 3: Network Access Maturity Levels 13

Table 4: Zero Trust Network Access Control Questions 16

Table 5: Threat Protection Maturity Levels 17

Table 6: Threat Protection Control Questions 18

Table 7: Identity and Access Management Maturity Levels 19

Table 8: Identity and Access Management Control Questions 22

Table 9: Change Management Maturity Levels 23

Table 10: Change Management Control Questions 25

Table 11: Asset Management Maturity Levels 26

Table 12: Asset Management Control Questions 29

Table 13: Incident Detection Maturity Levels 30

Table 14: Incident Detection 33

Table 15: Software Supply Chain Management Maturity Levels 34

Table 16: Software Supply Chain Management Control Questions 38

Table 17: Data Protection and Governance Maturity Levels 39

Table 18: Data Protection and Governance Control Questions 41

Table 19: Employee Awareness and Training Maturity Levels 42

Table 20: Employee Awareness and Training Control Questions 45

Table 21: Information Security Culture Maturity Levels 45

Table 22: Information Security Culture Control Questions 48

2 TECHNOLOGY

In the technology domain we will explore technical solutions and components in a modern Zero Trust architecture. The technology domain is at the core of any Zero Trust architecture.

2.1 Network Segmentation and Infrastructure

Focus area	Traditional	Advanced	Optimal
Network Segmentation and infrastructure	• Perimeter-based security • Legacy applications have no added security	• Limited egress and ingress traffic • The most critical internal services are micro-segmented	• Micro-segmentation based on application workflows • Full usage of micro-perimeters for ingress and egress traffic • Encapsulated legacy systems

Table 36: Network Segmentation and Infrastructure Maturity Levels

2.1.1 Background

Traditionally, network segmentation and infrastructure were configured around a perimeter-based model: Internal traffic is generally trustworthy and only egress or ingress traffic is decrypted if necessary and inspected as it passes through the perimeter. Further segmentation of the internal network happened on the macro-level, often through different VLANS (Virtual Local Area Networks) and VRFs (Virtual Routing and Forwarding) with varying attributes.

The perimeter-based model has proven ineffective at preventing threat actors already on the inside of perimeter from moving laterally (East-West movement). This is particularly noticeable in today’s ransomware and supply chain attacks, where an attacker can spread their payload to large parts of the internal network with relative ease and with little need for employing sophisticated methods. This is usually achieved through lateral movement, followed by a top-down deployment after obtaining domain admin access. Furthermore, due to the increased utilization of remote-working services, the perimeter is becoming increasingly difficult to define and control. One solution to these problems is micro-segmentation, where all traffic regardless of its origin and destination will be heavily scrutinized by a Policy Enforcement Point described in the chapter on Dynamic Network and Application Access.

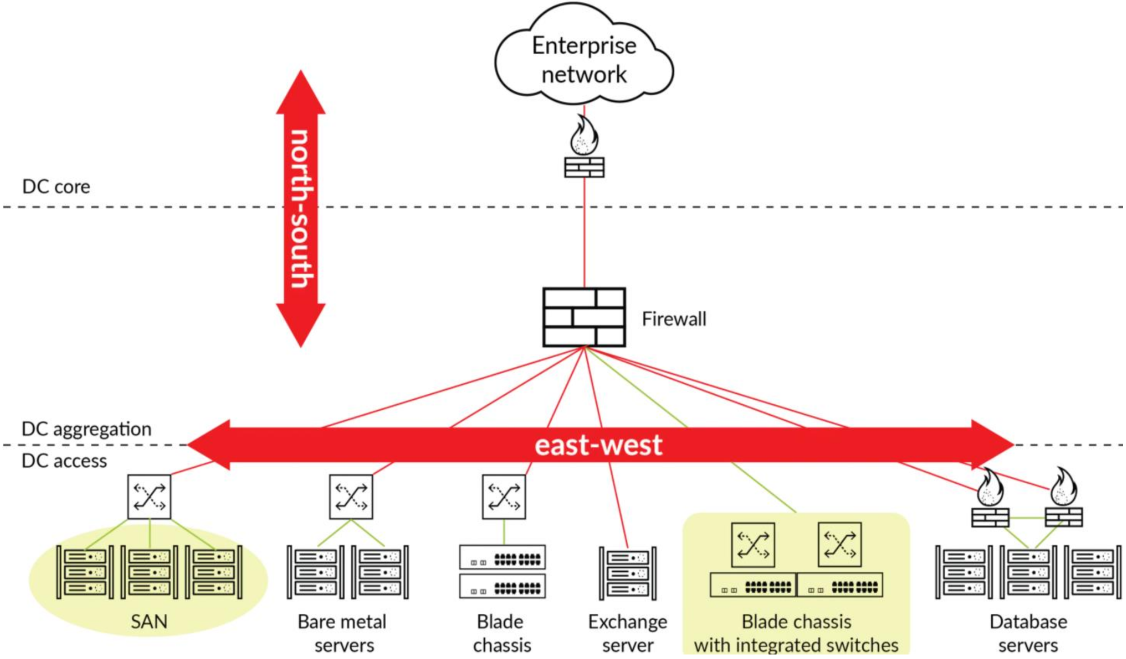

Figure 12: Micro-segmentation

Retrieved from https://www.paloaltonetworks.com/cyberpedia/what-is-microsegmentation

2.1.2 Traditional

Network segmentation is largely done at the perimeter, w
ith little to no micro-segmentation. Internal networks may be macro-segmented utilizing different VLANS, VRFs or additional hardware to suit the organization’s needs. An example of such macro-segmentation is the separation of development, test and production environments. Legacy applications have no added security and trust internal traffic implicitly.

2.1.3 Advanced

The organization's most critical assets are micro-segmented, with particular emphasis on internet-exposed services. The organization's less critical assets remain macro-segmented. The micro-segmentation of critical assets is handled by physical firewalls that will decrypt and analyze traffic if necessary.

2.1.4 Optimal

The entire network is micro-segmented utilizing virtual or physical firewalls for each computing resource or workload. The micro-segmentation is done based on application workflows, with full scrutiny of both internal and external traffic using Policy Enforcement Points. The application workflows are described through traffic analysis done by policy designers.

In cases where the network traffic is heavily encrypted, the organization has transitioned from traffic inspection on the firewall to using agent-based endpoint detection mechanisms. Since all encrypted traffic is terminated on the recipient, the bottleneck caused by firewalls having to decrypt all traffic that passes through them is removed.

Legacy systems are encapsulated, allowing modern access control and strong authentication to be used when accessing the devices. Encapsulation tailored for each legacy system depending on its dependencies on DNS, DHCP, NTP etc. Below is an example of legacy system encapsulation using forward and reverse proxies deployed on the device running the apps.

