ROYAL DANISH DEFENCE COLLEGE

DBI

SAVE

Social Vulnerability & Assessment Framework - A study on Social Engineering 2.0

Dennis Hansen

DANISH DEFENCE

SAVE

Social Vulnerability & Assessment Framework

Dennis Hansen (ed.)

2017

Royal Danish Defence College

Project SAVE - Social Vulnerability & Assessment Framework

Dennis Hansen (ed.)

Copenhagen, January 2017

Royal Danish Defence College
Ryvangs Allé 1
DK-2100 Copenhagen
Tel.: 728 17000
Editor in Chief: Danish Institute of Fire and Security Technology (DBI)
Printed in Denmark by Rosendahls A/S
Cover photo: Cyber criminal, © Igor Stevanovic, Dreamstime.com
Layout: Bent-Ole Kure
ISBN: 978-87-7147-175-5
Number printed: 350

Executive Summary

Introduction to the Study

Project SAVE is a study on social engineering supported by the Royal Danish Defence College, and developed by a consortium consisting of the Danish Institute of Fire and Security Technology (DBI), the Alexandra Institute and CenSec.

Social engineering is fundamentally the art of exploiting the human factor of cyber security, in an effort to compromise an organisation. Governments and private corporations can spend an indefinite amount of resources on security, yet social engineering will still remain a threat that poses a serious cyber security risk: the reason being that social engineering exploits people who already have access to the information that the attacker desires.

The aim of this study is to conduct an explorative investigation of the phenomenon of social engineering 2.0, through the following three objectives: (1) the execution of real-life simulations of social engineering 2.0 attacks against three Danish companies; (2) to raise awareness among key organisations and actors; (3) and to provide a framework to mitigate the associated risks of social engineering, on the basis of our findings.

Results

Three targets have been subjected to a total of 185 social engineering 2.0 attacks, applying various reconnaissance methods and attack vectors, which have been used to test the social vulnerability level of the involved parties, including spear-phishing, whaling, conventional phishing, smishing, PDF attacks and USB attacks.

Two out of three targets involved in the study had significant information available about them from open sources, which were successfully utilised in the attacks conducted. However, the one with the least information available – Target #2 – proved to be the one with the highest success rate of the three organisations, amounting to 77 pct. of the attacks being successful. For Target #1 and Target #3, the successful attacks amounted to 60 pct. and 43 pct., respectively.

The most deceptive attack vector was SMS, which had an overall success rate of 88 pct., while the least successful was the USB vector, which no participating individuals were deceived by. We believe the deceptiveness of the SMS vector relies on the trust that people generally have in their phones, combined with the lack of knowledge of SMS spoofing.

In a post-talk with the involved companies that were subjected to a total of 185 attacks, we discovered that only seven instances occurred, where an incident was reported by an employee to the responsible department or point of contact within the respective organisations. This is an alarmingly low number of incident reports, considering that the deception level of an attack decreases when more attacks against the same organisation is executed, as people become more aware.

Social Vulnerability Assessment Framework

The social vulnerability assessment (SVA) framework developed consists of a multi-level defence as countermeasure against the threat of social engineering. The SVA framework includes four levels: (1) Policy level: Procedures for employees to strictly adhere to; (2) Parameter level: Interactive awareness training of employees; (3) Persistence level: Frequent reminders of the threat of social engineering; and (4) Defensive level: Incidence response procedures and the establishment of a central point of contact.

The Consortium of SAVE

The consortium behind SAVE consists of the Danish Institute of Fire and Security Technology (DBI), The Alexandra Institute - both GTS institutes with international R&D activities and practical experience in the field of IT-security - and the Center for Defence, Space and Security (CenSec) - one of the leading industrial clusters in the defence and security industry in Denmark.

The Danish Institute of Fire and Security Technology

DBI is the leading Danish knowledge centre in the field of fire safety and security. They perform research and consultancy in relation to traditional, physical security solutions, and secure organisations against social engineering attacks and counterfeiting. We provide these services to private enterprises, institutions and public authorities.

Through their international network, DBI systematically collects and processes the latest information, and we actively contribute to, and participate in setting norms and standards within our key fields of activity, both nationally and internationally. Based on these activities, DBI develops and maintains a programme of innovative, valuable services that assists their customers in fulfilling their goals and obligations.

The Alexandra Institute

The Alexandra Institute (ALX) is a privately held, non-profit Research and Technology Organization (RTO) with approximately 100 employees. ALX is located in Aarhus and Copenhagen in Denmark, and is recognised by the Danish government as an advanced technology provider.

The Alexandra Institute focuses on applied research in computer science. The Alexandra Institute has (among other areas) strong expertise on IT security with a particular focus on applied cryptography. Their main competencies lie in evolving theoretical computer science results into practical solutions. AXL's security lab participates in several Danish and EU research projects in this field.

CenSec – Center for Defence, Space & Security

CenSec is an industrial cluster and a network centre for small and medium enterprises that already are, or wish to become suppliers to the defence, security and/or space industry.

CenSec's mission is: (1) to develop business networks among small and medium-sized sub-suppliers to the defence, security and space industry; and (2) to offer assistance to business members to improve market knowledge, competences and education, and thereby enabling them to participate in business networks.

Table of Contents

List of Figures

List of Tables

Abbreviations

AP – Access Point
BYOD – Bring Your Own Device
ccTLD – Country Code Top Level Domain
CNI – Critical National Infrastructure
cSE – Conventional Social Engineering
DDoS – Distributed Denial of Service
EMS – Electromagnetic Spectrum
EW – Electronic Warfare
FOCA – Fingerprinting Organisations with Collected Archives
HID – Human Interface Device
HUMINT – Human Intelligence
ICS – Industrial Control System
IDS – Intrusion Detection System
IoT – Internet of Things
IRP – Incident Response Plan
MITM – Man-In-The-Middle
OSINT – Open Source Intelligence
rSE – Reverse Social Engineering
SCADA – Supervision, Control and Acquisition Data
SDR – Software Defined Radio
SE – Social Engineering
SE 2.0 – Social Engineering 2.0
SET – Social Engineering Toolkit
SIGINT – Signals Intelligence
SMPP – Social Media Personality Profiling
SNA – Social Network Analysis
SOCMINT – Social Media Intelligence
SSL – Secure Sockets Layer
SVA – Social Vulnerability Assessment
TDoS – Telephone Denial of Service
TLD – Top Level Domain
URL – Uniform Resource Locator

Chapter 1:
Introduction

1. Introduction

The advent and popularity of social media has transformed the way we access and share information. As a result, accessing potential vital information has become increasingly easy, allowing for new and improved IT-security threats to emerge. Latest insight into security breaches reveals that approximately 90 pct. of recorded security incidents include the human element as a major component in cyber attacks[1]. Even global corporations, who have invested comprehensively in IT-security, have experienced attacks that exploit the human element[2]. Kaspersky Lab recognises that "*almost every type of [cyber] attack contains some kind of social engineering*"[3]. But what is social engineering? How has it developed in light of the emergence of new technology and social media platforms, and how can we use the methods developed for social engineering attacks to protect companies against these types of cyber threats?

In this study, we seek to answer these questions by exploring the phenomenon of social engineering and how it has developed from its traditional methods of obtaining and exploiting information by leveraging the human factor in security, to an advanced and highly developed skillset, which transcends the boundaries of the physical and cyber realms of security.

In the context of information security, social engineering refers to the psychological manipulation of people to perform specific actions that can compromise an organisation's security, or make them divulge confidential information for the purpose of information gathering, fraud, or system access.

Professional social engineers continue to devise new and novel approaches for attacking their targets. We have in the past decade witnessed social engineering moving into a new era - coined *Social Engineering 2.*0 *(SE 2.0)* - which combines advanced information gathering techniques with the use of social media, email and SMS as attack vectors. An attack vector is a path or means by which the attacker can gain access to a computer or network server in order to deliver a malicious outcome. Attack vectors enable the attacker to exploit system vulnerabilities, including the human element. Most notoriously known is perhaps the phishing attack, of which most are familiar. They can be very generic in terms of design and approach, or they can be very targeted and personal, making them ever more complex and difficult to detect. This deception is what makes modern social engineering thrive.

Social engineering remains an essential component in a hacker's arsenal, simply because it is a convenient and easy delivery method for malware, which requires little technical knowledge, yet a great understanding of the human mind[4].

The convenience of utilising social engineering can be found in the emerging and ever-increasing use of social media networks and electronic means of communication. Employees at every level, from small businesses to large corporations, have become increasingly exposed to social engineering attacks, as malicious hackers – often characterised as *Blackhats* - are now moving into the personal sphere, using social media platforms, e.g. Facebook, in an effort to trick people into performing actions that are against their interests.

According to the open source statistics bureau, *Statistica*[5], more than 1.5 billion people are using Facebook; 900 million people use WhatsApp; and LinkedIn report having more than 414 million users globally[6].

Each social media account constitutes a point of entry for social engineers to perform their skills and execute attacks, in an effort to gain the desired information and/or system access. If you have ever had a *friend request* on Facebook from a person, who you are certain you do not know, then you have likely been playing a part in a social engineering attempt. However, the vast majority of all recorded social engineering 2.0 attacks remain largely generic, making them more conspicuous and easy to detect, even for the untrained. Examples include emails, which are either written in English or poorly translated into your native language, asking you to click on a link, open an attached file or fill in a form.

Figure 1: Statistics on Social Media

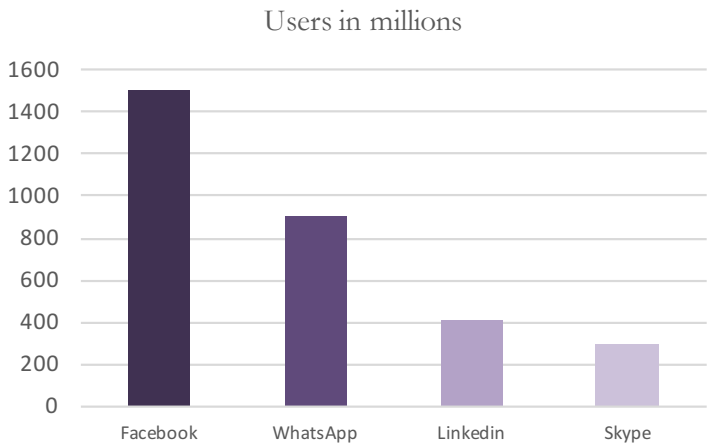

The motive for conducting cyber attacks may differ greatly, but can broadly be categorised as: political, economic fraud, prestige, state/industrial espionage, and cyber terrorism, of which the latter remains in terms of disruption rather than destruction[7], as no incidents involving a cyber attack has yet to directly cause any physical damages, i.e. hard kill capability.

Nonetheless, social engineering 2.0 is more complex than performing a series of cyber-attacks based a variety of improved attack vectors. It includes a vast array of new and renewed methods and technologies, namely analysis of information from open sources; social network analysis (SNA); psychological profiling (e.g. personality profiling with the purpose of identifying the most vulnerable individual within an organisation); memetics and sentiment analysis; and new trends in contextualising attacks on a one-to-one basis.[8]

Furthermore, advanced attempts of social engineering employ skills often acquired and utilised specifically in the intelligence sector, including human intelligence (HUMINT), signals intelligence (SIGINT) and open source intelligence (OSINT).

HUMINT refers to an intelligence discipline, which collects information via interpersonal contact from human sources. This directly relates to the classical understanding of social engineering, which encompasses the interaction between the social engineer and his/her target in an effort to elicitate information.

SIGINT refers to the collection and analysis of information from electronic signals, and includes both the collection of encrypted and unencrypted signals from the electromagnetic spectrum (EMS). For a social engineer this may include the analysis and interception of GSM, Wi-Fi or Bluetooth signals in an effort to either collect valuable information about a target or to directly identify a systemic vulnerability that could be used in a cyber attack.

OSINT relates to information collected from open sources. In a classical context, OSINT was collected from newspapers and old library records, while the modern OSINT analyst collects primarily from online sources, which include both indexed and unindexed information. As part of the OSINT collection process, social media intelligence (SOCMINT) can also be collected and analysed, with the purpose of evaluating targets and deciding on the appropriate deception tactics that is to be employed in an attack.

The information collected from the abovementioned intelligence disciplines is strategically analysed and tactically employed for targeting the social engineering attacks against specific individuals.

The modern social engineers therefore use a large and complex mix of different competences, covering a broad spectrum of sciences, including: technological, cyber-sociology, psychology, marketing and design. All of these combined are used in an effort to create a holistic, trustworthy and targeted cyber attack.

The evolution from traditional social engineering to social engineering 2.0 has been driven by the developments associated to our digital lives and society at large. The societal evolution of the last few decades has seen an increase in public exposure of information about individuals' personal details, making society more transparent, which is exploited by social engineers, who can collect valuable information about potential targets. This exposure is extensively reshaping the way people are conducting themselves - both personally and professionally.[9] For example, the information exposed on social networks is completely machine-readable and open to automatised web crawling processes and analysis, which can provide the social engineers insight into the interests of the targets, which in turn can be used to tailor the cyber attacks. Web *crawling* refers to the automatic collection of unstructured data or information from the Internet with the purpose if indexing it or for post-analysis.

Social engineering therefore becomes the enabling factor that paves the way to a large number of new infection scenarios and risks. This situation recently led Symantec to declare that standard defence systems, such as firewalls and antivirus software, are dead, since they do not provide the necessary security for the end users[10].

In this study on social engineering, we will explore how publicly accessible information on individuals can be exploited for constructing social engineering 2.0 attacks. The project of Social Engineering and Vulnerability Assessment Framework (SAVE) will address these issues and more, in an effort to understand not only the complex nature of social engineering 2.0, but also how vulnerable companies are and how this risk can be mitigated.

1.1 – The Purpose of the Study

In an effort to understand the phenomenon of social engineering 2.0 (SE 2.0), SAVE seeks to address the integration of human, technological and conceptual real-world issues, in a research field, which remains to be explored into greater depth on all levels, including the technical, psychological and user-experience aspects.

The overall aim and purpose of Project SAVE is to investigate the nature of the problem with social engineering attacks, and raise awareness among key Danish and international companies, who are either particularly exposed to the problem, or who constitute part of critical infrastructure (CI) in Denmark, where a successful attack can have severe consequences (cf. ch. 2 for examples).

1.1.1 – Objective

The objective of SAVE is three-fold and revolves around conducting an explorative study, with dissemination to relevant national and international actors, and providing recommendations on how to mitigate the associated risks of social engineering 2.0 attacks. The objective can be summarised as follows:

Explorative investigation of the problem with social engineering 2.0 by conducting simulated attacks against three Danish companies that either directly or indirectly constitute part of critical national infrastructure (CNI).

Raise awareness about the problem and danger of social engineering 2.0 among Danish and international companies and policy makers.

Provide recommendations on how to mitigate the associated risk by developing an evaluation framework to assess the organisational vulnerability of employees.

For achieving effective and useful results, a social vulnerability assessment (SVA) approach is utilised, which simulates attack patterns in order to measure the real vulnerability of the human security barrier in an organisation. An SVA approach is a new type of assessment, which proactively uses social engineering techniques to attack the enterprises, in an effort to evaluate their current social vulnerability level. Since it is a quite new area, there is still neither any established procedure nor suitable solution on the market able to simulate the whole cycle of a SE 2.0 attack for the purpose of conducting SVAs (cf. ch. 3 for more on the SVA).

In the following section, we will address social engineering as a phenomenon, which will give insight into the field, using the mind-set of an attacker to conduct a realistic social vulnerability assessment.

1.2 – Social Engineering

Social engineering (SE) is often considered the third oldest profession in the world, only succeeded by prostitution and espionage, as the first and second oldest, respectively. Nevertheless, it is often an abstract and misunderstood concept. The perception of SE commonly relates to *lying to people to get information, being a good actor,* or *tricking people*[11].

While this perception may contain some element of truth, it is far from the complete picture. A conceptual understanding of the basic concept of SE is therefore needed, in order to enhance our understanding of social engineering and further illustrate how it has evolved into its current state, employing vastly more advanced techniques than previously seen.

1.2.1 – Definition of Social Engineering

In an effort to get a basic understanding of social engineering, we break down the construction of the words, and separate 'social' from 'engineering' to understand each one respectively:

According to the Cambridge Dictionaries Online, the word *Social* is defined as: "*relating to activities in which you meet and spend time with other people...*", whereas the word *engineering* is defined as: *"the study of using scientific principles to design and build machines, structures [...] and buildings*"[12].

Combining 'social' and 'engineering' therefore leaves us with an understanding of the concept as an act, which through social interaction with people, and using scientific principles to design and construct this interaction, allows the practitioner to reach his/hers objective in a process as follows: (1) studying, (2) plan/design, (3) interact, (4) execute plan, (5) objective reached.

Kevin Mitnick (2003) is often perceived as the 'father' of modern social engineering. He defines social engineering as follows: "*Social Engineering uses influence and persuasion to deceive people by convincing them that the social engineer is someone he is not, or by manipulation. As a result, the social engineer is able to take advantage of people to obtain information with or without the use of technology*"[13].

Christopher Hadnagy (2010), a professional practitioner of social engineering, defines social engineering as: "*The act of manipulating a person to take an action that may or may not be in the [target's] best interest*[14].

Social engineering thus becomes a collection of skills mixed from the scientific realms of psychology, sociology, anthropology, social sciences and information technology, which are combined to construct a ruse that allows the social engineer to manipulate the target.

The following two sections will respectively cover (1) the human factor, and (2) the cycle of a traditional social engineering attack in an effort to elaborate further on the subject.

1.2.2 – The Human Factor

As established in the previous section, social engineering is fundamentally the art of exploiting the human factor, in an effort to compromise an organisation's security. Governments and private corporations can spend an indefinite amount of resources on security, yet social engineering remains a threat and a greater security risk than ever: The reason being that social engineering exploits individuals, who have direct access to the information the attacker is seeking. Rather than attempting to comprise the systemic

security barriers of the organisation, the social engineer goes directly to the source – the user.[15]

Albert Einstein is quoted as saying: "*Only two things are infinite; the universe and human stupidity, and I'm not sure about the former*"[16]. While people who fall subject to a social engineering attack might look back at the episode with a sense of foolishness - stupidity is not the reason behind the success of a social engineering attack. Rather, people fall victim to their own honesty, trust in others, and desire to help, which is exactly what a social engineer exploits in methods and tactics designed to put their victims at ease[17].

In essence, the social engineer simply facilitates the options for targets to be helpful and desirably leaves them with a feeling of having been so. This is why, in most cases, successful social engineers have strong people skills and are often charming, polite and likeable - all social traits needed for quickly establishing trust and rapport - leaving the human factor as the most vulnerable part of any security setup[18].

The human factor is essential in cyber security. Employees are not only the target of social engineering attacks because they are the ones who have access to the desired information; they also act as the first line of defence in mitigating the risk of an attack. That is the reason why many companies seek to educate and train employees on the subject of SE, in an effort to teach them how to identify and report suspected threats and attempts of social engineering attacks. In the following section we will cover conventional social engineering methods to get a basic understanding of the processes that takes place when a social engineer constructs attacks.

1.2.3 – Conventional Social Engineering (cSE) Methods

To further deepen our understanding of the concept of social engineering, we will review the traditional social engineering methods, which we have coined *conventional social engineering* (cSE).

Conventional social engineering refers to the methods, which have traditionally been used to compromise a target's physical or information security. cSE follows a simple linear, yet effective, pattern which consists of four phases: (1) Reconnaissance, (2) Hook, (3) Play and (4) Exit. The pattern can be summarised as illustrated in figure 2:

Figure 2: Process of a Conventional Social Engineering Attack

1.2.3.1 – Reconnaissance

Reconnaissance relates to the information gathering process, where research of the target is conducted, with the purpose of analysing the subject(s) and concludes by selecting the most desirable approach, based on the results of the analysis. For cSE, examples include *dumpster diving*, which is the collection of information from the trashed papers of a target[19], but can also include shadowing targets in an effort to understand their daily life in an effort to understand their interests, routines or habits.

Additionally, methods such as *shoulder surfing*, which is basically looking the targets over the shoulder as they are working, could be utilised in order to obtain information that can be used at a later stage in the attack, or it could simply be making a phone call to an employee, in an effort to gain insight into the *jargon* used at the company or in that particular line of business.

The reconnaissance phase is therefore the gathering of essential information about a target, with the purpose of understanding the target so that a plan can be devised for the next step, which is relationship development.

1.2.3.2 – Relationship Development

The 'hook' refers to relationship development with the target. Examples include interaction with the target, either in person or over the phone. Typically, some form of *pretexting* is utilised. Pretexting refers to pretending to be someone, who by the target is considered a trusted individual, in an effort to obtain information or build trust as part of an attack. Pretexting can be used to convince a target that the attacker works for the company, with the purpose of establishing a relationship that can lead to the divulgement of secret, sensitive or otherwise unauthorised information.

The relationship development phase is critical for the social engineer. Trust and rapport is built during this phase, and it typically relies on the information gathered from the reconnaissance phase. The attacker's use of deception tactics can greatly vary depending on the target at hand; examples include a collegial approach, where trust is established via common challenges in

the work place or via an authoritative approach, where trust is established via the deception of authority.

1.2.3.3 – Exploitation

The exploitation phase is essentially the attack phase, which can utilise a vast array of attack vectors. In the exploitation phase, the attacker attempts to *animate* the target to conduct a specific action that will benefit the attacker. For conventional social engineering attacks this could include *tailgating* into the building, by following employees as they access the building, thus compromising the organisation's physical security barrier. Obtaining access to the offices could potentially lead to sensitive information. Another example could be pretexting over the phone, with the purpose of getting the target to reveal otherwise unobtainable information that compromises the target's security, e.g. getting credentials by pretending to call from the IT-department, and thereby animating the target to reveal sensitive information.

1.2.3.4 – Exit

The exit phase refers to closing the communication with a target, after the social engineer has successfully obtained the desired information and/or system access, leaving the target unsuspicious of what has occurred - and often with a feeling of having done something good. The exit phase is considered an optional phase, depending on the needs of the attacker. If it is an on-going attack, the attacker might want the keep a good relationship with the target in an effort to maintain the opportunity to continue an attack cycle, or even start a new attack at a later point in time. However, if the attacker has obtained the necessary information and no longer requires access to the target, then the attack cycle ends and communication with the target stops.

In the following section, we will address social engineering 2.0, in an effort to understand how it differentiates from conventional social engineering.

1.3 – Social Engineering 2.0

In the context of security, social engineering has typically been understood as the art of *getting someone to do something they would not otherwise do, through psychological manipulation.* Thus social engineering remains a confidence trick with the purpose of gaining access to otherwise unobtainable information, through human interaction. However, conventional social

engineering has undergone an evolutionary process, following societal developments and trends, such as online trends and widespread use of social media, which has become one of the primary sources of information for the modern social engineer. Essentially, social engineering can be compared to the art of *elicitation*, where the attacks are conducted in person. Social engineering 2.0 attacks differentiate by being executed online - using email and social media, or any other platform, which can electronically facilitate the interaction between attacker and target.

Social engineering 2.0 involves many renewed methods from conventional social engineering, but differentiates by employing innovative techniques that include the use of cyber sociology, advanced marketing methods, and psychological manipulation through various communication channels, with the sole purpose of gaining access to otherwise restricted information or systems. The information gathering phase of a strategic SE attack has been vastly improved through systematic collection of open source intelligence (OSINT) on targets, Social Network Analysis (SNA) of employees as well as psychological profiling of individuals via sentiment analysis of social media content, with the purpose of evaluating how susceptible potential targets are to attacks, e.g. phishing attacks.

Figure 3: In-between Security Fields

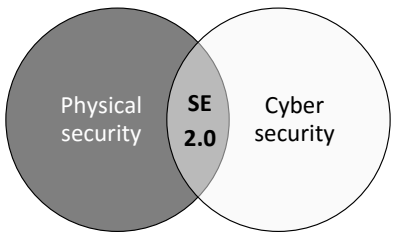

The utilised attack vectors in SE 2.0 have been greatly improved as well, and encompasses the possibility of attacking via email, SMS, USB, online chat and social media networks, thus increasing the complexity and level of sophistication, compared to cSE, by transcending the borders between the physical and virtual layers of security as illustrated in figure 3.

Although SE 2.0 utilises the same baseline pattern as with conventional social engineering, namely following the process of reconnaissance, hook, exploit and exit (cf. figure 2), SE 2.0 is more complex and dynamic as it has

more phases, resembling other advanced cyber attacks. Figure 4 illustrates this and takes us through the cycle of a SE 2.0 attack, illustrating the complexity of executing an advanced SE 2.0 attack and the many steps that is often needed to successfully exfiltrate the desired information from the target.

Figure 4: Process of a Social Engineering 2.0 Attack

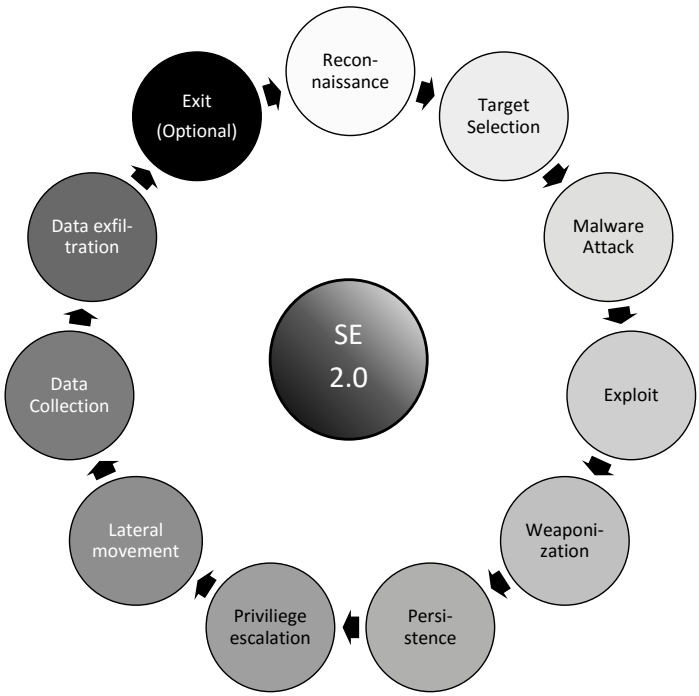

One noteworthy difference between the SE 2.0 cycle and the cSE process is the technical knowledge required to successfully execute a SE 2.0 attack. As illustrated in figure 4, the technical understanding of each step in the cycle is needed, whereas a cSE attack can be conducted with very little or no technical knowledge as it focuses on the human interaction.

However, in the following we will address three of the phases included in figure 4, and briefly discuss how these SE 2.0 phases differentiate from cSE. These three phases are relevant for this study, as they have been explored in-depth. Although the remaining phases are relevant to understand, the focus of the human factor is only represented in the first three phases. Ad-

ditionally, it will also become clearer how SE 2.0 differentiates from cSE by covering the same phases as covered in section 1.2.3 on cSE methods.

The phases covered in the following sections are: (1) the reconnaissance phase, (2) the target selection phase, and (3) the attack phase. The exit phase will not be addressed, as it remains fully unexplored in this study.

1.3.1 – Reconnaissance Phase

The reconnaissance phase in a social engineering 2.0 attack is critical for a successful attack, as it provides the baseline of information gathered on potential targets and provides the necessary insight into which deception tactics should applied in an effort to increase the success of the attacks and furthermore contributes to the identification of which individuals are the most susceptible individuals or valuable targets within an organisation.

The reconnaissance phase heavily relies on methods deriving from the field of open source intelligence (OSINT). Methods include crawling of available information from the Internet, which are then structured, filtered and analysed with the purpose of uncovering specific pieces of information relevant for the social engineer to devise a targeted attack.

The information that a social engineer will attempt to gain from open sources is often related to the company as a whole, including official documents available online in an effort to understand the organisational structure, the sector the company operates within, current business partners, investors, subcontractors and competitors. However, the reconnaissance phase is also focused on gathering information about the individual employees working at the company, spanning from their corporate email addresses to personal information available on social media networks, such as Facebook, Twitter or LinkedIn.

Automatising certain aspects of the OSINT gathering process can be utilised with the purpose of effectively uncovering new and potentially valuable information from the targets that would otherwise be too time-consuming to manually uncover and/or analyse. This could for example be *crawling* of documents (an automated gathering of specific information) from the corporate website, which may include documents that would be considered sensitive, but not easily available from the corporate website itself, e.g. to see if confidential information or credentials are leaked by mistake, which in turn could be utilised in a SE 2.0 attack.

1.3.2 – Target Selection

The reconnaissance phase results in a target selection, based on an evaluation of the potential targets. Depending on the need of the attacker, the target group is typically narrowed down to include only the ones, who are identified as being the most susceptible to SE 2.0 attacks, or alternatively to include targets who have privileged access to the desired information. Targets with privileged access to vital information are classified as *information gatekeepers*, and have either through their role in the organisation e.g. HR, legal, upper management, or through systemic access, e.g. the IT and security departments, access to information or data that the attacker wishes to obtain.

Which targets the social engineer select partially depends on the level of sophistication of the attack, and partially on the organisational structure and size of the organisation. Generic phishing emails will typically be sent to the entire organisation, whereas a targeted spear-phishing attack may only be designed for one or two employees that are deemed susceptible, or where the reconnaissance phase has established that there is ground for conducting attacks with a high level of deception.

1.3.3 – Attack Phase

When the targets have been selected, the initiation of contact and execution of the attacks can be initiated. The level of sophistication is often dependent on the importance of the target; the more important the target, the more effort is put into tailoring the attack, in an effort to maximise impact and enhance the chances of a successful SE attack. The reason being that the attacker might only have *one* opportunity of getting it right, and if the target is of vital importance, the attacker(s) can spend months trying to gain enough valuable information before evaluating the attack options. Furthermore, attacks can often be very elaborate, extending over longer periods of time, and are naturally covert in nature.

In the attack phase the social engineer can select one or more attack vectors. As briefly touched upon in the introduction, an attack vector is a mean by which the hacker can gain access to a computer or network server in order to deliver a payload to gain unauthorised access to the system. Attack vectors therefore enable the attacker to exploit system vulnerabilities, and for SE 2.0 this includes phishing emails, SMS, USB, to name a few.

Multiple attack vectors can be used against the same target - and at the same time - to increase the chances of a successful attack. In this sense, an attack resembles an asymmetrical (cyber) warfare situation, with the target attempting to defend itself from unknown attackers, and the attackers attempting to identify structural weaknesses within the human security layer of the designated target. However, employing multiple attack vectors on the same target might raise red flags with the target, and thus be counterproductive for the attacker.

In summary, social engineering 2.0 can be defined as the delivery method and deceptive layer of an advanced cyber attack, which exploits the human element, by utilising one or more attack vectors, with the purpose of compromising otherwise restricted information and/or gain physical access.

1.4 – Attacking Critical Infrastructure

As covered in section 1.1, this study on social engineering relates to attacking and compromising critical national infrastructure (CNI) with the purpose of uncovering and assessing current vulnerabilities in Danish companies that either directly or indirectly constitutes part of critical infrastructure. However, it is out of scope for this study to address the definition of critical infrastructure, which constitutes a study in itself. We have thus delimited it to rely on the definition by the European Council (EC):

[A]n asset, system or part thereof located in member states that is essential for the maintenance of vital societal functions, health, safety, security, economic or social well-being of people, and the disruption or destruction of which would have a significant impact on a member state as a results of the failure to maintain those functions.[20]

The only addition to the above definition is that we consider subcontractors to critical infrastructure as possible targets as well, yet only to the extent of subcontractors that have access to vital information about critical infrastructure in Denmark.

In conclusion, future reference to critical infrastructure thus constitutes either CNI itself or partners/subcontractors who have access to vital information about CNI - or can facilitate attacks directly at CNI in Denmark.

1.5 – Delimitation

Project SAVE has inherent structural limitations that have affected the scope of the project. This section will briefly cover the aspects of SE 2.0 that we will not be addressing, as well as other inherent limitations. The purpose of highlighting unexplored aspects of the project is to further illustrate and explain the actual aspects covered:

- In accordance with the project description, it is not the intention of this study to investigate the history or development of the phenomenon of social engineering prior to its current state, with the exception of when it contributes to the understanding of SE 2.0.

- The project has been subject to structural boundaries, i.e. a timeframe of eight months, limiting the possibility of conducting social vulnerability assessments of more targets during the field trial testing, where the simulated attacks took place. Multiple target assessments could have given a more in-depth understanding of SE 2.0 and the behaviour and weaknesses related to the human factor of cyber security as well contribute to more elaborate datasets and more precise results.

- Inherent limitations from legal and ethical considerations have also delimited the scope of the project in terms of utilising more aggressive reconnaissance and attack methods, which would otherwise have facilitated a closer reflection and representation of real-life SE 2.0 attacks. These limitations are covered in chapter 3, and have directly affected the use of specific methods in the reconnaissance phase, specifically pertaining to the collection of signals intelligence, as well as in the attack phase, in relation to the use of advanced malware tailored to the organisations that took part in this study.

- It is out of scope for this project to investigate and explore the underlying psychology behind what makes people fall victim to social engineering attacks; rather it is our goal to demonstrate that new methods can be applied to compromise organisations - critical infrastructure or otherwise - all companies and governments can fall victim to social engineering.

- It is not the intention of the project to develop new SE 2.0 methods, albeit some novel approaches have been the result of Project SAVE,

particularly relating to the OSINT methods applied in the reconnaissance phase.

- For the social vulnerability assessment carried out in the field trials, the author of the study has taken the point of view of an offensive attacker, in an effort to understand the attacker's perspective and create the most realistic attack pattern.

- It is not the purpose of this study to uncover the causes for conducting cyber attacks, including those of SE 2.0, though we recognise the various reasons covered in the introduction of this publication.

The overall aim and objective of the study is therefore to conduct an explorative study of the phenomenon of social engineering 2.0 by performing real-life, simulated SE 2.0 attacks against three Danish companies, raise awareness among key organisations and actors, and to provide a framework for mitigating the associated risks of SE 2.0 attacks on the basis of our findings.

Chapter 2:
Cases of Social Engineering

2. Cases of Social Engineering

In this chapter, we do not seek to engage in the discussion on cyber warfare, nor will state-sponsored cyber attacks be covered. Rather, we seek to uncover empirical evidence, which can put the concept of social engineering into perspective in relation to compromising critical infrastructure, and thereby support the notion of the methods applied in the social vulnerability assessment (SVA) conducted in this study (cf. chapter 3).

The essence of this chapter is to underline the severity of the applied methods, and how social engineering 2.0 constitutes a threat to CNI, and equally so to stress how effortlessly some of these methods can be utilised, which is a defining factor when assessing social engineering as a method for attacking; the easier it is to apply the method, the more common and frequent the attacks become, as more and more attackers apply SE to their repertoire.

The cases covered in this chapter include: (1) Attack on Ukrainian power grid, (2) Attack on Kiev airport, (3) Compromising a hospital, (4) Attacking the U.S. Department of Justice (DoJ), (5) The 2013 Target breach, and (6) Accidental Insider Attack using social media.

Each case will illustrate how various elements of social engineering 2.0 have been utilised and how it constitutes a threat for CNI, thereby providing the underlying basis and need for further investigation of the phenomenon of social engineering.

2.1 – Attack on Ukrainian Power Grid

2.1.1 – Introduction to the Incident

On December 23, 2015, the Prykarpattyaoblenergo electric utility in Ivano-Frankivsk Oblask, a region in Western Ukraine, reported that the power was out[21]. Of the 24 regions in Ukraine, up to 8 regions were affected by the power outage, lasting 6 hours and affecting more than 80,000 people, caused by a sophisticated and well-coordinated cyber attack[22]. It is believed to stem from the pro-Russian and state-sponsored group known as the *Sandworm Team,* whose interests are closely aligned with the Russian government. The attack has been believed to be ongoing since March, 2015, and culminated

in December, 2015. Experts have described the incident as the first known power outage caused by a cyber attack[23].

2.1.2 – Applied Methods

The applied methods, as uncovered by the Wired Magazine, points to an advanced and highly coordinated triple-tier attack, consisting of the following three elements:

1. A spear-phishing attack that allowed the hackers access to the operator's system, by using the malware known as *BlackEnergy3* as the payload, which was the source for the power outage caused by directing the industrial control systems (ICS) to disconnect power substations.

2. The use of *KillDisk*, which is a wiper virus that overwrites data in essential system files, causing computers to crash without the possibility of doing a reboot, as the virus overwrites the master boot record.

3. A Telephone Denial of Service (TDoS) attack, which essentially floods the centre's phone systems with calls in an effort to maintain disruption of phone services.

BlackEnergy3 was most likely used to manipulate systems to indicate that power was back in regions still affected by the attack. The high level of sophistication is not to be found in the technical aspects of the malware though, but instead in the application of several vectors for maintaining power disruption. It is important, however, to differentiate between destroying a power grid and merely causing temporary disruption, although they both cause severe national problems.

2.1.3 – Relevance to Social Engineering 2.0

The attack employed spear-phishing emails for delivering the BlackEnergy3 malware, which was executed on the target's system. The Ukrainian security company CyS Centrum has published screenshots of emails containing the BlackEnergy malware, which shows an Excel sheet.

CyS Centrum reports that the BlackEnergy campaigns were constructed and designed by attackers, who spoofed the sender address, so that it appeared to belong to RADA - the Ukrainian parliament. The document attached in the spear-phishing emails contained text which attempted to convince the target to run the macro in the attached Excel sheet[24].

This is an example of an attack on critical infrastructure, which employs social engineering methods, rather than exploiting vulnerabilities in software or online web applications. If the recipient of the spear-phishing email allows the macro to run, then it will download and execute the BlackEnergy malware and instantly infect the user's system.

2.2 – Attack on Kiev Airport

2.2.1 – Introduction to the Incident

In mid-January, 2016, Boryspil Airport in Kiev was subject to a cyber attack that utilised the same methods as the attack on the Ukrainian power grid, namely using the BlackEnergy trojan[25].

