Master's Thesis
Master of Computing Science


Radboud University Nederlandse Spoorwegen

# Securing R&D in OT

## Developing a Holistic Secure Operational Technology Development Life Cycle

*Author:*
D. Vonk, BSc
S4681533
david@vonk.frl

*NS supervisor (Treindigitalisering / OT Security Officer):*
K. Ooms-Geugies, BSc
klaasjan.ooms@ns.nl

*RU supervisor/assessor:*
Prof. dr. L. Batina
lejla.batina@ru.nl

*NS supervisor (ATO/Business):*
R. Van Pelt, MSc
rowan.vanpelt@ns.nl

*RU second supervisor/assessor:*
Prof. dr. E.R. Verheul
eric.verheul@cs.ru.nl

October 15, 2022
Version 4.2.1

**Published by** Wackelsteyn. Private publishing name for David Vonk.

**Editors** Klaasjan Ooms-Geugies; Rowan van Pelt; and Lejla Batina.

**ISBN: 978-90-9036624-1**



```
@book{Vonk2022,
 title    = {Securing R\&D in OT: Developing a Holistic Secure Operational Technology
              Development Life Cycle},
  author    = {Vonk, David},
  editor    = {Ooms-Geugies, Klaasjan and {van Pelt}, Rowan and Batina, Lejla},
  year      = {2022},
  month     = {10},
  publisher = {Wackelsteyn},
  isbn      = {9789090366241},
  url       = {https://wackelsteyn.org/securingrdinot},
  note      = {[Online; accessed November 1st 2021]
               Masters Thesis at Radboud University (Nijmegen, Netherlands) and
                 Nederlandse Spoorwegen (Dutch Railways).},
}
```

**I would love to hear from you** if you found this book useful. Don't hesitate to drop me an email at david@vonk.frl.

# Contents

# Chapter 1

# Introduction

With new times come new challenges: industry 4.0, the digitisation of industrial processes, introduces the need for new research and development processes. As systems become interdependant, cybersecurity matters become ever more complex. Novel perspectives are required, specified in frameworks and policy, to deal with the challenging nature of large developments in operational technology.

Currently OT security is often conducted in a one-dimensional or waterfall approach: security requirements are defined and implemented once (see ch. 4 and 8.2). This research adds a dimension to this: time. Risk, risk-acceptance and thus security requirements change constantly throughout R&D projects (in OT). These changes through time must be accounted for in models and controls to ensure adequate levels of security, without threatening progress. To do so, this research combines lessons from related fields to define a two-dimensional security application model called the Secure OT Development Life Cycle (SOTDLC).

## 1.1 Reading Guide

This document has been designed to be read in its entirety. The six pillars of the Secure Operational Technology Life Cycle (SOTDLC) (Risk, Risk Acceptance, Security Requirements, Purpose, Process, and People) build on each other. Hereafter you can an overview of interesting chapters for more specific audiences:

### 1.1.1 NS employees: management summary

Start with chapter 13 *Recommendations for NS*. If you are then seeking to understand the reasons for these recommendations, use chapter 12 for an overview of the outcomes of the interviews. Finally, chapter 14 gives a summarised overview of the entire SOTDLC: it describes how the framework can be implemented, and what steps should be taken throughout the project to make its security a success.

### 1.1.2 For implementing the SOTDLC

Start with chapter 14 *Implementing SOTDLC Into a Project*. When certain sub-processes or decisions are unclear, refer back to the relevant earlier chapter (when in doubt, consult chapter 9 *Process*).

## 1.2 Problem Definition

In recent years, there have been many reports of cyber-attacks on critical infrastructure such as the energy industry, water treatment and transportation [1]. The increasing risk for OT/ICS in industry 4.0 [2] progression, means that more resources are put into security to ensure reliable operation. Governments and infrastructure providers are hard at work to secure critical systems: security reporting [1] shows that incidents in OT context seem to be reducing slightly. Nonetheless, there is much work left to be done.

Securing OT is already a challenging task when the system architecture is clearly defined and when established security frameworks can be applied. The difficulty of securing a system increases dramatically when this system is still under development: when the architecture is changing often, this leads to an ever-changing need for security updates as well. Consequently, a different approach to risk-management is needed for such R&D contexts.

Security efforts in R&D currently do not align with process-architectures in projects, and are still often an afterthought. There is little scientific knowledge available about this subject, making it difficult for companies to improve. The articles and policies that are available are often adaptations of IT-security strategies that may be ill-suited for OT. Approaches with potential that do exist are little known, and are not often combined in complete and consistent frameworks.

Unfortunately, there is also a shortage in security experts and knowledge: a lot of the security personnel that is currently available is working on whatever is most effective in short-term, sometimes akin to sticking a plaster on a wooden leg. This means that security-policy is often lacking or unpractical: security efforts are consequently often inconsistent, ineffective or incomplete. There are many OT projects which could benefit from a clear strategy for establishing a security baseline in the system-under-development. Furthermore, these projects would benefit from an improved approach to education and awareness, where security is subtly integrated into core development and practices.

This research hopes to shape a bit of order in the chaos of securing an R&D system, by defining a framework for securing OT R&D; proposing policies based on the framework; and reinforcing the claims to effectiveness of the policies and frameworks through a case-study. Ultimately, security can be aligned with business, project and engineering needs through the Secure OT Development Life-Cycle as proposed in this thesis.

**The aim** of this project is to give a starting point and guidelines for securing research projects that are primarily based on OT systems, such as train-, track- and manufacturing-systems in the train industry. This means that people from various backgrounds have means to understand the basic elements that are involved in securing OT-R&D:

- *Managers:* should know in which parts of such a project it is most effective and cheapest to take security measures. This requires knowledge of business aspects and stakeholders, applied to security. They should also understand the difference between security theatre and measures that are actually effective.
- *Security specialists:* should be able to get an understanding of how OT-R&D within NS compares to other fields of security such as e.g. web-applications or IOT security and which approaches are most effective in this environment.
- *Risk Management:* should know which risks are and are not acceptable in OT-R&D and which measures have which effect.

- *Third parties:* such as external contractors should know what is important for the project and organisation, and what is the vision of the project and organisation with regards to security.

**The end product** of this project exists of three parts:

1. *This document* gives an in-depth overview of lessons-learned and the research process.
2. *Recommendations for NS* in a short form, describing what the next steps are for the NS to improve their security.
3. *Standalone policy drafts* describe individual subjects for specific audiences. These documents are based on this document, but are more digestible. These documents aim to concisely specify a specific part of the R&D project. It should be possible to read the guidelines as separate papers, with clear references and with a minimal amount of confidential information.

   The goal of these drafts is ultimately adoption into organisational policy, such that innovative projects are handled in a uniform way. As such, they should be supported by the organisational structure (e.g. cyber security management system, CSMS).

**The approach** of this research is to repeatedly use smaller building blocks to build a bigger structure. It therefore has the means to give a meaningful view on the organisational structure that constitutes the SOTDLC. Security models are used to substantiate a security procedure (sec. equation), which in turn is used throughout a project-process. This project is supported in its activities by the organisation, which is in essence people who work together in a social structure (i.e. in an organisation).

This is a quality-aspect oriented thesis: this means that our objective is to determine what is required from the organisation to support the social processes for such quality-aspects, such as security. Unlike organisational theory, this thesis looks at the organisation from a 'needs' perspective rather than a 'wants' perspective: it is focussed on the process more than the outcome, with the idea that a good process is required to facilitate a good outcome.

### 1.2.1 Vision

This thesis is striving for a pragmatic view on security where security is not over-done and everyone involved understands why the remaining security-processes and practices are necessary. Since security controls and threat analysis methods are comparatively well-defined, security engineers generally know what to do. On the other hand, non-security stakeholders still have a hard time understanding the details of what is needed for deploying security in their projects. By taking a life-cycle based approach in this research, this understanding is expanded by building on known management and risk management structures in OT. As such, the goals of other stakeholders can be aligned with security interests.

Consequently this thesis is not aiming to be a technically in-depth overview, rather a practical and comprehensible approach, which touches on all relevant aspects of security.

#### Structure

Where a security policy gives guidelines on how to deal with the risk (security measures), a theoretical framework describes how and why the measures affect security; how to determine the risk acceptance; and how to maintain and enforce the policy.

Security is all about risk: how much risk is an organisation willing to take, which risks are present, and how combinations of measures affect this risk. It is tempting to dive right into determining the levels of risk. After establishing risk, a nice overview can be made and presented to management: they can pick and choose which of the risks are and are not acceptable based on their understanding of the project, the organisation, and recommendations by the security specialist. This approach is, however, dangerous because it may lead to escalation.

#### Escalation

Underestimation of risk may lead to ignorance towards security. Overestimation of risk may lead to ignorance towards other important aspects, such as usability or innovative capability. Both of these behaviours threaten business continuity. This research thus aims to produce a framework that people are willing to follow, such that escalation is reduced.

Security is often ignored, because adherence to policies is often seen as a burden. This means that people are inclined to ignore procedures if they think another goal is more important. An example could be quickly setting up an alternative internet connection because the main one is too slow, with disregard for setting up security, in an effort to continue an experiment. The employee or team in question may think that they are doing the project a favour by continuing the valuable experiment. Risk is, however, well known to be underestimated: especially if there is no way to quickly assess the risk of such an action and if the person assessing the risk has no background in the area at risk. Therefore, this perception that other goals are more important is often wrong. Instead, this behaviour threatens the project.

At the same time, security must be proportional. In many cases, security overrules other interests such as usability. This is not only frustrating to engineers and maintenance, but also reduces the overall functionality and thus value-producing ability of a system.

### 1.2.2 Scope and Limitations

The goal of the literature review is to gain understanding about a possible framework for evaluating security in research and development for operational technology: the Secure Operational Technology Development Life-Cycle (SOTDLC). For such exploratory research it is not necessary to comprehensively review all available literature in related fields. This paper is thus limited to reviewing some authoritative sources and standardisation that give a broad, generalised, understanding of related fields.

Some rigid models are, on the other hand, imperative for understanding, discussing and modifying the current state of affairs. More effort is thus put into defining these.

This research aims to look only at fields that are directly related to the SOTDLC. Specifically we will look at OT security in general; the secure software development life-cycle (SSDLC); and security in the automotive industry (preferably in R&D context). An example of a subject that is out-of-scope is common vulnerabilities in web-applications: this is not relevant to a generalised framework for OT security.

Being more process-focussed, exploratory research, this thesis will not go into much technical detail with regard to controls or implementing technical solutions. A lot of research on this is already established and referred to throughout this thesis, so for the technically inclined-reader there will be enough to dive into.

This project will look at the security of the ATO-project and associated processes as a case-study, including security related reports of the past. Although this reinforces the ideas proposed in the theoretical framework, it does not replace in-depth effectiveness research. The case-study is primarily focused on the NS internally; third party suppliers and partners are not the subject of this report, except for procedurally. This research will look at comparing policy with practice in a broad sense, and is as such not a comprehensive analysis of the technical properties of ATO-technology or ATO-systems.

The NS is looking for a practical approach for innovative projects, which favours effectiveness over efficacy. From this perspective the scope of this project is different from the traditional academical approach of systematically checking and securing for everything.

This thesis contains a bit of both worlds: a more systematical approach to the risk modelling and risk acceptance strategy; but also an emphasis on practicality for the NS.

### Existing Architecture and API's

This thesis considers R&D projects as standalone entity which cannot influence different parts or systems in its context. For example the ATO-system, which is adding systems to a train to allow it to drive automatically, is assumed not to influence the existing train systems. Of course in reality the ATO-team could communicate with those responsible for existing systems and ask for a new feature or endpoint, but this is an impractical, slow and unreliable process. Especially in cases where hardware has to be changed, or where systems may impact safety, this is infeasible due to the often long procedures to guarantee the change does not negatively impact safety of existing systems.

Assuming the context is static requires that the R&D project in question is flexible in its (security) design. This is, however, not a problem because the whole point of such a project is that it finds out what does and does not work: it is therefore inherently flexible.

The consequence of this for security is that there has to be a clear designation of trust-boundaries and -assumptions between the system-under-development and its context. These boundaries should be considered in threat-, risk-, and impact-analyses and incident-response strategies. This means that the modelling and analysis of security in the research project is easier and more straightforward (i.e. the context need not be considered dynamic in the models), while not giving up any reliability for the analyses.

All together, the context of the system-under-development (SuD) is considered out-of-scope, except for the trust-boundaries between the SuD and its context.

### Security Analysis

Analysis of the security of any existing system or development project is explicitly not the primary goal of this project: rather it is focussed on policy and practices. Of course this may lead to incidental findings. Wherever this happens, the findings will naturally be reported to improve security. The in-depth analysis of specific findings or incidents is, however, out-of-scope.

### Existing Security Principles

OT in general, including rail infrastructure, is remarkable from a security standpoint, because well-known security principles from systems such as web-applications or office environments cannot be applied directly.

As with any consideration of security, one has to look at value in a broad sense: for infrastructure, safety is a very important value: this is often not in itself an important value to generic attackers, who are mostly interested in the monetary value [3]. For safety-critical systems, sabotage is the most common attack-type, often with a geo-political agenda [4, 3].

### 1.2.3 Research Questions

To answer the questions that are set forth in this section, the ATO-project is used as the subject for a case-study. This effectively means that the questions will first be answered for generic OT development. Then the ATO-project will be used to challenge the new framework and to fine-tune it.

As described in the problem definition 1.2, these questions are used to propose a set of recommendations and policy suggestions which are hoped to ultimately lead to a more advanced and whole security policy for innovative OT projects.

1. *What should a theoretical framework for reasoning about security in innovative OT projects look like?*
    (a) What does relevant literature prescribe for similar or related security subjects?
    (b) How could existing theoretical frameworks be adapted towards security in innovative projects?
    (c) How, and with what basis should security requirements in innovative projects be defined?
        i. How could an acceptable level of risk be determined in an innovative OT project?
        ii. What knowledge is needed to produce a vision on risk acceptance?
        iii. How can Risks and Risk acceptance be translated to suitable security requirements?
    (d) How are innovative projects organised?
        i. Which goals do stakeholders have, how are these goals prioritised, and what does this mean for security?
        ii. Which structure do innovative projects have?
        iii. Who are stakeholders in innovative projects, and what are their responsibilities?
        iv. What could stand in the way of security from the stakeholders' perspectives and what could be done to attenuate this?
2. *Are the findings for the theoretical framework supported in practice?* (Case study ATO-project)
    (a) What is currently done to implement security?
    (b) What do stakeholders involved in the security process think about current practices?
    (c) What are the wants and needs of people involved in security?
    (d) Do people agree with the findings from literature and the theoretical framework?
    (e) Would stakeholders be willing to accept the framework into new policy?

(f) What could be improved with regards to security implementation to become consistent with proposed policies?

(g) What is the impact of security implementation, in adherence to policy, on innovative projects?

(h) What lessons can be learned from the ATO-project for other projects (lessons-learned, best-practices)?

## 1.3 Justification

Operational technology (OT) and industrial control systems (ICS) are used to monitor and control critical industrial systems. OT systems are used in many facets of society, and their importance make the systems interesting for attackers.

Rail, for example, is increasingly seen as an alternative to cars in the EU, boasting advantages for the environment, safety, costs, capacity and reliability. Because of this increased interest, ProRail expects rail travellers to increase by 30% - 40% [5] in the Netherlands during the coming decade, despite Covid-19. The rail network is already used heavily so to meet increasing demands [5, 6], the Dutch rail sector is hard at work to increase capacity. Unfortunately, high urban density leaves limited room for physical expansion. Therefore, the sector is also looking at ways to increase capacity of existing infrastructure, by using modern technical means.

Within the train industry, OT systems are used everywhere:

- industrial workshops maintaining trains;
- track control systems responsible for signalling;
- controlling data-centers on wheels (i.e. trains);
- sensors for (real-time) monitoring of wear and tear;
- maintenance systems;
- and much more.

Many of these OT systems are not only critical for the business of train companies, but also for the safety of hundreds of thousands of travellers that use their services daily. Wherever safety is critical, security is naturally also a concern, for protection against malicious actors [7, 8].

Regrettably, there are also people who would rather not see such progress, or who want to enrich themselves from its value: cyber-terrorists and hackers threaten the safety of OT systems [3]. With the increasing application of IT in OT environments, industry 4.0 [2] is also expanding the reachable surface-area for adversaries.

The past decade, a lot of progress has been made in cybersecurity for OT. Nonetheless, it is still difficult to structure security efforts in non-standard projects or systems. Innovation - or Research and Development, R&D - is one of such areas where dealing with security is challenging: with the evolving nature of OT, its environment is also changing. Where traditionally OT is often separated from regular IT (and IP) systems, they are becoming ever-more interconnected.

OT in the R&D phase is even more complex: during R&D, development of new or improved systems is often done 'in the field'. Trains, for example, have to be tested on tracks

and with track-side systems that are in operation, often at night. The next day, these same trains may be used in normal operations again. This means that security concerns (risk) for the system-under-development are similar to those for operational systems, albeit on a shorter time-scale. Compared to operational OT, the requirements of security must thus be flexible, along with the ever-changing system. This must be considered when looking at the risks within the system: how should we discuss, model and quantify risks in the context of innovation?

### 1.3.1 Combined Approach with ATO-project Case Study

The research presented in this thesis focuses on applying security in Automatic Train Operation (ATO) technology in particular (more on the methodology in ch. 3). The ATO-project aims to, for example, investigate technical potential for increase in punctuality, reliability, energy-efficiency and more for train operations through (partial) automation of the train driving. There are several sub-projects focusing on different contexts and concepts in which ATO-technology could be applied. ATO and related concepts are explained in more detail in ch. 2, preliminaries.

The ATO-project is an interesting case in this regard, because of several reasons:

- ATO-technology at GoA level 4 (and arguably level 3) is safety critical. This means that security is crucially important for this project. Note that GoA2 systems are not strictly considered safety-critical by safety-specialists, because the systems can be physically overridden by the train-driver. For the purpose of this report, we do consider GoA2 systems as security-critical in a broad sense, because these systems can still be used by attackers to influence the train, or to propagate.
- ATO-technology is technically interesting: the ATO-project is very complex, with many facets. Most, if not all, common security challenges are represented in the ATO-system. This includes (but is not limited to):
  - internal and external communication of many sub-systems via all of direct wiring, data-busses, and IP-based systems;
  - remote control systems (SCADA) e.g. for traffic management on the rail network;
  - local control systems, e.g. for driving the train;
  - there are many internal and external stakeholders with different interests.
- Security in OT innovation in general is not fully mature: there is a lot of room for improvement. The ATO-project is actively looking to contribute to security policies and practices, in order to improve this. This is practical, because this means that people involved in the program are interested in working towards the best possible result, and are willing to invest time and resources into (usable outcomes of) this research.
- Many countries are currently working on ATO-projects and are very active in collaborating (see preliminaries, ch. 2): even findings related to just this particular case, can immediately be applied throughout the industry.

# Chapter 2

# Preliminaries

In this chapter, background knowledge needed to understand further parts of the research are introduced. As justified in ch. 1.3, Automatic Train Operation (ATO) is used as a case study for verifying and refining the theory developed in this thesis, so the relevant parts of rail and ATO will also be explained.

## 2.1 Definition of Terms

The definition of terms and models is limited to those that play a central role to the framework developed in this research.

### 2.1.1 System under Development (SuD, and SuC/CS)

Traditionally, for example, in context of TS50701 and IEC62443 [9] or [8], the term System under Consideration (SuC) may be used to describe the system that is subject to security scrutiny. In this report this term is, however, not suitable as is, because it is too broad.

ISO/IEC/IEEE 15288 [10], and NIST (800-160-1 and 800-37) [11, 12] define the System-Of-Interest (consisting of System Elements), Enabling Systems and Other Systems.

This research takes a similar approach, by more accurately describing the systems in the OT DLC (see ch. 8.2). Therefore, the SuC is split into: [7]

- *Contextual Systems (CS)* which are relevant for the project, but not the main focus of the development. These OT systems already exist. For example, the components that drive the trains with ATO-systems, are fit into existing trains: the interfaces of the components are made to fit the existing hardware.
- *System under Development (SuD)* is the main focus of the OT R&D project. This system is flexible and volatile. This system interfaces with the CS or is otherwise related to it. Parts of the SuD may be modified versions of existing systems.

Hardware-bound systems are difficult to change: designing new hardware, testing it, and finally deploying the physical product take a lot of engineering time. Security and R&D on the other hand, are fields where changes in insights and suddenly revealed vulnerabilities or faults require flexibility and rapid response. This means that hardware developments are slow and unreliable with respect to security and R&D. The consequence of this is that changing CS to fit a SuD is left as a last resort.

The CS may be entrance points for an attacker into the newly developed system. Likewise, the SuD is of concern to the existing systems, where it creates new entry points or paths to the CS. In an innovative OT project, CS are thus also being considered while looking at security, but their security status is by definition static. Any dynamic components of existing systems are considered automatically part of the SuD in this research: they are part of the research into the new system and its interface.

Summarising, the *System under Consideration (SuC)* is the combination of *Contextual Systems (CS)* and the *System under Development (SuD)*. The CS are static with respect to security, but they are no less relevant. The SuD is part of the development cycle, and is thus considered dynamic for security.

### 2.1.2 Operational Technology (OT)

We use the definition proposed by Williamson [13]: "Operational Technology (OT) refers to computing systems that are used to manage industrial operations as opposed to administrative operations. Operational systems include production line management, mining operations control, oil & gas monitoring etc.".

Particular to the rail industry, we consider systems on-board trains which manage or monitor physical systems (e.g. automatic climate control, traction controllers, braking computer, etc.) to be OT as well. The controls the train driver uses could be considered the SCADA (Supervisory Control And Data Acquisition system). Similarly, the switches and signals are also OT, with the train services controller as the operator of the SCADA.

In this project, particular attention is given to OT systems that have safety-aspects as well. This means that if the system fails in some particular way, the safety of people or the environment is impacted. Consequently, some failures in security may lead to significant safety concerns.

### 2.1.3 ATO

Automatic Train Operation can mean multiple things. Hereafter it is described which terms are used in this thesis for which specific things:

1. *ATO-system* - the technical systems used for automatic driving;
2. *ATO-technology* - the standardisation and generalised concepts of the related technology that are used internationally, and international or generalised implementations of ATO-systems;
3. *ATO-project* the R&D program (i.e. project) that is researching or developing ATO-systems or a specific ATO-concept for the NS (unless specified otherwise);
4. *ATO-team* the team conducting the ATO-project in the NS (unless specified otherwise); or
5. *ATO-driving* the actual automatic driving itself.

### 2.1.4 Security

*Some parts of this section have been adapted from earlier, unpublished, work by the same author:*[14]

*Security* can be loosely defined as the desire to protect *assets* from *violations*. Assets can be abstract: e.g. reputation or privacy, but also very concrete: e.g. a physical bike. Violations are events related to the asset which are considered negative to some *value*. The abstract properties that can be negatively influenced and pointers to assessing the risk thereof, have been defined in security models, which are introduced later in chapter 5.4.2.

In this research, privacy is not often discussed separately from security in general. Although there are many practical cases (especially in design) where discussing privacy as standalone topic, this thesis considers it as something of value that is sufficiently covered by our value-models (e.g. confidentiality): *security* thus includes privacy except for where it is mentioned explicitly.

A clear definition for threats, attacks and attackers is given by Bishop [15, Ch. 1.2]:

> "A *threat* is a potential violation of security. The violation need not actually occur for there to be a threat. The fact that the violation *might* occur means that those actions that could cause it to occur must be guarded against (or prepared for). Those actions are called *attacks*. Those who execute such actions, or cause them to be executed, are called *attackers*."

This definition can be extended by adding a definition for *vulnerabilities*:

*A technical construct that allows attackers to actually conduct an attack (with post-conditions) given certain assumptions (pre-conditions).*

A vulnerability is therewith by definition undesirable, as the consequence of an actual attack would be the violation occurring, with negative effects. A vulnerability has some pre-conditions which must be met before it can be exploited. For example: an input field which is prone to injection attacks can only be exploited if the attacker has access to the input field. The post-conditions of a vulnerability describe the assets (in the broadest sence of the word) that an attacker will gain after performing the attack. The combination of pre- and post-conditions of a vulnerability determine the value of the vulnerability, i.e. how much priority should be given to fixing a priority.

By describing the *Attackers*, *threats*, *vulnerabilities* and *attacks* a full image can be constructed of the *risks* within a system. Examples of common modi operandi of attackers combined with an attack-surface analysis produce a *Threat Model*. This threat model is used to prioritise the investigation into certain attack-vectors, and therewith the usage of specific vulnerabilities. These vulnerabilities are then listed in a vulnerability report. The attacks themselves are not of as much interest as the individual vulnerabilities, and mostly serve as example for possible consequences of the exploitation of certain vulnerabilities.

*Risk Acceptance* describes the degree of *risk* that an organisation is willing to take. Risk acceptance is determined by looking at the goals and values within an organisation. The capabilities of an organisation to deal with risk are *not* part of the risk acceptance analysis: if the organisation is unable to meet their security requirements and thus unable to reduce risk to acceptable levels, this should be explicitly noted. The organisation ought to readdress the risk as soon as more resources become available.

The *Security Equation* is defined as follows:

$$Risk - Risk\ Acceptance = Security\ Requirements$$

To know how much must be done to make a system sufficiently secure, first it must be known how much risk is acceptable and what this means for the project and organisation; then risk itself is assessed. Risk acceptance comes from a business- and technology-analysis. The amount of risk is determined by activity of attackers.

The security equation is based on the general risk management philosophy that there are 4 ways of dealing with risk [16, 8]:

1. Avoidance
2. Reduction/Mitigation
3. Transfer
4. Retention/Acceptance

Risk transference to third parties is not a common or easy thing in the cyber-security world, so we ignore this option. For the security equation we group avoidance and reduction in security requirements. Retention is represented by risk acceptance.

### 2.1.5 Quality

Quality is everything that does not directly contribute to producing value in and of itself. Quality can improve value production or prevent negative events from happening to it. Examples of quality aspects are legal, compliance, safety, security, reliability, environmental impact, etcetera. For those not intimately involved in security, but who are aware of operations in other quality aspects, it may help to think about security as similar to other quality aspects. Quality is used in the interviews (see ch. 3.3 and 11) over security, because it is more abstract and more palatable for people from other specialties, and believed to be analogous in procedure.

#### Management

For this thesis, we consider management as a quality aspect as well: management here is defined as the guidance of social processes within the (innovative) project.

The responsibility and ownership over decision-making is not considered to inherently be connected to management roles, especially in those cases where there may exist a conflict of interest. This will seem counter-intuitive to most readers: the idea here is to explore other options without bias towards tradition. It is important for businesses to make efficient use of their resources. The idea is therefore that decisions should be concentrated on those people who are best equipped to assess the outcomes of those decisions, and who can best determine what is a good direction. This research explores how existing approaches can be changed to improve in this regard: this viewpoint requires a different definition for management based on their subject-matter speciality, as is given here.

### 2.1.6 Projects and Programs

Within the NS, there is a difference between a program and a project. A project is roughly a team-effort towards a particular business objective of the company. A program is a combination of multiple similar projects into a single bigger project (i.e. the program), where knowledge and team-members are commonly shared.

In this research we use project and program synonymously, a program simply considered a large project.

## 2.2 Introduction to How Trains Operate

Many people will know what the signals alongside the train tracks look like, similar to traffic signals found on road intersections. In the Netherlands, these signals are part of the ATB system. ATB keeps track which parts of the tracks are occupied or safe to move through, and manages track systems such as switches.

Signalling systems, such as ATB, are needed because trains are very different from road traffic. Due to their length and weight, trains have a very high inertia. This means that where a car takes 50 meters to stop from a speed of 90 km/h, a passenger train might keep moving for over 600 metres. Worse still for freight trains, which could take to over a kilometre to come to a halt. If a train driver would thus have to respond to an obstacle, e.g. another train, they would have to be able to see ahead for at least their stopping distance. Since this is near to impossible at high speeds, trains operate with the help of signalling systems.

Current developments in Europe are the digitisation of these signals: e.g. track-side signalling (part of ATB) is currently being replaced by digital 'movement-authorities' which the train systems display to the driver in the cab. The new European Rail Traffic Management System (ERTMS) is the next step in train and track management. The aim is to have the same system everywhere in Europe, such that trains can more easily cross borders without having special locomotives or switching locomotives. ERTMS will also offer significant improvements over old systems, allowing for a greater overall capacity on existing tracks, and increasing safety and reliability.

## 2.3 Automatic Train Operation (ATO)

Many rail companies are currently experimenting with ATO-technology concepts. The research and development entailed in this, are combined into the various ATO-programs. The programs are hard at work with developing the knowledge and experience necessary to implement automatically-driving trains.

ATO-systems work in conjunction with the rail traffic management systems and on-board train control technology. This means that existing signals are used, verified and manipulated by ATO-systems to drive the train and to observe the surroundings of the train.

This section gives a general overview of ATO-technology and its general aims. This research is not part of the ATO-program itself, and thus cannot speak in its name. This chapter is a summary, solely based on the interpretation of the author; it cannot be used for technical or legal purposes, nor is it representative of NS policy because it is imprecise.

### 2.3.1 Goals and Justification

Ultimately the goal of using ATO-systems is to increase safety, quality and efficiency of rail operations. Automatically driving trains are able to precisely optimise energy usage and driving times. There are many more advantages to ATO-driving: which specifically and to what degree is often still subject of research. As discussed in the introduction, the Dutch rail network is expected to be near capacity in 2030 (without intervention): ATO-technology is expected to contribute to this capacity. Naturally, such innovations are interesting to companies, because ultimately it will increase customer satisfaction through reliability, shorter journey times, and environmental advantages.

### 2.3.2 Dutch Railways NS

The NS is also investing in ATO-technology, by means of an R&D programme. NS is looking at a few different flavours of ATO-systems, with current focus predominantly on GoA2.

The NS has found that the issue of security in ATO-systems and -projects is extremely important for guaranteeing safety. Increasingly, the NS is thus also researching improvements for security in R&D projects. Although security professionals have a good baseline of security procedures, they have found that default software procedures are unfit for one-to-one application in the physical environment of ATO-system OT.

### 2.3.3 International Developments

Many countries are currently working on some form of ATO. An overview can be seen in figure 2.1.

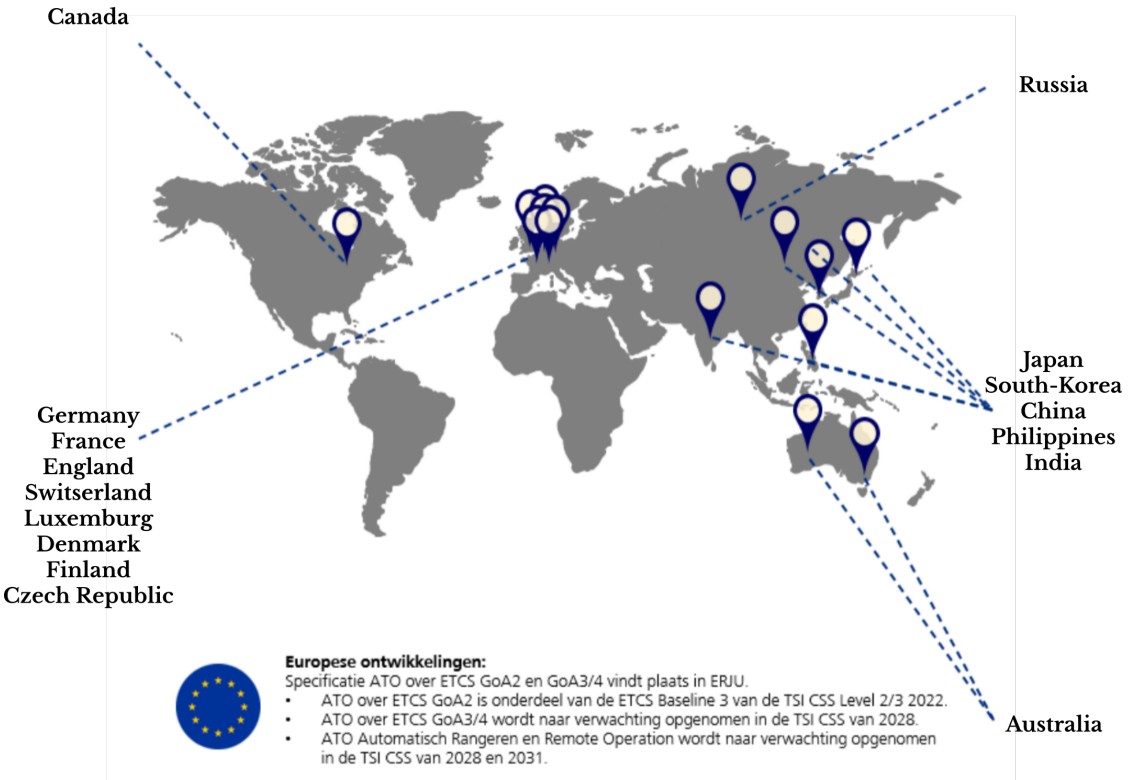


Figure 2.1: ATO throughout the world [17]

There is extensive collaboration between countries in sharing knowledge about ATO experiments. Team members of the various ATO-teams are also visiting progress meetings of other companies on a weekly basis.

An overview of articles about international developments is included in appendix B.

## 2.4 Automatic Vehicles: Grades of Automation

In the train industry, Grades of Automation (GoA) describe the degree to which a train can operate with or without human intervention. This is similar to the car industry, which is split

into 6 levels of automation.

- *GoA1* The train driver is fully driving the train, but is aided by the track- and signalling-systems. The driver may also be aided by driving-aid software that gives recommended amounts of breaking and accelerating (e.g. TimTim [18]).
- *GoA2* Accelerating, driving and braking are automated. The train driver will intervene to prevent accidents or if the system fails.
- *GoA3* The train is fully self driving (including door operations), and the train driver need not be in the driver's compartment. In case of calaminties, the train must be operated by a driver from the cab.
- *GoA4* Same as GoA3, but there does not need to be a train driver on the train. The train could be operated and monitored remotely in case of calamities if it is set up to do so, but this is optional.

GoA1 and GoA2 can also be combined with remote operation, where the train driver operates or monitors the train remotely. These are considered as two distinct categories of automation, because of the challenges that remote operation brings.

# Chapter 3

# Methodology

To achieve a good *Secure Operational Technology Development Life Cycle* (SOTDLC), a strong theoretical basis is needed. This chapter describes how this research was conducted, and introduces the overall structure.

The research questions (sect. 1.2.3) follow a *Design Science Research* (sect. 3.1) approach to research. DSR uses verified models to produce theoretical frameworks. These frameworks can then be applied for writing policy.

This chapter...

- first introduces DSR as our research approach;
- then discusses all the major model-components used for the SOTDLC; and
- finally shows how the SOTDLC framework was verified.

## 3.1 Design Science Research (DSR)

This research is based on the Design Science Research (DSR) method [33, 34]. The aim of this project, which was introduced briefly in the introduction, is the development of a set of guidelines for securing research and development in operational technology. In terms of DSR, the design artifact that is produced with this research is a *method* [34].

Before we can produce the method we need to consider the *constructs* and *models* [34] that are already available about the subject matter. These definitions are readily available in related fields, and can be adapted to suit our needs.

Besides having a goal that conforms to the design artifacts by March & Smith [34], this project is also grounded in reality as a real-world problem for the Dutch Railways. It thus fits the summary description of properties of typical DSR by Gregory [33].

While Gregory mentions that DSR is nowadays typically an iterative process, this project is structured in three distinct phases:

1. *Developing a Theoretical Framework*
2. *Verification of Theory*
3. *Results*

This approach was chosen, because each phase naturally follows from the previous, verifying and expanding it. The phases and the way they relate to each other are elaborated upon further in the following sections. This project is structured as a episodic or waterfall approach for readability. In practice, however, these phases are somewhat more intertwined or parallel.

Nonetheless, the generally episodic approach allows for creating a theoretical basis first with the perspective of the security engineer. This theoretical basis is then nuanced with insights from stakeholders with different perspectives (e.g. maintenance, safety, etc.). This ensures the rigidity and coherence of the developed methodology.

The first phase is about defining constructs and expanding existing models about similar subjects. These are then used to define a new framework that suits the needs of OT R&D. The method that is outlined in the first phase is then verified in the second phase by conducting a case study and involving people with experience. The goal of the second phase is thus to refine the developed method and to strengthen its claim to effectiveness. Finally, these findings are used in the third phase to distill key findings. Smaller chunks of the complete method are defined as standalone guidelines for certain parts of OT R&D projects. This furthers the readability and availability of useful information for specialists who are *instantiating* the method into a practical security implementation (for particular parts of the projects).

Each phase is represented by a part in this work. Each part will describe its own specific methodology for each of its sub-phases.

## 3.2 Developing a Theoretical Framework

The goal of this phase is ultimately to gain understanding of what security in OT R&D should look like from a theoretical perspective and how it should be applied in practice. The end product of this phase is a framework (in DSR terms methodology) that describes best practices and approaches to dealing with security in this particular context.

There are two distinct challenges in developing a SOTDLC:

1. Understanding relevant threats and designing appropriate countermeasures; and
2. Actually going through the processes developing the aforementioned understanding and subsequently actually implementing the countermeasures.

The first is an engineering challenge. It amounts to solving the *security equation* (as defined in ch. 2.1.4):

$$Risk - Risk\ Acceptance = Security\ Requirements$$

The second is a management challenge which comes down to performance. This means understanding the organisation: Purpose, Process, and People; and managing the people and process in such a way that the organisation fulfills its purpose.

For both challenges, the same research approach is used:

1. *Find a baseline* consisting of the current literature and standardisation. Concepts that are already well-established need not be reinvented. Literature consists of standards, scientific literature, and informal articles by news media and companies. Leads towards standards were found through reference by security specialists, but also through the

search functions of standardisation bodies. Whenever a relevant standard or article was found, its bibliography was also examined for further leads. Since the objective of this research is to give a pragmatic view on the SOTDLC, it does not use a formal method for finding literature: it is not necessary to comprehensively review all literature on topics to produce an accurate image of industry trends. What is clear, however, is that further interdisciplinary and comparative literature research would be useful.

2. *Identifying and defining required contextual models* for use in the DSR method. This is done by looking at the current understanding of concepts in literature, which gives an intuitive understanding of a subject. This intuition is formalised in more rigid theoretical models which can be used to make the subject matter and its properties, requirements and consequences more clear and unambiguous: this is needed to reason about practical properties (shortcomings) of the models. The identification is done by looking at the project structure and the identified challenges and steps. When discussing a challenge or step, models are needed to discuss them.

3. *Identifying and addressing shortcomings* in current strategies and models. The baseline is scrutinised and expanded using the models to identify current strengths and weaknesses with respect to the SOTDLC. By doing so, it can be determined what should be done to improve. Often, small changes or additions to theory are sufficient to make the model or strategy suitable for the intended purpose in the SOTDLC.

   To identify shortcomings, an open introduction and understanding of needs and peculiarities of OT R&D is applied to the models. If multiple models exist, they are compared, and fundamental properties are discussed. The combination of all could lead to novel insights or suggestions to take a slightly different approach. The aim of this research is not to develop new methods with scientific rigor, rather to adapt proven methods to the environment of the SOTDLC.

4. *Defining a solution for the challenge* based on reasoning using the models combined with understanding of the available literature.

### 3.2.1 Solving the Security Equation

$$Risk - Risk\ Acceptance = Security\ Requirements$$

To get to this end goal, we have to go through and thoroughly understand the individual variables and their relationships in the security equation. In order to do so, the variables must be dissected into their underlying components:

1. **Risk** = **Strategy** $*$ **Impact**
   (a) **Strategy** = $Threat\ Intelligence + Adversarial\ Model$
   (b) $Attack$ **Impact** = $Assets + Value$
2. **Risk Acceptance** = $\frac{\textbf{Impact}}{\textbf{Goals}}$
   (a) $Continuity$ **Impact** = $Assets + Value$
   (b) $Business$ **Goals** $\Leftarrow Vision$
3. **Security Requirements** $\Rightarrow Risk\ Prioritisation\ \&\ Classification \rightarrow$ **Controls**

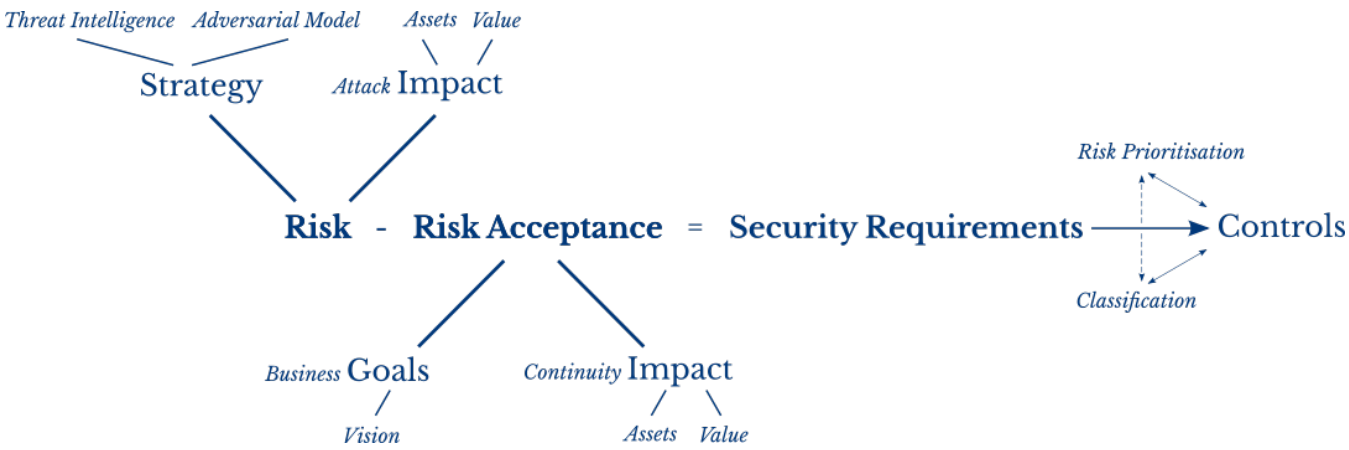


Figure 3.1: Risk Equation

The security requirements that are the result of solving the equation, lead the way towards security controls (see fig. 3.1). The act of implementing such controls in accordance with the requirements is outside of the scope of this research: this process is considered to be well-understood. The models and process definitions described in this research will nonetheless go a long way towards understanding and formalising the implementation of countermeasures based on the security requirements.