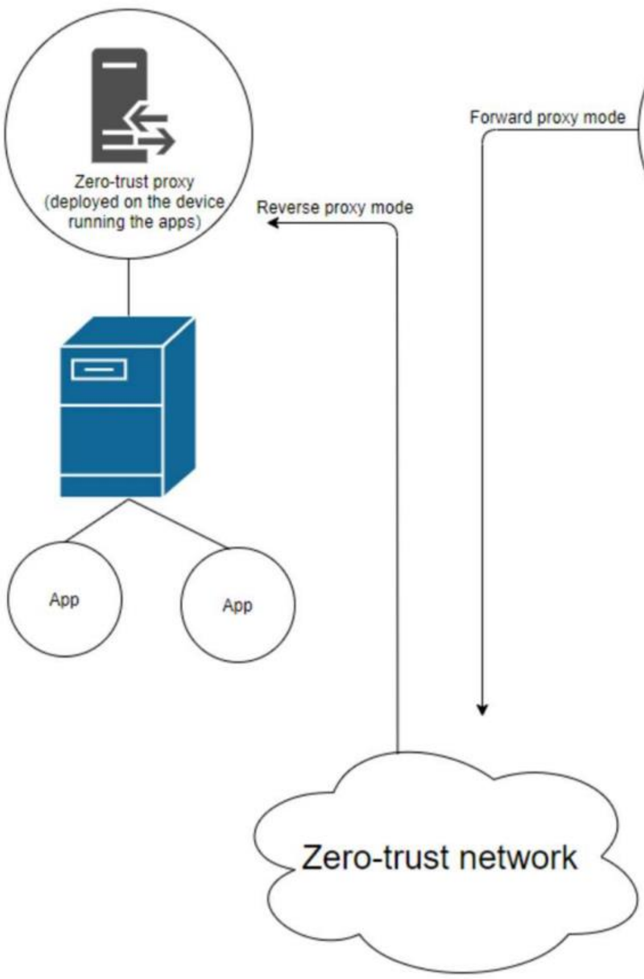

Figure 13: Legacy System encapsulation

Retrieved from Trust No One? A Framework for Assisting Healthcare Organizations in Transitioning to a Zero-Trust Network Architecture (Tyler and Viana 2021)

2.1.5 Control Questions

In the event of multiple networks with differing security measures, the control questions work best when applied to each network individually. Depending on each network's purpose and need for security, the organization can choose to either focus on improving the security of the lowest scoring network until they are all equally secure or consider different maturity levels as acceptable for different networks.

Level	Question
Traditional	1. Does your organization have a clearly defined perimeter and macro-segmentation of the network?
Advanced	1. Does your organization only allow the minimum required ingress and egress traffic? 2. Are internet exposed and critical services micro-segmented?
Optimal	4. Is micro-segmentation applied throughout the network based on application workflows? 5. Are micro-perimeters implemented for all ingress and egress traffic? 6. If the network is heavily encrypted, are agent-based endpoint detection mechanisms utilized? 7. Are legacy systems encapsulated, allowing modern access control and authentication?

Table 37: Network Segmentation and Infrastructure Control Questions

2.2 Dynamic Access

Focus area	Traditional	Advanced	Optimal
Dynamic access	• Access decisions are not centralized • Mainly authenticating identities with passwords • Most on-premises applications are accessible through VPN	• Centralized policy engine used to make access decisions • Multi-factor Authentication • Device compliance with specified security policy	• Access decisions are continuously reviewed and verified • Password-less authentication • Assume that users access all applications via untrusted networks, such as the Internet

Table 38: Network Access Maturity Levels

2.2.1 Background

NIST defines Zero Trust in the following way: “Zero Trust (ZT) provides a collection of concepts and ideas designed to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services in the face of a network viewed as compromised." It is clear from this definition that making good access decisions and removing uncertainty from this process is the main goal of Zero Trust. This was also the big selling point when Kindervag first introduced the term in his article "No More Chewy Centers: Introducing the Zero Trust Model of Information Security".

If we dive into the tools for removing uncertainty, the major components are authentication, authorization, and accounting. Authentication is the process of providing proof of the claimed identity and the most common form is when the claimer provides a username with an associated password. The main purpose of authorization is defining required permissions to access resources and then enforcing this for access requests. In Zero Trust, authorization is based upon the principle of least privilege, and access is granted on a need-to-know basis, which means permissions and access to data are limited to only what is required for someone or something to perform their operations. Accounting is storing information on each user's consumption of resources, and what actions they perform during access. It is the combination of authentication and authorization put into a system supported by other context-enriching tools that enable dynamic access decisions for every request. Access can be accepted or denied based on different factors such as identity and credentials, device compliance, geographic location, previous access information, and more. However, being able to fully leverage this technology requires organizations to have strong control of identity of users and applications both on-premises and in their cloud solutions.

2.2.2 Traditional

Traditional organizations often do not have any central system for authorization, but instead, hand this responsibility over to the individual applications. The applications perform a check against some static values and the decision for access is made only once. Identities are stored and managed in identity providers on-premises. Identities are authenticated using weak authentication methods, usually username and password. Some external endpoints like Virtual Private Networks (VPNs) might require additional factors when authenticating, however, very few applications are exposed on the internet. This means remote users must connect to the on-premises network through a VPN to access most applications.

2.2.3 Advanced

A centralized policy engine is the heart of the process of granting access. The system should take multiple signals into consideration when making an access decision. Most important is authentication, authorization, and device security compliance. However, access policy compliance is only enforced on the first access, and not verified continuously. Access to sensitive data is only granted to managed devices to maintain sufficient security on the connecting devices. Security health checks may consider patch level of device and existence of an endpoint detection and response tool. The figure below shows the concept of access requests going through a policy engine, becoming trusted by evaluation and verification.

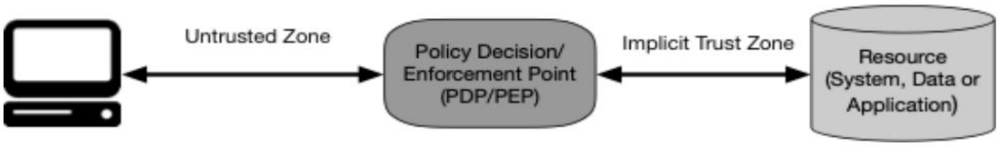
Untrusted Zone
Policy Decision/ Enforcement Point (PDP/PEP)
Implicit Trust Zone
Resource (System, Data or Application)

Figure 14: Using a PEP for determining access

Retrieved from NIST Special Publication 800-207: Zero Trust Architecture (NIST 2020)

Multi-factor authentication is enforced for all services, but Single Sign-On (SSO) can be enabled to improve workflow. Privileged roles should be limited, and just-in-time access should be utilized to follow the principle of least privilege. The organization federates user and application identities with cloud and on-premises solutions. Most of the applications on-premises are internet-facing but some are still only accessible through VPN.

2.2.4 Optimal

Advanced policy engines consider additional context-based signals in the decision-making process, including data on previous behavior collected for accounting purposes and threat information. Instead of only granting access at first request, continuous evaluation and verification are performed. This is also true for device compliance and security health, meaning devices are checked for compliance in real-time for every session. If there is a change of context during the connection access may be revoked. Users prove their identities with multi-factor password-less authentication methods. It is assumed that users access all applications and services from untrusted networks, such as the

Internet. Emphasis is therefore placed on keeping applications up to date and behind defensive mechanisms such as web application firewalls and application proxies.

2.2.5 *Control Questions*

Level	Question
Traditional	1. Is authorization required for application access? 2. Is at least one authentication factor used to authenticate users? 3. Is it possible for remote users to access internal services securely, for example via a VPN?
Advanced	1. Is a centralized policy enforcement engine used to make access decisions? 2. Is MFA used for identity authentication? 3. Is the health and antivirus status of the requesting device considered in the access decision?
Optimal	1. Are access decisions being logged and reviewed continuously in real-time? 2. Have usernames and passwords been replaced with password-less authentication methods? 3. Are directly available to users regardless of logical and physical location in a secure manner?

Table 39: Zero Trust Network Access Control Questions

2.3 Threat Protection

Focus area	Traditional	Advanced	Optimal
Threat Protection	• Static traffic filtering • Known threats	• Basic analytics identify new and unknown threats • Use of End Point Detection and Response agents for critical assets	• Machine learning used for threat identification • Dynamic traffic filtering based on context • Use of End Point Detection and Response agents for all assets

Table 40: Threat Protection Maturity Levels

2.3.1 Background

Threat protection has traditionally been about protecting the network from threats, often using databases containing known threat signatures or static traffic filtering configured on the firewalls by network technicians. As the threat landscape becomes ever more changing and complex, reactive threat protection is no longer adequate. This has given rise to threat analytics and the use of machine learning to identify new threats on the fly and dynamic traffic filtering based on context.