2.2.2 – Applied Methods

The malware was sent using a spear-phishing email with the attached Excel sheet, which if executed would run a macro and download the BlackEnergy malware. The infected workstation at the Boryspil Airport was connected to the airport's main IT network, which included the air traffic control centre. According to Wes Widner, director of threat intelligence and machine learning at the cyber security company Norse, targeting an airport's IT network can potentially lead to lasting damage, because airplanes are "fly-by-wire", and disruption that affects the air control system could lead to accidents during take-off or landing, or even cause mid-air collisions[26]. However, the incident was quickly discovered, and no harm was done.

2.2.3 – Relevance to Social Engineering 2.0

As with the former example, this illustrates that a simple spear-phishing attack and an advanced payload can compromise critical infrastructure, by using SE 2.0 to convince and animate the recipient to execute the attached file and thereby leveraging the cyber security of an airport.

2.3 – Compromising a Hospital

2.3.1 – Introduction to the Incident

On February 2016, Sergey Lozhkin, Expert at Kaspersky Lab, demonstrated in a real-life environment how he could hack a hospital and gain access to pre-determined and fictitious personal data, by using nothing more than his personal computer. The hack gave him access to personal information

as well as access to medical equipment hooked onto the network of the hospital, including a tomographic scanner[27]. Lozhkin used a similar approach as the vulnerability assessment employed in the SAVE study, whereupon a real-life field trial testing is employed to discover otherwise unknown vulnerabilities - be it social or cyber vulnerabilities.

2.3.2 – Applied Methods

The first method applied was an unsuccessful attempt, by remotely trying to access medical equipment connected to the Internet, using the Internet of Things (IoT) search engine, known as *Shodan*. The next approach was going on-location and cracking the Wi-Fi network key to gain access, which ultimately leveraged the security of the hospital and gave him access to the desired information.

2.3.3 – Relevance to Social Engineering 2.0

An important part of a SE 2.0 attack is the reconnaissance phase - gathering and assessing available information that can facilitate or build the foundation for a targeted attack. In this particular case, the attacker started with remotely assessing systems and devices that the hospital had connected to the Internet, and then moved on to an on-location approach, hacking the Wi-Fi from his car, in an attempt to compromise the hospital's cyber security. Though this was a test-setup, it clearly illustrates how easily one could gain access from an unauthorised computer.

Without segmentation of the networks, the attacker was able to identify the desired access point (AP), crack the network key, and move from there to extract, manipulate or delete any of the data. The data could either be what he was looking for or it could be vital information used in other scenarios: ranging from other SE attacks to extortion, based on the identified individuals in the hospital's records.

2.4 – Attacking the U.S. Department of Justice (DoJ)

2.4.1 – Introduction to the Incident

On February 7, 2016, a hacker with the twitter handle *@DotGovs* was in contact with the online tech news magazine *Motherboard*, claiming to have hacked the U.S. Department of Justice (DoJ) by using social engineering. The hacker later followed up on his claim by releasing personal details (e.g.

employee name, title, phone number, email, country) of 9,000 employees from the Department of Homeland Security (DHS)[28].

At a later point, he released details on an additional 20,000 employees at the Federal Bureau of Investigation (FBI), and claimed to have had access to 1TB of data, of which he extracted 200GB, which included detailed forensic reports. The hacker is possibly driven by a political cause, based on previously tweeted content[29].

2.4.2 – Applied Methods

The hacker compromised the email account of a DoJ employee, and attempted to log on the DoJ web portal, but was unsuccessful in his attempt. He then moved on to attempt a conventional social engineering attack by calling the DoJ, in an effort to manipulate them into providing him with access to the portal:

"*So I called up, told them I was new and I didn't understand how to get past* [the portal]," the hacker told Motherboard. "*They asked if I had a token code, I said no, they said that's fine—just use our one.*"[30]

He then proceeded to logon to the portal, using the same credentials as with the DoJ email account he had compromised, and thereby gained access to the DoJ web portal. In an effort to prove that he had access to the portal, besides the 200GB of collected data, he took a screenshot of his access, which was provided to Motherboard. The screenshot shows the compromised user being logged in with full access to navigate the webportal.

2.4.3 – Relevance to Social Engineering 2.0

This is an interesting case, as it employs a simple *vishing* method - voice phishing method - that basically consists of elicitating information from a target over the phone. Although the hacker has not publicised how he initially gained access to the DoJ email account, it is well-known that leaked data can be found online, which can quickly be accessed and therefrom a vishing attack can be utilised to escalate access, as this case clearly illustrates.

2.5 – The 2013 Target Breach

2.5.1 – Introduction to the Incident

In 2013, attackers managed to extract the credit card details of 40 million customers, as well as the names and addresses of 70 million customers from the customer database of the American retail chain *Target*.

2.5.2 – Applied Methods

The attackers initiated the attack by sending malware-infected phishing emails to one of Target's subcontractors, who handle heating, ventilation and air conditioning. The subcontractor's IT-security measures were not sufficiently sophisticated to detect the malware-infected email and the attackers therefore managed to compromise their IT security and gain access.[31]

The subcontractor's network was connected to Target's systems for providing electronic billing services, and the attackers used this connection to hop onto Target's payment network, where they extracted credit card details of millions of users.[32]

Some experts believe that the subcontractor was not targeted directly by the attacker. It is more likely that the phishing campaign was sent to a wide range of companies and once some of these were infected, the attackers determined whether any of them were of interest.[33]

2.5.3 – Relevance to Social Engineering 2.0

Besides the use of phishing emails as an attack vector, the attackers may have employed the use of open source intelligence (OSINT) for providing the basis for the tactics employed in the attack. In a post-investigation conducted it has become known that a list of suppliers was available online through OSINT, so it would have been an easy task for the attackers to conclude that they might be able to leverage the subcontractor's system in an effort to launch an attack on Target.

Furthermore, a lot of information about Target's internal network structure was available through OSINT – all which would be very useful for an attacker[34]. The cost of the breach has been estimated to be at least $252 million[35].

2.6 – Accidental Insider Attack using Social Media

2.6.1 – Introduction to Incident

This example does not include any particular incident, but rather a severe breach of security, caused by something as simple as a *selfie*. It might seem common sense, that when one works at a power plant or other sectors that are part of critical national infrastructure, it is not in the interest of the enterprise to have photos of the SCADA systems (Supervision, Control and Acquisition of Data) leaked on social media networks, such as Facebook, Instagram and Twitter, where hackers can easily obtain these as part of their reconnaissance of potential targets. This constitutes a severe security breach for these organisations, as helps a potential attacker to identify the systems that are in use.

In February, 2016, Panda Security wrote about one such incident, where an employee from a power plant took a photo of himself with the SCADA systems in the background running on his computer. The SCADA systems are used for monitoring critical components in e.g. a power plant. Additionally, the photo shows schedules hanging on the wall as well as a dozen other documents, which could arguably be internal documents not intended for public exposure on social media[36].

2.6.2 – Relevance to Social Engineering 2.0

In the above example, insights into the user's identity, interests, line of work, what system they use at his work place, and so forth, constitute the basis for any good SE 2.0 attack and is part of the reconnaissance phase.

This is an issue that has been widely discussed in dedicated IT security forums, with the focus on the sharing of selfies with valuable information on SCADA systems. Selfies, however, are not the only information IT security experts have discovered; publicly available virtual tours of control rooms from critical infrastructure have also been discovered.

The German security expert, Ralph Langner, uncovered an image of the Natanz nuclear plant in Iran, which was distributed by president Ahmadinejad's own press office[37]. Langer commented that:

"*While no Western plant manager would have cleared such photographic material for publication, Iran didn't seem to bother to hide [it] from the media.*"

Panda Security has stated that the released photo from the Iranian nuclear plant was used in the infamous Stuxnet malware attack of 2007, which arguably constitutes the most targeted and well-coordinated cyber attack to date, which used no less than four zero-day exploits - an unprecedented number of zero-day exploits for a single attack[38]. This example illustrates the power of open source intelligence and the lack of insight into the security threat a single photo can pose, when in the hands of a social engineer.

2.7 – Summary of Social Engineering Cases

The instances covered in this chapter serves as exemplary cases where social engineering 2.0 has been employed in cyber attacks on critical infrastructure or other larger corporations.

While many companies take precautionary measures to protect their networks, both in form of systemic countermeasures and alert systems, e.g. Intrusion Detection Systems (IDS), as well as awareness training of employees, dedicated attackers will continue to test the faith of both the systemic and human element of cyber security, in an effort to gain access or deprive the company from accessing their data.

The use of phishing emails, tailored malware, and open source intelligence for the reconnaissance of companies and individuals, are all aspects of social engineering, which we will investigate further in the following chapter. Here we will focus on the Social Vulnerability Assessment approach employed to test three Danish companies that have taken part in this study.

Chapter 3:
Social Vulnerability Assessment

3. Social Vulnerability Assessment

As part of the study on the phenomenon of social engineering 2.0, the SAVE project has incorporated a *field trial testing* as part of the social vulnerability assessment, which is intended to simulate actual attacks on three Danish companies. *Simulated attacks* refer to conducting actual attacks in a real-life setting, but without compromising critical information. The field trial testing is carried out in cooperation with the companies involved, and the main goal is to penetrate the human barrier of their cyber security. The employees from each company involved in the field trial testing have been pre-selected by the companies themselves, which has delimited our focus to only those, as part of the agreement with each respective party.

Although we have maintained full autonomy on the design and execution of the attacks, each company has been involved in the process to validate the utilised attack vectors. In respect of, and upon request, we have kept the identity of the involved companies anonymous. This chapter will focus on the matters relating to the boundaries set forth by ethical considerations and legal limitations, as well as the various methods applied for the field trial testing as part of the social vulnerability assessment.

3.1 – Legal Limitations

The first task in the study has been to investigate the legal complications involved with certain aspects of carrying out simulated SE 2.0 attacks. The *pretexting* and *method acting* - where you pretend to be someone you are not, in order to gain physical access or extract information from targets - has been deemed in conflict with national criminal law, should the company and/or individual exists in real life.

Originally, we were planning to pretext someone from an established news agency, with the purpose of being invited into the office of the companies involved in the study with the purpose of planting GSM audio-visual surveillance equipment for the reconnaissance phase. However, this could potentially be in conflict with national criminal law, despite having formal legal agreements with the company we wished to impose as being a representative of. Legally speaking, Danish criminal law supersedes any written or verbal agreements, thus making any agreement with a company void.

This came to our attention after receiving legal counsel on the subject, and we were advised not to go forward with some of the attacks, which consequently resulted in a minor setback in terms of planned attack vectors.

As a concluding remark on the legal limitations, which we have operated under in this study, it is crucial to stress that a rouge hacker group would not honour the law, and would hence execute these cyber attacks without any regard to ethical considerations or legal implications. This is an inherent limitation of the study which needs to be recognised, as this *can* have affected the overall results of the study if compared to real cyber attacks in the same context.

3.2 – Ethical Considerations

The ethical aspect of using simulated attacks, when conducting the social vulnerability assessment in the field trials, restricts the scope of the utilised attack vectors, as well as the measures for handling crawled data and information from the reconnaissance phase.

Specifically, for the reconnaissance phase, we have implemented end-to-end privacy measures to maintain the confidence of the involved organisations. Although open source intelligence (OSINT) and social media intelligence (SOCMINT) derives from publicly available information that can be collected and analysed by anyone with an Internet connection, certain ethical aspects have been incorporated as to protect personal relations, third parties and irrelevant personal information, which goes beyond the scope of this study.

As such, all information that has been collected on individuals have been processed anonymously, and for the most part, with automated processing, which entails that we have only seen the outcome of the analysis conducted (the end results) and not the individuals' actual published content from e.g. Facebook. All of our crawlers have been designed with privacy and anonymity in mind to protect the identity of the involved parties, both pre- and post-analysis. To make sure we maintained a strict anonymity policy, we decided to implement a complete systemic segmentation of data from and about the users and companies involved in the study. This meant that from the information gathering phase was initiated to the finished intelligence product was produced, the data underwent a process of anonymisation,

segmented on three systems, so that if any one system was compromised, it would be impossible to identify the users or companies involved in the study.

Additionally, we decided to implement policies for only including individuals, who were part of the designated target group, which the involved companies decided upon themselves. To accommodate this, we had to implement measures to make sure that information was aggregated in the reporting of the results to the involved companies, since the companies would be given insight into their own data after the social vulnerability assessment was completed. In an effort to accommodate this, we aggregated the results of the SVA, so that the company would not be able to single out one individual from the rest of the designated target group.

In relation to the attack phase, none of the attacks conducted in the study included the execution of malicious scripts on the targets' systems. Each script, site, web form or other was validated and verified by the Alexandra Institute, who will confirm that the scripts, sites or web forms did not contain any malicious code that might compromise information or gain unauthorised access - now or at a later point in time - to the target computer. Each company was additionally given the option of inspecting the scripts themselves, and we provided them the option of appointing an observer from their company, which could oversee the process and be involved in the development of the attack phase, to make sure that SVA was conducted in a manner in compliance with their respective company policies.

Furthermore, our ethical considerations include to not targeting *Bring your own device* (BYOD), as this would constitute the targeting of personal devices. We have also decided not to engage in contact with targets on social media networks, as this could lead to otherwise inappropriate conversations, e.g. by impersonating as a recruiter from an attractive competitor and offering the target a new job with a highly competitive salary package. However, we deemed this ethically unjustifiable due to unforeseen consequences for the targets in the aftermath of the social vulnerability assessment.

Finally, we have decided not to target individuals at third-party locations, e.g. cafés, simply because it goes beyond the scope of the project and would include surveillance and shadowing of targets, which is not ethically justifiable in this particular context.

The overall ethical considerations are therefore strongly focused on maintaining full anonymity at all time for the sake of the companies, their employees and the involved parties of Project SAVE. For all steps apply that names and other identifiable information, have been anonymised as to eliminate all professional complications it could otherwise impose on the people involved - now or in the future. No identifiable information will thus be published to any parties involved in the project, which would not even be a possibility as all records are anonymised and all identifiable information has been deleted from each respective system upon finalising the analysis of the data.

3.3 – Targets

To simplify the levels of a target, a social engineer will typically operate with various target classifications. The targets involved in Project SAVE are divided into three subcategories: (1) The target organisations, (2) the target groups, and (3) the target individuals, which will all be covered respectively.

Table 1: Target Classifications

Target	**Definition**
Target organisation	Company/Organisation
Target group	Office/Department
Target individual(s)	Specific employee(s)

3.3.1 – Target Organisations

Three organisations have agreed to take part in Project SAVE, of which one directly constitutes part of critical national infrastructure, and the other two are organisations that support critical infrastructure in Denmark with services and/or products. We have evaluated the two latter on the basis of whether the organisations have access to critical information relating to other companies that are part of CNI.

3.3.2 – Target Groups

The target groups have been selected in cooperation with the organisations involved. If we had been given full autonomy, all employees would have been selected to take part in the field trial testing. The target groups thus constitute pre-selected offices, groups and/or individuals that we have been allowed to test by the respective companies.

3.3.3 – Target Individuals

The target individuals constitute the specific employees who took part in the field trial testing. They ranged from 18-65 years of age, male and female, on every organisational- and employment level, as well as with differing academic, professional and social backgrounds. The target individuals are subject to full anonymity and only if they identify themselves as having taken part in the field trial testing, will others know of their involvement and/or whether they were animated to conduct a task they should not otherwise have done.

3.4 – Utilised Social Engineering 2.0 Methods

3.4.1 – Introduction

The following section covers the social engineering 2.0 methods utilised in the social vulnerability assessment conducted in this study. Figure 5 illustrates the cycle of a social engineering 2.0 attack as it progresses. This is an elaboration of the model covered in the introduction (cf. ch. 1), and differs by giving insight into the elaborate reconnaissance phase that has taken place in the study. Each of these phases is respectively covered in the following sections.

Figure 5: The Reconnaissance Phase in thee SE 2.0 Cycle

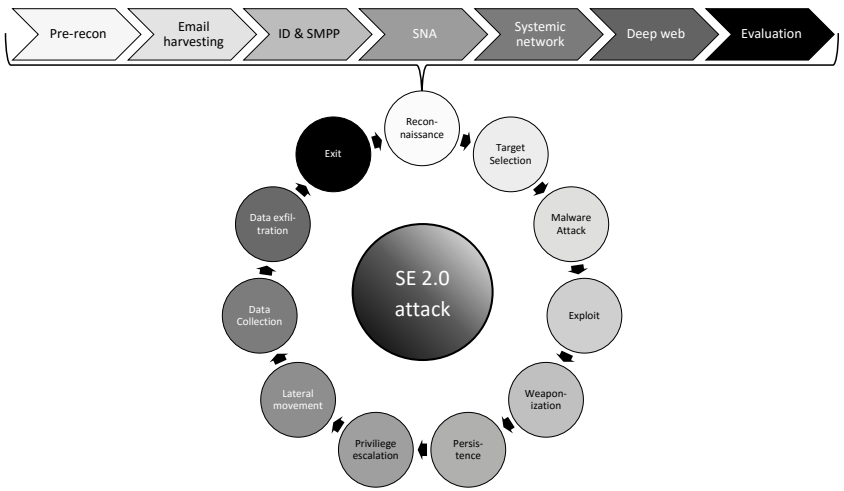

The following sections will cover the applied process of the social vulnerability assessment and also indicate which inherent limitations exist in regard to each respective phase. For the overall process of an attack, our general limitation has been the attack phase, from which we have been limited, because we were unable to execute advanced forms of malware.

This section is divided into each respective phase, starting with: (1) The reconnaissance phase, (2) the target selection phase, and finalised by (3) the attack phase.

3.4.2 – Phase One: Reconnaissance

The reconnaissance phase is a fundamental part of SE 2.0. It mostly relates to open source intelligence (OSINT) and social media intelligence (SOCMINT), but also utilises other advanced techniques for collection, processing and analysing data – all of which can create the basis for a targeted SE 2.0 attack. The following sections will cover the entire applied reconnaissance phase used on the target organisations in the social vulnerability assessment. The sections included are:

- Pre-reconnaissance
- Advanced Google searches
- Robot Exclusion Protocol
- Metadata analysis
- Systemic infrastructure analysis
- Email crawling
- ID of employees' social media accounts
- Sentiment analysis and personality profiling from social media
- Social Network Analysis (SNA)
- Deep web & Darknet Investigation

3.4.2.1 – Pre-Reconnaissance

The pre-reconnaissance phase requires no immediate preparation, as it relates to OSINT from traditional online sources. These sources includes, but are not limited to: (1) Physical publications, (2) online publications, (3) online articles, (3) the target's website, (4) citations from former and current employees.

The pre-reconnaissance phase merely serves the purpose of introducing the *target organisation* and/or *target individuals* to the attacker. The information acquired from the pre-reconnaissance does not necessarily have a direct

influence on the execution of a social engineering 2.0 attack, but instead gives the attacker insight into:

- The organisational structure
- Products and services
- Business model
- Current and potential partners
- Competitors
- Financial records
- Current and former employees

For a high-level targeted attack to be successful, it is important for the attacker to have a thorough understanding of the target organisation, as it increases the success rate of an attack. Often the attacker may only have one opportunity at a given attack vector, and while timing is crucial, doing the necessary reconnaissance is equally so – after all, 90 pct. of a targeted attack relies on the reconnaissance phase (Graves, Kimberly: 2010). In addition to collecting and disseminating OSINT on the target organisation, identification of known employees and their presence on social media will be carried out.

3.4.2.2 – Advanced Google Searches

Processing of available OSINT from indexed websites via advanced Google searches, employing a variety of search operators, reverse image lookups, and Google dorks for uncovering information about the targets are important steps.

Many sources exist in regard to advanced search operators, including Google's own documentation[39] and the MIT Library[40], and some of the Google dork queries applied have been found via the Google Hacking Database[41]. *Google dorking* refers to using advanced search parameters to return information that is difficult to find using simple search queries. It is sometimes referred to as *Google hacking* as it can be used to identify information that was not intended for public viewing, but which has not been subjected to adequate protection.

In relation to Google dorks, we required a specific one for identifying email addresses of current and former employees of the companies involved, though we were unable to find a suitable one via known resources. After several attempts, we discovered a Google dork of our own, which we used

for the advanced Google searches carried out in this study, and we also applied it in one of the scripts that was coded for the email crawling part of reconnaissance phase. In a post-test of the Google dork, conducted after finalising the project, we have determined that the dork is no longer applicable on the Google search engine, though a similar dork works equally well on the Bing search engine, which remains functional at the time of writing.

These methods contribute to a more thorough research process of the targets at hand. It is our experience that companies often are unaware of their *digital shadow* – the information that can be uncovered about them online – and as such, this information can become valuable for the social engineer in regard to designing an attack.

3.4.2.3 – Robot Exclusion Protocol

The Robot Exclusion Protocol (robots.txt) is a standard used by web crawlers[42], e.g. search engines, which specifies the information a search engine is *allowed* to collect from a website. As most search engines respect robots.txt, we have accessed the file directly from the target organisation's website with the purpose of uncovering potential sites/files that the targets did not want the search engines to index. This can be done by accessing the desired website, e.g. http://www.[site].dk/ and then add *robots.txt* after the forward slash, i.e. http://www.[site].dk/robots.txt.

3.4.2.4 – Metadata Analysis

The National Information Standards Organization (NISO) defines metadata as: "*structured information that describes [...] or otherwise makes it easier to retrieve, use or manage an information resource*"[43]. Metadata is furthermore often characterised as 'data about data'. It is the underlying information that supports the actual information and can contain various sorts of information, e.g. author name, operating system when the file was created and GPS coordinates, if available in the respective files.

To extract metadata from OSINT in relation to the three target organisations, we employed the use of a free, publicly available software solution known as *Fingerprinting Organisations with Collected Archives* (FOCA).

Essentially, FOCA scans popular search engines (Google, Bing and Exalead) for files relating to the web domain of the targets' websites. It then *crawls* the files from the Internet and then performs locally analyses them for metadata. As previously covered, web *crawling* refers to the automatic collection of

unstructured data from the Internet with the purpose if indexing it or for post-analysis. When conducting the analysis in FOCA, the results are presented in a structured, easily accessible manner, which allows for immediate evaluation. With all of the metadata details extracted from the files, FOCA furthermore matches information with the purpose of identifying, which documents have been created by the same team, including what servers/clients may be inferred from them[44].

3.4.2.5 – Systemic Infrastructure Analysis

To further deepen our understanding of the targets at hand, we decided to use Paterva's software package *Maltego*, which provides additional insight into the targets organisations. Maltego provides an overview of the systemic internet protocol (IP) infrastructure based on the web domain used by the target organisations, which can provide identification of people, exchange of information, DNS information, and additionally help us identify structural weaknesses within the information network.

Maltego essentially helps to identify additional metadata, employees, email addresses, related web domains, mail servers and documents available, as well as accounts on social media, e.g. Twitter handles, and provides a structured and accessible overview of the information gathered. In Maltego, the results of the analysis can be adjusted to endless number of levels, which gives the operator the possibility to consider whether targeted information collection is required, or whether as much information as possible will provide the best insight into the organisation. This allows for an attacker to conduct both quick and thorough assessments of a potential target organisation. Additionally, Maltego presents the results in a similar fashion as a social network analysis, where nodes are illustrated with icons of the respective information collected. This means that when, e.g. Twitter accounts are identified, a Twitter icon for each node is shown in the network, and the details about the account can be analysed further within Maltego. Finally, plenty of third party plugins exists, which can expand the usability of Maltego.

3.4.2.6 – Email Crawling

One of the most important information relating to the reconnaissance phase is the uncovering of the email addresses of potential targets. We have therefore developed two scripts, which scan the Internet for email addresses with a specified domain name (TLD). It utilises the Google search engine, from

which it crawls all results from both indexed websites and documents (incl. PDFs), based on the parameters: ([a-z+0-9+._-]@[Domain name].[ccTLD]).

The email crawling script was coded in Python, and it utilised a Google dork that we discovered in assessing already-existing dorks in an effort to provide the script with the best possible search pattern, prior to conduct a vast amount of searches. However, since the Google search engine uses anti-crawling measures to counter automated bots that are trying to crawl their content, we realised that we had to implement modules, which the script could use to emulate the natural, human behaviour of browsing a website. We therefore decided to implement the following:

- The module *selenium*, which emulates a human browsing experience on the website, as well as *Scrapy*, which makes it possible to quickly crawl information from many sites at the same time.
- We have furthermore utilised *user agents* (cf. Appendix A), which emulates browsers (e.g. Firefox), confusing the counter crawling algorithms to believe that the machine running is in fact a browser.
- Finally we implemented simple measures, such as *delays* in the crawling process, which were incorporated to emulate a more natural browsing experiencing in an effort to disguise the crawling bot from the anti-crawling measures implemented in the Google search engine.

The purpose is to collect as many relevant email addresses as possible from current and former employees at the three target organisations, as the email attack vector is the most utilised vector in social engineering attacks. Identifying email addresses is therefore crucial for the social engineer to direct the SE 2.0 attacks. However, more often than not, it is possible to identify the desired email addresses by simply knowing one email address from a desired target organisation. Some companies, for example, use the first two letters of the first name and the first two letters of the surname of an employee, along with the web domain of the company, to construct the employees corporate email addresses. By having a predictable email syntax, it makes it very easy for the attacker to reverse engineer the email addresses of other employees in the organisation, by simply knowing their names. One very easy way to acquire the names of employees is simply to look them up on the social media network LinkedIn, where a list of employees from the organisation that the attacker wishes to target, are likely to be represented.

Additionally, by knowing the syntax of the email addresses, the attacker can construct a simple script, which automatises the creation of every single possible outcome of an email address containing, e.g. four letters, if the attacker wanted to pursue with targeting the entire organisation. By doing so, the attacker makes sure that every possible combination of the email addresses will be subjected to the attack. This is mostly used in cases of spam emails used for either fraud or marketing purposes, but it can easily be applied for both generic and organisationally targeted social engineering attacks.

3.4.2.7 – ID of Employees' Social Media Accounts

As described in the former section, our email crawling script would crawl through all of the Google search engine results to identify every email address known that was associated with the web domains of the target organisations involved in the study. In an effort to verify whether or not the email address was still in use, the information could be correlated with popular social media networks, e.g. Facebook or LinkedIn, in an effort to verify whether or not the individual associated with the email was still employed at the target organisation. Additionally, this process would contribute to the identification of their social media accounts. In this study we decided upon identifying three social media networks of the target individuals: Facebook, LinkedIn and Twitter.

The task can either be performed manually, which can be tremendously time consuming, or it can be automated by using open-source scripts, e.g. *Scythe*, which scans popular social media networks for accounts relating to an email address. It is, however, our experience that very few people use their corporate email addresses for social media accounts, which makes it more difficult to identify social media accounts using the targets' corporate email addresses. This would instead require the identification of their private email addresses, which is beyond the scope of this study, unless it relates to the upper management level of the target organisations.

The purpose of identifying employees' social media accounts is two-fold: (1) It allows for phishing attacks to be conducting via the social media networks; and (2) it allows an attacker to identify whether information is available on their social media accounts, e.g. public posts on Facebook, which in turn can be analysed and used to provide information about the target - information that can be used in an attack scenario, which is covered in the following section.

3.4.2.8 – Sentiment Analysis and Personality Profiling

We have developed a script in Java, which crawls content from social media networks, specifically Facebook and Twitter, and then it analyses the content based on predefined parameters and instructions. The script conducts a sentiment analysis, which basically analyses the actual words used in posts and tweets. The sentiment analysis is based on whether positive, neutral or negative words are used in a sentence, and uses the *bag-of-words* approach, which is then used for making a simple – yet effective – personality profiling of target individuals, based on the content they publicly publish on their social media accounts. The bag-of-words approach entails having a long list of words, which are predefined as either being positive, neutral or negative. The method is very effective for conducting *quick-and-dirty* sentiment analyses, which is what we expect a social engineer or a hacking group to use. The words analysed would then be categorised into five sections, with the purpose of conducting a personality profiling of the targets. This method will be described later in this section.

First, however, we have to address inherent limitations that exist for this particular scientific method for conducting sentiment analyses, which needs to be recognised. Most notably is that the method entails that specific words are always positive or always negative, which is not always the case as the following will demonstrate. Consider the word '*torturously*', which can arguably be deemed a word with very negative connotations. However, in the context of the following sentence, the meaning changes from a negative to a positive: '*The chocolate was torturously good*'. In this context it is used to emphasise that something is extremely positive, rather than extremely negative. Our method did not take this into account, as this would require a tremendous effort to categorise each negative word that can be used to emphasise a positive one. We believe that an offensive social engineer would not consider this either, but would instead focus on gaining the best possible outcome in the shortest timeframe possible, in an effort to explore other reconnaissance methods, while still maintaining an interest in the baseline results of a simpler approach, rather than pursuing the most academically valid approach.

The purpose of the sentiment analysis is to categorise each employee into five different sections, based on the content they have publicly published on their social media network profiles. The method we applied for deciding upon the categorisation is the personality profiling principles known as the *Big Five* Framework. The framework provides five different categories,

which employees can be assigned to, based on the results of the sentiment analysis. The five categories are: (1) Openness, (2) conscientiousness, (3) extraversion, (4) agreeableness, and (5) neuroticism, as illustrated in figure 6. Current research in phishing studies indicate that two of the five categories are particular susceptible to phishing attacks, and have sufficient curiosity to have a higher chance of being actively be manipulated[45]. The categories in question are openness and neuroticism.

Figure 6: 'Big Five' Personality Profiling Model

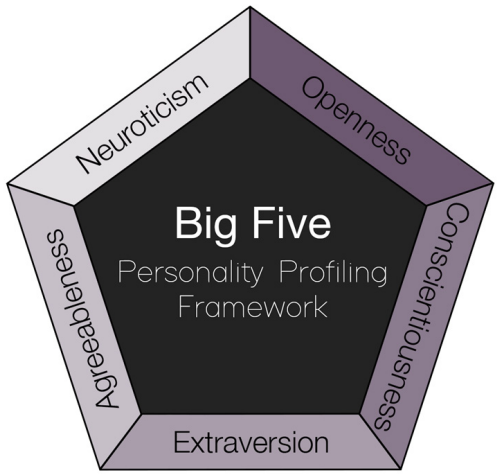

According to the paper "*Phishing, Personality traits and Facebook*", a high neuroticism score is correlated with how receptive an individual is to phishing attacks. Other factors, which are relevant in assessing a target, include the gender of the individual, as well as how much (quantity) content is shared, and how often (frequency) an individual shares posts. Additionally, we know from other studies conducted in the field that people who have a tendency to share many links are more likely to click on links themselves, phishing or otherwise (Schwartz & Eichstaedt et al.: 2013).

From an attacker's perspective, it therefore becomes relevant to analyse the results of these parameters of potential targets, in an effort to narrow the scope of the attack to only (or at least initially) include target individuals, who are deemed the most susceptible to social engineering attacks. From the paper "Phishing, Personality traits and Facebook", the authors note the following:

"*Our research shows that when using a prize phishing email, we find a strong correlation between gender and the response to the phishing email. In addition, we find that the neuroticism is the factor most correlated to responding to this email. Our study also found that people who score high on the openness factor tend to both post more information on Facebook as well as have less strict privacy settings, which may cause them to be susceptible to privacy attacks.*"[46]

The purpose of our sentiment analysis and personality profiling of target individuals can therefore be summarised to the narrowing of the scope of our attacks, in an effort to target individuals that are deemed more susceptible to social engineering attacks, based on the results of our analysis of their published content on social media networks.

The applied method for conducting our analysis is verified by academic research in the field, and is widely used in marketing research in relation to understanding specific market segments. Specifically the Big Five framework is the dominant personality profiling model applied for conducting research in this field[47].

3.4.2.9 – Social Network Analysis (SNA)

For conducting the social network analysis (SNA), we have developed a script that crawls relational data between employees, using social media as the point of departure. The script collects all available data about an employee's social network, and correlates it with other employees from the same organisation, which is outputted in a CSV file that can be imported into IBM i2 Analyst's Notebook - a powerful piece of software for conducting social network analysis (SNA). The software calculates the centrality of actors within the network and can visually illustrate how the constellation of the employee's social network. These actors within the organisational network of a company are often referred to as *nodes.*

It is important to stress that all connections crawled from social media networks that are not on predefined lists, provided to us by the respective target organisations in this study, will not be collected as to comply with the agreed upon ethical considerations.

Thus, based on the advanced algorithms of IBM i2 Analyst's Notebook, we are able to calculate and identify critical nodes – or target individuals - within the organisational social network of the target organisations, allowing us to identify individuals, who are highly interconnected and who holds

a professional position or high social status within the respective organisations. This allows for greater access to knowledge or other *information gatekeepers* – that is, employees who due to their function within an organisation hold a significant role with privileged access to information and/or system access, e.g. the IT department or human resources.

The purpose of the SNA is to identify structural weaknesses within the constellation of the social network of the target organisation (more specifically the target group, which is the group of individuals that we have been allowed to target with SE attacks). The goal of our social network analysis can therefore be divided into the following:

- Identification of information gatekeepers
- Identification of nodes within the network, who holds a privileged position, due their social role within the target group
- Identification of target individuals, who have one or more nodes in common, but who themselves are not connected directly

By assessing the structural properties of the employees within the organisational social network, we can start identifying central subgroups and, perhaps more importantly, the elite actors, which in the context of conducting SE 2.0 attacks consist of the information gatekeepers, who either directly have access to the desired information, or indirectly have access to the people who do.

Figure 7 exemplifies how larger cell subgroups can be identified, and how the cell core and cell periphery relates to the larger network.

Figure 7: Sub-Groups within a Network

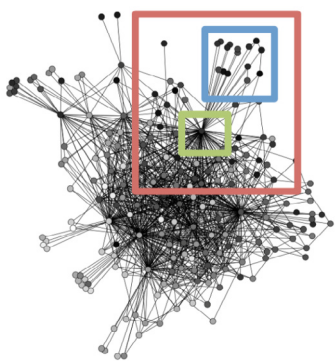

The network mapping software calculates various key metrics for the network and its nodes. These key metrics are concerned with centrality, including: *Degree centrality, Betweenness centrality, Closeness centrality,* and *Eigenvector measures.* They all contribute to identifying the key actors within the organisation - actors who are interesting for the social engineer (Please cf. Appendix B for further information on this matter).

In summary, we have developed a script that crawls relational data between employees from their public social media network accounts. The script collects all available data about an employee's network, correlates it with other employees within the same company and outputs a CSV file, which can be imported into IBM i2 Analyst's Notebook. This script can then visually illustrate how the social network of the employees are structured, which in turn is analysed by using the above methods, based on the algorithms of the toolset in IBM i2 Analyst's Notebook.

3.4.2.10 – Deep Web & Darknet Investigation

Deep web and darknet[48] investigation methods have been applied in an effort to uncover possible leaked information on unindexed sites or information sold on darknet market places. The deep web investigation includes using online services as Internet Archive's 'Wayback Machine', which indexes sites with the possibility for users to see how the sites looked weeks, months and even years ago[49].

Not *all* websites are stored on the Wayback Machine over the years, often only one or two examples per. However, as the WayBack Machine currently has 466 billion sites stored, there is a chance of finding some relevant information, which in turn can be used for the design of a SE 2.0 attack. Additional deep web methods include scouring sites like *Pastebin*, where people anonymously can post any text-based information they wish. This is a common place for leaked passwords to appear.

In an effort to uncover information about the target organisations from the Darknet, several darknet marketplaces have been searched for content, in an effort to uncover any information about the target organisations. Although, we did not expect to uncover any valuable results from the darknet investigation of the targets, since it would be rare to find information about the three target organisations that are taking part in this study. However, information that is being sold on darknet marketplaces is often valuable

information that can potentially constitute the basis for constructing a targeting SE 2.0 attack.

Often, the information sold on the darknet market places relates to highly sensitive information about employees at specific companies or governmental institutions, and can include information such as credentials, private details, schematics and blueprints of buildings or offices - all of which are invaluable information, if the target is important enough for the social engineer to attack.

We have deliberately decided not to include information about the darknet market places investigated in this study, as we do not wish to promote the use of these places. We are confident that the reader will be able to find any relevant information about these market places from a simple online search on the topic.

3.4.2.11 – Summary of Reconnaissance Methods

A summary of the reconnaissance methods covered above is provided in table 2 below, which is included with the purpose of giving the reader an overview of the applied methods, and is intended for future reference:

Table 2: Utilised Reconnaissance Methods

Method	Tool(s)	Description
Pre-Reconnaissance	• Search engines • Public databases • Target's website • Social media	Applying traditional searches using open sources to understand the target organisation. Pre-reconnaissance includes investigating: • the target organisation's structure • its products and services • its business model • current and potential partners • competitors • financial records • current and former employees
Advanced Google Searches	• Adv. search parameters & • Google dorks	Advanced searches using Google Search Engine to uncover otherwise hidden information.
Robots.txt	• Accessing robots.txt directly	While most search engines respect robots.txt, accessing the robots file directly might uncover, what the target organisation is attempting to hide, e.g.: http://www.[target_website].com/robots.txt

Metadata analysis	• Fingerprinting Organisations with Collected Archives (FOCA)	FOCA scans Google, Bing and Exalead for documents on the target domain, and crawls the document and locally analyses the metadata.
Systemic Infrastructure analysis	• Maltego	Maltego provides an systemic overview of the IP infrastructure, with the purpose of identifying structural weaknesses within the information network of an organisation.
Email harvesting scripts	• Custom scripts	We have developed two scripts in Python, which scans the internet for email addresses with a specified domain name (TLD). It utilises the Google search engine, from which it crawls all results from both indexed websites and documents (incl. PDFs), based on the parameters: ([a-z+0-9+._-]@[Domain name].[ccTLD]).
ID of social media accounts	• Facebook • Twitter • LinkedIn	ID of the following social media accounts: • Facebook • Twitter • LinkedIn
Sentiment analysis and personality profiling (SMPP)	• Custom script	We have developed a script in Java which crawls content from social media sources (Facebook and Twitter), and analyses the content based on predefined parameters. The sentiment analysis contributes to making a simple but effective personality profiling of target individuals based on what they publicly post, and how much information and how many links they share with others.
Social Network Analysis	• Custom script • IBM i2 Analyst's Notebook	We have developed a script that crawls relational data between employees, by using social media as the point of departure. The script collects all available data about an employee's network from Facebook, correlates it with other employees and outputs a CSV file, which can be read into IBM i2 Analyst's Notebook that can visually illustrate how the social network of the employees are structured.
Deep Web & Darknet	• Specialised search engines • Web.archive.org • Pastebin	Using special search engines to investigate leaks on the deep web and darknet have been utilised with the purpose of uncovering both intentionally and unintentionally leaks as well as sold information about the target organisations.