### Impact

The sub-equations of the risk-equation use *impact* twice. In this research, impact is also discussed twice, from different fundamental views. Impact in terms of risk is determined by the actions of attackers: the impact is the consequence of those. For risk acceptance, impact is discussed from a business-continuity perspective. Naturally, these two views are congruent to some degree (think of ransomware targeting business assets), but given the peculiarities of R&D (where assets may not exist yet, or might be volatile) it is still meaningful to consider both perspectives.

$$\textit{Attack}\ \textbf{Impact} \neq \textit{Continuity}\ \textbf{Impact}$$

$$\textit{Attacker}\ \textbf{Value} \neq \textit{Business}\ \textbf{Value}$$

What does hold is:

$$\textit{Attack}\ \textbf{Impact} \leftarrow \textit{Continuity}\ \textbf{Impact}$$

$$\textit{Attacker}\ \textbf{Value} \leftarrow \textit{Business}\ \textbf{Value}$$

Another note here, is that the *perceived* assets for the attacker is likely to be a subset of the actual assets. Generally, the attacker does not have complete knowledge of the system.

### Solving the Equation

As stated in the introduction of this chapter, this security equation has to be solved multiple times throughout the project. For R&D projects, which tend to favour quick cycles of experimentation and learning, it is imperative that the solving of the equation can be done quickly and efficiently without compromising on risk-acceptance.

This leads to the next challenge for the SOTDLC: actually integrating the security equation into the development process.

### 3.2.2 Integrating the SOTDLC

The dynamic nature of R&D projects changes the way security is engineered: it asks for a different preparation and some steps which are usually performed once, now become dynamic processes. It is thus imperative that those in charge of managing the project, and thus security, know about this and are able to deal with it. The perspective of process management is hence extremely important for this research.

Because dealing with security can be so different in R&D, it is important to extensively analyse the models used to manage it. Some previous work has been done on this already, both specific to security and in more general business research. Overall, the goal is to give management a structured approach to security, with clear and feasible steps throughout the project life-span.

Before the security equation can be integrated into a SOTDLC project, first the properties of that project must be understood:

1. *Purpose* Security goals in alignment with project objectives;
2. *Process* Project structure through time; and
3. *People* Stakeholders and their dynamics in the (social-)process.

By looking at the project from this perspective, it can be determined how security can give support and increase performance of the project as a whole. It is hoped that this approach will also address the common feeling that security is a burden, with the only purpose of slowing the project down. Security is then no longer a post-mortem, but a strength of the project.

The order of these three factors is important: before a process can be defined, it must be known what the purpose of that process is. Likewise, the purpose dictates who (specialists) are required to solve specific problems. Finally, the whole project has to be organised and supported by organisational people (support staff, e.g. management, HR, etc.). Who exactly are necessary is also dependant on the project structure, i.e. process.

From that the following structure follows: $Purpose \Rightarrow Process \Rightarrow People$.

## 3.3 Verification of Theory

To ensure the objective neutrality, completeness and effectiveness of the policies produced by this research, it is important to have a solid theoretical basis, both in public consensus (literature) and the organisational structure. The policy must be based on the experience and priorities of the organisation as a whole, and employees should be able to see the reasoning behind guidelines to better understand and support their importance.

This means that we must verify that our theoretical view of security matches reality. The verification is done by means of case study: by discussing existing security architecture, incidents and experience with various people who are stakeholders in the ATO-project.

Through in-depth interviews we can compare the complementary views that stakeholders may have with the security-framework, and challenge the objectivity of the framework. Depth interviews are also a precursor to acceptance research: people can comment on whether

they think the theory and its guidelines are useful, and whether they would consider following the guidelines without fear for consequence.

Taking a closer look at actual stakeholders also yields a better view of who should bear responsibility for what and refines the general project structure and stakeholder evaluations for our case study.

The goal of this part of the research is explicitly not to compare the quality of this theoretical approach to other approaches, or to give a strong proof to claims. Our practical approach is aimed at quickly using available insights to suggest a better framework than what has been available so far, and is known to be inadequate for this particular context. As such, it should be seen as exploratory research. Researchers and industry specialists are warmly invited to criticise and improve this framework to better substantiate (or disprove) it scientifically!

## 3.4 Results: Writing Policy

The verified theory is now applied in a policy document. The policy specifications should follow naturally from the theoretical framework:

$$Risk - Risk\ Acceptance = Security\ Requirements$$

The security policy formalises this equation into meaningful statements about security measures. The security policy also formalises practical considerations such as who is responsible (based on the stakeholder analyses and interviews with stakeholders) for enforcing the policy, for updating the policy, and for implementing the policy. As a general basis, NIST 800-12 [20, Ch. 5] will be used for the policy structure.

The resulting policy will consist of multiple documents, each describing an aspect of the security theory. This new level of granularity will make it easier for experts to find the requirements to their specific work, and for managers who are trying to oversee the entire process.

# Chapter 4

# Previous- and Related Work

## 4.1 Introduction

This chapter contains an overview of literature which deserves separate mention, or which is not used elsewhere in this thesis. This chapter sets the stage for the rest of the research: what has already been done before, and what can be learnt from that.

It is difficult to discuss relevant literature without having suitable models or even knowing which ones to discuss them with. Such discussions, as introduced in ch. 5, 6, 7, 8, 9 and 10, require these models to determine what direction to look towards in the first place. It is difficult to know which models are required to discuss a subject, before understanding the subject at a basic level first. Therefore, the later sections on contextual models and this chapter will reference each other quite a few times. Both chapters are a combined effort to dissect the problem and get to the bottom of it.

### 4.1.1 Summary

Rather than performing security as an after-thought, it should be addressed early on by the engineers who design and implement systems. To achieve this, the organisation has to train and support these engineers/developers with an organisational support-structure with security-coaches and -specialists. The desired level of quality (i.e. security) must be formalised and motivated: upper management should clearly convey this message and make sure that all stakeholders (including management) know, share, support and understand their responsibility. Simultaneously, this means that things should not be over-done: a system need not be completely locked down, simply because. Security should be functional and proportional.

Overall, the current trend in literature is that a one dimensional security approach is stretched and translated to fit the OT development life-cyle: the product development is iterative, whereas the security is still an incremental approach. For more complicated and larger projects this does not suffice: the security must be designed iteratively as well, to accommodate for changes in Contextual Systems (CS) and requirements of the System under Development (SuD). Furthermore, common evaluation techniques are ill-suited for OT, and do not give the desired results in practice. Better alternatives are available, though not specifically for OT (innovation).

## 4.2 The Security Development Life Cycle

The primary source for dealing with security in development projects is "The Security Development Lifecycle" by Howard and Lipner [35] (Microsoft). This work describes the approach taken by Microsoft in developing secure products. The book is mostly focussed on project management and secondarily on designers and architects.

In large part, standardisation or research improving on secure soft- and hardware development use this book as a basis (e.g. IEC 62443-4-1 [36], as discussed later). Therefore, it is useful to discuss this book in more detail here, since most of current approaches are based on it. Naturally, the book will be referenced more often throughout this research, citing specific sections wherever relevant.

### 4.2.1 Important notions

This book is overall a great resource for anyone involved in security, but some particular things are highlighted.

#### "It's Really About Quality"

Howard and Lipner [35, Ch 1, p10] describe the role of security in relation to privacy and reliability: they are partly overlapping. The takeaway is that security is often really about quality: if security fails, the product will fail and customers do not like that. Therefore, security must be on-par with the level of quality that is expected of a product. The same could be said about privacy.

For OT, one could add safety, similar to privacy. Privacy might seem less of an issue, since OT is managing physical systems and not data, but in fact some industries, such as traveller transport on rail, do have to think about it a lot (e.g. location/speed of a train, in relation to the travellers within the train).

Safety, however, is almost always very important to OT because of the potential physical consequences. Like privacy, safety is often a qualitative property of a system. If the system is able to do a task but is not safe, it is unlikely to perform its duty satisfactory. Clearly, safety also has overlap with both security (e.g. digital access control for remote SCADA operations) and reliability (If a train does not brake when it is instructed to do so, this impacts safety because the train may not stop in time for a hazard).

Howard and Lipner make the point "Only when you start to think about security holistically—as the intersection of privacy, reliability, and quality—does it start to make business-value sense.". This statement is completely true for OT as well when adding safety.

#### Security is not a one-man job

The core of the story is that security should be addressed as early as possible: this means during the development by the developers. This need not be perfect, but all security bugs that are prevented are much, much cheaper than fixing them later.

Therefore, security requires an aware, educated and motivated team! This also means that the team must have access to an expert who has more in-depth knowledge of security for difficult challenges. This security advisor (or security coach) should be someone skilled in management and security, who actively looks for, and helps solve security problems.

One important caveat here is that (upper-)management must come along for the ride. Management must thoroughly understand the value of security to the final product and the

organisation as a whole. They must explicitly motivate teams to take the necessary action: this means that projects must be structured differently (e.g. include security advisor) and that time and resources must be allotted to security.

### 4.2.2 Possible Improvements for SOTDLC

Although the SDL by Howard and Lipner is a great start for any development project, there are some important parts of the work that are not quite applicable to OT R&D projects:

#### Single-phase project structure

Similar to most other works discussed in this chapter, the SDL is not quite applicable to multi-phase projects (see ch. 8.2). Although agile variants are supported, these still consider the security goals or purpose (see ch. 5.4.2) as statically defined based on the application of the product.

The application of agile methods to the SDL [35, Ch. 18, p234], is nonetheless a great step in the right direction for making security an iterative process instead of incremental. This means that any intermediate version is mostly secure, albeit not quite completely tested.

#### IT vs OT

OT is different from IT, even at a fundamental level in attacker strategy (see ch. 5.3). A lot of the best practices and secure coding policies are not applicable to hardware. The idea of having secure engineering practices (as a replacement for secure coding practices) could be a great solution for OT.

## 4.3 A Threat-Driven Approach to Cyber Security by Muckin and Fitch

Muckin and Fitch [37] (Lockheed Martin) propose a very interesting approach to cyber-security: a threat-driven approach. This methodology is not specifically directed at IT or OT, but it is clear that Lockheed Martin is quite involved in hardware and OT: the methodology is quite applicable to cyber security in a broader context.

Since this approach is a slightly different perspective from traditional approaches (e.g. SDL or as defined in standardisation), it is useful to discuss this in some more detail here.

### 4.3.1 Important Notions

Reading the freely available paper is definitely recommended, but hereafter some important notions are compiled.

#### Cyber security should be threat-driven

Compliance, control, asset and vulnerability driven approaches are counterproductive according to Muckin and Fitch. These approaches implement controls for the sake of adhering to some kind of policy. In complicated systems, where multiple assets are involved in attacks, merely looking at independent assets and applying security controls is ineffective. Rather, security should look at architecture, and how the would-be-attacker would likely propagate through the system.

#### Quality

Muckin and Fitch agree with Howard and Lipner that security is really best viewed as a qualitative component of a system [37, p. 7]:

> Classic Systems Engineering practices do not effectively translate to cyber security practices. Development of secure systems – per the threat-driven approach – is very closely related to FMEA/FMECA (failure mode effects analysis/failure mode effects and criticality analysis) and other fault analysis practices used for quality and reliability engineering [10]. This supports the belief that highly secure systems are a corollary indicator of high-quality systems, a viewpoint the authors of this paper advocate.

#### Collaboration between operations/engineering and operations/analyst

Particular emphasis is put on the isolation of operations/engineering and operations/analyst [37, ch. 2]: Muckin and Fitch describe how this isolation manifests itself and how these two groups ought to collaborate.

#### Preparation is key

Before the design and development security can already be prepared. Threat intelligence can be gathered, attackers modelled and existing architecture is inventoried. The outcomes of this work are used in the requirements and design phases to refine security goals.

#### Threat modelling approach

The goal of the paper is to define an overarching threat-modelling approach with emphasis on looking at the system from the perspective of attacker operations (cyber kill chain, see also [38]). More on this in ch. 5.2.

## 4.4 Standardisation works

Quite a few standards are available for security. Especially for OT, standards are important works for industry and for scientific reference. Standards are often established by representatives of varying backgrounds, with extensive public review. It should be noted that standards are not always unbiased: the involvement of organisations which may have interests which are opposing security makes this so. Standards are nonetheless vital to the security industry, and must be considered when proposing novel frameworks or approaches.

In this section, a few important security standards are considered. Furthermore, some subject-specific sources are discussed in chapters and sections specifically related to that subject (e.g. threat moddeling).

This is by no means a comprehensive review, rather an introduction to broad perspectives.

### 4.4.1 IEC 62443 (Industrial communication networks - Network and system security)

The IEC62443 [9] is a series of standards for OT. It contains guidelines for policies and procedures on an organisational level, as well as controls that can be applied to OT. With this

standard, IEC was aiming to take "a risk-based approach to cyber security, which is based on the concept that it is neither efficient nor sustainable to try to protect all assets in equal measure. Instead, users must identify what is most valuable and requires the greatest protection and identify vulnerabilities."[9]

Especially for existing OT projects, 62443 should be a primary source of information. For R&D projects, IEC added a specific standard: IEC 62443-4-1: Secure product development lifecycle measurements [36]. At first glance, this standard seems to specify exactly what is needed: an approach for dealing with R&D in OT. At closer inspection, this is unfortunately only partly true. The development lifecycle (SDL in 62443 terminology, based on the work by Howard and Lipner [35]) is defined very similarly to a single-phase project (see ch. 8.2.1). Furthermore, the standard seems to be focussed on Type 2 projects [39] (see ch. 8.1) - although this is unclear. The requirements proposed by 62443-4-1 [36] remain a very valid starting point for other types of OT R&D projects, but as is made clear in chapters 8 and 9, a more solid understanding of project structure and associated goals is necessary to make a balanced security assessment. If the interests of other stakeholders in the various stages of an R&D process are not observed, a stringent and ill-fitting security regime may unnecessarily impede progress and threaten business-goals. Security could then seriously threaten the continuity of the project.

The core practices in 62443-4-1 are:

1. Security management
2. Specification of security requirements
3. Secure by design
4. Secure implementation
5. Security verification and validation testing
6. Management of Security-related issues
7. Secure update management
8. Security guidelines

Some of these practices will be discussed in later sections of this report where relevant.

### 4.4.2 NIST

The National Institute of Standards and Technology has published a plethora of standards that are relevant for OT security. NIST standards tend to be more pragmatic in nature, guiding the reader towards improving security, without laying a heavy burden of striving for perfection. Nonetheless, there is more than sufficient depth for organisations who have reached higher levels of maturity. An important advantage of NIST standards is that they are freely accessible.

Hereafter follows a selection of standards that were found interesting for this research. For readability, they are categorised as follows (sources that are particularly interesting are *emphasized*):

- General and System Development

  - *Framework for Improving Critical Infrastructure in Cybersecurity [19]*
  - 800-12 An Introduction to Information Security [20]
  - 800-53 Security and Privacy Controls for Information Systems and Organizations [24]
  - 800-82 Guide to Industrial Control Systems (ICS) Security [27]
  - *800-160 and Systems Security Engineering (SSE) Project [31, 11, 32]*

- Management
  - 800-18 Guide for Developing Security Plans for Federal Information Systems [21]
  - 800-37 Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy [12]
  - 800-60(-1) Guide for Mapping Types of Information and Information Systems to Security Categories [25]
  - 800-100 Information Security Handbook: A Guide for Managers [30]

- Toughness and Resilience (See ch. 7.3.2 and 7.3.3)
  - 800-34 Contingency Planning Guide for Federal Information Systems [23]
  - 800-61 Computer Security Incident Handling Guide [26]
  - 800-86 Guide to Integrating Forensic Techniques into Incident Response [28]
  - 800-92 Guide to Computer Security Log Management [29]

Hereafter, some of the NIST standards are emphasized:

### Framework for Improving Critical Infrastructure in Cybersecurity

This framework [19] is a great source for any organisation that is aiming to increase their level of cybersecurity. This (not-)standard takes a very pragmatic approach, aiming for maximum impact with a minimal amount of unnecessary bureaucracy and security-theatre. The improvement framework makes the organisational stance towards cybersecurity two-dimensional by adding a review and update process. The approach has five basic steps for organisations:

1. Describe their current cybersecurity posture;
2. Describe their target state for cybersecurity;
3. Identify and prioritize opportunities for improvement within the context of a continuous and repeatable process;
4. Assess progress toward target state;
5. Communicate among internal and external stakeholders about cybersecurity risk.

The core of the framework, the cybersecurity outcomes, are five important pillars of organisational cybersecurity: Identify, Protect, Detect, Respond, Recover. Furthermore, a basic maturatiy model is added in the form of implementation tiers.

Although this research is focussed on cybersecurity in projects, looking at this organisational process is still valuable due to its two-dimensional nature. The approach of describing target states, working towards these and meanwhile communicating the goals with stakeholders is a basic process-philosophy that is applicable to two-dimensional projects as well.

#### 800-160 and Systems Security Engineering (SSE) Project

The SSE project is focussed on establishing standardised foundations of systems security engineering [31] So far, the two most important publications of the project have been the current two volumes of NIST 800-160 [11, 32]. Currently, two further volumes are planned.

The 800-160 standard takes a holistic approach: in 800-160-1 [11], the systems security engineer is described as a generalist who is aware of all aspects of the system. Central to this idea is the Critical Systems Thinking [40] (CST) approach.

This perspective on security seems to be somewhat unique in the security world, where engineers often think in terms of risk and controls. In practice, however, the risk assessment and control implementation strategies are well-defined, whereas the actual implementation and execution of security knowledge in projects remains a problem. The result of this is that risk is assessed, and that controls are recommended, but that ultimately many systems remain insecure. Audits and compliance processes ensure some mandatory level of security, but such approaches can hardly ever be called efficient.

Some of these processes can be solved by involving engineering generalists who focus on security: their holistic view allows for sharp, project specific, analyses with a focus on effectiveness over efficacy. Even though this may seem like extra work, the efficiency gained leads to easier solutions which in the end are cheaper to implement. For this reason this approach is expected to be cost-effective, although this is yet to be proven in practice.

What is great about this standard, are its definitions of system life-cycle processes, general approach and clear correspondence to ISO ISO/IEC/IEEE 15288 [10]. De facto, it is an introduction of a holistic approach to security in systems engineering. This is very much applicable in OT as well. Nonetheless, this standard cannot be applied to the goal of this project directly, because it assumes single-phase projects (see 8.2.1 with one-dimensional security approach. This means that systems are potentially insecure during testing and experimentation.

Part two of these standards, 800-160-2, describes Resilience for e.g. new systems [32, ch. 1.1]. This is described in sections 7.3.2 and 7.3.3.

### 4.4.3 ISO 2700x

The 27000 series of ISO/IEC standards [41] describes many aspects and controls for security. This series is predominantly focussed on the traditional perspective on security, based on software and the management thereof within in an organisational context. The 27000 series is particularly strong in displaying what an organisation should do to ensure security from a broad management perspective. This describes what structural elements (policies, processes, etc.) the organisation should maintain to achieve a certain base-level of security. These structures are at a very different level than those at project-level: although very much

applicable to all projects and activities within a company, they are out-of-scope for this research. This research looks more at the actual product that is being developed. The 2700x series is a good start for any company, but the general ideas are also represented in the standards by NIST and other cited works. Because the 62443 [9] standard (also by IEC), is more applicable to project-management specifically in OT, and because NIST standardisation is more easily accessible, this research will not actively use the 2700x series as a reference.

## 4.5 Internet of Things (IoT)

Where OT is often focussed on systems at the larger scale of industrial processes, IoT often looks at smaller, consumer-grade devices in and around homes and offices. With the advent of industry 4.0 principles, OT increasingly starts to look like IoT. Both industries use sensors and communications systems to automate parts of existing processes: there are many similarities between these industries which make an interesting comparison of security principles.

Depending on the exact definition of IoT, OT systems could be considered a subset of IoT (e.g. see [42]).

### 4.5.1 Comparison of OT and IoT

Some of the similarities are as follows:

- Both OT and IoT use sensors to monitor processes or surroundings.
- Functionality is generally physically distributed over a facility, i.e. not in a single device.
- IoT systems are often remotely accessible: this is also increasingly the case for OT.
- Both IoT and OT make use of a combination of hardware and software (e.g. device firmware). Both thus have to deal with issues such as hardware or firmware attacks.

There are also some important differences:

- in OT safety and availability are primary concerns; in IoT this is often less so.
- OT systems are often expensive due to the high development costs to ensure safety and availability. IoT systems are generally cheap(er) and more easily replacible.
- OT systems often use proprietary protocols whereas IoT generally uses some implementation of the 802.11 standard (WiFi, Bluetooth, zigbee, thread, etc.)

Besides this comparison, there is another reason for interest in IoT. The modular nature of IoT allows for continual development and expansion of the possibilities. This is also reflected in some of the literature (e.g. the OWASP ISVS [43, Chapter V1] which is currently under development), where system development and development life cycle are important pillars. Of note here is also the approach for supply-chain security, and hardware-[43, Chapter V5] and communication-requirements[43, Chapter V4].

### 4.5.2 Applying IoT insights and literature to OT

Because of the popularity of IoT, more research and best-practices are available. Given the similarities between OT and IoT, it can only be useful to look at OT from an IoT perspective.

For example, when looking at the OWASP top 10 for IoT of 2018 [44], the only point that may not be directly applicable to OT is #6: "Insufficient Privacy Protection". But this is only the case because OT is not often concerned with processing private data. Therewith, the top 10 for IoT is a great reference for OT engineers.

Other, similar sources [45] are also available giving insight into common issues that have a high impact on security.

Where IoT started to become more popular some 15 years ago [42], OT has been gaining awareness for some 5 years. As Wind River [46] puts it 'Unfortunately, there is no "silver bullet" that can effectively mitigate every possible cyberthreat': the same is true for OT. Nonetheless, we can try to look (5-10 years) ahead to see what progress IoT has made so far, and reuse existing controls and strategies "that have evolved over the past 25 years"[46].

For example, Riahi et al. [47] model the relationship between the "Intelligent Object" (i.e. System under Development, see ch. 2.1.1) to "Technological Ecosystem" (i.e. Contextual Systems, again ch. 2.1.1), the "Person" and the "Process". Riahi also shows her model in relation to the current state of research in theoretical security principles. This approach is interesting, because it is a refreshing view on security complementing the standard CIA model by McCumber [48, 49], who also discusses the importance of viewing the security subject in relation to "information states" - related to "processing" for Riahi. See appendix A for more on McCumber's and Riahi's models with a review of other related literature and approaches.

### 4.5.3 Security by Design

A repeating principle in the above sources [42, 44, 43] is security by design in some form or another: starting with security as early as possible to integrate security in the design process of a new system. By doing so, the quality of the security, and thus the final product, increases, whereas costs and required effort go down. This emphasises the need for a similar approach for OT, where security is addressed early.

### 4.5.4 Hardware

For the more technically inclined reader, it may be worth it to look into physically unclonable functions or PUF's (e.g. [50]). Cryptographic systems can use PUF's, i.e. use physical properties, to more securely communicate with others. Because OT is also often based on hardware, it is worth exploring IoT hardware abilities for security sensitive systems.

For this thesis, however, this is too technical and thus out-of-scope.

## 4.6 Automotive

Only a few years ago, the automotive industry was bombarded with negative press [51, 52] about its state of security. Digitisation of cars is an interesting frontier of development, also from a security perspective and at the surface [53] car-manufacturers make it seem like there has been rapid and substantial improvement. In practice, the security of cars is still often based on old safety standards [54], where security is an aside. Last year, progress has been made on the standardisation front in a new automotive security standard [55].

Cars are products that the consumer takes home: willing car-owners can tamper with their vehicle all they want. This all seems very interesting, but a quick glance at some available literature (e.g. [56, 51, 52]) does not yield especially progressive or advanced generalisable insights on security, especially when comparing this to the rail industry. Projects such as EVITA [56] do conduct a lot of valuable in-depth research focussed on cars, but this is not easily extensible to other industries.

## 4.7 Evaluation

Based on reviewing literature that is currently available, we can determine what works well and where improvement is needed to reach a satisfying SOTDLC.

### 4.7.1 Holistic Approach

A common issue is that security is a burden to projects. When the project is well on its way, there is some security specialist who demands that all is changed because of some obscure security problem. This is time-consuming, annoying and expensive. The problem here is that security is an after-thought, and that the original design simply did not account for it.

One of the ideas proposed in literature to prevent this, is that security should be a holistic approach [35, 11]. This idea is also supported in [20, 21, 37] as secondary sources. In short, this approach mandates that the actual engineers/developers who design and implement the system are the most important to security, by preventing security problems from getting into the system in the first place. This means that activities such as threat modelling, security reviews and security testing are primarily done by those who engineer the systems, and within the development teams themselves. This way, security can be weighed against other interests at early stages and by those who understand the system best. An important theoretical basis for this approach (notably in [11]) is the Critical Systems Thinking [40] (CST) framework for dealing with complex problems with contrasting interests.

Needless to say, this also means that the developers must receive basic training in (OT-) security and how this is most effectively applied in their industry.

Simply hiring a security architect or security engineer to do everything in a project that involves security is insufficient, because this person will have a hard time understanding and scrutinising all the design decisions without knowledge of context. Furthermore, if a certain design choice turns out to be impossible because of security constraints, it is expensive to go back and change it. This approach is thus more dependant on generalists.

#### Security Coach

Of course, it cannot be expected from system engineers that they are experts at security as well as all the other factors that they already have to consider. Especially if there are conflicting interests between the other engineering subjects and security. Therefore, there must be oversight from someone who is actually an expert in security, that can keep oversight over the security process and final implementation. This person is primarily a manager, who coordinates between project management, security specialists, and the development team.

In this research we will call this person the 'Security Coach', but other terms have been used in literature for the same concept:

- Security Advisor (and also security coach for agile methods) [35]

- Systems Security Engineer [11]
- Information System Security Officer [21, ch. 1.7.5]
- Security Master [57]

The security coach oversees all activities on security and makes sure that an adequate level is reached. While the development team is primarily self-supporting, they can be connected to security specialists through the security coach. The security coach initiates important activities for security (e.g. design and threat-modeling reviews), keeps track of the current state of affairs (known issues), and is responsible for dealing with security incidents. These activities are also the reason why this thesis uses 'security coach' instead of advisor or engineer: this person has an active role in the development process, and as such is not just an advisor; this person is also not really an engineer, because they help the engineers with their work, but are not actively engaged in engineering themselves.

As Howard and Lipner emphasize: "do not lose track of the high-level goal of the security advisor: it is to help product teams become self-sufficient and "good at security."" [35, ch. 6, p69].

#### Other Security functions

The above approach requires the support of an organisational security team that specialises in security at a technical level. This could include pen-testers, security analysts, incident response team(s), and much more. The exact organisational support structure is out-of-scope for this research.

The core idea here, is that the system-developers must have access to quality information and advise whenever they have issues or questions. The security coach should know how to reach these people and make sure that their knowledge is used effectively.

### 4.7.2 Single-Phase Approach as Default

As seen in this chapter, a common assumption is that projects are single-phase. This is also addressed briefly in section 8.2. A basic single-phase methodology suggests an incremental approach to security, where more controls are implemented as the system matures. It is relevant to consider multi-stage projects separately, because they introduce possibilities as well as constraints. In large development programs and projects, the distinct design phases may pose challenges for security: for example, during the design and experimentation phase, it is not desirable to have very stringent security constraints on all prototypes. It is difficult to continually assess risk-acceptance and risk to develop security controls for said prototypes, especially if prototype throughput is high (e.g. multiple versions per day). Nonetheless, contextual systems (CS) or the prototypes themselves may have need for security insight and controls, albeit in a different way, i.e. a special multi-phase approach.

### 4.7.3 Motivations

Implicit to a lot of standards is an incremental (waterfall) security approach with the aim of getting a grip on security at an early stage, to predict and manage security risks early on. This is up to 150 times cheaper (in IT [58], OT is likely even more dramatic) than waiting until the product is done. This is a good reason for starting with security from the outset.

The crucial pitfall to consider is the intrinsic assumption that it is also an option to secure a project after the product was finished, even if at a much greater cost. For the systems which are subject to this research, this is not the case: trains must be reasonably secure before prototypes are tested, because the environment is practically the same as for the end-product. Similarly, self driving cars which are tested on the road must be sufficiently secure. The question is of course: what does this 'reasonably secure' mean during each phase of the SOTDLC?

Even when it is not strictly necessary to secure the system during development, it is still prudent to adapt the security implementation strategy (see ch. 7) in such a way that the system security is developed iteratively in a strategic fashion, prioritising effectiveness (NIST also describes this on an organisational level in [19]). By considering the system through a more accurate DLC model, and describing the desired security at each stage more specifically, security can be designed to fit even better to the entire SuC. This potentially saves even more time and effort, thus increasing quality while saving money.

### 4.7.4 Agile Based SDL

Although the dominant approach to security during development is a one-dimensional incremental approach, there have been some publications applying agile methods to the security development life cycle, e.g. [35, Ch. 18] (see also ch. 4.2.2) and [57].

Howard and Lipner [35, Ch. 18] mostly describe how their SDL can be applied to projects using agile. This means that some agile processes are modified to include security; they also suggest adding a security sprint "security push" and a "final security review" to remove most of the remaining security bugs before deployment.

#### Applying Agile Methods to OT

The overall challenge with agile as described by Howard and Lipner [35, Ch. 18] seems to be the wicked and interdependent nature of security: the consequence of this is that security has some distinct steps that do not necessarily clearly fit the agile philosophy, for example in documentation, preparation and final security review before publication. Howard and Lipner also comment on the suitability of agile to large projects: "What sets (...)[agile-] projects apart from most Microsoft projects is that (...)[these] are not huge development efforts (...). Complex to a degree, they have an important goal: rapidly developed small releases". Mohino et al. [57] confirm this in their reading of the agile manifesto [59]: "[this] generally implies not considering security activities during the life cycle (...), in addition to not taking into account the support and supervision by an expert or security team". If anything, OT projects are not generally simple or rapid. This begs the question whether agile interpretations are suitable at all for OT or whether OT development could use some more agility.

While Mohino et al. [57] recognise the shortcomings of agile with respect to security, they also describe an important pitfall for traditional incremental security approaches: "it is increasingly necessary to immerse in agile methodologies to take advantage of the flexible response to functional requirement changes and reach a high degree of quality in software development projects". This is exactly a problem as described in this research with respect to OT R&D, showing the conflict between the traditional approach to security and current needs.

### Crystalising agile

Ultimately, the project-structure does not really matter, provided that stakeholders realise that security is a bit odd. With an eye on the needs and wants of security and with flexibility towards overly strict organisational paradigms, it should be possible to adapt any project to include security successfully. In this research, a multi-phase DLC approach is used (see ch. 8.2.2). This approach could also be agile, where each phase has several sub-projects that each have their own agile team which works in parallel to- and in collaboration with the other sub-projects.

What ultimately matters most for the SOTDLC is the way agile deals with the wickedness of problems. For the SOTDLC, the following principles are crucial, regardless of whether an agile approach is actually used (based on [35, 57, 59]):

- Frequent review of requirements and goals, welcoming changes (as is done with a DLC, see 8.2);
- Shared responsibility for quality;
- Collaboration between business people and engineers (e.g. with a security coach, see ch. 4.7.1;
- Promote sustainable development ecosystems, where the feedback-cycle can be repeated indefinitely; and
- Simplicity is essential.

Of course there are more agile principles, supplemented by Mohino et al. [57] with their principles.

# Part I

# Theoretical Framework

# Chapter 5

# Risk

## 5.1 Overview

This chapter looks at risk by splitting it into two components: *strategy* and *impact*:

$$\textbf{Risk} = \textbf{Strategy} * \textbf{Impact}$$

1. **Strategy** = $Threat\ Intelligence + Adversarial\ Model$
2. $Attack$ **Impact** = $Assets + Value$

This way of looking at risk is uncommon: usually the equation $likelihood * impact$ is used, which allows for visualisation of risks in colorful (green, yellow, red; denoting risk acceptance, see ch. 6) risk matrices that are easy to understand. The difficulty for security is, however, that the likelihood of attack attempts on a system is 100%, as long as there is anything of value in the system, which is always the case for OT. Likelihood is thus not a meaningful way to look at security. The question is not *if* an attack will occur, rather *how* those attacks will occur [37, p. 40]. How will the inevitable attacker react to the system and adapt to its challenges? What is their *strategy*?

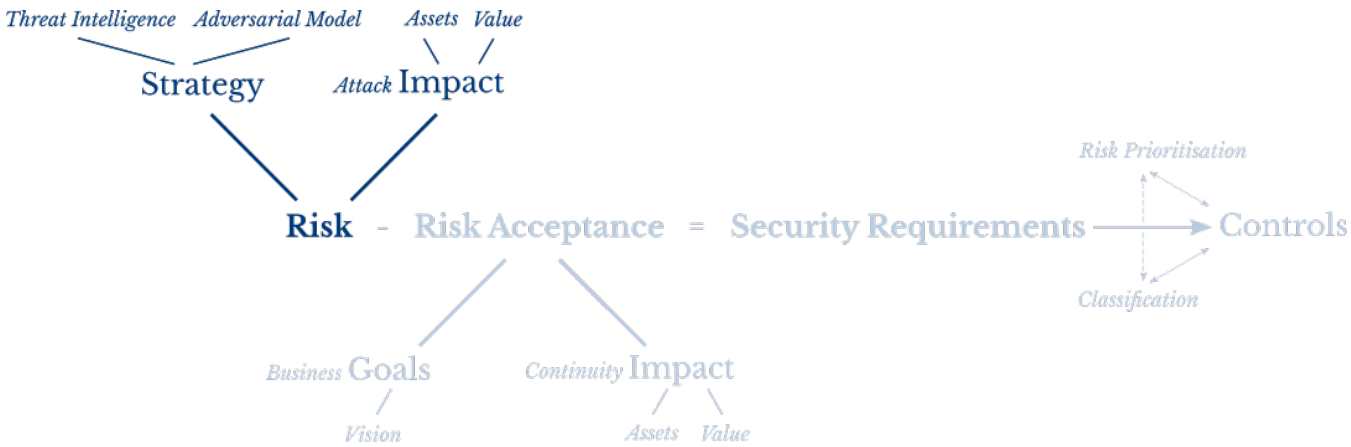


Figure 5.1: Risk Equation: Risk

**The approach** to threat assessments as taken in this research is different than what many would expect based on common practice. This approach is more suitable here, because dissecting risk into its sub-components *strategy* and *impact* allows careful scrutinising of these sub-components. This is necessary because the interests and practical considerations for OT are occasionally fundamentally different from IT, especially in innovations: IT based practices must therefore be evaluated at a fundamental level as well.

The core idea of understanding risk, is to understand what is causing that risk. For cybersecurity, an important part of this is looking at the attacker. The discipline of understanding, predicting, and preventing attackers is (cyber-)criminology. This field is thus of paramount importance for accurate threat assessment and risk analyses. Figure 5.2 shows criminological research areas that are relevant for the security equation in yellow.

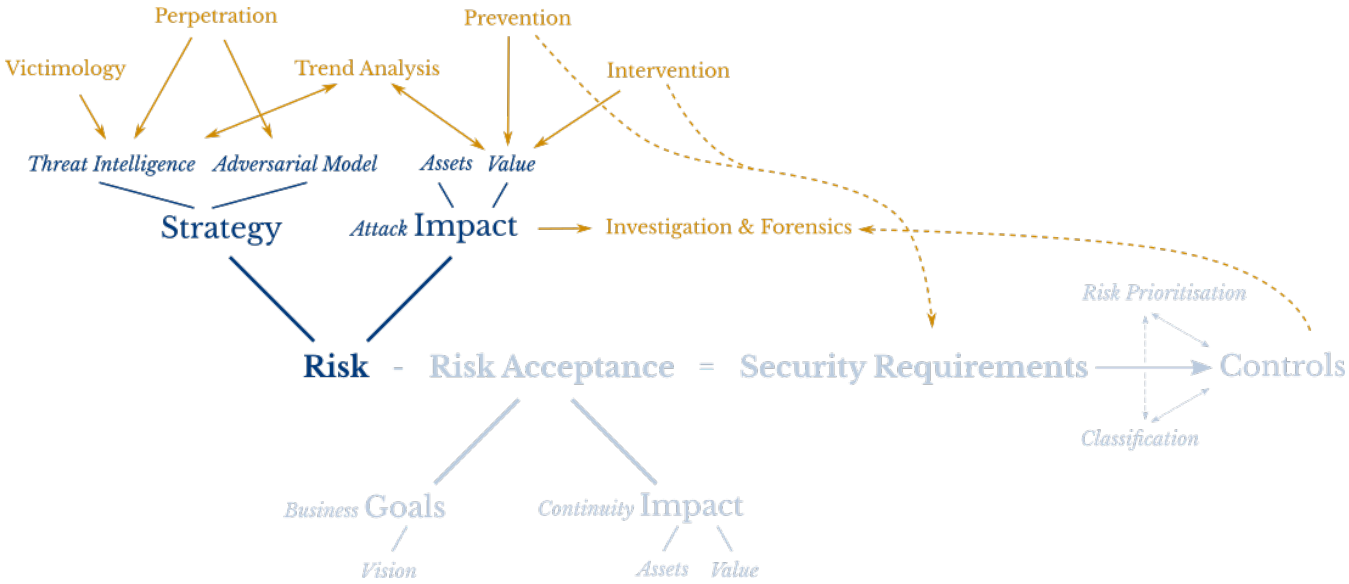


Figure 5.2: Risk Equation: Working together with Criminologists. Criminological aspects in yellow, Security in Blue.

After describing strategy and impact individually, they are combined into a risk analysis strategy specific to OT at the end of this chapter. This strategy is later used in the SOTDLC process, in chapter 9.

### 5.1.1 Further steps

This chapter introduces ways to talk about value, assets, attackers and risk. OT is complex and requires detailed approaches to business goals and threat-modelling. Based on the models in this chapter it is known what should be considered when looking at threat-modelling approaches in the SOTDLC. What remains is applying this knowledge to the structure of the SOTDLC process. Chapter 9 goes into depth about the project model, stakeholders, and various interests. Here the threat-modelling steps will be made more concrete for the various stages of the project. Chapter 10 will later give insight into stakeholders in the threat-assessment.

## 5.2 Threat Assessment

This section describes how strategy (sect. 5.3) and impact (sect. 5.4) can be used in threat-assessment.

A *threat-model* is the result of the threat-assessment. It is a structured description and valuation of security related threats to a SuC. Threat-assessment (or -modelling) is a common activity in both IT and OT, where usually a threat-model is made first, whereafter security controls are implemented.

Quite a bit of research has already been done for threat modelling for OT in particular: Hollerer et al. [8] have recently listed quite a few threat modelling techniques for OT. Unfortunately, none of these techniques are quite suitable for the SOTDLC, as described in section 5.2.2.

### 5.2.1 Differences Betweeen OT and IT

OT is complex and Expensive. Projects are difficult and large. The advantage is that to attack such systems, specialist understanding of the architecture and hardware are required. This knowledge used to be quite rare, but is becoming ever-more common now that attackers are catching on to the value of these systems.

There are quite a few things that clearly distinguish OT from IT. In the risk equation, this most heavily influences the determination of risk. Therefore, important differences are discussed here, to substantiate later choices in the threat-modelling approach.

#### Safety and Regulation

A major difference between traditional IT and OT is that IT manages data, whereas OT manages physical processes. Kaspersky warns about this in their report about the state of industrial cyber security [60]: "Physical assets can be manipulated or even destroyed by cyberattacks. Criminal organizations are now exploiting these possibilities as a business model". Because of the physical nature of OT, safety is an important concern: physical consequences can harm people, property and the environment.

Safety has been an important topic for OT for a long time: as of now, safety has been well-regulated in all aspects of industrial policies, standardisation and law. This means that changes to existing OT can be very difficult due to bureaucracy. Furthermore, OT is generally developed with long lifetimes in mind: commonly between 5-30 years compared to 2-8 in IT [61]. This means that technology is updated infrequently and does not change much.

Kaspersky did some research into what are typical bariers/delays in the implementations of ICS security projects[60, p. 17]. The three most reported problems were:

1. Approval takes too long (31%)

2. Too many decision-makers delay execution (23%)

3. Production stop for implementation not accepted (34%)

The first two illustrate that bureaucracy is really an issue for OT security. The last barrier ties in to the next section:

#### Availability has priority

Production processes and factories are often expected to be online 24/7/365 for multiple years on end without interruptions. Some production processes require multiple days to start and stop. With the small financial margins that these systems are often operated at, it is relatively very expensive to shut down a production process if there is no immediate need to do so. Security is, of course, generally preventative, and is thus not deemed important enough to shut down the factory for. Patches have to wait until the next maintenance interval, and are even then heavily scrutinised.