2.3.2 Traditional

Threat protection is largely based on static traffic filtering and known threats. The approach is purely reactive, and any newly discovered threats will need to be added to the database of known threats. The traffic filtering is done manually through firewall configuration.

2.3.3 Advanced

Basic analytics is deployed to proactively discover new threats. The combination of proactive threat analytics and reactive static traffic filtering and known threat databases provides significantly better coverage than a purely reactive approach. The dark web is regularly checked for any signs of data leaks that may relate to the organization.

Threat protection is further enhanced by deploying End Point Detection and Response (EDR) agents for critical assets.

2.3.4 *Optimal*

Telemetry is heavily utilized in the threat detection process. The emphasis on telemetry allows organizations to detect potential threats as they establish connections to potential Command and Control (C&C) servers or try to download malicious code. Furthermore, usage data and other contextual data such as the time of access, the location the request originates from or the account that is used to make the request is analyzed to make a decision on whether to let the traffic through or not.

All compute resources that support it now have EDR agents with the exception of the employees' personal mobile phones (if they are provided by the organization). Machine learning is leveraged to have significantly more accurate proactive threat detection. The machine learning algorithms continuously improve their detection rate by analyzing data, learning threat signatures and predicting other similar threats. Traffic filtering now considers context-based signals such as application workflows.

2.3.5 *Control Questions*

Level	Question
Traditional	1. Does your organization perform traffic filtering? 2. Does your organization utilize a database of known threats for threat protection?
Advanced	1. Does your organization deploy analytics to proactively discover new threats? 2. Does your critical assets have EDR agents installed?
Optimal	1. Does your organization utilize telemetry for threat detection? 2. Does your organization utilize machine learning to improve threat analytics? 3. Is the traffic filtering based on contextual data such as application workflows, telemetry and usage patterns? 4. Do all your assets except the employees' personal mobile phones have EDR agents installed?

Table 41: Threat Protection Control Questions

3 PROCESSES

The processes domain explores the organizational processes where Zero Trust principles can be applied to enhance information security.

3.1 Identity and Access Management

Focus area	Traditional	Advanced	Optimal
Identity and access management	• Central directory of identities • Manual process for granting access and permissions • Separate accounts for administrative tasks	• Permissions granted following principle of least privilege • Time-limited roles and permissions • Manual access reviews	• Automated access reviews performed periodically • Account activity is monitored and inactive accounts are deactivated • Penetration tests to harden the Identity and Access Management solution • Just-in-time activation of privileged roles

Table 42: Identity and Access Management Maturity Levels

3.1.1 Background

Confidently granting access dynamically to resources in the network requires good Identity and Access Management (IAM). This includes having full insight into all existing identities in the organization, be it employees, applications, or machines. Additionally, each of these identities has its own rights and permissions that must be managed. An identity and access management solution is often leveraged to manage this process. The IAM solution is preferably fed data from an HR (Human Resources) system, serving as a source of truth. The reason for using the HR system is its’ detailed information on all employees, their full name, department, role, start and stop date, and so on. Having information on roles for all users enables easier granting of permissions and privileges in the IAM solution.

After compromising legitimate user accounts, attackers will start examining what permissions they have and look for servers and systems accessible to them. The principle of least privilege is central in the prevention of this and can severely reduce the impact of such compromises. Following this principle, identities should only have access to what is required at minimum to perform their purposed operations. A database administrator should have access to work on databases, and employees in HR have the right to administrate the HR systems. These permissions should however not be mixed. A challenge often occurring is when individuals stay at an organization for many years and take on different roles. In these cases, it can happen that permissions are kept from previous roles, and they end up having very privileged accounts. To prevent this, organizations can establish frequent access and permission reviews where employees must justify the further need for their current permissions and roles.

A possible threat is former employees who are not offboarded properly and thus still have access to the organization's systems and resources. The employee may use this access to cause damage to assets or steal information. A proper offboarding process can be used to mitigate this threat. It is also recommended to track accounts for inactivity to detect employees who are no longer with the organization or service accounts no longer in use.

Assuming accounts will be compromised is also a good reason for having separate accounts for performing administrative tasks. Accounts that are used for business purposes such as e-mail and browsing the internet have a higher likelihood of compromise and should therefore not be given privileged roles. However, using separate accounts for admin tasks is no guarantee against compromise, and privileged roles should therefore not be given permanently. Instead, roles should be available for activation by using multi-factor authentication.

3.1.2 Traditional

The organization keeps a record of all users, applications, and machines in a central directory service. There is a manual process for granting access and permissions to identities. Instead of limiting permissions to only what is required, extensive permissions are granted to guarantee sufficiency. Users perform all tasks using the same account. Lacking off-boarding processes may enable previous employees the ability to access internal systems after leaving the organization. Employees changing roles keep their permissions, eventually ending up with highly privileged accounts. Separate accounts are used for administrative tasks.

3.1.3 Advanced

The IAM solution is used for federating identity and authorization for applications, making it central in the process of granting dynamic access to applications and systems. Permissions are granted to identities through a formal and automated process where approval from one or more parties is required. Project-related roles are time-limited, meaning they will expire when the project ends. Business justification must also be provided when requesting new permissions. There is a strong focus on only providing identities with the minimum required permissions for performing purposed operations. Permissions are automatically granted to employees based on their role, which is information collected from the HR systems. Automated processes are established for employees changing positions, joining, or leaving the organization.
Manual access reviews are performed regularly with the objective to remove permissions that are no longer necessary based on the employee's role and tasks.

3.1.4 Optimal

Automated access reviews are performed periodically where the employee or owner of a service account must justify why their roles and permissions are still needed. Accounts are continuously monitored for inactivity, and dormant accounts are deactivated.

Privileged roles are not standing, meaning they must be activated when needed, which may require additional authentication. This process if often referred to as just-in-time activation. If the role is required for a project or task, eligibility should be set for a specified period.

Both automated and manual penetration tests are performed regularly to investigate the available paths for attackers after successfully compromising an account. Automated tests are great for finding weaknesses that can easily be identified by scanning tools, while manual tests are performed by domain experts to identify weaknesses with higher complexity. The results are used to harden and improve the security of the IAM solution.

3.1.5 Control Questions

Level	Question
Traditional	1. Is a central directory service used to manage identities? 2. Is there an existing process for granting identities permissions and roles? 3. Are separate accounts used for administrative tasks?
Advanced	1. Are permissions granted following the principle of least privilege? 2. Does the process for granting permissions require business justification and approval from manager and/or system owner? 3. Are permissions time-limited? 4. Are manual access reviews performed regularly?
Optimal	1. Are access reviews automated and run regularly? 2. Are automated and manual penetration tests performed and used to harden the IAM solution? 3. Are just-in-time activation leveraged for permissions and roles?

Table 43: Identity and Access Management Control Questions

3.2 Change Management

Focus area	Traditional	Advanced	Optimal
Change Management	- Formally defined change management process for regular and urgent changes - The people allowed to request changes are clearly defined	- Criteria for compliance with Zero Trust principles and security requirements are defined for changes, non-compliant changes reviewed by security architects - Functional accounts are used to make changes	- Changes are handled using Configuration or Infrastructure as Code pipelines - The strict pipeline review process replaces the architect reviews in previous maturity levels.

Table 44: Change Management Maturity Levels

3.2.1 Background

Managing changes in the organization is about controlling the changes and making sure that they go through the correct approval processes. Changes can be initiated both internally (An employee suggests an improvement to the existing system architecture or a new component) and externally (A customer requests a change to infrastructure that your organization operates for them).

3.2.2 Traditional

A formal change management process is defined, along with a set of people who are allowed to initiate and request changes. Both the internal employees able to make changes and any customer representatives allowed to request changes are clearly defined. A special process is defined for handling urgent changes, bypassing or escalating some of the testing and quality checks of the change.