3.4.3 – Phase Two: Target Selection

The target selection procedure consists of: (1) analysis of results from the reconnaissance phase; (2) identification of susceptible target individuals; and (3) identification of information gatekeepers. The criterion for this phase relies on simple - yet effective - parameters that directly relate to the reconnaissance phase, and which are illustrated in layers in figure 8. The purpose of having a target selection procedure is to methodologically follow a process for determining, which target individuals are of the greatest importance to the attacker. The target selection procedure is as follows:

- Identification of the employees in the designated target group of the organisation
- Identification of the target groups' email addresses and cell phone numbers
- Identification of the target groups' social media accounts
- Identification of the target individuals' publicly accessible social media accounts
- Based on the Social Media Personality Profiling (SMPP), who from the target group are most susceptible to social engineering 2.0 attacks?
- From the social network analysis of target individuals' social media accounts, who can be identified as having a large group of colleagues as connections (social status)?
- Which target individuals within the respective target organisations hold a significant role with either access to people of interest, direct system access, or access to other relevant information of interest (information gatekeepers)?

Figure 8: Target Selection Process

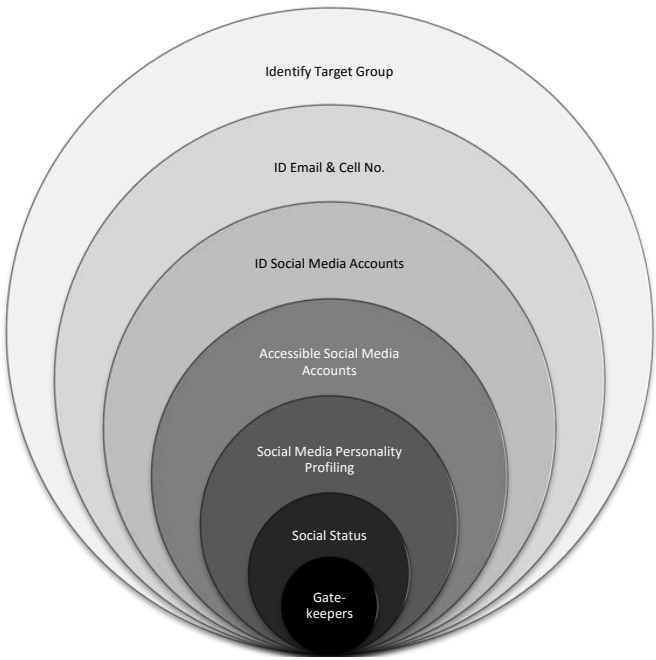

Mainly three types of target individuals are intended to be identified: (1) target individuals who are deemed susceptible to phishing attacks; (2) target individuals with a high social status within the target group; and (3) information gatekeepers, who have direct access to critical information and/or systems.

3.4.4 – Phase Three: Attack & Exploit

This section covers the applied attack methods used in the social vulnerability assessment (SVA) in the field trials of Project SAVE. All methods presented in this section constitute social engineering 2.0 attacks. The utilised attack vectors have been selected on the basis of what each company involved would accept, and include: (1) Conventional phishing, (2) whaling attacks, (3) spear-phishing, (4) smishing, (5) PDF attacks, and (6) USB attacks. This section is intended to give a brief introduction to the various methods and further deepen the reader's understanding of the terminology within the field of social engineering. The section is finalised by a summary of the attack vectors, including methods that have been excluded from the SVA for various reasons.

3.4.4.1 – Conventional Phishing

Conventional phishing attacks refer to the classical phishing emails, which are often generic in nature, designed purely to cause deception by attempting to trick users into performing clicks on a malicious link, downloading a file and executing it, or getting them to fill in a rogue web form, in an effort to harvest their security credentials, i.e. their username and password.

Figure 9: Generic Facebook Phishing Example

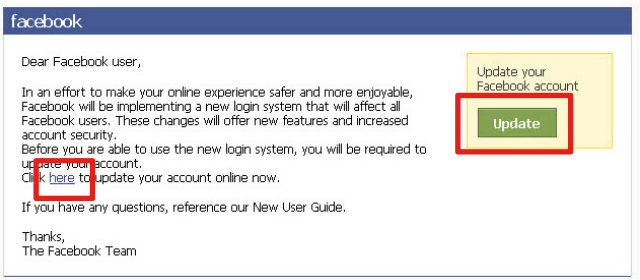

In figure 9 is seen an example that illustrates the content of a generic phishing email, identifiable by its generic introduction: '*Dear Facebook user*', rather than being specified for the actual owner's name of the Facebook account. In this particular example, users are tricked into clicking the link, which will then proceed to open a web form that requires the users to update their security credentials. All of the hyperlinks in such phishing email will typically direct the user to the same destination, namely to the rouge web form that will harvest the credentials once entered into the form. The phishing attempt employs deception tactics based on the trust that most users have in Facebook, by emulating a legitimate request, when in fact it is not.

However, in this study we have refrained from attacking via social media networks, as they constitute private accounts. We have also refrained from using generic content in the context of sending phishing emails without any relevance for the user. Instead, we have personalised each phishing email by *spoofing* the email address of someone within the target organisation, which were identified during the reconnaissance phase (cf. ch. 5 for further information on this). Email spoofing refers to the creation of email messages with a forged sender address. Spoofing is intended to mislead the recipients about the origin of the messages, in an attempt to get them to conduct actions they would not otherwise do, i.e. it is an attempt to social engineer them using manipulation and subterfuge.

The applied methods utilised in the social vulnerability assessment for spoofing the sender email addresses varied. During the tests, we experienced both internal and external difficulties with spam filters, mail servers and the technical solution that we used. For some attacks we used the *Social Engineering Toolkit* (SET), which is an open source framework for conducting social engineering 2.0 attacks. However, in other attacks, we relied purely on purchased domain names that resembled legitimate domain names – often referred to as Typosquatting – which are domain names that closely resembles the actual domain name of a company, in an effort to deceive the user into trusting the legitimacy of the malicious website, for example: If the link in the phishing email directed the users to a malicious website, then this would be an option for disguising the malicious purpose of the website[50].

3.4.4.2 – Whaling

Whaling constitutes a phishing attempt specifically targeted for upper level management within an organisation. It is thus the professional status within the target organisation that determines which targets are considered relevant whaling targets. Most often, a whaling attack will target upper level management, e.g. the CEO, CFO or CTO, as they are key individuals within an organisation, who most likely have direct access to sensitive information. Additionally, if an email account from upper level management is compromised, e.g. the CEOs email account, then it can be exploited for attacks on employees within the organisation, or even against other targets outside the organisation that are relevant for the attacker to target.

Whaling does not have to take the form of a targeted spear-phishing attack, but it certainly *can*, if relevant OSINT has been uncovered to substantiate such an attack. Whaling thus simply refers to targeting an individual at the decision making level of an organisation[51]. The most commonly known form of whaling is perhaps *CEO Fraud*, where the attacker attempts to get a company's accounting department to transfer funds to a specified bank account under the attacker's control. The attacker does so by spoofing the email address of the company's CEO, in an attempt to disguise the fraud attempt as a legitimate request. This particular method has not been explored in Project SAVE.

3.4.4.3 – Spear-Phishing

Spear-phishing attacks resemble regular phishing attacks, but differentiate by being highly targeted, containing an element of personalisation, in an effort to raise the deception level of the attack. Spear-phishing attacks are

therefore not generic in nature, but are rather addressing the recipient by name, and can contain information relevant to the recipient, or relevant files, which the attacker has discovered in the reconnaissance phase.

The spear-phishing attempts executed in this study were constructed on the same basis as the aforementioned whaling and phishing attempts. However, what signifies a spear-phishing attack is the level of personalisation and targeted nature of the phishing email. This requires an in-depth understanding and insight into the target, the target's operational environment and the owner of the spoofed email address, as well as the relationship between the spoofed sender and the recipient of the spear-phishing email. As previously mentioned, email spoofing refers to manipulating the header details of the email, making it look as if the email originates from someone the target knows and trusts[52]. This can be accomplished in several ways, ranging from low-level manipulation of the sender's name, to high-level spoofing exploiting knowledge of the mail server[53].

A spear-phishing email often contains a hyperlink, which directs the recipients to a malicious website, or a web form that requires the recipients to enter their credentials, either in the form of a link, or directly implemented into the phishing email, or can simply contain a malicious file -malware – that the sender will attempt to animate the recipient into opening.

3.4.4.4 – Smishing

A *Smishing attack* resembles a phishing attack, with the only difference being that the attack vector is SMS rather than email. Figure 10 illustrates an example of an actual smishing attempt, which was recorded during the social vulnerability assessment of one of the target organisations involved in Project SAVE. The SMS is spoofed using a SMS toolkit, which alters the identifying data of the text message, so that it seems as if the SMS derives from a contact of the recipient[54].

However, there exists an inherent problem with performing smishing attacks in this regard; namely that it is practically impossible for the attacker to know how the target has stored the contact details of the person the attacker intends to spoof in the target's cell phone.

However, we overcame this challenge in the social vulnerability assessment by not trying to spoof the name of a sender, but instead spoof the actual phone number of the person that we wanted to be disguised as. This me-

ant that when the text messages were received by our targets, they would naturally be shown on the targets' phones under the correct names of the senders we spoofed, and would fall into any previous correspondence the targets would have had with the spoofed senders.

Figure 10 illustrates an example of this, where a correspondence between two individuals was ongoing. However, the last text message in the figure is one of the messages we spoofed. We were therefore neither compromising the phone of the sender, nor the recipient of the spoofed text messages, but rather we were manipulating the SMS to seem as if it derived from a particular contact that we expected the target to know, based on information derived from the reconnaissance phase.

Figure 10: Smishing Attack

3.4.4.5 – PDF Attack

The PDF attacks used in Project SAVE is a simple PDF file, which contains an integrated hyperlink that when clicked will direct the user to our website, where we can record that the user has in fact opened the file. The PDF attacks conducted in the SVA of the three target organisations are intended to resemble malware, as the PDF file resembles a typical file type, which is often associated with actual attempts of malware infection.

The PDF attack was designed to be a follow-up attack to the various sorts of phishing attacks, since we anticipated that one or more target individu-

als would reply to our phishing attempts and request more information or instructions. The PDF attack thus constitutes a secondary attack, which requires the target to have an initial trust in the original phishing emails sent at an earlier stage in the SVA. If the target replies to a received phishing email, we argue that the recipient will have sufficient trust in the legitimacy of the sender of the email for us to initiate the PDF attack, and consequently the target individual is more inclined to be animated into executing an attached file.

3.4.4.6 – USB Attack

The USB attack vector directly intersects the physical and cyber security spheres, as the USB is a physical device with the purpose of compromising the recipient's cyber security. The USB drive could potentially execute any number of malware types within mere seconds, then retrieve the desired data, and email it back to the attacker's email address. This would have been a quick and clean attack, if the goal of the attacker was to retrieve passwords, execute malware for persistent access, or simply monitor all activity on the target's system, including that of microphone and webcam for audio/visual surveillance of the target.

For the USB attacks utilised in this study, we used an USB thumb drive with an applied automated keystroke injection platform. It has a 32-bit on-board processor to process the injection and execution of the payload. The USB platform uses a covert design with an inconspicuous custom-fit casing made out of plastic, resembling any normal USB drive. However, this particular USB drive is far from a normal USB drive. It includes a MicroSD card, which contains the actual payload that can be either custom-coded, or which can use pre-programmed scripts.

As the USB device emulates a *human interface device* (HID), it is recognised by any computer as a keyboard or a computer mouse, thus overcoming any traditional countermeasures to protect against USB attacks, which are easily bypassed by the HID feature of the drive. This means that even if a company has implemented protective policies that restrict the system from recognising any USB drive that is inserted into the corporate computers, the USB drive will still be able to inject the computers with malware, since it registers as a HID. In addition, this allows the device to have multi operating system (OS) injection capability, as all operating systems recognise keyboards as legitimate devices to connect to the computer.

The utilised USB drive was custom-coded with a payload that was executed as a *macro,* which is short for macroinstruction, referring to an instruction of pre-specified sequences of actions that the computer will perform. It was intended to emulate actual malware in the study.

Upon plugging the USB drive to a computer, it would auto-run the macro, which would instruct the computer to open the *Run* command on the Windows computers, and auto-type a pre-defined website that we wanted to the target individual to visit. In this particular case, they were directed to our web server, so we could record that the target had used the USB drive, which would be logged on our web server log.

The time from plugging in the USB drive until the incident has been recorded on the web server log is approximately 4 seconds. This is an extremely fast process, as the processing capability of the on-board processor far exceeds 1,000 keystrokes per minute.

This is an extremely fast process, as the processing capability of the device far exceeds 1,000 words per minute. It is important to note, that in this study, only one company agreed to the USB attack test.

3.4.4.7 – Summary of Attack Vectors

Table 3 summarises the SE 2.0 attack vectors utilised in this study. The grey fields indicate utilised attacks, while the blue field indicates a vector that was explored, but not used in the SVA, and the purple fields indicate vectors that were excluded.

Table 3: Utilised SE 2.0 Attack Vectors

Method	Tool(s)	Description
Conventional Phishing	• Social Engineering Toolkit	The conventional phishing methods targeted the entire target group with a generic attack.
Whaling	• Social Engineering Toolkit	Whaling refers to targeting a high-valued individual within the target organisation. Often this refers to CEO, CTO, COO or similar. It could also be a board member, with privileged access to sensitive information.

Spear-Phishing	• Social Engineering Toolkit • Own clients	The spear-phishing attacks are targeted attacks, often personalised and based on information from the reconnaissance phase.
Smishing	• SMS Toolkit	Smishing, or SMS phishing, were conducted against one or more individuals, where we have been able to identify phone numbers.
PDF Attack (follow-up attack)	• Spoofed email • Custom PDF	The contact via email is intended as a follow-up attack on target individuals, who responded to the phishing attempts. The aim was to convince them to execute a file locally on their computer. Our example used a PDF file with an integrated link.
USB Attack (Evil USB)	• USB Toolkit • Custom coding	The USB attack involves an evil USB that acts as a keyboard, which executes a macro that will visit the webserver, when plugged into a computer.
Targeted Ads (cf. Appendix C)	• Social media	Via social media networks, we can in some instances create highly targeted adverts, which are so specific that they are only exposed to one or two individuals (cf. Appendix C for additional details on this). REASON FOR EXLUSION: We excluded attacks via social media.
Physical Access	• Pretexting • Domains, website & background story	REASON FOR EXCLUSION: The attack refers to a conventional social engineering (cSE) attack, and thus falls out of scope in terms of addressing SE 2.0.
Vishing	• Phone	REASON FOR EXCLUSION: Vishing would allow for extraction and/or tricking employees over the phone into giving information or accessing downloaded files. However, it was excluded as it does not constitute a 2.0 attack, but rather conventional social engineering.
On-location drop of payload	• GSM audio surveillance equipment (bug)	REASON FOR EXCLUSION: Method was deemed out-of-scope, and thus merely servers as an extreme example of what could be used in a scenario of industrial espionage.
Contact via Social Media	• Fake profile on LinkedIn • Fake connections	REASON FOR EXCLUSION: We completed a fake profile and gathered fake connections with the intention of gaining the trust of target individuals. However, the method had ethical implications, which is why we excluded it from the SVA.

Pharming[55]	• Small script, which redirects the user	REASON FOR EXCLUSION: Essentially, this requires getting the targets to execute malware.

3.4.5 – Phase Four: Exit Strategy

The exit strategy represents an important part of any social engineering attack, because it enables the social engineer to perform other attacks at later point in time, if the trust of the target is maintained. However, due to limitations in the scope of the study, we have decided to delimit our focus from this particular phase, as we have been more concerned with executing and analysing the results of the reconnaissance and the attack phase[56].

If we were to explore and utilise an exit strategy, we would have limited our attacks to one or two vectors, which would only have targeted one or two individuals at each organisation, in an effort not to raise suspicion. Instead, we have focused on attacking as many target individuals as we were allowed to, with the purpose of conducting a more accurate and holistic social vulnerability assessment of the companies involved in this study.

3.5 – Expected Results

We expect the ethical considerations and legal limitations to affect the final results of the social vulnerability assessments conducted of the involved companies, as these limit our operational capability to conduct more aggressive attacks that could otherwise improve the results.

In addition, we expect that the tests conducted on individuals, who were informed of the social vulnerability assessment, will result in an increased awareness level of the targets, thereby making them less likely to fall victim to the attacks executed in the SVA. Nevertheless, we argue that this will not affect the aggregated results greatly, as only a few people from each company were informed of the SVA.

Despite these obvious limitations, we believe that the various attack vectors will still be successful for each participating organisation. In particular, we expect our spear-phishing attacks to be more successful than averagely calculated in other phishing studies, where the main focus was on manually constructed spear-phishing attacks. The average success rate for spear-phishing attacks is approximately 45 pct., which we expect to exceed.[57]

The success rate of Project SAVE is rated on the number of individuals who fall victim to the social engineering attacks, which can successfully recorded on our web server log for documentation.

Chapter 4:
Operationalisation

4. Operationalisation

This chapter will cover the operationalisation of the social vulnerability assessment of the field trials in the project, merging the reconnaissance phase with the attack phase by incorporating the methods used for measuring a successful attack, as well as how the ethical considerations have affected the information recorded from the executed attacks. The chapter will cover the following: (1) the strategy of attack, which will give insight into the *game plan* behind the SVA; (2) Level of analysis, which cover the various levels measured, and describe why some methods are more suitable for an SE 2.0 attack than others; and (3) the criteria for measuring a successful attack.

4.1 – Strategy of Attack

The strategy of attack is the applied *game plan*, which has been developed upon finalising the reconnaissance phase and the target selection procedure. It would have been possible to establish this prior to the reconnaissance phase. However, because the reconnaissance phase is an integral part of the attack cycle of social engineering, we have opted for a post-reconnaissance decision in regard to the utilised attack vectors, in order to first establish which useful information could be uncovered using the applied OSINT methods, prior to selecting the vectors. The attack vectors are listed in table 4, which also includes the total number of attacks executed for each type of attack, against each respective target organisation.

Table 4: Attack Matrix on Targets

Method	Target #1	Target #2	Target #3
Spear-Phishing	3	1	3
Whaling	1	1	1
Conv. Phishing	2	4	146
Smishing	3	5	9
USB Attack	0	0	3
PDF Attack	1	2	0

In order to be able to differentiate between the three target organisations, while still maintaining the anonymity of them, we have decided to number

each organisation as Target #1, #2, and #3. The total number of executed social engineering 2.0 attacks amount to 185, with a total of 10 attacks performed against Target #1; a total of 13 attacks performed against Target #2; and a total of 162 attacks performed against Target #3.

Figure 11: Statistical Overview of Executed Attacks

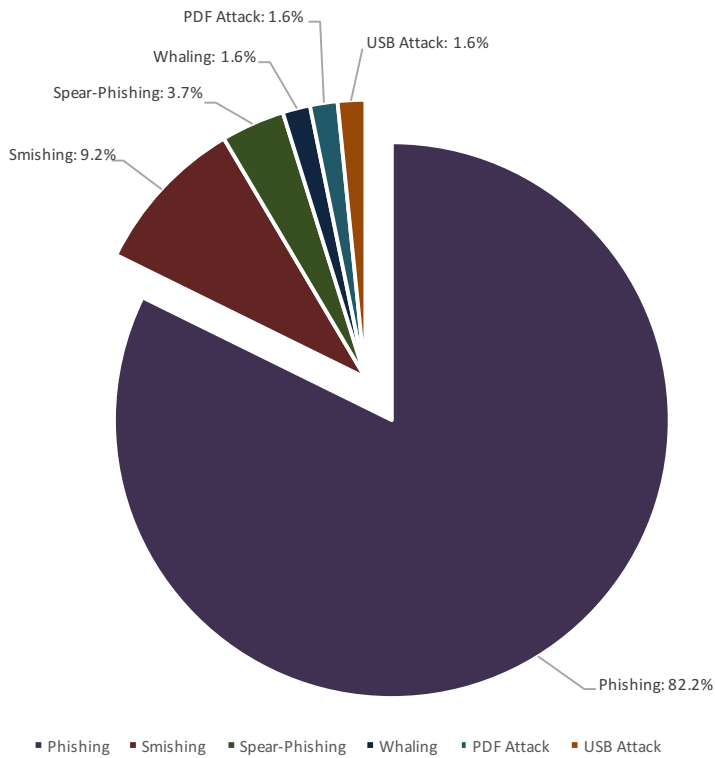

As illustrated in figure 11, the attacks are dominated by conventional phishing emails. This is primarily due to the high amount of phishing emails sent to Target #3, which received a total of 146 phishing emails in the test.

However, all attacks – even the generic phishing emails – were crafted on the basis of information gathered from the reconnaissance phase.

The difference between the regular phishing emails and the spear-phishing emails, which were sent to the targets, is that the latter were personalised and tailored specifically for the recipients, whereas the former were designed for an entire target group – or an entire target organisation - as in the case of Target #3. This makes the conventional phishing attacks more generic in nature, while the spear-phishing attempts were highly personalised.

The smishing attacks - carried out via text messages - are the second-most applied method amounting to 9.2 pct. of the total attacks executed, followed by the spear-phishing attempts amounting to 3.7 pct. The spear-phishing attempts relied on OSINT from the reconnaissance phase, and on whether or not it was possible to identify information that could be utilised in a SE 2.0 attack. Hence, the spear-phishing attempts were excluded in the cases where no relevant information was uncovered from the reconnaissance phase, which was needed to establish relevant content in this spear-phishing email.

Whaling, PDF attacks and USB attacks all constituted the least applied attack methods because of their inherent limitations: Whaling only targeted upper management; PDF attacks were only used as follow-up attacks and required a response from the recipient from a previously sent phishing email; and USB attacks were limited by the companies agreeing for this type of attack and the availability of this type of USBs.

4.2 – Level of Deception

The level of deception required for conducting a successful social engineering 2.0 attack is dependent on the *action* that the target is required to do. The greater the physical action required of a target, the greater the level of deception needs to be, in an effort to make the action seem legitimate.

Although we recognise that this is inherently subjective to every target included in the study, we believe that people will be more inclined to make the required cognitive shortcuts for the impulsive decision to conduct the required action, if the action to be performed has the least amount of physical requirements. In other words, the easier the task is to perform, the greater chance there is for the target to perform the action that the attacker wants the target to perform.

As illustrated in table 5, the four levels that we will be operating with in this study are: (1) Clicking on a link, (2) entering credentials into a web form, (3) executing a file, and (4) plugging a USB into a computer.

Table 5: Level of Deception

Level	Type	Description
Level 1	Link	Target clicks on a malicious link sent by email or SMS.
Level 2	Web form	Target enters credentials, i.e. username and password, into a malicious web form.
Level 3	File	Target executes malicious file, which *could* contain malware.
Level 4	USB	Target connects an unknown and malicious USB drive to the computer.

4.2.1 – Level 1: Link

The first level is concerned with getting a user to click on a link. This level is the easiest of the five actions that are covered in table 5. For recording when a target clicks on a link, in for example a phishing email, we have set up a web server, which records various sorts of information about the user, who clicks on the link. The information recorded about the users include their IP addresses, which can contribute to differentiating one user from another, and also if they tried to access the link from more than one device, namely from smartphones, tablets or other computers. Finally the web server log records the timestamps of when the link was clicked. Each link is constructed with a *unique identifier*, which allows us to separate one phishing attempt from another.

4.2.2 – Level 2: Web Form

The web form constitutes the second level on our list, as it requires people to click on a link and enter personal information in a web form, which requires the target to perform a greater action than the previous level, where only a link needed to be clicked by the user.

For the second level, we have configured different web forms to redirect the entered data to our web server, so that we can observe users, who have entered their credentials. However, due to the established ethical considerations (cf. ch. 3), we have anonymised the data by hashing (MD5 encoding) the data and reduced the hashed values to mere eight characters (out of a total of 64 characters). In doing so, we have maintained full anonymity of the users, while still being able to differentiate one user from another, based purely on the information entered in the web forms.

No passwords were recorded from the information entered in the web forms used in this study – we merely recorded whether or not characters were entered in the password field. This was done with the purpose of verifying that the user had the intention of updating the credentials and/or attempted to login using the web form.

4.2.3 – Level 3: File

Opening a file constitutes a level 3 action, as most people are aware of not executing files from suspicious senders, although it requires less of a physical action than for example entering credentials into a web form. The common public awareness related to executing files makes it a higher level, which can facilitate an array of more advanced attacks and persistent threats, based on the utilised payload.

As previously mentioned we have prepared PDF files that are intended to resemble malicious files. The PDF files are part of the follow-up attacks, which are initiated once: (1) a correspondence with a target is established via email; (2) the target has confirmed that he or she believes the sender to be legitimate by replying to the email; and (3) when it has been established that the target is inclined to open the attached file - thereby allowing us to animate the target to do so.

Although we are restricted from using actual malware, we did integrate a link in each PDF file, which would be recorded on the web server log, once the target clicked on it, allowing us to determine when a file had been opened. Consequently, it was imperative that we managed to convince the target to click on the link in an effort to confirm that the PDF file had been executed by the target.

4.2.4 – Level 4: USB Drive

Finally, the fourth level is the act of connecting a USB drive to a computer. For this study we have prepared three malicious USB drives that were intended to be used by employees at Target #3, as the other targets were excluded from this part of the SVA.

The USB drives acts as a human interface device (HID) and emulates the actions of a keyboard, allowing us to execute macros as soon as the USB drive is plugged into the target's computer. The physical actions associated with plugging an unknown USB drive into once computer, combined with the high level of awareness for using unauthorised USB drives, makes it the

highest level of attack in our SVA. Even people who have not previously received awareness training will be more cautious about plugging in an unknown USB drive to their computer. The aim of the social engineer is therefore to tricker the curiosity of the target, which is sometimes done by applying a label to the USB, e.g. with the word 'Private'.

4.2.5 – Summary of Deception Levels

Figure 12 below illustrates the social engineering process covered in ch. 1, and is a visual representation of the levels involved in the applied attack vectors.

Figure 12: Overview of Levels in the Attack Process

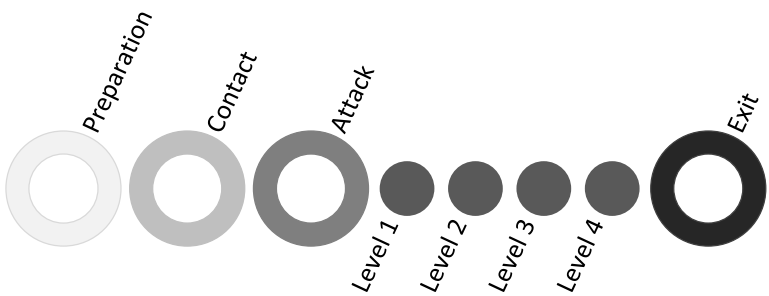

It is important to note that each attack only has one associated level of deception. Therefore, the figure above does not represent a series of attacks conducted against any given target. Rather, it exemplifies that any attack is represented by at least one of the above levels, and that the higher the level, the more difficult it is for the attacker to animate the target to perform the desired action.

4.3 – Criteria for Successful Attacks

Based on the covered levels in the previous section, we have established the criteria for the difficulties of succeeding with an attack. Following table 6 gives a descriptive overview of the success criteria:

Table 6: Criteria for Success

Type	Criteria for Success
Link	When the target clicks on the malicious link and the click is recorded on the web server log.
Web form	When the user enters credentials and clicks on the button, sending the information to the web server log.
File	When the user executes a file, however, only when users click the link inside the executed file, as we would otherwise be unable to confirm that the file has in fact been opened.
USB	When a USB is plugged into the target's computer and the script (payload) is auto-executed, making the browser visit our website, so that we can record the action on our web server log.

However, it is only when the above actions are recorded and confirmed on our web server log that the attempts count as successful. A successful attempt using a web form would show the following on the web server log:

```
|Wed Dec 03 11:16:23 UTC 2015 (from T1SPa,+aUgFE5e): USER-PASS
```

The above can be described as:

[Day] [Date] [Timestamp] [Year] [(Attack type, 8 characters of hash of user's IP):] [User-Pass]

The *attack type* as denoted above, *T1SPa*, refers to *Target #1*, *Spear-Phishing* attempt on *user A*. Please see Appendix D for complete list of recordings from the web server log.

It is important to note that we did not collect passwords, but merely recorded that characters were entered in the password field on the web form. The web server log would therefore only register it, if information was entered in both the username and password fields of the rogue web form.

Chapter 5:
Analysis of Results

5. Analysis of Results

The current chapter will address the results from the social vulnerability assessment (SVA) for each involved target organisations. The chapter is divided into three sections, covering the results from the reconnaissance phase and the attack phase of Target #1, #2 and #3, and will address the information uncovered from OSINT and SOCMINT, as well the executed social engineering 2.0 attacks. It will also portray why some types of attacks were successful, while others remained unsuccessful. Due to the ethical considerations (covered in section 3.2), some of the information that is covered in this chapter has been altered in an effort to maintain the anonymity of the employees and companies involved in the SVA. Finally, the chapter will provide a comparative overview of the uncovered information and the executed attacks, along with a timeline that illustrates the progression of the three SVAs.

The chapter will cover: (1) introduction to the respective targets, (2) results of the reconnaissance phase, including what information has been uncovered in each step of the phase, (3) the results of the attack phase for each applied attack vector against the respective target organisation, and (4) the chapter will be finalised with an overview of the aggregated results.

5.1 – Results for Target #1

5.1.1 – Brief Introduction to Target #1

Target #1 directly constitutes part of critical national infrastructure (CNI) in Denmark, and therefore maintains a vital societal function that - if compromised - constitutes a direct threat towards national interests. The following sections will include information uncovered during the reconnaissance phase, as well as the results from the attack phase.

5.1.2 – Results of Reconnaissance on Target #1

5.1.2.1 – Pre-Reconnaissance

The pre-reconnaissance of Target #1 revealed key business partners and customers, which could potentially be utilised in social engineering 2.0 attacks. More importantly, employees' email addresses, their phone numbers and a complete organisational overview was uncovered directly from their corporate website. This information made it very easy to target specific

individuals within the organisation. This is common practice for many companies, and is therefore not unusual. However, increased transparency and exposure online leaves the company with increased risk of receiving not only harmless spam emails, but also targeted SE 2.0 attacks.

5.1.2.2 – Advanced Google Searches

Google results allowed for identification of the used acquired web domains, CMS systems, corporate sister websites, and it also revealed key partners that the target organisation shares their webserver with. This knowledge makes an attack all the more interesting, as more companies can be targeted in an effort to reach the same goal of compromising their network, as they share the same servers.

5.1.2.3 – Robot Exclusion Protocol

No relevant information was discovered.

5.1.2.4 – Metadata Analysis

Using *Fingerprinting Organizations with Collected Archives* (FOCA), we were able to identify very little about the metadata itself, due to structural challenges of the files stored at Target #1's webserver. Nevertheless, a qualitative assessment of the collected documents and files revealed the following valuable information:

- Written hand signatures of upper management
- A complete standard design layout used by the target organisation to make financial deals with their key partners
- A complete guide to a critical database system
- A complete list of stakeholders and their exact share of options and their voting rights within the organisation

The guide for the database system was of particular interest to us, as it seemed that the target organisation and their strategic partners used a database containing key values to report important figures, relevant for their particular industry. Having access to the guide, ensured that we were familiar with the system prior to gaining access, and allowed us to get intimate with the expected system, how to access the data, and perhaps most importantly, how to delete or manipulate it.

5.1.2.5 – Systemic Infrastructure Analysis

Using Paterva's Maltego, we were able to uncover some interesting and valuable information. Since we were not fully aware of the entire scope of the available content on their website, we used Maltego on Target #1 to gain an in-depth understanding of the content and an overview of the systemic infrastructure, based on their web domains.

Figure 13: Maltego Results for Target #1

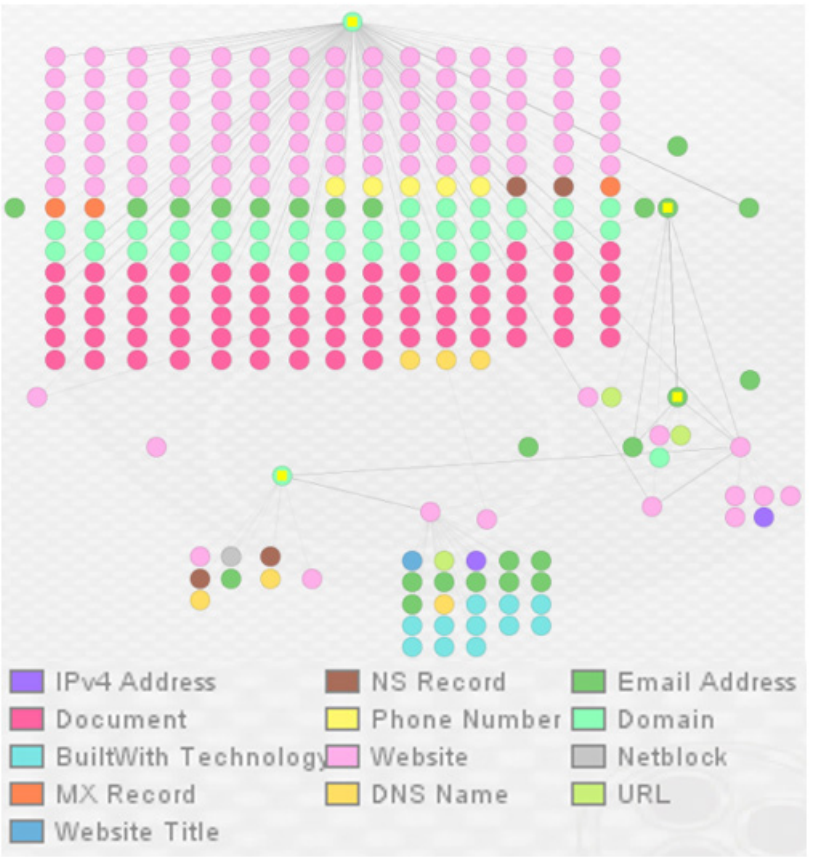

Initially the results seemed inadequate for contributing to an appropriate SE 2.0 attack. Only a single additional email address was uncovered, and the documents uncovered, using Maltego, had already been scanned and analysed by using FOCA during the metadata analysis, and had undergone a thorough qualitative assessment.

However, one particularly interesting piece of information in the network was revealed. We identified how to access the database system used for regulation and maintenance of critically important records. In the metadata analysis we uncovered the guide for the database, but we found additional information about the database using Maltego. Perhaps most troubling was that the database was accessible online by using a standard login form. If a real attacker had knowledge of this particular system used by Target #1, and he or she knew which actors within the company had access to the database system, it would be an easy task to design appropriate targeted attacks, in an effort to gain unauthorised access to the system. Several attack strategies and deception tactics could be employed by the attacker to leverage the target's cyber security and gain access to the company database, of which we decided to move forward with one particular strategy, which will be explored further in later sections.

Perhaps even more troubling was that the database did not use the secure SSL-protocol. Figure 14 shows a partial URL, and reveals the use of *HTTP*, rather than the secure SSL-protocol (which would be shown as *HTTPS* in the URL). This allows for easy access using a simple *Man-in-the-middle* (MITM) attack to collect the credentials needed to access the system.

Figure 14: Partial URL of Online Database over HTTP

Even if their website used SSL to secure traffic, the security could still be leveraged by more advanced MITM-techniques, which are beyond the scope of this particular study. Essentially, having a way of gaining access to a database's login form, which stores critical records for the company involved and its partners, is far from ideal, from any security perspective.

Besides identifying the database and how to access it, we were able to identify the exact users, who have access to the database by using a *trial/error* approach, which essentially is attempting to decode an error message by performing random login attempts on the system, and then interpret the error message that is outputted. By doing so, we were able to identify which employees have access to the system, as exemplified in figure 15 and 16.

Figure 15 illustrates the login form's error message, which reveals that the account (associated email address of an employee) will be suspended if continuous errors occur: "*You have 4 more attempts before your account is suspended*".

Figure 15: Failed User-spcific Login Attempt

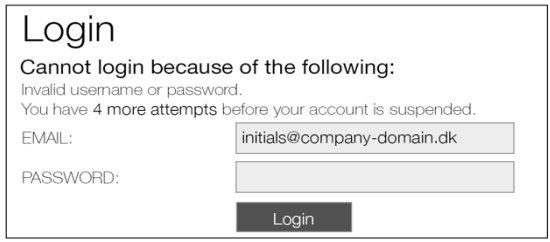

It was not our intention to conduct dictionary attacks or brute force the password in an effort to gain access. Rather, we were looking to identify users with access to the system, with the purpose of determining which users should be our primary targets for the SE 2.0 attack phase for this target organisation.