The age old addage 'if it ain't broke, don't fix it' is often a holy mantra in OT: changing things that are perfectly fine and have been so for time immemorial should not be changed, because they might break after the change.

Security is quite different, it quickly develops and guarantees about security, especially through time, do not exist. Luckily for this framework, the conservatism of OT has not quite adulted to its final form yet during the R&D process, where tests and changes are still allowed. The SOTDLC can conveniently disregard the production-stops, because they are not part, or an accepted part of the development process anyway.

### 5.2.2 Requirements to Threat-assessment Strategy

To produce a suitable threat-assessment strategy, it is important to understand its goals. Hollerer [8] describes an overall threat-assessment approach:

1. Model the system to build, deploy or change
2. Find threats based on that model
3. Address the identified threats
4. validate the completeness and effectiveness of the countermeasures

Muckin and Fitch use the mnemonic "There are no idle (IDDIL) threats - they attack (ATC)" [37, p. 7], see fig. 5.3.

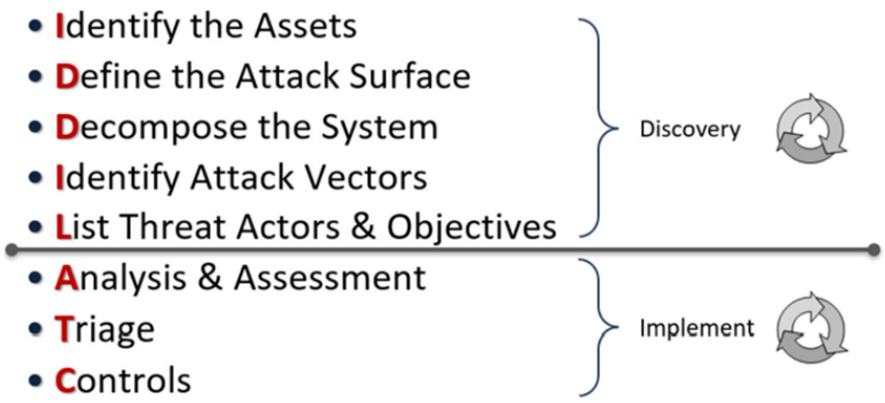


Figure 5.3: IDDIL-ATC, [37, Figure 4, p. 7]

These methods are quite similar to what is found in software development, e.g. at OWASP. The exact execution of these approaches (e.g. using fault- or attack-trees, or UML diagrams;

or modelling the system in a DFD) is quite well researched and defined. These methods are not really different for OT than for IT, and are well-known for both. Therefore, they are not discussed further here.

### Triangulation Approach

The difficulty of both the above approaches to threat-assessment, is that the system architecture must be known to quite a degree of detail to make meaningful assertions about value, impact and security. It helps to define a rudimentary system architecture to reason about general threats and potential security challenges later in the design process, but this can hardly be called a robust approach to security.

Although Muckin and Fitch [37] argue against looking at threats at asset-level, this does provide a useful supplement to threat modelling during early stages of the development project. Combined with other techniques, such as the black-box asset modelling approach (ch. 5.4.1), and strategy assessment (see ch. 5.3), this can be used to build a triangulation methodology. Through triangulation a reasonably accurate estimation of the true threats can be made.

Triangulation here, is defined as the usage of multiple independent strategies for estimating the same hidden variable. The degree to which the different strategies give similar results or insights is a measure of expected accuracy.

### Adaptive Approach

Not only the degree of knowledge about the system varies throughout the R&D process, but also the speed at which changes are applied. At the beginning of a project, it is likely that many directions and concepts are tried and tested. Later in the project, core concepts and architecture are known, and big changes become decreasingly likely. Security benefits from certainty, so where it should allow for flexibility in the beginning, to not impede progress, it should evolve with the project. As the system matures, so does its security: over time adding security at deepening levels of depth.

Overall there are two realms of adaptability that the method must provide:

1. dealing with different amounts of knowledge about the system architecture; and
2. dealing with varying speeds at which changes are applied on the system.

### Dynamic Determination

Because the SuD in the SOTDLC is prone to rapid changes, especially during tests or experiments, project teams must have a way to equally rapidly assess risks to determine the acceptability of sudden changes or engineering choices. For example, when a sub-component is added, removed, or replaced, the engineer must decide whether that is acceptable: they must understand the security requirements of the component at hand and be able to interpret these in real-time.

## 5.3 Strategy

In general for new OT concepts, subsystems could be tested or simulated in a controlled, secure, environment. Such an environment effectively eliminates most of the opportunity

for an attack, and even if the attack succeeds the consequences are small. The developers can choose between security-by-design or retrofitting security measures. Especially for projects where a lot of fundamental changes are made as lessons are learned, it may be tempting to skip the security until the end, when it is sure that the component will actually be used.

When developing an improvement to existing systems or even some new systems, such an approach may, however, not be feasible: such systems will eventually need to be combined into an actual working test-system for testing of its entirety. Other systems must be tested in practice, because simulation is not practical. For ATO-systems this also holds: the test trains drive on the same tracks as regular trains, albeit at night, and use the same rail traffic management systems as regular trains. Furthermore, trains that are used at night for ATO-system or -driving tests, may be used the next day in normal operations. The SuD is such a case in a 'hostile' environment, where threats are also very real.

This research is specifically focussed on this latter category of projects: systems so large, complex or specific, that they need to be developed in a hostile environment.

**The Cyber Kill Chain®** by Lockheed Martin [37, 38] describes the steps taken by advanced attackers. The idea is that the attack-path (the ckc) can be interrupted at any stage to stop an attack. Ideally, measures are taken to prevent the success of an attack at multiple stages. The expected strategy of an attacker is important here, because this determines which assets are likely to be targeted within the various stages of an attack.

The strategy of an attack in a hostile environment is dependent on that which causes the threat. In our definition of security (see ch. 2.1.2), security is defined in terms of assets and violations. Here, impact is the consequence of violations to assets; and strategy is determined by the attacker's desire to violate this particular asset or set of assets in an attack path. Therefore, to understand the strategy of an attack, we must understand how the attacker and the asset relate to each other, and to their context in the world.

### 5.3.1 Adversarial Model

First, we must understand the adversary: the attacker. What is their motive, how capable are they, and what are their means? Ultimately the goal for the SOTDLC is to understand the possible approaches that an attacker would use to influence or access the system. This allows for a tailored approach with regard for priority and effectiveness over efficacy.

Currently in security, mostly simple models of attacker properties are used to give rough insight into attackers (engineering perspective). Unfortunately, these models are inaccurate and not sufficiently effective to optimise security (and thus save costs). Therefore, a more criminological view could give more precision and insight (see figure 5.2). Unfortunately, not a lot of research is currently done on cybercrime in organisations, let alone OT.

#### Engineering Perspective

Rocchetto and Tippenhauer [62] have written a great comprehensive (freely available) article on attacker profiles in cyber physical systems. Based on extensive literature review, they define and motivate terminology, fundamental attacker properties and attacker models. They even provide an open-source tool to customise and better understand incorporated models.

An important shortcoming of this article, and of most of the articles it cites, is that they are entirely a theoretical exercise, based on (mathematical) reasoning and consensus about probable properties and behaviours of attacks. A common basis for this is taxonomy of known

common attacks combined with hard-earned experience and common sense. In general, such models are a good start and do a decent job. If looked at critically, however, this is a case of *survivorship-bias* (and by extension selection bias): mostly successful attack strategies and defensive experience with attackers are known. Failed attacks, and more importantly the elements that make such attacks fail are thus excluded from the models. In essence the attacker models by Rocchetto and Tippenhauer are great at attacker classifications, after attacks have manifested. They are not really great at predicting attacks.

---

An interesting exception to the former paragraph is the paper by Matusitz [63] (as cited by Rocchetto and Tippenhauer [62]), which takes a philosophical approach to examining cyberterrorism through post-modernistic chaos-theory and game theory; an interesting, though slightly pessimistic, read. This paper confirms the view that defining cybersecurity in static, linear measures does not work:

> "Without a doubt, conventional strategies - tools, weapons and tactics as we know them in the real world - lead nowhere against opponents who carry out attacks in a postmodern [complex, multi-faceted and uncertain] manner. The cyber attackers and the people trying to protect networks or systems not only engage in real-time game play; they also use strategies that are not conceivable in conventional conflict."[63, p. 181]

It should be noted that Matusitz' claims in this paper are highly debatable, as my colleagues and I discussed in an essay [64]. Furthermore, Matusitz seems to claim that in the post-modern, chaotic attackers we see nowadays do not allow modeling, which would contradict the whole premise of Rocchetto and Tippenhauer that tooling can be made to choose the most appropriate model: it is not possible to choose an appropriate model, if you support the claim that modelling is not possible in the first place. It thus seems strange to me that Rocchetto and Tippenhauer would cite Matusitz in their paper, even if it is just for substantiating their definition of a cyber-terrorist.

---

The common attacker-models in security engineering do not show the dynamic human perspective [65] of the attacker: in a sense, the attacker is seen as a rational and passive agent, who makes decisions and behaves based on rules. In practice, however, many more factors play a role such as prevention, external reputation of a company, ease of access, etcetera.

The field of criminology covers all of these additional elements, by looking at actual manifestations of attacks and attackers in both qualitative and quantitative research and is able to more accurately predict behaviour and associated consequences.

### Criminology

The criminological perspective can be broadly split into qualitative and quantitative research. Qualitative research is more about the attacker and their modus operandi on a case-specific level, whereas quantitative research is about statistical evaluation of many attacks and commonality of various modi operandi (i.e. threat intelligence, see ch. 5.3.2).

The disadvantage of the engineering approach is that it only provides insight into the passive capabilities of an attacker. Research suggests [65], that by looking at behavioral factors such as decision-making processes and psychological effects, a more accurate, reliable and usable understanding of attackers can be built. This understanding could be used to influence decision-making of attackers [66].

Currently, there is often little substantiation to claims about the likelihood of attacks: generally such claims are based on the technical difficulty of an attack. This does not tell the whole story. criminological analysis can help motivate risk-assessment: some attacks are unlikely, simply because nobody is interested in performing them or because there are easier and more obvious ways to achieve the same result.

---

Appendix C contains a list of URL's pointing to criminological articles that may be useful. Unfortunately, there are also still a lot of open questions in cyber-criminology [67], especially when looking at businesses, and moreso in OT. This is considered a very important open wound to threat-modelling in both OT and IT: it is highly advised to read into this further before defining a threat modelling approach for your organisation. It is hoped that more organisations will work together to employ and encourage criminologists to further this research.

Cybercriminology offers potential for a large amount of new approaches to dealing with cybersecurity. Figure 5.2 shows many interesting research areas in relation to security risk-assessment. The research agenda "The human factor in cybercrime and cybersecurity" [67] gives an overview of open questions and potential research areas that could be useful to businesses also.

---

For this research we consider the important point to be that a common way of discussing attackers is necessary, especially in a threat driven approach, such as described by Muckin and Fitch (Lockheed Martin) [37]. Ultimately, it does not matter much which model (engineering or criminological) is used at the level of implementation. At a strategic level, however, this does matter (a lot), because more precise models allow for greater efficiency in security implementation; efficiency is of vital importance for any viable business-model. Nonetheless, it is out-of-scope for this research to dive in any further: this should be taken up by specialists at an organisational level, e.g. as addition to the threat-intelligence process (see ch. 5.3.2).

### 5.3.2 Threat Intelligence

Gathering meaningful knowledge about potential threats to a particular organisation or asset is a complex and time-consuming. In general, this is assumed to be an activity conducted on an organisational level, supervised by the CISO/CIO. The production of the threat intelligence reports is thus largely out-of-scope for project security and thus this research. We do discuss some methods of looking at, or gathering threat-intel, because this may be relevant for consideration at a project-level.

Despite the high-level nature of it, the threat intelligence should be combined with the attacker model. A purely vulnerability driven approach is bound to be inefficient or even ineffective [37, p. 3]. With both models combined, statements can be made about how likely some modi-operandi for attackers are. By extension this gives knowledge about the types of attacks that can be expected, specifically about the effort and means an attacker might be willing to invest. This in turn gives direction to security efforts.

#### Criminology

For most organisations, the threats are more likely to be using common attack-strategies and attackers: these are modelled in quantitative criminology through statistics about occurrence and attack-patterns.

Organisations and systems that are deemed critical to society are more likely to be targeted by attacks tailored to specific goals and interests of attackers. These organisations benefit from the support of good and recent criminological research. A good intelligence office in government could support this with qualitative research based on their secret sources, but cannot give sufficient insight alone. Trend reports and analysis is currently mostly done by private companies, who have a direct interest in selling products that are related to the trends they analyse. The security industry could really use independent criminological research, giving quantitative insight to make reasonable and accurate assessments about the reality of such threats.

It should be considered that security cannot be 100% great all the time and budgets are not endless. Therefore, this latter category of organisations should be aware that accurate threat intelligence is paramount to securing the right assets and architectures sufficiently and effectively. More average organisations will likely benefit more from following mainstream security trends and advice based on known attack-trends.

#### Getting hands on threat-intel through Honey-Pots

One method of getting high quality data that is highly applicable to a particular organisation is by using honey-pots. Currently, honey-pots are mostly used for detecting attackers [68] [69, p741], but they can also be used to gain understanding about attacker behaviour. Honey-pots are especially useful in environments where attacks are frequent and sophisticated (unsophisticated attacks are generally already prevented and detected and thus well-known, eliminating the need for further investigation).

This research method is called out here, because it is something companies can do themselves to gain precise insight into their specific risk-profile (although it would be great if they shared their findings for the better of others). Honey pots are already used in practice to gain insight into the ways attackers react to security measures such as deterrence and prevention [66]. At the same time, some work still needs to be done to refine and increase reliability of this new research tool [70, 69].

## 5.4 Impact Analysis

This section is primarily focussed on the perspective of an attacker. From a business continuity perspective in R&D, assets and value are different than from a risk perspective. In practice, an attacker has limited knowledge, and they must react to what is found in a system. Furthermore, an attacker has a broader view of value than the business itself.

If impact is considered from the perspective of the attacker, it could be described as the potential revenue or profit for the attacker. As stated this is heavily based on the impact on the business, i.e. continuity impact. What the attacker is able to gain (or perceives as attainable), depends on the SuC architecture and its value to the organisation. Note that progress towards value (i.e. a step in an attack) is also considered valuable in and of itself.

### 5.4.1 Modelling Assets

There are two useful perspectives on assets and value for determining risk:

1. Organisational perspective, which includes all of the available knowledge; and
2. Attacker perspective, which is limited by what is visible to them.

The classical assumption is that the attacker knows everything about the system, and that the attacker is able to accurately value assets. For most purposes this is good, because this stimulates toughness (see 7.3.2) which prevents attackers from moving further into the system once they gain a foothold. This position is equivalent to the first perspective, and is mostly considered in chapter 6.2.

#### Time

In deployed systems, which are not changed often, attackers have more than enough time to discover the system and do reconnaissance. A stealthy attacker is able to alternate gathering knowledge with making another move. Effectively this means that the attacker knowledge is equivalent to knowing all about the system architecture.

During early stages of development, the system changes quite frequently. At this stage, the system does not often have a major impact on business continuity yet. These two things combined make the SuD both more difficult and less interesting to attack. Of course such effects are variants of *security by obscurity*, which should be heeded. Nonetheless, when used strategically and carefully, such effects do help in practice by increasing required investments, efforts and skill-levels.

### 5.4.2 Modelling Value

Ultimately, an attacker wants to achieve their goals. To make meaningful comparisons between various attacks and their consequences to an organisation, the 'value' abstraction is used. Here the presupposition is that:

- either the attacker wants to gain as much value as possible; or
- it is intrinsically valuable to an attacker that the target looses value, e.g. for (cyber-) terrorism or vandalism.

If there is no value to be gained for an attacker, it is less likely that an attacker will put in the effort required for an attack. Of course, this is limited by the what the attacker is able to perceive.

On the opposing side, the organisation does not want to loose their value: they would rather increase it. It is important to note that the meaning of value to the attacker and the organisation can be, but are not necessarily the same (see ch. 6.2). At the same time a common method for attackers to transform business-value to attacker-value is ransomware.

#### Properties of Assets: Time-Adjusted McCumber

Various security models for describing valuable properties of assets have been proposed, typically focussed on software. The properties described in these models are valuable to the organisation, and are known for their ability to be transformed into *Attacker Value*.

The general consensus is that *Confidentiality*, *Integrity* and *Availability* (CIA) are the core properties that must be protected. Appendix A discusses many of the available models. There it is concluded that overall value of assets is accurately modelled using the following three dimensions:

1. *CIA*, which models value in terms of security properties;

2. *Storage, transmission, and processing*, describing the stages in which the information and its value can manifest itself; and

3. *Unrealised, Ongoing, Contained*, describing the temporal phases for violations to value.

This new model, *Time-Adjusted McCumber*, addresses and improves shortcomings of existing systems, while maintaining the implicit trend towards a more general value model. The disadvantage of this model is that it is more complex and less intuitive. For this reason, this method is unlikely to be suitable for usage by engineers and management. Nonetheless, such a more complex value model seems to be necessary to deal with the similarly complicated nature of the SOTDLC: similar to what was stated in section 5.3.1, this value model should be considered by specialists (e.g. security coach) who support engineers in their implementation of security and help management with making decisions.

This model is relevant, because it opens up the option for additional precision when this is needed. It gives users a tool to verbalise the relationships between value, risk, and causes of risk. Everyone is involved in security, so stakeholders must be able to precisely explain goals, interests, and consequences when it matters the most, especially in comparisons with other aspects of a system (e.g. usability, cost, computing power).

While appendix A describes the value-model from the organisation's perspective, it also gives insight into the attacker:

1. How can the attacker profit? e.g.:

   - *Confidentiality* yields value by selling data or publishing it.
   - *Integrity* yields value from the effects of modifying or manipulating some data
   - *Availability* yields value through ransom, or competitive advantage when the targeted organisation is unable to function.

2. Where is the value?

3. How much control over the value does the attacker need to profit from it?

#### Criminology: Prevention and Intervention

Analysing what business value can be converted to attacker value is only one side of the coin. There are also values to the attacker, that organisations can use to their advantage. For example, an attacker would likely wish to stay anonymous, such that it is difficult for authorities to find them. Attackers also value their time. If such aspects are understood, the organisation could also implement measures to increase the risk and required investments for attackers.

Unfortunately, little is known about perpetrators of cybercrime. Especially specialised attackers are little understood. Nonetheless, based on some research into perpetrators, their age seems to be similar to conventional crime, i.e. most perpetrators are younger than 20 [71, p116, ch 4.4.1]. There has been further research among young hackers into their cybercriminal activities, and there are indications that effective preventative measures can be taken [71, p 272-274 ch 8.4.2.]

# Chapter 6

# Risk Acceptance

The second of the variables in the security equation is Risk Acceptance. For security specialists, it is easy to forget context and think in terms of decreasing risk only. On the business side, people may be tempted to seek zero risk. These are both not, however, viable solutions: business priorities do play a role in selecting security controls; and systems are simply too complex to guarantee security. As Muckin and Finch also discuss in their threat-driven approach:

> "Risk acceptance and risk management criteria must be determined per each scenario, and shift appropriately as the business/mission objectives, assets, threats and risk variables vary over time. Empowered decisionmakers need to have the best information available to make educated decisions concerning risk acceptance and management." [37, ch. 5 p. 41]

Resources are generally scarce, and thus risk acceptance should be the chosen solution wherever reasonable to free up resources for those things that have a higher priority. Before decision-makers can make informed decisions, they must have the right tools.

Risk acceptance is split into two components:

$$\textbf{Risk Acceptance} = \frac{\textbf{Impact}}{\textbf{Goals}}$$

1. *Continuity* **Impact** = *Assets* + *Value*
2. *Business* **Goals** = *People* + *Purpose* + *Process*

Because of the nature of risk acceptance, its boundaries can be mostly considered long before any actual prototype of the SuD is realised. As will be shown in this chapter, both impact and goals are defined based on organisational vision and strategy, independent of the project. This also means that much of the work done to produce a risk-acceptance strategy can be reused: it is very similar across projects. Clearly, each project will need to develop this strategy into project-level risk-acceptance based on the SuD architecture as it emerges.

## Further steps

This chapter introduces ways to talk about impact and business goals. OT is complex and requires detailed approaches to business goals and threat-modelling. This chapter gives an


Figure 6.1: Risk Equation: Risk Acceptance

idea about what organisational facilities are necessary for good security practices in the SOTDLC. What remains is applying this knowledge to the structure of the SOTDLC process. Chapter 9 goes into depth about the project model, stakeholders, and various interests. Here the exact steps will be made more concrete for the various stages of the project. Chapter 10 will later give insight into stakeholders in the risk acceptance analysis.

## 6.1 Risk Acceptance Analysis

This section describes the overall process determining Risk Acceptance. Ultimately, a Business Impact Analysis (BIA) and business Goals are combined into a risk acceptance strategy for the project, through a risk acceptance analysis.

### 6.1.1 Requirements to the risk acceptance strategy

The risk acceptance strategy is effectively a common language between the business vision and security specialists. It describes a border between when security is good enough, and when it is not: it facilitates measurement and evaluation. The risk acceptance strategy is shared between multiple stakeholders, and should thus be written in a common language (figuratively), which is shared by all stakeholders.

#### Security Coach

The risk acceptance strategy should give the security coach insight into what the business needs from the security department. The security coach can translate the strategy into technical terms: what does it concretely mean for the SuD and thus the design team? The security coach can (more accurately) select and promote security controls that are promoted and taught to engineers and architects, increasing overall performance in the realm of 'security-by-design'.

This means that the strategy must reference system architecture: which assets have which value, and what impacts are and are not acceptable. The security coach should have enough information to cross-reference particular considerations with SuC specialists for more details.

#### Project Management

The risk acceptance strategy should be specific enough that project management can use it as a (crude) project-maturity metric. This means that core strategic concepts should be measurable and explainable to people without technical backgrounds.

#### Business Specialists

These people must understand how the project relates to the organisation as a whole. Upper management should be able to uniformly weigh the significance of projects for strategic decisions.

## 6.2 Continuity Impact

Impact on Continuity is analysed in a Business Impact Analysis (BIA). By definition, a company must remain operational - i.e. it must ensure business continuity - to ensure its longevity and achieve its goals. Anything that impacts said continuity is thus a fundamental threat to the organisation.

**In a BIA** , Assets (sect. 6.2.1) and value (sect. 6.2.2) are combined to give insight into their relative importance. Overall the BIA can be done as usual, and even before starting the project: the SOTDLC does not require special methodology at that stage. Later in the project, as a company becomes increasingly dependant on the success and function of the project, assets produced by the SOTDLC should be added to the BIA in a component-level analysis.

Similar to threat intelligence, the BIA is considered a part of normal security activities, directed on an organisational level. Furthermore, the BIA is useful for other quality-related fields as well. Analyses tactics for the BIA are a scientific field in their own right, far beyond what can be discussed in this document. It is important that an organisation or project has a security specialist/coach who is able to value the BIA for security, and apply it to the context at hand.

### 6.2.1 Assets

Value relates to assets; assets can be organised through architecture in systems. Half a wheel does not do half as good as a whole wheel: systems, in their combination of assets, produce (much) more value than the sum-total of the individual assets. Moreover, if a single asset fails, the entire system could fail. As the SuDis developed, it becomes more complete and more valuable to the business. Therefore, risk acceptance analysis is partially based on the system architecture.

In the SOTDLC, the SuD is, as the name suggests, under development: the exact assets are still unclear, and architecture is being designed. In early stages of the project, it is thus not meaningful yet to discuss the SuD as an independent architecture. Given that an architecture-based risk-assessment approach is not feasible in most of the project-stages, the use of the remaining possibilities is unavoidable:

- *Architecture abstractions* Use what is generally known about the system to determine its impact on context. By extension this gives constraints to the SuD to protect the CS.

- *Asset-level evaluation* Dynamically analyse individual assets and their direct interfaces/-communications as the components are built.

#### Black Boxes

The more information is available, the more detailed impact assessments can be, and the better security can be brought into alignment with business goals. Therefore, it is advisable to use reliable information that is known, to construct a non-volatile architecture overview, modelling components that will likely be in the final system. These components can be considered black-boxes while little is known about their contents, but where meaningful reasoning can be done about value and impact (i.e. considering a component a high-level asset).

#### The SuC as a whole

For contextual systems (CS), more is assumed to be known in terms of threat assessments. Through the black-box-architecture, it can be assessed what impact the current SuD in its developing state has on the CS. This way, a lot can be said about what is and is not acceptable in early stages, giving guidelines for engineers and architects.

#### Asset based risk evaluation

Building components with security in mind, given the constraints that are already available. A rudimentary assessment of the value within the asset should be made at design-time. Based on that, security controls can be designed into the asset from the start.

Early on, a business is not yet dependant on the functionality SuD in itself, although it may need the development project to succeed for continuity. This means that during testing, the SuD could fail without significant impact on the company. The CS, on the other hand, may be very important to the business: especially in cases of enhancement projects (type 2, see ch. 8.1). Asset modelling for determining impact on continuity is thus more feasible, because architecture modelling is possible: the CS already exist.

#### Knowledge

As the development project progresses, the SuD may start to have impact on the CS. Strictly speaking, the SuD still does not have much business-value at this stage: it mostly has value to the attacker as an entry-point or through persistence. Contrarily, knowledge gained in the R&D project may be very much of value. Especially if the organisation is dependant on the project, it is important to also secure the body of knowledge. Fortunately, this is not a case of OT security: this is an IT endeavor, which is well-described in literature and experience, and thus it is out-of-scope for this research.

### 6.2.2 Value

The determination of value to estimate impact from the perspective of a company is very different to the perspective of an attacker. As stated in chapter 6.2, whatever is valuable to the organisation may be of value to the attacker, if the attacker can convert 'business-value' to 'attacker-value', e.g. through ransomware or disruption.

Business-value, in the context of security, is ultimately based on the degree in which the asset contributes to the overall value-production of the company. If failure of an asset means

that the company cannot effectively produce anymore, this is a problem. If the same failure does not impact daily operations, then it is less a problem. Of the three dimensions used for modelling attacker-value (see ch. 5.4.2), the temporal (time) dimension of the asset plays an important role here: a common approach is to describe the time an asset can be offline, or compromised before this becomes a problem for the company.

## 6.3 Business Goals

Ultimately, the aim of a business is to produce value. Often this value is considered financial, but for governmental or non-profit organisations this could be very different. The type of value the organisation produces dictates the goals of the business and by extension its projects: the inner workings of the business should be in alignment with these goals. The vision of the organisation describes how the organisation wants to achieve its mission (value). Business ventures (goals) should be in alignment with the vision.

In some sense, the business goals are already taken into account in the BIA. Fundamentally though, the BIA is more concerned with the *present* state of the business, whereas the business goals are about the *future*. This is relevant when considering security: security is also time-dependant. Attackers need time to develop a position on a system, and abilities of attackers evolve through time. In some cases, some threat may be an issue now, but not later: if security controls to mitigate the threat take long to implement or are expensive, it may be more realistic to simply do nothing, or to look at impact reduction e.g. by implementing a sub-optimal solution that is good enough.

The analysis of project-related business-goals is divided into *people, purpose,* and *process* (ppp) in the rest of this paper. This is not a rigorously defined methodology in any way, but it shows how the business goals can manifest themselves in the various aspects of the process as well. This is good to consider later in this report when people, purpose and process are analysed more in-depth with regard to the project structure.

Clearly there is also a feedback loop going on here:

- Analysing ppp
- Applying insight into goals in security equation
- Update ppp with new goals, also learn from external security trends.

This means that there is organisational learning cycle going on, where goals are redefined based on lessons learned about the target-system: management should be aware of this.

The exact manifestation of the ppp triad is highly context-dependant: the analysis is thus left to the reader. Wherever relevant for security goals, the related chapters will address this. The precise definition of how to understand and analyse the goals of a business is a science in and of itself, far beyond the scope of this research. The focus for this research is mostly the application of the information in organising project security for R&D.

### Escalation

If the system-at-risk is critical to the safety of an OT system and the organisation deems it unacceptable that the system fails, then the conclusion must be that the installation of an unsecured connection to this system is unacceptable. If, on the other hand, the experiment is critical - *it must not fail* -, the only solution is to do contingency planning to guarantee

secure success of the project. Either way, security controls and incidents should not come as a surprise.

From the outset of this project, an important concern has been the practicality of the proposed standards: given that the whole point of R&D is the rapid development of knowledge, security should not prevent the project from progressing. Nonetheless, an organisation does have limits on how much risk they can take, even in an R&D project.

In general, at every moment in the project there are two choices:

1. Implement measures to bring security (and thus risk) to an acceptable level; or
2. Discontinue (part of the) project because the risk is unacceptable.

When the boundaries of acceptability are breached, the project is *escalating* in terms of security. Such situations are precarious: gambling beyond your capabilities to bear consequences is a dire thing to do. Nonetheless, many projects suffer from escalatory behaviour in projects, where decisions are continually made beyond risk acceptance. This is sometimes because of particular situational context that leaves no other option. Sometimes it is simply a matter of incompetence (e.g. following the fallacy of sunken cost [72]). Regardless of cause, it is wise to consider such situations because they are acutely threatening to the organisation. It is thus important to be aware of coping strategies to prevent (further) bad things from happening (de-escalation [73]).

(De-)escalation is especially important in the context of R&D: because of rapid and free development, escalation can happen very quickly and profoundly. Fortunately, de-escalation is also easier because of the flexibility of the project-organisation. Nonetheless, this means that the SOTDLC ought to consider anti-escalation strategies (e.g. [73]). The risk-assessment strategy should be flexible enough to deal with tendencies toward escalation. This could be done by building a strong case for certain security requirements or processes, and communicating it clearly, understandably and digestibly.

# Chapter 7

# Security Requirements

The previous chapters (5 and 6) have introduced the concepts of risk (engineering) and risk acceptance (business). Now we have a ways to discuss risks on all relevant levels of abstraction and granularity, with every stakeholder.

The *Threat Model* (ch. 5) describes threats to value within the project, and the *Risk Acceptance Strategy* (ch. 6) gives a threshold for what is, and is not, acceptable. Whatever is not acceptable must be mitigated through security measures (i.e. controls).

This chapter describes how the threat model and risk acceptance strategy can be valued accurately. To do so, the following steps can be used:

- Risk prioritisation
- Classification
- Controls


Figure 7.1: Risk Equation: Requirements

**Alignment with Business Objectives** Although architects and engineers know which controls exist, and how to implement them, they do not have infinite time and resources: it is not possible to protect everything optimally. A clearly defined direction is needed to direct

efforts the most effectively. This chapter discusses how to use the threat model and risk acceptance strategy to develop suitable requirements. These requirements should respect business interests (notably security), while being maintainable and realistic.

**Further steps** Together with chapters 5 and 6, this chapter gives broad pointers and substance to risk assessment for OT in dynamic, innovative environments. What remains, is applying the risk-equation in a process-architecture for projects in the following chapters:

- Chapter 8 *Purpose*: setting security goals based upon the project business objectives: this scales the outcomes of requirements definition based on what the customer/organisation wants.
- Chapter 9 *Process*: integrating solving of the security equation into regular process architecture.
- Chapter 10 *People*: analysing responsibility and ownership for determining variables in the security equation, and eventually solving the equation throughout the project process.

## 7.1 Risk Prioritisation

Because of limited resources, threats have to be ordered by priority. The natural, although difficult, way to do this is to use the strategy based ordering of the threat model corrected for business impact. The theoretical challenge here is that the security equation has non-linear feedback loops [63] (difficult to predict or even observe) where attackers respond and technology develops. This non-linearity makes it difficult to make accurate comparisons between risks. This feedback is not so much an issue in practice, since a rough ordering is sufficient to get a reasonable level of effectiveness.

Besides weighing risk, other factors can be the financial cost or time-investment needed to address a certain risk. The security coach should be equipped to balance these factors and choose final priorities. Any reasonable philosophical motivation should suffice here. Doing something is generally better than nothing, so if a choice has to be made between two seemingly equal options, we recommend choosing one randomly and giving it a shot. Efforts can always be redirected later if needed: keep the *fallacy of sunken cost* [72] in mind.

Ultimately, the goal of prioritisation is that it serves as a formal motivation to explain why certain choices were made with regard to security. Especially when another aspect (e.g. usability) is hindered by the security requirements, it must be clear why the loss of value in that other aspect is proportional. This motivation is an important part of clear communication within the project team and with other stakeholders.

## 7.2 Classification

Clearly, there is a need to discuss threats, risk, and business impact among people with mixed backgrounds. Such concepts can be quite complex to understand, especially for people without a technical or security background. Therefore, it may help to use classification strategies.

To understand particular threats, for example, it might help to classify certain attack types using a system such as STRIDE [74, 75, 76] (see appendix A). To visualise risks, risk matrices can be used.

Classification should be used with care, and only if this serves a direct purpose in communication. Blind, purposeless classification of threats, risk, and business-impact, has many consequences, many of which are unwanted.

Essentially, classification means that similar things are grouped and given a name, i.e. they are labeled. There are dangers to labeling threats or risks:

- Labeling may lead to blind application of security controls, because of bias. This is dangerous because variations in context or the manifestation of a threat may change which controls are effective. This is similar to stigmatisation in the analysis of mental illness, e.g. [77], which is a highly debated topic.
- Labeling may lead to selection bias, for example if labels do not cover all possibilities. This means that some threats may be left out of consideration simply because they do not fit in the groups that are used.
- Labeling could create the illusion that a threat was addressed, creating a false sense of security.

  This is an effect akin to placebo. The lone act of labeling psychologically feels like effective progress towards a solution, while this is not the case. Nonetheless, the perceived risk is thus reduced. Such underestimation is naturally dangerous, because the actual risk may then be larger than what is deemed acceptable.

There are many more potential consequences, for example as described by Pohl [78]. Although there is no direct scientific proof that these effects play a role in security, this is likely based on its proven prevalence in other contexts.

### 7.2.1 Risk Matrices

Because the risk of an attack is 100%, we consider *strategy* instead. This means that risk matrices are not a suitable modelling tool in their traditional form. Nonetheless, such matrices are useful when the nuances of looking at strategy are considered. Often, risk matrices are already based on prevalence instead of likelihood, i.e. they describe the expected interval at which an event will occur [79]. Based on strategy, meaningful statements can be made about prevalence, because it can be determined how much time attacks take (related to impact), and how many actors are reasonably (i.e. with potential for success) attempting to target a particular system.

The aim of a risk matrix is to give insight into the degree of risk in a technology (knowledge) independent manner. This also enables comparisons across different disciplines. Therefore, the definition of the dimensions of the risk matrix is crucial for decision outcomes.

### 7.2.2 Maturity Models

A type of classification that can *sometimes* be counterproductive, are maturity models. These models are used to describe the degree of proficiency an organisation has in e.g. security aspects. When used responsibly for abstraction to benefit specific purposes, maturity models are great tools. Typically, however, there is a reverse relation back to the labelled object through compliance efforts, which is not so great. Compliance in such a reverse relation is not an efficient use of resources, if it is blindly (i.e. no knowledge of details) enforced.

**Good Applications** Undeniably, maturity models are essential to evaluate business performance (auditability), compare organisations, and certify. In these contexts maturity labels serve a clear purpose for internal abstract purposes such as strategic discussions, performance evaluation (comparing apples to apples), and risk treatment on an organisational level.

In negotiations with third parties (e.g. suppliers, or regulators), measures of performance are also indispensable for contract enforcement or to show adherence to law and safety standards. In these cases compliance and maturity descriptions are the best we can currently do to proof certain standards. It is simply infeasible to give (and externally evaluate) a complete, nuanced overview of all argumentation and efforts within an organisation: abstraction is thus a valid tool.

**Bad Applications** Enforcing technical controls for compliance to a certain level of maturity is likely to be ineffective [37, p26, ch3]. As stated, labelling techniques have all kinds of risk. Many of these risks are related to the fact that labels are abstractions that reduce concepts to a set of core properties. When discussing security controls on a technical level, the omitted details are highly relevant. In common problems, standardised solutions (with which compliance is enforced) may help. OT is often not, however, uniform or standard in such a way.

For security efforts internal to an organisation or project, compliance enforcement in OT is therefore highly likely to distract from the issues at hand. Within a single organisation, everyone ultimately is in the same boat, wanting the organisation or project to succeed. Now, opinions may vary on which course is the best, but it is best to discuss this based on charts and compass (subject matter, SuD), and not based on a dowsing rod (obscure maturity models which are ill-suited for the SuC).

## 7.3 Control strategies

As described in chapters 4.7.2, 4.7.3, and 5.2, security in innovative projects must not be incremental but iterative. In R&D, time is an important dimension: the system changes frequently, and thus conventional security hardening techniques are unsuitable. On the other hand, traditional OT values such as availability are different as well: this provides opportunity to look at security more in terms of resilience and incident management than prevention.

Controls can be implemented with different goals with respect to threats, extending the definition of value (ch. 5.4.2) which describes the three stages of violations to value (threats):

- *Prevention* - controls implemented before a violation to value has been realised;
- *Toughness* - controls to reduce impact of ongoing violations to value, and to allow for easier, quicker or more effective containment of these violations;
- *Incident Management* - indirect controls (e.g. handbooks or incident practice) to enable the organisation to quickly and effectively deal with unexpected situations.

**On the specification of Controls,** many a standard and research paper has been written (e.g. [24, 80] or [37, p26]). Overall, security engineers and architects are well equipped to design countermeasures for particular threats. Therefore, this research need not reinvent the wheel: the definition of concrete controls is considered out-of-scope.

The definition of requirements leads to the selection of controls with respect to threats. Requirements can be specified on three levels:

1. *Administrative (high-level)* requirements directly relating to organisational or project-wide strategy or vision. These requirements are broadly and openly formulated to give general direction.

2. *System/component (mid-level)* requirements relating to an entire (sub-)system or component. The internals of the system are considered a black-box for these requirements: they do not contain technical detail, but do give clear boundaries for system performance. These are used, for example, when negotiating technical requirements for contracts.

3. *Technical (low-level)* requirements specify exact performance or design criteria. These requirements have a large impact on development, improvement, and maintenance; they should therefore only be used as a last resort. In these cases, the requirement must be accompanied by a rigorous explanation for why it is imperative to follow.

### 7.3.1 Prevention

Preventative controls attempt to prevent an attacker from gaining access to the system in the first place. The classic example is a firewall with rules specifying which network packets are, and are not allowed.

In the SOTDLC, it is important to stay flexible. An important pitfall is the specification of many deep technical requirements that are highly dependant on context. Such specifications are annoying if the system changes, as is often the case. It is therefore always preferred to specify requirements in terms of value and value-loss (see ch. 5.4.2 and 6.2.2). Such requirements are equally verifiable and enforceable, while leaving room for engineers to build and change architectures.

### 7.3.2 Toughness

Reality has taught us that an outer shell of security measures is insufficient for most systems. Moreover, SuD's in the SOTDLC are changing so often, that it is infeasible to guarantee a complete outer shell, without security holes. Due to the difficulty of considering all attack options for preventative controls, it also takes a lot of time to fully develop and adapt them.

Toughness is the ability to absorb attacks up to the point of failure. These controls address *ongoing* violations to value by looking at ways of mitigating impact and registering ongoing attacks. Examples of toughness controls are logs; real-time monitoring; backup systems or data; and physical overrides (see e.g [23, 29, 32]).

In the SOTDLC such approaches are advantageous, because of their inherent flexibility. Often detailed logging and software backups are already available during development, simply because these are convenient for engineering. With small adaptations, such systems can be made suitable for security too.

#### Defense in Depth

A modern idea for security is defense in depth, where security is not only an outer shell to the system, but also creates trust-boundaries inside the system to prevent attackers from moving through it.

For innovative systems in early stages of development, this is not really relevant yet for the SuD. For the CS, however, it is often very important that security is ensured. Defense in

depth principles can be applied to make sure that attackers cannot use the SuD to gain access to other, more critical systems. In essence, the SuD is seen as an external, untrustworthy system.

#### Maturing R&D systems

In later stages of development, where the system is becoming more stable, it is beneficial to keep this architecture: there is no need to remove extra layers of security, if they are not in the way of functionality. A transition of function is also possible: where logging systems are initially heavily used for system-diagnostics for engineering purposes, these can be transitioned to more security oriented logging and monitoring, with respect for privacy.

### 7.3.3 Incident Management

Incident management concerns threats that are either *ongoing* or *contained* (see ch. 5.4.2). Many organisations use the term crisis management instead of incident management. This suggests an unexpected state of escalated chaos where decisions have to be made hastily and frequently. In a competent organisation this is not an accurate term: security incidents are likely to occur, so to be surprised by them is irresponsible. As explained in this section, much can be prepared in advance, such as communications strategies (including drafts); ad hoc organisational structures; and cause/effect diagnosis tools.

Incident management encompasses resilience: much of the requisite planning and preparation is done in advance, and controls are applied to the system to make sure that it is possible to quickly get insights into ongoing incidents. During this phase, parts of the toughness controls can be used to speed up the identification of problems and recovery from them.