3.2.3 Advanced

Criteria for compliance with Zero Trust principles and the organization's security requirements are defined. Whenever a change fails one of these requirements, the change must be examined and verified by security architects to ensure that it is adequately secure before deployment. If a change request fails to adhere to the Zero Trust principles, or will otherwise result in reduced security, the request is denied or put on hold pending a workshop to improve the request and make it conform to the security requirements. Furthermore, the number of people both internally and externally who can issue a change request is as low as possible, following the principle of least privilege.
When making changes, the administrators use functional administrator accounts specifically designed for making that type of change, not their personal accounts. An administrator's personal account has no elevated privileges.

3.2.4 Optimal

Changes are now handled using either Configuration as Code (CaC) or Infrastructure as Code (IaC) depending on the change. This means that any changes to the infrastructure or configuration goes through a pipeline that is defined with code. This pipeline is coded such that the proposed changes will have to adhere to Zero Trust principles and the organization's security requirements to be implemented. Only the code defining the pipeline is reviewed by the architects, freeing up resources and making it possible to implement changes faster and more securely.

CaC and IaC automate the security review process of each change, allowing it to be used even for emergency changes. However, based on the organization's needs, a more lenient review pipeline for emergency changes can be coded to ensure an even faster implementation of the change. If this is the case, the security of the change is further improved to meet the stricter requirements in the main pipeline as soon as possible after implementation.

3.2.5 Control Questions

Level	Question
Traditional	1. Does your organization have a defined change management process for both regular and urgent changes? 2. Are the employees allowed to make changes and the customer representatives (if applicable) allowed to request changes clearly defined?
Advanced	1. Are security criteria for changes clearly defined and based on Zero Trust principles, as well as the organization's security needs? 2. Are internal and external change requests reviewed and verified by security architects when not conforming to the requirements in question 1? 3. Are the people allowed to request changes both inside and outside the organization limited according to the principle of least privilege? 4. Are functional administrator accounts used?
Optimal	1. Is IaC and CaC utilized for automated security reviews of changes? 2. If a less strict IaC or CaC pipeline for emergency changes is implemented, are these changes required to adhere to the stricter requirements of the main pipeline as soon as possible after implementation?

Table 45: Change Management Control Questions

3.3 Asset Management

Focus area	Traditional	Advanced	Optimal
Asset Management	- Manually updated asset inventory	- Automated asset inventory - Asset classification based on criticality - Separate patching regimes for critical assets - Limited access to Asset Management tools	- Asset Management tools are only accessible using privileged access workstations or secure administrator workstations - Changes to assets are verified automatically using a pipeline - Asset classification based on abstraction levels - Use of red teams

Table 46: Asset Management Maturity Levels

3.3.1 Background

IT asset management is the process of ensuring an organization's assets are accounted for, deployed, maintained, upgraded, and disposed of when the time comes (What is IT asset management (ITAM)?, 2022). Without having a complete overview of an organization's assets, it is impossible to define protect surfaces and perform the necessary network segmentation a Zero Trust architecture requires. However, many aspects of Zero Trust such as MFA can still be implemented without a complete server and software inventory. Implementing multiple improvements in parallel using a piecemeal approach is important for implementing a Zero Trust architecture in a reasonable timeframe.

Attack surface management can be seen as an extension of asset management where the organization approaches security from the attacker's perspective. Organizations generally employ red teams for this purpose. The red team utilizes various attack surface management tools to quickly find and close potential attack vectors that a real threat actor could exploit. Since the red team is employed by the organization and has easier access to internal information and a good overview of the infrastructure, they have a noticeable advantage over an external threat actor.

3.3.2 Traditional

In small to medium environments, asset lists or inventories are often maintained and created manually. One example of this is to use Microsoft Excel or similar spreadsheet tools to create a structured list containing information such as IP addresses, server names, FQDNs (Fully Qualified Domain Name) and operating system. This manual document is used as input for patch management. Strict and rigid documentation processes are implemented to avoid the asset inventory being out of sync with the real deployment. An asset list must be completely accurate for company to know exactly which assets need to be protected. Any inaccuracies in the asset list can lead to vulnerabilities not being patched properly and zombie servers.

3.3.3 Advanced

Asset management is done using an automated asset management tool in all environments. All newly created servers and decommissioned servers are automatically updated via the asset management tool. The risk or employees making mistakes or not following the strict documentation processes of the traditional level is removed. The inventory generated by the asset management tool is used as input for patch management, which are essential to keeping the organization’s systems updated, reducing exposure to known and unknown exploits. Gathering data through agents on each server or client and storing them in a Configuration Management Database (CMDB) is a common way of doing this. Any access to the asset management system requires MFA.

Critical assets are tagged in the asset management tool to allow employees to immediately identify them. The top ranked critical assets can be defined as the organization’s crown jewels. Any downtime on the crown jewels is highly detrimental to the organization and should be avoided at all costs. These critical assets are enrolled in a separate patching regime, with pilot testing being done prior to the rollout, removing the implicit trust in the software developers supplying the patches.

Following the principle of least privilege, only configuration managers and other employees who need to have access are granted access. Some employees may only need partial access to a specific system or group of systems, and others may only need read permissions for reporting purposes. The level of access each employee has to the asset management tool is reviewed regularly based on the organization's security requirements. Any permissions that are not strictly necessary are removed, and any employees changing roles within the organization or being offboarded will have their permissions reviewed.

3.3.4 Optimal

Any access to the asset management system is done through Privileged Access Workstations / Secure Admin Workstations (PAWs/SAWs). The number of people who can access an organization's asset management system is limited because of the high value such information will provide for a potential attacker.

The organization always assumes breach for all changes made through the asset management system. Therefore, the changes are all processed in a defined pipeline as described in chapter 3.2.4. Any changes to this pipeline must be reviewed and approved before they take effect. The pipeline ensures the integrity of any changes made.

In an effort to perform gap analysis and resolve any gaps, assets are further classified in abstraction levels such as data, component, system, value chain.

Furthermore, the organization employs red teams and to spot vulnerabilities and insecure configurations, unpatched systems/applications and other vulnerabilities from the attacker's perspective. These red teams work closely together with the asset management teams to remedy any potential attack vectors that are discovered.

3.3.5 Control Questions

Level	Question
Traditional	1. Does your organization have at least a manually updated asset inventory that is fully accurate? 2. Is the asset inventory used as input for patch management?
Advanced	1. Does your organization utilize asset management tools and to automatically keep the asset inventory updated and minimize the risk of human errors? 2. Has your organization implemented a CMDB to keep track of all assets? 3. Does your organization classify and rank assets based on criticality or other criteria? 4. Are separate patching regimes and piloting employed to secure the most critical assets? 5. Does your organization implement MFA for configuration managers? 6. Does your organization limit the number of people with access to the asset management system according to the principle of least privilege?
Optimal	1. Is access to the organization's asset management tool only granted when the request originates from a PAW/SAW? 2. Does your organization handle asset management using IaC and CaC? 3. Does your organization classify assets in abstraction levels to facilitate gap analyses? 4. Does your organization utilize red teams for finding insecure configurations, unpatched systems/applications and other vulnerabilities?

Table 47: Asset Management Control Questions

3.4 Incident Management

Sub-Focus area	Traditional	Advanced	Optimal
Security incident detection	- Establish logging and monitoring	- Assume breach detection capabilities.	- Red herring defenses.
Verification of detection	- Manual testing of detection capabilities	- Attack simulation	- Red and purple team exercises
Governance and information control	- Incident management team with clearly defined roles and access	- Restricted physical areas - Audit logging in security tools - Out-of-band communication during incidents	- Justification for accessing data seemingly not related to incident detection.