To confirm that the error message from figure 15 was correctly interpreted, we attempted the same login procedure - but this time by using a random email address with Target #1's TLD ([company-name].dk) – one that is completely fictitious.

Figure 16 illustrates the error message received when using a fictive email address, which completely differs from the former. This allowed us to single out specific users with access to the database, by comparing the different error messages received.

Figure 16: Failed Random Login Atteempt

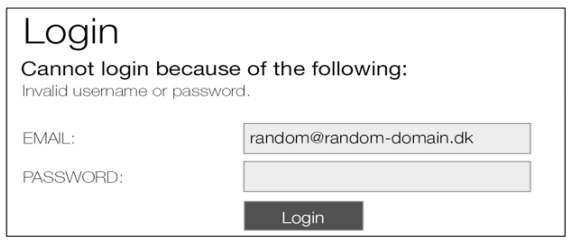

Attempts to access the system were done using all the gathered email addresses recovered during the reconnaissance phase, allowing us to identify all of the employees, who had access to the system. This allowed us to tailor targeted spear-phishing attacks, based on the uncovered information and the results of the conducted tests on the database's login form.

Finally, we uncovered that people from other companies, including key partners, also use this system. To fully understand which individuals from other companies use the system too, we could have written a small script, e.g. in Python or Java, which would try each email address uncovered from key business partners, and thereby reveal some of the employees from other companies who have access. As a result, one could have expanded the perimeter of the attack, if the end goal of the attacker was to gain access to the system.

5.1.2.6 – Email Crawling

Our email harvesting script did not uncover any additional, relevant emails that were not already available from Target #1's website or that were otherwise uncovered in a previous step.

5.1.2.7 – ID of Social Media Accounts

Target #1's use of social media is rather limited, and very little useful information was uncovered. Only a few of the employees from the target group with for example Facebook accounts were accessible, limiting our capability for conducting more in-depth analysis of potential targets.

Figure 17: Social Media Accounts for Target #1

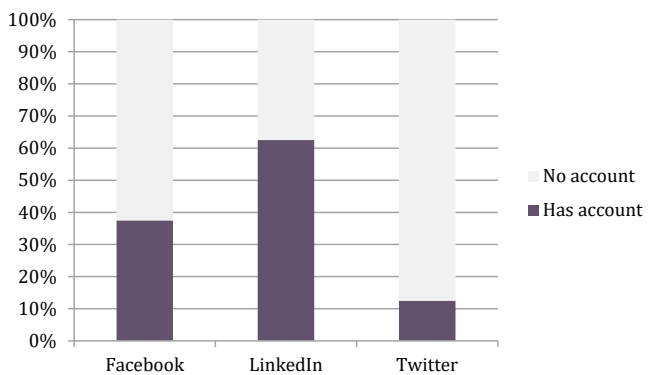

5.1.2.8 – Sentiment Analysis and Personality Profiling

Due to lack of data from the identified employees' social media accounts of Target #1, we had very little basis for conducting a sentiment analysis and personality profiling of the employees.

5.1.2.9 – Social Network Analysis (SNA)

Due to lack of data from the identified employees' social media accounts of Target #1, we found little need for a social network analysis, as it would not provide additional insight into the social constellation of the work place.

5.1.2.10 – Deep Web & Darknet Investigation

No relevant information was discovered.

Summary of Reconnaissance Results for TARGET #1

- Gathered complete list of emails and phone numbers
- Gathered complete list of members
- Gathered Signatures of all board members
- Identified which online platform (CMS) they use
- Identified layout for the deals (incl. letter head and design)
- Identified complete list of stakeholders and their voting rights within the organisation
- Identified some social media accounts, though nothing noteworthy
- Identified a critical database system and how to access it (accessible online)
- Identified user guide for the database
- Identified which specific users have access to the database

5.1.3 – Results of Attacks on Target #1

The social vulnerability assessment of Target #1 included a total of 10 conducted SE 2.0 attacks, with five utilised attack vectors (and the number of executed attacks for each respective attack vector), as shown in table 7.

Table 7: Utilised Attack Vectors on Target #1

Attack vector	Number of attacks	Pct. of total
Spear-phishing	3	30
Whaling	1	10
Conv. Phishing	2	20
Smishing	3	30
PDF Attack	1	10

5.1.3.1 – Spear-Phishing

The spear-phishing attack consisted of a true copy of Target #1's website, as well as a true copy of the web form used to access their critical database system, which was identified during the reconnaissance phase.

The individuals identified to have access to the system from the reconnaissance phase, were contacted by email, where we spoofed the CEOs email address and name, by using the Social Engineering Toolkit (SET) - an open source toolkit for conducting various sorts of SE 2.0 attacks.

The email instructed the target individuals to click on the supplied link to the rouge website and login by using their credentials. The individuals were informed that the system was causing an unknown error, and required the target individuals to confirm that they had the same issue with the system. This was an attempt to make it an urgent matter, by having the CEO request the employees whether or not the system was working as intended. The tone-of-voice of the content in the spear-phishing email was formulated to have a sense of an urgent matter, yet still with the necessary authority expected from a CEO.

The login form on the malicious website would log any entry into its email and password fields. However, it would not log the characters being entered; it would merely confirm that something had been entered in the username and password fields. When the target individuals had entered the credentials and clicked on the *login* button, they would be redirected to the original website and web form, which would display an error message stating that the wrong credentials had been typed. This would mask that any illegitimate action had taken place, as the target individuals would merely see it as a typo on their part. From that point on, if the targets were to attempt another login, they would be gained access to the system as usual, because they would have been redirected to the original login page, after we had harvested the credentials from the attempt on the malicious site. In total, three spear-phishing attacks were executed against the target organisation; however, all attempts to compromise their system proved unsuccessful as illustrated in table 8.

Table 8: Spear-Phishing on Target #1

Attack vector	Total attacks	Successful attacks	Success rate %
Spear-phishing	3	0	0

After consulting with Target #1, they revealed that the sender we were spoofing would never use his Gmail account (which was utilised by using the Social Engineering Toolkit). That made the targets aware that protocol had been broken, and that it constituted an issue, since the actual sender would never use a Gmail account for contacting and requiring them to log into the system.

Essentially, the implementation and strict use of protocol for communication amongst employees, for example only using a corporate email addresses, proved to be an important factor in defending against the spear-phishing attack.

5.1.3.2 – Whaling

The whaling attack targeted the chief operating officer (COO) of Target #1. We constructed the attack to derive from a woman with a familiar name, representing an organisation with a familiar name.

The whaling attempt was spoofed using SET and proved unsuccessful. One reason for this might be that we were restricted from using names of real people and organisations, which might otherwise have given a different outcome in the whaling attempt.

Table 9: Whaling on Target #1

Attack vector	Total attacks	Successful attacks	Success rate %
Whaling	1	0	0

5.1.3.3 – Phishing

The phishing attacks were constructed on the basis of spoofing the email address of the COO of Target #1, and were personalised to random employees, who had not taken part of the spear-phishing and whaling attempts.

We were spoofing the COOs email address, asking employees if they were aware of a specific topic that relates to their key business area.

Table 10: Phishing on Target #1

Attack vector	Total attacks	Successful attacks	Success rate %
Phishing	2	2	100

Although the tone-of-voice of the email was one of urgency, the tone was proper (as to be expected from a COO) as the objective was not to cause unnecessary panic within the organisation.

Our goal was merely to provoke either a confirmative click on the malicious link included in the phishing email, or alternatively a response to the email, so we could initiate the PDF follow-up attacks that we had planned. Both phishing attempts proved successful in getting the recipients to click on the phishing link.

5.1.3.4 – Smishing

The smishing attacks were constructed on the same basis as the phishing attacks. We spoofed text messages sent to the targets, though, instead of utilising email as an attack vector, we used SMS.

The SMS messages included a malicious link that when clicked on would redirect them to our web server, where we could record the action on the web server log.

We used the COO as the sender of the spoofed SMS. The tone-of-voice of the message could on the one hand be interpreted as an urgent matter, and on the other hand as a required matter for the employees to react on or be informed of.

The SMS contained the text: "*Are you aware of this?: [LINK]*". The link looked as if it was directed to an article relating to their key business area, and with a headline that could be argued to be of relevance for all of the employees that received the SMS. In conclusion, all of the smishing attacks were successful, as illustrated in table 11.

Table 11: Smishing on Target #1

Attack vector	Total attacks	Successful attacks	Success rate %
Smishing	3	3	100

5.1.3.5 – PDF Attack

From an attacker's perspective, we were aware of the fact that some employees of Target #1 might respond to the initial phishing emails sent - either with questions or with comments on why the link for the phishing email did not work, as the link merely directed them to a blank website. Because

these target individuals would reply to the email address we used, we anticipated that they believed the sender was legitimate, and we therefore felt comfortable in animating the targets to open a file. For this purpose, we created a PDF file that was sent to the individuals who responded.

Our task was to get the targets to execute the PDF file. However, without using malware, we could only confirm that they had in fact opened the PDF file, if they clicked on the integrated link. Consequently, we instructed the target to click on this link inside the file, which the individual was told would lead to the information previously attempted to be accessed in the previously sent phishing email that simply redirected them to a blank website.

For Target #1, this was only attempted on one individual of the target group, although it had an instant, positive result, which confirmed that the target trusted the sender, and therefore perceived us as legitimate.

The target fully believed that we were who we pretended to be, allowing us to continue the email correspondence with the individual. As illustrated in table 12, the PDF attack was successful.

Table 12: PDF Attack on Target #1

Attack vector	Total attacks	Successful attacks	Success rate %
PDF Attack	1	1	100

The entire email correspondence with the target individual lasted six email exchanges over a 45-minute period, after which we decided to terminate communication, utilising an exit strategy. The target had been animated to click on the integrated link in the PDF file, allowing us to confirm that the file had been executed, which concluded this particular test.

The link in the PDF file still redirected the individual to a blank website page, which the target individual addressed in the email correspondence. We used an exit strategy by lowering the importance of the content sent, implying that it was of less importance, contradicting what we had initially indicated, thus closing the conversation for the time being.

The target individual is the CEO of the organisation, and though we believe we could have convinced the individual into revealing sensitive information, we did not want to compromise the integrity of the company - both for ethical and legal reasons. Nonetheless, this would have been a defining

example of a social engineering 2.0 attack, namely using cyber space as a platform for the elicitation of sensitive information.

5.1.3.6 – Summary

As can be seen in table 13, which is visualised in figure 18, the results of the social vulnerability assessment of Target #1 revealed mixed results. Both the complex spear-phishing attempts and the whaling attempt proved unsuccessful on both counts; whereas the phishing, smishing and PDF attacks were all successful in each instance.

Table 13: Attack Results on Target #1

Attack vector	Total attacks	Successful attacks	Success rate %
Spear-Phishing	3	0	0
Whaling	1	0	0
Phishing	2	2	100
Smishing	3	3	100
PDF Attack	1	1	100
Total	10	6	60

There were a total of 10 attacks executed against Target #1, and of the collective attacks carried out 60 pct. proved successful.

Figure 18: Results on Target #1

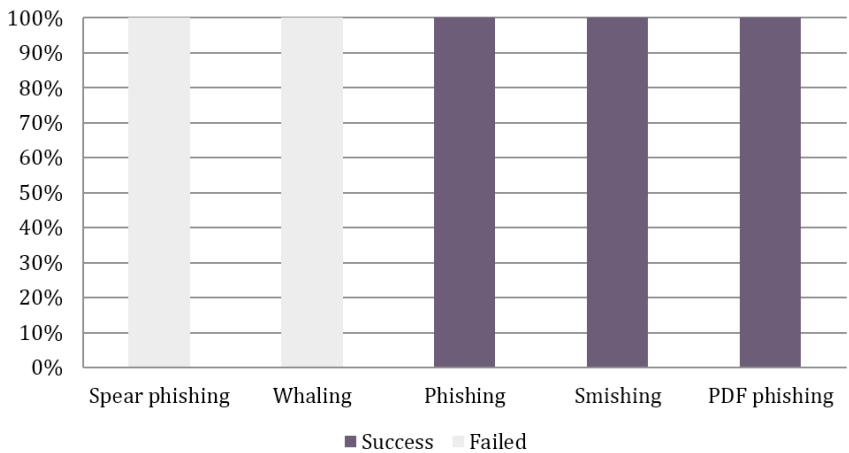

Upon finalising the social vulnerability assessment of Target #1, the participating organisation was briefed on the results of the SVA, as well as informed of all the information that was uncovered during the reconnaissance phase.

5.2 – Results for Target #2

5.2.1 – Introduction to Target #2

Target #2 does not directly constitute part of critical infrastructure in Denmark. However, this company is a subcontractor to critical infrastructure and delivers solutions of critical importance for the operations of a specific segment of critical infrastructure. Target #2 therefore constitutes a different approach to attacking critical infrastructure, as exemplified in the Target breach from 2013, covered in chapter 2, where an attack on a subcontractor can lead to compromising the security of another company.

5.2.2 – Results of Reconnaissance on Target #2

5.2.2.1 – Pre-Reconnaissance

Pre-reconnaissance for Target #2 did not reveal anything of relevance. The nature of the business that this particular company is engaged in affects the transparency of the company, which in turn makes it difficult to uncover the 'digital shadow' of the company. In conclusion, very little information about the company can be retrieved from online sources in general.

5.2.2.2 – Advanced Google Searches

Only in a single instance, during the reconnaissance phase, was it found to be relevant to use Google. This instance was in regard to other domain names owned by the company. The information about the domain ownership revealed insight into business areas otherwise unknown to us, which provided us with a clearer understanding of the business area this company is involved in.

The information could be used for the construction of spear-phishing attacks. However, in this particular case, it would require insight into a highly technical field, which we could not have possessed without prior knowledge in the field the company operates within.

5.2.2.3 – Robot Exclusion Protocol

Robots revealed former business areas, which the company had previously been engaged in. This information would otherwise have been unknown to us. However, they did not play a role in the construction of social engineering attacks for Target #2, as these business areas were no longer pursued by the target organisation.

5.2.2.4 – Metadata Analysis
No relevant results were discovered.

5.2.2.5 – Systemic Infrastructure Analysis
No relevant results were discovered.

5.2.2.6 – Email Crawling
From the company website, only a few email addresses were identified, and we were particularly surprised that our crawling scripts did not identify additional email addresses. This indicates that the target organisation is very conscious about their online exposure and their cyber security in general. This is was an important factor that had an impact on the attack phase designed for Target #2, because our attacks were only going to be performed on the basis of the information that we are able to retrieve during the reconnaissance phase. As a result, we were limited to only perform attacks against Target #2 based on the actual email addresses and phone numbers that we were able to identify.

5.2.2.7 – ID of Social Media Accounts
While employees at Target #2 use social media networks, very little useful information was uncovered. Only a few of the employees in the designated target group had accessible accounts on Facebook, while others did not have any social media profiles.

Figure 19 provides an overview of the social media networks and the percentage of identified accounts for the employees in the target group from Target #2.

Figure 19: Social Media Accounts for Target #2

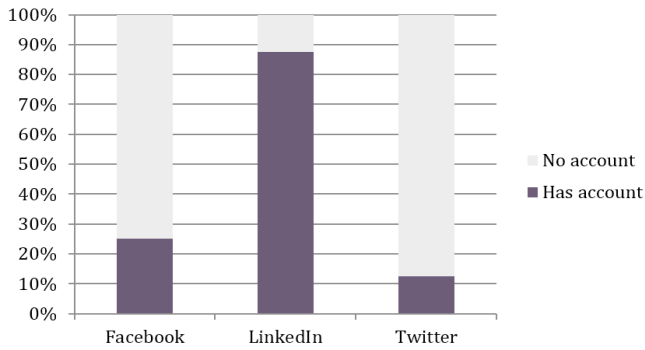

5.2.2.8 – Sentiment Analysis and Personality Profiling

Due to lack of data from the identified employees' social media accounts of Target #2, we had very little basis for conducting a sentiment analysis and personality profiling of the employees.

5.2.2.9 – Social Network Analysis (SNA)

Due to lack of data from the identified employees' social media accounts of Target #2, we found little need for a social network analysis, as it would not provide additional insight into the social constellation of the work place.

5.2.2.10 – Deep Web & Darknet Investigation

No relevant information was discovered.

Summary of Reconnaissance Results for TARGET #2

Surprisingly little information from open sources was available. We were only able to:

- Gather a partial list of emails and phone numbers
- Identify social media accounts, though nothing noteworthy
- Identified the CEOs private Gmail account

There are two reasons for this:

1. The company has extremely tight security measures with very little available information from open sources.
2. The company works within a security-related field, which may have made them more conscious of their online presence.

5.2.3 – Results of Attacks on Target #2

The social vulnerability assessment of Target #2 included a total of 13 executed SE 2.0 attacks, using five different attack vectors. In the table below you can see the utilised attack vectors and the number of attacks executed:

Table 14: Utilised Attack Vectors for Target #2

Attack vector	Number of attacks	Pct. of total
Spear-phishing	1	7.7
Whaling	1	7.7
Conv. Phishing	4	30.7
Smishing	5	38.5
PDF Attack	2	15.4

5.2.3.1 – Spear-Phishing

Due to the lack of information from the reconnaissance phase of Target #2, very little was known about the individual, who was targeted in the spear-phishing attack. The reason for selecting this particular individual, however, was due to information uncovered from the target's LinkedIn account, which were directly relational to information uncovered from the company website, making it the most substantial basis for a spear-phishing attack on an individual in the target group.

The general lack of information on Target #2 influenced the design of the spear-phishing email, which took form based on the collected intelligence, namely their key business area that was correlated with the little information available on social media.

To overcome the human barrier of security in this case, we created a news agency that showed interest in their key business area. In doing so, we expressed a desire to write a news article about the company, with the promise of including the article in an upcoming issue of the newsletter from the fake news agency. Free marketing and advertisement remains an attractive feature, and our goal was to exploit this from a business perspective, in an effort to peak the interest of the target individual.

We acquired a web domain for the purpose of this attack. We then used a female character as the sender, since the opposite sex is often more appealing[58], and we targeted a decision-maker at Target #2. The spear-phishing attack proved successful.

Table 15: Spear-Phishing on Target #2

Attack vector	Total attacks	Successful attacks	Success rate %
Spear-phishing	1	1	100

5.2.3.2 – Whaling

The whaling attempt specifically targeted the CEO of Target #2, and was constructed on the same basis as the spear-phishing attack. The attack differed by targeting the C-suite level of the company. The attempt, however, proved unsuccessful, and after concluding the field trial testing, we consulted the CEO, who explained that the email looked suspicious, because he did not know the organisation of the sender. Furthermore, the CEO of Target #2

was informed of all the attacks conducted, which might have influenced his susceptibility to phishing emails in general, due to the heightened awareness.

Table 16: Whaling on Target #2

Attack vector	Total attacks	Successful attacks	Success rate %
Whaling	1	0	0

Additionally, we were restricted from using real names and organisations due to legal concerns – a factor that might have influenced the outcome of the whaling attempt. An example could have been to identify his connections on LinkedIn, spoof their email addresses, and establish contact with the CEO via email, and animate the target to conduct the desired action.

5.2.3.3 – Phishing

During the reconnaissance phase, we uncovered the CEO's private Gmail account, which created the basis for the phishing attacks conducted against employees in the target group of Target #2. We registered an email account on Gmail that closely resembled that of the CEO's personal account. This was made possible due to the construction of the individual's name. For example, if the CEO's name was *Rasmus Kjer* and he had a personal email account named *rasmuskjer@[mail].com,* we could register the email *rasmuskier@[mail].com*, using the letter *i* rather than the letter *j*, which the original mail account used. This is a minor detail that is often difficult for recipients to notice, making it a popular method for masking the true intent of the email.

Next, we merely had to register with the CEO's actual name and prepare the server setup. We sent the phishing mail to four employees from the target group, instructing them to catch up on a business related issue of critical importance – an issue, which they could obtain access to via the link included in the phishing email. As shown in table 17, the results had a 100 pct. success rate.

Table 17: Phishing on Target #2

Attack vector	Total attacks	Successful attacks	Success rate %
Phishing	4	4	100

5.2.3.4 – Smishing

The smishing attacks were constructed on the same basis as the phishing attacks, where we spoofed the sender's email address to resemble the CEO's. However, this time we were spoofing the phone number, which we had identified during the reconnaissance phase.

The attack was a two-step process, as the attack evolved parallel with the collection of new information and data of the targets, which in turn was used in a new phase of attacks.

The smishing attacks were initially only meant for three target individuals, but due to employees responding to the former phishing emails sent, we received more phone numbers than initially collected during the reconnaissance phase as their phone numbers were listed in their email signature. The total amount of smishing attacks thus increased from three to five attacks, for the designated target group. Four out of five target individuals were susceptible to the attacks, amounting to a success rate of 80 pct., as shown in table 18 below.

Table 18: Smishing on Target #2

Attack vector	Total attacks	Successful attacks	Success rate %
Smishing	5	4	80

5.2.3.5 – PDF Attack

As stated in the previous section on smishing, two individuals initially responded to the phishing attacks, thereby leaving room for various follow-up attacks. In the initial phishing email, the recipients were directed to a blank website, when they clicked on the phishing link in the email. Consequently, we anticipated that some employees might respond to the email, although we were aware that the entire operation could be compromised, if just *one* employee would call the CEO to ask about the content of the email. However, we were confident that the standard operating procedure of the company would be to reply to an email with an email, instead of a phone call. We therefore anticipated email replies from some of the employees, which we could use in the follow-up attacks.

Two target individuals replied to our initial phishing email. We sent them both an email with a PDF file attached, in an attempt to animate them to

open the file and click on the integrated link. It is worth mentioning that both individuals have highly technical backgrounds.

While the link merely redirected the employees to a website without content, we could continue the email correspondence with the victims without raising any suspicion, as they merely reported back that the link still did not work. During one of the email correspondences, the employee suggested that we contacted other employees within the organisation. This could have created basis for executing more attacks against Target #2, although we decided to apply an exit strategy to stop the process of the attacks. The PDF attacks had a 50 pct. success rate, as only one of the two employees clicked on the integrated link in the attached PDF file.

Table 19: PDF Attack on Target #2

Attack vector	Total attacks	Successful attacks	Success rate %
PDF Attack	2	1	50

5.2.3.6 – Summary

As shown in table 20, the results of the social vulnerability assessment of Target #2 were overwhelmingly positive. Only the whaling attempt proved unsuccessful, whereas the spear-phishing, phishing, smishing and PDF attempts all had successful counts of deceiving the target individuals into performing malicious actions. There was a total of 13 attacks executed against Target #2, and of the collective attacks carried out, 77 pct. proved to be successful. A particularly interesting point for Target #2 was the lack of information uncovered during the reconnaissance phase, which indicated a higher than average level of organisational security consciousness. However, the results reflect that the deception tactics employed worked well, and the employees of the target organisation were quick to react to the instructions provided during the attack phase.

Table 20: Attack Results on Target #2

Attack vector	Total attacks	Successful attacks	Success rate %
Spear-Phishing	1	1	100
Whaling	1	0	0
Phishing	4	4	100
Smishing	5	4	80
PDF Attack	2	1	50
Total	13	10	77

The results are illustrated in figure 20 below, indicating the high level of success. What is crucial to note about Target #2 is that they work in a highly technical field, which means that the majority of the employees represented in this target group cannot be labelled as laymen.

Figure 20: Reesults on Target #2

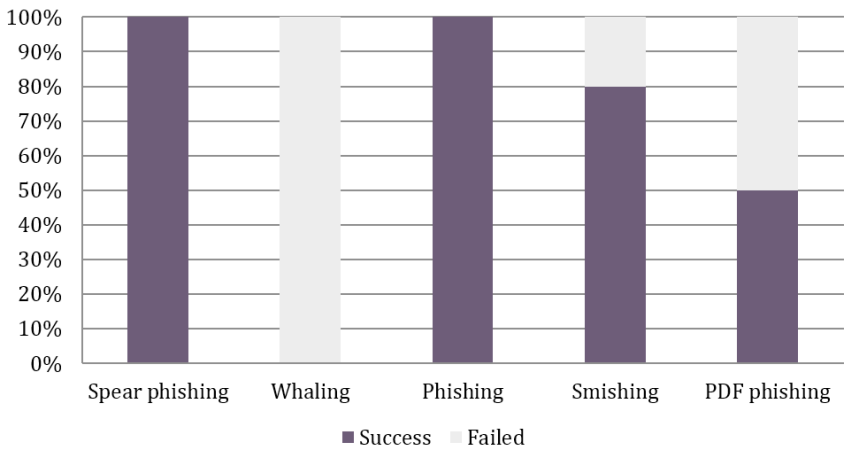

Target #2 has been briefed on all of the information uncovered during the reconnaissance phase, as well as on the results of the social vulnerability assessment. During the debriefing, we gained insight into the thought processes of the involved target individuals, when they realised that they might have been compromised. Thankfully, the CEO was informed of all activities and could thus communicate internally, so that unintended panic was avoided.

5.3 – Results for Target #3

5.3.1 – Introduction to Target #3

Target #3 does not directly constitute part of critical infrastructure in Denmark. However, based on their customer base, they carry out significant assignments for several companies that are part of CNI. Prior to conducting the SVA on Target #3, it was confirmed that they do in fact have critical information on companies that constitute part of critical infrastructure. These results made them highly interesting to us, and hence created the basis for their inclusion in this study.

5.3.2 – Results of Reconnaissance on Target #3

5.3.2.1 – Pre-Reconnaissance

The pre-reconnaissance phase revealed some of their key business partners, as well as a partial customer base that could be utilised in for social engineering 2.0 attacks. More importantly, emails, phone numbers and job positions of 99 pct. of the employees were retrieved directly from the corporate website of Target #3, making it very easy to target individuals within the organisation with social engineering attacks.

5.3.2.2 – Advanced Google Searches

Advanced Google search results revealed a vast amount of information - most of lesser or no importance. However, one interesting finding was the complete design of a guest card used for visitors entering the target organisation.

The design of the guest card was discovered on in a post on Facebook. It was uploaded by a user, who had recently visited the organisation, and thereby made the design of the guest card publicly available online. It would take very little effort to confirm the validity of the design, as a visit to the target organisation in many instances is possible. Knowing the design of the company's guest card makes it easy to forge copies, which can be used to substantiate validity as a legitimate guest, when moving around inside the perimeter of the company.

Furthermore, as we know from the pre-reconnaissance phase, this particular organisation has many visitors with various professional backgrounds. Hence, knowing the actual design of a guest ID card can be used for a physical security penetration test. By using social engineering, the potential attacker could move freely, using methods like pretexting and tailgating into otherwise restricted areas in an effort to retrieve valuable information.

Although this was deemed out of scope for the social vulnerability assessment of this study, it remains a highly possible scenario, as seen in the case of the diamond theft from 2007, where thieves used a similar approach to gain access to AMN Amro Bank in Antwerp, Belgium[59].

5.3.2.3 – Robot Exclusion Protocol

No relevant results were discovered.

5.3.2.4 – Metadata Analysis

The metadata analysis, using FOCA, revealed some interesting results. The most notable results from the metadata analysis were the additional authors, emails and contact details that could be extracted from publicly available documents. Of particular importance was the list of utilised software, which historically speaking, gave an in-depth understanding of the software that had been in use in the past – as well as which ones were currently being used - based on the *creation date* of the documents, for the files that contained this specific metadata information.

Table 21: Sample of Software Versions Found on Target #3

Microsoft Office 2008 for Mac
eDocPrinter PDF Pro (W08Svr x64) Ver 6.76 Build 5953-5949
eDocPrinter PDF Pro (W08Svr x64) Ver 6.48 Build 5428-5426
eDocPrinter PDF Pro Ver 6.10 Build 3390-3388
Adobe PDF Library 10.0.1
Adobe InDesign CS6 (Macintosh)
GPL Ghostscript 8.15
Adobe InDesign CS4 (6.0.4)
Adobe InDesign CC 2014 (Windows)
rqeDocPrinter PDF Pro (W08R2 x64) Ver 6.81 Build 6095-6091
Adobe PDF Library 15.0
Adobe InDesign CC 2015 (Windows)
Acrobat 5.0 Image Conversion Plug-in for Windows
Adobe InDesign CS2 (4.0.4)
Adobe InDesign CC (Macintosh)
iText® 5.3.2 ©2000-2012 1T3XT BVBA (AGPL-version)
Acrobat Distillier 7.0.5
Grafikhuset Publi Service on PDF til PJ Schmidt
Adobe InDesign CC 2014 (Macintosh)
Adobe InDesign CS6 (Windows)
Adobe InDesign CS5.5 (7.5)
Nitro Pro 7 (7. 4. 1. 4)
Acrobat Distillier 6.0
PScript5.dll Version 5.2
Acrobat PDFWriter 3.03
Adobe InDesign CS2 (4.0.2)
Adobe InDesign CS5 (7.0.4)
ABBYY FineReader 8.0 Professional Edition
Adobe Acrobat 5.0 Paper Capture Plug-in
Acrobat Distillier 4.0
QuarkXPress 1.0
CorelDRAW 12.0
Corel PDF Engine 1.0.0.458
Acrobat PDFWriter 4.05

Microsoft Office 2000
Adobe InDesign CS5.5 (7.5.3)
Acrobat Distillier 4.05
Adobe PDF Library 8.0
Adobe InDesign CS3 (5.0)
Acrobat Distillier 8.1.0
PScript5.dll Version 5.2.2
Adobe Photoshop 6.0
Adobe Photoshop
Microsoft Office
Adobe Photoshop CS
Microsoft Office 97
1-Step RoboPDF
Microsoft Office XP
Adobe PDF Library 9.0
Adobe InDesign CS4 (6.0)
Adobe PDF Library 7.0
Adobe InDesign CS2 (4.0.5)
Acrobat Distillier 5.0.5
Adobe InDesign CS4 (6.0.6)
Microsoft Office 95
QuarkXPress 4.11
Acrobat Distillier 5.0
Adobe PDF Library 9.9
Adobe InDesign CS5 (7.0.3)
Adobe PDF Library 11.0
GPL Ghostscript 9.0
PDFCreator 1.1.0Windows
PageMaker 6.5
Acrobat PDFWriter 3.02
Adobe InDesign CS3 (5.0.4)
Adobe PDF Library 6.66
Illustrator
Adobe InDesign CS2 (4.0)
Grafikhuset Publi Service
Acrobat Distillier 6.0.1
Microsoft Office 2007
GFI FAXmaker
PDFlib 2

Although it is not relevant for the scope of this study, the results could reveal software in use with known critical vulnerabilities. This could facilitate the creation of highly targeted attacks, as we would have been able to determine the author of the documents, who was using the vulnerable software. All of this would be accessible information based purely on the metadata analysis.

A similar case was revealed in the RSA 2011 attack:

> *The email was crafted well enough to trick one of the employees to retrieve it from their junk mail folder, and open the attached excel file. It was a spreadsheet titled "2011 Recruitment plan.xls. The spreadsheet contained a zero-day exploit that installs a backdoor through an Adobe Flash vulnerability (CVE-2011-0609).*[60]

The 2011 RSA attack was constructed by using spear-phishing techniques, which were sent over two days to two different low-level groups within the target organisation. The subject line of the email read "2011 Recruitment Plan", and was successful in a single instance, allowing the attacker(s) to gain full access via the trojan payload.

5.3.2.5 – Systemic Infrastructure Analysis

Due to the sheer size of Target #3, and the information available via open source intelligence, the collection of information had to be supported by Paterva's Maltego, where the investigation was divided into three groups: (1) collection of email addresses, (2) systemic network analysis, and (3) additional metadata analysis.

1. Collected email addresses

One of the most important information needed for a SE 2.0 attack is an email address. Without it, it becomes difficult to know where the attack should be directed.

Therefore, much work goes into uncovering email addresses, which is the baseline for any attack revolving around a phishing attempt. Thus, we found it necessary to visualise the volume of employees at the target organisation, as illustrated in figure 21.

Figure 21: Emails on Target #3

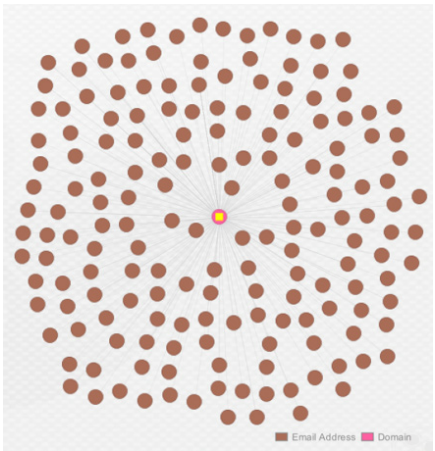

2. Systemic network
The systemic network reveals all the domain names owned by the target organisation, and reveals operations in more than its host country.

For a complete reconnaissance of the target organisation, each domain would be investigated separately, the files would be stored in a *case management system*, and the information would be both automatically and manually assessed and evaluated, depending on the stamina of the attackers.

In figure 22 is illustrated a sample of some of the results uncovered from the systemic network analysis.

Figure 22: Systemic Overview of Target #3

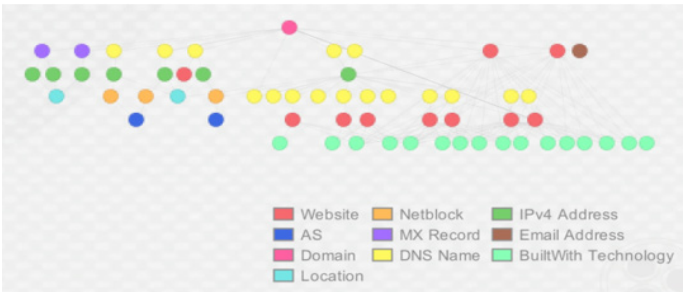

3. Maltego - metadata

Although we had already conducted metadata analysis by using FOCA, Maltego provided investigative insight into where the files were stored and could potentially add to the already-collected information on Target #3.

The network of files shown in figure 23 illustrates the volume of the publicly available files from the target. While we excluded the processing of additional files identified, we include this section to support the use of tools like FOCA and Maltego for the collection and structuring of publicly available information of a target.

The network illustrated in figure 23 contains more than 800 nodes. Each node represents a document or a person; the former dominates the type of nodes represented in the network for Target #3. The network could be extended to identify associated social media networks as well.

Figure 23: Maltego Metadata on Target #3

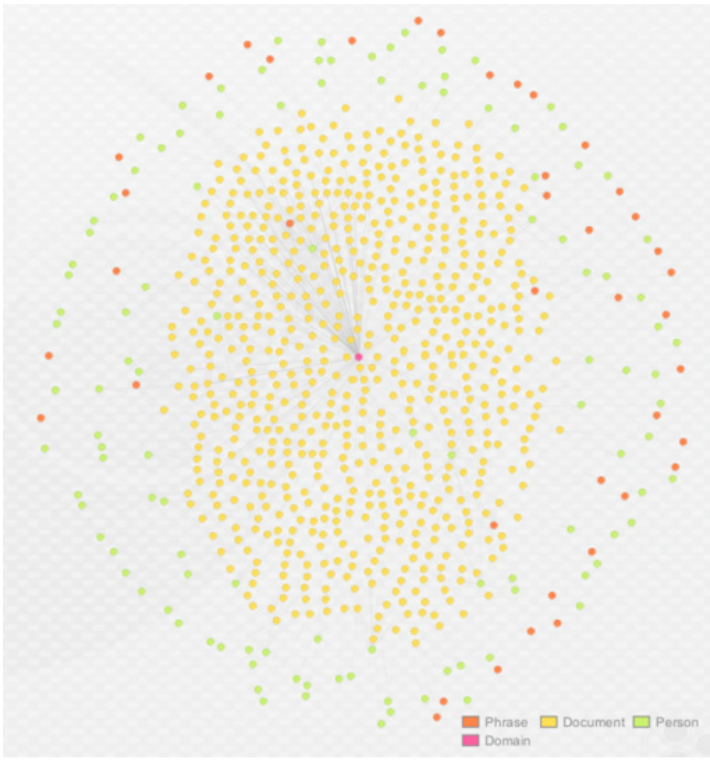

5.3.2.6 – Email Crawling

99 pct. of all active email addresses were available on Target #3's corporate website, though we still used the custom coded script to harvest additional email addresses. The results show that two additional individuals, still employed at the organisation, were discovered using our script. This included one individual, who were of particular interest, due to the person's highly dense network within the organisation, which will be covered in section 5.3.2.9 on the social network analysis. However, without the crawling script, we would most likely have missed this particular target individual.

5.3.2.7 – ID of Social Media Accounts

The employees at Target #3 are both present and active on social media networks. Because Target #3 had many accessible and open social media accounts, we reduced the scope to only address Facebook accounts of the employees, and thus excluded the identification of Twitter and LinkedIn accounts. As a result, the sentiment analysis and personality profiling rely solely on the data crawled from Facebook. This choice was preferred, as it would reduce the time involved in identifying social media accounts, which is a qualitative, manual process that is very time consuming.

The reason for selecting Facebook over other social media networks was that people in general tend to be more active on Facebook than on other networks. Since the sentiment analysis and personality profiling relied on written content that was posted on the target individuals' social media profiles, Facebook created the best basis for conducting a thorough analysis.

Additionally, it is more widely accepted to have spelling-mistakes on Twitter than on Facebook, as tweets are limited to 140 characters, and this is an important factor to consider, since the devised sentiment analysis would only analyse pre-defined words that are spelled correctly. Consequently, Twitter might constitute an inherent bias, if sentences are purposely written in a manner that reduces the size of the words (and the overall tweet), without compromising the interpretation.

42 pct. of the employees at Target #3 had identifiable accounts on Facebook, as illustrated in figure 24. While this number might seem relatively low, many accounts were public and their content accessible, which created a sound basis for conducting the personality profiling based on the devised sentiment analysis of the written content from the target individuals published Facebook posts.