Quite a bit has been written about incident handling. A lot depends on a rigid and practiced organisational structure which allows for ad-hoc response. This means that incident response is primarily organised at an organisational level, and utilised at project/asset level.

Nist standards:

- 800-34 Contingency Planning Guide for Federal Information Systems [23]
- 800-61 Computer Security Incident Handling Guide [26]
- 800-86 Guide to Integrating Forensic Techniques into Incident Response [28]
- 800-92 Guide to Computer Security Log Management [29]

Another interesting resource is the guide by Public Safety Canada on "Developing an Operational Technology and Information Technology Incident Response Plan"[81].

#### Communications Strategy

Although research shows clear benefits to openly and proactively sharing about cyber incidents [82, 83], a common phenomenon among organisations is the fear of public scrutiny after a cyber-incident. Many companies do not dare share about incidents that have occurred, unless it is absolutely mandatory to do so (e.g. due to privacy violations).

Often in practice, speculation about what might have happened is many times worse than what actually happened, provided that an organisation was not actually seriously negligent or reckless. If rumours are going around, or if people must be informed eventually, it is best to do so soon and as complete as possible, to prevent such speculations. Similar to [82],

communication should involve clear information about what is (broadly) being done to solve the issue, and the likely impact is on stakeholders.

Of course incident disclosure should not help the attacker in their intelligence position: organisations must be aware of the non-linearity of security in their communications strategy. Especially when communicating with the public or third parties, it is not necessary to disclose what the company has done to mitigate future to show competence.

### 7.3.4 Control Design Without Implementation

It is worth noting that it may be useful to design controls that mitigate a certain risk, without actually implementing the controls. The following could be reasons for this:

- The likelihood of a risk occurring is unknown.
- There are serious objections against permanent implementation of the controls. E.g. they would have a lot of negative impact on other important aspects; or the controls would be very expensive.

In such a case, monitoring of the system and of threat-intelligence is extra important.

Specific for R&D projects, this approach may be useful when the project wants to innovate/move quickly without the burden of a specific security-control. The control could be prepared and fully designed, ready to go in case of an incident. Effectively, increased impact is traded for flexibility and efficiency here, while still reducing potential impact sufficiently to stay within risk-acceptance limits.

---

For example: there is a risk that motivated attackers (e.g. nation states) could target a crucial planning system. Nobody would die if it fails, but if it is offline or malfunctioning for more than a few hours, people would start to notice unacceptable problems and delays. Adding strong encryption and authentication to the systems would reduce its value output by 30%, and consequently cost a million euro's per year in added labour.

In this case, you might decide not to implement the encryption and authentication to save costs. Especially if there are no imminent threats from such motivated attackers. To still be prepared, a secondary version of the system is developed with the authentication and encryption. If the risk of attack rises above an unacceptable level (intelligence updates), a switch is made to the more secure system. If an attack has actually occurred, the recovery procedure could add a default switch to the more secure system for the time needed to fix security bugs.

---

# Chapter 8

# Purpose

The purpose of a project is determined by the business department in business goals (ch. 6.3): they define a certain need for a project. For development projects, the SuD ideally fulfills this need, and in doing so it produces value. The type, amount and relevance of this value is later determined in the Business Impact Analysis (ch. 6.2) and applied to the threat-model to get a good understanding of security risks (threats) towards this value goal.

Before we can start applying the SOTDLC to R&D projects, it has to be determined if it is applicable. Not all projects require the SOTDLC approach: sometimes a basic 62443 [9] (or another) approach is more suitable. The type of project (sect. 8.1) says a lot about whether CS exist, and what the base-type of project architecture is. Furthermore, the project strategy dictates whether an incremental or iterative approach to security is more appropriate. The project architecture in turn shows whether the SOTDLC is appropriate. This chapter will explore the properties that make these frameworks and standards unsuitable for general application, by discussing the fundamental models they are implicitly based on.

First the general types of R&D projects are discussed. Then the development life-cycle is analysed for a basic project-structure.

**Further Steps** After understanding the basic purpose, type, and structure of a project, the next step is to start solving the security equation. Chapter 9 *Process* describes how to do this effectively, as such describing the beating heart of the SOTDLC. After that, chapter 10 *People* discusses how the SOTDLC occurs within the organisational ecosystem, through responsibility, ownership, education and awareness.

## 8.1 Research and Development project types in OT

For R&D, we focus on projects with the purpose of significantly developing a new or improved OT system. For this project we do not consider ongoing maintenance and minor improvements to existing, operational, systems. Therefore, the project must be making changes to the system that inherently impact the ongoing security and therewith the safety of the system as a whole; or developing an entirely new system.

Coombs et al. [39] describe three different approaches that are suitable for different types of development:

1. New products or processes for major business impact;

2. Enhancement of core products or technologies to defend market position; and
3. Creation or development of new product technology platforms.

Type 1 and 3 projects involve a higher risk to business and safety, and are thus most interesting. Our case-study, the ATO-project, is a typical example of a type 1 project.

*Type 1 projects* depend on existing systems and processes. In case of OT, these systems are often complex and expensive and testing and experimentation may thus need to be done in a practical setting. This means that for every test and experiment, the security of the SuD must be up to an acceptable level to ensure secure execution of the test and to guarantee the security of the SuC for regular operation even after the experimental components of the SuD have been removed.

*In type 2 projects*, applied to OT, we assume that a basic level of security is present in the existing system. New features or parts must simply be brought up to security specifications before using them in practice. This can be done in accordance to existing OT security frameworks. Projects within this category are thus out-of-scope for this research.

*Type 3 projects* are only really interesting with respect to this research insofar as there is security risk during the development phase. If a new OT product is developed in an enclosed sandbox environment, the risk is low: the consequences of a security breach are minor, because the system is not important to any existing industry or business processes. In this case, security can also be developed iteratively according to existing OT security frameworks. If the research is not sandboxed, security concerns increase: depending on the exact implementation of the project architecture, security may need to be at a similar level to type 1 projects, where security must be considered before each test or experiment.

### Iterative Requirement

For this project, we assume the approach to be iterative or cyclic. This could be called Agile or spiral, but the exact term or approach is not directly relevant for this research. A non-iterative (i.e. waterfall) approach is not interesting for this research, since this would mean that regular (non R&D) OT-security practices could be implemented without problems. The dynamic nature of a iterative development approach is the principal reason for the need for a tailored security-framework. The iterative approach is as such used in most, if not all, complex and multi-faceted research projects.

Besides the scope of this research, type 1 and 3 OT research and development are likely to be cyclic anyway, since there are many unknowns and complexities: hence it is not possible to completely plan and analyse the problem at the beginning. Consequently, there must be at least one stage where the lessons learned from the first implementations and tests are analysed and fed back into an update of the problem analysis and the design and implementation of improved features based thereupon. A project that does not cycle at least once is considered simply an implementation project (type 2) which needs no further research findings to achieve a viable end-result, fit for production.

## 8.2 Development Life Cycle (DLC)

After the purpose and type of the project have been determined, a suitable process can be chosen. The DLC is the beating heart of any research and development project. Regardless

whether it is a single- or multi-phase (funnel) project, the DLC makes sure that lessons are learned, fed back to, and applied in the end-product. This research aims to be broadly applicable to most types of OT R&D, including most sizes of projects: to this end, the DLC is considered in both single- and multi-phase projects.

The basic structure of a R&D project was defined by Coombs et al. [39]. This structure can be taken as-is for structuring a single-phase project. For larger projects, the structure is adapted slightly and individual steps are cycled to facilitate multiple distinct phases.

1. *Scoping*
2. *Project Specification*
3. *Detailed Planning*
4. *Action / Review cycle*
5. *Completion and Delivery*
6. *Post Project Evaluation*

The basic structure changes depending on the type of project. For most of this research it is sufficient to consider the general case, but the implications for OT were discussed in section 8.1.

### 8.2.1 Single-Phase Projects

Here we describe a common approach [84, 85, 86, 87] to defining single-phase projects in software development. Overall, the structure is comparable to the structure by Coombs et al.

1. Planning
2. Requirements Definition

---

3. Design
4. Implementation (Adding functionality)
5. Testing (review and feedback)

---

6. Deployment
7. Maintenance

The parts 3-5 are the actual cycle.

Planning and requirements definition are predominantly done at the beginning of the project (or phase in multi-phase projects). Sometimes a re-definition of requirements or an intermediate deployment may occur as part of a cycle. The specifications and architecture are updated based on the findings after each development cycle through feedback.

Some steps of this cycle may be performed multiple times in a single cycle (e.g. design and implementation). Additional project-specific steps might also be added, but are out-of-scope for this research.

### 8.2.2 Multi-Phase Projects (Development Funnels)

As stated, there are many ways of managing a R&D project. A common way for orchestrating complex research projects, is through a development funnel. The development funnel is especially useful for larger projects that run for many years and involve a large team of specialists: in these projects multiple distinct moments of thorough evaluation; re-planning; and subsequent development are necessary. Such projects thus have multiple phases of iterative design, each following its own complete development life-cycle (DLC, see section 8.2.1) constituting of multiple cycles. For smaller projects a single-phase approach could be used, following the format in section 8.2.1.

There are many flavours of development funnels [88], but for this research we consider a funnel similar to the one proposed by McGrath [89], consisting of 6 phases (see figure 8.1):

0. *Concept Evaluation* focussed on Brainstorming and generating ideas.

---

1. *Planning and Specification* discovery process looking at viability of ideas.
2. *Design & Experimentation* developing specifications for what the end-product should look like in a broad sense.
3. *Test & Evaluation* refining the specifications and building reliability of prototypes. First specifications of maintenance and production processes.

---

4. *Pilot* practicing deployment of the final product. Refining maintenance and production processes.
5. *Implementation & Roll-out* executing production and maintenance processes.

The project can be discontinued at any stage, so the later stages need not always be performed. Between each phase there is a moment of review, where the project is evaluated. This helps to prevent escalation (ch. 6.3). Higher management has to decide whether they want to:

- *Continue* if the project is viable and valuable to the organisation;
- *Redirect* if a drastic change in course is desirable; or
- *Discontinue* if the project is not viable.

The phases of the funnel look very similar to the single-phase DLC as introduced in section 8.2.1. In this case, however, the phases each consist of their own complete DLC, with a specific focus. Furthermore, the phases of the funnel are not cycled: under normal circumstances, each phase is executed at most once. This means that each phase of the funnel is, in effect, an independant R&D project with its own DLC.

#### Phases

Phase 0 is out of the scope of this project. In this phase there is no project yet and there is no actual product to speak of yet, and thus no significant security risk.

Phase 1-3 are the actual research and development: here the System under Development (SuD) is still changing a lot. At this stage, the project is conducted by a special R&D team

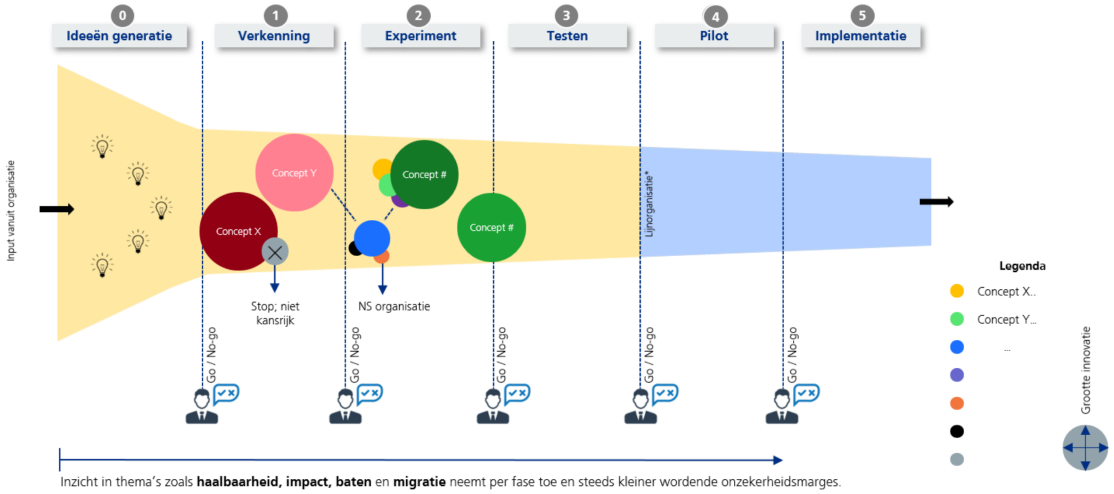


Figure 8.1: Example development funnel, as used by NS at the time of writing [17]

with various specialists. End-users and customers are not able to use the SuD yet, and there are no dedicated provisions or teams for maintenance.

Phase 4-5 are the deployment phase. Here the research and development is done, and the project is rolled out into actual use. To prevent teething problems phase 4 is used to gain knowledge about operating the newly developed system. In these phases maintenance is done and the system is in its final configuration.

### Generality

This way of looking at large R&D projects is considered sound, because multi-phase R&D projects can generally be reduced to a basic model of three stages:

1. Gathering ideas;
2. Developing/implementing ideas;
3. Deploying the products of the development.

In practice, however, most companies are interested in reducing losses to unviable projects: one way of doing this is by adding a larger number of distinct phases to produce a development funnel. This research aims to be broadly applicable and comprehensible for securing OT research. Therefore, a more frequently used model is more suitable. Nonetheless, this research can be generalised back to the basic three-stage model, by taking the strictest approach of each phase as the general approach for the entire stage.

Dealing with special flavours of two-stage R&D projects is left as an exercise for the reader.

## 8.3 Cyber Security Management System (CSMS)

Although the CSMS plays an important role in many organisations, it is only briefly addressed in this thesis. The reason for this is the fundamental view on it as used in this thesis: the CSMS is regarded as a support structure for security on project-level and only considered

as such. This means that security activities are decided and driven on a project-local level. These decisions are supported by security and business specialists, based on the risk acceptance strategy. This strategy is in turn defined on a more broad, possibly organisation-wide, level.

Ultimately, this means that security operations are not driven by the CSMS, rather project-level needs drive the CSMS! Consequently, security personnel such as ISO's, OT-SO's and the CISO are not directly responsible for decisions on security, but they observe what is needed for successful security and adapt (policy-)recommendations accordingly. They must be intricately aware of the state of security within their respective parts of the organisation, and should give expert advise to decision-makers; all to facilitate the organic production of good security.

**Further reading** Chapter 9.1 discusses what can be done at an organisational level to support project-level security. More about responsibilities and ownership is discussed in chapter 10 *People*. Here the logic behind the chain of responsibility is explained, reinforcing the ideas proposed in this section.

# Chapter 9

# Process

This chapter and chapter 10 *People* follow a structure inspired by the one used by Howard and Lipner in their SDL [35] (see ch. 4.2 for a review of this book). This approach to projects has been widely accepted and thus serves as a solid basis for the adaptations for OT, as proposed in this research. Many practical tips for actually implementing the SDL are given in their book.

**The differences between the SDL and the SOTDLC** are mostly in the difference between IT and OT. The most important differences are as follows:

- The approach to threat-assessment (Risk Analysis in SDL):
  - *A triangulation approach* is recommended here, which combines (early) designs with component level threat assessment. The SDL uses a more traditional approach with initial surveys to select priorities, expanded with detailed architecture models; this approach is not possible in OT R&D as discussed in chapter 5.
  - *A multi-phase threat assessment approach* is necessary. The SDL only supports single-phase projects.
  - *Criminology based adversarial model and threat intelligence approach*. The SDL does not address threat intelligence or attackers, except for general trend observations.
  - *3-Dimensional impact modelling* is used, which covers changes of value through time and information states. This replaces STRIDE, as used by Howard and Lipner. This is not only better for OT, but we would recommend using such an approach in large IT projects as well.
- Risk Acceptance Strategy as basis for security:
  Howard and Lipner do not address risk acceptance except for the paragraph "Determine the "Bug Bar"" (p. 74). This does not satisfy the needs of business and management and leads to strain, as described by Muckin and Fitch [37] (see ch. 6).
- Overall this thesis uses the *risk equation* to structure the threat assessment, and interface it with business and engineering needs.

**Building on the SDL** is done by discussing the whole process, and tweaking it to OT. This is done using the differences between IT and OT that were identified earlier. This thesis uses

a different structure that is more suitable for the security equation, and which translates to the development funnel model (ch. 8.2).

Table 9.1 describes how this structure relates to the 13 phases described by Howard and Lipner [35], and the project architecture (funnel). This table clearly shows that the mapping of SDL phases to the funnel-model is not straightforward. For example, the "Security Response Planning" and "Security Response Execution" were translated to more logical locations on the timeline for multi-phase projects.

**Further Steps** With a known baseline for security, the security equation must be resolved as the SuC changes. Based on the *what* in this chapter, more of the *how*, and the *who* are defined in the chapter *People* (10).

## 9.1 Organisational Aspects

A lot is done even before a project starts to organise security. Many of these aspects are described in the CSMS (policy). Howard and Lipner [35] devote an entire stage to education awareness: this underlines the importance of these topics in preparation.

### 9.1.1 CSMS

Cyber-security operations on an organisational are described in the CSMS. The role of the CSMS in relation to projects is discussed in section 8.3: it should describe facilities and activities undertaken at an organisational level. Furthermore, the CSMS contains policies and frameworks for project level activities to give uniform direction to security management within the organisation.

Examples of organisational facilities and activities that should be addressed in detail in the CSMS are:

- Threat intelligence (gathering and publication throughout the organisation).
- Adversarial models, if these are uniform across the organisation (otherwise, projects should produce their own appropriate models).
- Incident response and monitoring (incl. security operations center, 24/7 incident response teams if applicable).
- Specialist security team(s), which are available for consultation on specific questions.
- Basic Awareness and Education. Advanced Education should be made to measure for specific projects or security needs.
- Project level *recommendations* for security management (e.g. such as this framework), preferably adapted for the various project types.

Many of these things are often too complex for a single project to set up and manage. Moreover, these things are easily scalable and therefore suitable for setup at an organisational level.

The CSMS should fit logically into the organisational structure, and take into account other interests the organisation may have. As stated before, this is described in ch. 10.

The CSMS does not set in stone what security design should look like at implementation and project level. Rather it supports projects in their needs by giving guidelines on what has

| Phase Type | 0 | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|---|
| Funnel (Ch. 8.2.2) | Concept Evaluation | Planning & Specification | Design & Experimentation | Test & Evaluation | Pilot | Implementation & Roll-out |
| Security | Organisational Aspects | Baseline | Cruising | | Transition | |
| SDL (Stage) [35] | Education and Awareness (0) | Inception (1) | Best Practices & Product Risk Assessment (2,3) | Coding and Testing Policies (6,7) | Push (8) | Final Review & Release (9,11) |
| | | | Risk Analysis (4) | | | |
| | | | | Customer Prescriptions (5) | | |
| | Response Planning (10) | | Response Execution (12) | | | |

Table 9.1: Timeline of the multi-phase project

worked in the past. It does not prescribe responsibility and ownership of security activities that interfere with the project management structure, rather it observes, raises awareness, educates and recommends. More on this in chapter 10 *People*.

This section does not go into more detail. For project-level security, it is only relevant to know what supporting services and facilities are expected from an organisational level. The exact implementation of these things is another topic, which is out-of-scope for this research.

## 9.2 Establishing Baseline Security

Preparation for security operations and establishing a baseline security implementation, is conducted over two phases in the project: *Planning & Specification* and *Design & Experimentation*.

It all starts with raising awareness among project initiators and management. As soon as they are on board, a security coach (advisor in the SDL) should be assigned who starts working right away. The advisor starts making social connections with experts and team-members, and if needed assigns others to work on the project as well. This phase is finished by building an up-to-date inventory of all available (security) knowledge and assessing the current security status of the SuC.

### 9.2.1 [Planning & Specification] Initial Evaluation

As soon as a purpose and goal for the project have been established, and a project-architecture defined, activities to initiate the project are undertaken. For security this also means that preparations can start to be made. During project inception, things that are known about relevant systems (CS) are used to build an initial image of the situation, project needs, and scale. Here an applicable framework is chosen for security management, which could be this very one.

The first goal of the newly appointed security coach is to understand the goals and business side of the project. The coach uses this knowledge to make project management aware of is needed in terms of security to support these project goals.

At the start of any project, new team-members need to figure out what is going on. The security coach does this too by looking at known parts of the security equation:

- *Business Goals* should be quite well-defined at this point. The organisational (or even departmental) vision and business model should give insight into the greater context and direction.
- A *rough system architecture* can be made (with many black-box parts), to get a general understanding of *assets*, *value*, and resulting *Attack* and *Continuity Impact*.
- *Threat Intelligence* should be readily available within the organisation.
- The *Adversarial model* can be refined based on insights in criminology and the type of system that is being developed.

Together, this knowledge gives a rough idea what will be needed in terms of security and high level requirements.

### 9.2.2 [Planning & Specification] Build a team

Now that the security coach knows the system and its security properties, they can convey this message to management and stakeholders. This includes needs and wants for security, and seeks collaboration with other stakeholders. Based on how much work has to be done, a security team is assembled at this stage. For practical tips refer to howard and lipner, [35, Ch. 6]. For stakeholders, chapter 10.

### 9.2.3 [Planning & Specification] Education

Now we know how to consider security on an engineering level (risk equation); we know what the purpose is of security; and we know who is involved. The last thing that remains is to tell this story to stakeholders, and to have them understand and apply these concepts. A big part of this is education of people within the organisation

#### Complaints

A common approach is to send stakeholders (or security specialists for that matter) to expensive courses about some security related standard. Much is quickly forgotten, and never applied in practice. Educating an engineer on IEC62443 [9] does not teach them much more about security in OT than the tables of contents of the standards. A short internal meeting about risk management and threat-modelling in a specific context is much cheaper, and just as effective for a project.

There are some common complaints about security education practices:

- Security courses are boring, especially to people who do not specialise in it;
- Security courses are expensive;
- Security is complex and multi-faceted: formal courses that strive to introduce some kind of comprehensive framework lie about their comprehensiveness. In basic rules they may teach a lot, but it requires experience to develop an intuition about intricate security concepts and relations.
- Because it is too complex to give people enough tools in a short course to effectively deal with security, the alternative is often to teach about default security controls as standardised. This supports a compliance and vulnerability driven culture, which is not as effective as wanted.

#### Then what does work?

Howard and Lipner make a good point aboud education [35, p55]: knowledge about all the details of a cryptographic algorithm does not equate to knowledge about making software secure. In a broader sense, it could be said that knowledge about controls does not guarantee that these are implemented in a good security architecture.

Education within organisations should be focussed on what is useful for engineering activities. This fits into the critical systems thinking approach [31, 11, 32], where specialists work together to help decision-makers the best they can. In education this means selecting, editing and curating a custom knowledge base of relevant knowledge. This knowledge is than transferred to said decision-makers, architects and engineers, who then use this in their broader decisions about system design and requirements.

The security coach, with help of organisational specialists, develops an understanding of likely attacker strategies, and curates this into defense strategies. Later in the process, experiences and incidents are added to this. Preferably with help of an expert on didactics, the knowledge compilation is then transformed in educative content: for example into short weekly workshops, lectures or articles. The goal her is to make the discussion practical: it is not about some theoretical abstract system, it is about a specific system, as maintained or developed by this particular company, in this particular project by this exact team: this means concrete examples, and concrete collaboration to fix security problems in our system.

This is the core of aspect management by the security coach: not to convince others of their viewpoint by constantly saying how important security is, but to help people in making their own decisions by giving them the tools they need. Sometimes such decisions will positively impact security, other times they will not. This is both OK. The security coach should not judge about this, rather they should observe and update their risk assessment and view on the system, which will eventually be reflected indirectly in the educational material again.

More about this process in section 9.3.

### 9.2.4 [Planning & Specification] Solve Security Equation for CS

Contextual systems (CS), e.g. existing systems that are improved by this project, will be influenced by changes in their context. The system under development (SuD) may increase the risk to the CS. Of course, business continuity and value must stay protected (assuming CS security was up to date): therefore, the security equation for these systems must be updated and, if necessary, countermeasures as well.

Before experimentation starts in the next phase, the security of the CS must be dealt with. It does not have to be in its final form yet: at such an early stage, it may be enough to have a (manual) switching system to detach the SuD, to return to the original state of the CS. Nonetheless, an attacker may propagate from SuD to CS, before the switching system kicks in. If this is a threat to the CS, countermeasures tailored to this connection must be implemented.

As the SuD grows, it may have an increasing impact on CS as well. This means that the CS should also be periodically re-evaluated during the cruising phase.

### 9.2.5 [Planning & Specification] External Stakeholders

If third parties are going to be involved, it should be clear how their security should relate to the project. An important factor here is the education and involvement of internal actors in the security processes: these people are ultimately responsible for weighing the different interests of the organisation in negotiations with third parties.

The security coach must be prepared for such cases: this means they have to become aware of parts of the developments that may be done externally. They should research and understand strategies to deal with expected challenges. It might be necessary to involve market-experts and to deal with unknowns.

Beware that in OT, third parties do not often achieve a desired quality of security in their products. If this is the case, the project must be prepared to take measures internally.

### 9.2.6 [Design & Experimentation] Triangulation

As the project team starts with designing the actual system, the SuD starts to take more shape. At this point, security should already be considered by the designers (thanks to the support

from the security coach). Now is the time that the security coach starts modeling progress in the system and its security, to decide on a way forward.

This is the first time the new system (SuC) will be evaluated: this is different from other security equation solving evaluations. As of this moment, no experimentation has occurred yet. The system is still *sandboxed*. In practice components are already being experimented with individually. These parts are not yet physically part of the SuD yet, although they may be part of its design.

This stage is remarkable because of the sandboxing and component level analysis in a few ways:

1. There are risks on a component level, but these have very limited impact on the CS, because the attacker does not have easy access.

2. Risk acceptance is high, because the SuD has little direct value to the organisation yet; value is concentrated in intrinsic component value and valuable knowledge.

3. The CS is being adapted to facilitate the new SuD, which changes the game for security.

4. Knowledgeable attackers may be willing to invest in attacking the SuC now, to get a better position for future attacks or to use the SuC as stepping stone to attack a supplier.

In this phase, focus is on resilience of CS, because the value in the SuD is still comparatively small. Knowledge can be very valuable, especially in a competitive business environment, but is part of regular organisational administrative security procedures.

To complement these remarks, security analysis is conducted in two ways, to intersect (i.e. triangulation with two sources) the actual risk for the current SuD.

### Early System Architecture

Architecture analysis is conducted using system sketches and fixed CS structures. This can be used to discuss likely attack strategies and how these might impact the CS. As stated, harm to the SuD only, would be only a minor event from the perspective of business continuity: the business does not yet rely on it. The priority is thus to prevent or reduce damage to existing systems that do have such organisational value. Architecture analysis shows how the attacker could use early versions of the SuD to reach the CS, and thus what potential entry-points are that must be protected.

### Component Analysis

Commercial Of The Shelf (COTS) components are often used in OT as building blocks for new, larger systems. Such components are developed, maintained and secured by separate companies (third parties) and are often built for function and not for security. Therefore, they must be analysed carefully to get insight into trust-boundaries and attackable surface.

## 9.3 Cruising

The goal of this phase is to keep updating the *security equation* as the project progresses. This could be done by re-evaluating it wholly, but this is too cumbersome. Therefore, this section analyses the aspects that remain constant, and which properties can be predicted in advance

(see ch. 5, 6 and 7 for more details on these variables). All these components can be optimised: this leads to a straightforward iterative partial-solving approach. An overview of the categories can be seen in figure 9.1.

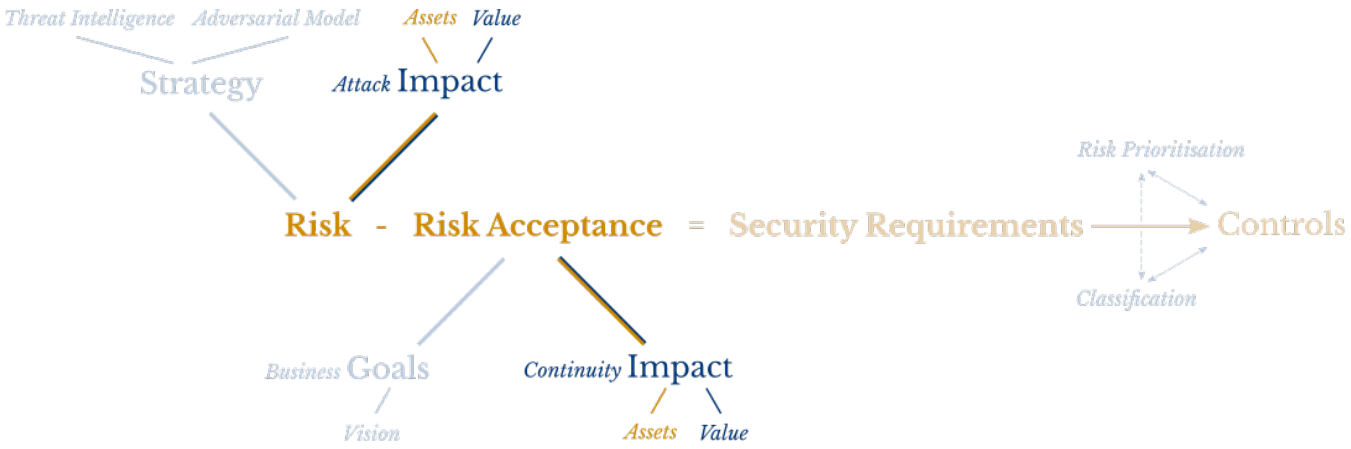


Figure 9.1: Risk Equation: Constants (light-blue), Predictable Variables (dark-blue) and Unpredictable Variables (Yellow).

### 9.3.1 Constants

**Business goals** are fairly constant. Organisational learning will lead to updates in business model and consequently to changing goals. On the scale of a project this is of little concern: if impactful changes occur, these can be dealt with on an ad-hoc basis.

**Strategy** of attackers is also stable on the scale of projects. As the system develops, the dominating value structure may change, but this does not modify the overarching adversary abstraction that describes these values.

The same stability is true for threat intelligence. At the time of writing of this report, threat intelligence is mostly gathered on a yearly basis by companies specialising in cybersecurity and governmental organisations. Since threat intelligence is a more organisationally driven endeavour, the project itself need not involve in this actively.

**Future** As developments in criminology driven threat intelligence continue, future SOTDLC's may need to consider ad-hoc response to finer changes in threat-actors and -intelligence. Such developments are, however, not expected soon. It is now difficult to predict what will be possible in terms of *threat-directed* security for process-efficiency, where knowledge about the current threat-landscape drives security.

### 9.3.2 Predictable Variables

It is not necessary to reassess these variables with every change to the SuC. Throughout the project, predictions can be refined as needed (e.g. at the start of each phase, or after major design updates). This saves a lot of energy, while retaining almost the same degree of accuracy.

**Continuity impact**

is determined by combining value and assets. Both of these can be predicted to a fair degree.

**Value** is correlated to the project-timeline: with each phase that passes, the decision to continue the project implies that the current system is sufficiently valuable to the organisation to warrant further and expanded investment. The fact that project progress is related to system value, means that dominating value structures can be determined ahead of the project.

Table 9.2 shows an example of an abstract project prediction. This prediction can be made more specific for sub-systems of the SuD, or specific parts of the attack-surface of the CS.

| Project phase | Focus | Dominating Value / Security Target |
|---|---|---|
| 0 | CS | Organisational facilities and operations |
| 1 | CS | Preventative wrt CS; Confidentiality of knowledge |
| 2 | CS | Preventative and Resilience wrt CS; Resilience wrt SuD; safety-driven |
| 3 | SuD | Resilience for whole SuC; Adding preventative measures for SuD |
| 4 | SuD | Reinforcing preventative measures for SuD to improve availability |
| 5 | SuC | Availability (reliability) for SuC |

Table 9.2: Value related to project phases

**Assets** are predicted by looking at abstract-design documents and determining the volatility of expectations for components.

For example for the ATO-system, it is known that

- it must have some type of vision of the environment;
- there must be interaction with the systems controlling the train; and
- it is known that there must be a central control unit that directs it all.

**Impact Analysis** Such a basic structure can already be used for component-analysis (done by engineers), e.g.:

- Vision has to deal with privacy of observed subjects (confidentiality); images must be reliable (i.e. maintain integrity) and must give an update regularly (availability).
- the train must not receive faulty instructions (integrity)
- the control unit has to give instructions to the train regularly (availability) and it must function accurately and reliably (integrity)

An architecture analysis can also be initiated by looking at communications and storage of information that is inevitably required in these components.

In reality this assessment should be much more accurate, with information about value through time attached, resulting in potential security-requirements and -controls. This gives a guideline to the security coach to prepare for security efforts throughout the project.

#### Attack impact

roughly describes the profit an attacker is bound to gain by successfully attacking the system. As stated in chapter 5.4, this impact is mostly dependant on continuity impact. Nonetheless, the way the SuC is interesting to an attacker changes as the system matures: it is thus also dependant on the project phase. This process is described in table 9.3.

| Project phase | Example adversary interest |
|---|---|
| 0 | - |
| 1 | Stealing ideas, code or component designs |
| 2 | All of the above; establishing a internal position for later use; using components as stepping stone to CS or supplier |
| 3 | All of the above; |
| 4 | All of the above; disrupting developments for competative advantage; exploiting SuD |
| 5 | All of the above; exploiting business continuity value of SuC |

Table 9.3: Adversary interests related to project phases

Such tables show the need for good criminological models to predict behaviour of attackers. With current knowledge it is difficult to make motivated statements about what an attacker is and is not likely to find important. This means that we must understand both the attacker, and the threat-landscape (i.e. context) of the organisation.

### 9.3.3 Solving What Remains

As the SuD is designed and developed, asset level design and configuration is also conducted. As new components are designed, implemented, or changed, security needs to be continually reassessed. This is done in the following processes:

1. Refining System architecture, done by the security coach based on what they learn from the system design. Based on this, the security coach directs the development of architecture-level security controls.

2. Conducting Asset level impact assessments (both from business and attacker perspective), informally determining security requirements, and dynamically selecting component-level controls. This is done while designing or implementing the component, by the responsible architect or engineer.

While engineering is in the lead for asset level security measures, they may not have the oversight needed for assessing architecture level security. The security coach helps with this by keeping track of the complete system, and making sure that architecture level requirements are kept up to date with engineering and business.

#### Delay

While component-level countermeasures can be implemented while the system-design is ongoing; it takes some time before the security coach learns about new components or configuration changes. During periods where few experiments or tests are conducted, this is not

much of an issue. In these moments, the system is back in *sandboxed* mode, and thus difficult to access externally. It is unlikely that an attacker will manage to coordinate a directed attack without direct access to the system. Component-level experiments may be lead to vulnerability of the component, which is caught by the design-time impact and security assessment.

When there are tests and experiments going on, engineers may need to swap configurations and components quickly due to time constraints. In these cases, it is necessary that the security implications of all potential swaps and changes that are planned during the test are evaluated before the test. This requires that all updates to the architecture are known and evaluated by the security coach beforehand. Naturally, security must be brought up to spec in advance too.

Sometimes, unexpected events may require sudden changes. Especially in these cases the pressure of making a quick-fix is prone to lead to vulnerabilities (e.g. swapping an internet connected component, and forgetting to configure the firewall). To reduce the risk of incidents, the security coach could train one of the testers for scrutinising ad-hoc changes, or maintain a presence during experiments themselves.

## 9.4 Transition

As the SuD matures into a complete and reliable system, its importance to the business grows. Off course, the goal of developing the system is to enable the organisation to produce more value in the long run. During previous phases the focus of security was more oriented towards protecting CS (see tables 9.1 and 9.2). Now that the SuD is becoming more valuable itself, there will be a clear shift in focus towards the SuD.

Throughout the transition phase, the SuC is being prepared for permanent deployment and maintenance. Security grows along with business impact of the SuD, into a robust and complete state. In the pilot phase, the system can be tested for security in the real world for prolonged periods of time, in essence proving the operational effectiveness of security. Later, during the final implementation and roll-out of the final-product, any open security problems and concerns must have been addressed.

### 9.4.1 Pilot

During pilots, the SuD is still under (minor) development, while the systems are simultaneously already used for their intended purpose. In this phase, it is examined whether the system is capable of delivering the desired value-production and whether its integration into regular operation is feasible: the objective is to get the system ready for long-term use and maintenance.

#### Testing Efficacy in Practice

The system is now completely being tested in a real-world environment, and only small changes are being made to optimise and perfect the new system. Since the system is used in a 'hostile' environment for prolonged periods of time, the operational effectiveness of the security is being proven. Because of the still limited dependency on the SuD, and the previously gained experience with the security of this particular system, this is a great opportunity to learn in practice, without unacceptable risk to the organisation. If any substantial vulnerabilities still exist in the system which have been overlooked, they may surface now. Although

this will impact the organisation somewhat, it is not yet too late to respond and mitigate these vulnerabilities, if they are exploited now.

#### Final Security Push

This is the time to resolve outstanding security concerns and bugs, and to start performing more rigorous and detailed security testing (e.g. black-box audits). Tools and systems that are part of the SuC, e.g. used for debugging or development, are a special concern here, since these must be disabled or removed before deployment. Howard and Lipner [35, ch. 13] provide directions on how to organise such a security push well.

### 9.4.2 Implementation & Roll-out

The final stage of a project is the actual deployment of the new SuC at scale. At this point, it is presumed that security is finished and up to specifications. To substantiate this presumption, a final review is needed to verify this.

After the quality of the security has been signed of on, there is a final move to maintenance. Systems and responsibility are handed over to the maintenance team in a seamless manner.

#### Final Review

The final review could be described as an audit of the security architecture produced in the SOTDLC. It is an expanded version of the inter-phasal (go/no-go) review (see section 9.5.1), which gives the final go/no-go decision before actual deployment. This review is performed either by organisational or by external security specialists, who have not been involved in the project before.

Based on what Howard and Lipner describe [35, Ch. 14], the final review should check the following things:

- *Process execution*. What was done in practice during the project, and does this give confidence in the resulting security? Particular focus on contracts and agreements with third parties.
- *Architecture Level Review*. Evaluating architecture level threat models and architecture overviews.
- *Component Level Review*. Evaluating component level threat models.
- *Remaining Concerns*. As described in "Unfixed Security Bugs Review"

The outcome of the final review should be a comprehensive review of the quality of the security in the SuC. If the quality is insufficient, the product cannot be deployed yet. Sometimes, there may be minor concerns that can be addressed later, after the system has been deployed in a patch or update. The facilitation of such patching processes should be made to measure: there is no (known) set way to deal with these things.

### Towards Maintenance

With the finalisation of a development project, often many team-members move on to do other projects. This also means that a lot of knowledge about the intricacies of the system move away with them. Such dilution of knowledge is a big challenge for maintenance, because they will inevitably encounter issues that are a culmination of architecture decisions. Information is needed to deal with said problems.

Unforeseen security problems may come to light during the use of the SuC. This means that the systems have to be patched. Also, if the systems are updated and improved as part of the maintenance process, security must be observed. Such processes are out-of-scope for this project, but no less important.

**Transferring knowledge.** Before handing over the project to the maintenance team, it has to made sure that documentation is organised. It costs a lot of work to do this, but it is much more expensive to have maintenance staff reverse-engineering the system to understand the system.

Clearly, it is too late now to start documenting components and architecture decisions from scratch now. The task here is mostly to gather all knowledge in one place, and to make a complete, coherent, and consistent documentation. It is probably a good idea to go over this documentation with the intended users, to make sure understand this structure. Make sure that people are available for answering questions even for a period after deployment.

## 9.5 Parallel processes

Besides the regular intra-phasal activities, as described before, there are some relevant inter- and extra-phasal activities that must be considered, i.e. things to be done in between phases and things that must be done outside (in parallel) of regular phasal-structure.

Between phases (*inter*-phasal), the multi-phase development life-cycle model (see ch. 8.2.2) prescribes a go/no-go decision, where the project is evaluated. The state of security is a part of this as well.

Outside of regular operations (*extra*-phasal), there are multiple things that need their own parallel processes:

- *Incident management* is by far the most important for the project itself.
- *Organisational Learning*, such that future projects, CSMS and process architecture versions can be improved. This is a management process, and should be specified as such by management specialists, specific to the organisation.
- *Observing Security*, as a process to assimilate a complete view on security throughout the organisation to inform higher-level decision-makers such as senior management. This includes distribution of aggregate knowledge to other projects that are running at the same time. This is an organisational process, conducted by high-level security specialists such as ISO's and the CISO. This is also out-of-scope for this project architecture research.

### 9.5.1 Periodic and Inter-Phasal Review

Between each phase of the development life-cycle there is an inter-phasal review, led by an information security officer. For security this process is about checking and verifying earlier work, to (in)validate current views about the quality of security in the SuC.

It is recommended to involve external parties such as an organisational security team, or external organisations for a brief audit of security processes and documentation. Business impact assessments, both on architectural and component level, should be reassessed and, if necessary, adjusted.

The goal of this review is not to do a comprehensive security evaluation, i.e. it should not go into too much details. After all, such details are likely to change anyway, especially if this review is conducted at an early stage. The review is mostly oriented on verifying whether policies have been brought into practice effectively, and whether security products and documentation are plausibly sufficient.

Based on the results of this review, a security push (See Howard & Lipner [35, Ch. 13]) may be necessary before continuation to the next phase, to bring security to acceptable levels. A course-adjustment may also be used to redirect security efforts in the next phase, to steer towards a higher level of security. Examples of such adjustments are:

- Changing emphasis from one activity to another.
- Revising education or awareness strategies.
- Escalating security concerns to a higher level of management in case of mis-management within the project.
- Adjusting the size of the project's security team.