Table 48: Incident Detection Maturity Levels

3.4.1 Background

The ability to detect security incidents is an important part of any digital defense against cyberthreats. Failing to detect a security breach could make the impact of the incident a lot worse. When Kindervag (Kindervag, No More Chewy Centers: Introducing The Zero Trust Model Of Information Security, 2010) first introduced the term Zero Trust, one of the foundational concepts was to inspect and log all traffic. According to Verizon's 2021 Data Breach Investigations Report (Verizon, 2021) 20% of breaches analyzed were not detected before months had gone by. The zero-trust assumption of compromise requires us to think differently about how we develop our detection mechanisms. Many organizations spend large on security products and services and trust them to detect security incidents in their own environment. Unfortunately, there is no one-size-fits-all solution for security detection which is why trust in detection capabilities should not be implicit but gained with testing and confirmation. Also, security products often have access to large amounts of systems and sensitive data. This is a gold mine for attackers and malicious insiders, and access should be protected.

3.4.2 Traditional

Organizations have established security monitoring both for network traffic and endpoints of different kinds. Organizations collect and forward logs to a central location where security analysis can be performed. Security tools for endpoint detection and response are leveraged.

Detection capabilities are tested regularly. Manual testing is performed for some basic verification where adversarial behavior is simulated.

Security tools for detection and response have access to vast amounts of log data and end-client systems. Strict access control must be enforced, which requires a defined team of incident managers. Only members of this team will have access to the security portals but the principles of least privilege and need-to-know still apply. So even though security tools are very powerful and provide broad access, the security team only have access to the data and information that is required to do their job.

3.4.3 Advanced

Applying the assume breach mindset to the detection development is a game-changer. An organization that assumes that clients or network have been compromised is required to shift their focus away from trying to detect if someone is trying to get in. Instead, they are focusing on detecting behavior deeper into the attack chain. Typical behavior to look for is attackers performing internal reconnaissance, moving laterally, dumping credentials on endpoints, communicating with command-and-control servers, or already completing their objectives. This could be preparing and exfiltrating data or encrypting files as part of a ransomware attack.

Simulated attacks in a lab environment can be used to further increase confidence and reduce implicit trust in the detection capabilities. There are several open-source solutions that deploy virtualized environments combined with attack simulation test frameworks.

Communication between team members is key amid a security incident but secure communication is hard to achieve in an environment assumed to be compromised. Staying one step ahead of the adversary is hard if they can tap into the investigation. Severe incidents like a ransomware attack may also take down services used for communication, crippling teams' cooperation capabilities. Out-of-band services for communication and cooperation are therefore established and leveraged when deemed necessary by the incident response team.

Incident management teams work with sensitive information and often have a need for a visual representation of data. During incidents or handling of sensitive information, the incident detection and response team relocates to a physical area where access is controlled. Audit logs are stored to keep control of who accesses the physical area, who accesses which logs, and actions are performed in the security tools and services.

3.4.4 Optimal

The assumption is breached networks, and it is used as an advantage against the attackers. Red herring defenses can be used to distract attackers from their actual objectives, but also to make them step into a trap and set off the alarms. Honeypots are decoy systems in the network that appear legit, but their entire purpose is to lure attackers. Honeypots are deployed and whoever interacts with them is detected. The same principle can be used for sensitive files or files that appear sensitive before you open them. A benefit of using files is that they can also be used to expose malicious insiders.

Reaching the final stage of maturity requires regularly performing red team exercises to test detection capabilities. Red teams will simulate real threat actor activity and the ability to detect a red team can be used as an indicator of your capability's efficiency. It is important from a Zero Trust perspective that the red team tests the organization's ability to detect activity that implies compromise.

To gain better control of who accesses what data, analysts must provide justification on why access to data is required. This should only be applied to data that is seemingly not related to cybersecurity.

3.4.5 *Control questions*

Level	Question
Tradi-tional	1. Are systems established to monitor and detect cybersecurity incidents? 2. Have detection capabilities been tested and verified? 3. Is there a clearly defined team working with incident detection and response?
Advanced	1. Are detection capabilities developed with the assumption of a breached network? 2. Has attack simulation been used to verify detection capabilities? 3. Are out-of-band services established and used in security incidents? 4. Is the incident detection and response environment physically separated from the rest of the organization? 5. Is audit logging enabled for security tools?
Optimal	1. Are red herring defenses part of the detection capabilities? 2. Have red team exercises been performed to test detection capabilities? 3. Is justification required to access data not related to security?

Table 49: Incident Detection

3.5 Supply Chain Management

Focus area	Traditional	Advanced	Optimal
Software development and dependencies	- Software dependency inventory	- Separate software environments - Vulnerability scanning of software dependencies. - Remove unnecessary dependencies	- Review software dependencies - Separate application by services.
Vendor purchased software	- Identify software in use - Download software from vendor using HTTPS.	- Only use third-party certified software - Security requirements for software vendors. - "Software bill of materials"	- Isolate and monitor all software
Third-party services	- Inventory of third-party services	- Security requirements for service providers - Assess security of service provider	- Continuous assessment of service providers - Cloud access security broker

Table 50: Software Supply Chain Management Maturity Levels

3.5.1 Background

Buying and developing software and services or outsourcing a part of the IT environment to a third party are both common practices in modern IT strategy. While making any of these decisions you are making a choice to trust one or more parties. If you are developing your own software, it is likely that you are using libraries or dependencies developed by others. When buying software, you trust a third party by running their code on your systems. In both cases, you must trust that their intentions are pure and that their set of security standards matches yours. Making use of third parties means increased risk because your attack surface is growing. If the company offering you software or services gets compromised, it could potentially mean you getting compromised. An attacker might alter the source code of legitimate software to gain unauthorized access and detecting this is usually a lot harder than detecting regular malware. Many software development companies maintain a great level of security, but that does not mean they cannot be compromised.

The research paper "Software Supply Chain Attacks, a Threat to Global Cybersecurity" (Durán & Jeferson, 2021) suggests the reuse of code being the main problem in software supply chain attacks. The author provides the following explanation: "... from 85% to 97% of the code currently used in the software development industry comes from the reuse of open-source code frameworks, repositories of third-party software and APIs, creating potential vulnerabilities in the development cycle of a software product". Very often developers tend to import code written by others to perform simple tasks they could have written on their own or import large blocks of code while only using a small part of its functionality. Recent attacks and high amounts of code reuse show that this risk is real. The following section shows how applying the Zero Trust principles and mindset when working with cyber security supply chains can reduce this risk significantly.