Figure 24: Social Media Accounts for Target #3

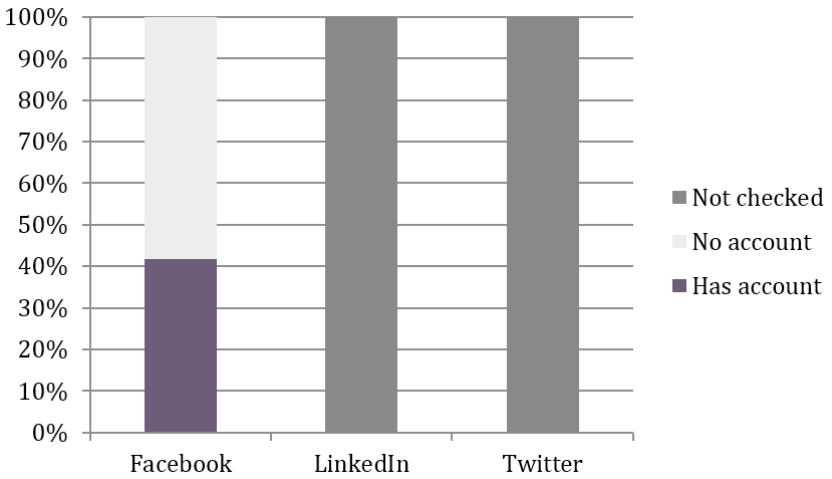

From a qualitative assessment, we could determine that there existed enough substantiated data to be used in our sentiment analysis and personality profiling of target individuals, as covered in the following section. Furthermore, we also had access to their network of *friends* - their connections on Facebook - which provided the basis for conducting our social network analysis.

There was a clear lack of information available from target individuals of Target #1 and #2. However, for Target #3 we saw an exciting development, as the employees of Target #3 had a more open approach to the use of social media networks, which worked to the benefit of this study.

5.3.2.8 – Sentiment Analysis and Personality Profiling (SMPP)

As covered in chapter 3, a high level of neuroticism correlates to the level of receptiveness of an individual, thus indicating how inclined that individual is to social engineering attacks. When running the personality profiling script that was based on the target individuals' Facebook accounts, we were able to identify employees with a high neuroticism score at Target #3. According empirical evidence these individuals should be more susceptible to phishing attacks in particular.

According to our analysis, one individual in particular scored very high in the Big Five model (see figure 25 below). The analysis was based on how

often certain words occurred in the target individuals' Facebook posts, by utilising the *bag of words* approach previously described.

Additional factors that were taken into account included: (1) how often people posted (the frequency of the targets' wall posts); (2) how much they posted (the quantity of the posts, meaning how many characters they post in total); (3) how many links the target individuals shared. The latter was particularly interesting for our study, as current research in the field of phishing studies demonstrates that extremely active FB users who share a great number of links on their walls have a greater tendency to click on links that are shared with them; and (4) we analysed whether or not they shared particular sensitive information, including political or religious conviction and/or sexual orientation or relationship status, as sharing these types of information gives an indication of how 'open' the target individuals are.

Figure 25: Personality Profiles of Users with > 2500 Characters in Public Posts on Facebook

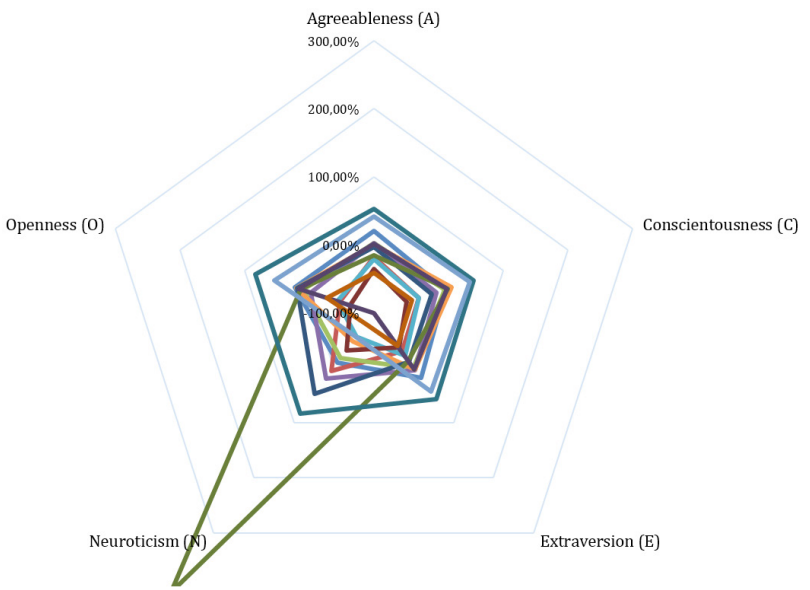

Figure 25 illustrates the individual, who had an extremely high neuroticism score according to the Big Five personality-profiling model. In conclusion, that particular individual constituted an interesting target for the attack phase on Target #3.

5.3.2.9 – Social Network Analysis (SNA)

The social network analysis of Target #3 revealed that certain employees had larger social networks than others within the organisation; these numbers were deduced based on the target individuals' relational networks that were crawled from Facebook. It is not uncommon for employees to be connected on professional networks, e.g. LinkedIn. However, when the private and professional spheres are merged and transcended, i.e. when employees have professional connections on their personal social media networks, it reveals target individuals who are highly engaged on social media, and who are experienced in having - and dealing with - many connections across the professional and personal realms.

Based on the same assumption and findings from the research on personality profiling of social media, we expect highly interconnected individuals to have a greater susceptibility to phishing attacks than less interconnected individuals (cf. section 5.3.2.8).

In figure 26 the social network of Target #3 is illustrated, and it reveals several individuals with a high degree of interconnectedness. This is indicated by the size of the blue spherical objects in the figure. Hence, the larger the sphere, the greater is the density of connections for the target individual.

People with a large social network often have easier access to information as well as access to other individuals within the same organisation. This makes an interesting target easier to compromise, because the attacks can be escalated from their account when it has been leveraged.

Figure 26: Social Network Analysis (SNA) of Target #3

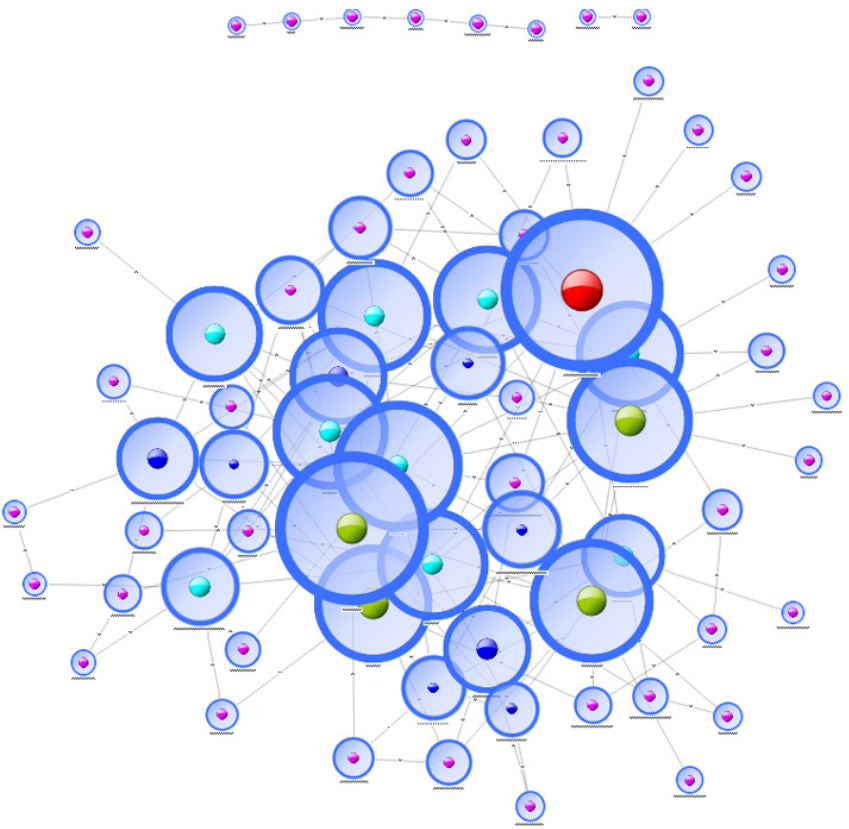

One particular individual had a greater connection density than others, indicated with a red dot in the largest blue sphere in figure 26. As the USB attacks were deemed the most difficult to attempt of all of the SE 2.0 attacks, this specific individual was selected for one of these attacks (cf. section 4.2 on the deception levels). However, this particular target individual did also take part in the broad phishing campaign conducted on Target #3.

5.3.2.10 – Deep Web & Darknet Investigation

No relevant information was discovered.

Summary of Reconnaissance Results for TARGET #3

More information was found than was possible to process within a reasonable time frame. However, from assessed results, we have found:

- Design of Guest ID Card
- Full list of emails and phone numbers
- Identified social media accounts
- Identified useful information based on the metadata analysis of documents
- Analysed data from publicly accessible Facebook accounts, resulting in useful findings
- Conducted a social network analysis, which resulted in useful findings

5.3.3 – Results of Attacks on Target #3

The social vulnerability assessment of Target #3 included a total of 162 executed social engineering 2.0 attacks, with five utilised attack vectors. The table below shows the number of executed attacks for each attack vector:

Table 22: Utilised Attack Vectors for Target #3

Attack vector	Number of attacks	Pct. of total
Spear-phishing	3	1.9
Whaling	1	0.6
Conv. Phishing	146	90.1
Smishing	9	5.5
USB Attack	3	1.9

5.3.3.1 – Spear-Phishing

The spear-phishing attacks constructed for Target #3 were targeted against the management level of the organisation. We spoofed the email address of the CEO as the sender, as we believed this would provide us with the best results.

The tone-of-voice of the written content of the email was in a rather provocative tone, which was intended to provoke the mid-level management personnel into performing a quick reaction e.g. clicking on the malicious

link in the spear-phishing emails. It was intended to project a sense of frustration from the CEO towards the management level.

The content of the spear-phishing email was purely fictional and structured around an issue that did not in really exist, but which related to their business area. The three counts of spear-phishing attempts that were executed against the target individuals proved to be successful in all instances, as shown in table 23.

Table 23: Spear-Phishing on Target #3

Attack vector	Total attacks	Successful attacks	Success rate %
Spear-phishing	3	3	100

5.3.3.2 – Whaling

In the whaling attempt, we targeted the CFO of the organisation, by asking the target to create a presentation on the company's results for an upcoming event - a scenario that was purely fictional. In this attack, we used the same email and sender as with the spear-phishing attempt, namely the CEO.

The whaling attempt included a malicious link that had the purpose of directing the subject to a PDF file, e.g. *http://www.[website].dk/?=Invite.pdf*, and the purpose was to get the CFO to click on the link, so we could record it on the web server log.

The email request was sent later in the day around 7PM. The timing here was important, as we could not anticipate whether or not the CEO and CFO would be at the office at the same time. We therefore timed the attack at a later time during the day, where we expected the CFO to be at home. This proved to be an important strategy, as we later discovered that both are seated on the same floor, with offices next to one another, and this could have ruined the attempt and alarmed the rest of the target organisation by informing them that tests were being conducted, which would otherwise have compromised the SVA being conducted on Target #3. The whaling attempt, however, proved successful.

Table 24: Whaling on Target #3

Attack vector	Total attacks	Successful attacks	Success rate %
Whaling	1	1	100

5.3.3.3 – Phishing

The phishing attempt against Target #3 was planned, timed and designed specifically for the employees at the target organisation. However, the content of the phishing email was broad enough to be sent to a larger part of the organisation. Though classified as a conventional phishing attack, the construction could reasonably be argued to be a spear-phishing attempt, though none of the phishing mails were personalised, and the attack in question was mass-mailed to many employees at Target #3, which is the reason why we have classified it as a conventional phishing attack.

The method applied was a complete copy of the website of the target organisation, where we modified a login form in the middle of the website to record any credentials that would be entered. The purpose of the attack was to get the target individuals to *update* their passwords, and the login form would therefore require them to enter:

1. Their username (we suspected this would be their email address)
2. Their current password
3. And their new password

The website was uploaded to a domain containing a subdomain with the name of the target organisation (www.[company-name].[our-domain].dk) to further substantiate the reliability of the source. Finally, an email address from the attack-domain was spoofed, and the deception tactics applied was to pretext the internal IT-department. The content of the phishing email required all employees to update their password from the web form, which could be accessed from the link included that we included in the email, which directed the target individuals to our rogue web form. We anticipated that some employees might wonder the sudden need for updating their passwords, so we argued in the email that it was due to server updates over the holidays.

For Target #3 we were allowed to conduct the test on a larger target group within the target organisation, resulting in the phishing emails being sent to a total of 146 employees, of which at least 7 were on vacation (which was reported back via automatic email responses).

It is important to notice that during the reconnaissance phase, we had identified the necessary means to utilise the information from this phishing campaign to an actual attack. We knew how to access the mail server, but

needed the necessary credentials to do so, and that could have been used to conduct internal attacks within the organisation to leverage more accounts. We therefore only needed to compromise one account, so we could potentially use that for additional attacks.

A total of 146 phishing attacks were executed against employees at the target organisation, and a total of 58 unique employees entered their credentials, amounting to a success rate of 39.7 pct. (which also included employees on vacation, as have been evident from a post-talk with the target organisation).

Table 25: Phishing on Target #3

Attack vector	Total attacks	Successful attacks	Success rate %
Phishing	146	58	40

In order for us to differentiate one user from another, we recorded the hashes of their username, which were automatically manipulated to maintain full anonymity; yet it allowed us to differentiate one user from another, so we could count the total number of unique users who fell victim to the phishing campaign. This was done by hashing the entire username to a 64-digit string, which was then reduced to containing only the first 8 characters. This process was automatic, and meant that we were unable to decode the hashed text string, while still maintaining the ability to differentiate between unique users.

Neither old nor new passwords were recorded in the test. We did, however, record whether or not information was entered into the password fields to make sure that we would have recorded something from the attempt. Had we collected the actual passwords, we would not only have the basis for attacking the individuals at their work place, but could as easily start running dictionary attacks on private accounts that had previously been identified. Insight into how a person constructs passwords makes it easier to conduct dictionary attacks, rather than a never-ending brute force attack.

Email correspondences with five target individuals were necessary to convince them to fill in the web form, as they were questioning whether or not it was relevant for them to do so. We do not know how many individuals followed protocol for changing their password locally on their computers (and not via the rogue web form that we sent to them). Nevertheless, we

do know that some employees followed the established company policy for changing their passwords.

Finally, three out of the total number of 146 target individuals were observant enough to contact the IT-department with a suspicion of the attack being a phishing attempt. In an effort to avoid unnecessary panic for Target #3's IT-department, it was agreed with our contact person that they would be notified of the social vulnerability assessment and that we would notify them when the phishing campaign began.

5.3.3.4 – Smishing

The smishing attacks were designed to convince select target individuals to click on a malicious link, which was sent via SMS. The attacks consisted of spoofing the sender as being the directors of four separate departments in Target #3. The smishing attacks were sent to nine victims in total, ranging from regular staff to middle management.

The tone-of-voice of the SMS was not of an urgent matter, but rather one that could be interpreted as something the recipient should be aware of, and which the respective spoofed directors wanted to inform their employees of. As demonstrated in table 26, eight out of the nine smishing attempts were successful, amounting to a success rate of 89 pct.

Table 26: Smishing on Target #3

Attack vector	Total attacks	Successful attacks	Success rate %
Smishing	9	8	89

Only one target individual did not click on the link. However, it should be noted that the individual in question, who did not click on the link, was informed of the attacks. It could therefore be argued that because the target individual was bias, due to this prior knowledge of the smishing campaign his results should not be included in the test results. Nevertheless, we decided to include the failure in the results, as this could represent an individual with a highly increased security consciousness.

The smishing campaign was carried out around 7PM, which is relevant to note, as this means the target individuals were most likely at home when the attacks occurred. In relation to that, there was a noteworthy development in the aftermath of the smishing attacks executed on Target #3, as the target

individuals began intercommunicating and quickly realised that they were under attack.

One target individual from middle management copied the entire text message - incl. the malicious link - into an email and sent it to all employees of the target organisation, warning them not to click on the link if they had received an SMS with this particular link. Despite the alerting email, an additional three individuals, not part of the smishing attacks clicked on the malicious link in the email. Most likely they clicked on the link without first reading the content if the email warning them not to click on the link. Nevertheless, this was interesting insight into the underpinning psychological motive for proactively engaging in shared content from trusted sources, prior to first evaluating the content. The subject of the email even read 'Red Alert' and the message was written partially in capital letters warning all users not to click on the link, were they to receive such SMS.

In a post-talk with the person who sent the alerting email with the malicious link included, the person realised that he should have taken a screen dump of his phone rather than copy the content into an email.

5.3.3.5 – USB Attack

The targets for the USB attack were selected exclusively on the results of the personality profiling and the social network analysis. A total of three USBs were configured to emulate a human interface device (HID), e.g. a keyboard, and were programmed to execute the malicious script as soon as it was plugged into the targets' computers. The script would open the *Run* command in Windows and type in the URL we wanted them to visit, in an effort to record the instance on the web server log. However, this proved more difficult than initially thought, due to unforeseen countermeasures at Target #3 in terms of operating procedures.

The initial plan was to establish an on-location meeting at Target #3, with the purpose of dropping the payload (the USB drive) at desirable spots at the targets' desks. However, as the environment cannot be controlled, and we cannot determine whether or not the designated targets would be the actual users of the USB drive, we opted for a different solution. The issue was that other employees could pick up the USB drive and insert them into their computers, which would be recorded as a legitimate attempt, but we would be unable to verify whether or not that would be the correct recipient – without compromising ethical boundaries. Additionally, there was

a fear that clients of the target organisation could potentially find the USB drives and insert them into their computers. This scenario would be beyond ethical and legal boundaries of the study. In sum, we needed a method that was highly targeted, and which ensured that *only* the intended recipients would be exposed to the malicious USB drives.

Therefore, we opted for an alternative scenario, where we would send a letter to each of the three target individuals at the target organisation, so that the attack would be strategically targeted to those particular individuals; thereby, eliminating most of the unknown factors. However, this proved to be an obstacle in the end, which resulted in three failed attempts. We enclosed letter in the envelope with a specific relevant logo and content constructed for the department of the respective targets, which differentiated from one target to the other.

For two of the three targets, the attempt was, however, quickly identified as coming from a fake organisation. The final attempt carried some mistrust due to the logo of the sender, which was a very close resemblance to an actual organisation that they know of in that particular department of the target organisation. Furthermore, the final attempt was intercepted by the supervisor of the individual we were targeting. The target in this case therefore never got a chance to review the enclosed letter and USB drive.

After consulting the individuals involved, we identified the errors in our attempt:

- We had to send the USB drive by mail
- We used fake organisations, which we were limited to, due to legal concerns
- We did not include phone numbers in any of the three attempts
- Awareness was heightened due to the recent phishing attempts

Table 27: USB Attack on Target #3

Attack vector	Total attacks	Successful attacks	Success rate %
USB Attack	3	0	0

5.3.3.6 – Summary

As can be seen in table 28, the results for the social vulnerability assessment of Target #3 were mixed. Only the USB attacks proved completely unsuc-

cessful, whereas the spear-phishing, phishing, smishing and whaling attacks all had successful counts.

In general, we were allowed to attack a larger target group at Target #3, compared to the other organisations who took part in this study.

Table 28: Attack Results on Target #3

Attack vector	Total attacks	Successful attacks	Success rate %
Spear-Phishing	3	3	100
Whaling	1	1	100
Phishing	146	58	40
Smishing	9	8	89
USB Attack	3	0	0
Total	162	70	43

The results of the SVA are listed in table 28 and illustrated in figure 27, which shows a high level of success for the spear-phishing, whaling and smishing attacks. There were a total number of 162 attacks executed against Target #3, and of the collective attacks that were carried out, a 43 pct. success rate was recorded.

Figure 27: Results on Target #3

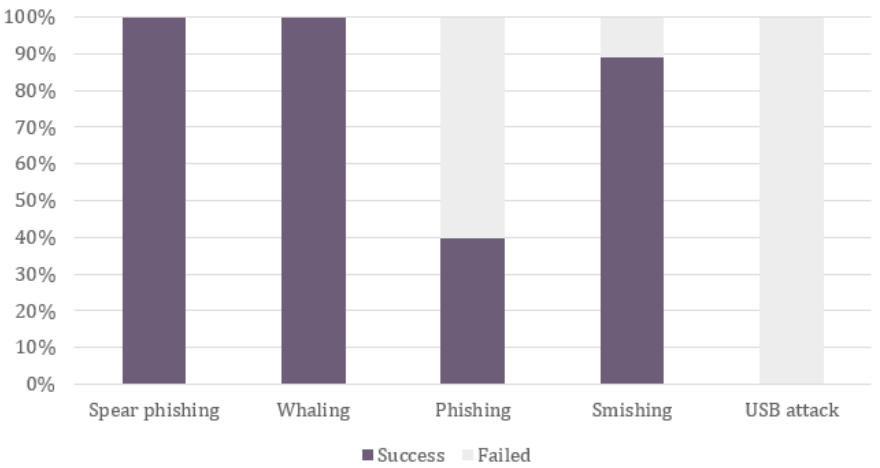

Target #3 has been briefed on all the information uncovered from OSINT during the reconnaissance phase, as well as made aware of the results of the social vulnerability assessment. After having finalised the study, we have been informed that Target #3 has implemented new security procedures and are currently working on an awareness training plan for its employees, as well as an incident response plan for mitigating the risk of social engineering attacks and cyber attacks in general.

5.4 – Comparative Overview of Results

5.4.1 – Overview

In table 29 is provided an overview of the executed attacks for all three participating target organisations. As it is clearly indicated in the table, the size of the designated target groups varied greatly for Target #1 and #2, in comparison with Target #3, where we had a greater target group in the social vulnerability assessment.

Table 29: Comparative Overview of Executed Attacks

Attack vector	Target #1	Target #2	Target #3	Total
Spear-Phishing	3	1	3	7
Whaling	1	1	1	3
Phishing	2	4	146	152
Smishing	3	5	9	17
PDF Attack	1	2	0	3
USB Attack	0	0	3	3
Total	10	13	162	185

As shown in figure 28 below, SMS was the overall most successful attack vector. A potential explanation for this could is the way people use and trust their smartphones which is arguably the most personal device we own, as we always have our smartphone with us.

Figure 28: Aggregated Results

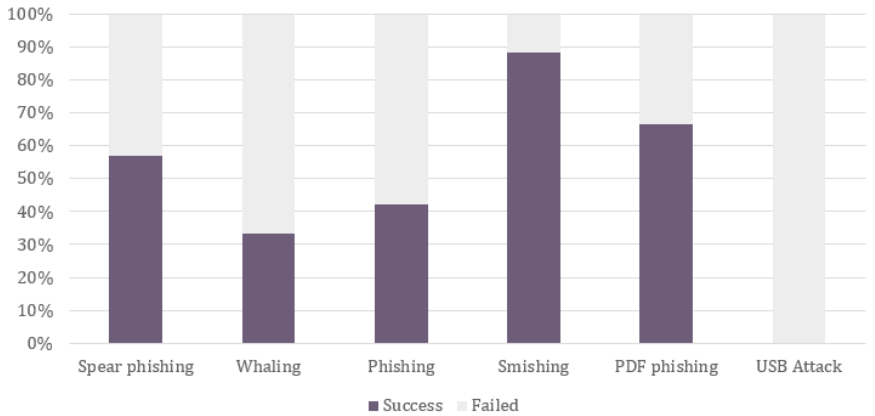

In general, users interact differently with their mobile devices than with their laptop and desktop computers. Often, a user will check emails while being at a meeting at work or watching TV at home. As a consequence, the user does not pay the same amount of attention, and will be more susceptible to a phishing attack while using the smartphone. Additionally, the user interface of many mobile devices does not allow hovering over a link in order to show the URL; this makes it harder for the users to detect whether or not an email is a phishing attempt.

Some users have multiple email accounts connected to their smartphones, some of them private and hence not protected by corporate email filters. So attackers might choose to attack employees through their personal email account instead of their corporate one. Additionally, mobile devices can receive both emails and SMS messages, and contrary to emails, SMS messages are usually not filtered by a spam filter, making it easier for smishing attempts to reach the user.

The amount of successful PDF attacks is also notable, as these require the target to not only be convinced of the legitimacy of the sender by using a spoofed email address, but it also requires the target to be convinced of the content, the context, the wording and the required of the target.

Finally, out of the collective 185 executed attacks, a mere seven individuals reported the attempt to their supervisor and/or responsible department, e.g. the IT-department. This is noteworthy as it gives an indication of how people in this study reacted to receiving phishing emails or alike. Either

they did not realise that they were subject to a cyber attack, or they simply did not bother to report it. This can be a problem on two levels: firstly, from a strategic level with the implemented cyber security policy of the respective organisations; and secondly, on the operational level, where it is not enforced that employees are required to report suspicious cyber activity or identified cyber attacks.

5.4.2 – Progression of Attacks

We will finalise the chapter with a visual overview of the progression of the attacks, as illustrated in figure 29. The purpose is to provide insights into the order of the executed attacks. The colouring of each spherical object signifies how successful each attack vector was - grey represents the failed attempts, while purple represents the successful attempts, from the entire field trials of the study.

Figure 29: Progression of SVA

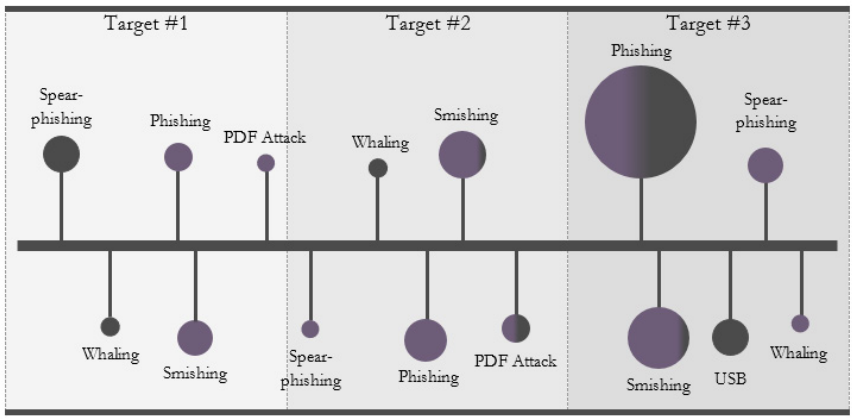

Chapter 6:
Dissemination

6. Dissemination

As outlined in the introduction, one of the objectives with Project SAVE was to disseminate the results at two workshops, which were held at the Danish Institute of Fire and Security Technology (DBI): (1) A national workshop consisting of invitees from the Danish security sector, primarily from defence and police, and (2) an international workshop, which was more broadly focused on industry and academia as well as international actors in the security sector.

The primary objective of the workshops was to disseminate the results, in an effort to engage government, industry and academia in a discussion on the phenomenon of social engineering 2.0. The secondary objective was to uncover the collective perception of social engineering 2.0 by the respective sectors, as to gain a deeper understanding of the *perceived* threat of social engineering.

In this chapter, we will cover the results of a survey sent to participants of the national and international workshops, and it will be supplied by results from a questionnaire regarding delivery methods for awareness training. Particularly the latter seeks to uncover the opinions of the participants in regard to various types of learning methods that can be applied for effective awareness training to counter social engineering attacks and heighten overall security consciousness of employees.

6.1 – Survey Results on Social Engineering

6.1.1 – General information

A total of 42 respondents participated in the survey that was distributed after the development of both the national and international workshop. A total of 32 respondents completed the survey by answering all of the questions.

Figure 30: Response Status

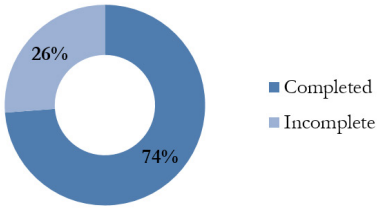

Figure 31: Type of Users

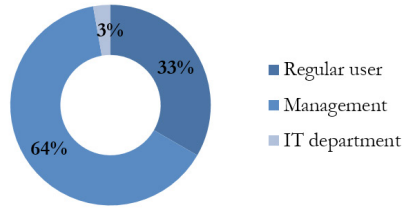

According to the initial answers of the survey, the following results are based predominantly on answers from management level, which are represented by 64 pct. of the respondents, as shown figure 31 above. The remaining users were regular users, amounting to 33 pct., and people from the participating organisations' IT-departments, represented by 3 pct. of the respondents from the national and international workshop.

6.1.2 – Experience with Social Engineering

54 pct. of the respondents answered that their company or institution had previously been subject to targeted social engineering attacks, while 20 pct. did not believe that they had been subject to attacks at all.

Figure 32: Experience with SE Attacks

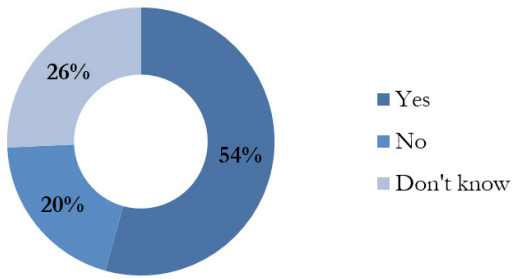

The respondents were also asked to answer whether or not they perceived social engineering as a threat to their organisation. Approximately 28 pct. of the respondents believed it was a high risk and 63 pct. answered a medium risk. A vast majority of the respondent thus recognised that social engineering attacks could pose a serious threat to their organisation, while merely 3 pct. of the respondent perceived it to have a low risk.

Figure 33 Perceived Threat of SE

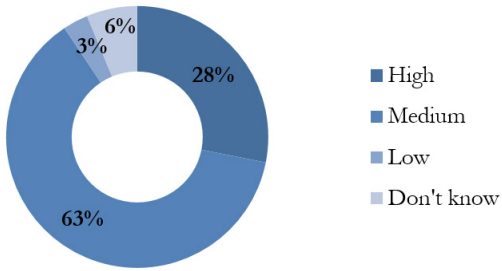

We believe that the data indicates that organisations are becoming more aware of the threat, and that the results constitute a level of recognition of social engineering as a risk factor.

We additionally wished to gain insight into whether or not companies and institutions specifically address social engineering in their IT-policy. 39 pct. of the respondents answered that social engineering was addressed in their IT-policy; whereas a majority of 46 pct. answered it was not addressed. 15 pct. was unaware of whether their IT-policy addressed it.

Figure 34: Does your IT-policy Address SE?

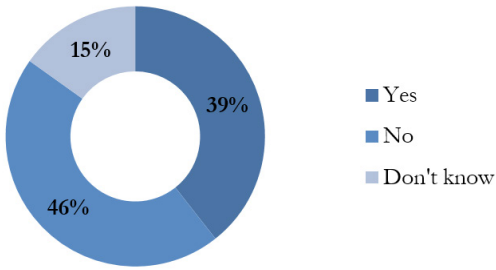

In summary, while social engineering is perceived to constitute a medium to high risk for the companies represented in the survey, and more than half of the respondents have had direct experience with being targeted, only 39 pct. to 54 pct. address social engineering directly in their IT-policy.

6.1.3 – Interest in Countermeasures

The final set of questions in the survey addressed the willingness of companies to uncover information about themselves, and their interest in testing and training their employees to counter social engineering attacks.

Figure 35: Interest in OSINT for Corporate SVA

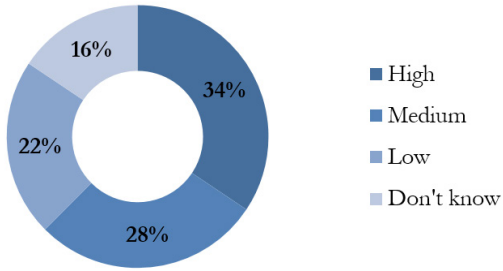

62 pct. of the respondents had a medium to high interest in uncovering information from open sources about their organisation, while only 22 pct. had little interest in the matter.

From this we can conclude that while most organisations do have an interest in uncovering information from OSINT, with the purpose of gaining insight into what type of information is publicly available about them, not all companies find this necessary. Correlating the results with the answers from whether or not organisations have experienced SE attacks, we see a close resemblance in the results: 20 pct. have no experience with SE attacks (cf. figure 32), while 22 pct. have little or no interest in gaining insight into information available about their organisation from open sources.

The respondents were also asked to answer, if they were interested in having their employees tested against common social engineering attacks, including phishing, smishing and vishing. 47 pct. had a medium/high interest in having their employees tested, while 28 pct. had low interest in the matter.

Figure 36: SE Teest of Employees

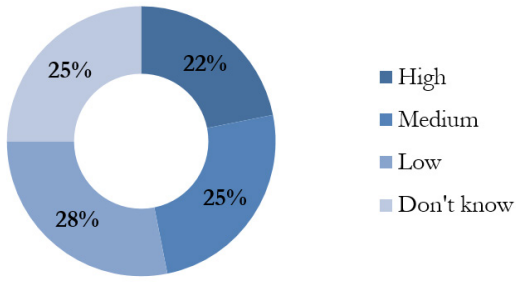

Most interesting is perhaps that 25 pct. of the respondents answered that they did not know whether or not they were interested in having their employees tested (cf. figure 36). The results from this are almost identical to the respondents' answers of whether or not they were interested in having their employees subjected to awareness training (cf. figure 37).

Figure 37: Interest in Awareness Training

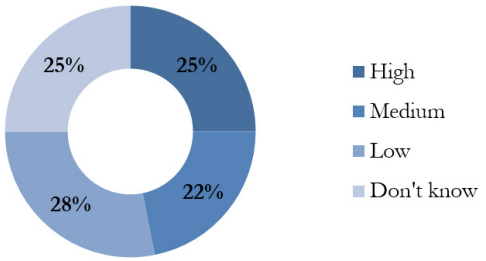

6.2 – Questionnaire on Awareness Training Methods

At the international workshop, the attendees were asked to answer a questionnaire relating to awareness training methods. In total, 21 respondents answered the questionnaire, of which 28 pct. were from management level, 28 pct. were from the organisations' IT-departments, and 42 pct. were regular users, as illustrated in figure 38.

Figure 38: Type of Users

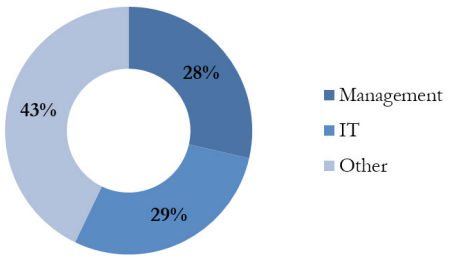

The median age of the respondents was 39,7 and 86 pct. of the respondents were male, while the remaining 14 pct. were female.

Of the 21 respondents, 48 pct. answered that they had previously received awareness training, while 52 pct. had not.

Figure 39: Previously Received Awareness Training

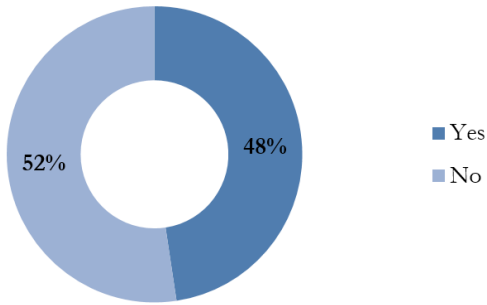

Of those who had received awareness training, the majority had received training in the form of leaflet/poster, video training or classroom training as illustrated in the results presented in figure 40.

Figure 40: Type of Awareness Training

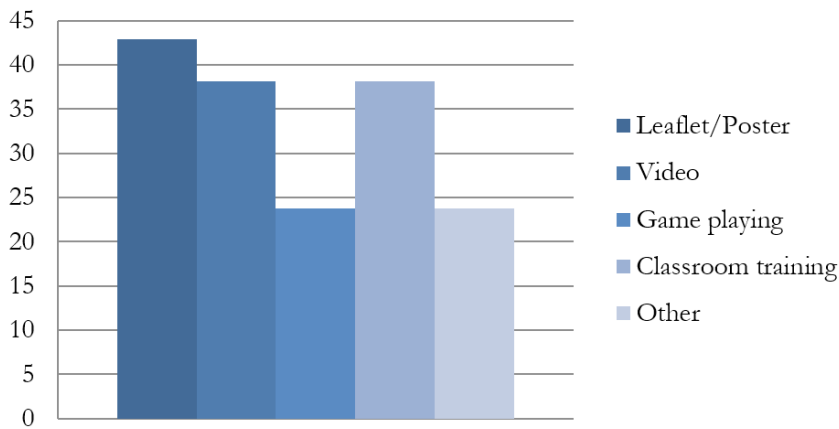

Some respondents have experience with serious games, digital or other creative delivery method for receiving awareness training, in relation to social engineering. As such, almost 43 pct. have had awareness training from leaflets, 38 pct. has received either training by video (e.g. webinar or otherwise) or classroom training.

Only 23 pct. have received other forms of training or have used serious games for *incidental learning*. Serious games typically employ some form

of incidental learning, i.e. where the players learn relevant information by playing an interactive game. Respondents were also asked, which delivery methods they perceived to have the biggest impact on their organisation, in relation to awareness training methods. More than 52 pct. of the respondents identified serious games as the delivery method with the biggest impact on their organisation, followed by classroom training amounting to 43 pct., as answered by the respondents. In addition, respondents were asked to consider, which delivery methods for awareness training they would personally prefer, and the results closely resembled the results on which delivery methods they perceive to have the biggest impact on their organisation, as illustrated in figure 41 and 42, respectively.