### 9.5.2 Toughness and Incident Management

As described in chapters 7.3.2, 9.2.6 and 9.3, toughness and resilience is an indispensable part of security. Part of being resilient against attacks as organisation (and project), is knowing how to act when something unexpected occurs: being prepared to do incident management.

Incident management is mostly an organisational process, described by the CSMS and dealt with by security and crisis-management specialists through preparation and ad-hoc response processes. It is thus mostly out-of-scope from a project-level perspective.

Nonetheless, it is important to realise that incident response must be prepared to act at any moment. Sometimes, responses are needed then and there, as the incident is observed. Other times, response can wait until the next day. For R&D projects, this may mean that OT specialists must be on stand-by during experiments or tests, in case of unexpected security breaches.

In any case, make sure to read up on incident management and plan ahead as an organisation (e.g. start reading with [81, 23, 26, 28, 29, 41, 30, 19]).

# Chapter 10

# People

In this chapter, *purpose* (ch. 8) and *process* (ch. 9) come together. *People* give purpose and execute the process. People work together in social processes to achieve common goals. The social constructs needed to regulate this, are not self-evident; rather they are complicated and often vague. This chapter discusses what is needed with regard to people to make the SOTDLC a success: hierarchies, interaction, and social constructs organise the collaborative process of making a system secure.

Based on the SOTDLC process, three core social prerequisites were selected. These factors must be adequately present to make the SOTDLC a success:

1. *Clear decision structure*. Someone must have ownership of and responsibility for decisions. This person must have sufficient knowledge and mandate to make such decisions, while simultaneously there must be remedial processes to prevent escalation.

2. *Ownership and Association of advisors*. People who are not making decisions (advisors, specialists, managers) should still feel like they can make an impact on the final product. Division of labour ought not lead to professional dissociation (see 10.2).

3. *Communication*. It is important that people understand each other, for them to be able to accurately discuss issues. This may seem obvious, but in practice people from different backgrounds have different ways of understanding and articulating a problem and this may lead to communicative difficulty.

If these aspects are not present, problems will occur that will greatly diminish the effectiveness and efficiency of security, and even of the project as a whole. Such problems, and potential solutions are discussed in the sections hereafter.

## 10.1 Organisational Hierarchy

A clear decisional and ownership structure can be achieved through clear organisational hierarchy. Care should be taken to ensure that people within the same hierarchy do not have conflicting interests, e.g. a manager observing and evaluating the performance of policy. If there are structured hierarchies for core issues, people can easily find who they need (ownership or advise), hold someone responsible, or attach new structures. There are multiple hierarchies that must be specified to deal with operational decisions, here scoped for relevancy in project-level analysis:

1. *Policy/Performance*. Policy specifying the direction of the organisation based on vision and strategy. Understanding the business model. There is a reverse Performance hierarchy that observes and evaluates the performance with regard to the policy.

2. *Advisory/Oversight*. The advisory part describing who selects and prioritises knowledge that should be used for decisions. This hierarchy gives ownership over knowledge and advise to specialists and advisors. The reverse of this is Oversight, where these specialists keep an eye on the overall state of affairs in their speciality and speak up if action is needed.

3. *Decisional* consisting of operational people that actually decide what happens at a granular level to achieve the goals described by policy. They use the information fed through the advisory hierarchy.

4. *Remedial* hierarchy is the feedback architecture used to learn from decisions, and which is used to de-escalate if necessary. This hierarchy also gives room for *Big-Red-Buttons* (BRB's), as discussed in section 10.2. Action in the remedial realm is generally taken based on the *Performance* and *Oversight* information flows.

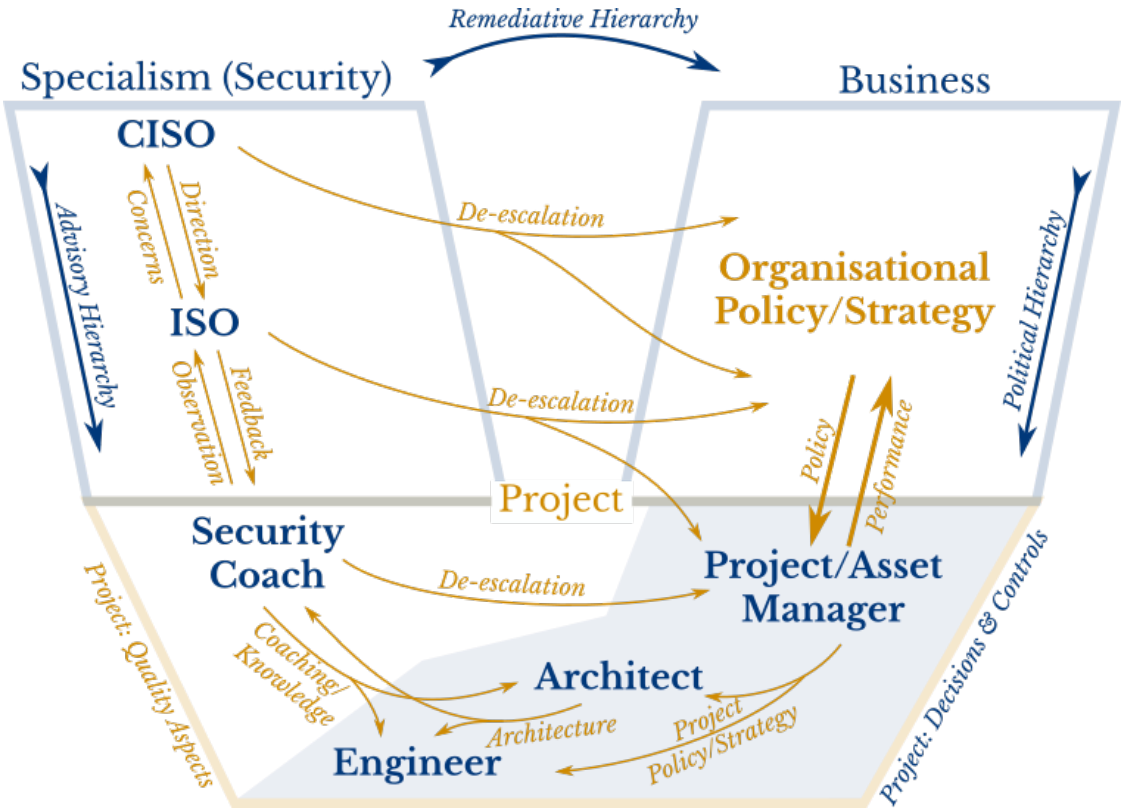


Figure 10.1: Project level hierarchy for decisional support structure in OT Security. The four hierarchies are displayed: Policy/Performance (right column); Advisory/Oversight (left column); Decisional (blue shaded area); and Remedial (horizontal, left to right). The bottom trapezium shows the project; the left column the quality speciality (security); and the right column symbolises the business side of projects (all the way up to CEO).

The hierarchies are interwoven with each other. Nonetheless, it should be clear who has which responsibility in each regard. Ownership and responsibility about specific tasks in these hierarchies should be synchronous (i.e. in the same person or place). If these are not in the same place, this will lead to ignorance and dissociation. As a simple example, if ownership and responsibility are not in the same place: when a project-manager keeps deciding against security controls, and a security specialist is consequently reprimanded for the poor performance of security, this will likely upset the security specialist, who is unable to change the situation due to a lack of mandate. The security specialist will then (at best) stop caring (dissociation) about that project and focus on other projects which they are able to positively influence (ignorance of the initial project).

As seen in figure 10.1 in the blue shaded area, the decisional hierarchy is concentrated in engineers and architects: they try to adhere to the project-strategy (and by extension the organisational strategy) and get specialist knowledge from the advisory hierarchy.

Bounded rationality [90] and capability lead to conflicts between opposing interests en route towards a common objective. In other words, the engineering problems that need to be solved are so complex that no one person can fully comprehend the problem. And even if this was possible, it would be impossible to make a rational decision based on the information because the relative weight of interests is subjective.

Bounded rationality and capability challenge centralised decision making. Organisations must nonetheless make decisions, so they have to make do with the means they have, i.e. optimising decision processes, at the risk of escalation. The next two sub-sections will discuss why the centralised strategy depicted in figure 10.1 is a suitable and optimal solution to these problems. This approach to hierarchy within a project does emphasize issues surrounding dissociation and communication. These are addressed in sections 10.2 and 10.3.

### 10.1.1 Bounded Rationality

A problem with centralising decisions, is that taking into account and optimising all factors is a difficult, time consuming, and costly process. A 'good enough' solution is frankly good enough [91], and people are not inclined to seek more information than they think they need [90]. Unfortunately, this is also prone to mistakes: good enough may turn out to be not good enough and lead to escalation.

People who make decisions are performing a complex task. They need to:

1. Understand the current condition the SuC is in.
2. Use policy to produce a view on what the SuD should behave like, while keeping CS in mind, to achieve business objectives.
3. Make design decisions components such that they function effectively, while limiting the risk that anything significant goes wrong (quality: legal, safety, security, etc.).

Designing a system is complex: it is difficult to prioritise objectives and actions[91]. When looking at the variables in the design process, there are some factors which we do understand, and which we can optimise into a strategic social process (see figure 10.2):

1. Having adequate models of the current state of the SuC, and agreeing on their meaning. This agreement is fundamental: ultimately everyone wants the company and the product to succeed, but different people may have different interpretations of success. Only if people agree on what is reality now, can they start looking towards future scenario's.

*The only way to reach agreement about the 'now' is through communication and stakeholder participation.*

2. Using this common understanding, discuss what should change to make it better: what should the SuD become. Therefore, policy is made based on discussion, mutual understanding and common agreement on business objectives in relation to present reality in technical terms. Scenario's incorporating uncertain and risky known variables can be used to make these policy objectives more robust.

3. Deciding and designing technical objectives is then the primary process (e.g. through multiple criteria analysis [91]), where engineers convert concepts into practical reality. These are then augmented into concrete requirements based on more detailed understanding of risk and uncertainty: i.e. solving risk-equations for qualitative aspects of the SuC, among which is security. This is where the design cycle with feedback is actually occurring.

This process is similar to how organisational competency is converted into strategy ("link competencies to aspirations") [92]. Eden and Ackermann also describe participative methods to better understand organisations. Here the objective is also to look at the current situation the organisation is in; identify things that the organisation is good at, and not so good at (competencies); and produce a business strategy based on that.

While it could be argued that businesses are much more abstract than OT systems; OT systems and associated risks are sufficiently complex that abstract models are needed - requiring agreement and prioritisation. Furthermore, project-level decisions are steered by organisational policy, which are based on macro-level competencies; in the end, this is based on what the organisation is able to do on a micro-level (project) as well, through a feedback cycle. Understanding the relationships between organisational policy, project policy, and project decisions, and defining them in similar fashions, is thus relevant.

In essence, the above description (see also figure 10.2) is the same as what was described in chapter 9 process. The difference here, is that group processes are emphasized: team-work. Complexity is addressed by focusing on a shared understanding of the now. People may have different interpretations or visions on what the future should be like, but at least the starting point should be the same. In group discussions about the now, team-members can challenge each others' understanding of the system, current approaches, and overall consistency.

Furthermore, a common model of the SuD, gives quality specialists a foundation to base their advise on (e.g. BIA, 6.2). This prevents inefficiency. As discussed further in the next section, engineers (who make decisions) have a limited capability of understanding all details of all relevant specialities. The knowledge hierarchy helps with that, since specialist advisors therein can use the shared understanding of the system to select and curate their advise to the needs of the deciding engineers.

### 10.1.2 Bounded Capability

Part of bounded rationality, is that the deciding person can only grasp a limited model of reality. Even if highly precise and complex models exist of the system, its subsystems, all of their risks, and all their futures, it is still too much to combine this into one whole. Even if it were possible to make one such super-model, a decision is yet to be made after understanding and thinking about it all. In practice, this means that the decision will be an optimisation of an

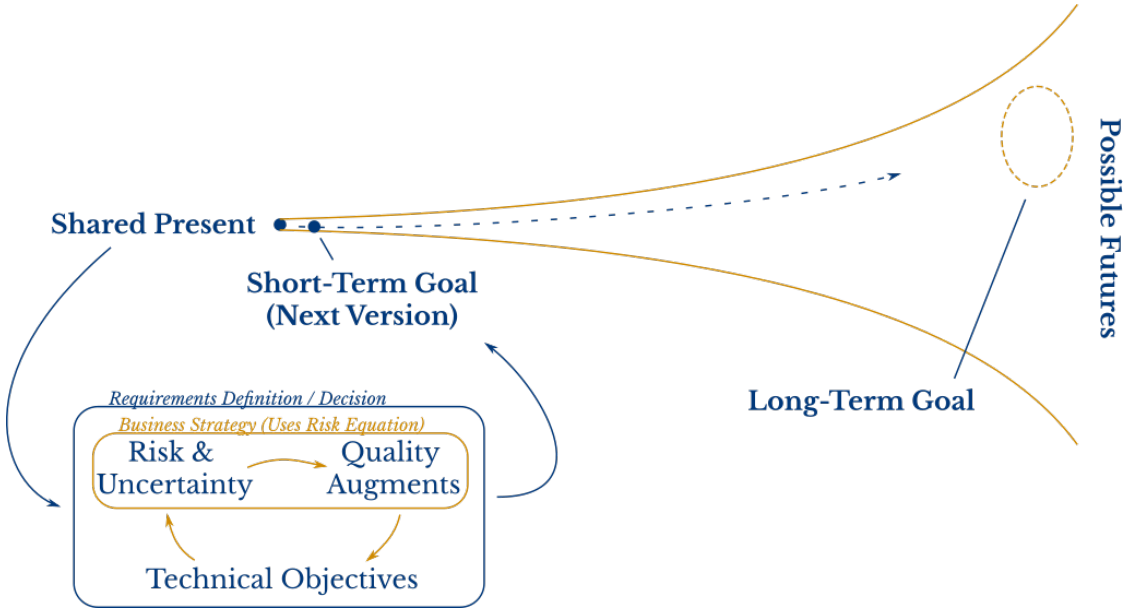


Figure 10.2: Decision process in teamwork. Knowing the present to model the future: optimising robustness against possible scenario's, and choosing the next step based on quality augments that address risk and uncertainty. Here the next step is clear, because of the relatively clear and limiting boundaries given by the diverging scenario's.

interpretation of the model: not a holistic consideration, leading to the best possible solution. Humans simply have limits to the amount of information they can compare and process.

**Information streams** One approach to solve this, is clarifying objectives, as discussed earlier. This in itself is not enough. The other side to the solution, is to simplify and structure relevant knowledge: to include only the most relevant information. Taking control over the positioning of the *bounds* of rationality and capability. This means optimising what a person knows and can do, specifically for the decisions they have to make. A pre-processing step is thus required, in which available information is assessed and compiled by supporting specialists (e.g. security coach). This process is facilitated in the knowledge structure of the project, where advisory specialists:

1. Look at the SuC from their own specialty (solving risk equation);
2. Select the most important risks (prioritisation); and
3. Supplement these risks with potential solutions (cost v. impact and important considerations).

Absolutely paramount here, is that specialists only select and serve the most crucial information for the deciding engineer. The idea is to limit information as much as possible, such that decisions are based only on the most relevant information from many specialties. Sometimes, this means that some interests way heavier than others, and that these others are subsequently (temporarily) ignored. Other times, this would ideally even mean that a specialist expresses no significant concerns, leaving room for other risks to be addressed.

If particular risks are ignored too often, concerns among related specialists may start rising, and the project may eventually escalate beyond acceptable bounds. Not only does this threaten the success of the project, but it is also demotivating and frustrating to everyone involved. This problem is discussed in the next section on ownership and interest, 10.2.

## 10.2 Ownership and Interest

The complexity of modern OT and IT systems makes specialisation indispensable. An important criticism of the hierarchies proposed in the previous section is that people could loose their sense of ownership over and interest in their work.

Similar negative consequences of specialisation were already discussed in the 18th century by Adam Smith [93], in the context of factory labour. Division of labour itself is much older, discussed even by Ancient Greek philosophers. Ferguson, Marx, Durkheim and others also discussed these concepts at length and warn for the possible social and psychological consequences of division of labour for workers [94].

Even today the consequences of specialisation are still an active topic of discussion [95]. It is, therefore, important to discuss what the expected effects of the proposed hierarchies are, and how they facilitate individual ownership and interest.

### 10.2.1 Professional Dissociation

The most important problem with respect to the SOTDLC and the proposed project hierarchy, is what we will here call *professional dissociation caused by specialisation*, meaning that people mentally loose a connection with their work, because they (feel like they) cannot see or influence outcomes or the bigger picture anymore: they do not feel ownership over their work.

An example of a cause for this is repeated *bounded capability* (see sect. 10.1.2), where decision makers seriously undervalue certain interests. An advisor specialising in this subject may dissociate with the consequent issues and the underlying system, because they feel like they are not listened to; that the current approach is reckless; or that important issues are ignored.

Given the policy and decisional hierarchies of the previously introduced system, this problem of dissociation deepens further.

### 10.2.2 Strategies to deal with professional dissociation

The idea is that dissociation is predominantly caused by the perceived lack of influence. This feeling is compounded by occurrence of actual escalation with regard to business-goals. The solution is to either address perception, or increase actual influence. Therefore, one could accentuate flows, direction and effect of influence (i.e. the remedial hierarchy); or institute ways to guarantee de-escalation or the reduction of perception of escalation. A core assumption her is that within one organisation, people may seem to have conflicting interests, but in reality they have conflicting interpretations of the same fundamental interests. Hence, the goal is to address this underlying problem by developing a common understanding of proper implementation of business objectives.

### Big Red Button

To address important risks that are left untreated, balancing of the purely hierarchical decisional system is required: a way to push a 'big red button' (BRB), to stop the escalation. Designing and implementing such a BRB is up to projects and organisations themselves, because this depends largely on organisational culture and project-specific priorities. Whatever the BRB is in practice, it should adhere to the remedial hierarchy described in section 10.1, and is a de-escalation process (see 6.3)

An example BRB could be a mandatory de-escalation meeting where responsible decision makers, project or asset management, and high-level advisors (e.g. ISO) must be present to discuss current developments (recent decisions) with respect to the project- and organisational-strategy, and business objectives. Ultimately, everyone should agree on what is in the best interest of the organisation, based on factual (quantitative) analysis.

A more radical strategy would be to implement the BRB as a veto (or voting) power for advisors: if they disagree with decisions they can use this to prevent business-escalation.

### Calibrating common understanding of the SuC

Another way of addressing the dissociation can be looking at how the relevant stakeholders perceive (interpret) system properties and objectives. Where the BRB is about adjusting the course of the project to reduce escalation, looking at perception is about taking a closer look to determine whether escalation is actually occurring. From the perspective of a single specialty, escalation may seem obvious, whereas from a business oriented overview, the decisions are perfectly reasonable.

To keep everyone happy and satisfied, it is important to go through the shared understanding of the current system and to refine or reinforce important objectives and policies. Having such a 'maintenance' meeting should prevent actual escalation, and take away concerns in the bigger picture. Dealing with situations where no consensus can be reached, is the beautiful magic of managers, and out of scope of this research.

The core idea here is to clearly communicate the reasoning behind decisions and strategy, and to give people the opportunity to question this.

## 10.3 Communication

Ultimately, solving the risk equation; making decisions; implementing controls; and dealing with actual threats is a collaborative effort between technicians and business. Each party has their own responsibility and ownership of a piece of this puzzle. This also means that parties ought to and must seek to understand each other, and to translate challenges and needs to a common language.

*Difficult communication*, is where people with different specialities do not speak the same 'language' i.e. they mean different things with the same words. An engineer, for example, will consider boundaries (e.g. maximum risk-acceptance) as concrete mathematical facts that must not be exceeded. A manager who is focussed on making political decisions is more likely to consider such boundaries as flexible lines, that are subject to interpretation and framing. If these two groups of people have to coordinate and agree with each other, quite a bit of translation work must be done.

In reality, hard mathematical boundaries do not always reflect available choices, while on the other hand political choices also need concrete justification. Nonetheless, there must be

some place where this middle ground is facilitated. The hierarchy diagrams help with this: they show where the different hierarchies meet, and thus where these should interface with each other: i.e. the places where mutual understanding is needed. By using these hierarchies and clearly describing them, the focus points for managers (i.e. the people who guide these social processes) are more clearly visible.

### 10.3.1 Advisory Hierarchy and Communication

There is a difference between what the security coach and security specialists must understand (in-depth considerations such as value models and adversarial models) and the level of understanding required for implementing effective security (engineers, supported by security coach in regular education, awareness and intervision sessions). Advisors seem to speak the same language, but decision makers need tools to precisely understand the meaning of policy and associated knowledge/advise to properly weigh decisions.

#### Conflict of Interest

When management judges the performance of their own decisions/policy, a conflict of interest may arise. A manager is considered a specialist in social processes and policy: of those, policy (politics) could conflict with making rational decisions based on available information in a particular (specialised) aspect of a project. For example ignoring security in a project to speed it up: this may make it seem like the project is performing very well; but on the larger scale of the organisation, risk may rise to unacceptable levels. Such a contradiction (mismanagement) is dangerous to the organisation as a whole, and thus such effects should be minimised by having independent specialists comment on the performance of the project in that particular perspective.

#### Safe haven

Within a knowledge support structure, the same language is spoken about a topic. The various layers of this hierarchy will use appropriate levels of abstraction to discuss the topic, but people are unlikely to have conflicts of interest. In cases of escalation, or other forms of conflicting interests between groups (e.g. legal and human factors disagree; or safety and security), the decision maker is the neutral factor that decides. The knowledge hierarchy facilitates a place where people can find peers with similar backgrounds to test and refine ideas with.

A down side of such a safe haven is isolation between different specialities: they might develop an atmosphere of rivalry.

## 10.4 People in Practice

Bringing these social techniques into practice is a challenge. Internal to the project, there is a lot of control over how things are done. But a project has very limited control over the structure of the organisation as a whole. Furthermore, the project may sometimes have to work with external customers or suppliers, who use different methods or are not as organised or proficient.

### 10.4.1 Stakeholders

In terms of change management to a different organisational structure that facilitates development, responsibility, and good decision making, internal stakeholders can be influenced more easily than connected stakeholders (e.g. customers and suppliers), or even external stakeholders (government). Although out of scope of this report, it is noted that having a radically different organisational structure than others in the immediate context, may cause practical difficulties. Going with small steps is recommended.

### 10.4.2 Attaching the Project Hierarchies to Organisational Context

Organisational integration is difficult. What if the organisation uses a dualistic system (management and advisors)? Somewhere, this must interface with a project structure where advise, policy and decisions are separated. Where do responsibilities and ownership align between the larger organisation and a project.

**Managers** , for example, take decisions, and it is the generally accepted way of doing things. But ultimately managers specialise in guiding social processes and providing policy based on business interests: they are generally not equipped to take practical (technical) decisions about implementation of this policy and the implications of such decisions. Even worse is taking decisions while also judging the performance of said decisions themselves: it is like a fox guarding the henhouse. In general, it would be better to prevent such conflict of interest: e.g. by making the engineer responsible, who only cares about staying within as many boundaries as possible, including policy boundaries. In practice, such a story is a hard sale: organisations are not willing to make such a drastic change to their project organisation, because they have always done it this way.

**Hybrid approach and Convergence of hierarchies** Nonetheless, it might be worth the effort to try a hybrid approach, where the organisation at large is traditionally organised, while individual projects use a more bottom up decision approach.

Ultimately, all responsibility and ownership must converge to the CEO and the board. To minimise the changes to conservative organisational structures, it is recommended to limit the more complex hierarchical system to projects which are more dynamical. Projects benefit from it, because they need a flexible and forgiving system to deal with short-term uncertainty and supply of policy and knowledge to decision makers.

This means that the first hierarchical layer outside the project, is where the other hierarchies converge into a more limited set of people. Because of the level of abstraction associated with that level in the hierarchy, such convergence is not a problem; i.e. it causes a manageable degree of conflicting interests.

### 10.4.3 Customers and Suppliers

Dealing with customers and suppliers is difficult to generalise: it is difficult to say what is go or no-go. Such stakeholders are very industry specific.

In any case, communication is key: make sure that there is at least a practical common understanding of relevant systems, responsibility and ownership. On the other hand, specifying everything in legal documents is ineffective (except in legal context); therefore choose efficient and clear models. Strive for open and clear communication in terms of risks and

quality. Of course, there will be competition with competitors: good luck with that, there is no silver bullet (yet).

# Part II

# Verification of Theory

# Chapter 11

# Framework Validation - ATO-project Case-study

The theoretical framework is quite new in several ways: this means it has not been proven to work in practice. This is risky, since organisations need to be able to rely on such frameworks to assure sufficient levels of security. As a step in this direction, the case study of the Dutch Railways and its Automatic Train Operation project tries to establish some credibility to the claims made within this thesis. If there are fundamental shortcomings, this group of specialists is interested to find them such that they can improve their own processes, and better achieve their goals.

## 11.1 Goal

The goal of these interviews is twofold:

1. Gain understanding of wants and needs of people involved in security.

2. Test, challenge, and verify the theory in the theoretical framework.

Essentially, we are asking specialists "What do you think?" and "Does this make sense?" in a systematic way.

This research was conducted in the context of the Dutch Railways (NS), with the organisational goal to produce effective policy for dealing with security in OT R&D projects. Before some theoretical framework is introduced as policy, it first needs support from within the organisation. The case-study NS is thus not only an exercise in proving validity of the framework draft, but it also probes the feasibility of actually implementing it as policy.

The goal of this part of the research is explicitly not to compare the quality of this theoretical approach to other approaches, or to give a rigorous substantiation to claims. Our practical approach is aimed at quickly using available insights to produce a better framework than what has been available so far, and is known to be inadequate: e.g. in the 62443 series of standards [9]. As such, it should be seen as exploratory research. Researchers and industry specialists are cordially invited to criticise and improve this framework to better substantiate (or disprove) it scientifically!

### 11.1.1 Outcomes

The overall outcome that is hoped for, is a semi-formal, multidisciplinary insight into OT security, specifically for R&D projects. It would be great if this research as a whole sparks discussion in two areas:

1. First, the aim is that the framework contributes to cybersecurity at the NS, possibly even by being accepted into policy itself. The case-study is hoped to lead to internal discussion and evaluation of current practices and opportunities for improvement.

2. Secondly, industry specialists may find value in the ideas proposed in this thesis. By asking various specialists to scrutinise the ideas, the framework is strengthened. With some luck, this will lead to evolution in standardisation and OT security industry practices.

## 11.2 Overview

Given that we want to verify the legitimacy of the framework, input from other people is needed. Due to the nature and scope of this research, these people were found within the NS, and notably (about half) in relation to the ATO-project case-study.

There are a few general ways of generating feedback from these people (participative approaches):

- *Questionnaire* A more quantitative approach with comparable input. The advantage here is potentially more response overall, which is more easily compared. A big disadvantage is that questionnaires lack depth and abstraction.

- *Group discussion* A qualitative approach where people spend time together discussing a subject along certain themes. Advantage is that this can stimulate creative thinking among group members. Disadvantages are that this takes a long time for everyone involved, and that people could easily end up off-topic or in a fundamental disagreement. (see also [92])

- *Interviewing one on one* Another qualitative approach where the interviewer and interviewee have a conversation about the topic. The advantage here is that the interviewer can more easily control the conversation and direct it towards interesting topics, but this is simultaneously a disadvantage due to bias.

- *Written request for comments* is a more informal way of asking for feedback. The advantage here is detailed on-topic in-depth feedback, but the disadvantage is that this document is much too long, so participants will most likely not read it.

Based on this consideration, interviews were chosen as method. The rest of this section will expand challenges, pitfalls, and practical considerations for the interviews.

### 11.2.1 Challenges and Pitfalls

#### Diverse Backgrounds/Specialities

To rule out speciality based bias, most respondents will come from outside the security 'bubble'. This makes it more likely that strange customs in the security world are identified and

criticised. The consequence of this is that most respondents will thus also find it hard to relate to security concepts. All interviewee's will, however, have at least some experience with quality aspects in general. Therefore, this quality-based perspective will be used to discuss the abstract concepts related to security, using the individual expertise of interviewees as basis. This angle helps the respondents to come to the strongest and most consistent version of an opinion that is possible within the available time.

To further challenge the respondents, they will be asked to not describe their current experience with the subject matter, but to instead describe their opinion about an idealised version of those processes. This approach has two advantages:

1. By discussing an improved version of current practice, this forces the respondent to reflect on the performance and shortcomings of the current system; and implicitly they must thus consider their own experience and opinion about the practical functionality of the methods. The idea is that this makes it easier to later discuss performance of their system and relationship to context such as team and organisation.

2. By discussing an ideal situation, the interviewee and interviewer can more effectively work together towards the strongest possible version of that particular approach. This is great for the next interview, where this approach will be used to challenge the ideas at the basis of the SOTDLC:

    - if the interviewee designed an approach that is largely similar to the SOTDLC, that is great, because it shows support for the ideas and concepts within it.
    - if the interviewee designed a different approach, but can be convinced that the SOTDLC is a better way of doing things, that is also good, because then the ideas in the SOTDLC are sufficiently convincing to overcome strong counter-arguments, and it is thus more likely to garner wide support when put into practice.
    - if the interviewee designed a different approach, and is able to successfully defend a position of approximate equality or superiority against the SOTDLC, this shows that the SOTDLC can be improved. Because the criticism is based on a concrete and experience based model, it is likely to also give concrete direction to what can be improved in the SOTDLC. Improvement to the SOTDLC is, of course, also a great outcome.

**Terminology** The way in which different groups communicate about certain processes or issues can lead to misunderstandings. For example, many engineers who deal with tight tolerances in design have a different definition for the words 'precise', 'exact', or 'maximum' than e.g. business consultants who deal with deadlines, people, and policy.

In a multidisciplinary field such as security, this could lead to misunderstandings. In the interviews this should be taken into account when interviewing from different backgrounds. One of the ways to deal with this is to summarise the interviewees words in other words, to verify if the meaning came across. Another strategy is extrapolation: "if you say $A$, and I know $A \rightarrow B$, then you also mean $B$, is that correct?". If it turns out $\neg B$, then the interviewee was likely misunderstood (assuming $A \rightarrow B$ is valid).

### Confidentiality of Case Specific Details

Discussions of opinions of the interviewee are most likely based on their experience within this organisation, or others. This may mean that they use concrete examples to substantiate their point, leading to confidentiality issues. This makes it difficult to completely publish the interviews, and thus to verify their findings. Two strategies were used to minimise related effects:

1. Focus on scientific basis over case-specific details. Meaning that scientific papers or non-confidential information are used as prompts for interview sections. This will nudge interviewee's towards using this as a basis for their arguments, while still being neutral.

2. The interviews are summarised, and this summary is checked by the interviewee to make sure that it accurately reflects their opinion. This way, interviewer induced biases in summaries are minimised. The interviewee is asked to make sure that the emphasis in the summaries is also accurate to their opinion.

### Prior Knowledge of Interviewees

Most of the interviewees are people that work in or near the same department as the interviewer. Coffee-machine talks have sometimes also discussed work on this research, and thus some people may not come into the interview completely neutral. Where this was the case this was noted in the interview summary.

Prior knowledge (no more than coffee-machine depth) is considered only a minor issue for the first interview: it may influence the opinion of the interviewee, but overall the goal of this verification is to understand interviewees opinion about the framework and challenge the framework: for that, knowledge about the framework is necessary anyway. Furthermore, the interviewees will be given a general introduction to the topic at the start of the interview anyway.

### Bias

There is a risk of bias in this interviewing approach.

Firstly, there is a risk of confirmation bias through framing the topic and the questions. The effect of this is minimised by having two interviews, where the first one is about the opinion of the interviewee only. This way, the interviewee is able to express their personal opinion about the topic without having been influenced by ideas about the framework yet. It is hoped that this will also stimulate interviewees to be more critical of the proposed framework in the second interview. Nonetheless, confirmation bias is still a substantial risk which must be taken into account.

Secondly, the population used for this research is likely fairly uniform, which reduces the reliability of the results. Because all interviewees are employees of the NS, or closely related to NS, they are quite likely to have a similar experience with organisational processes such as security. Therefore, a certain degree of similarity in ideas and critiques is to be expected. It would have been better to interview people from multiple organisations in different OT sectors, but unfortunately this was not possible with the means and time available.

### 11.2.2 Validation

To increase the reliability and verifiability of these interviews, some validation measures were taken:

- *Interview Structure.* By using two separate interviews, where the first one introduces very little about the topic, the interviewee is able to express their opinion free of influences from the research project. The interviews are structured like this to first find an unbiased opinion about the topic, and then use this in a structured way to jumpstart critical review of the framework. By letting the interviewees think about it themselves first, it is tried to reduce their agreeableness towards the legitimacy of proposed frameworks.
- *Recording.* The interviews are audio-recorded. Furthermore, a standard summary-form is filled out during the interview to make a draft overview of important topics or points. After the interview the recordings are used to make a more detailed summary of the interview.
- *Verified Summaries.* The interviewee is asked to read the summary and to reflect on how accurately it represents their views. They are also asked to grade how important they think each of these points are in their overall opinion.
- *Sample Size.* Due to the qualitative nature of this research methodology, only a relatively small number of people is interviewed. To make up for this, a varied group (many different professions) of people is chosen, such that as many different perspectives are included as possible.

### 11.2.3 Interview Plans

Based on all the considerations, two interview plans were made. These plans can be found in appendix D.

## 11.3 Results

Verified summaries of all the interviews can be found in appendix E. These results have been compared and compiled into the general outcomes discussed in this section.

### 11.3.1 Important subjects

At the end of each interview, the respondents would be asked what they think are the three most important topics that were discussed in the interview, or related to it. Table 11.1 shows the topics that respondents picked at the end of their first interview. Table 11.2 shows the same for the second interview.

It should be noted that the grouping of certain priorities into larger categories is subjective: this means that these tables should be read by an appropriate level of scepticism. These tables are used mostly to show general direction in the subjects that were important to people.

| resp. № | Formal v. Informal | Proportionality | Risk Analysis | Governance | IT vs. OT | Support & Oversight | Others |
|---|---|---|---|---|---|---|---|
| 1 | | | 3 | | | 5 | (4) Everyone is involved, but sufficient knowledge is missing among many |
| 2 | 3 | | 4 | 3 | | | |
| 3 | | | 4 | 7 | | | (Governance: (3) Governance + (4) Responsibility) |
| 4 | | | | 3 | | | (5) Maturity; (5) Managers with personal interests |
| 5 | | 4 | | 7 | | | (Governance: (3) Decision making + (4) mandate) |
| 6 | | | | | | | (-) Involvement and determination; (-) collaboration & discussion leading to a socially safe work environment, improvement, and quality; (-) The NS is in the middle of society and has an important role. |
| 7 | | | 3 | 5 | | | (3) Definition of the term 'Innovation' |
| 8 | | | | | | | (5) Flexibility in long-running projects; (4) up-to-date overview and insight into threats |
| 9 | 4 | | 2 | 5 | | | |
| 10 | | | 5 | | 2 | | |
| 11 | | 3 | | 4 | 4 | | |
| 12 | | | | 3 | | | |
| 13 | | | 5 | | | 4 | (5) Test v. goal |
| 14 | | | | 3 | 2 | 5 | (4) Maturity |
| 15 | | | | | | | |
| 16 | | | | | | | (5) Everyone is involved in security; (5) phased project approach; (4) incident response |
| 17 | 4 | 3 | | 5 | | | |
| 18 | | | | | | | (5) Getting the basics in order! |
| Total: | 11 | 10 | 26 | 45 | 8 | 14 | 54 |

Table 11.1: Interview 1: Subjects that respondents selected as 3 most important subjects.

Interviewees scored the importance of their three subjects from 1 to 5 (Noteworthy to Crucial). The subjects that were mentioned the most have a specific column, whereas less frequent subjects are grouped under 'Others'. Note that this 'importance' score does not signify whether the respondent had a positive or negative opinion about the subject, nor the specificity of the original topic-description (topics were subjectively grouped).

| resp. № | Formal v. Informal | Proportionality | Risk Analysis | Governance | Phased Approach | Others |
|---|---|---|---|---|---|---|
| 1 | | *5* | | | | (3) Relationship between Architects & Engineers |
| 2 | | | 2 | | | (4) Resiliency; (1) ATO should be considered as independent sub-projects |
| 3 | | | | | | |
| 4 | | | 3 | | | (4) Mitigation is not always prevention; (2) information sharing |
| *5* | | | 4 | 4 | 4 | |
| 6 | | | | | | - No second interview - |
| 7 | | *5* | 4 | | 5 | |
| 8 | | | | *5* | | (4) Threat intelligence; (3) sharing information for analysis |
| 9 | | | *5* | | 4 | (5) Make value-flow in an organisation visible |
| 10 | | | 4 | | 4 | (2) 3 different types of innovative projects |
| 11 | | | | *5* | 5 | (5) Relationship between technology and its context |
| 12 | | | 3 | | | (4) Test v. Goal solution; (4) IT vs. OT, what is the difference? |
| 13 | | *5* | 2 | | 4 | |
| 14 | | | | | 4 | (5) Reporting back: knowledge & Oversight; (5) Make re-definition of roles and responsibilities explicit in phase-transitions |
| 15 | *5* | 3 | | | 4 | |
| 16 | | | | | 5 | (3) Difference between really developing a product and abstract innovation |
| 17 | *5* | | | | 2 | (4) Everyone is involved in security |
| 18 | | | | | | - No second interview - |
| Total: | 10 | 18 | 27 | 14 | 41 | 58 |

Table 11.2: Interview 2: Subjects that respondents selected as 3 most important subjects.

Interviewees scored the importance of their three subjects from 1 to 5 (Noteworthy to Crucial). The subjects that were mentioned the most have a specific column, whereas less frequent subjects are grouped under 'Others'. Note that this 'importance' score does not signify whether the respondent had a positive or negative opinion about the subject, nor the specificity of the original topic-description (topics were subjectively grouped).

# Chapter 12

# Discussion

## 12.1 Overview

The results described in tables 11.1 and 11.2, show important themes in the interviews: some topics really stand out from the tables. This section gives a quick overview of what was discussed.

### Interview 1

In the first interview (table 11.1), many respondents expressed a need for a clear and structured approach to cybersecurity: this is reflected by the *Formal v. Informal*, *Governance*, and *Support & Oversight* columns. These topics relate to the abstract concepts described in chapter 10 *People*.

Many respondents are interested in the best strategies to assess risk (column *Risk Analysis*), and to compare risks with other aspects to make proportional decisions on risk-acceptance and security requirements (column *Proportionality*). This is to be expected, given that this is what cybersecurity is all about.

**Formal v. Informal** concerns horizontal relationships between team-members. Some respondents noted that it is very important to have healthy informal relationships within projects, because these are where most of the work happens, and where decisions are made quickly and effectively. These processes need the right balance between bureaucracy (formal, generally slow but accurate) and discussions (informal, often subjective yet fast).

**Governance** Describing the organisational structure surrounding the security process. This regards responsibility as well as ownership for parts of the process.

This topic also relates to observed conflicts of interest as discussed further in the discussion, sect 12.4.1. And also addresses the issue of risk abstraction (part of oversight), as is explained further in sect. 12.4.2.

**Support and Oversight** This is about the concept of a Friendly Neighbourhood Police Officer, which was brought up by respondent 1. This neighborhood police officer is someone that is approachable, who gives direction, tips, help and stimulates good behaviour. Someone who observes how everything is going, calls out minor mischief, and knows how to call in support when bigger problems arise.

The idea is that there is someone independent who checks that the project is doing well with regard to cyber-security. This fits into the 'three lines of defense' model, where the 2nd line should verify the quality of the people in the 1st line of defense.

**Everyone is involved** Although just a few noted this as an important topic, a majority of the respondents expressed an opinion along these lines: everyone has a responsibility in security. This means that everyone should have the right knowledge and access to tools, information and specialists to do their part. This expression *everyone is involved* is a clear and easy to explain the concepts from chapter 10 *People*, specifically the issues around *Ownership and Interest* and *Communication* (sections 10.2 and 10.3).

Although only briefly mentioned, this issue seems to sum up all that is important about the social process: people need to work together on security. This is done by communicating clearly about goals, proportionality, controls, and so on; by having clear distinction of who does what; working together both formally and informally; and much more.

### Interview 2

The second interview (table 11.2) shows a clear shift away from the *governance* column, towards *phased approach*. Many respondents liked the ideas about a phased approach for security, with clear security-goals in each phase. People seem to accept the SOTDLC framework as a suitable solution for the previously reported problems with governance. Some respondents were already starting to think of concrete ways to implement the SOTDLC; others were giving tips on what parts should be emphasized (Column *Others*).

**Risk Analysis** There were mixed feelings about using criminological insights to improve risk analysis methods. Most people agree that risk-analysis methods leave substantial room for improvement. Nonetheless, roughly half of the respondents think that current methods are sufficiently accurate.

In general, people are positive about looking into criminology, but improving risk-assessment methods is not considered a high priority, as shown by some of the lower scores in the *Risk Analysis* column.

Part of the *Risk Analysis* column was also the prioritisation of risks (e.g. in risk matrices; related to risk abstraction 12.4.2). Respondents who are specialists inside a project seem to report that this is not going well, whereas respondents outside of projects and management seem to report that this is going well.