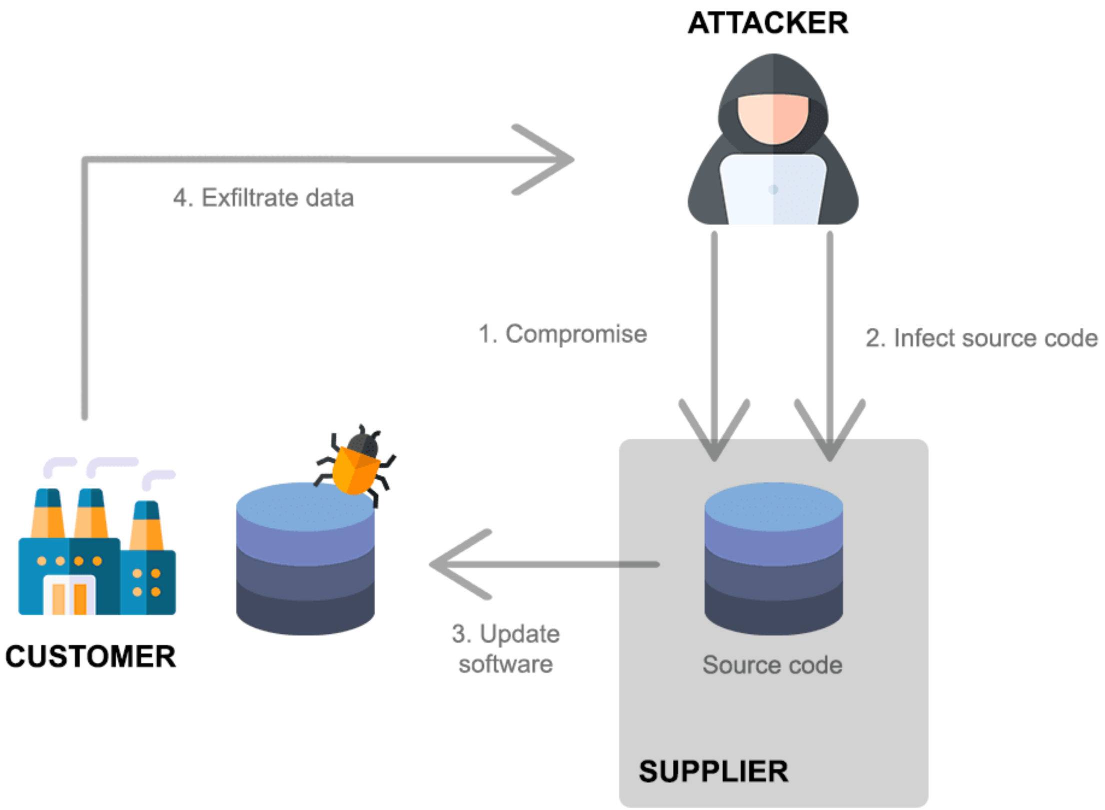

Figure 15: Example of a Supply Chain Attack

Recreated based on Threat Landscape for Supply Chain Attacks (ENISA, 2021)

3.5.2 Traditional

An inventory of software dependencies is established to identify existing vulnerabilities and reduce the risk of trust in the software supply chain. This is often done by storing source code on a central location preferably in a source code management service like GitHub or Bitbucket.

The risk of supply chain attacks is also present in software and services bought from external vendors. Organizations keep track of service providers and software in use, as a first step to mitigate this risk. It is also made sure that software is downloaded for official sources using HTTPS. Manual comparison of hashes is used to verify the integrity of downloaded software.

3.5.3 Advanced

Changes in dependencies can have a major impact on your systems. If a developer decides to delete a software package that you depend on, or one of your dependencies depend on, it could break the application. To mitigate this risk, organizations have established separate environments for development, testing and production. A set of different types of tests are ran in the testing environment revealing potential security and stability issues. It is great to combine these environments with a software composition analysis tool that can scan imported packages for known vulnerabilities and often includes different tests to reveal breaking changes to the application.

Using libraries and dependencies developed by others introduce trust to another party. This can in itself be problematic, but the real issue occurs when dependencies include dependencies of their own. This causes the supply chain to grow, as well as the number of trusted parties. To reduce this risk, organizations advocate the practice of removing unnecessary dependencies in code to their developers.

Buying software and third-party services from external vendors involve a certain amount of trust. With a Zero Trust mindset, we want to transform that trust from being given implicitly to something that is gained. Organizations may use two different practices to achieve this. The first practice is a control suggested in a security framework published by the Norwegian National Security Authority. It suggests that organizations “aim to only use software evaluated and certified by a third party. An example of such a certification regime is Common Criteria.” The second practice is defining a set of security requirements that external vendors must comply with when acquiring new software or services. CISA (CISA, 2021) suggests including some of the following requirements: Description of a software development lifecycle, vulnerability program, patch management capabilities and details on management of supplier lists. Reviewing providers' standardized assessment reports such as Service Organization Control 2 (SOC2) or using custom security questionnaires are also performed when acquiring new services. Purchased software should also include a “Software bill of materials” which is similar to a nutritional list and describes all the software components that make up the software.

3.5.4 Optimal

Organizations at the highest level of maturity may review the actual source code of their dependencies. As this is a very costly operation it might only be possible to do for security critical functionality like access control and encryption. Taking the source code and including it into their own code enable testing with their own tools and allow manual reviewing. An article titled “Secure Your Software Supply Chain – Threats and Mitigations” published by Truesec refers to this mitigation as “Vendoring” and claims it can reduce the risk of malicious publishers and supply chain attacks. The same article suggests separating applications into different services. Dependencies are often used to solve tasks in one part of an application. Separating these different parts into services reduces the risk of the dependency being compromised or deleted by the developer.

Assuming compromise is one of the foundational principles of Zero Trust and it can also be applied for software supply chains. Organizations that assume purchased software is or will be compromised implement controls for both prevention and detection. The software is allowed to run as intended, but least privilege principles are applied to prevent non-legitimate connections. This includes preventing outgoing connections except for destinations required to receive updates from the vendor and other expected traffic. A baseline of how information is flowing between the software and other systems is established to detect anomalies.

Similarly, a fully mature organization assumes its third-party service providers are compromised. Cloud access security brokers sit between the users and the cloud services to detect unauthorized exposure of information and non-compliant behavior. Organizations also continuously assess their service providers' security compliance, read release notes, and monitor the dark web for related leaks.

3.5.5 *Control questions*

Level	Question
Traditional	1. Is source code located in a central source code management service? 2. Has there been established a software inventory? 3. Is software downloaded from official sources using HTTPS?
Advanced	1. Are there different software environments for development, testing, and production? 2. Are software dependencies being scanned for known vulnerabilities? 3. Is there a goal of only using software certified by third parties? 4. Are security requirements set for software vendors? 5. Have unnecessary dependencies been removed?
Optimal	1. Is the separation of services in developed applications used when possible? 2. Is software being monitored and isolated like it has been compromised? 3. Is a Cloud access security broker in use?

Table 51: Software Supply Chain Management Control Questions

3.6 Data Governance and Protection

Focus area	Traditional	Advanced	Optimal
Encryption	- End-user client encryption	- Encrypt data at rest	- Encrypt all data at rest and in transit
Data inventory, classification and access	- Data inventory and file classification system - Data access control lists	- Sensitive data categorized and protected - Dynamic access control	- All data is inventoried and access to sensitive data is continuously monitored
Data recovery	- Automated backups	- Isolated instances of recovery data - Restore capability tests	- Immutable backups - Multi-user authentication for modification

Table 52: Data Protection and Governance Maturity Levels

3.6.1 Background

Data is a valuable asset for any organization and protecting its confidentiality, integrity, and availability is of high importance. Encryption is an efficient tool to protect the confidentiality of data, especially in a network we assume to be compromised. However, being able to protect data sufficiently requires us to know what data exists in the organization, where it resides, and its level of sensitivity. To be able to answer these questions an organization should establish a data inventory keeping track of their data. Knowing what data you need to protect, the next step is to control who accesses it, and categorize it based on sensitivity. Having these tools and processes in place enables us to grant access to data and applications following the least-privilege principle and on a need-to-know basis.

Availability of systems and data is a critical part of information security. Adversaries may perform unwanted changes to applications and systems, or even destroy data as part of a ransomware attack. There is also the risk of employees making mistakes or performing sabotage with intent. Data recovery measures are recommended to reduce said risk.

3.6.2 Traditional

With remote work on the rise and a shift from office desktops to laptops (Gartner, 2021) it is safe to say that devices are leaving the physical perimeter of the enterprise more than ever. Organizations trust their employees to protect the security of their devices, but controls are implemented to reduce the necessary amount of trust. Encryption of disks on end-user devices are leveraged to mitigate the risk of them being lost or stolen. Devices containing sensitive data are prioritized.

Organizations have established a basic data inventory and file classification system. Data inventories describe where to find what data. Rule and keyword-based methods are used to discover sensitive data. A file classification system is used to categorize data by labels. “Public”, “Internal” and “Confidential” are labels often used by organizations. Labeling of data is performed manually at this level of maturity. Having these labels makes it easier to treat data the same way regardless of location. Access control lists are used to ensure only authorized users can access data.