Figure 41: Methods with the Biggest Impact

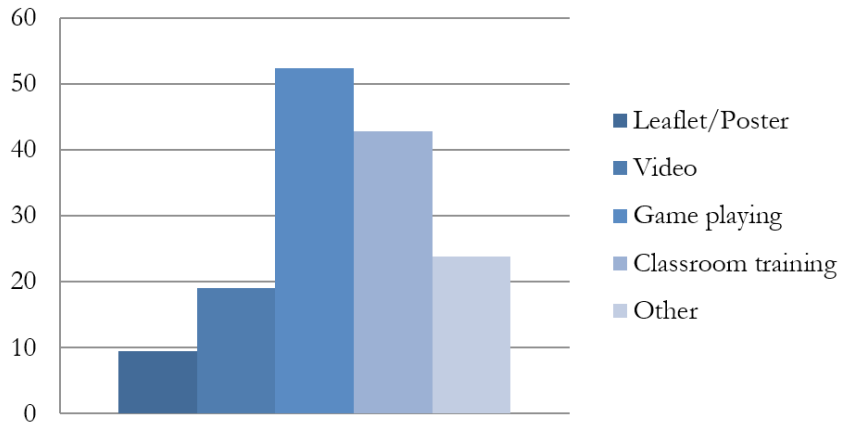

Finally, the respondents were asked to consider which teaching methods for awareness training they believed were the easiest to implement in their respective organisation.

Surprisingly, none answered serious games, which initially was perceived to have the biggest impact, as well as of personal preference to most of the respondents. Instead, classical methods, e.g. frontal classroom teaching, was believed to be the easiest to implement, which could be the result of known benefits from this type of training.

Figure 42: Personal Preference on Awareness

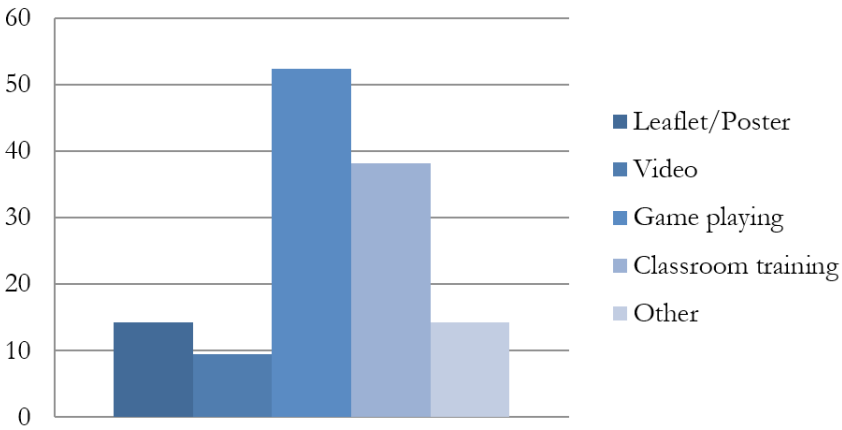

Leaflets/posters were the second highest method on the list. We believe the respondents answered that the classroom teaching and the leaflet/poster methods were the easiest to implement, as these are methods that the respondents have some form of experience with, making it easier for them to rely on the traditional methods over more innovative methods such as games, which for some may constitute an entirely new concept for receiving awareness training.

Figure 43: Easiest Methods to Implement

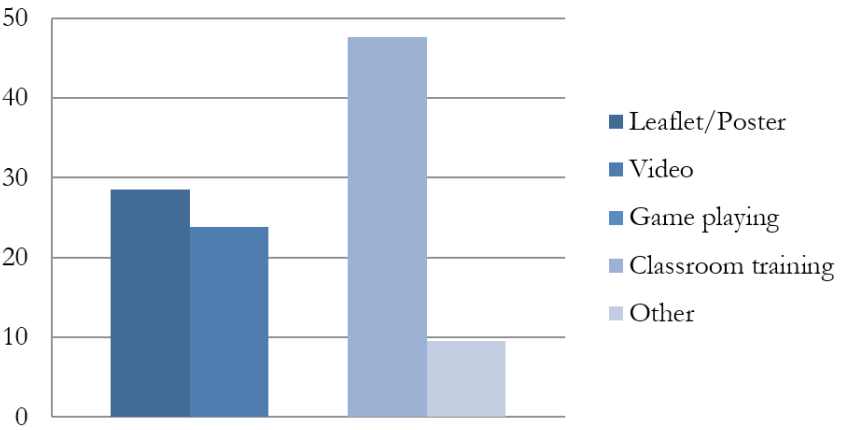

This concludes the chapter on the dissemination action associated with Project SAVE. In the following chapter we will address concluding remarks on the experiences and results gained from the study.

Chapter 7:
Conclusion

7. Conclusion

As outlined in the introduction, the objective of Project SAVE is three-fold: (1) To conduct an explorative investigation of social engineering 2.0, by performing simulated attacks in a real-life setting; (2) raise awareness, by disseminating the results to key stakeholders in industry, government and academia; and (3) to provide recommendations on how to mitigate the associated risks, by developing a social vulnerability assessment framework for assessing the organisational vulnerabilities relating to the human element of cyber security (cf. chapter 1).

The conclusion of this study will start by summarising the results of the simulated attacks, followed by the proposed framework for mitigating risks. The chapter will be finalised by providing recommendations for further studies on the phenomenon of social engineering, more specifically on topics that have not been covered in this study.

7.1 – Summary of Results

Three organisations have been targeted with a total of 185 social engineering 2.0 attacks, where various reconnaissance methods and attack vectors have been applied, which have been used to test their social vulnerability level. All three organisations were either directly part of critical infrastructure in Denmark, or have a supporting function to critical infrastructure.

Two out of three targets that were involved in the study had significant public information available about them from open sources, which were successfully utilised in the attacks conducted. However, the one with the least information available, namely Target #2, proved to be the one with the highest success rate during the attack phase, with an amount of 77 pct. successful attacks (cf. table 20). For Target #1 and Target #3 the number of successful attacks amounted to 60 pct. and 43 pct., respectively.

As covered in section 5.4, the most deceptive attack vector was SMS with an aggregated success rate of 88 pct., while the least successful was the USB vector, which did not deceive any of the recipients of the USB drives. We believe the deceptiveness of the SMS vector relies on the trust that people generally have for their smartphones, which can be considered a more personal device than a computer, since most people carry their phones with

them everywhere. We believe this to be the main reason for why people in general fall for smishing attempts. Additionally, most are unaware of the concept of SMS spoofing, which makes it easier to successfully trick people into conducting certain actions.

Of the 185 attacks executed against the three participating target organisations, only seven instances occurred, where an incident was reported to the responsible department or employee, e.g. the internal IT-department (cf. section 5.4.1). This illustrates how deceptive social engineering can be – even more so when considering that multiple attack vectors were applied simultaneously and against multiple targets. This should have raised the awareness level of employees, or perhaps raised a few red flags within the respective organisations that participated in the social vulnerability assessment. A real social engineering attack would only utilise one or two attack vectors, and only target a few individuals from each organisation, in order to avoid raising suspicion amongst the employees. It is important to note that an attacker would, in many cases, only need to compromise a single user in order to gain access, or in order to be able to escalate the attack from within the organisation's network.

The results of the study were disseminated at a national and an international workshop and were well received. From dialogues with participants at the workshops, we have determined that people in general were expecting the attempts to work, yet the overall success rate was alarming to most of the organisational representatives who were present at the workshops.

7.2 – Framework for Social Vulnerability Assessment

Based on the experiences from Project SAVE, we have constructed a framework for assessing the social vulnerabilities of an organisation – an assessment that we believe can provide a basis for understanding how vulnerable the human factor of a company is.

The framework is inspired by (1) the approach utilised in the SAVE study, and (2) by a report from the SANS Institute, named *A Multi-Level Defense Against Social Engineering (2003)*[61]. These frameworks both provide an in-depth strategy for securing against social engineering attacks, a strategy which we believe to be one of the most consistent compilation of useful advices on how to mitigate the risk of social engineering.

7.2.1 – A Multi-Level Defence Against Social Engineering

The following multi-level defence against the threat of social engineering covers four levels: (1) policy level, (2) parameter level, (3) persistence level, and (4) defensive level. These four levels, combined with the proactive SVA approach, establish a solid foundation for mitigating the risks of social engineering attacks (cf. section 7.2.2 for overview of the complete proposed framework). In other words, we believe the proposed framework will provide the necessary measures for countering most social engineering attacks.

7.2.1.1 – Policy Level

The foundation of cyber security in an organisation is the strategic level, of which the policies set the standards and level of security. This extends to more than mere IT-policies, and can include how to handle questionable request from customers calling the organisation's support lines[62].

A study in metacognition has established that increasing the confidence of employees, by providing clear policies, decreases the chance of an attacker influencing an employee to divulge otherwise restricted information[63].

The security policies must address several areas, including access, both physical and cyber-wise, procedures for setting up accounts, password changes, shredding of papers, locks, escorting of visitors, internal and external email policy and phone policies for partners, customers, and so forth. Most importantly is that the policies must be enforced, which requires organisational discipline[64].

7.2.1.2 – Parameter Level

Once the policy level has been implemented, the parameter level should be established. The parameter level includes security awareness training of all users[65]. This can effectively be divided into training of three different groups of employees: (1) management, (2) IT and HR, and (3) other users.

The reason for dividing the employees into groups is that each group typically has a different level of access to information. While management might have access to sensitive information, the IT-department might have systemic access, and HR might have access to confidential information about employees. Finally, other users such as customer support or sales representatives are faced with other attacks than the former groups, which creates basis for tailored training of each respective group of employees.

Good resistance training will help prevent employees from being talked into divulging critical or confidential information, and awareness training will teach them to identify social engineering attempts[66].

7.2.1.3 – Persistence Level

A complete defence against social engineering attacks requires regular reminders to *prime* and maintain a high level of security consciousness[67]. As such, we recommend not only continuous awareness training but also proactive testing, applying actual SE attempts, to maintain a required level of awareness of the methods and the threat it constitutes. Creativity sets the limit on the reminder implemented in an organisation.

A novel approach could be to apply serious games as a method to maintain a high level of awareness. Employees could be required to play a serious game for 10 minutes every week, in an effort to remind them of the dangers of SE and the various scenarios and attack vectors used in an attack, rather than receiving six hours of awareness training twice a year.

7.2.1.4 – Defensive Level

The final level is the defensive level, which revolves around an established incident response plan (IRP). This is a critical level, as a defined procedure for handling an identified attempt will let the employee know how to deal with suspicious behaviour or attempts from third parties[68]. There is a need for a well-defined procedure that can be set forth as soon as an attempt is identified.

If an organisation lacks a clear and well-defined incident response procedure, it can either cause panic (if everyone is notified of an incident) or it will leave it up to each individual employee to define an approach for handling the situation, which will decrease the effectiveness, as opposed to having a defined procedure that everyone strictly adhere to, with a central actor being the central point of contact for identified or suspected social engineering incidents.

An incident response plan typically includes six steps: (1) Preparation, which includes policies, (2) detection and identification of the incident, (3) containment of incident, (4) remediation of malware, (5) recovery from the attack, which could include resetting system before reintroducing them to the company network, and (6) reporting of the experience gained from the incident.

7.2.2 – Social Vulnerability Assessment Framework

Based on the methods covered in the previous section 7.2.1, the following framework presented in table 30 can be constructed, which acts as an overview of the guidelines:

Table 30: SVA Framework

Level	Description	Methods	Frequency
1. Assess Vulnerability Level	Using the same SVA methods applied in this study, could provide the basis for initial assessment of the vulnerability level in an organisation, using simulated attacks as the baseline for assessment.	• Reconnaissance • Phishing • Smishing • PDF attack • USB attack	Initial testing
2. Policy Level	Strategic policies that dictates how employees are to act in a given situation, providing the tools for the employee to counter social engineering attempts through policies and procedures.	• IT policy • Procedure guidelines • Business ethics	Update when needed
3. Parameter level	Awareness training of employees tailored to the specific user group they represent: (1) management, (2) IT & HR, and (3) other users. The awareness training should address each aspect of the interaction between people, both internally and externally, and take its departure from the policies implemented, acting as a guide for the awareness training.	• Classroom training • Webinar • Serious games • Leaflets • Video training	Every 3-6 months
4. Persistence Level	A complete defence against social engineering attacks requires regular reminders to maintain a high level of security consciousness. Serious games could be an option for this.	• Proactive testing • Awareness training • Reminders	Daily / weekly / monthly / quarterly

5. Defensive Level	It is critically important to have an incident response plan in place, so that employees know who they should contact, when an SE attempt has been identified, and the IT or security department knows how to handle the incident.	• Incident response procedure • Point of contact	Updated when required
6. Reassess Vulnerability Level	Reassessment of the vulnerability level of the organisation allows for measuring the effect of the implemented policies and the effect of the awareness training. The results can be compared with the initial testing conducted prior to the implementation.	• Reconnaissance • Phishing • Smishing • PDF attack • USB attack	Continuous testing

7.3 – Recommendations for Additional Research

As a results of the entire process and results showcased in this study, we recommend further studies to be conducted on the phenomenon of social engineering, both considering the conventional methods (cSE) as well as the evolved SE 2.0 methods and techniques. In appendix E, we have compiled a list of studies that relates to social engineering, which can provide the reader of this study additional material on the subject.

Our recommendations for additional research in the field are to further investigate three aspects of social engineering: (1) advanced reconnaissance methods, (2) the insider threat, and (3) reverse social engineering (rSE). Each represents an aspect of social engineering, which has not been covered in the current study, and which are complex enough to constitute a study on its own.

7.3.1 – Advanced Reconnaissance Methods

We believe further studies into the applied reconnaissance methods are relevant, in terms of gaining a greater understanding of how and where social engineers collect their intelligence on targets. Specifically, various methods of both tactical applications, e.g. electronic warfare (EW), as well as strategic, e.g. signals intelligence (SIGINT), have become increasingly available options for attackers. The technological development of software-defined radios (SDR) at affordable prices means that most are able to acquire

technology that can provide advanced signals intelligence, which in return can be applied in more advanced social engineering attacks.

Additionally, the availability of drones for commercial means are increasingly becoming a privacy concern, as they can be used in connection with the SDR for aerial SIGINT, which can fly over military bases, government buildings, airports, large corporations, etc., and collect vital data on cellular handsets, geo-mapping the whereabouts of personnel and conducting MITM-attacks. The capabilities that were previously reserved for the defence sector are increasingly becoming available for civilians, who can apply them for criminal activities.

Commercially available hardware can therefore be applied for tactical operations, both by rogue hacking teams as well as by state-sponsored groups. For these reasons and more, we believe further research into the vast array of possible reconnaissance methods should be subjected to continuously investigation, which can provide a dual purpose of both understanding reconnaissance methods applied for criminal activities, as well as enhance current methods for the intelligence and defence sector in an operational capacity.

7.3.2 – The Insider Threat

A demonstrator project could provide insight into whether or not it would be possible to forge a candidate's educational papers and professional record, in an effort to apply for a job in a pre-defined organisation/company that constitutes part of critical infrastructure, from where the candidate's objective would be to extract data from the target organisation. Measurable parameters could include how long (*time*) a candidate was able to continue the exfiltration, and how sensitive information (*lateral movement*) the candidate was able to acquire from various critical locations, offices or departments, within the target organisation.

7.3.3 – Reverse Social Engineering (rSE)

Perhaps one of the most complex methods that a social engineer can attempt is reverse social engineering (rSE). Reverse social engineering refers to the attacker setting the stage, so that instead of a social engineer being the one to approach the target, the target is the one who approaches the social engineer. This could provide basis for yet another interesting study on the subject of social engineering. A study on rSE is recommended to focus on the psychology of social engineering, including mapping of psychological

profiles, concepts of nudging, subliminal manipulation of people, and methods for deceptive stage setting.

7.4 – Final Comments

Social engineering remains a very real threat and risk for all layers of society, ranging from private individuals to corporations, to critical infrastructure and public institutions and governmental bodies.

Currently, social engineering seems to have free reign, particularly the rise in ransomware attacks is considered an indication of this, as the attacks almost doubled from 2013 to 2014[69]. However, when businesses start taking social engineering seriously and thus start to implement smart – not complex – defensive measures, to protect themselves against these methods, potential attacks will become more difficult for attackers to execute.

In relation to this, we believe the proposed framework can inspire for a smart implementation of effective counter-measures to mitigate the risk of social engineering attacks, and call for a change in the current business culture, with the purpose of raising the awareness level of employees, thus providing the necessary basis to stay ahead of the attackers.

8. Bibliography

[1] Mann, G. (2014, March 20th). forget the horse, this is the year of the F[ph]ish and the RAT. (G. Mann, Performer) London, United Kingdom.

[2] Granville, K. (2015, February 5). The New York Times. Retrieved March 29, 2016, from 9 Recent Cyberattacks Against Big Businesses: http://www.nytimes.com/interactive/2015/02/05/technology/recent-cyberattacks.html?_r=0

[3] Kaspersky. (n.d.). What is Social Engineering? Retrieved March 29, 2016, from Kaspersky Lab: http://usa.kaspersky.com/internet-security-center/definitions/social-engineering#.VsGf6zZZtHg

[4] Social Engineer. (n.d.). Security though education. Retrieved March 29, 2016, from The Social Engineering Framework: http://www.social-engineer.org/framework/psychological-principles/human-buffer-overflow/

[5] Statista. (n.d.). Leading social networks worldwide as of January 2016, ranked by number of active users (in millions). Retrieved March 29, 2016, from Staista - The statistics Portal: http://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/

[6] LinkedIn. (n.d.). LinkedIn Newsroom. Retrieved March 29, 2016, from About Us: https://press.linkedin.com/about-linkedin

[7] For further insight into the discussion on cyber warfare, please cf. "The Weaponization of Social Media" by T.E. Nissen, The Royal Danish Defence College, 2015

[8] Symantec Corporation. (2014, April). INTERNET SECURITY THREAT REPORT 2014. Retrieved 31 03, 2016, from http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v19_21291018.en-us.pdf

[9] Federal Trade Commission. (2014, May). Data Brokers - A Call for Transparency and Accountability. Retrieved from https://www.ftc.gov/system/files/documents/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014/140527databrokerreport.pdf

[10] Yadron, D. (2014, May 4). The Wall Street Journal. Retrieved from Symantec Develops New Attack on Cyberhacking: http://www.wsj.com/news/article_email/SB10001424052702303417104579542140235850578-lMyQjAxMTA0MDAwNTEwNDUyWj

[11] Hadnagy, C. J., & Kelly, P. F. (2011). Social Engineering: The Art Of Human Hacking, p. 9

[12] Cambridge Online Dictionary:
http://dictionary.cambridge.org/dictionary/english/social
http://dictionary.cambridge.org/dictionary/english/engineering

[13] Mitnick, K. D. (2003). The Art of Deception - Controlling the Human Element of Security. Indiana: Wiley Publishing, p. i

[14] Hadnagy, C. J., & Kelly, P. F. (2011). Social Engineering: The Art Of Human Hacking, p. 10

[15] Mitnick, K. D. (2003). The Art of Deception - Controlling the Human Element of Security. Indiana: Wiley Publishing, p. 16-17

[16] Byme, R. (1990). 637 Best Things Anybody Ever Said (1 ed.). Atheneum.

[17] Hadnagy, C. J., & Kelly, P. F. (2011). Social Engineering: The Art Of Human Hacking, p. xv

[18] Mitnick, K. D. (2003). The Art of Deception - Controlling the Human Element of Security. Indiana: Wiley Publishing, p. 16

[19] Goodchild, J. (2009, March 18). CSO Online. Retrieved March 29, 2016, from A Real Dumpster Dive: Bank Tosses Personal Data, Checks, Laptops: http://www.csoonline.com/article/2123810/identity-theft-prevention/a-real-dumpster-dive--bank-tosses-personal-data--checks--laptops.html

[20] The identification and designation of European critical infrastructures and the assessment of the, COUNCIL DIRECTIVE 2008/114/EC of 8 December 2008 (EU Directive December 8, 2008): http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2008:345:0075:0082:EN:PDF

[21] Zetter, K. (2016, January 1). Everything We Know About Ukraine's Power Plant Hack. Retrieved March 29, 2016, from Wired: http://www.wired.com/2016/01/everything-we-know-about-ukraines-power-plant-hack/

[22] Robertson, J., & Riley, M. (2016, January 14). How Hackers Took Down a Power Grid. Retrieved March 29, 2016, from Bloomberg Businessweek: http://www.bloomberg.com/news/articles/2016-01-14/how-hackers-took-down-a-power-grid

[23] BBC News. (2016, January 16). Hackers caused power cut in western Ukraine. Retrieved March 29, 2016, from BBC News - Technology: http://www.bbc.com/news/technology-35297464

[24] Lipovsky, R., & Cherepanov, A. (2016, January 4). BlackEnergy trojan strikes again: Attacks Ukrainian electric power industry. Retrieved March 29, 2016, from We Live Security - Security News, Views and Insight from the ESET experts: http://www.welivesecurity.com/2016/01/04/blackenergy-trojan-strikes-again-attacks-ukrainian-electric-power-industry/

[25] Bission, D. (2016, January 18). BlackEnergy Involved in Targeted Attack Against Boryspil Airport, Says Ukraine. Retrieved March 29, 2016, from The State Of Security: http://www.tripwire.com/state-of-security/latest-security-news/blackenergy-involved-in-targeted-attack-against-boryspil-airport-says-ukraine/

[26] Jones, D. (2016, January 20th). Tech NEWS World. Retrieved March 29, 2016, from Ukraine Mounts Investigation of Kiev Airport Cyberattack: http://www.technewsworld.com/story/83005.html

[27] Snow, J. (2016, February 11th). Kasper Sky. Retrieved March 29, 2016, from Medicine under fire: how to hack a hospital: https://blog.kaspersky.com/hacked-hospital/11296/

[28] Cox, J. (2016, February 12th). Motherboard. Retrieved March 29, 2016, from DOJ Hacker Also Accessed Forensic Reports and State Department Emails: http://motherboard.vice.com/read/doj-hacker-also-accessed-forensic-reports-and-state-department-emails?trk_source=recommended

[29] Shekhar, A. (2016, February 09). Hacker Leaks The Personal Information Of 20,000 FBI Agents. Retrieved March 29, 2016, from Fossbytes: http://

fossbytes.com/hackers-leaks-information-of-thousands-of-fbi-and-dhs-employees/

[30] Cox, J. (2016, February 7). Motherboard. Retrieved March 29, 2016, from Hacker Plans to Dump Alleged Details of 20,000 FBI, 9,000 DHS Employees: http://motherboard.vice.com/read/hacker-plans-to-dump-alleged-details-of-20000-fbi-9000-dhs-employees

[31] Krebs on Security. (2014, February 12th). Krebs On Security - In-depth security news and investigation. Retrieved March 29, 2016, from Posts Tagged: Fazio Mechanical Services: http://krebsonsecurity.com/tag/fazio-mechanical-services/

[32] Munson, L. (2014, March). ComputerWeekly.com. Retrieved March 29, 2016, from Target data breach: Why UK business needs to pay attention: http://www.computerweekly.com/feature/Target-data-breach-Why-UK-business-needs-to-pay-attention

[33] Krebs on Security. (2014, February 12th). Krebs on Security - In-depth security news and investigation. Retrieved from Email Attack on Vendor Set Up Breach at Target: http://krebsonsecurity.com/2014/02/email-attack-on-vendor-set-up-breach-at-target/#more-24313

[34] Krebs on Security. (2014, February 12th). Krebs on Security - In-depth security news and investigation. Retrieved from Email Attack on Vendor Set Up Breach at Target: http://krebsonsecurity.com/2014/02/email-attack-on-vendor-set-up-breach-at-target/#more-24313

[35] McGinty, K. M. (2015, February 26). Privacy & Security Matters. Retrieved March 29, 2016, from Target Data Breach Price Tag: $252 Million and Counting: https://www.privacyandsecuritymatters.com/2015/02/target-data-breach-price-tag-252-million-and-counting/

[36] Panda Security. (2016, February 3). Panda Mediacenter. Retrieved March 29, 2016, from Employees' selfies and the dangers of cybercrime for critical infrastructures: http://www.pandasecurity.com/mediacenter/security/employees-selfies-dangers/

[37] Langner, R. (2013, November). Langner. Retrieved March 29, 2016, from To Kill a Centrifuge - A Technical Analysis of What Stuxnet's Creators Tried

to Achieve, p. 6: http://www.langner.com/en/wp-content/uploads/2013/11/To-kill-a-centrifuge.pdf

[38] Keizer, G. (2010, September 16). ComputerWorld. Retrieved March 29, 2016, from Is Stuxnet the 'best' malware ever?: http://www.computerworld.com/article/2515757/malware-vulnerabilities/is-stuxnet-the--best--malware-ever-.html

[39] Google Support. (n.d.). Google. Retrieved March 29, 2016, from Advanced Search: https://support.google.com/websearch/answer/35890?hl=en&ref_topic=3081620

[40] MIT Libraries. (n.d.). MIT. Retrieved March 29, 2016, from Google Search Tips: Search Operators: http://libguides.mit.edu/c.php?g=176061&p=1159512

[41] Exploit Databse. (n.d.). Exploit Databse. Retrieved March 29, 2016, from Google Hacking Database (GHDB): https://www.exploit-db.com/google-hacking-database/

[42] Robotstxt. (n.d.). Robots Txt. Retrieved March 29, 2016, from About / robots.txt: http://www.robotstxt.org/robotstxt.html

[43] NISO. (2004). NISO. Retrieved March 29, 2016, from Understanding MetaData: http://www.niso.org/publications/press/UnderstandingMetadata.pdf

[44] Eleven Paths. (2015, May 22). Eleven Paths. Retrieved March 29, 2016, from FOCA: https://www.elevenpaths.com/labstools/foca/

[45] Halevi, T., Lewis, J., & Memon, N. (2013, February 8). Retrieved March 29, 2016, from Phishing, Personality Traits and Facebook: http://arxiv.org/pdf/1301.7643.pdf

[46] Halevi, T., Lewis, J., & Memon, N. (2013, February 8). Retrieved March 29, 2016, from Phishing, Personality Traits and Facebook: http://arxiv.org/pdf/1301.7643.pdf

[47] Please note that deep web refers to information not indexed on common search engines, whereas darknet refers to an online overlay network

that cannot be accessed using regular web browsers, and thus constitutes a 'hidden' network.

[48] Hauge, P. N., Hauge, N., & Morgan, C.-A. (2013). Market Research in Practice: How to Get Greater Insight From Your Market (Vol. 2). Kogan Page Publishers, p. 224

[49] WaybackMachine. (n.d.). archive.org. Retrieved March 29, 2016, from Internet Archive: https://archive.org/web/

[50] Cluley, G. (2016, March 16). Intego. Retrieved March 29, 2016, from Type a URL Wrong, and You Might End up with Malware on Your Mac: http://www.intego.com/mac-security-blog/type-a-url-wrong-and-you-might-end-up-with-malware-on-your-mac/

[51] Techopedia - Security. (n.d.). Techopedia; Dictionary - Whaling. Retrieved March 29, 2016, from Dictionay - Whaling: https://www.techopedia.com/definition/28643/whaling

[52] Techopedia - Security. (n.d.). Techopedia; Dictionary - Spoofing. Retrieved March 29, 2016, from Dictionary - Spoofing: https://www.techopedia.com/definition/5398/spoofing

[53] DePaul, N. (n.d.). Veracode. Retrieved March 29, 2016, from Spoofing Attack: IP, DNS & ARP: http://www.veracode.com/security/spoofing-attack

[54] Techopedia - Security. (n.d.). Techopedia; Dictionary - SMS Phishing. Retrieved March 29, 2016, from Dictionary - SMS Phishing: https://www.techopedia.com/definition/24898/sms-phishing

[55] Mohammed, A. (2013, August 1). InfoSec. Retrieved March 29, 2016, from Pharming Attack: http://resources.infosecinstitute.com/pharming-attack/

[56] Hadnagy, C. J., & Kelly, P. F. (2011). Social Engineering: The Art Of Human Hacking, p. 148

[57] Seymour, J. & Tully, P. (2016). Weaponizing data sciences for social engineering: Automated E2E spear phishing on Twitter, p. 7: https://www.blackhat.com/docs/us-16/materials/us-16-Seymour-Tully-Weaponizing-

Data-Science-For-Social-Engineering-Automated-E2E-Spear-Phishing-On-Twitter-wp.pdf

[58] Schmitt, D. P. (2002). British Journal of Social Psychology. Retrieved March 29, 2016, from A meta-analysis of sex differences in romantic attraction: Do rating contexts moderate tactic effectiveness judgments?: http://www.bradley.edu/dotAsset/165835.pdf

[59] Castle, S. (2007, March 18). The Independent. Retrieved March 29, 2016, from Thief woos bank staff with chocolates - then steals diamonds worth £14.5 mio.: http://www.independent.co.uk/news/world/europe/thief-woos-bank-staff-with-chocolates-then-steals-diamonds-worth-16314m-5332414.html

[60] RSA Fraud Action Research Lab (2011, April 1), Anatomy of an attack, online article, Retrieved March 29, 2016: https://blogs.rsa.com/anatomy-of-an-attack/

[61] Gragg, David (2003), A Multi-Level Defense Against Social Engineering, SANS Institute, InfoSec Reading Room

[62] Gragg, David (2003), A Multi-Level Defense Against Social Engineering, SANS Institute, InfoSec Reading Room, p. 10

[63] Petty et al. (2002), Thought Confidence as a Determinant of Persuasion: The Self-Validation Hypothesis, in Journal of Personality & Social Psychology, Vol. 82(5), May 2002, pp. 722-741, p. 722

[64] Gragg, David (2003), A Multi-Level Defense Against Social Engineering, SANS Institute, InfoSec Reading Room, p. 10

[65] Gragg, David (2003), A Multi-Level Defense Against Social Engineering, SANS Institute, InfoSec Reading Room, p. 11

[66] Gragg, David (2003), A Multi-Level Defense Against Social Engineering, SANS Institute, InfoSec Reading Room, p. 13

[67] Gragg, David (2003), A Multi-Level Defense Against Social Engineering, SANS Institute, InfoSec Reading Room, p. 15

[68] Gragg, David (2003), A Multi-Level Defense Against Social Engineering, SANS Institute, InfoSec Reading Room, p. 18

[69] Symantec (2015), Internet Security Threat Report (ISTR20), April 2015, Vol. 20: p. 92-93: https://www4.symantec.com/mktginfo/whitepaper/ISTR/21347932_GA-internet-security-threat-report-volume-20-2015-social_v2.pdf

9. Appendices

Appendix A: User Agents

Following explains the various sections of a user agent:

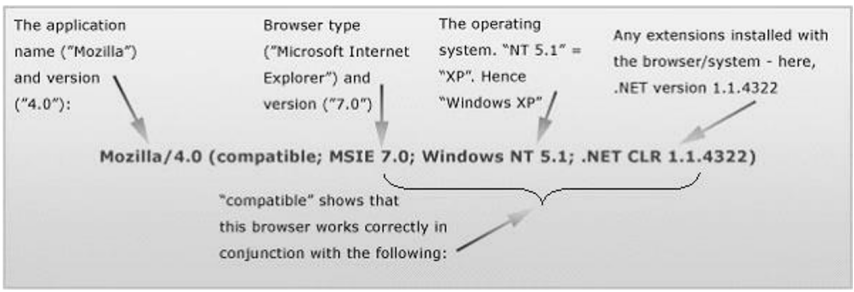

Following explains the process of how a website interprets the data in the user agent, which is to be used to circumvent countermeasures implemented in the Google Search Engine:

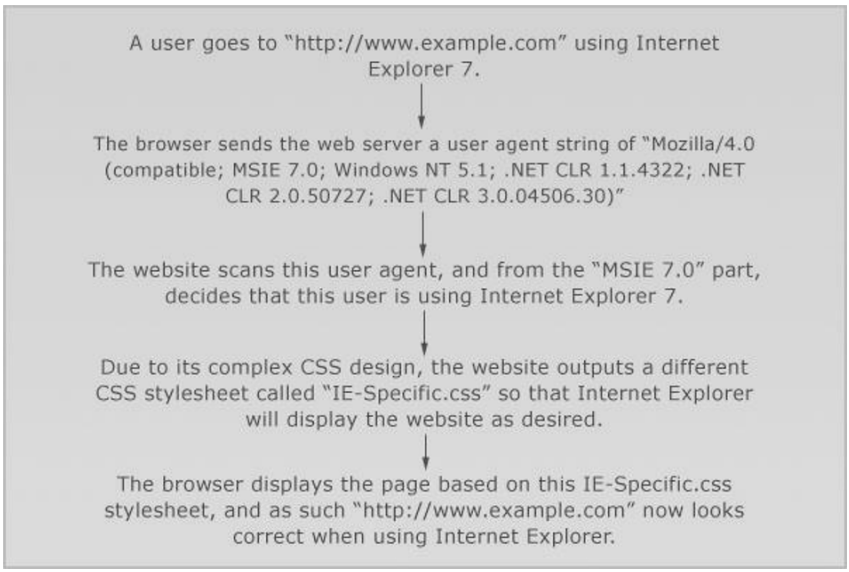

Appendix B: SNA Centrality

The following describes the various degree of *centrality* relating to a social network analysis (SNA). While this information is crucial for the analyst, we found it less relevant to discuss in the context of Project SAVE, and hence included it in the appendices.

Betweenness centrality measures the number of paths that pass through each entity. This can identify entities with the ability to control information flow between different parts of the network. These are sometimes referred to as *gatekeepers* (Carley (2005): 7). Gatekeepers might have many paths that run through them, which allow them to channel information to most of the other entities within the network. Alternatively they might have few paths, but still be a powerful facilitator, if they e.g. are positioned between different structurally important and key network clusters or subgroups (Carley (2005): 14).

Closeness centrality measures the proximity of an entity to the other entities in the network. An entity with a high measure of closeness centrality has the shortest paths to the other entities, allowing them to pass on and receive communication more quickly than anybody else in the network. Information travels further to and from an entity on the edge of a network that is attached to few other entities. These will have a lower measure of closeness centrality. Closeness centrality measures both direct and indirect closeness, where direct closeness is when two entities are connected by a link, and indirect closeness when information can pass only from one entity to another via a path that runs through one or more entities (IBM i2 Analyst Notebook).

Degree centrality is a key concept within SNA and denotes how centralised the network is. A highly centralised network is dominated by one (or few) person(s) who controls the information flow, and can become a single point of communication failure. A less centralised network has no single point of failure as more people have access to the same information and the same communication channels (Ibid; Carley (2005): 14).

Eigenvector measures how connected an entity is, and how much direct influence it might have over other connected entities, by considering the eigenvector scores of the entities that the initial agent is connected to, e.g.: a person with a high eigenvector score is likely to be at the center of a cluster

of key entities that themselves have high eigenvector scores. That person can communicate directly with those key entities compared with a person with a low eigenvector score located in the periphery of the network (IBM i2 Analyst Notebook).

Appendix C: Targeted Ads

Introduction
Social networks make much of their profit though advertisement, and very often they give advertisers the possibility to target specific audiences based on attributes such as age, gender and location. However, some social networks also allow advertisers to upload lists of existing customers that they want to target, making it possible to narrow the audience down to a few users', or just a single individual.

This was exploited by the American blogger, Brian Swichkow, in a prank he played on his roommate in 2014. Brian Swichkow created Facebook ads, and uploaded a list of target user's containing only one user: his roommate. He also made sure that the ads contained personal information about his roommate that only very few people knew. Brian Swichkow's blog contains more details on how this was done.

The Targeted ads-attack exploits the possibility of very narrow targeting of ads to launch a phishing attack against an individual or a small group of individuals. The goal is to show the targets ads containing a link, which (if clicked) will direct the target to a malicious website to phish his or her credentials.

The benefit of using ads on social networks for the attack compared to other channels is that the ad will be shown embedded in the layout of a site the user trusts, namely the social network. We believe that this will make the attack more credible than if the link to the malicious website was presented to the user e.g. in a phishing e-mail.

Description
First the attacker obtains e-mail addresses and/or social network profiles for the users that should be targeted. The attacker also prepares a malicious website that can phish credentials from users. This will most likely be done by imitating a well-known site that the target users are already using. Now, the attacker creates ads on a social network which links to the malicious website. When targeting the ads to users, use the different features of the social network's advertisement system to target the ad to the users desired targets, e.g. by just using a target list feature.

If uploading a list of targets is not possible on the social network, other features such as narrowing the target audience it to specific ages, locations, interests etc. could be used. However, this will probably result in the phishing ad being shown to more users than desired, which might make the attack more likely to get noticed, and will also make it more expensive, since the advertisers usually pay per time an ad is shown to a user.

The attack is successful if the user clicks the link in the ad and provides his credentials on the malicious website. However, we do also count the number of visitors that visited the website but did not provide their credentials.

Social Network

The social network used could be any social network that allows advertisers to target the ads narrowly enough such that they are only shown to very few users. The prank described in the preceding section shows that it is possible to do with Facebook since it is possible here to upload a list of users to target using the Power Editor. It is also possible to upload such a list when advertising via Twitter.

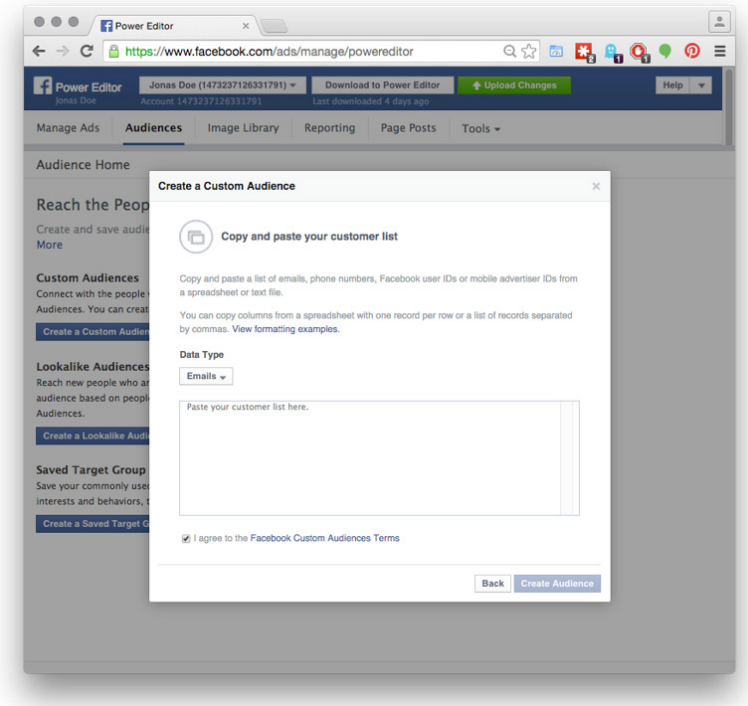

We also investigated whether it is possible to do the same in LinkedIn, which seems not to be the case. We have not investigated any other social networks.