**Incident Management** In the SOTDLC, it is assumed that incident management in innovation is similar to non-innovative projects, and that this is thus a straightforward implementation of existing frameworks and standards. Some respondents noted that many colleagues are not aware of the importance of incident management, and this concerns them. Respondents report that security officers have recently started an improvement initiative on incident response.

Given that this is apparently not common knowledge, this is an interesting find.

## 12.2 Definitions

Some respondents noted that they found some of the terminology confusing. Therefore, it is important to clarify these terms in the theoretical substantiation of the framework:

- *Escalation* In this thesis we use escalation to relate to the risks to business goals getting out of control, thus escalating out of acceptable bounds. In business context escalation is often used to describe a process of asking higher levels of management to give an opinion or decision about controversial or high-risk matters.
- *Innovation* There are different kinds of innovation: innovation for the business; innovation in the industry; or innovation for the whole world. Also making something completely new; vs. an addition to an existing system; vs. updating an existing system.
- *Architecture* To people involved in the operational technology in the projects, architecture denotes the structure of the system that they are working on. To people who work at a more abstract level, architecture is also something more abstract, sometimes denoting the relative structure of systems within the organisation; other times architecture would denote the structure of the organisation itself, or even structure of a process within the organisation.
- Risk = *chance* * impact.

  Chance defined by people with a background in...

  - *Operational cybersecurity* is often implicitly defined as the difficulty of exploiting a vulnerability or attack path.
  - *Organisational cybersecurity* use both difficulty of exploitation, as well as expert-opinion about qualitative threat intelligence.
  - *Safety* seem to use statistical incidence of an event based on previous experience
  - *Other areas* often seem to assume that security specialists can substantiate their claims about 'chance' with quantitative motivation, and thus seem to overestimate the reliability of claims.

## 12.3 Proportionality and Auditability

If security is over-done, this may be a disadvantage to other aspects. Although this trade-off is often legitimate, given business interests, some respondents report frustration about a lack of explanation about such choices.

### 12.3.1 Conflicts between aspects

When asked about conflicts between security and other aspects such as usability or maintainability, most respondents would answer that they could not recall of any. Conversely, these same people note that security often comes at the cost of other aspects: more security often means a less usable or maintainable system. Apparently, people do observe a conflict between security and other aspects, but they are not consciously aware of any reconciliation process that is felt as 'conflict'.

At a project-level, I think security is often blindly favoured over other aspects that are implicitly deemed less important than it, such as usability. This means that such decisions are not explained or weighed. In such a case, there is no conflict, because there has been no discussion about other possibilities. Naturally, a lack of substantiation to (thus implicit) priority decisions has a high risk of leading to disproportionality.

#### Security-By-Design or Design-By-Security?

Security-conscious people generally agree that security-by-design is the way to go: it is the most cheapest and easiest approach. Security-by-design is interpreted by some as adding as much security as possible at the earliest possible stage of a project. The interviews suggest that this leads to a phenomenon that could be described as *design-by-security*, where security is implemented over everything else without thinking twice, mandating other design decisions. This is, however, inefficient, especially in OT innovation, because the objectives of the project and its systems change over time. The security targets thus also change over time (i.e. table 9.1). Security should evolve along with the project, implementing (and sometimes removing) measures in a manner appropriate to the context.

There was one security specialist who went as far as to say that most goal-oriented security requirements should be enforced early on in the project, despite yielding little benefits, because this would make it easier to enforce the requirements in projects where it would yield benefits. Doing this design-by-security would thus only be for image. Such an approach clearly leads to frustration among stakeholders: quite a few other respondents explicitly said that they feel needlessly inhibited by security, and that they would love to work together towards effective and efficient solutions, but to do so they report a need for clear, concise and comprehensive explanations for security choices. The conclusion can only be that design-by-security leads to a loss of trust and a loss of willingness from stakeholders to collaborate.

### 12.3.2 Comparing Security to Other Aspects: Risk Matrices

Opinions differ on whether security should be considered fundamentally different due to feedback mechanisms from controls to risk. There are roughly two ways of looking at this issue:

1. The attacker changes their behaviour when controls are implemented, and thus addressing one risk, may negatively impact another. This is different from e.g. safety, because there risks are independent.

2. The total amount of 'pressure' on the system (caused by potential attackers) is roughly constant: addressing a risk through controls could reduce this pressure or address the effect this pressure has locally. This is similar to e.g. safety, where sometimes the cause of risk can be taken away, and other times the potential impact is reduced.

A non-deterministic feedback mechanism (1) would mean that risk matrices are useless in their current form. Consequently, the risk matrices would be inadequate models to compare strategies/controls or base decisions on.

I think the first interpretation is logical in context of this case-study, because the ATO-project is mostly focussed on an impact-reduction based approach in order to gain flexibility in the R&D process. This means that the attacker is not clearly explained in risk appreciations (matrices). It is therefore not strange that people would consider the attacker black-boxes.

This thesis is more aligned with the second explanation: through the risk equation, relationships between attackers, impact and controls can be better explained. Security controls can be directed towards the attacker, or towards reducing impact of an attack.

Such nuance is important for making security choices understandable and auditable. While it might not directly change the level of security itself, it does help with awareness and quality of security implementation.

## 12.4 Conflict of Interest & Risk Abstraction

### 12.4.1 Project level: Security Engineer

One of the topics discussed often was the potential conflict of interest concentrated on a security engineer if they have both responsibility towards project management (responsible for project policy/performance) and towards the Information Security Officer (ISO, responsible for knowledge and oversight).

In the second interview, the suggestion is made by the interviewer to make a more clear distinction between these performance- and oversight-responsibilities. It is difficult for one person to have both, without an internal conflict. This would mean that the security engineer becomes a security coach, who only has a knowledge and oversight responsibility. The performance-responsibility is kept in the project-manager, architect, engineer line.

Respondents mostly agree that the described potential conflict of interest is sub-optimal. Opinions differ on how important this phenomenon is in practice and whether a solution is thus needed.

### 12.4.2 Organisational level: Risk Abstraction

Organisations do not only make risk assessments at a project level: at an organisational level the overall risk must not exceed the risk appetite. This means that information about risks on project, program and department level must flow up to higher management and the CISO. Taking into account that higher management has to assess not only security, but also many other aspects, it is not possible for these people to weigh all individual risks. Therefore, a process which we will call *Risk Abstraction* has to happen, where risks from different aspects (e.g. security, safety, usability, legal, etc.) are put into comparable risk matrices and risk registers. Figure 12.1 depicts the risk-abstraction step as the transition from project to cluster. In the risk-equation (e.g. figure 7.1), risk-abstraction can be partially found in *classification* and *prioritisation* leading to controls. The dimensions and scales of these matrices are predefined based on the organisational risk-appetite, and make comparisons across domains a lot easier to comprehend.

Risk-abstraction is part of a complex social process. Although it is very necessary for understanding the bigger picture of an organisation, it muddles the waters for others who work on very specific issues. Abstracted information is used to determine strategy and to make decisions on departmental and organisation-wide levels: this has a lot of impact on engineers, architects, project-managers and others who have to implement it all. A poorly performing risk-abstraction process quickly leads to professional dissociation (ch. 10.2.1), leaving people unwilling to contribute to security.

#### Pitfalls

During the interviews, it has become clear that the current approach to this process of risk abstraction has some inherent pitfalls, which mostly people higher in the organisation do not seem to notice. Risks are interpreted and put into an appropriate risk matrix at a project level. This interpretation is for a large part based on expert opinion: given the limitations of the dimensions of the risk matrix, it is difficult to include the nuances of this opinion. After risk classification, management becomes risk-owner, and risk-abstraction is further performed through increasing levels of management.

There are two pitfalls here:

1. *Management* does not have the expertise to understand and appreciate the risks they own. Furthermore, there may be a conflict of interest between the oversight and performance that management wants. The risk management department (incl. CISO), who keeps oversight over the total risk profile, mainly receives their information through this management barrier.

2. *The quality* of the abstraction can be improved, because it is based on informal expert-opinion and discussion, and the rationale for the outcomes is scarcely or not written down. This means that it is difficult to trace causes; trace decisions; have a uniform framework; or learn from past mistakes.

These problems lead to figure 12.1 which depicts the level of chaos/uncertainty introduced at any given level of the abstraction process.

The process of sharing information at departmental level seems to go very well: all respondents who are part of this process report that they are happy. The problem is, however, that the experts at project level complain about the appropriateness of the risk-matrices, and the lack of substantiation for decisions and strategy. Ultimately the conclusion is that the quality of the documentation of risk-classification and -abstraction is low, meaning that quality of the organisational and departmental risk assessments is also inherently low (garbage in, garbage out), despite robust procedures at departmental level. This problem is compounded because management is unable to oversee and understand technical nuances, especially without documentation; and because outcomes (targets, requirements, controls, etc.) of the risk-assessment process cannot be traced.

### 12.4.3 Recommendations

Given that people at project level are reporting a lack of explanation for security claims and risk appreciation anyway, improving in this regard would also help with improving the quality of organisational risk assessment.

A further recommendation would be to do one or both of two things:

1. *Parallel information structure.* Improve oversight by using a parallel information structure to verify and intervene if the project is under-performing on security.

2. *Improve transparency.* Tying in to the concept of *'everyone is involved'*, here the project-team is better able to track information to decision-making, allowing a self-governing structure of internal oversight, and making external oversight much easier as well.

In both cases, at a project level, the security engineer becomes a *coach* who is not directly responsible for delivering an advice or requirements for security, but who has a only a knowledge and oversight responsibility. The security coach helps engineers to assess risks and make appropriate requirements, while keeping an eye on the performance with respect to security targets.

#### Parallel information structure

Based on the hierarchies proposed in chapter 10.1, the idea is to let security specialists at project level directly report to their ISO, who then reports to the CISO. This hierarchy is shown in the left branch of figure 10.1. At every level of the organisation, the (security) specialist can then communicate in a horizontal and downward relationship with the relevant management:

- Security coach - project/program management
- ISO - department management & project/program management
- CISO - C-level management & department management

This way, the knowledge/oversight process is completely decoupled from the performance/policy process. The two main advantages are:

1. There is less potential for conflicting interests;
2. Malice and incompetence are less likely to cause problems; and
3. Specialists are better able to distinguish what is and is not important to relay to the next level of abstraction.

### Improved Transparency

A simpler solution is to clearly document each processing step from the threat-analysis through risk-abstraction until the definition of security requirements. This is based on the concept *'everyone is involved'* from the second interview. Everyone should be able to understand why security is an important part of the system, and why it designed in this or that particular way. If security is in the way of other interests, people can trace requirements back to the original risks it was designed for: if possible and necessary security can be changed to facilitate other needs, while still covering the original risks.

Advantages of this approach are that it:

1. Improves the ability to do security-by-design (instead of design-by-security);
2. Gives technical staff the information they need to get involved in security discussions;
3. Better prepares the organisation for later re-design/patching, because design-rationale is better documented; and
4. Better prepares the organisation for auditing and compliance, when legislation for OT-security is implemented/enforced.

## 12.5 Social Process

In the 'people' part of the framework, an important point is that the managers with the most responsibility for security, have the least knowledge and understanding of it. This lead to the initial conclusion that it might be better to concentrate decision-powers lower in the organisation, e.g. at the architect or engineers, because these would have a much better technical understanding of the consequences of certain approaches. In practice these people play an important role in decisions already, and managers mostly follow their advise. Many respondents found the current approach much better, because according to them decisions and business interests should be in the same place (person). Nonetheless, many people reported that they would like a more direct link with specialists, and that they favoured a strong informal network over indirect communication through a manager with centralised responsibility.

This leads to the conclusion that the initial approach of the framework was too much, but that it was in the right direction: information about security should flow freely on a horizontal

level in a project. This also means that security specialists must be available to explain choices and to help people work with security. The rest of this section discusses concrete ways this can be done better, which were brought up during the interviews.

### 12.5.1 Neighbourhood Police Officer

The idea that everyone is involved in security is widely emphasized by respondents. This means that security cannot only be the responsibility of specialists, but others must contribute as well to make it a success. Some (non-security) respondents report that they find it difficult to do their part due to one or multiple of the following reasons:

- They do not know where to report security concerns about system design or implementations;
- It is unclear where they could get information; and
- Security problems are not often obvious, and non-specialists may overlook potential problems without knowing it.

The issues above are most apparent in smaller projects, e.g. small updates or revisions to existing systems.

Respondent 1 reported that he would love to see some kind of 'friendly neighbourhood police officer' ("wijkagent" in Dutch). This person would be a visible person who helps through providing knowledge, but who also keeps an eye out for potential problems. They would also have the connections to the security department for calling in backup if needed. Multiple respondents mentioned the importance of a socially-safe work-environment, with high regard for learning, improvement, and quality: it seems like the neighbourhood police officer captures this idea.

This 'neighbourhood police officer' is the Information Security Officer for smaller projects. Clearly, the security department is not visible enough, especially for smaller projects. I think this is caused by the following two phenomena:

1. There is a shortage of security personnel, and thus smaller projects with a relatively low security risk profile and projects with a low priority and are ignored.
2. The security department is not visible and approachable enough: it should be clearer who is the primary contact for information and feedback. The security department should also be able to keep an eye on what (small) projects are going on, to keep an eye on potential risks.

For larger R&D projects, as targeted by this framework, it is assumed that a 'security coach' is directly involved in the project, who should address these problems directly as a primary responsibility. The security coach is thus this neighbourhood police officer. They keep up positive informal relations with the project, while making sure to select and present relevant technical knowledge from the rest of the organisation or security industry. They also keep track of the status of security and guide the process along. Furthermore, the security coach knows who to ask for technical help when needed (e.g. specialists or organisational facilities such as SOC, pen-testers, legal, privacy). Finally, the security coach will be the first to notice if a project is not doing well on security: they will call in the troops, e.g. (C)ISO, higher management or even others if things get out of hand.

### 12.5.2 Three Lines of Defense model

There are different views on the exact application of this model in the NS. The most discussion is about the role of the departmental information security officer (ISO): whether it is 1st or 2nd line. Part of this, is that there is currently insufficient qualified personnel available to keep a strict separation. Nonetheless, it is good to discuss what would be the ideal approach.

Earlier in this chapter (12.4) conflicts of interest were discussed. If there is no clear distinction between who is part of the 1st line and who is part of the 2nd line, it will also be unclear what loyalties a particular person should have. This will lead to conflicts of interest in the same person, regardless of whether they are part of the same knowledge hierarchy or not, including ISO's.

Given the recommendations about involving everyone in security, and that projects should have a security coach to help them, the following structure seems logical (from a purely security perspective):

1. *Security Coach:* is the first line. They advise the project, keep an eye on the state of affairs and report back to the ISO. One security coach can be involved in multiple projects for efficiency reasons, but they are explicitly not responsible for making more abstract risk assessments on a departmental level. The security coach works together with enterprise-architects, solution architects, engineers and business to understand the issues specific to that project, and gives independent advice about security, which they do not implement themselves. The security coach is thus moreso a security solution architect.
   *Enterprise- & solution architect for security* is also first line. They translate business needs to suitable architectures for departmental and project levels. They verify the understanding and usage of these architectures as well, but the ISO is ultimately responsible for looking at the bigger picture and determining whether enough was done for that department or project.

2. *ISO:* is part of the second line. The ISO keeps track of all projects in a department or cluster. They check the work of security coaches, and make sure that projects are on the right track with regard to higher level objectives. The ISO may advise the security coach about particular questions, but is not involved in the project itself. If escalation occurs, the ISO can coordinate with the product-owner, project lead or even with the cluster- or department-head that the ISO is assigned to. In principle, it is not the ISO's task to make the business happy at lower levels of the organisation, but moreso to seek alignment with higher-level security targets by verifying the validity of low-level decisions.
   *CISO (office):* ISO's can have different levels of seniority, depending on how much of the organisation is their responsibility. This means that ISO's together form a hierarchy, all the way up to the CISO who has the overview over the whole organisation. THis is all still 2nd line of defense

3. *Audit:* Audits are done by an independent department or organisation. This is the third line. Most of this is out-of-scope for this research.

It is crucial to give people in the 1st line and 2nd line different function titles, even if they do similar work. Giving them the same title leads to a lot of confusion. The function title thus shows others what the primary responsibility of a person is: achieving business goals effectively by helping with and advising about security (1st line); or verifying the validity

of an implementation and critically assessing risk-acceptance (2nd line). Again, these two responsibilities *do not go together*, without conflict of interest.

For example, Local Information Security Officer (LISO) would thus be a misnomer for a security-coach-like function, both because this person is not an officer (rather operational, 1st line staff), and because this person should be an advisor and not making engineering decisions.

Under this model, it practical to make the 2nd line hierarchically part of organisational risk-management structure. The 1st line is then part of the relevant department. Despite the hierarchical decoupling of the 1st and 2nd line, ties between them should be short with a healthy dose of informality. It is important that 1st and 2nd line know how to find each other to ask questions, discuss issues, share information, and provide help. There should be high amounts of mutual trust, but also the opportunity to informally and formally question decisions (in both directions) if this is needed.

## 12.6 Project Model: Phases

### 12.6.1 Dynamic Phases

The framework should be flexible with regard to the amount of phases, and their goal. Especially innovative projects are unpredictable in their size and amount of phases. Therefore, the security process should be able to scale along with a project.

### 12.6.2 Reliability Phase

An interesting note by one of the respondents (14): there is another phase after transitioning to maintenance, that could still be seen as part of the learning curve. It is the time immediately after roll-out, where learning about common malfunctions and the character of the system is the goal.

Thinking in terms of security, a relevant learning-objective for in this phase is developing understanding of new types of attack-vectors that were previously unforeseen. This is mostly related to cyber-criminology in broad sense [71], i.e. where the system is misused as a tool to achieve a goal that is not directly related to the system itself. An example of this is an attacker using a train to trap someone, e.g. to intimidate or hinder them, by keeping the doors of the train closed. A practical example of this is regimes who intimidate opponents or journalists who live abroad to silence their critiques.

Such a perspective on the security process is very relevant, because new attack strategies are a logical consequence of developing new OT systems. Given the intricate understanding of the SuC, it is comparatively unlikely that an attacker will find a very surprising attack vector in it. Attacks where the SuC is used as a tool (i.e. cybercrime in broad sense [71]) are difficult to predict and understand well. It is thus interesting to add a distinct learning phase to the overview, and add security weighting to the relevant phasal overview tables (e.g. 9.3 and 9.2).

### 12.6.3 Test vs. Target

A few people expressed during the first interview that it is important to have different approaches for security for testing and for the intended target system. This stems from the different objectives of the business for an unfinished project, and the innovative goals that are

put at risk when security is too constrictive; in contrast to the value that may be lost when the system is under-protected.

This idea resonates with the proposed phase progression for security targets. Only one respondent said that they think security should be fully implemented from the outset; whereas most of the other respondents were explicitly positive about this model.

## 12.7 Incident Management

Existing incident response plans must be updated when R&D projects introduce new processes and systems. Currently, it is unclear who is responsible for making the adjustments or for bringing them into practice during development. Some people reported a desire for clarity and communication in this regard.

In the SOTDLC, the assumption was that incident-management was well-developed, and that organisations would have sufficient means to develop their own plan. Not all scenarios can be foreseen, so whatever things were overlooked should be effectively contained. Given that there is room for improvement in this case-study, it might be a good idea to develop this concept further. It seems a good idea to make incident management a prominent part of the SOTDLC process, as well as looking at specific recommendations for R&D. Furthermore, the concept of toughness could be expanded to become a clearer design principle.

## 12.8 Conclusion: Revisions to the Framework

The interviews, and above discussion have lead to new insights for the framework. Based on the results of the interviews, the following amendments are recommended for the original framework:

### 12.8.1 Hierarchies to Horizontal Relationships

Originally the focus of the hierarchies (see ch. 10.1) were focussed moreso on making sure that technical knowledge and decisions would stay in the same spot. Based on the interviews, this emphasis turned away from hierarchies and more towards informal horizontal relationships.

This means that it is not so much about who makes the final decisions, but moreso at how the team arrives at this decisions together. Everyone is involved, and the team has to work together to weigh all aspects and interests related to a decision, to make sure that no single discipline takes over.

### 12.8.2 Add Discovery Phase

An distinct discovery phase is added to the phase-diagram, where security aims to discover system misuse in the broad cybercrime definition [71].

Because OT systems have impact on their physical surroundings, attackers may start misusing these systems for attacking other targets. Such attacks are difficult to predict, so monitoring seems like a better idea. This ties into monitoring for changes in the risk landscape, as part of the maintenance process in section 12.8.3.

### 12.8.3 Include Maintenance

Security goals regarding maintainability should be added to the phase-tables. Maintenance is an important part of OT systems, especially if they have to last multiple decades. Security will have to evolve along with the rest of the system, and along with newly discovered vulnerabilities and threats (see sect. 12.8.2 as well).

In developing the security requirements for the final SuC, evolution of these requirements throughout its lifetime will have to be considered. This means that a maintenance process, along with a clear description of who is responsible for it or owner of it should be one of the deliverables of the SOTDLC.

### 12.8.4 Expand Incident Management

Especially in innovations, where systems and risks change quickly, it is difficult to keep up with security. It is difficult to guarantee prevention security-incidents. The SOTDLC already describes an approach where impact-reduction and strategic risk-acceptance are corner-stones, but incident management is still a small part. Given that incident management is a crucial part of good security in innovations, its importance should be more clearly expressed.

It may also be worth it to specifically study incident-management in R&D projects, in further research.

### 12.8.5 Explaining the Social Process

The interviews have been very insightful about how the framework can be explained in simpler language to make it easier to understand and implement. Chapter 10 *People* in particular was written in a very abstract style. Concrete concepts such as the 'neighbourhood police officer' (wijkagent); 'everyone is involved'; and 'risk abstraction' have worked very well already in better explaining and discussing the framework, its concepts and the findings of the interviews in the NS.

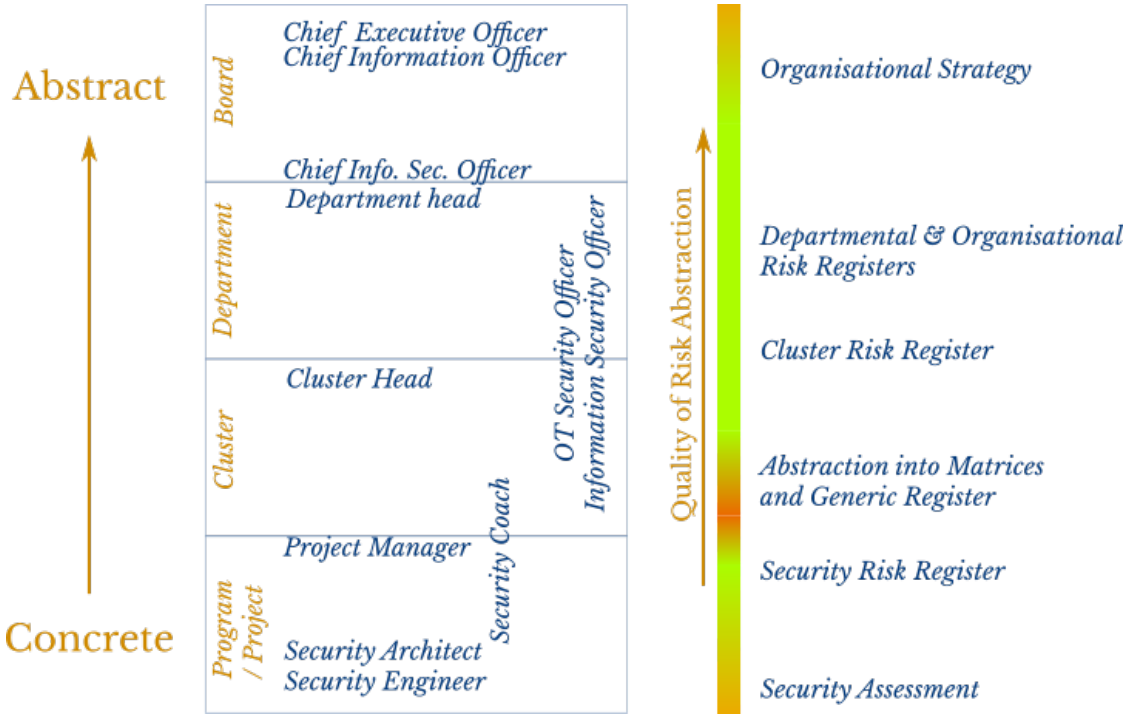


Figure 12.1: Quality of Risk Processing Throughout the Organisation. The color of the gradient-bar denotes the reported degree to which it is possible to trace decisions through this part of the process in the NS, based on the interviews.

For example, decisions based on data high in the organisation, e.g. at department-head level, can be clearly followed and understood by cluster-heads and the board but less so by a security architect. The architect cannot 'see' through the orange part, because the abstracted form cannot be traced back to the underlying risk: clear documentation is lacking in the orange part.
The bottom is yellow because here the risks have not been formalised into a risk register yet. Therefore, it is still difficult to trace, and any (implementation) decisions are made on an informal level.
The gradient is also yellow at the top, simply because there is no layer above there: any abstractions made by the board are for themselves, and are made by generalists without particular expertise in cybersecurity. The strict abstraction-quality is probably somewhat lower, but this does not matter because it is goal- and audience-oriented.

The effect of this is that decisions made from cluster level and up are difficult to concretely motivate to people at program/project level. This is undesirable, because it could lead to frustration and ultimately professional dissociation.

# Part III

# Way Forward

# Chapter 13

# Recommendations for NS

This chapter gives a brief overview of recommendations for NS based on the case study. These recommendations are broadly supported by the respondents of the interviews.

## 13.1 Redacted

In this version the recommendations for NS have been redacted to protect the interests of Nederlandse Spoorwegen.

If you work at NS, contact Klaasjan Ooms-Geugies to request the version for internal distribution: that version does contain this chapter.

# Chapter 14

# Implementing SOTDLC Into a Project

This chapter gives a first draft of a more compressed version of the framework, which could be used for developing (future) policy.

## 14.1 Does this framework apply to my project? - Decision Tree

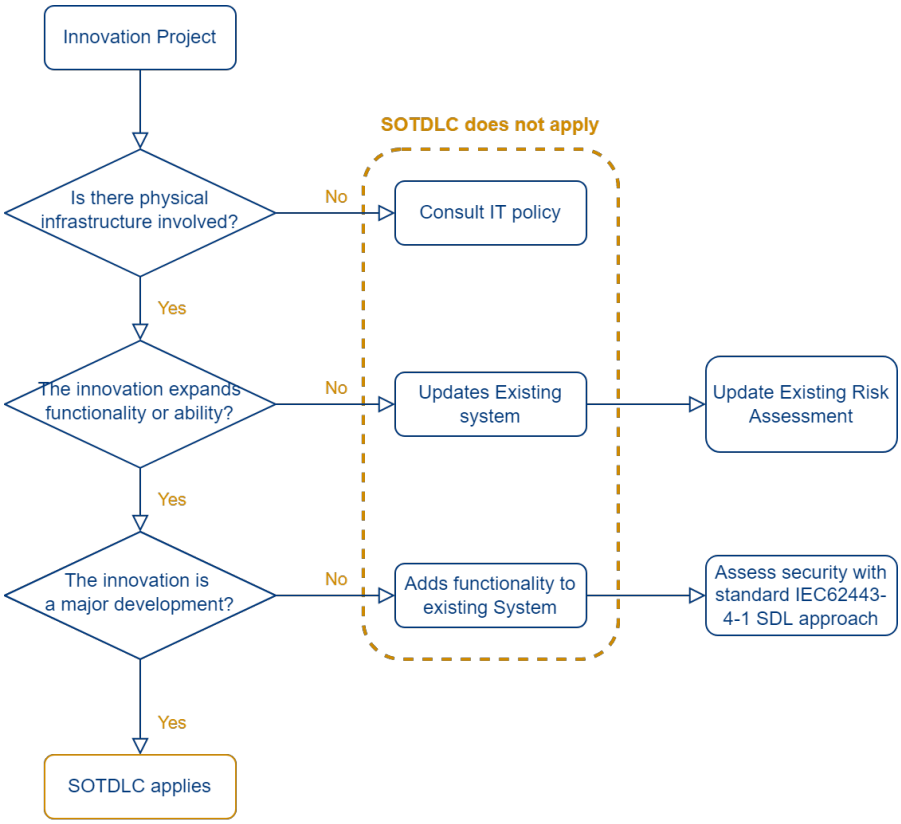


Figure 14.1: Decision tree: does this framework apply to my project?

This framework is directed at OT innovation projects that really develop new technology.

For more information, see chapter 8. Such projects are typically larger projects, spanning over more than a single year, and with multiple distinct phases of maturity.

Figure 14.1 shows the basic questions you have to ask to decide what kind of approach to take.

## 14.2 Overview of the Cybersecurity Process

The Secure Operational Technology Development LifeCycle (SOTDLC), aims to guide the security process in relation to an OT innovation project. The framework is based on, and goes hand-in-hand with the IEC62443 standard [9]. The process is explained through two tools:

**The risk Equation** Describes technical processes for assessing and accepting risk. In the end this leads to risk classification and requirements. This segment is *focussed on security specialists* working with or in innovative projects.


Figure 14.2: Risk Equation

The core procedures in the equation are in chronological order:

1. *Risk-acceptance analysis*, determining what kind of impact is and is not acceptable given business objectives and vision.

2. *Threat assessment*, where an inventory is made of the risks (risk-register) based on attacker-properties and the SuD.

3. *Requirements Definition*. After risk-abstraction (i.e. classification and prioritisation), unacceptable risks are mitigated through (requirements for) security controls, in order of priority.

**Process Model** Describing which steps must be taken throughout the project to keep up with security. Organisational needs are the basis for this, such that security is proportional and effective. The process also gives guidance on social interaction and responsibilities in the team. This segment is *focussed on management* working with or in innovative projects.

### 14.2.1 Timeline

The project is defined in chronological order in this policy; starting with formation of a team and roles, and ending with deployment.

| Funnel (Ch. 8.2.2) | Phase | Security (SOTDLC) |
|---|---|---|
| Concept Evaluation | 0 | Organisational Aspects |
| Planning & Specification | 1 | Baseline |
| Design & Experimentation | 2 | Cruising |
| Test & evaluation | 3 | |
| Pilot | 4 | Transition |
| Implementation & Roll-out | 5 | |
| Maintenance | 6 | Discovery |

Table 14.1: Timeline of the project

The main idea in this project model, is that security grows along with the goals of the project: moving from reactive security to preventive. Early in the project, innovation needs more freedom. Later, as value and risk to the business increase, security is expanded, until it reaches production-level. The nature of security also changes throughout the phases, as shown in table 14.2.

| Project phase | Focus | Dominating Value / Security Target |
|---|---|---|
| 0 | CS | Organisational facilities and operations: starting up the project and gathering existing documentation about CS |
| 1 | CS | Preventative wrt CS; Confidentiality of knowledge |
| 2 | CS | Preventative and Resilience wrt CS; Resilience wrt SuD; safety-driven |
| 3 | SuD | Resilience for whole SuC; Adding preventative measures for SuD |
| 4 | SuD | Reinforcing preventative measures for SuD to improve availability |
| 5 | SuC | Availability (reliability) for SuC |
| 6 | SuC | Discovery of new misuse or attack scenarios; monitoring; and adaptation |

Table 14.2: Value related to project phases

## 14.3 Security Phase 0: Organisational Aspects

This phase is all about getting prepared. As manager, now is the time to contact a relevant Information Security Officer (ISO), who can help you get started.

Security is cheaper and better if it is done as early as possible. This means involving security specialists from the first moments a project is started.

After involving the right people, these specialists will start to assess the impact of the new system on its context. They will ask questions about the objectives of the project, what would be potential problems and the time-line of the project.

### 14.3.1 The team

The following roles are likely to occur in a project. For each, their responsibilities and needs are listed.

**Project Manager / product owner** This person is responsible for decisions in the project. This includes cybersecurity. Business-specialists do not have sufficient knowledge to make decisions or set requirements by themselves; specialists are thus needed. Given the complexity of OT innovations, and the impact on other systems, multiple security specialists will likely be required. The coordination of this, is overseen by the security coach.

It is important to realise that everyone is involved in security: it is not just a thing for specialists. This means that everyone in the team will need to be aware of their responsibilities in this regard. It is important that the project manager is aware of the exact implications of this, and that they express this to the team as well.

Examples of things that everyone will need to do:

1. Use secure (company) devices (e.g. laptops) and software (e.g. mail)
2. Have open discussions about security: ask for explanation if security is annoying, but also notify the security coach with concerns.
3. Design systems with security in mind (security-by-design)

Business and management should be aware of problems related to risk-abstraction. Shortly put, this concerns a loss of accuracy due to the categorisation of risks, for example through a risk-matrix. Security risk matrices, and other risk-classification tools are not sufficiently mature yet: this means that very serious risks sometimes score low in the classifications, whereas low risks may sometimes be scored very high. Therefore, when making important decisions about risk acceptance or security strategy, it is also important to listen to specialist appreciations of the risk-registers.

(See ch. 12.4.2)

**Security Coach** This person is central to the cybersecurity process in the first line. The security is also called the (friendly) neighborhood police(wo)man (wijkagent in Dutch) Responsibilities include:

- Connecting the right people with the project
- Supplying technical security knowledge to team-members (e.g. engineers, architects, management)
- Observing state-of-affairs for security in the project, and reporting it back to ISO (2nd line)
- Raising awareness in the team about security, and explaining design choices wrt security
- Serving as a visible point-of-contact for team-members if they have questions, comments, or feedback
- Making Risk-assessments, and ensuring risk-registration

Every project under this policy must involve a security coach. This does not have to be full-time, depending on the size and scope of the project.

Beware conflicts of interest: if the security coach is involved in defining requirements or engineering, they will likely encounter conflicts-of-interest between having adequate security and protecting the interests of the project. Preferably, the security coach is not in the lead for such a process. Furthermore, it is advised to actively discuss this issue among each other, to prevent unwanted pressure on this person and to ensure the accuracy of security-assessments.

(See ch. 12.4 and 12.5)

**Security Engineer** Is responsible for defining technical requirements in terms of security, and ensuring proper implementation of security controls.

This role is sometimes combined with the security architect or -coach.

**Solution-Architect Security** Is responsible for a security solution-architecture (describing e.g. zoning, information flows). This includes functional requirements for security.

This role is sometimes combined with the security engineer or -coach.

**Enterprise Architect Security** Makes a reference-architecture for security that can be used for the solution-architecture. The enterprise architect works on many projects and looks at the bigger picture. They can give advice based on experience in other projects.

**Information Security Officer** This person is responsible for overseeing a portfolio of projects from the second line of defence. The ISO will make sure that all projects adhere to company-policy. They do so through documentation and discussions with the security-coach. The ISO will also feed back the state-of-affairs to upper management and senior ISO's or the CISO.

If the risks in the project become too great or if the project is no longer compliant with laws or, the ISO will also try to de-escalate this behaviour directly, through discussions with the project-manager and cluster- or department management. The ISO will make sure that organisational goals are met.

**Other specialists; Safety, Reliability, Maintenance** There will be overlap between some of the other aspects and security. Management should make a plan on who is in the lead with risks or issues that concern multiple specialisms. Especially safety is likely to require such cooperation in OT.

### 14.3.2 Pre-requisites

After building a team, the security coach needs to build a picture of the systems that already exist (*Contextual systems*, CS), and which the new system (*System under Development*, SuD) will attach to. Together the CS and SuD form the *System under Consideration* SuC. In this phase 0, the security coach will gather all the necessary documentation, and contact necessary people: this organisational preparation is required for phase 1 (Baseline), where all the knowledge is combined into a risk-assessment for the CS.

The security coach needs the following:

- Previous risk-assessments and design/decision-documentation for the CS
- Business Impact Analyses related to CS, and if available for the SuC

- Incident response plans for CS; contact details of people responsible for handling incidents
- Project structure details for the innovation project itself: which phases; project goals; team members
- Details about previous incidents and lessons learned with the CS and similar systems: threat intelligence

If any of the above are missing, the security coach will need to make them: this requires extra time and manpower ahead of the project. Management should expect and facilitate this.

### 14.3.3 Understanding Risk Acceptance

Ahead of the project, the business will need to decide what risk-acceptance strategy they will use. The security coach can give guidance to a risk-acceptance definition session based on table 14.2. It is important to start this before actual risks are discussed, because risk-acceptance-escalation is a common psychological pitfall. Business should give a clear boundary of what is, and what is not acceptable for business impact.

This process starts with taking the BIA (business impact analysis) for the CS. Business should already have a general idea of its objectives for the project. Based on this, a rudimentary value-driven BIA can be made for the project, which forms the basis of an explicit risk-acceptance boundary. This boundary shows technicians which consequences are and are not acceptable, along with documented example-cases and a short motivation.

(See chapter 6)

## 14.4 Security Phase 1: Baseline

The goal of this phase is to make sure that the CS and existing value-creation remain protected during the developments. This means that attackers cannot use insecure prototypes to attack other systems.

The second goal of this phase is to make sure that the knowledge produced by the research and development is protected. For example industrial espionage, or preparation for sabotage should be considered.

Both of these topics include concerns about third parties, e.g. suppliers, and which knowledge or system details they may see.

Finally, this phase looks ahead to experimentation in the next phase. Preparations are made to deal with these experiments. This includes incident response procedures and safety-considerations.

### 14.4.1 Understanding Risk

Here an initial risk assessment is made. All the information gathered so far is combined with a rough initial system-architecture. The risk-equation guides this process.

This risk-assessment is weighed against the risk-acceptance agreements: this leads to an initial set of security controls. These controls will at this stage focus on keeping attackers away from the CS and on incident management for the SuD.

(See chapter 5)

### 14.4.2 Incident management & Resilience

The SuD prototypes themselves have little value for the business at this point: most of the value is concentrated in the knowledge about them. This is reflected in the way security deals with the whole SuC.

To make sure that innovation can keep going, the goal is to reduce the lasting impact of successful attacks. This is partly done through easy, basic preventive measures, but mostly through strong incident management capabilities.

See chapters 7.3.3 and 9.5.2.

### 14.4.3 Transition: inter-phasal review

Before the project can start with experimentation (project phase 2), there should be an inter-phasal security review. The ISO and optionally independent specialists should be involved to check the status of the project.

- Are the right people involved?
- Are the remaining risks for the CS accepted? What are high priorities?
- Is the incident management functional and effective?
- Are security procedures and rules followed? Why (not)?

A security push may be necessary before continuation to the next phase, to bring security to acceptable levels. A course-adjustment may also be used to redirect security efforts in the next phase.

See chapter 9.5.1.

## 14.5 Security Phase 2: Cruising

This phase is about keeping the security equation up-to-date with the project. This is done through keeping an eye on architectural developments for the SuD, and by updating the business impact analyses to reflect the size and shape of prototypes.

As the prototypes grow, their value increases to the business and to targeting attackers. This means that the focus of security shifts away from the CS, which should be well-protected by now, towards the SuD itself. Previously, incident-management and resilience were used to guarantee progress in innovation. Now security moves towards building walls around the prototypes, to increase their stability and reliability.

(See ch. 9.3)

### 14.5.1 Updating Risk Assessment

The updates of the risk-assessment are primarily driven by the maturing system-architecture; the level of integration of prototype components; and the degree of coupling between the SuD and the CS.

This means that the security coach must regularly keep an eye on these things, and adjust risk-registers accordingly.

(See ch. 9.3.3)

### 14.5.2 Prevention

As the SuD is expanded and development progresses, components that were initially developed independently are combined into a larger system. This means that the possibilities for attackers expand dramatically.

Furthermore, the system will start performing more of the desired functions, meaning that it will likely communicate more with CS. This increases the risk, since the attacker will have more possible paths or tools to work with. To counter this, more preventative measures for the SuD must be taken.

(See ch. 9.3.2)

### 14.5.3 Defining Target-Architecture

Now that is is becoming clear what the final architecture of the system will be like, the security coach should start looking ahead to this target-architecture. It has to be decided what the final security architecture should look like.

Now that the final prototypes have not been fully implemented yet, changes can still relatively easily be made to increase security: this security-by-design ensures cost-effectiveness and efficiency of security efforts.

Looking at target-architecture also includes provisions for maintenance: for example key-management; documentation; or patching should be planned.

(See ch. 12.6.3)

## 14.6 Security Phase 3: Transition

During the transition phase, the SuD is fully integrated into business processes. This means that security of the SuC is brought up to its final form, both in prevention and resilience.

With respect to security, the *transition* means that the SOTDLC process is wrapped up and that all documentation; the team; and processes are made to comply with common security standards (e.g. 62443 and 2700x). After the transition, these standards are used to measure and maintain security.

(See ch. 9.4

### 14.6.1 Maintenance Processes

Since the SuD is new, it will likely have its own peculiarities in terms of security maintenance. This is the moment to implement the maintenance; key-management; and patching processes as they were defined in the previous phase.

(See ch. 9.4.2, )

### 14.6.2 Final Push

The final push is about finishing up. This includes documentation, final security debugging, and making sure that all incident-response, monitoring, and patching works.

### 14.6.3 Final Review

Before the system is ready for widespread use, there should be a comprehensive review of system-security. At least this includes an independent review by the organisational auditing

department, and preferably by an external party as well. This is also the time for extra pen-tests and documentation reviews.

The topics for review are at least *process execution; architecture threat models; Component level threat models; and left-over concerns*.