Organizations have established regular automated backups, especially for assets considered sensitive. In case of an incident the organization can roll back to the last known good state.

3.6.3 Advanced

Encryption is taken one step further by encrypting all data at rest, also counting data on removable devices. Data at rest is defined as data not currently being used, or in a state of transit. Organizations at the advanced stage of maturity account for, categorize and protect all information of value to the organization. Access to data is governed by a policy enforcement engine considering the context of the request. Modern information protection tools can protect your data with encryption and authorization policies. The applied protection measures will follow the data, so it is protected regardless of location. Many security incidents are caused by human error where an e-mail is sent to the wrong recipient or uploaded to a cloud service outside of the organization’s control. Protection measures as such are applied to data categorized as sensitive to prevent the occurrence of these types of incidents.

Assuming compromise implies greater risk of both incoming ransomware attacks and tampering of systems’ configuration and data. Organizations store backup data in isolated environments such as cold storage, separate cloud solutions or completely separated sites. Recovering the backed-up data is tested in regular intervals to make sure that both the backup and restore process is properly functioning. Soft delete functionality is used to protect against unintentional deletes and human error. Retention for the backup data is then set for a given period, in which deleted backups can be restored.

3.6.4 *Optimal*

All data is encrypted both at rest and in transit. That way data is protected from attackers performing man-in-the-middle attacks or gaining physical access to equipment. Files and databases are encrypted and masked, especially those containing sensitive data. Fully mature organizations classify all data regardless of sensitivity using machine learning and automated processes. Data that is classified as sensitive is continuously monitored to detect unauthorized access.

Immutable backups are used to protect critical data from ransomware, accidental deletion, data corruption and more. For less critical data backups are protected with multi-user authentication, requiring more than one administrator to make changes to configuration or on the backup data directly.

3.6.5 *Control questions*

Level	Question
Traditional	1. Is disk-encryption enabled for end-user clients? 2. Has there been established a basic inventory of data? 3. Is manual file-classification being performed? 4. Is access to data controlled with static or dynamic lists? 5. Is there a process in place for automated backups?
Advanced	1. Is data at rest encrypted? 2. Is sensitive data categorized and protected? 3. Is access to data granted dynamically? 4. Is backup data stored in isolated environments, and restore functionality tested regularly?
Optimal	1. Are data both at rest and in transit protected with encryption? 2. Is all data accounted for, categorized and protected regardless of sensitivity? 3. Is access to sensitive data being monitored? 4. Are Immutable backups used to protect critical data?

Table 53: Data Protection and Governance Control Questions

4 PEOPLE

The People domain explores how an organization's security posture can be further improved by providing training and increasing the awareness of its employees. Not all security risks can be adequately mitigated by technical measures alone, making this an important, but often overlooked domain from a Zero Trust perspective.

4.1 Employee Awareness and Training

Focus area	Traditional	Advanced	Optimal
Employee Awareness and Training	• Routines, policies and controls • Software aids	• General awareness training of employees to evaluate internal and external requests for information based on ZT principles • Targeted training that focuses on underperforming departments or individuals • Appointed security champions	• Risk and strategy-based training, where topics focus are tailored to the organization's largest threats. • Specialized awareness training of employees and departments to evaluate internal and external requests for information based on ZT principles

Table 54: Employee Awareness and Training Maturity Levels

4.1.1 Background

Employee awareness and training is essential when attempting to remove implicit trust from an organization: Most employees interact with their colleagues and exchange information many times a day. With the current pandemic and prominence of working from home, a lot of these exchanges happen digitally instead of face-to-face. Furthermore, employees often need to communicate with people outside the organization itself. This type of communication has led to an increase in successful phishing attacks and several data breaches. As a response, many companies employ phishing training tools and awareness campaigns to alert their employees to potential attacks from the outside. This training is a good complement to existing technical measures designed to prevent security incidents.

4.1.2 Traditional

The organization has good routines and policies for working securely and handling internal and external communication. New employees receive basic security training as a part of the onboarding process, and that all employees receive training whenever the policies or routines change. Controls are implemented to ensure that employees follow these routines and policies. Software aids such as alerting the employee whenever they are about to send an email to an external actor, or when they receive an email from such an actor assist in making the employee more alert when dealing with external actors.

4.1.3 Advanced

Awareness is continuously maintained through regular security training sessions and workshops for all employees. Here the employees learn the importance of assuming breach and never trusting, always verifying when receiving requests for information. This training applies to both internal and external requests for information. They learn to carefully consider who the requester is, and why they would need the information in their requested role. If the request seems suspicious, they are instructed to confirm the request by contacting the requester face-to-face or via video call to confirm that their identity has not been stolen and used to extract information. This applies regardless how the request was made (email, instant message, etc.). The employees are taught to ask themselves questions such as:

- Where does the requester work, and does it make sense for them to request this information from me?
- Why are they requesting this information from me?
- How can I be sure that the requester’s identity has not been compromised and used to exfiltrate information

Organizations utilize tools to visualize the progress made from awareness campaigns and for tracking the success rate of confirmed phishing attack against employees and departments. This data is used for extra training and awareness sessions targeted at specific employees or departments.

Furthermore, security champions are appointed for each department. These champions receive additional security training and are instructed to keep an eye out for suspicious employee behavior and malicious insiders. They will also provide additional training and advice based on the employees' awareness campaign results.

4.1.4 Optimal

The organization realizes that phishing requests targeted at the HR or economy departments are different when compared to requests targeted at the IT department. Furthermore, the departments mentioned need different kinds of training. Therefore, the regular workshops from the advanced level are now specialized for each department: These workshops provide clever phishing examples that are tailored to each department, along with learnings from previous phishing attempts and other security incidents of relevance for that department.

Risk and strategy-based training is employed, meaning that the training given to employees is tailored to the largest threats the organization faces. Security champions are responsible for action-based training, meaning that all employees whose actions led to a security incident will receive specific training to prevent the incident from reoccurring.

The organization also assumes that their employees have been phished. There are no repercussions for employees who admit to being phished, and any successful phishing attempts against employees result in the reason why the attack was successful being assessed and specific training being developed if needed.

4.1.5 Control Questions

Level	Question
Traditional	1. Are routines and policies for working securely implemented? 2. Does your organization have initial information security training for new employees? 3. Are software aids such as warnings when sending or receiving external emails implemented?
Advanced	1. Does your organization have regular, mandatory information security awareness sessions? 2. Does your organization utilize tools to maintain employee security awareness? 3. Is the statistical data from the tools used for targeted training of employees? 4. Does your organization have security champions in place for all departments?
Optimal	1. Does your organization have regular, specialized, mandatory information security awareness sessions for all departments? 2. Does your organization leverage risk and strategy-based training? 3. Is action-based training performed as a result of security incidents?