Concerns

The social networks have a great interest in not directing users to malicious websites since this could damage the network's reputation. So it is likely that the networks will have some kind of automatic checking that the websites advertised on their site are not malicious. So, a successful attack will have to somehow circumvent this protection. Also, since an attack is successful only when a user clicks the link and trusts it the website enough that he provides his credentials on it, the website and the ad should be made such that the user wants to click it and provide his credentials on the site. How this is done depends on the targeted users and what credentials the attack is seeking to obtain.

References

- http://mysocialsherpa.com/the-ultimate-retaliation-pranking-my-roommate-with-targeted-facebook-ads/
- https://www.facebook.com/business/a/online-sales/custom-audiences.
- https://blog.twitter.com/2014/new-ways-to-create-and-use-tailored-audiences.

Appendix D: Web Server Log

The following is the documentation for the field trial testing with the results for each recorded attack. The first part of the line is a timestamp. The part of the form *T1PHa* or similar is an identifier - each attack was given a unique identifier in order to be able to distinguish one from another. The string following the identifier is 1/8 of the hash of the user's IP address, which is used to see if a user was compromised on several platforms. The last part is an optional data-parameter, which for most part was *null*, except for the phishing attacks where the user is redirected to a form to provide his credentials - here it is a hash of the username provided by the user.

```
Thu Dec 03 15:07:07 UTC 2015 (from T2PHa, f0OwhKs2): null
Thu Dec 03 15:07:11 UTC 2015 (from T2PHb, f0OwhKs2): null
Thu Dec 03 15:07:14 UTC 2015 (from T2PHc, f0OwhKs2): null
Thu Dec 03 15:07:17 UTC 2015 (from T2PHd, f0OwhKs2): null
Thu Dec 03 15:08:25 UTC 2015 (from T2SM, f0OwhKs2): null
Thu Dec 03 15:09:02 UTC 2015 (from T2SP, f0OwhKs2): null
Thu Dec 03 15:09:06 UTC 2015 (from T2W, f0OwhKs2): null
Thu Dec 03 15:09:22 UTC 2015 (from T2USB, f0OwhKs2): null
Thu Dec 03 15:51:59 UTC 2015 (from T1PHb, Ilzz8w0q): null
Thu Dec 03 15:52:24 UTC 2015 (from T1PHb, Ilzz8w0q): null
Thu Dec 03 15:54:06 UTC 2015 (from T1PHb, Ilzz8w0q): null
Thu Dec 03 16:43:36 UTC 2015 (from T1PHb, Ilzz8w0q): null
Thu Dec 03 16:43:51 UTC 2015 (from T1PHb, Ilzz8w0q): null
Thu Dec 03 18:11:45 UTC 2015 (from T1SMa, SWFH5M0H): null
Thu Dec 03 18:13:02 UTC 2015 (from T1SMb, 1BTGOVKQ): null
Thu Dec 03 18:13:23 UTC 2015 (from T1SMb, 1BTGOVKQ): null
Thu Dec 03 18:16:22 UTC 2015 (from T1SMc, RbYQVuk2): null
Thu Dec 03 18:18:43 UTC 2015 (from T1SMc, RbYQVuk2): null
Thu Dec 03 18:18:45 UTC 2015 (from T1SMc, RbYQVuk2): null
Thu Dec 03 18:18:46 UTC 2015 (from T1SMc, RbYQVuk2): null
Thu Dec 03 18:20:37 UTC 2015 (from T1SMc, RbYQVuk2): null
Thu Dec 03 18:20:56 UTC 2015 (from T1SMc, RbYQVuk2): null
```

```
Thu Dec 03 18:22:49 UTC 2015 (from T1SMc, RbYQVuk2): null
Thu Dec 03 18:22:52 UTC 2015 (from T1SMc, RbYQVuk2): null
Thu Dec 03 18:52:06 UTC 2015 (from T1SMa, SWFH5M0H): null
Thu Dec 03 21:15:59 UTC 2015 (from T1SMb, 1BTGOVKQ): null
Thu Dec 03 21:19:25 UTC 2015 (from T1SMb, 1BTGOVKQ): null
Fri Dec 04 07:55:36 UTC 2015 (from T1PHa, jCu0Rf4E): null
Fri Dec 04 07:55:36 UTC 2015 (from T1PHa, jCu0Rf4E): null
Fri Dec 04 11:50:23 UTC 2015 (from T1SMc, eHY0gS2j): null
Mon Dec 07 12:08:37 UTC 2015 (from T2SMe, lQ+fUoXv): null
Mon Dec 07 12:08:40 UTC 2015 (from T2SMf, lQ+fUoXv): null
Mon Dec 07 12:08:52 UTC 2015 (from T2SMg, lQ+fUoXv): null
Mon Dec 07 12:09:08 UTC 2015 (from T2SMg, lQ+fUoXv): null
Mon Dec 07 12:39:37 UTC 2015 (from T2PHc, 6aeen7jK): null
Mon Dec 07 12:40:38 UTC 2015 (from T2PHc, 6aeen7jK): null
Mon Dec 07 12:41:21 UTC 2015 (from T2PHb, Eao8aVnd): null
Mon Dec 07 12:52:15 UTC 2015 (from T2SMg, 3WKgnPbV): null
Mon Dec 07 12:54:00 UTC 2015 (from T2SP, 3WKgnPbV): null
Mon Dec 07 12:55:43 UTC 2015 (from T2SMf, 3WKgnPbV): null
Mon Dec 07 12:55:57 UTC 2015 (from T2SMf, 3WKgnPbV): null
Mon Dec 07 12:56:43 UTC 2015 (from T2SMf, woV65EyJ): null
Mon Dec 07 12:57:44 UTC 2015 (from T2SP, 3WKgnPbV): null
Mon Dec 07 12:59:57 UTC 2015 (from T2SP, 3WKgnPbV): null
Mon Dec 07 13:00:46 UTC 2015 (from T2SP, 3WKgnPbV): null
Mon Dec 07 13:12:44 UTC 2015 (from T2SMe, 6TNXXC+A): null
Mon Dec 07 14:01:15 UTC 2015 (from T2SMe, eHY0gS2j): null
Mon Dec 07 14:01:18 UTC 2015 (from T2SMe, eHY0gS2j): null
Mon Dec 07 14:01:19 UTC 2015 (from T2SMe, eHY0gS2j): null
Mon Dec 07 14:01:20 UTC 2015 (from T2SMe, eHY0gS2j): null
Mon Dec 07 14:01:20 UTC 2015 (from T2SMe, eHY0gS2j): null
Mon Dec 07 14:01:20 UTC 2015 (from T2SMe, eHY0gS2j): null
Mon Dec 07 14:01:21 UTC 2015 (from T2SMe, eHY0gS2j): null
Mon Dec 07 14:01:21 UTC 2015 (from T2SMe, eHY0gS2j): null
Mon Dec 07 14:01:21 UTC 2015 (from T2SMe, eHY0gS2j): null
```

```
Mon Dec 07 14:01:22 UTC 2015 (from T2SMe, eHY0gS2j): null
Mon Dec 07 14:01:22 UTC 2015 (from T2SMe, eHY0gS2j): null
Mon Dec 07 14:18:15 UTC 2015 (from T2PHd, KGtGNqE9): null
Mon Dec 07 14:18:54 UTC 2015 (from T2PHd, KGtGNqE9): null
Mon Dec 07 14:19:44 UTC 2015 (from T2PHd, KGtGNqE9): null
Mon Dec 07 14:20:16 UTC 2015 (from T2PHd, KGtGNqE9): null
Mon Dec 07 14:44:18 UTC 2015 (from T2PHa, 6aeen7jK): null
Mon Dec 07 14:44:26 UTC 2015 (from T2PHa, 6aeen7jK): null
Mon Dec 07 14:53:07 UTC 2015 (from T2SMe, lQ+fUoXv): null
Mon Dec 07 14:53:10 UTC 2015 (from T2SMf, lQ+fUoXv): null
Mon Dec 07 14:53:13 UTC 2015 (from T2SMg, lQ+fUoXv): null
Mon Dec 07 14:57:23 UTC 2015 (from T2SMh, lQ+fUoXv): null
Mon Dec 07 14:57:30 UTC 2015 (from T2SMi, lQ+fUoXv): null
Mon Dec 07 15:05:44 UTC 2015 (from T2PDF, lQ+fUoXv): null
Mon Dec 07 15:11:00 UTC 2015 (from T2PDF, lQ+fUoXv): null
Mon Dec 07 15:13:50 UTC 2015 (from T2PDF, +sDcojJs): null
Mon Dec 07 15:14:10 UTC 2015 (from T2PDF, +sDcojJs): null
Mon Dec 07 15:14:12 UTC 2015 (from T2PDF, +sDcojJs): null
Mon Dec 07 16:20:25 UTC 2015 (from T2SMh, +sDcojJs): null
Mon Dec 07 17:05:18 UTC 2015 (from T2SMe, 6TNXXC+A): null
Mon Dec 07 18:41:26 UTC 2015 (from T2PHd, kvauQhfQ): null
Mon Dec 07 18:43:12 UTC 2015 (from T2PDF, kvauQhfQ): null
Mon Dec 07 18:44:35 UTC 2015 (from T2PHd, kvauQhfQ): null
Mon Dec 07 18:47:37 UTC 2015 (from T2SMh, KGtGNqE9): null
Mon Dec 07 18:59:47 UTC 2015 (from T2PDF, kvauQhfQ): null
Tue Dec 08 07:32:42 UTC 2015 (from T2PHb, 6aeen7jK): null
Tue Dec 08 13:55:46 UTC 2015 (from T3SP, f0OwhKs2): null
Tue Dec 08 14:25:13 UTC 2015 (from T3SP, f0OwhKs2): null
Tue Dec 08 14:34:08 UTC 2015 (from T3SP, f0OwhKs2): null
Tue Dec 08 23:26:26 UTC 2015 (from T3SP, +aUfFE5e): null
Tue Dec 08 23:47:45 UTC 2015 (from T3SP, +aUfFE5e): null
Wed Dec 09 08:54:47 UTC 2015 (from T3USB, f0OwhKs2): null
Wed Dec 09 09:06:55 UTC 2015 (from T3USBa, f0OwhKs2): null
```

Wed Dec 09 09:07:06 UTC 2015 (from T3USBb, f0OwhKs2): null
Wed Dec 09 09:07:10 UTC 2015 (from T3USBc, f0OwhKs2): null
Wed Dec 09 15:33:53 UTC 2015 (from T3SP, f0OwhKs2): 688787d8
Wed Dec 09 15:34:10 UTC 2015 (from T3SP, f0OwhKs2): 688787d8
Wed Dec 09 15:34:23 UTC 2015 (from T3SP, f0OwhKs2): e797c0013
Wed Dec 09 15:39:45 UTC 2015 (from T3SP, f0OwhKs2): 688787d8
Wed Dec 09 15:40:14 UTC 2015 (from T3SP, f0OwhKs2): cb8379ac
Sun Dec 13 10:13:01 UTC 2015 (from T1SMa, ZEh50bC2): null
Mon Dec 14 20:28:24 UTC 2015 (from T3SP, +aUfFE5e): 707a3c64
Mon Dec 14 20:30:25 UTC 2015 (from T3SP, +aUfFE5e): 707a3c64
Mon Dec 14 20:31:02 UTC 2015 (from T3SP, +aUfFE5e): 9f584196
Mon Dec 14 20:35:50 UTC 2015 (from T3SP, +aUfFE5e): 707a3c64
Mon Dec 14 20:36:17 UTC 2015 (from T3SP, +aUfFE5e): 707a3c64c
Mon Dec 14 20:37:13 UTC 2015 (from T3SP, +aUfFE5e): 57917bceb
Mon Dec 14 20:37:39 UTC 2015 (from T3SP, +aUfFE5e): 688787d8
Mon Dec 14 22:22:00 UTC 2015 (from T3USBa, +aUfFE5e): null
Mon Dec 14 22:22:24 UTC 2015 (from T3USBa, +aUfFE5e): null
Mon Dec 14 22:22:26 UTC 2015 (from T3USBb, +aUfFE5e): null
Mon Dec 14 22:22:29 UTC 2015 (from T3USBa, +aUfFE5e): null
Mon Dec 14 22:22:30 UTC 2015 (from T3USBc, +aUfFE5e): null
Mon Dec 14 22:22:49 UTC 2015 (from T3USBa, +aUfFE5e): null
Mon Dec 14 22:22:50 UTC 2015 (from T3USBb, +aUfFE5e): null
Mon Dec 14 22:23:15 UTC 2015 (from T3USBa, +aUfFE5e): null
Mon Dec 14 22:23:16 UTC 2015 (from T3USBb, +aUfFE5e): null
Mon Dec 14 22:23:34 UTC 2015 (from T3USBa, +aUfFE5e): null
Mon Dec 14 22:23:36 UTC 2015 (from T3USBa, +aUfFE5e): null
Mon Dec 14 23:02:02 UTC 2015 (from T3USBa, +aUfFE5e): null
Mon Dec 14 23:02:51 UTC 2015 (from T3USBa, +aUfFE5e): null
Mon Dec 14 23:09:04 UTC 2015 (from T3USBb, +aUfFE5e): null
Mon Dec 14 23:12:53 UTC 2015 (from T3USBc, +aUfFE5e): null
Tue Dec 15 06:43:48 UTC 2015 (from T3SP, +aUfFE5e): ad3e69e9
Tue Dec 15 06:48:32 UTC 2015 (from T3PH, +aUfFE5e): 707a3c64c
Tue Dec 15 07:00:53 UTC 2015 (from T3SM1a, lQ+fUoXv): null

```
Tue Dec 15 07:01:01 UTC 2015 (from T3SM2a, lQ+fUoXv): null
Tue Dec 15 07:01:05 UTC 2015 (from T3SM2b, lQ+fUoXv): null
Tue Dec 15 07:01:09 UTC 2015 (from T3SM2c, lQ+fUoXv): null
Tue Dec 15 07:01:18 UTC 2015 (from T3SM3a, lQ+fUoXv): null
Tue Dec 15 07:01:21 UTC 2015 (from T3SM3b, lQ+fUoXv): null
Tue Dec 15 07:01:23 UTC 2015 (from T3SM3c, lQ+fUoXv): null
Tue Dec 15 07:01:32 UTC 2015 (from T3SM4a, lQ+fUoXv): null
Tue Dec 15 07:01:34 UTC 2015 (from T3SM4b, lQ+fUoXv): null
Tue Dec 15 07:02:10 UTC 2015 (from T3SPa, lQ+fUoXv): null
Tue Dec 15 07:02:14 UTC 2015 (from T3SPb, lQ+fUoXv): null
Tue Dec 15 07:02:16 UTC 2015 (from T3SPc, lQ+fUoXv): null
Tue Dec 15 07:02:26 UTC 2015 (from T3W, lQ+fUoXv): null
Tue Dec 15 08:17:07 UTC 2015 (from T3SM1a, +sDcojJs): null
Tue Dec 15 08:17:15 UTC 2015 (from T3SM2a, +sDcojJs): null
Tue Dec 15 08:17:19 UTC 2015 (from T3SM2b, +sDcojJs): null
Tue Dec 15 08:17:21 UTC 2015 (from T3SM2c, +sDcojJs): null
Tue Dec 15 08:17:25 UTC 2015 (from T3SM3a, +sDcojJs): null
Tue Dec 15 08:17:31 UTC 2015 (from T3SM3b, +sDcojJs): null
Tue Dec 15 08:17:35 UTC 2015 (from T3SM3c, +sDcojJs): null
Tue Dec 15 08:17:39 UTC 2015 (from T3SM4a, +sDcojJs): null
Tue Dec 15 08:17:43 UTC 2015 (from T3SM4b, +sDcojJs): null
Tue Dec 15 08:21:12 UTC 2015 (from T3SPa, lQ+fUoXv): null
Tue Dec 15 08:24:26 UTC 2015 (from T3W, +sDcojJs): null
Tue Dec 15 08:25:04 UTC 2015 (from T3SPa, lQ+fUoXv): null
Tue Dec 15 08:26:27 UTC 2015 (from T3W, +sDcojJs): null
Tue Dec 15 08:28:18 UTC 2015 (from T3SPb, +sDcojJs): null
Tue Dec 15 08:28:18 UTC 2015 (from T3SPb, +sDcojJs): null
Tue Dec 15 08:28:22 UTC 2015 (from T3SPc, +sDcojJs): null
Tue Dec 15 08:28:26 UTC 2015 (from T3SPa, +sDcojJs): null
Tue Dec 15 08:35:12 UTC 2015 (from T3SPa, +sDcojJs): null
Tue Dec 15 08:35:16 UTC 2015 (from T3SPb, +sDcojJs): null
Tue Dec 15 08:35:16 UTC 2015 (from T3SPb, +sDcojJs): null
Tue Dec 15 08:35:19 UTC 2015 (from T3SPc, +sDcojJs): null
```

```
Tue Dec 15 08:35:21 UTC 2015 (from T3SPc, +sDcojJs): null
Tue Dec 15 08:35:26 UTC 2015 (from T3W, +sDcojJs): null
Tue Dec 15 08:58:50 UTC 2015 (from T3PH, 2+niBKop): 707a3c64c
Tue Dec 15 09:57:32 UTC 2015 (from T3W, +sDcojJs): null
Tue Dec 15 09:59:34 UTC 2015 (from T3W, 2+niBKop): null
Tue Dec 15 10:00:25 UTC 2015 (from T3PH, 2+niBKop): 707a3c64c
Tue Dec 15 10:11:56 UTC 2015 (from T3SPa, 2+niBKop): null
Tue Dec 15 10:12:00 UTC 2015 (from T3SPb, 2+niBKop): null
Tue Dec 15 10:12:00 UTC 2015 (from T3SPb, 2+niBKop): null
Tue Dec 15 10:12:02 UTC 2015 (from T3SPa, 2+niBKop): null
Tue Dec 15 10:12:04 UTC 2015 (from T3SPc, 2+niBKop): null
Tue Dec 15 10:12:04 UTC 2015 (from T3SPc, 2+niBKop): null
Tue Dec 15 10:12:40 UTC 2015 (from T3SPc, 2+niBKop): null
Tue Dec 15 10:13:16 UTC 2015 (from T3W, 2+niBKop): null
Tue Dec 15 12:49:08 UTC 2015 (from T3PH, +sDcojJs): 707a3c64c
Tue Dec 15 12:53:16 UTC 2015 (from T3PH, +sDcojJs): 707a3c64c
Tue Dec 15 14:44:12 UTC 2015 (from T3PH, +sDcojJs): 2a517c2f6d
Tue Dec 15 14:45:05 UTC 2015 (from T3PH, xfyyq6L7): f51403d4
Tue Dec 15 14:45:18 UTC 2015 (from T3PH, +sDcojJs): 9f5997d4
Tue Dec 15 14:45:40 UTC 2015 (from T3PH, uljlpwIV): c18c0e43
Tue Dec 15 14:46:08 UTC 2015 (from T3PH, +sDcojJs): 5bd1269b
Tue Dec 15 14:46:11 UTC 2015 (from T3PH, +sDcojJs): 9d68f3c
Tue Dec 15 14:47:17 UTC 2015 (from T3PH, +sDcojJs): 29d9489
Tue Dec 15 14:49:07 UTC 2015 (from T3PH, TVCU1cAC): 18889a721
Tue Dec 15 14:49:53 UTC 2015 (from T3PH, +sDcojJs): f16eef839
Tue Dec 15 14:50:52 UTC 2015 (from T3PH, JXMOLZAe): 569ec613
Tue Dec 15 14:51:55 UTC 2015 (from T3PH, +sDcojJs): fb1edcf28
Tue Dec 15 14:52:19 UTC 2015 (from T3PH, 5NRqZMWo): a1a3c4e4c
Tue Dec 15 14:53:55 UTC 2015 (from T3PH, +sDcojJs): bfb2f0821
Tue Dec 15 14:54:23 UTC 2015 (from T3PH, +sDcojJs): 84d10d96
Tue Dec 15 14:56:08 UTC 2015 (from T3PH, +sDcojJs): e91f904c
Tue Dec 15 14:57:25 UTC 2015 (from T3PH, +sDcojJs): dc80b8c0
Tue Dec 15 14:57:51 UTC 2015 (from T3PH, +sDcojJs): 1b18bf1f3
```

```
Tue Dec 15 15:00:45 UTC 2015 (from T3PH, +sDcojJs): 58d4532d42
Tue Dec 15 15:04:08 UTC 2015 (from T3PH, +sDcojJs): cd7fd95e0
Tue Dec 15 15:11:08 UTC 2015 (from T3PH, +sDcojJs): 8ee6a3d6
Tue Dec 15 15:13:30 UTC 2015 (from T3PH, +sDcojJs): 652742992
Tue Dec 15 15:23:18 UTC 2015 (from T3PH, xHZ4jiku): 6b9cd7872
Tue Dec 15 15:23:43 UTC 2015 (from T3PH, +sDcojJs): 593f2d04aa
Tue Dec 15 15:25:02 UTC 2015 (from T3PH, LH8/lXBj): 3cfbea0930
Tue Dec 15 15:26:04 UTC 2015 (from T3PH, KJuCYI5y): d6e82fbe4a
Tue Dec 15 15:37:19 UTC 2015 (from T3PH, +sDcojJs): d9fd4a51a4
Tue Dec 15 15:50:24 UTC 2015 (from T3PH, +sDcojJs): 49de23cab
Tue Dec 15 16:13:52 UTC 2015 (from T3PH, +sDcojJs): 250fccd5f5f
Tue Dec 15 16:17:06 UTC 2015 (from T3PH, 4IDjqV6M): 266294f7
Tue Dec 15 16:18:49 UTC 2015 (from T3PH, +sDcojJs): 10b1df1db
Tue Dec 15 16:29:38 UTC 2015 (from T3PH, +sDcojJs): d182efaf7e
Tue Dec 15 16:36:31 UTC 2015 (from T3PH, +sDcojJs): 1d5c72881c
Tue Dec 15 18:25:47 UTC 2015 (from T3SM2b, 6bqw3m84): null
Tue Dec 15 18:26:21 UTC 2015 (from T3SM2c, Ani275NC): null
Tue Dec 15 18:26:43 UTC 2015 (from T3SM2b, 6bqw3m84): null
Tue Dec 15 18:26:50 UTC 2015 (from T3SM2c, Ani275NC): null
Tue Dec 15 18:27:03 UTC 2015 (from T3SM2b, 6bqw3m84): null
Tue Dec 15 18:27:06 UTC 2015 (from T3SM3c, dPjzG62L): null
Tue Dec 15 18:27:24 UTC 2015 (from T3SM2c, Ani275NC): null
Tue Dec 15 18:28:29 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:30:32 UTC 2015 (from T3SM2c, Ani275NC): null
Tue Dec 15 18:31:05 UTC 2015 (from T3SM2c, Ani275NC): null
Tue Dec 15 18:33:24 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:33:47 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:33:50 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:33:58 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:34:10 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:34:29 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:34:35 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:34:52 UTC 2015 (from T3SM3a, TTXY1U2U): null
```

```
Tue Dec 15 18:35:13 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:35:53 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:35:56 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:36:39 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:36:42 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:37:05 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:37:34 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:37:37 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:38:12 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:38:14 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:38:15 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:38:16 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:38:16 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:38:17 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:39:16 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:40:25 UTC 2015 (from T3PH, +sDcojJs): b08f256307
Tue Dec 15 18:44:24 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:44:38 UTC 2015 (from T3SM3a, TTXY1U2U): null
Tue Dec 15 18:47:47 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:48:17 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:50:10 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 18:59:09 UTC 2015 (from T3SM2c, Ani275NC): null
Tue Dec 15 19:04:15 UTC 2015 (from T3SM2c, Ani275NC): null
Tue Dec 15 19:05:50 UTC 2015 (from T3SM2c, Ani275NC): null
Tue Dec 15 19:06:14 UTC 2015 (from T3SM2c, Ani275NC): null
Tue Dec 15 19:06:30 UTC 2015 (from T3SM2c, Ani275NC): null
Tue Dec 15 19:08:32 UTC 2015 (from T3SM2a, fKU3CF9f): null
Tue Dec 15 19:15:15 UTC 2015 (from T3SM4b, 1xrGjVmo): null
Tue Dec 15 19:20:50 UTC 2015 (from T3PH, +sDcojJs): 3a8584e70
Tue Dec 15 19:22:22 UTC 2015 (from T3W, +aUfFE5e): null
Tue Dec 15 19:23:16 UTC 2015 (from T3W, dG5Sk0po): null
Tue Dec 15 19:24:43 UTC 2015 (from T3W, +aUfFE5e): null
Tue Dec 15 19:25:29 UTC 2015 (from T3PH, +sDcojJs): bf1b6501ca
```

```
Tue Dec 15 19:25:40 UTC 2015 (from T3W, dG5Sk0po): null
Tue Dec 15 19:31:11 UTC 2015 (from T3SM2b, 4IDjqV6M): null
Tue Dec 15 19:53:54 UTC 2015 (from T3PH, Lxvdiy0E): 6a0ac0fd972c
Tue Dec 15 19:56:04 UTC 2015 (from T3SM4a, 6CO0M++2): null
Tue Dec 15 19:57:41 UTC 2015 (from T3PH, Lxvdiy0E): 6a0ac0fd9
Tue Dec 15 20:02:40 UTC 2015 (from T3PH, aVrP65VC): a083ca2d0c
Tue Dec 15 20:53:18 UTC 2015 (from T3SM2b, xfyyq6L7): null
Tue Dec 15 21:49:50 UTC 2015 (from T3PH, mdtjC+MD): 5b98c02a
Tue Dec 15 22:46:05 UTC 2015 (from T3SM1a, CqvTXDS5): null
Wed Dec 16 05:33:50 UTC 2015 (from T3PH, +sDcojJs): 6beca1ae
Wed Dec 16 05:52:24 UTC 2015 (from T3PH, +sDcojJs): e021d7d48
Wed Dec 16 06:20:59 UTC 2015 (from T3PH, +sDcojJs): 3beb1abb
Wed Dec 16 06:29:54 UTC 2015 (from T3SM2b, vanxlBWD): null
Wed Dec 16 06:31:33 UTC 2015 (from T3PH, +sDcojJs): 7b7d41876
Wed Dec 16 06:31:52 UTC 2015 (from T3PH, +sDcojJs): 8c962f45f
Wed Dec 16 06:33:11 UTC 2015 (from T3PH, +sDcojJs): 8c962f45f
Wed Dec 16 06:43:39 UTC 2015 (from T3PH, TVCU1cAC): 7cfb92e6ce
Wed Dec 16 06:48:43 UTC 2015 (from T3PH, +sDcojJs): a7a94faf02
Wed Dec 16 06:53:48 UTC 2015 (from T3PH, +sDcojJs): f541ab1317
Wed Dec 16 07:05:18 UTC 2015 (from T3PH, +sDcojJs): 4ca310908a
Wed Dec 16 07:08:52 UTC 2015 (from T3PH, /RFddoAL): 8d64666b3
Wed Dec 16 07:09:51 UTC 2015 (from T3PH, +sDcojJs): 2157db6d
Wed Dec 16 07:20:05 UTC 2015 (from T3PH, +sDcojJs): 10e4be30e
Wed Dec 16 07:38:29 UTC 2015 (from T3PH, lValcQzK): c2e7e7b1d
Wed Dec 16 07:48:04 UTC 2015 (from T3PH, v1cGPs/i): f51403d4b0
Wed Dec 16 08:02:39 UTC 2015 (from T3PH, +sDcojJs): f0d2be946
Wed Dec 16 08:08:05 UTC 2015 (from T3PH, +sDcojJs): b6616a3e97
Wed Dec 16 08:40:58 UTC 2015 (from T3PH, +sDcojJs): c1199d80c
Wed Dec 16 09:17:32 UTC 2015 (from T3PH, +sDcojJs): 1f34503f65
Wed Dec 16 09:40:46 UTC 2015 (from T3PH, +sDcojJs): cf3abc8326
Wed Dec 16 10:56:08 UTC 2015 (from T3PH, +sDcojJs): 37fb527b7
Wed Dec 16 11:43:20 UTC 2015 (from T3PH, +sDcojJs): 0023a085
Thu Jan 07 08:43:29 UTC 2016 (from T3PH, +sDcojJs): 78dba6d45
```

```
Thu Jan 07 08:44:36 UTC 2016 (from T3SM2b, +sDcojJs): null
Wed Jan 13 12:47:38 UTC 2016 (from T3USBa, +sDcojJs): null
Wed Jan 13 12:49:51 UTC 2016 (from T3USBb, +sDcojJs): null
Wed Jan 13 12:50:14 UTC 2016 (from T3USBc, +sDcojJs): null
Wed Jan 13 14:54:26 UTC 2016 (from T3USBa, +sDcojJs): null
Wed Jan 13 15:21:48 UTC 2016 (from T3USBb, +sDcojJs): null
Wed Jan 13 15:30:37 UTC 2016 (from T3USBc, +sDcojJs): null
Mon Jan 18 11:10:05 UTC 2016 (from T3SM2a, p6/empje): null
Tue Jan 19 10:12:59 UTC 2016 (from T3SM1a, 3380bpWF): null
Tue Jan 19 10:13:33 UTC 2016 (from T3W, 3380bpWF): null
Tue Jan 19 10:13:50 UTC 2016 (from T3SPa, 3380bpWF): null
Tue Jan 19 10:13:55 UTC 2016 (from T3SPb, 3380bpWF): null
Tue Jan 19 10:13:55 UTC 2016 (from T3SPb, 3380bpWF): null
Tue Jan 19 10:14:00 UTC 2016 (from T3SPc, 3380bpWF): null
Tue Jan 19 10:15:26 UTC 2016 (from T3W, 3380bpWF): null
Thu Jan 21 15:51:58 UTC 2016 (from T3SPa, goZ7NGSE): null
Thu Jan 21 16:11:22 UTC 2016 (from T3W, goZ7NGSE): null
Sun Jan 24 18:26:08 UTC 2016 (from T3W, +aUfFE5e): null
Sun Jan 24 18:26:29 UTC 2016 (from T3SPa, +aUfFE5e): null
Sun Jan 24 18:26:30 UTC 2016 (from T3SPa, +aUfFE5e): null
Sun Jan 24 18:26:41 UTC 2016 (from T3SPb, +aUfFE5e): null
Sun Jan 24 18:26:45 UTC 2016 (from T3SPc, +aUfFE5e): null
Sun Jan 24 18:31:46 UTC 2016 (from T3W, +aUfFE5e): null
Sun Jan 24 18:43:55 UTC 2016 (from T3SPa, +aUfFE5e): null
Sun Jan 24 18:43:57 UTC 2016 (from T3SPb, +aUfFE5e): null
Sun Jan 24 18:44:00 UTC 2016 (from T3SPc, +aUfFE5e): null
Sun Jan 24 18:47:40 UTC 2016 (from T3W, +aUfFE5e): null
Sun Jan 24 18:49:17 UTC 2016 (from T3W, +aUfFE5e): null
Sun Jan 24 18:53:34 UTC 2016 (from T3SPa, +aUfFE5e): null
Sun Jan 24 18:57:22 UTC 2016 (from T3SPb, +aUfFE5e): null
Sun Jan 24 18:57:26 UTC 2016 (from T3SPc, +aUfFE5e): null
Sun Jan 24 18:57:29 UTC 2016 (from T3W, +aUfFE5e): null
Sun Jan 24 18:57:43 UTC 2016 (from T3W, +aUfFE5e): null
```

```
Sun Jan 24 18:57:45 UTC 2016 (from T3W, +aUfFE5e): null
Sun Jan 24 18:57:46 UTC 2016 (from T3W, +aUfFE5e): null
Sun Jan 24 19:40:49 UTC 2016 (from T3SPa, 6oDgfsbb): null
Sun Jan 24 19:41:09 UTC 2016 (from T3SPa, 6oDgfsbb): null
Sun Jan 24 19:41:44 UTC 2016 (from T3SPa, 6oDgfsbb): null
Sun Jan 24 20:49:57 UTC 2016 (from T3SPc, N/RIs3o+): null
Sun Jan 24 21:41:57 UTC 2016 (from T3SPb, ism4cGiy): null
Sun Jan 24 21:48:07 UTC 2016 (from T3SPb, ism4cGiy): null
Mon Jan 25 06:48:57 UTC 2016 (from T3W, +sDcojJs): null
Tue Jan 26 05:10:00 UTC 2016 (from T3SPc, LhGjv6M7): null
Wed Feb 03 21:16:23 UTC 2016 (from T1SPa, +aUfFE5e): USER-PASS
```

Appendix E: Additional Literature

[1] W. Pieters, D. Hadziosmanovi´c, and F. Dechesne, "Security-by-experiment: Lessons from responsible deployment in cyberspace," Science and Engineering Ethics, vol. N/A, no. N/A, 2016.

[2] J. H. Bullee, A. L. M. Morales, M. Junger, and P. H. Hartel, "Telephone-based social engineering attacks: An experiment testing the success and time decay of an intervention," in Singapore Cyber Security R&D Conference (SG-CRC), Singapore, Singapore, vol. 1 of Cryptology and Information Security Series, (Amsterdam), p. 1–6, IOS Press, IOS Press, 2016.

[3] E. E. H. Lastdrager, P. H. Hartel, and M. Junger, "Apate: Anti-phishing analysing and triaging environment (poster)," in 36th IEEE Symposium on Security and Privacy, San Jose, CA, USA, (USA), IEEE Computer Society, IEEE Computer Society, May 2015.

[4] M. G. Ivanova, C. W. Probst, R. R. Hansen, and F. Kammueller, "Attack tree generation by policy invalidation," in 9th IFIP WG 11.2 International Conference on Information Security Theory and Practice, WISTP 2015, Heraklion, Crete, Greece (R. N. Akram and S. Jajodia, eds.), vol. 9311 of Lecture Notes in Computer Science, (Berlin), p. 249–259, Springer Verlag, Springer Verlag, August 2015.

[5] R. Jhawar, B. Kordy, S. Mauw, S. Radomirovi´c, and R. Trujillo-Rasua, "Attack trees with sequential conjunction," in International Conference on ICT Systems Security and Privacy Protection (IFIPSEC), Hamburg, Germany, IFIP, IFIP, May 2015.

[6] W. Pieters and M. Davarynejad, "Calculating adversarial risk from attack trees: Control strength and probabilistic attackers," in 9th International Workshop on Data Privacy Management, Autonomous Spontaneous Security, and Security Assurance (DPM), Wroclaw, Poland, vol. 8872 of Lecture Notes in Computer Science, (Berlin), p. 201–215, Springer, Springer, March 2015.

[7] W. Pieters, J. Padget, F. Dechesne, V. Dignum, and H. Aldewereld, "Effectiveness of qualitative and quantitative security obligations," Journal of Information Security and Applications, vol. 22, p. 3–16, June 2015.

[8] P. A. Hall, C. P. Heath, L. Coles-Kemp, and A. Tanner, "Examining the contribution of critical visualisation to information security," in New Security Paradigm Workshop (NSPW), Twente, The Netherlands, (New York), p. 1–14, ACM, ACM, September 2015.

[9] W. van der Wagen and W. Pieters, "From cybercrime to cyborg crime: Botnets as hybrid criminal actor-networks," British journal of Criminology, vol. 55, p. 1–18, March 2015.

[10] C. Herley and W. Pieters, ""if you were attacked, you'd be sorry": Counterfactuals as security arguments," in New Security Paradigm Workshop (NSPW), Twente, Netherlands, (New York), p. 1–12, ACM, ACM, 2015.

[11] D. Ionita, R. J. Wieringa, J. H. Bullee, and A. Vasenev, "Investigating the usability and utility of tangible modelling of socio-technical architectures," Enschede, May 2015.

[12] Z. Benenson, G. Lenzini, D. Oliveira, S. Parkin, and S. Uebelacker, "Maybe poor johnny really cannot encrypt - the case for a complexity theory for usable security," in New Security Paradigm Workshop (NSPW), Twente, Netherlands, (New York), p. 1–15, ACM, ACM, 2015.

[13] Z. Aslanyan, M. G. Ivanova, F. Nielson, and C. W. Probst, "Modeling and analysing sociotechnical systems," in 1st International Workshop on Socio-Technical Perspective in IS development (STPIS), Stockholm, Sweden, vol. 1374 of CEUR Workshop Proceedings, p. 121–124, CEUR, CEUR, June 2015.

[14] F. Kammu¨ller and C. W. Probst, "Modeling and verification of insider threats using logical analysis," IEEE Systems Journal, vol. 99, p. 1–12, August 2015.

[15] N. David, A. David, R. R. Hansen, K. G. Larsen, A. Legay, M. C. Olesen, and C. W. Probst, "Modelling social-technical attacks with timed automata," in Proceedings of the 7th ACM CCS International Workshop on Managing Insider Security Threats (MIST), Denver, Colorado, US, (New York), p. 21–28, ACM, ACM, October 2015.