(See ch. 9.4.2 and 12.3)

## 14.7 Security Phase 4: Discovery

During the discovery phase, the main security processes and quality are managed based on common standardisation, such as the 62443 standard. There are, however, some additional concerns that stem from the innovative nature of the project.

It is important to keep an extra keen eye on the SuD, because it may be misused as a weapon for secondary goals. Furthermore, there may be new ways of attacking the system that are a consequence of its innovative nature. These effects are both more concerning than implementing an established type of system, because it is less predictable.

It takes some time to get to know the system and the way it manifests itself in the big bad world: to facilitate this process and to ensure caution, these topics are put into this explicit phase, which can last for up to a few years after pilots with mature prototypes.

(See chapters 12.6.2 and 12.8.2)

### 14.7.1 Updating Risk Assessment & Security Patching

To deal with the increased uncertainty, more emphasis should be put on re-evaluation of risks and the possibility of security patching. It is recommended to regularly schedule moments of review after deployment, to make sure that risk-assessments are accurate and that controls are sufficient.

# Chapter 15

# Conclusion

With new times come new challenges: the need for new ways to deal with complex innovations in OT have lead to the framework and policies proposed in this research. The SOTDLC gives a phased strategy that fits large innovative projects. It provides a risk-assessment method using the *risk equation*, which is suitable for use with complex OT systems in dynamic environments, even if system architecture is still incomplete. The SOTDLC seeks alignment with business objectives, and ensures proportionality of security to other aspects. Finally, it gives guidance on the social constructs needed to make the security process a success.

The current approach to modelling R&D projects in standardisation does not match what is done in practice. It does not work as desired in large innovations. Better models are available and can be used to more accurately work with stakeholders throughout the phases of a development project. The SOTDLC uses these models, making it suitable for use in multi-phase innovation projects.

Besides aligning security with the rest of the project, the social process is also very important. Simply deploying security staff and having them write threat models and recommend controls is insufficient: a holistic approach is needed, where everyone is involved. This can be achieved through clearer management using a critical systems thinking approach. The security coach is the guide of this in projects: they facilitate education, awareness and clear goals. Informal processes are important for technicians in projects: this allows them to move quickly and effectively to really innovate. At the same time the security coach can guard the quality and reliability of security solutions for prototypes.

Industry guidelines on threat-assessment are vague and non-specific. This can be improved by taking a threat-intelligence approach based in criminology. The first step towards this is supporting criminological research, and taking initiatives on sharing incident data both inside and outside the organisation. Unfortunately, high-quality criminological support for cybersecurity is not a reality yet, but the SOTDLC is future-proof by showing how it can seamlessly be fit into the process. It gives an overview of what kinds of criminological support could be useful for risk-assessments.

This thesis provides a pragmatic and holistic framework for reasoning about OT security in research and development. The resulting policy gives guidance to all stakeholders, including those without a technical background. The framework uses an extensive basis in literature, combining insights in multiple fields and numerous specialities into a single, complete framework.

Not only does the framework find support in literature, an extensive qualitative study with 18 respondents has shown that professionals from many different backgrounds support this

framework and its main ideas. The SOTDLC addresses shortcomings in current approaches, strengthening the precision and effectiveness of risk-assessment and -treatment. Moreover, it does so without disrupting established processes, principles or practices. This paves the way for implementation in organisational policy for NS and in industry.

## 15.1 Further Research

Despite the many aspects that the SOTDLC covers, there are quite a few open questions left where the framework can be improved.

### 15.1.1 External Alignment

This research has mostly looked at literature that is directly relevant for development processes or risk assessments. It has given less attention to other processes within an (OT oriented) organisation, that may be relevant.

Such alignment can, for example, be sought by making:

- secure engineering practices for OT in general;
- lists of common mistakes similar to the OWASP top 10;
- collaborations with related fields such as physical security and safety; or
- the relationship between project-level security processes and organisational policy such as the CSMS.

**CSMS** This research has taken the stance that project- and operational-level security policy should be designed with precedence over the CSMS, because this is where value is produced for the organisation.

This does not degrade the importance of the CSMS, nor the ability of high-level management to override project-level decisions. Nonetheless, the organisational policy and CSMS should be designed with the flexibility in mind, that is required for the success of projects.

It should not be forgotten that it is the operational level of the organisation whose task it is to generate value; the rest of the business enables this.

### 15.1.2 Methodical Variation

This research has tried to include perspectives from business, IT, OT, criminology, and more; but there are other angles that may yield valuable insights into how organisations should deal with SOTDLC-like processes, and organisational security structure.

An example is using applied mathematics (e.g. game theory) to make quantitative models. This would make the framework more robust. Here fields such as operations research seem relevant.

### 15.1.3 Additions

During the interviews two security-related methods were named that were not taken into account for the first framework. There may be more (proprietary) methods that are relevant; these could be added to the research to make the review more complete.

**IRAM2** is a proprietary risk assessment method that has a bigger focus on business objectives and fitting security into this. This seems like a better, more rounded approach than what has been analysed for this thesis. It is thus worth checking out and adding to the literature, but unfortunately it is not public. A starting point for looking into this could be Dehkhoda's thesis [96].

**SABSA** (Sherwood Applied Business Security Architecture [97]) is a method for developing both enterprise and solution architectures for security. In general, the framework is light on an architecture perspective. It may improve the applicability of the framework to include a review of architecture methods in general.

### 15.1.4 Incident Response

Incident response approaches specific for OT innovation contexts (based on respondent 15 and 2). These should be integrated into the SOTDLC, and adjusted to its unique objectives.

### 15.1.5 Sector-specific Application

Sector-specific (e.g. Rail-specific) literature may have a different perspective on the theoretical framework and addresses many practical matters that are left out of general standards for generalisation purposes. By comparing this literature to the findings in the theoretical framework, its applicability and feasibility could be further reinforced.

These are some examples of standards that could be used:

- ISO/IEC/IEEE 15288 Systems and software engineering — System life cycle processes [98]
- CLC/TS 50701 Railway applications - Cybersecurity [99]
- EN 50126 Railway Applications. The Specification and Demonstration of Reliability, Availability, Maintainability and Safety (RAMS) Generic RAMS Process [100]. This can be applied in security- and classifcation models, such as in Appendix A.
- Project SECRET for European Rail Traffic Management System (ERTMS) [101]
- Project Shift2Rail [102]
- Guidelines for cybersecurity in Railway (International Union of Railways) [103]
- Railway Cybersecurity - Good Practices in Cyber Risk Management (ENISA) [104]
- Zoning and Conduits for Railways (ENISA and ER-ISAC) [105]
- Cordis project, cybersecurity in the railway

### 15.1.6 Criminology

This topic has been discussed extensively in this report and in the interviews. There are many sub-fields of cyber-criminology that are of great interest to cybersecurity, especially in the dynamic and groundbreaking environment of research and development. There is a big need for cyer-criminological research of all kinds.

It is hoped that companies interested in cybersecurity will increase their investment and efforts in this regard, employing criminologists, and opening up research and internship positions directed at cyber-crimninology.

# Bibliography


[1] Barbara Filkins and Doug Wylie. *SANS 2019 State of OT/ICS Cybersecurity Survey*. Tech. rep. Accessed on Sept 29, 2021. SANS, June 2019. URL: https://industrialcyber.co/wp-content/uploads/2020/05/Survey_ICS-2019_Radiflow.pdf.

[2] Bernard Marr. *What is Industry 4.0? Here's A Super Easy Explanation For Anyone*. [Online; accessed November 1st 2021]. Forbes, Sept. 2018. URL: https://www.forbes.com/sites/bernardmarr/2018/09/02/what-is-industry-4-0-heres-a-super-easy-explanation-for-anyone/?sh=3f3cb07c9788.

[3] Unknown. *Cyber attacks on critical infrastructure*. Tech. rep. Accessed on Sept 29, 2021. Allianz. URL: https://www.agcs.allianz.com/news-and-insights/expert-risk-articles/cyber-attacks-on-critical-infrastructure.html.

[4] Unknown. "Cybersecuritybeeld Nederland 2021". In: (June 2021). URL: https://www.ncsc.nl/onderwerpen/cyber-security-beeld-nederland/documenten/publicaties/2021/juni/28/csbn-2021.

[5] Unknown. *Programma Hoogfrequent Spoorvervoer*. Tech. rep. Accessed on Sept 29, 2021. ProRail. URL: https://www.prorail.nl/programmas/programma-hoogfrequent-spoorvervoer.

[6] Unknown. *Meer en snellere treinen*. Tech. rep. Accessed on Sept 29, 2021. ProRail, Aug. 2020. URL: https://www.prorail.nl/nieuws/meer-en-snellere-treinen.

[7] Unknown. *What is OT Security*. Tech. rep. Accessed on Sept 29, 2021. Infradata. URL: https://www.infradata.nl/en/resources/what-is-ot-security/.

[8] Siegfried Hollerer, Wolfgang Kastner, and Thilo Sauter. "Towards a Threat Modeling Approach Addressing Security and Safety in OT Environments". In: *2021 17th IEEE International Conference on Factory Communication Systems (WFCS)*. IEEE, 2021, pp. 37–40. DOI: 10.1109/WFCS46889.2021.9483591.

[9] *IEC 62443 Industrial communication networks - Network and system security*. Standard. Accessed November 3d, 2021. International Electrotechnical Commission. URL: https://www.iec.ch/blog/understanding-iec-62443.

[10] "ISO/IEC/IEEE International Standard - Systems and software engineering – System life cycle processes". In: *ISO/IEC/IEEE 15288 First edition 2015-05-15* (2015), pp. 1–118. DOI: 10.1109/IEEESTD.2015.7106435.

[11] Ron Ross, Michael McEvilley, and Janet Oren. *SP 800-160 Vol. 1 Systems Security Engineering: Considerations for a Multidisciplinary Approach in the Engineering of Trustworthy Secure Systems*. Tech. rep. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Mar. 2018. DOI: https://doi.org/10.6028/NIST.SP.800-160v1. URL: https://csrc.nist.gov/publications/detail/sp/800-160/vol-1/final.

[12] Joint Task Force. *SP 800-37 Rev. 2 Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy*. Tech. rep. Version 2. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Dec. 2018. DOI: https://doi.org/10.6028/NIST.SP.800-37r2. URL: https://csrc.nist.gov/publications/detail/sp/800-37/rev-2/final.

[13] Graham Williamson. *OT, ICS, SCADA – What's the difference?* Tech. rep. Accessed on Oct 4, 2021. KuppingerCole Analysts, July 2015. URL: https://www.kuppingercole.com/blog/williamson/ot-ics-scada-whats-the-difference.

[14] David Vonk. "Security Analysis: OV-Fiets Connect". [Confidential]. Bachelor thesis. Nijmegen, Netherlands: Radboud University and Nederlandse Spoorwegen (NS), Jan. 2020.

[15] Matt Bishop. *Introduction to Computer Security*. ISBN: 0-321-24744-2. Addison-Wesley, 2003.

[16] Unknown. *Four Ways to Manage Risk*. [Online; accessed December 6th 2021]. Capstone Capital. URL: https://capstonecap.com/four-ways-to-manage-risk/.

[17] Unknown. *ATO stand van zaken paper [ATO current state of affairs paper]*. Tech. rep. [Confidential] Version 0.6; reporting on the years 2019-2021. NS, 2021.

[18] Unknown. *Realtime route-informatiesysteem wint innovatie-award NS*. [Online; accessed November 1st 2021]. 2017. URL: https://www.infosupport.com/referenties/timtim-wint-ns-innovatie-award/.

[19] Matthew Barrett. *Framework for Improving Critical Infrastructure Cybersecurity Version 1.1*. Tech. rep. Version 1.1. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Apr. 2018. DOI: https://doi.org/10.6028/NIST.CSWP.04162018. URL: https://www.nist.gov/publications/framework-improving-critical-infrastructure-cybersecurity-version-11.

[20] Michael Nieles, Kelley Dempsey, and Victoria Pillitteri. *Special Publication 800-12 Rev. 1, An Introduction to Information Security*. Tech. rep. Version 1. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, June 2017. DOI: https://doi.org/10.6028/NIST.SP.800-12r1. URL: https://csrc.nist.gov/publications/detail/sp/800-12/rev-1/final.

[21] Marianne Swanson, Joan Hash, and Pauline Bowen. *SP 800-18 Rev. 1 Guide for Developing Security Plans for Federal Information Systems*. Tech. rep. Version 1. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Feb. 2006. DOI: https://doi.org/10.6028/NIST.SP.800-18r1. URL: https://csrc.nist.gov/publications/detail/sp/800-18/rev-1/final.

[22] Joint Task Force Transformation Initiative. *SP 800-30 Rev. 1 Guide for Conducting Risk Assessments*. Tech. rep. Version 1. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Sept. 2012. DOI: https://doi.org/10.6028/NIST.SP.800-30r1. URL: https://csrc.nist.gov/publications/detail/sp/800-30/rev-1/final.

[23] Marianne Swanson et al. *SP 800-34 Rev. 1 Contingency Planning Guide for Federal Information Systems*. Tech. rep. Version 1. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Nov. 2010. DOI: https://doi.org/10.6028/NIST.SP.800-34r1. URL: https://csrc.nist.gov/publications/detail/sp/800-34/rev-1/final.

[24] Joint Task Force. *SP 800-53 Rev. 5 Security and Privacy Controls for Information Systems and Organizations*. Tech. rep. Version 5. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Dec. 2020. DOI: https://doi.org/10.6028/NIST.SP.800-53r5. URL: https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final.

[25] Kevin Stine et al. *SP 800-60 Vol. 1 Rev. 1 Guide for Mapping Types of Information and Information Systems to Security Categories*. Tech. rep. Version 1. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Aug. 2008. DOI: https://doi.org/10.6028/NIST.SP.800-60v1r1. URL: https://csrc.nist.gov/publications/detail/sp/800-60/vol-1-rev-1/final.

[26] Paul Cichonski et al. *SP 800-61 Rev. 2 Computer Security Incident Handling Guide*. Tech. rep. Version 2. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Aug. 2012. DOI: https://doi.org/10.6028/NIST.SP.800-61r2. URL: https://csrc.nist.gov/publications/detail/sp/800-61/rev-2/final.

[27] Keith Stouffer et al. *SP 800-82 Rev. 2 Guide to Industrial Control Systems (ICS) Security*. Tech. rep. Version 2. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, May 205. DOI: https://doi.org/10.6028/NIST.SP.800-82r2. URL: https://csrc.nist.gov/publications/detail/sp/800-82/rev-2/final.

[28] Karen Kent et al. *SP 800-86 Guide to Integrating Forensic Techniques into Incident Response*. Tech. rep. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Aug. 2006. DOI: https://doi.org/10.6028/NIST.SP.800-86. URL: https://csrc.nist.gov/publications/detail/sp/800-86/final.

[29] Karen Kent and Murugiah Souppaya. *SP 800-92 Guide to Computer Security Log Management*. Tech. rep. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Sept. 2006. DOI: https://doi.org/10.6028/NIST.SP.800-92. URL: https://csrc.nist.gov/publications/detail/sp/800-92/final.

[30] Pauline Bowen, Joan Hash, and Mark Wilson. *SP 800-100 Information Security Handbook: A Guide for Managers*. Tech. rep. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Mar. 2007. DOI: https://doi.org/10.6028/NIST.SP.800-100. URL: https://csrc.nist.gov/publications/detail/sp/800-100/final.

[31] *Systems Security Engineering (SSE) Project*. Standard. Accessed November 3d, 2021. National Institute of Standards and Technology. URL: https://csrc.nist.gov/projects/systems-security-engineering-project.

[32] Ron Ross et al. *SP 800-160 Vol. 2 Developing Cyber Resilient Systems: A Systems Security Engineering Approach*. Tech. rep. Accessed on Nov 3d, 2021. National Institute of Standards and Technology, Nov. 2019. DOI: https://doi.org/10.6028/NIST.SP.800-160v2. URL: https://csrc.nist.gov/publications/detail/sp/800-160/vol-2/final.

[33] Robert Gregory. "Design Science Research and the Grounded Theory Method: Characteristics, Differences, and Complementary Uses". In: Jan. 2010. ISBN: 978-3-7908-2780-4. DOI: 10.1007/978-3-7908-2781-1_6.

[34] Salvatore T. March and Gerald F. Smith. "Design and natural science research on information technology". In: *Decision Support Systems* 15.4 (1995), pp. 251–266. ISSN: 0167-9236. DOI: https://doi.org/10.1016/0167-9236(94)00041-2. URL: https://www.sciencedirect.com/science/article/pii/0167923694000412.

[35] Michael Howard and Steve Lipner. *The Security Development Lifecycle*. The url for the pdf was last accessed on Jan 27th, 2022 (end of url broken, ...Microsoft_Press_eBook_TheSecurityDevelopmentLifecycle_PDF.pdf). Microsoft Press, June 2006. ISBN: 9780735622142. DOI: 10.1007/s11623-010-0021-7. URL: https://download.microsoft.com/download/8/1/6/816C597A-5592-4867-A0A6-A0181703CD59/Microsoft_Press_eBook_TheSecurityDevelopmentLifecycle_PDF.pdf.

[36] *IEC 62443-4-1:2018 Security for industrial automation and control systems –Part 4-1: Secure product development lifecycle requirements*. Standard. International Electrotechnical Commission, Apr. 2018. URL: https://webstore.iec.ch/publication/33615.

[37] Michael Muckin and Scott C. Fitch. "A Threat-Driven Approach to Cyber Security". In: (2019). [Online; accessed December 7th 2021]. URL: https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Threat-Driven-Approach.pdf.

[38] Unknown. *The Cyber Kill Chain ®*. [Online; accessed December 7th 2021]. Lockheed Martin Corporation. URL: https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html.

[39] Rod Coombs, Andrew McMeekin, and Roger Pybus. "Toward the development of benchmarking tools for R&D project management". In: *R&D Management* 28.3 (1998), pp. 175–186. DOI: https://doi.org/10.1111/1467-9310.00094. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-9310.00094.

[40] M. C. Jackson. "The origins and nature of critical systems thinking". In: *Systems practice* 4.2 (Apr. 1991), pp. 131–149. ISSN: 1573-9295. DOI: 10.1007/BF01068246. URL: https://doi.org/10.1007/BF01068246.

[41] *ISO/IEC 27000:2018 Information technology — Security techniques — Information security management systems — Overview and vocabulary*. Standard. International Standardisation Organisation / International Electrotechnical Commission, Feb. 2018. URL: https://www.iso.org/standard/73906.html.

[42] Colin Tankard. "The security issues of the Internet of Things". In: *Computer Fraud & Security* 2015.9 (2015), pp. 11–14. ISSN: 1361-3723. DOI: https://doi.org/10.1016/S1361-3723(15)30084-1. URL: https://www.sciencedirect.com/science/article/pii/S1361372315300841.

[43] Aaron Guzman and Cedric Bassem. *IoT Security Verification Standard (ISVS)*. Tech. rep. Version SNAPSHOT December 22, 2020. Snapshot version, accessed on Nov 1st, 2021. This was not a finished version of an OWASP standard. OWASP, Dec. 2020. URL: https://github.com/OWASP/IoT-Security-Verification-Standard-ISVS.

[44] Daniel Miessler et al. *OWASP Top 10 Internet of Things 2018*. Tech. rep. Accessed on Nov 1st, 2021. OWASP, 2018. URL: https://owasp.org/www-project-internet-of-things/.

[45] Unknown. *IOT SECURITY ISSUES IN 2021: A BUSINESS PERSPECTIVE*. Tech. rep. Accessed on Nov 1st, 2021. Apr. 2021. URL: https://www.thalesgroup.com/en/markets/digital-identity-and-security/iot/magazine/internet-threats.

[46] Unknown. *Security in the Internet of Things: Lessons from the Past for the Connected Future*. Tech. rep. Accessed on Nov 1st, 2021. Aug. 2016. URL: https://events.windriver.com/wrcd01/wrcm/2016/08/WP-IoT-security-in-the-internet-of-things.pdf.

[47] Arbia Riahi et al. "A Systemic Approach for IoT Security". In: *2013 IEEE International Conference on Distributed Computing in Sensor Systems*. 2013, pp. 351–355. ISBN: 978-1-4799-0206-4. DOI: 10.1109/DCOSS.2013.78.

[48] John R McCumber. "Information systems security: a comprehensive model". In: *14TH NATIONAL COMPUTER SECURITY CONFERENCE: Information Systems Security Requirements & Practices*. Vol. 1. NATIONAL INSTITUTE OF STANDARDS and TECHNOLOGY / NATIONAL COMPUTER SECURITY CENTER. Oct. 1991, pp. 328–337. URL: https://csrc.nist.gov/csrc/media/publications/conference-paper/1991/10/01/proceedings-14th-national-computer-security-conference-1991/documents/1991-14th-ncsc-proceedings-vol-1.pdf.

[49] John McCumber. *Assessing and Managing Security Risk in IT Systems: A Structured Methodology*. 1st ed. ISBN: 9780429208409. Auerbach Publications, Aug. 2004.

[50] Teng Xu, James B. Wendt, and Miodrag Potkonjak. "Security of IoT systems: Design challenges and opportunities". In: *2014 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. 2014, pp. 417–423. DOI: 10.1109/ICCAD.2014.7001385.

[51] Ondrej Burkacky et al. *Cybersecurity in automotive: Mastering the challenge*. Tech. rep. Accessed on Nov 1st, 2021. McKinsey & Company, Inc., Mar. 2020. URL: https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/cybersecurity-in-automotive-mastering-the-challenge.

[52] Remy Spaan. "Secure updates in automotive systems". Accessed on Nov 1st, 2021. Masters thesis. Nijmegen, Netherlands: Radboud University, May 2016. URL: https://www.ru.nl/publish/pages/769526/z_remy_spaan.pdf.

[53] Unknown. *Position on automotive security*. [Online; accessed November 1st 2021]. Verband der Automobilindustrie VDA, Nov. 2017. URL: https://en.vda.de/en/topics/innovation-and-technology/data-security/automotive-security.html.

[54] *ISO 26262:2018 Road vehicles — Functional safety*. Standard. Geneva, CH: International Organization for Standardization, Dec. 2018.

[55] *ISO 21434:2021 Road vehicles — Cybersecurity engineering*. Standard. Geneva, CH: International Organization for Standardization, Aug. 2021.

[56] *EVITA: E-safety vehicle intrusion protected applications*. [Online; accessed November 1st 2021]. URL: https://www.evita-project.org/.

[57] Juan de Vicente Mohino et al. "The Application of a New Secure Software Development Life Cycle (S-SDLC) with Agile Methodologies". In: *Electronics* 8.11 (2019). ISSN: 2079-9292. DOI: 10.3390/electronics8111218. URL: https://www.mdpi.com/2079-9292/8/11/1218.

[58] Mohammed M. Olama and James Nutaro. "Secure it now or secure it later: the benefits of addressing cyber-security from the outset". In: *Cyber Sensing 2013*. Ed. by Igor V. Ternovskiy and Peter Chin. Vol. 8757. International Society for Optics and Photonics. SPIE, 2013, pp. 113–118. DOI: 10.1117/12.2015465. URL: https://doi.org/10.1117/12.2015465.

[59] Kent Beck et al. *Manifesto for Agile Software Development*. [Online; accessed November 25th, 2021]. Feb. 2001. URL: http://agilemanifesto.org/iso/en/manifesto.html.

[60] Thomas Menze. *The State of Industrial Cybersecurity 2020*. Tech. rep. Accessed on December 6th, 2021. Sept. 2020. URL: https://ics.kaspersky.com/the-state-of-industrial-cybersecurity-2020/.

[61] John Lindström et al. "Oh, no – not another policy! Oh, yes - an OT-policy!" In: *Procedia CIRP* 81 (2019). 52nd CIRP Conference on Manufacturing Systems (CMS), Ljubljana, Slovenia, June 12-14, 2019, pp. 582–587. ISSN: 2212-8271. DOI: https://doi.org/10.1016/j.procir.2019.03.159. URL: https://www.sciencedirect.com/science/article/pii/S2212827119304640.

[62] Marco Rocchetto and Nils Ole Tippenhauer. "On Attacker Models and Profiles for Cyber-Physical Systems". In: *Computer Security – ESORICS 2016*. Ed. by Ioannis Askoxylakis et al. Cham: Springer International Publishing, 2016, pp. 427–449. ISBN: 978-3-319-45741-3.

[63] Jonathan Matusitz. "Cyberterrorism: Postmodern State of Chaos". In: *Information Security Journal: A Global Perspective* 17.4 (2008), pp. 179–187. DOI: 10.1080/19393550802397033. URL: https://doi.org/10.1080/19393550802397033.

[64] David Vonk et al. "Evaluating post-modern chaotic attackers as described by Matusitz". Jan. 2022. URL: https://drive.google.com/file/d/1vEpHU3av1HljxwrUBlrkL9uyebAVMqo/view?usp=drivesdk.

[65] Aunshul Rege. "Factors Impacting Attacker Decision-Making in Power Grid Cyber Attacks". In: *Critical Infrastructure Protection VII*. Ed. by Jonathan Butts and Sujeet Shenoi. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 125–138. ISBN: 978-3-642-45330-4. URL: https://link.springer.com/chapter/10.1007/978-3-319-45741-3_22.

[66] Alexander Testa et al. "Illegal Roaming and File Manipulation on Target Computers: Assessing the Effect of Sanction Threats on System Trespassers' Online Behaviors Sanction Threats on Online Behaviors". eng. In: *Criminology and Public Policy* 16.3 (2017), pp. 687–724. URL: https://heinonline.org/HOL/PrintRequest?handle=hein.journals/crpp16%5C&collection=journals%5C&div=57%5C&id=689%5C&print=section%5C&sction=57.

[67] Rutger Leukfeldt, ed. *RESEARCH AGENDA: THE HUMAN FACTOR IN CYBERCRIME AND CYBERSECURITY*. Eleven International Publishing, 2017. ISBN: 9789462367531. URL: https://www.academia.edu/es/45580594/Individual_cybercrime_offenders_Weulen_Kranenbarg_M_van_der_Laan_A_de_Poot_C_J_Verhoeven_M_van_der_Wagen_W_and_Weijters_G_2017_Research_Agenda_The_Human_Factor_in_Cybercrime_and_Cybersecurity_Leukfeldt_R_ed_Den_Haag_Eleven_International_Publishing_p_23_32.

[68] Unknown. *What is a honeypot?* [Online; accessed May 29th, 2022]. Kaspersky. URL: https://www.kaspersky.com/resource-center/threats/what-is-a-honeypot.

[69] Thomas J. Holt. "On the Value of Honeypots to Produce Policy Recommendations Sanction Threats on Online Behaviors: Policy Essay". eng. In: *Criminology and Public Policy* 16.3 (2017), pp. 737–746. URL: https://heinonline.org/HOL/PrintRequest?handle=hein.journals/crpp16&collection=journals&div=59&id=739&print=section&sction=59.

[70] Kevin F. Steinmetz. "Ruminations on Warning Banners, Deterrence, and System Intrusion Research Sanction Threats on Online Behaviors: Policy Essay". eng. In: *Criminology and Public Policy* 16.3 (2017), pp. 725–736. URL: https://heinonline.org/HOL/PrintRequest?handle=hein.journals/crpp16&collection=journals&div=58&id=727&print=section&sction=58.

[71] Wytske van der Wagen, J. Oerlemans, and Marleen Weulen Kranenbarg, eds. *Basisboek Cybercriminaliteit: Een criminologisch overzicht voor studie en praktijk*. Dutch. Studieboeken criminologie & veiligheid. Boom Criminologie, Nov. 2020. ISBN: 9789462361478.

[72] Unknown. *Why are we likely to continue with an investment even if it would be rational to give it up? The Sunk Cost Fallacy, explained.* [Online; accessed January 19th, 2022] Some readers of early drafts of this thesis did not know the fallacy of sunken cost. So I added a link here to safe people from the effort of googling a good source themselves. The Decision Lab. URL: https://thedecisionlab.com/biases/the-sunk-cost-fallacy/#section-10.

[73] Donal Flynn et al. "De-escalating IT projects: the DMM model". In: *Commun. ACM* 52 (Oct. 2009), pp. 131–134. DOI: 10.1145/1562764.1562797.

[74] Loren Kohnfelder and Praerit Garg. *The threats to our products*. Tech. rep. [Accessed October 22nd 2021] Found through: https://en.wikipedia.org/wiki/STRIDE_(security). Microsoft Corporation, 2009. URL: https://adam.shostack.org/microsoft/The-Threats-To-Our-Products.docx.

[75] Unknown. *The STRIDE Threat Model*. [Online; accessed October 22nd 2021]. Dec. 2009. URL: https://docs.microsoft.com/en-us/previous-versions/commerce-server/ee823878(v=cs.20)?redirectedfrom=MSDN.

[76] Unknown. *Applying STRIDE*. [Online; accessed October 22nd 2021]. Dec. 2009. URL: https://docs.microsoft.com/en-us/previous-versions/commerce-server/ee798544%28v%3dcs.20%29.

[77] Sarah Rosenfield. "Labeling Mental Illness: The Effects of Received Services and Perceived Stigma on Life Satisfaction". In: *American Sociological Review* 62.4 (1997), pp. 660–672. ISSN: 00031224. URL: http://www.jstor.org/stable/2657432.

[78] Rüdiger F Pohl. *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*. 1st ed. Psychology Press, Oct. 2004. Chap. 18. ISBN: 9780203720615. DOI: https://doi.org/10.4324/9780203720615. URL: https://www.taylorfrancis.com/chapters/edit/10.4324/9780203720615-26/effects-labelling-r%C3%BCdiger-pohl.

[79] Addie Cormier and Christopher Ng. "Integrating cybersecurity in hazard and risk analyses". In: *Journal of Loss Prevention in the Process Industries* 64 (2020), p. 104044. ISSN: 0950-4230. DOI: https://doi.org/10.1016/j.jlp.2020.104044. URL: https://www.sciencedirect.com/science/article/pii/S0950423019307995.

[80] *IEC 62443-3-1:2009 Security for industrial automation and control systems –Part 3-1*. Standard. International Electrotechnical Commission, July 2009. URL: https://webstore.iec.ch/publication/7031.

[81] Unknown. *Developing an Operational Technology and Information Technology Incident Response Plan*. [Online; accessed January 14th 2022]. Public Safety Canada, 2020. URL: https://www.publicsafety.gc.ca/cnt/rsrcs/pblctns/dvlpng-ndnt-rspns-pln/index-en.aspx.

[82] Tomomi Aoyama et al. "On the Importance of Agility, Transparency, and Positive Reinforcement in Cyber Incident Crisis Communication". In: *Critical Information Infrastructures Security*. Ed. by Simin Nadjm-Tehrani. Cham: Springer International Publishing, 2020, pp. 163–168. ISBN: 978-3-030-37670-3.

[83] Olga Kulikova et al. "Cyber Crisis Management: A Decision-Support Framework for Disclosing Security Incident Information". In: *2012 International Conference on Cyber Security*. 2012, pp. 103–112. DOI: 10.1109/CyberSecurity.2012.20.

[84] Ernest Mougoue. *SSDLC 101: What Is the Secure Software Development Life Cycle?* Tech. rep. Accessed on Oct 4, 2021. DZone, Security Zone, July 2017. URL: https://dzone.com/articles/ssdlc-101-what-is-the-secure-software-development.

[85] Alexandra Altvater. *What Is SDLC? Understand the Software Development Life Cycle*. Tech. rep. Accessed on Oct 20, 2021. Stackify, Apr. 2020. URL: https://stackify.com/what-is-sdlc/.

[86] Goran Jevtic. *What is SDLC? Phases of Software Development, Models, & Best Practices*. Tech. rep. Accessed on Oct 21, 2021. Phoenix NAP, May 2019. URL: https://phoenixnap.com/blog/software-development-life-cycle.

[87] Unknown. *6 Phases Of The Software Engineering Life Cycle*. Tech. rep. Accessed on Oct 21, 2021. Erbis, May 2020. URL: https://erbis.com/blog/6-phases-of-the-software-development-life-cycle/.

[88] Gerry Katz. "Rethinking the Product Development Funnel". In: (June 2016). Accessed on Oct 14, 2021. URL: https://ams-insights.com/wp-content/uploads/2016/06/Rethinking-the-Product-Development-Funnel.pdf.

[89] Michael E McGrath. *Setting the PACE in product development*. 1st ed. Routledge, 1996. ISBN: 9780080517056. DOI: https://doi.org/10.4324/9780080517056.

[90] Unknown. *Why are we satisfied by "good enough"?* [Online; accessed June 17th, 2022]. The Decision Lab. URL: https://thedecisionlab.com/biases/bounded-rationality.

[91] Warren E. Walker. "Policy analysis: a systematic approach to supporting policymaking in the public sector". In: *Journal of Multi-Criteria Decision Analysis* 9.1-3 (2000), pp. 11–27. DOI: https://doi.org/10.1002/1099-1360(200001/05)9:1/3<11::AID-MCDA264>3.0.CO;2-3. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/1099-1360%28200001/05%299%3A1/3%3C11%3A%3AAID-MCDA264%3E3.0.CO%3B2-3.

[92] C. Eden and F. Ackermann. "Mapping Distinctive Competencies: A Systemic Approach". In: *The Journal of the Operational Research Society* 51.1 (2000), pp. 12–20. ISSN: 01605682, 14769360. URL: http://www.jstor.org/stable/253943 (visited on 04/05/2022).

[93] Adam Smith. *An Inquiry into the Nature and Causes of the Wealth of Nations [The wealth of nations]*. London: W. Strahan and T. Cadell, Mar. 1776.

[94] Lisa Hill. "Adam Smith, Adam Ferguson and Karl Marx on the Division of Labour". In: *Journal of Classical Sociology* 7.3 (2007), pp. 339–366. DOI: 10.1177/1468795X07082086. URL: https://doi.org/10.1177/1468795X07082086.

[95] Y. Joseph Lin. "Division of Labor in Teams". In: *Journal of Economics & Management Strategy* 6.1 (1997), pp. 403–423. DOI: https://doi.org/10.1111/j.1430-9134.1997.00403.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1430-9134.1997.00403.x.

[96] Dorna Dehkhoda. "Combining IRAM2 with Cost-Benefit Analysis for Risk Management: Creating a hybrid method with traditional and economic aspects". [Online; accessed July 12th 2022]. Masters thesis. Luleå University of Technology: Department of Computer Science, Electrical and Space Engineering, 2018. URL: http://www.diva-portal.org/smash/get/diva2:1180133/FULLTEXT01.pdf.

[97] Unknown. *SABSA, Sherwood Applied Business Security Architecture*. [Online; accessed July 12th 2022]. SABSA Institute. URL: https://sabsa.org/.

[98] *ISO/IEC/IEEE15288:2015 Systems and software engineering — System life cycle processes*. Standard. International Standardisation Organisation / International Electrotechnical Commission / Institute of Electrical and Electronics Engineers, May 2015. URL: https://www.iso.org/standard/63711.html.

[99] *CLC/TS50701:2021 Railway applications - Cybersecurity*. Standard. European Standard, 2021. URL: https://www.en-standard.eu/clc/ts-50701-2021-railway-applications-cybersecurity/.

[100] *EN50126:2017 Railway Applications. The Specification and Demonstration of Reliability, Availability, Maintainability and Safety (RAMS) Generic RAMS Process*. Standard. European Standard, Dec. 2017. URL: https://www.en-standard.eu/bs-en-50126-1-2017-railway-applications-the-specification-and-demonstration-of-reliability-availability-maintainability-and-safety-rams-generic-rams-process/.

[101] Unknown. *Project SECRET: security of railways against electromagnetic attacks*. [Online; accessed July 14th, 2022]. 2012. URL: https://secret-project.eu/About.

[102] Unknown. *Europe's Rail*. [Online; accessed July 14th, 2022]. Nov. 2012. URL: https://rail-research.europa.eu/about-europes-rail/.

[103] UIC - Rail System Department. *Guidelines for Cyber-Security in Railways*. UIC - Rail System Department, June 2018. ISBN: 9782746127326. URL: https://www.shop-etf.com/en/guidelines-for-cyber-security-in-railways.

[104] Marianthi Theocharidou et al. *Railway Cybersecurity - Good Practices in Cyber Risk Management*. Tech. rep. Accessed on July 14th, 2022. ENISA, Nov. 2021. DOI: https://doi.org/10.2824/92259. URL: https://www.enisa.europa.eu/publications/railway-cybersecurity-good-practices-in-cyber-risk-management.

[105] Helmut Klarer and Christian Schlehuber. *Zoning and Conduits for Railways*. Tech. rep. Accessed on July 14th, 2022. ENISA and ER-ISAC, Feb. 2022. DOI: https://doi.org/10.2824/761090. URL: https://www.enisa.europa.eu/publications/zoning-and-conduits-for-railways.

[106] U.E. Gattiker. *The Information Security Dictionary: Defining the Terms That Define Security for E-Business, Internet, Information, and Wireless Technology*. Kluwer International Series in. Springer US, 2004. ISBN: 9781402078897. URL: https://books.google.nl/books?id=Z63j0hWZ1W0C.

[107] Mark Wijkhuizen. "Security analysis of the iTasks framework". [Structure] Section on Network Security Requirements. Bachelor thesis. Nijmegen, Netherlands: Radboud University, June 2018.

[108] Wm. Arthur Conklin. "IT vs. OT Security: A Time to Consider a Change in CIA to Include Resilience". In: *2016 49th Hawaii International Conference on System Sciences (HICSS)*. 2016, pp. 2642–2647. DOI: 10.1109/HICSS.2016.331.

# Part IV

# Appendix

# Appendix A

# Security Models - In-depth discussion

Traditionally, security models are used to discuss value and related security goals in the context of assets. This chapter features an in-depth discussion of some security models. It is sought to understand the function and thus effectiveness of these models in various applications.

Although this thesis is adapting existing techniques to the dynamic innovation environment, the security models are relevant: these models are not software-bound by definition and can thus be used as they are, without modifications, in most fields including OT. What remains is a question of effectiveness: this is answered by looking at the function and role of the security model in the threat-assessment process.

Throughout the security models, it becomes clear that there is a desire to better model security. We hope to help in this discussion by combining some of the fundamental views on value of assets, and by putting this security model into perspective in the security equation for the SOTDLC.

**Functions of a security model** Although the security model is primarily meant to *assess value in relation to assets*, it must be recognised that these models are also often used very effectively

- in raising awareness and education of security;
- expand the notion of merely assessing value by also looking outward to mitigation strategies; and
- risk classification.

**Proposed model** Overall value of assets is accurately modelled using the following three dimensions:

1. *CIA*, which models value in terms of security properties;
2. *Storage, transmission, and processing*, describing the stages in which the information and its value can manifest itself; and
3. *Unrealised, Ongoing, Contained*, describing the temporal phases for violations to value.

This new model addresses and improves shortcomings of existing systems, while maintaining the implicit trend towards a more general value model. The disadvantage of this

model is that it is more complex and less intuitive. For this reason, this method is unlikely to be suitable for introductory awareness courses. Nonetheless, such a more complex value model seems to be necessary to deal with the similarly complicated nature of the SOTDLC. Furthermore, the security equation is thought to be able to perform this duty.

## A.1 Traditional CIA

The CIA model (Confidentiality; Integrity; Availability) [15, Ch. 1.1] describes primary security properties for assets. If these properties are breached by an attacker, the asset is considered compromised. This model is however debated for its completeness, as not all types of attacks nor all security-values to an asset strictly fall within the boundaries of CIA. The advantage of taking this basic strategy, is its applicability to both soft- and hardware. Nonetheless, some the extensions of CIA have been developed specifically with software in mind. Therefore, particular care must be taken when analysing problems that are very specific to hardware.

## A.2 CIA-UAA

The CIA-UAA model [106, 107], adds three basic principles:

- *User accountability* The person or entity accessing certain information is reliably logged and can therewith be held accountable for this access in case of an audit.
- *Authentication* Only authorised users should be able to access specific data.
- *Audit* The system should keep logs of security-relevant events. These logs can then be audited and used to trace certain data events.

This model essentially adds a layer of access control and logging. This is captured in what ostensibly are basic properties. This model is however subject to the criticism that access control and logging are part of a correct implementation of the original CIA principles.

- User accountability fits into confidentiality and integrity, and considers what happens if these properties are broken: who should be held accountable? This implicitly expresses the need for some framework to deal with risks through time: what happens if the manifestation of a risk was not prevented, and it happened? More on this in section A.5.
- Authentication is a security control: it enforces accountability through forcible controls. Authentication is really not a property of information or architecture, rather a security control applied to distinguish attackers from legitimate users when they may impact one of the CIA properties of one or more assets.
- Audit, in the sense of logging, is an implementation of a control for accountability.

Although UAA is not really a useful fundamental addition, it is good to note that correct application of access control and logging are imperative to a mature security implementation.

## A.3 Stride

STRIDE [75, 74] (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege) is the default system used by the NS. This model is used to categorise threats by attacker goals. STRIDE extends the traditional CIA model, by adding three additional principles (table A.1) (similar to CIA-UAA, see above). By giving a more complete overview of security properties for assets, the model aids in finding such vulnerabilities.

| | | |
|---|---|---|
| Spoofing | S | Authentication |
| Tampering | T | Integrity |
| Repudiation | R | Non-Repudiation |
| Information disclosure | I | Confidentiality |
| Denial of Service | D | Availability |
| Elevation of privilege | E | Exploitation |

Table A.1: Security Principles of STRIDE (CIA-ANE)

STRIDE is commonly seen as threat modeling technique; it is however more closely related to the above security models. More accurately, it is a strategy towards adhering to such a model: a classification method. Microsoft, the author of the STRIDE model, has documented the intended usage of the STRIDE model [76]: they recommend categorising the threats *after* listing them.