Table 55: Employee Awareness and Training Control Questions

4.2 Information Security Culture

Focus area	Traditional	Advanced	Optimal
Information Security Culture	• Defined secure processes and policies	• Regular reviews of secure processes • External audits	• Shared security-first mindset • Effective controls • Targeted responses to non-compliance

Table 56: Information Security Culture Maturity Levels

4.2.1 Background

Technical solutions struggle with preventing security incidents that occur because of a conversation or phone call between colleagues in a public place. This is an argument for why Zero Trust principles are so important, also in organizational aspects such as information security culture. Furthermore, this heightened awareness among employees is yet another layer an attacker must breach before gaining access to an organization's secrets. Information security culture can be viewed as a sub-culture of the corporate culture, and consists of four levels: Artifacts, Espoused Values, Shared Tacit Assumptions and Knowledge (Van Niekerk & Von Solms, 2010)

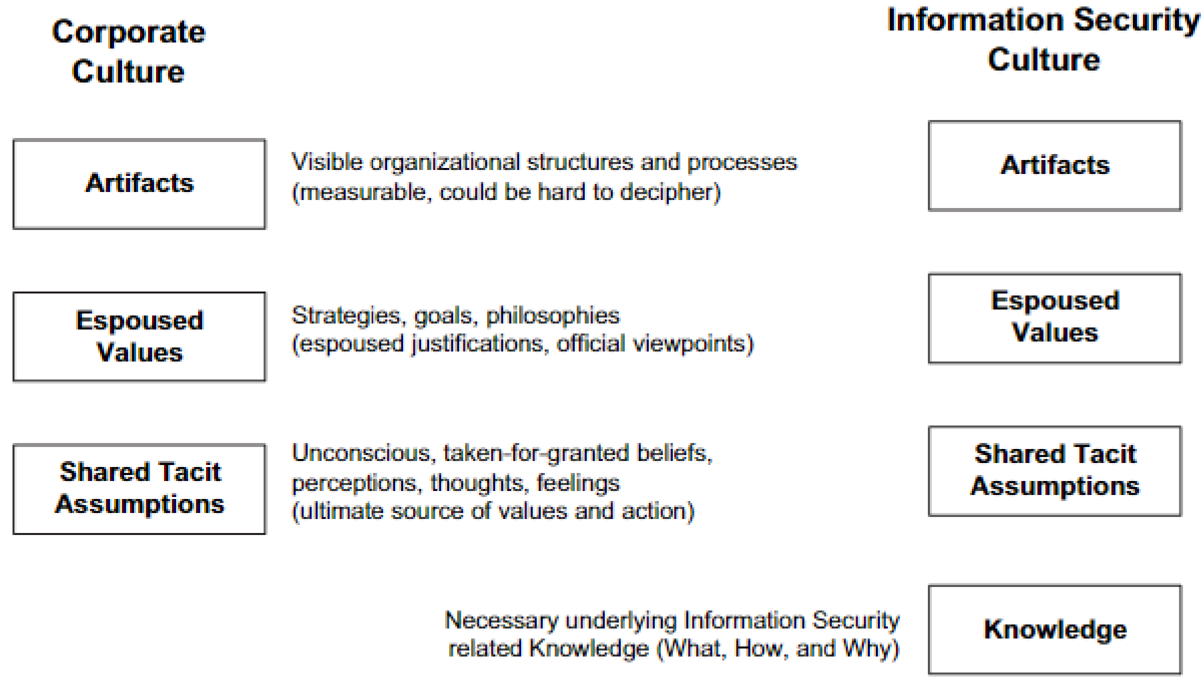
Corporate Culture
Information Security Culture
Artifacts
Visible organizational structures and processes (measurable, could be hard to decipher)
Artifacts
Espoused Values
Strategies, goals, philosophies (espoused justifications, official viewpoints)
Espoused Values
Shared Tacit Assumptions
Unconscious, taken-for-granted beliefs, perceptions, thoughts, feelings (ultimate source of values and action)
Shared Tacit Assumptions
Necessary underlying Information Security related Knowledge (What, How, and Why)
Knowledge

Figure 16: Information Security Culture Illustration

Retrieved from Information security culture: A management perspective (Van Niekerk & Von Solms, 2010)

4.2.2 Traditional

The organization has implemented policies, processes and procedures with at least some emphasis on information security. Input was gathered from internal security experts in this process. The policies detail the reasons **why** employees should adhere to the organization's processes and procedures. The processes detail **what** the employees must do and **who** is responsible for doing it. Lastly, the procedures explain exactly **how** the task shall be carried out.

4.2.3 Advanced

Information security best practices are ever-changing, and the organization addresses this by doing regular reviews of its policies, processes and procedures. This updated knowledge is immediately shared with the employees, raising awareness around the organization's policies, processes and procedures. Knowledge sharing around information security contributes to building Shared Tacit Assumptions among the employees, assuring compliance with the organization's policies and processes. Following the principle of never trust, always verify, the organization's established Shared Tacit Assumptions are complemented with controls that prevent non-compliance with the organization's policies where possible.

Part of the review process described in the paragraph above involves limiting employee access according to the principle of least privilege: Employees should only have access to the data and tools that they require to perform their tasks. The reasoning behind the privilege removal must be communicated to the employees in such a way that they understand the necessity of following the principle of least privilege.

Consulting external security specialists when reviewing the secure processes or performing regular third-party audits can contribute to an even higher level of security and uncover weaknesses or faults in existing processes or controls.

4.2.4 Optimal

The majority of employees share the same security mindset and are updated on the latest security practices. They are very likely to comply with the secure policies, processes and procedures defined by the organization. Incentives are used to further increase compliance and create a sense of ownership to the organization's security posture both on an individual level and the department level. Furthermore, the incentives encourage even the employees that may not be completely aligned with the organization's security mindset to comply.

The organization assumes breach and knows that the possibility for malicious insiders exists. Therefore, the organization has implemented security controls designed to detect non-compliance with the policies, routines and processes. These controls make it harder for malicious insiders to exfiltrate information or cause damage to the organization, while also allowing for the detection of mistakes made by employees that result in non-compliance. The organization responds to these mistakes with targeted training as described in the chapter 4.1.3 and 4.1.4.

4.2.5 *Control Questions*

Level	Question
Traditional	1. Has your organization developed security policies and secure processes for all tasks considered important to the organization?
Advanced	1. Does your organization regularly review and improve the secure policies and processes? 2. Are external security experts consulted during the review process? 3. Are security audits carried out regularly?
Optimal	1. Are employees kept updated on the latest information security trends? 2. Are incentives leveraged? 3. Are controls implemented to detect and prevent non-compliance with the secure processes and policies? 4. Is non-compliance responded to?

Table 57: Information Security Culture Control Questions

5 REFERENCES

CISA. (2021). *Defending Against Software Supply Chain Attacks* . CISA.

Durán, J., & Jeferson, M. (2021). *Software Supply Chain Attacks, a Threat to Global Cybersecurity: SolarWinds' Case Study.* IETA.

Gartner. (2021, April 1). *Gartner Forecasts Global Devices Installed Base to Reach 6.2 Billion Units in 2021*. Retrieved from Press Release Newsroom: https://www.gartner.com/en/newsroom/press-releases/2021-04-01-gartner-forecasts-global-devices-installed-base-to-reach-6-2-billion-units-in-2021

Kindervag, J. (2010). No More Chewy Centers: Introducing The Zero Trust Model Of Information Security. *For Security & Risk Professionals*.

Mehravari, P. D. (2015). Evaluating and Improving Cybersecurity Capabilities of the Energy Critical Infrastructure. Waltham, MA, USA: IEEE International Symposium on Technologies for Homeland Security.

SABSA. (2022, 04 20). *The Attributer's Blog – Zero Trusted*. Retrieved from SABSA Enterprise Security Architecture: https://sabsa.org/the-attributers-blog-zero-trusted/

Symantec. (2019). *Internet Security Threat Report.* Symantec.

Van Niekerk, J., & Von Solms, R. (2010). Information security culture: A management perspective. *Computers & Security*.

Verizon. (2021). *2021 Data Breach Investigations Report.* Verizon.

What is IT asset management (ITAM)? (2022). Retrieved from Atlassian: https://www.atlassian.com/itsm/it-asset-management

CHANGE LOG

Version	Date	Change
0.5	13/02/2022	Initial Draft
0.6	31/03/2022	Second Draft: First round of comments from respondents addressed, changes and additions from V 0.5 highlighted in yellow (Color-coded version only).
0.7	24/04/2022	Third Draft: Second round of comments from respondents addressed, changes and additions from V0.6 highlighted in green (Color-coded version only).
0.8	30/04/2022	Review of identities with subject matter experts, changes and additions from V0.7 highlighted in red (Color-coded version only).
1.0	29/05/2022	First release: Case study feedback implemented, changes and additions from V0.8 highlighted in blue (Color-coded version only).