[16] J. H. Bullee, A. L. M. Morales, W. Pieters, M. Junger, and P. H. Hartel, "The persuasion and security awareness experiment: reducing the success

of social engineering attacks," Journal of Experimental Criminology, vol. 11, p. 97–115, March 2015.

[17] T. Chen, F. Kammueller, I. Nemli, and C. W. Probst, "A probabilistic analysis framework for malicious insider threats," in Third International Conference on Human Aspects of Information Security, Privacy, and Trust (HAS), Los Angeles, US, vol. 9190 of Lecture Notes in Computer Science, (Berlin), p. 178–189, Springer Verlag, Springer Verlag, July 2015.

[18] R. Kumar, E. J. J. Ruijters, and M. I. A. Stoelinga, "Quantitative attack tree analysis via priced timed automata," in Proceedings of the 13th International Conference on Formal Modeling and Analysis of Timed Systems, FORMATS 2015, Madrid, Spain, vol. 9268 of Lecture Notes in Computer Science, p. 156–171, Springer Verlag, Springer Verlag, August 2015.

[19] J. H. Bullee, A. L. M. Morales, W. Pieters, M. Junger, and P. H. Hartel, "Regression nodes: Extending attack trees with data from social sciences," in Workshop on Socio-Technical Aspects in Security and Trust (STAST), Verona, Italy, (USA), IEEE Computer Society, IEEE Computer Society, July 2015.

[20] G. Lenzini, S. Mauw, and S. Ouchani, "Security analysis of socio-technical physical systems," Computers & Electrical Engineering, vol. online, 2015.

[21] F. Arnold, D. Guck, R. Kumar, and M. I. A. Stoelinga, "Sequential and parallel attack tree modelling," in Computer Safety, Reliability, and Security - Proceedings of the SAFECOM 2015 Workshops, ASSURE, DECSoS. ISSE, ReSA4CI, and SASSUR, Delft, The Netherlands (F. Koornneef and C. V. Gulijk, eds.), vol. 9338 of Lecture Notes in Computer Science, (Zurich), p. 291–299, Springer Verlag, Springer Verlag, September 2015.

[22] D. Gollmann, C. Herley, V. Koenig, W. Pieters, and M. A. Sasse, "Socio-technical security metrics (dagstuhl seminar 14491)," Dagstuhl Reports, vol. 4, p. 1–28, March 2015.

[23] D. Ionita, R. J. Wieringa, J. H. Bullee, and A. Vasenev, "Tangible modelling to elicit domain knowledge: An experiment and focus group," in 34th International Conference, ER 2015, Stockholm, Sweden (P. Johannesson, M. L. Lee, S. W. Liddle, A. L. Opdahl, and Lo´pez, eds.), vol. 9381 of Lecture Notes in Computer Science, (Berlin), p. 558–565, Springer Verlag, Springer Verlag, October 2015.

[24] R. Kumar, D. Guck, and M. I. A. Stoelinga, "Time dependent analysis with dynamic counter measure trees," in Proceedings of the 13th Workshop on Quantitative Aspects of Programming Languages and Systems (QAPL 2015), London, England, (France), Inria, Inria, April 2015.

[25] M. Nidd, M. G. Ivanova, C. W. Probst, and A. Tanner, Tool-based Risk Assessment of Cloud Infrastructures as Socio-Technical Systems, p. 495–517. Elsevier Science Direct, Amsterdam: Elsevier, Syngress, June 2015.

[26] M. G. Ivanova, C. W. Probst, R. R. Hansen, and F. Kammueller, "Transforming graphical system models to graphical attack models," in The Second International Workshop on Graphical Models for Security (GraMSec 2015), Verona, Italy (S. Mauw and B. Kordy, eds.), Lecture Notes in Computer Science, (London), p. 1–15, Springer Verlag, Springer Verlag, July 2015.

[27] D. Ionita, R. J. Wieringa, L. Wolos, J. Gordijn, and W. Pieters, "Using value models for business risk analysis in e-service networks," in 8th IFIP WG 8.1. Working Conference, PoEM 2015, Valencia, Spain (R. Ralyt´e, S. Espan˜a, and O. Pastor, eds.), vol. 235 of Lecture Notes in Business Information Processing, (Berlin), p. 239–253, Springer Verlag, Springer Verlag, November 2015.

[28] E. E. H. Lastdrager, "Achieving a consensual definition of phishing based on a systematic review of the literature," Crime Science, vol. 3, pp. 9:1–9:16, June 2014.

[29] D. Ionita, J. H. Bullee, and R. J. Wieringa, "Argumentation-based security requirements elicitation: The next round," in Proceedings of the 2014 IEEE 1st International Workshop on Evolving Security and Privacy Requirements Engineering (ESPRE), Karlskrona, Sweden, p. 7–12, IEEE Computer Society, IEEE Computer Society, August 2014.

[30] D. P. Sari, "Attacker profiling in quantitative security assessment," masters, Tallinn University of Technology, January 2014.

[31] A. Lenin, J. Willemson, and D. Sari, "Attacker profiling in quantitative security assessment based on attack trees," in 19th Nordic Conference on Secure IT (NordSec), Troms?, Norway, vol. 8788 of Lecture Notes in Computer Science, (Berlin), Springer, Springer, October 2014.

[32] S. Bleikertz, C. Vogel, and T. Gross, "Cloud radar: Near real-time detection of security failures in dynamic virtualized infrastructures," in Annual Computer Security Applications Conference (ACSAC), New Orleans, Louisiana, (New York), ACM, ACM, December 2014.

[33] F. Kammu¨ller and C. W. Probst, "Combining generated data models with formal invalidation for insider threat analysis," in IEEE Security and Privacy Workshops (SPW), San Jose, California, p. 229–235, IEEE Computer Society, IEEE Computer Society, May 2014.

[34] D. Ionita, "Context-sensitive information security risk identification and evaluation techniques," in 22nd IEEE International Requirements Engineering Conference (RE14), Karlskrona, Sweden, (USA), p. 485–488, IEEE Computer Society, IEEE Computer Society, August 2014.

[35] W. Pieters, C. W. Probst, S. Lukszo, and A. L. M. Morales, Cost-effectiveness of Security Measures: A model-based Framework, p. 139–156. Hershey, PA, USA: IGI Global, 2014.

[36] W. Pieters, D. Hadziosmanovi´c, and F. Dechesne, "Cyber security as social experiment," in NSPW '14 Proceedings of the 2014 workshop on New Security Paradigms, NSPW 2014, Victoria, BC, Canada, (New York), p. 15–24, ACM, ACM, September 2014.

[37] B. K. Kordy, L. Pi`etre-Cambac´ed`es, and P. Schweitzer, "Dag-based attack and defense modeling: Don't miss the forest for the attack trees," Computer Science Review, vol. 13-14, p. 1–38, 2014.

[38] M. Sytema, A. F. E. Belinfante, M. I. A. Stoelinga, and L. Marinelli, "Experiences with formal engineering: model-based specification, implementation and testing of a software bus at neopost," Science of computer programming, vol. 80, p. 188–209, February 2014.

[39] F. Dechesne, D. Hadziosmanovi´c, and W. Pieters, "Experimenting with incentives: Security in pilots for future grids," IEEE Security & Privacy, vol. 12, p. 59–66, November 2014.

[40] F. Kammu¨ller and C. W. Probst, "Invalidating policies using structural information," Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA), vol. 5, p. 59–79, June 2014.

[41] A. Lenin and A. Buldas, "Limiting adversarial budget in quantitative security assessment," in 5th International Conference on Decision and Game Theory for Security (GameSec), Los Angeles, CA, USA, vol. 8840 of Lecture Notes in Computer Science, (Berlin), p. 155–174, Springer, Springer, November 2014.

[42] C. P. Heath, L. Coles-Kemp, and P. A. Hall, "Logical lego? co-constructed perspectives on service design," in Proceedings of NordDesign 2014, Melbourne, Australia, p. 416–425, Aalto Design Factory, Aalto Design Factory, August 2014.

[43] M. Huisman and M. I. A. Stoelinga, "Meer vrouwen in de ict, waarom eigenlijk?," Bits en chips, vol. 9, p. 20–21, November 2014.

[44] C. W. Probst and R. R. Hansen, "Model-based abstraction of data provenance," in 6th USENIX Workshop on the Theory and Practice of Provenance, Cologne, Germany, p. Article 3, Usenix Association, Usenix Association, June 2014.

[45] J. Boender, M. G. Ivanova, F. Kammu¨ller, and G. Primierio, "Modeling human behaviour with higher order logic: Insider threats," in 4th Workshop on Socio-Technical Aspects in Security and Trust (STAST), Vienna, Austria, p. 31–39, IEEE, IEEE, July 2014.

[46] D. Guck, M. Timmer, H. Hatefi, E. J. J. Ruijters, and M. I. A. Stoelinga, "Modelling and analysis of markov reward automata," in Proceedings of the 12th International Symposium on Automated Technology for Verification and Analysis, ATVA 2014, Sydney, NSW, Australia, vol. 8837 of Lecture Notes in Computer Science, (Berlin), p. 168–184, Springer Verlag, Springer Verlag, November 2014.

[47] D. Ionita, S. K. Koenen, and R. J. Wieringa, "Modelling telecom fraud with e3value," Enschede, October 2014.

[48] B. Kordy, M. Pouly, and P. Schweizer, "A probabilistic framework for security scenarios with dependent actions," in 11th International Conference on Integrated Formal Methods, IFM 2014, Bertinoro, Italy (E. Albert and E. Sekereinsk, eds.), vol. 8739 of Lecture Notes in Computer Science, p. 256–271, Springer, Springer, September 2014.

[49] Proceedings First International Workshop on Graphical Models for Security, GraMSec 2014, Grenoble, France, 12th April, 2014, vol. 148, EPTCS.ORG, April 2014.

[50] F. Arnold, W. Pieters, and M. I. A. Stoelinga, "Quantitative penetration testing with item response theory," Journal of Information Assurance and Security, vol. 9, no. 3, p. 118–127, 2014.

[51] W. Pieters, Z. Lukszo, D. Hadziosmanovi´c, and J. van den Berg, "Reconciling malicious and accidental risk in cyber security," Journal of Internet Services and Information Security, vol. 4, p. 4–26, May 2014.

[52] G. Schaff, C. Harpes, M. Aubigny, M. Junger, and R. Martin, "Risk-det: Ict security awareness aspect combining education and cognitive sciences," in Ninth International MultiConference on Computing in the Global Information Technology, ICCGI 2014, Seville, Spain, IARIA, IARIA, June 2014.

[53] S. Uebelacker and S. Quiel, "The social engineering personality framework," in 4th Workshop on Socio-Technical Aspects in Security and Trust (STAST), Vienna, Austria, p. 24–30, IEEE, IEEE, July 2014.

[54] A. K. I. Remke and M. I. A. Stoelinga, Stochastic Model Checking: Rigorous Dependability Analysis Using Model Checking Techniques for Stochastic Systems, vol. 8453 of Lecture Notes in Computer Science. London: Springer Verlag, October 2014.

[55] F. Arnold, H. Hermanns, R. Pulungan, and M. I. A. Stoelinga, "Time-dependent analysis of attacks," in Proceedings of the Third International Conference on Principles and Security of Trust, POST 2014, Grenoble, France, vol. 8414 of Lecture Notes in Computer Science, (Berlin), p. 285–305, Springer Verlag, Springer Verlag, April 2014.

[56] C. Poskitt, M. Dodds, R. F. Paige, and A. Rensink, "Towards rigorously faking bidirectional model transformations," in Proceedings of the Workshop on Analysis of Model Transformations, AMT 2014, Valencia, Spain (J. Dingel, J. D. Lara, L. Lu´cio, and H. Vangheluwe, eds.), vol. 1277 of CEUR-WS, (Aachen), p. 70–75, RWTH Aachen, Germany, RWTH Aachen, Germany, September 2014.

[57] W. Pieters, D. Hadziosmanovi´c, A. Lenin, A. L. M. Morales, and J. Willemson, "Trespass: Plug-and-play attacker profiles for security risk analysis (poster)," in 35th IEEE Symposium on Security and Privacy, San Jose, California, (USA), IEEE Computer Society, IEEE Computer Society, May 2014.

[58] F. Arnold, D. Gebler, D. Guck, and H. Hatefi, "A tutorial on interactive markov chains," in Stochastic Model Checking. Rigorous Dependability Analysis Using Model Checking Techniques for Stochastic Systems, Vahrn, Italy (A. K. I. Remke and M. I. A. Stoelinga, eds.), vol. 8453 of Lecture Notes in Computer Science, (Berlin), p. 26–66, Springer Verlag, Springer Verlag, 2014.

[59] B. Kordy, P. Kordy, S. Mauw, and P. Schweitzer, "Adtool: Security analysis with attackdefense trees," in 10th International Conference on Quantitative Evaluation of Systems (QEST), Buenos Aires, Argentina, vol. 8054 of Lecture Notes in Computer Science, p. 173–176, Springer, Springer, August 2013.

[60] G. Schaff, C. Harpes, R. Martin, and M. Junger, "An application to estimate the cyber-risk detection skill of mobile device users (idea)," in Sixth International Conference on Advances in Human oriented and Personalized Mechanisms, Technologies, and Services (CENTRIC), Venice, Italy, p. Article 7, IARIA, IARIA, October 2013.

[61] E. E. H. Lastdrager, A. L. M. Morales, P. H. Hartel, and M. Junger, "Applying the lost-letter technique to assess it risk behaviour," in Proceedings of the 3rd Workshop on Socio-Technical Aspects in Security and Trust, New Orleans, USA, (USA), p. 2–9, IEEE Computer Society, IEEE Computer Society, June 2013.

[62] M. I. A. Stoelinga and W. Pieters, "Attack navigator vindt en verhelpt zwakke plekken," Bits en chips, vol. 4, April 2013.

[63] M. Timmer, J. C. van de Pol, and M. I. A. Stoelinga, "Confluence reduction for markov automata," in Proceedings of the 11th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS), Buenos Aires, Argentina (V. A. Braberman and L. Fribourg, eds.), vol. 8053 of Lecture Notes in Computer Science, (Berlin), p. 243–257, Springer Verlag, Springer Verlag, August 2013.

[64] M. Timmer, J. C. van de Pol, and M. I. A. Stoelinga, "Confluence reduction for markov automata (extended version)," Enschede, June 2013.

[65] D. Ionita, P. H. Hartel, W. Pieters, and R. J. Wieringa, "Current established risk assessment methodologies and tools," Enschede, September 2013.

[66] S. Bleikertz, T. Mastelic, S. Pape, W. Pieters, and T. Dimkov, "Defining the cloud battlefield supporting security assessments by cloud customers," in International Conference on Cloud Engineering (IC2E 2013), Redwood City, CA, (USA), p. 78–87, IEEE Computer Society, IEEE Computer Society, March 2013.

[67] W. Pieters, "Defining "the weakest link" comparative security in complex systems of systems," in 2013 IEEE 5th International Conference on Cloud Computing Technology and Science, CloudCom, Bristol, United Kingdom, (USA), p. 39–44, IEEE Computer Society, IEEE Computer Society, December 2013.

[68] F. Arnold, A. F. E. Belinfante, F. I. V. der Berg, D. Guck, and M. I. A. Stoelinga, "Dftcalc: a tool for effcient fault tree analysis," in Proceedings of the 32nd International Conference on Computer Safety, Reliability, and Security (SAFECOMP), Toulouse, France, vol. 8153 of Lecture Notes in Computer Science, (Berlin), p. 293–301, Springer Verlag, Springer Verlag, September 2013.

[69] F. Arnold, A. F. E. Belinfante, F. I. V. der Berg, D. Guck, and M. I. A. Stoelinga, "Dftcalc: a tool for effcient fault tree analysis (extended version)," Enschede, June 2013.

[70] M. G. Ivanova, C. W. Probst, R. R. Hansen, and F. Kammu¨ller, "Externalizing behaviour for analysing system models," Journal of Internet Services and Information Security, vol. 3, p. 52–62, November 2013.

[71] F. Kammu¨ller and C. W. Probst, "Invalidating policies using structural information," in IEEE Security and Privacy Workshops (SPW 2013), San Francisco, CA, p. 76–81, IEEE Computer Society, IEEE Computer Society, May 2013.

[72] A. Buldas and A. Lenin, "New effcient utility upper bounds for the fully adaptive model of attack trees," in 4th International Conference on

Decision and Game Theory for Security (GameSec), Fort Worth, TX, vol. 8252 of Lecture Notes in Computer Science, (Berlin), p. 192–205, Springer, Springer, November 2013.

[73] W. Pieters, J. Padget, F. Dechesne, V. Dignum, and H. Aldewereld, "Obligations to enforce prohibitions: on the adequacy of security policies," in SIN '13 - Proceedings of the 6th International Conference on Security of Information and Networks, Aksaray, Turkey, Proceeding, (New York), p. 54–61, ACM, ACM, November 2013.

[74] F. Arnold, W. Pieters, and M. I. A. Stoelinga, "Quantitative penetration testing with item response theory," in 9th International Conference on Information Assurance and Security, IAS 2013, Gammarth, Tunisia, Information Assurance and Security (IAS), 2013 9th International Conference on, (USA), p. 49–54, IEEE, IEEE, December 2013.

[75] F. Arnold, W. Pieters, and M. I. A. Stoelinga, "Quantitative penetration testing with item response theory (extended version)," Enschede, October 2013.

[76] C. W. Probst and R. R. Hansen, "Reachability-based impact as a measure for insiderness," Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, vol. 4, p. 38–48, December 2013.

[77] H. Prakken, D. Ionita, and R. J. Wieringa, "Risk assessment as an argumentation game," in 14th International Workshop on Computational Logic in Multi-Agent Systems, CLIMA XIV, Corunna, Spain (J. Leite, T. C. Son, P. Torrini, L. V. D. Torre, and S. Woltran, eds.), vol. 8143 of Lecture Notes in Computer Science, (London), p. 357–373, Springer Verlag, Springer Verlag, September 2013.

[78] S. Uebelacker, "Security-aware organisational cultures as a starting point for mitigating socio-technical risks," in Informatik 2013, University of Koblenz-Landau, Koblenz, Germany (M. Horbach, ed.), vol. P-220 of Lecture Notes in Informatics (LNI), (Bonn), p. 2046–2057, Gesellschaft fuer Informatik e.V, Gesellschaft fuer Informatik e.V, September 2013.

[79] W. Pieters, "On thinging things and serving services: technological mediation and inseparable goods," Ethics and information technology, vol. 15, p. 195–208, September 2013.

[80] Y. Feng and L. Zhang, "A tighter bound for the self-stabilization time in herman's algorithm," Information processing letters, vol. 113, p. 486–488, July 2013.

[81] A. L. M. Morales, "The trespass project," in ICTOpen2013, Eindhoven, (Netherlands), p. 1–1, ICTopen, ICTopen, October 2013.

10. Glossary

Access Point (AP) – In computer networking, it refers to a network access point, e.g. wireless access point that allows Wi-Fi compliant devices to connect to the network.

Accidental Insider Attack – Refers to an unintentional insider attack, which is caused by an employee from within the organisation, e.g. by taking a selfie, revealing company confidential documents in the background of the photo.

AP – cf. Access Point.

Attack Vector – An attack vector is a path or means by which the attacker can gain access to a computer or network server in order to deliver a malicious outcome. Attack vectors enable the attacker to exploit system vulnerabilities, including the human element. Most notoriously known is perhaps the phishing attack.

Awareness Training – Education and training of employees with the purpose of raising their security consciousness, i.e. make people more aware of potential cyber threats.

Big Five Framework – A psychological framework based on common language descriptors of personality, which is used for conducting a personality profiling of people, by dividing them into one of five categories: Openness; Conscientiousness; Extraversion; Agreeableness; Neuroticism.

BlackEnergy3 – Refers to a specific type of malware (cf. malware).

Blackhat – A blackhat is a hacker with criminal intent, who attempts to gain unauthorised access to systems. Can also refer to the annual BlackHat cyber security conference.

Bring Your Own Device (BYOD) – Refers to the policy of permitting employees to bring personally owned devices (laptops, tablets, and smart phones) to their workplace, and to use those devices to access privileged company information and applications.

Brute Force Attack – Refers to the mathematical guarantee for breaking a password by exhausting every possible combination, which is usually a time-consuming process.

BYOD – cf. Bring Your Own Device.

ccTLD – cf. Country Code Top Level Domain.

Centrality – cf. Appendix B.

CEO Fraud – When an attacker attempts to get the accounting department to transfer funds to the attacker's bank account by convincing the target that the attacker is the CEO of the company.

CNI – cf. Critical National Infrastructure.

Conventional Social Engineering – Attacks conducted via the interaction between two or more individuals with the purpose of elicitation of information from the target.

Country Code Top Level Domain (ccTLD) – Country-specific Internet top level domain name generally reserved for a sovereign states, e.g. Danish websites end in *.dk*, whereas British end in *co.uk*.

Cracker – A cracker is someone who breaks into someone else's computer system, often on a network; bypasses passwords or licenses in computer programs; or in other ways intentionally breaches computer security.

Crawling – Often referred to as *web crawling*, is the act of systematically collecting information from the Internet. This is a common act for search engines, which index the crawled websites.

Credential Harvesting – The act of deceiving individuals into divulging their username and password in a rouge web form, either online or implemented into an email.

Critical Infrastructure (CI) – Public institutions and private companies that are vital for the functioning of a society, which can be on a national, regional or global level. Though it differs from country to country, CI more often than not includes public/private institutions responsible for: Electricity generation, telecommunication, water supply, public health, transportation, financial services and security services.

Critical National Infrastructure (CNI) – cf. Critical Infrastructure.

cSE – cf. Conventional Social Engineering.

Darknet – Refers to an overlay network that can only be accessed with specific software, configurations, or authorisation, often using non-standard communications protocols and ports. A typical darknet type is *Tor*.

DDoS – cf. Distributed Denial of Service.

Deep Web – Refers to parts of the Internet, which are not indexed on popular search engines and therefore more difficult to find. This can include websites, documents, databases, but also web mail and online banking is part of the deep web.

Dictionary Attack – Refers to a technique for defeating a cipher or authentication mechanism by trying to determine its decryption key or passphrase by trying hundreds or sometimes millions of likely possibilities, such as words in a dictionary.

Distributed Denial of Service (DDoS) – Refers to an explicit attempt by attackers to prevent legitimate users use of a service. In a DDoS attack, the incoming traffic flooding the victim originates from many different sources – potentially hundreds of thousands. This effectively makes it impossible to stop the attack simply by blocking a single IP address; plus, it is very difficult to distinguish legitimate user traffic from attack traffic when spread across so many points of origin.

Dumpster Diving – In information technology, dumpster diving refers to a technique used to retrieve information that could be used to carry out an attack on a computer network. It is not limited to searching through the trash for obvious treasures like access codes or passwords written down on sticky notes. Seemingly innocent information like a phone list, calendar, or organisational charts can be used to assist an attacker using social engineering techniques to gain access to the computer network.

Electromagnetic Spectrum (EMS) – The electro-magnetic spectrum is the collective term used for all known frequencies and their linked wavelengths of the known photons. The electromagnetic spectrum extends from below the low frequencies used for modern radio communication to gamma radiation at the short-wavelength (high-frequency) end.

Electronic Warfare (EW) – Refers to a military and technological discipline involving actions in battlespace that utilise the electromagnetic spectrum or directed energy.

EMS – cf. Electromagnetic Spectrum.

EW – cf. Electronic Warfare.

Field Trial Testing – The part of Project SAVE for which the social vulnerability assessment was conducted.

Fingerprinting Organisations with Collected Archives (FOCA) – FOCA is a piece of software that can scan popular search engines (Google, Bing and Exalead) for files relating to the web domain of interest. It then *crawls* the files from the Internet and then performs local analyses of the metadata.

FOCA – cf. Fingerprinting Organisations with Collected Archives.

Google Dorking – Refers to a technique that uses Google Search and other Google applications to find security holes in the configuration and computer code that websites use by using advanced search operators to locate specific strings of information within the search results.

Google Hacking Database – A large collection of Google dorks that is continuously updated by the Google Hacking community.

Hacker – Refers to any highly skilled computer expert capable of breaking into computer systems and networks using bugs and exploits. Depending on the field of computing it has slightly different meanings, and in some contexts has controversial moral and ethical connotations.

Hashing – Refers to the transformation of a string of characters into a usually shorter fixed-length value or key that represents the original string. Hashing is used to index and retrieve items in a database because it is faster to find the item using the shorter hashed key than to find it using the original value. It is also used in many encryption algorithms.

HID – cf. Human Interface Device.

Human Intelligence (HUMINT) – The collection of intelligence via interpersonal contact. The collected information therefore derives form human sources.

Human Interface Device (HID) – Refers to a type of computer device that interacts directly with, and most often takes input from, humans and may deliver output to humans, e.g. a computer mouse or keyboard.

HUMINT – cf. Human Intelligence.

IBM i2 Analyst's Notebook – Refers to a software package developed by IBM. i2 Analyst's Notebook is a data analysis environment that allows analysts to quickly collate, analyse and visualise data from disparate sources visually.

ICS – cf. Industrial Control System.

IDS – cf. Intrusion Detection System.

Incident Response Plan (IRP) – Refers to a set of written instructions for detecting, responding and limiting the effects of a cyber incident. It acts as a guideline for how to respond to and record incidents.

Industrial Control Systems (ICS) – ICS is a general term that encompasses several types of control systems and associated instrumentation used in industrial production, including supervisory control and data acquisition (SCADA) systems. ICS are typically used in industries suck as electrical, water, oil, gas and data.

Information Gatekeepers – Refers to employees who due to their function within an organisation hold a significant role with privileged access to information and/or system access, e.g. the IT department or human resources.

Insider Attack/Threat – Refers to a malicious threat to an organisation that comes from people within the organisation, such as employees, former employees, contractors or business associates, who have inside information concerning the organisation's security practices, data and computer systems. The threat may involve fraud, the theft

of confidential or commercially valuable information, the theft of intellectual property, or the sabotage of computer systems.

Intrusion Detection System (IDS) – Refers to a device or software application that monitors a network or systems for malicious activity or policy violations.

IRP – cf. Incident Response Plan.

KillDisk – Refers to a wiper virus that overwrites data in essential system files, causing the computer to crash without the possibility of doing a reboot, since the virus overwrites the master boot record.

Macro – A macro allows short sequences of keystrokes and mouse actions to be transformed into other, usually more time-consuming, sequences of keystrokes and mouse actions. In this way, frequently used or repetitive sequences of keystrokes and mouse movements can be automated.

Maltego – Refers to a software application, which provides an overview of the systemic Internet protocol (IP) infrastructure based on a web domain, which can provide identification of people, exchange of information, DNS information, metadata, email addresses, social media accounts and much more.

Malware – Short for *malicious software*. Malware is any software used to disrupt computer or mobile operations, gather sensitive information, gain access to private computer systems, or display unwanted advertising. Before the term malware was coined, malicious software was referred to as computer viruses.

Man-In-The-Middle (MITM) – Refers to an attack where the attacker secretly relays and possibly alters the communication between two parties who believe they are directly communicating with each other. A man-in-the-middle attack can be used against many cryptographic protocols and is used to either manipulate or eavesdrop on the information exchange.

MD5 – Refers to a widely used hash function producing a 128-bit hash value.

Metadata – Often characterised as data about data, or information about information. Metadata is defined as the data providing information about one or more aspects of the data; it is used to summarise basic information about data which can make tracking and working with specific data easier, e.g. file size, author of the data, timestamp, software used for creating the data, etc.

Method Acting – In a social engineering context, method acting refers the act of becoming the person you are pretending to be, including clothing, body language, ID badge, jargon, etc.

MITM – cf. Man-In-The-Middle.

Network Nodes – In the context of a social network analysis, a network node refers to a connection point in a network, which can consist of, e.g. an actor, a place or an object like a phone, email or other.

Open Source Intelligence (OSINT) – Refers to intelligence collected from publicly available sources. In the intelligence community, the term "open" refers to overt, publicly available sources (as opposed to covert or clandestine sources). Most OSINT today is collected from the Internet, but has traditionally been collected from public archives or libraries.

OSINT – cf. Open Source Intelligence.

Pastebin – Refers to an online web service where you can store text for a certain period of time. The website is mainly used by programmers to store pieces of sources code or configuration information, but anyone is more than welcome to paste any type of text. The idea behind the site is to make it more convenient for people to share large amounts of text online. Leaked information is often shared on Pastebin as well.

Paterva – Refers to the company that has developed the software application *Maltego.*

Payload – In computer security, the payload is the part of malware, such as worms or viruses, which performs the malicious action, e.g. deleting data, sending spam or encrypting data.

PDF Attack – Refers to a specific type of attack used in Project SAVE, consisting of a PDF-file with an integrated link, which when clicked confirms that the user has opened the file.

Personality Profiling – Refers to a psychometric testing that measures an individual's personality based on predefined parameters. It is therefore not a measure of intelligence or ability, but rather a measure of behaviour.

Phishing – Phishing is an example of social engineering techniques used to deceive users, and exploits weaknesses in current web security. Phishing is typically carried out by email spoofing or instant messaging, and it often directs users to enter personal information at a fake website whose look and feel are almost identical to the legitimate one.

Pretexting – Refers to an act of creating and using an invented scenario (the pretext) to engage a targeted victim in a manner that increases the chance the victim will divulge information or perform actions that would be unlikely in ordinary circumstances.

Ransomware – Refers to a specific type of malware that installs covertly on a victim's computer, executes a cryptovirology attack that adversely affects it, and demands a ransom payment to decrypt the data.

Reconnaissance – In the context of social engineering, reconnaissance refers to the information gathering phase of an attack, where the attacker studies the target(s).

Reverse Social Engineering (rSE) – Refers to a specific form of social engineering attack, where the attacker sets the stage, so that instead of a social engineer being the one to approach the target, the target is the one who approaches the social engineer. This can be accomplished by creating an incident that requires the help of the social engineer.

rSE – cf. Reverse Social Engineering.

SCADA – cf. Supervision, Control and Acquisition Data.

Scythe – Refers to a script that acts as an account enumerator, which makes it easy to identify social media accounts across a vast array of social media networks, based on the same email account.

SDR – cf. Software Defined Radio.

SE – cf. Social Engineering.

SE 2.0 – cf. Social Engineering 2.0.

Secure Socket Layer (SSL) – A cryptographic protocol that provides communication security over a computer network.

Sentiment Analysis – Refers to the use of natural language processing, text analysis and computational linguistics to identify and extract subjective information in source materials. Sentiment analysis is widely applied to reviews and social media for a variety of applications, ranging from marketing to customer service to social engineering attacks.

SET – cf. Social Engineering Toolkit.

Shodan – Refers to a search engine that lets the user find specific types of computers or IoT devices (webcams, routers, servers, etc.) connected to the internet using a variety of filters.

Shoulder Surfing – In computer security, shoulder surfing is a type of social engineering technique used to obtain information such as personal identification number, password and other confidential data by looking over the victim's shoulder.

SIGINT – cf. Signals Intelligence.

Signals Intelligence (SIGINT) – Refers to the collection of intelligence by interception of signals, whether communications between people or from electronic signals not directly used in communication.

Simulated Attacks – In the context of Project SAVE, simulated attacks refer to the attacks that closely resemble an actual cyber-attack, though without the use of malware, since it is intended to focus on the human element of cyber security.

Smishing – Smishing, short for SMS Phishing, is a social engineering technique that attempts to trick a recipient into divulging personal information, such as passwords, and/or perform actions, by masquerading as a trustworthy entity in a SMS.

SMPP – cf. Social Media Personality Profiling.

SNA – cf. Social Network Analysis.

Social Engineering (SE) – Refers to psychological manipulation of people into performing actions they would not otherwise perform, such as divulging confidential information.

Social Engineering 2.0 (SE 2.0) – Refers to the evolution of conventional social engineering attacks, which includes new and novel approaches. SE 2.0 combines advanced information gathering techniques with the modern attack vectors, including social media, email, USB and SMS.

Social Engineering Toolkit (SET) – Refers to an open source script that can be used for conducting a vast array of social engineering 2.0 attacks.

Social Media Intelligence (SOCMINT) – Refers to the collective tools and solutions that allow organisations to monitor social channels and conversations, respond to social signals and synthesise social data points into meaningful trends and analysis. Social media intelligence allows one to collect intelligence gathering from social media sites, using both intrusive and non-intrusive means, from open and closed social networks.

Social Media Networks – Refers to computer-mediated virtual communities and networks that allow the creating and sharing of information, ideas, career interests or other forms of expression, e.g. Facebook, Twitter, LinkedIn.

Social Media Personality Profiling (SMPP) – Refers to a script developed within Project SAVE that crawls information about users from open Facebook accounts. The script the automatically conducts a sentiment analysis of the users' content, and then conduct as personality profiling, in an effort to estimate which behavioural patterns can be expected from the user when engaged.

Social Network Analysis (SNA) – Refers to the process of investigating social structures through the use of network and graph theories. SNA characterises networked structures in terms of nodes (individual actors, people, or things within the network) and the ties, edges, or links (relationships or interactions) that connect them. Examples of social structures commonly visualised through SNA include social media networks.

Social Vulnerability Assessment (SVA) – Refers to an approach utilised in Project SAVE, which simulates attack patterns in order to measure the real vulnerability of the human security barrier

in an organisation. An SVA approach is a new type of assessment, which, in the context of cyber security, proactively uses social engineering techniques to attack the enterprises, in an effort to evaluate their current social vulnerability level.

SOCMINT – cf. Social Media Intelligence.

Software Defined Radio (SDR) – Refers to a radio communication system in which some or all of the physical layer functions are software defined, instead of being hardware defined.

Spear-Phishing – A technique that fraudulently obtains private information by sending highly customised emails. The main difference between phishing and spear-phishing is that phishing campaigns focus on sending out high volumes of generalised emails with the expectation that only a few people will respond, while spear-phishing emails require the attacker to perform additional research on their targets in order to "trick" end users into performing requested activities and are only sent to a few.

Spoofing – Refers to the creation of emails or SMS messages with a forged sender address, which is used to trick the recipient into believing the request is legitimate.

SSL – cf. Secure Socket Layer.

Stuxnet – Refers to a malicious computer worm that was used to sabotage Iran's nuclear program.

Supervision, Control and Acquisition Data (SCADA) – Refers to systems that includes both hardware and software components for process control, gathering of data in real time from remote locations in order to control equipment and conditions. SCADA is used in power plants, oil and gas refining, telecommunications, transportation, and water and waste control.

SVA – cf. Social Vulnerability Assessment.

Tailgating – An action performed by an attacker seeking entry to a restricted area secured by unattended, electronic access control, e.g. by RFID card, who then simply walks in behind a person who has legitimate access. Following common courtesy, the legitimate person will *usually* hold the door open for the attackers or the attackers may ask the employee to hold it open for them.

TDoS – cf. Telephony Denial of Service.

Telephony Denial of Service (TDoS) – Telephony Denial of Service is a flood of unwanted, malicious inbound calls, blocking the service from functioning and allowing other calls to come through. The calls are usually into a contact center or other part of an enterprise, which depends heavily on voice service.

TLD – cf. Top-Level Domain.

Top-Level Domain (TLD) – Refers to the last segment of a domain name. The TLD is the letters immediately following the final dot in an Internet address, e.g. .com, .net, .org.

Typosquatting – Refers to a form of Internet cybersquatting, based on the probability that a certain number of Internet users will mistype the name of a Web site (or actually its URL) when surfing the Web.

Uniform Resource Locator (URL) – Refers to a web resource that specifies its location on a computer network. It is commonly used interchangeably with the term *web address*.

URL – cf. Uniform Resource Locator.

USB Attack – In the context of Project SAVE, an USB attack refers to a HID spoofing USB, which looks like a USB device, emulates a keyboard, and executes a malicious script and injects predefined keystrokes when plugged into a computer.

Vishing – Vishing (voice or VoIP phishing) is an electronic deception tactic in which individuals are tricked into revealing critical financial or personal information to unauthorised entities. Vishing works like phishing but does not always occur over the Internet and is carried out using voice technology. A vishing attack can be conducted by voice email, VoIP (voice over IP), or landline or cellular telephone.

Wayback Machine – Refers to a website that enables anyone to see what a particular site looked like at some time in the past - from 1996 to the present. This enormous archive of the Internet's past requires over 100 terabytes of storage and contains 10 billion webpages.

Whaling – Characterised as a type of fraud that specifically targets high-profile end users such as C-level corporate executives, politicians and celebrities.

Zero-Day Exploit – Refers to an exploit that takes advantage of a security vulnerability on the same day that the vulnerability becomes generally known. There are zero days between the time the vulnerability is discovered and the first attack. A hacker may be the first to discover the vulnerability, and since the vulnerability is not known in advance, there is no way to guard against the exploit before it happens.

About the book

Project SAVE - Social Vulnerability & Assessment Framework - is an explorative study on cyber attacks using Social Engineering 2.0 and advanced open source intelligence (OSINT). Social engineering involves the exploitation of the human element of cyber security. The study is an investigation into the phenomenon of Social Engineering 2.0 with the aim of developing a social vulnerability assessment (SVA) framework for conducting cyber reconnaissance to uncover critical information and using deception tactics to create social engineering attacks, which will help identify the human security vulnerabilities within an organisation. In Project SAVE, a total of 185 social engineering attacks were executed against three organisations that are part of critical infrastructure in Denmark.

About the author

Dennis Hansen, M.Sc., has since 2015 worked as an open source intelligence (OSINT) consultant at the Danish Institute of Fire and Security Technology (DBI), where he designs and executes social engineering attacks and develops mitigation solutions. He is involved in both Danish and European projects on offensive social engineering, where his expertise in online deception tactics and OSINT is used for the development of next-generation social engineering attacks. He has a professional background from the defence industry, working in the fields of electronic warfare (EW) and signals intelligence (SIGINT), and has strategic experience in political intelligence from the Danish Parliament on defence and foreign policy.