STRIDE's dangers are similar to those of the engineering-approaches to attacker models in ch. 5.3.1: they lead to survivorship bias. STRIDE uses known, commonly successful, attacks to determine likely areas of interest. It does not, however, give any insight into what exactly made other attacks unsuccessful. Those things that are actually effective at stopping attacks are thus (to some degree) excluded from consideration, while these are the most interesting things to apply elsewhere to prevent other attacks. Therefore, the usage of STRIDE is discouraged, since it gives a false sense of security, more so than the other techniques discussed in this chapter.

### A.3.1 Additions to CIA

STRIDE essentially adds the Authentication, non-Repudiation, and Exploitation. For each of these we look at their added value critically:

- Authentication is an implementation of a security control. This was also described for CIA-UAA in ch. A.2: Authentication is really not a property of information or architecture, rather a security control applied to distinguish attackers from legitimate users when they may impact one of the CIA properties of one or more assets.
- Repudiation is a generalisation of audit in CIA-UAA (ch. A.2). Here too, it is a manifestation of the need for looking at threats in a broader sense. (see CIA-R in ch. A.5).
- Exploitation, in elevation of privilege, really comes down to whether or not security mechanisms are considered assets with their own CIA properties. If security is seen as part of the system, which also needs its own appropriate countermeasures, the need to

secure against elevation of privilege arises as well because this impacts the integrity or confidentiality of a security control.

### A.3.2 Checklist Behaviour

In development stages, STRIDE introduces a risk towards the threat assessment process: in practice, it is easy for security testers and developers to quickly glance over their application and then select one or two examples for each STRIDE attack category. The resulting threat assessment is by no means exhaustive, thus providing a false sense of security. The benefit of the categorisation of the attacks (over simple CIA) in this way, therewith does not outweigh the false sense of security. To make the STRIDE model effective, it could be used from an asset perspective similar to the security models above. A listing of assets is a prerequisite for such a STRIDE model, as this allows for a thorough review of the six security properties for every asset.

## A.4 McCumber Cube

McCumber states the following about his model [48, p. 328]:

> "This model not only addresses the threat, it functions as an assessment, systems development, and evaluation tool. (...) Its application is universal and is not constrained by organizational differences."

McCumber [48, 49] argues that systems and their data-assets can be very abstract. This inherently makes it difficult to reason about properties of these assets. Furthermore, these assets do not only just exist (*storage*), but they are transmitted to other parts of the system (*transmission*) and processed by programs or hardware (*processing*). As an attacker, the often overlooked transmission and processing states are also interesting and thus a target. From that it follows that the security principles defined for the system should also be applied to these states.

The application of the CIA-triad on the three data-asset states can be done in three domains:

1. *Technology* (actual security measures in both soft- and hardware);
2. *Policy and Practice*: security should be expressed in policies which clearly describe and outline strategies to mitigate risks surrounding assets. McCumber uses the example of data multiplication as a security-breach which is effectively prevented by clear and enforced policy.;
3. *Education, Training and Awareness* if one is not actually aware of a thing called 'security', it cannot effectively be applied.

McCumber describes the usage of the model as follows:

> "The model has several significant applications. Initially, the two-dimensional matrix is used to identify information states and system vulnerabilities. Then, the three layers of security measures can be employed to minimize these vulnerabilities based on a knowledge of the threat to the information asset."

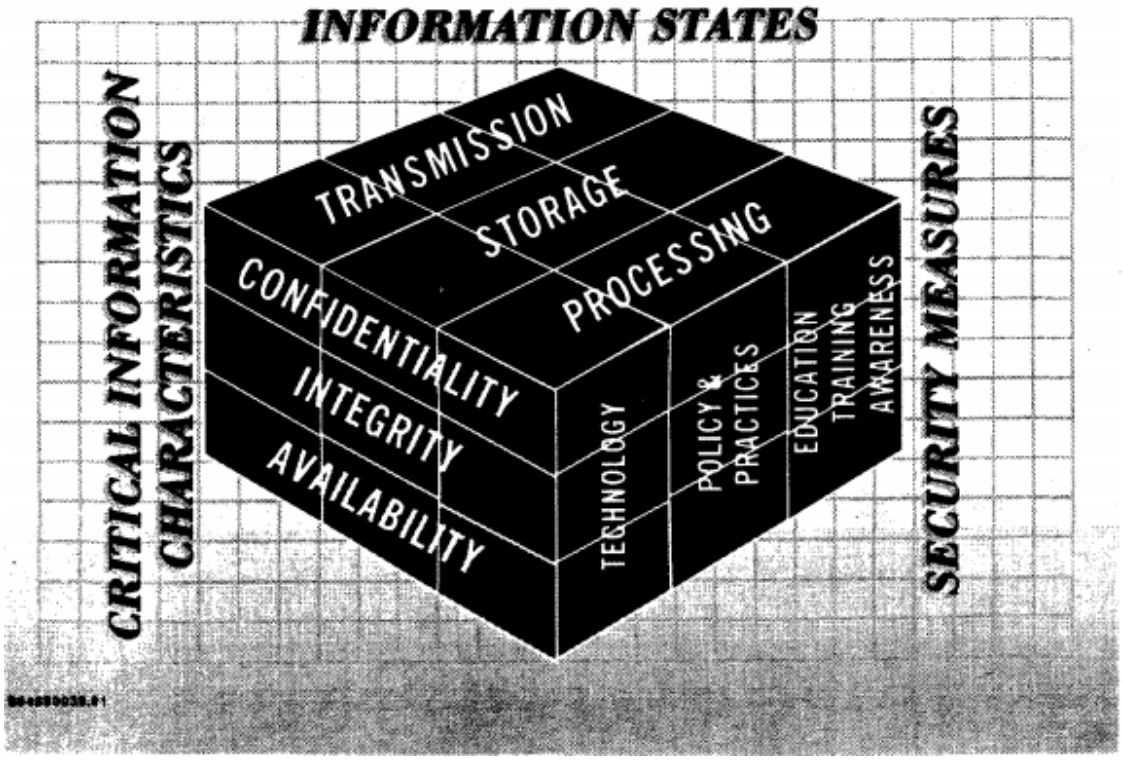


Figure A.1: McCumber Cube [48, p. 334 - Figure 1]

Overall this model gives an abstract approach to security, applicable to many contexts. Although the model has a very generalised approach, it is nonetheless on point and applicable in modern and diverse applications as seen OT.

### A.4.1 Value and Assets

The McCumber cube is not really a security model anymore, but more of a full security methodology: it is not meant to be just a security model. To simplify the discussion, we therefore only look at two of the three dimensions, the CIA-triad in the three information states (storage, transmission, and processing).

McCumber adds a very interesting dimension to the story. He recognises that information, whatever that may be, is dynamic: it is moved around and used. Not only could the security goals (CIA) vary between these information states, but the security controls could vary too. Fundamentally this changes the game, because the system is no longer viewed as standalone assets, but as assets with one or multiple functions with relation to the information. For example a data-bus could be considered from the perspective of transmission, whereas the devices connected to it are parsing and thus processing and possibly storing (e.g.) the information (or meta-data thereof) sent on the bus. This gives more mental-tools to the engineer to reason about the system.

## A.5 CIA-R

The security models discussed so far, do not completely cover the extent value may have. Time tends to play a role as well: some systems can be offline for weeks before there is a significant impact on the business, whereas other systems are concerning in any case of misbehaviour. This is especially the case in OT. Value here, is not definitively lost, if the security properties pertaining to the value of the asset are lost for a limited amount of time. The CIA-R model by Conklin [108] provides some structure to this idea.

Conklin [108] considers the needs for OT specifically, and finds several unique challenges to OT that distinguish it from IT. Ultimately, security is an aspect that should be in alignment with existing business objectives, safety and compliance.

Conklin defines four objectives with respect to risk to protect value [108, ch. 5]:

1. *Prevention* of risk realisation;
2. *Sustenance* of services when a risk is realised;
3. *Recovery* from realised risks, including minimisation of consequences to the organisation (if the organisation is severely damaged by the consequences, its ability to recover is also compromised); and
4. *Optimisation* of cost and system effectiveness by optimising prevention, sustenance and recovery.

It could be said that in essence, Conklin considers three stages in which a risk could manifest itself through time:

1. *Unrealised*: the risk is still only theoretical, the event has not occurred.
2. *Ongoing*: the risk has (partially) occurred and is in the process of producing consequences, where the consequences can be new risks or directly negative effects for the organisation.
3. *Contained*: the risk realisation has been contained and no additional consequences are developing. The consequences that have already occurred are being dealt with on their own.

This approach does not necessarily mean that the system can be brought back to its original state, even if the threat is contained. For example, if confidentiality has been broken, this cannot be undone; but the negative consequences of this to the company can be limited. This also addresses some of the concerns voiced in previous sections: the need for remedies against attackers. Sustenance and recovery involve the management of existing impact and the prevention of further consequences: this could also (partially) be achieved by catching the perpetrator. Controls such as logging, access control and monitoring could be implemented (before the incident) to aid in the sustenance and recovery phases. Especially if the impact of a risk is low, and prevention is difficult, it may be more cost-effective to focus on sustenance or recovery.

By looking at risks in this multi-stage manner, security controls can be fine-tuned to deal with the risk in the most (cost-)effective manner and stage. Consequently, it is more feasible to find appropriate countermeasures that fit into existing OT constraints. On a side-note,

Conklin develops this model purely from the business-objectives perspective of OT, but the consideration of risk-manifestation stages could be quite useful for complex IT environments as well.

## A.6 Riahi-model

Riahi et al. [47] describe a complex model to look at security properties of IOT. This approach is very different from the CIA security model, in describing the asset in context to the user/person, process and technological ecosystem. This approach is a more holistic approach, similar to the one developed in this research. This model shows the relationships between each of the nodes [47, Figure 1] and lists the formal techniques that existed in 2013 to deal with modelling these relationships. Riahi et al. describe the many areas that are still in need of more research.

The complicated and incomplete nature of this model, make it unsuitable for application in industry at this time. Nonetheless, it provides a very interesting different view on security, through formal methods. Note that we did not look into continuations of this research or more recent overviews of this area, so more recent work may very well exist.

## A.7 Concerns

The way the above models are defined raises some concerns.

### A.7.1 About Mitigation in Security Models

It is tempting to assess risk and immediately jump ahead and suggest important countermeasures at the same time. We see this in STRIDE, through authentication, repudiation and exploitation; in the UAA of CIA-UAA; and in one of the three dimensions of the McCumber-cube. Nonetheless, this is counterproductive when looking at the system at design-time, because it reflects only the perspective of security, and none other.

Sure, when a system already exists and security is implemented retroactively, it is good to think in terms of the system architecture holistically and to suggest a broad security approach based on possible attack vectors. This is what such security models are great at: security engineers who need only suggest countermeasures to make a system secure, need only consider basic approaches to countermeasures. Such an approach is, however, likely to be perceived as hostile and a burden by others. Those who have worked very hard on making a system work, have to change it all because of some security rules that are imposed on them, without any regard for the choices they have made. Furthermore, the countermeasures are likely to be disproportionate to the actual threat, and inefficient because the context and actual expected threat level is not yet known.

At design-time, however, the architecture is neither complete nor documented: only designers and engineers who are intricately aware of the subtleties of design-choices could know what appropriate countermeasures are, as the system evolves. These engineers cannot solely consider the SuD, but also have to take into account CS and business goals. Consequently, the interests they have to balance are much more complex than simply assuming CIA security goals and applying default countermeasures. Here, security has to be considered in context, and has to be weighed appropriately and kneaded with the other ingredients of the system.

This is exactly what the security equation is about. This approach is about developing a thorough understanding of the business-objectives, and applying these proportionally in every facet of the project and for each risk throughout time.

Because of this, it is argued that the security model should refrain from suggesting countermeasures. Countermeasures are proposed only after other interests have been considered by people who understand what is at stake.

### A.7.2 Awareness and Education

Acronyms such as CIA and STRIDE are easily remembered by and explained to beginners in security. This makes it easier to explain the necessity for security in less technical terms. This is a great start! The danger is, however, that those who have had such training start to apply it to their projects and think they are done. Awareness of risks and unstructured deployment of countermeasures are unlikely to be effective, and prove to be so in practice. In the end, we have achieved nothing or even worse, because the system remains insecure and meanwhile the waters have been troubled because stakeholders are wrongly under the impression that the system has been secured adequately.

A better approach for management and business-specialists would be to use a concept that is still easy to understand and remember, which does justice to project structure and business-objectives. For example, the risk-equation: risk and business-objectives are already well-understood concepts by business-specialists and managers.

Engineers and architects must be more intricately aware of the implications of security in their implementations. This means that they must understand what kinds of controls are available, and how they impact CIA. Engineers and architects know how to build systems that comply with the goals set out by business. They should be trained in recognising which information should remain C, I or A, and what could generally be done to ensure this. They should look further as well: security components and countermeasures become new assets in their own right with associated information. This must be assessed and treated as well. For example, when a data connection is encrypted, new assets are the encryption processor/program and keys. Assuming the encryption is suitable for the task, encryption programs must maintain integrity and possibly availability (depending on the data), and keys must remain confidential and should be replaceable for resiliency. In summary, the understanding of engineers should go much deeper than simply the security model: they are at the forefront of securing the system, and doing so (cost-)effectively.

## A.8 Conclusion: Time-Adjusted McCumber

Various security models for describing valuable properties of assets have been proposed, typically focussed on software. The general consensus is that *Confidentiality, Integrity* and *Availability* (CIA) are the core properties that must be protected. Based on the available models, it is concluded that none of them currently give a complete overview of the relevant considerations. Therefore they are combined into a new system, with high regard for the expressed concerns. Overall value of assets is accurately modelled using the following three dimensions:

1. *CIA*, which models value in terms of security properties;
2. *Storage, transmission, and processing*, describing the stages in which the information and its value can manifest itself; and

3. *Unrealised, Ongoing, Contained*, describing the temporal phases for violations to value.

This new model, *Time-Adjusted McCumber*, addresses and improves shortcomings of existing systems, while maintaining the implicit trend towards a more general value model. Such a more detailed value model seems to be necessary to deal with the more complicated nature of the SOTDLC.

# Appendix B

# Overview of International Developments

For those interested in further reading, I have included some links to public sources (based on NS internal document called "ATO ontwikkelingen om ons heen"):

https://www.spoorpro.nl/innovatie/2021/04/23/video-interoperabiliteit-van-ato-over-etcs-met-succes-getest-in-engeland/
https://www.linkedin.com/pulse/introduction-mainline-automatic-train-operation-bastian-simoni/
https://www.railwaypro.com/wp/philippines-invites-japanese-suppliers-for-nscr-tender/
https://www.globalrailwayreview.com/news/123711/france-autonomous-regional-train-testing/
https://www.sciencedirect.com/book/9780128228302/unmanned-driving-systems-for-smart-trains#book-description
https://www.uic.org/com/enews/article/autonomous-technologies-in-rail-discussed-at-a-successful-dedicated-webinar?page=modal_enews
https://www.spoorpro.nl/innovatie/2021/06/22/stadler-rondt-test-met-ato-af-op-zwitsers-spoor/
https://www.railway-technology.com/news/alstom-germany-s-bahn-trains/
https://projects.shift2rail.org/download.aspx?id=f3f711e9-79c0-4560-826a-f6652b45529f
https://www.railtech.com/rolling-stock/2021/07/01/first-regional-trains-with-ato-to-be-in-service-in-luxembourg-by-end-of-2021/?gdpr=accept
https://www.spoorpro.nl/spoorbouw/2021/02/22/rusland-zet-dit-jaar-eerste-autonome-reizigerstrein-in-op-reguliere-dienstregeling/
https://www.railjournal.com/regions/asia/jr-east-to-introduce-ato-on-the-joban-line-next-month/
https://www.railjournal.com/technology/sncf-and-db-sign-new-digital-cooperation-agreement/
https://www.sciencedirect.com/science/article/abs/pii/S0263224121002529
https://www.prnewswire.com/news-releases/lilee-systems-announces-saferail-the-first-hi-rail-collision-avoidance-system-on-the-market-301249107.html
https://railway-news.com/siemens-mobility-wins-sydney-rail-modernisation-

contracts/
https://www.railjournal.com/in_depth/unlocking-network-capacity
https://www.railjournal.com/technology/autonomous-technologies-for-rail-on-test-in-toronto/
https://www.youtube.com/watch?v=5vzoMk4Sjo4
https://www.transport.nsw.gov.au/projects/current-projects/digital-systems-program
https://www.spoorpro.nl/spoorbouw/2021/03/31/liveblog-railtech-europe-dag-2-ertms-ato-en-duurzaamheid/
https://www.spoorpro.nl/innovatie/2020/11/05/consortium-voor-ontwikkeling-communicatiesysteem-frmcs-officieel-van-start/
https://www.globalrailwayreview.com/news/112804/5grail-project-frmcs-officially-launched/
https://www.railjournal.com/signalling/jr-east-to-trial-ato-and-5g-on-joetsu-shinkansen/
https://www.railtech.com/infrastructure/2020/11/12/project-launched-to-improve-exact-train-position/
https://www.mckinsey.com/featured-insights/europe/digitizing-europes-railways-a-call-to-action
https://www.youtube.com/watch?v=gTYocE6jDHk
https://www.spoorpro.nl/innovatie/2020/11/26/twee-nieuwe-ato-onderzoeken-van-start-in-duitsland/
https://www.railjournal.com/technology/german-mainline-ato-research-to-identify-safety-parameters/
https://www.globalrailwayreview.com/news/115920/sncf-semi-autonomous-train-french-rail-network/
https://www.railtech.com/digitalisation/2021/01/14/germany-200-million-state-aid-to-promote-ertms-and-ato/
https://www.tno.nl/nl/over-tno/nieuws/2021/1/ato-ontwikkelingen-mogelijkheden-automatische-treinbesturing/
https://www.railjournal.com/infrastructure/alstom-awarded-delhi-meerut-rrts-signalling-contract/

# Appendix C

# Criminological Pointers

Unfortunately, we did not have time to give a more in-depth review of criminological research. Nonetheless, here is a list of sources that we read which seemed useful in this context:
https://ieeexplore.ieee.org/abstract/document/6149195?casa_token=z5rrJpSgAhQAAAAA:ZxRpPyxmd1RLZBNOSYuINSATdv1kN5g8KiitaP-zkkaI8jlaZmaE4DygtsnazIURnWwtq3j4OCXDGts
https://www.emerald.com/insight/content/doi/10.1108/13685201211266015/full/html?casa_token=FIpepBQLn9kAAAAA:dr4XE8x69dy5CalJ72o6zN_-c-ril35wJJ9cVc8GTRhfyr4UxFQAEIiYedtGEDVImerT5Fd7r3GdPtrprOxUS3NcGCZQMMJTOIsGqna1xwfuTCLN3k3dmg
https://link.springer.com/chapter/10.1007/978-3-319-68711-7_22
https://link.springer.com/content/pdf/10.1007%2F978-3-642-45330-4_9.pdf
https://www.researchgate.net/profile/Hamdi-Kavak-2/publication/305764162_Towards_Modeling_Factors_that_Enable_an_Attacker/links/579f6a0108aece1c721562b6/Towards-Modeling-Factors-that-Enable-an-Attacker.pdf
https://heinonline.org/HOL/Page?collection=journals&handle=hein.journals/crpp16&id=740&men_tab=srchresults
http://dspace.stellamariscollege.edu.in:8080/xmlui/bitstream/handle/123456789/5928/cyber.pdf?sequence=1&isAllowed=y
https://www.sciencedirect.com/science/article/pii/S2212567115010771
https://www.researchgate.net/profile/Hugo-Barbosa/publication/315351300_SOCIAL_ENGINEERING_AND_CYBER_SECURITY/links/599c43430f7e9b892bafc0df/SOCIAL-ENGINEERING-AND-CYBER-SECURITY.pdf
https://www.timreview.ca/article/861
https://link.springer.com/book/10.1007%2F978-3-319-97181-0
https://www.ideals.illinois.edu/handle/2142/103052

# Appendix D

# Interview Plans

## D.1 Interview 1: Free Exploration of Ideas and Perspective

The objective of this interview is observing the opinion of interviewee. The interviewee is encouraged to describe their point of view from the perspective of their own expertise, without influences from literature. The ideas of the interviewee are challenged and discussed in such a way that the strongest possible variant of them is the outcome. The interviewer will try to address inconsistencies and pitfalls in a positive attempt to legitimise the ideas of the interviewee. The interviewer will also attempt to prevent bias towards the framework.

In the second interview, the resulting models are compared to the framework result of this research.

Hereafter follows the outline of the interview. This layout is also used to make notes during the interview itself.

### D.1.1 Introduction (5 minutes)

1. **Date & time:** ____ - ____ - 2022, Start: ___ : ___ End: ___ : ___
2. **Location/room:** _____
3. **Respondent No:** _____
4. **Interviewer name:** David Vonk

#### Interviewee

5. **Name:** _____
6. **Department:** _____
7. **Job Desc:** _____
8. **Years experience:** _____

9. **Prior knowledge about SOTDLC framework?** ○No ○Heard about it ○In depth:

_____

_____

_____

#### Privacy

Note that audio of this interview is preferably recorded. Besides audio recording, the interviewer will be making notes throughout the interview. The recording will not be published, but summaries of the interview would be.

10. **Is it okay to continue recording audio throughout the interview?** ○Yes ○No

**Are you okay with publication of...** (You can reconsider later if you want)

11a. **... your name?** ○Yes ○No ○Maybe

11b. **... your job title?** ○Yes ○No ○Maybe

11c. **... your department?** ○Yes ○No ○Maybe

11d. **... years experience?** ○Yes ○No ○Only approximate ○Maybe

### D.1.2 Security and Quality (20 minutes)

Definition: quality is everything that does not directly contribute to producing value in and of itself. Quality can improve value production or prevent risk to it. Examples of quality aspects are legal, management, compliance, safety, security, reliability, environmental impact, etcetera. We use quality in this interview over security, because it is more abstract and more palatable for people from other specialties, and believed to be analogous in procedure.

12. **How should it be determined what level of security (i.e. quality) the SuD should have? What kind of criteria would you use?** Prompts about risk and ways of dealing with risk could be used, if needed. Also e.g. business model, strategy.
13. **What kind of knowledge or (organisational) support do you think is necessary to make such an assessment?**
14. **In what way would you prioritise which actions should, and should not be taken with regard to security (or quality aspects) in the SuD?** E.g. statically or dynamically; are some quality subjects inherently more important than others?

### D.1.3 Project Structure (10 minutes)

Within the NS, there is a difference between a program and a project. A project is roughly a team-effort towards a particular business objective of the company. A program is a combination of multiple similar projects into a single bigger project (i.e. the program), where knowledge and team-members are commonly shared.

In this research we use project and program synonymously, a program simply considered a large project.

15. **Do you think that security (quality) in general is similar or dissimilar to security (quality) in research and development projects (esp. in OT)?**
    ○Similar ○Dissimilar ○Depends
16. **How would you describe the stages a research and development project goes through?** Prompts about basic structure (startup, development, transfer to maintenance) if needed.
17. **How do these project-stages impact the way you would look at quality aspects (security) throughout the project?**

### D.1.4 Teamwork (20 minutes)

Gathering knowledge and insight is one thing, making decisions is what transforms it into progress.

18. **Who should make final decisions on security (quality) requirements, and how should they do that?** I.e. how would they weigh all the different interests that may play a role in such a decision? What could be potential problems?
19. **Who else would have a say in or influence on such decisions, and in what way?**
20. **What should the relationship between a project and the rest of the organisation look like?**
21. **How would you deal with disagreement between stakeholders or team members?**

### D.1.5 Wrapping up and summary (5 minutes)

**Most important topics or points:**

**22a. Topic:** _____

**Importance:** Noteworthy ○—○—○—○—○ Crucial

**22b. Topic:** _____

**Importance:** Noteworthy ○—○—○—○—○ Crucial

**22c. Topic:** _____

**Importance:** Noteworthy ○—○—○—○—○ Crucial

**23. Final Remarks**

_____

_____

_____

_____

_____

## D.2 Interview 2: Discussion of Framework Draft

### D.2.1 Introduction (5 minutes)

The result of the first interview was a collection of framework fragments. These fragments make up more or less all major aspects of a policy framework for security and quality. These framework fragments can be compared with the framework produced by this research.

The objective of this interview is to find out how robust the draft SOTDLC is: to test, challenge and verify the theory in the theoretical framework. We do this by looking at the intuition of the interviewee and comparing that to the in-depth framework philosophy: if the framework does not make intuitive sense, it may be too complicated and difficult to introduce into an organisation.

To achieve this objective, we will compare the framework fragments from the last interview to the SOTDLC, and discuss the opinion of the interviewee: based on the ideas in the SOTDLC, would they stick to their original views, or change their mind? To structure this, the interview has a base structure, describing two basic aspects about projects and security. After that, the idea is to choose 1-3 additional subjects to discuss in-depth, depending on the expertise and preference of the interviewee.

#### Basic subjects

- Project-model based on phases, figure D.2. Throughout this project, security priorities shift D.2.
- General Risk Equation, figure D.1

#### Elective subjects

- Modelling value - STRIDE and time-adjusted McCumber model
  - CIA
  - Storage, Transmission, Processing
  - Unrealised, Ongoing, Contained
- Risk Equation and Criminology - usage of criminology for quantifying risk (figure D.1)
- Organisational Context and Project Hierarchy
  - Organisational Support
  - Project Hierarchy (figure D.3)
- Project models in current practice and literature (table D.2)

### D.2.2 Admin

1. **Date & time:** ____ - ____ - 2022, Start: ___ : ___ End: ___ : ___
2. **Location/room:** _____
3. **Respondent Name:** _____
4. **Respondent No:** _____

#### Privacy

Note that audio of this interview is preferably recorded. Besides audio recording, the interviewer will be making notes throughout the interview. The recording will not be published, but summaries of the interview would be.

5. **Is it okay to continue recording audio throughout the interview?** ○Yes ○No

### D.2.3 Wrapping up and summary (5 minutes)

**Most important topics or points:**

6a. **Topic:** _____
**Importance:** Noteworthy ○—○—○—○—○ Crucial

6b. **Topic:** _____
**Importance:** Noteworthy ○—○—○—○—○ Crucial

6c. **Topic:** _____
**Importance:** Noteworthy ○—○—○—○—○ Crucial

7. **Final Remarks**

_____
_____
_____
_____
_____

### D.2.4 Base project-model and security-equation

#### Project model

Project-model based on phases, figure D.2. Throughout this project, security priorities shift D.2. Here the focus lies with knowing the now (assets) and finding out what value needs protection for every stage of the project.

8. **What are your thoughts about this way of modelling an innovative project?**
9. **Is this model sufficiently flexible to model innovation in OT in general?**
10. **Do you think this way of looking at priorities produces security by design?**
11. **Is this prioritisation feasible to implement in practice?**

#### Security Equation

General Risk Equation, figure D.1. Quantifying chance by looking at the attacker in more detail. Impact assessment is different, because the attacker has to do a value conversion. Attackers are a very varied population, so the risk equation offers an approach to better understand effective prevention and remediation through controls.

12. **Is this a complete representation of all inputs to the security process?**
13. **Do you think the criminological inputs are significant/meaningful as compared to BIA and current practices?**

### D.2.5 Modelling value - STRIDE and time-adjusted McCumber model

Time-adjusted McCumber is:

1. *CIA*, which models value in terms of security properties;
2. *Storage, transmission, and processing*, describing the stages in which the information and its value can manifest itself; and
3. *Unrealised, Ongoing, Contained*, describing the temporal phases for violations to value.

STRIDE is a classification method, with the risk of selection bias. Time-adjusted McCumber gives the flexibility to deal with all kinds of scenario's while still accurately looking at actual value in the organisation.

### D.2.6 Risk Equation and Criminology

The risk equation (figure D.1) estimates security risk based on knowledge about the attacker (criminology). It then considers two ways of dealing with that risk: simply accepting it; and mitigating it through controls (defined by requirements). Mitigation can be either preventive or remediative.

#### Core ideas: Risk

Give (rough) quantification of risks, but not to exact.

- How does an attacker operate?
- What is important to the attacker (value), and how can we exploit that?

#### Core ideas: Risk acceptance

Risk acceptance analysis is the current approach of the NS.

- Triangulation: architecture and component level analysis
- Seek clear alignment with business goals. Advantage in communication and awareness.
- Look at mccumber cube with time dimension based on system-model (STRIDE is only a classification method, and unsuitable for risk-assessment)

#### Core ideas: Security Requirements

Make explicit that classification is intermediate step between analysis and prioritisation.

- Classification is dark magic: be careful
- Prioritisation is guesswork because comparable quantification is impossible.
- Differentiate between prevention, resilience and incident management.

#### Questions

14. **Do you consider this way of assessing security (or quality) complete?** What is missing?
15. **Do you consider this way of assessing security (or quality) sound?** What would you remove from this? Which aspects are particularly important?
16. **What kind of relationship do projects and the organisation have with regard to such considerations?**
17. **Do you think that incorporating criminology and organisational support structure is reasonable and realistic?**

### D.2.7 Organisational Context and Project Hierarchy

**People** benefit from clear social cohesion and direction. This means:

- Have a clear hierarchy
- Work together to stimulate team-cohesion and feelings of ownership
- Speaking the same language
- Transparent communication

**The Organisation** is implied give support in a very different way than currently is done, e.g.:

- organisational learning,
- CSMS,
- SOC,

- incident management, and
- observing security state of affairs.

18. **Do you think this way of dealing with social relations would be usable in projects?** Think competency, ownership, and responsibility alignment.
19. **Do you think this system solves existing problems? If not, what are shortcomings/improvements?**
20. **Do you think NS (or any organisation) could adopt such an organisational model?** Cost v. benefit of adoption?

### D.2.8 Project models

This subject is less of a priority for these interviews unless the interviewee disagrees with the models described in the base-subjects of the interview.

The idea is to compare the 62443 [36] approach, which is based on the SDL model [35], to the multi-phase SOTDLC model using figure D.2.

21. **Is this a sound and complete view on projects? What is missing, and what should be added in the model?**
22. **What do you think about this process-model for OT innovations?** Agile, quirks, formality?

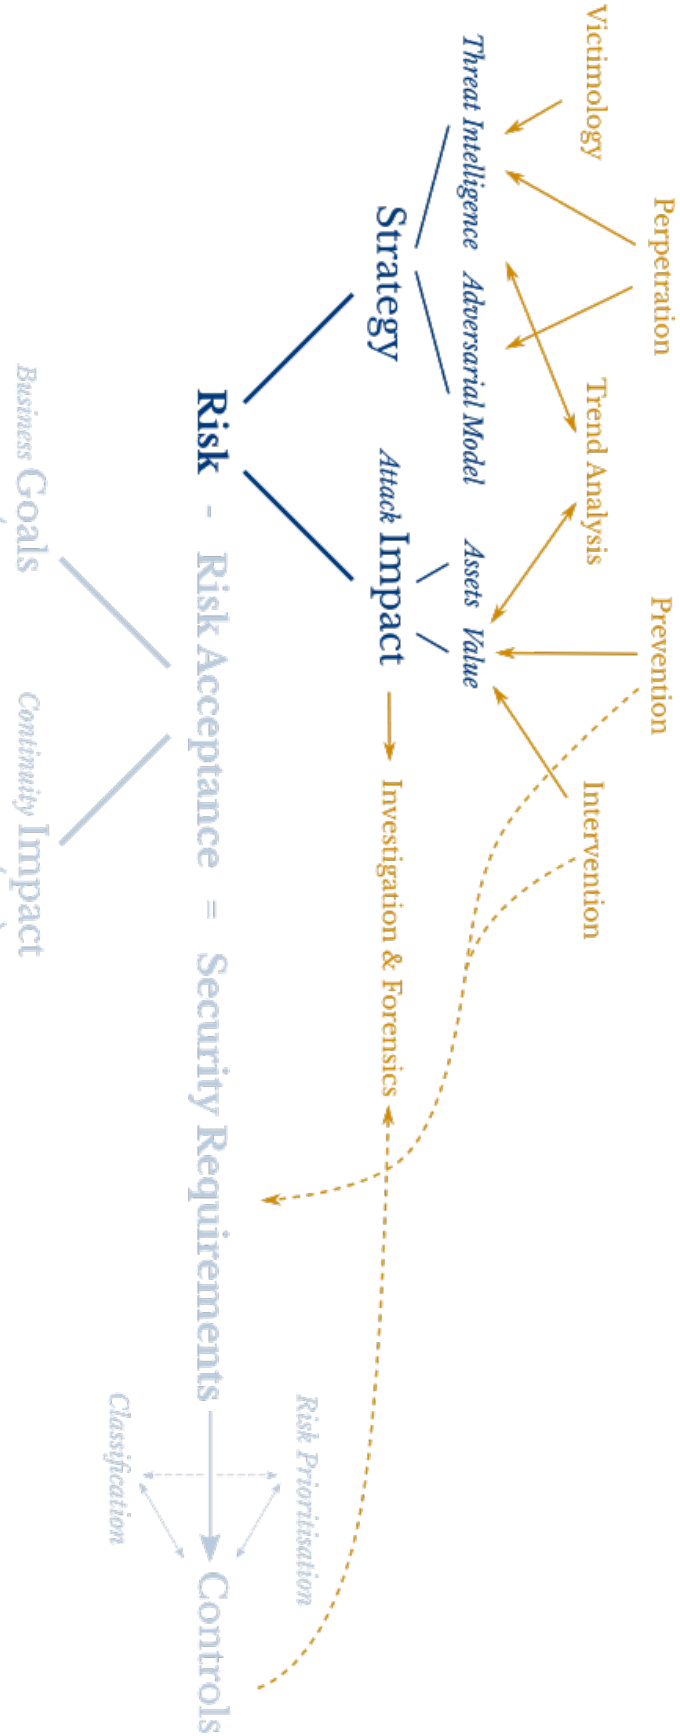


Figure D.1: Risk Equation with criminological context

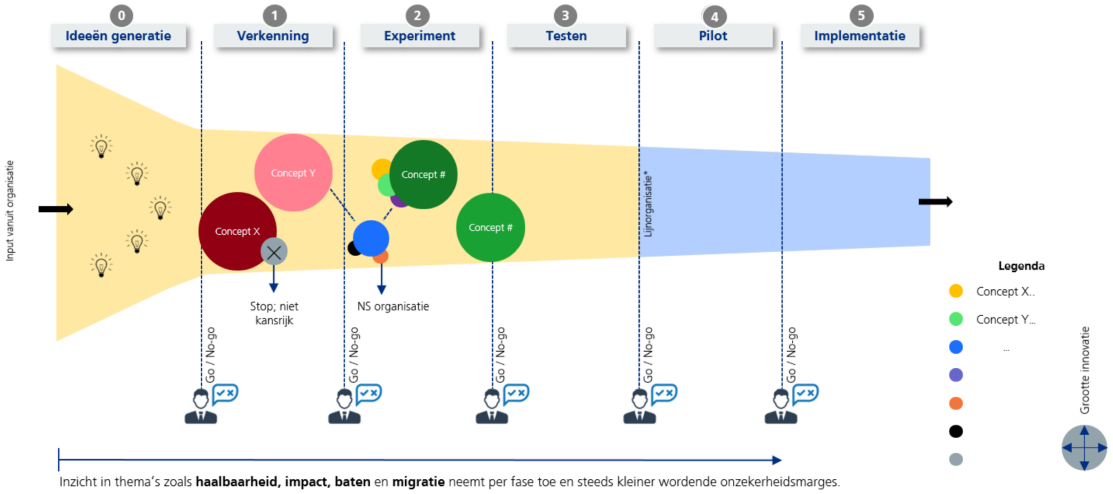


Figure D.2: Example development funnel, as used by NS [17]

| Project phase | Focus | Dominating Value |
|---|---|---|
| 0 | CS | Organisational facilities and operations |
| 1 | CS | Preventative wrt CS; Confidentiality of knowledge |
| 2 | CS | Preventative and Resilience wrt CS; Resilience wrt SuD; safety-driven |
| 3 | SuD | Resilience for whole SuC; Adding preventative measures for SuD |
| 4 | SuD | Reinforcing preventative measures for SuD to improve availability |
| 5 | SuC | Availability (reliability) for SuC |

Table D.1: Value related to project phases

| Phase Type | 0 | 1 | 2 | 3 | 4 | 5 |
|---|---|---|---|---|---|---|
| Funnel (Ch. 8.2.2) | Concept Evaluation | Planning & Specification | Design & Experimentation | Test & Evaluation | Pilot | Implementation & Roll-out |
| Security | Organisational Aspects | Baseline | Cruising | | Transition | |
| SDL (Stage) [35] | Education and Awareness (0) | Inception (1) | Best Practices & Product Risk Assessment (2,3) | Coding and Testing Policies (6,7) | Push (8) | Final Review & Release (9,11) |
| | | | Risk Analysis (4) | | | |
| | | | | Customer Prescriptions (5) | | |
| | Response Planning (10) | | Response Execution (12) | | | |

Table D.2: Timeline of the multi-phase project as described in chapter 9

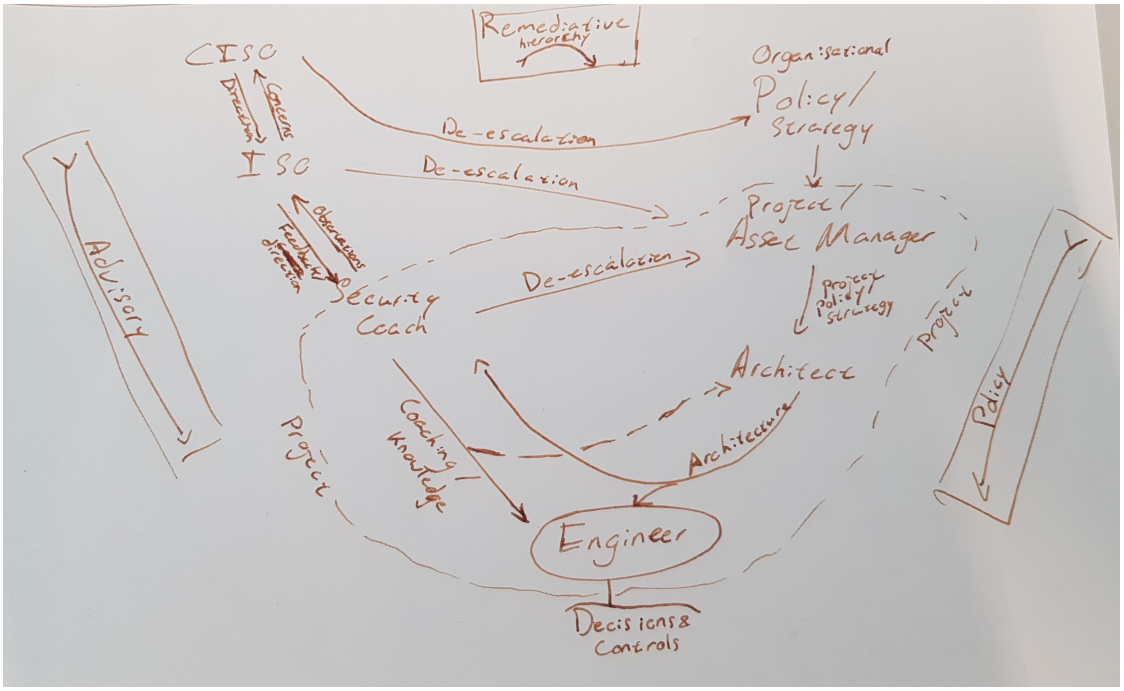


Figure D.3: Project level hierarchy for decisional support structure in OT Security. This sketch was used during all the interviews for uniformity, because a digital version was not yet available during the first interviews.

# Appendix E

# Interview Summaries

This chapter contains summaries of all interviews that were conducted. These summaries were checked and verified by the interviewee's.

### E.0.1 Respondents

#### Summaries

All interviews (except resp. 15.2) were recorded. These recordings were used to write extensive summaries for part of the interviews. For the other interviews, only the three most important topics, chosen by the interviewee, are reported along with organisational information about the interviews.

Unfortunately the recording was forgotten for the second interview of respondent 15. This interview-report was written on the same day, such that the report was as accurate as possible. This report was also verified by the respondent.

Due to time-constraints, not all interviews were given an extensive summary. The interviews with extended summaries give a good overview of what people thought, and what the main ideas were. Most of the respondents agreed with each other. Given the open, exploratory and subjective nature of these interviews, this does not impact the quality of the conclusions much.

All interviews were done by the same interviewer, David Vonk.

### E.0.2 Reflection on Interview Format

For the first interview, the interview-design worked well. It gave both sufficient structure and freedom for it to function well. Depending on the speciality of the respondent, the focus of the interview would lie on one of the three subjects, giving room for them to reason from their own perspective.

For the first respondent who had the second interview (resp. 3), a different interview-format was used. This design had the same content, but felt quite chaotic due to the lack of structure, and did not give sufficient structure to build on ideas and explore the robustness of the overall framework. Based on this initial experience, the interview-design was changed to the described strategy using primary and optional subjects. This worked well for the rest of the respondents.

## E.1 Interview Reports

In this version the interview reports have been redacted for the privacy of the interviewees and to protect the interests of Nederlandse Spoorwegen.

For scientific purposes only, the interviews can be requested by contacting the author.