Audun Einangen

Go Phish!

Educating Users Not to Bite on Phishing

Master's thesis in Information Security, Experience Based
Supervisor: Patrick Bours
June 2021

Master's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Dept. of Information Security and Communication Technology

Hammerfest

NTNU
Norwegian University of
Science and Technology

Audun Einangen

Go Phish!

Educating Users Not to Bite on Phishing

Master's thesis in Information Security, Experience Based
Supervisor: Patrick Bours
June 2021

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Dept. of Information Security and Communication Technology

NTNU
Norwegian University of
Science and Technology

Abstract

In a world where we are constantly being bombarded with information from every angle, some people try to take advantage of that, and sneaks harmless-looking email that can wreak havoc along with the legitimate flow of emails. Some of these tries to lure personal information from an unsuspecting user. These emails are called phishing emails, and have been around for quite some time.

Technical countermeasures for weeding out these phishing emails are often not enough, and it comes down to the users themselves to identify these harmful messages in order not to "bite". Many have therefore tried to find good training methods for the users, so that they can protect themselves.

This project will try to identify any differences in the way different training techniques affect different kinds of users, so that people that arrange training can choose the most effective technique based on the type of user they are educating. By also comparing user groups that have not been focused on before, we may find differences in susceptibility between user groups who interact with emails in different ways. The training methods chosen have been selected based on other studies finding them effective in educating users - Video based instructions, immediate feedback with explanations, web based instructions and gamification. This study tries to compare these methods against each other and against simple textual information with the same basic content.

Sammendrag

I dagens samfunn hvor vi hele tiden blir bombardert med informasjon fra alle kanter, finnes det noen som prøver å utnytte dette, og sniker harmløst utseende eposter som kan lage kaos inn sammen med den normale strømmen av epost. Noen av disse prøver å lure personlig informasjon fra intetanende brukere. Disse epostene kalles phish, eller fiskeepost, og har vært med oss ganske lenge.

Tekniske innretninger for å luke ut disse fiskeepostene er ofte ikke nok, og brukerne selv må kunne identifisere disse skadelige epostene for ikke å "bite på kroken". Mange har derfor prøvd å finne gode treningsmetoder for brukerne, slik at de kan beskytte seg selv.

Dette prosjektet skal prøve å identifisere forskjeller i måten forskjellige treningsteknikker påvirker forskjellige typer brukere, slik at de som setter sammen slik trening kan velge den mest effektive teknikken, basert på hvilken type bruker som skal motta den. Ved også å sammenligne brukergrupper det ikke har vært fokusert på tidligere, kan vi kanskje finne forskjeller i hvor mottakelige grupper som bruker epost forskjellig er for denne type svindel. Treningsmetodene som har blitt valgt, er basert på at andre studier har funnet dem effektive i å oppdra brukere - Videobasert trening (I mangel av klasseromsundervisning i forbindelse med CoVid-19), umiddelbar tilbakemelding med forklaring på hva brukeren gjorde feil, webbasert trening, og trening i form av spill eller konkurranse. Denne studien prøver å sammenligne disse metodene, i tillegg til å sammenligne med enkel tekstbasert informasjon med samme basisinnhold.

Acknowledgements

I would like to thank my wife for her patience while I have been working on my studies. The last two years have been stressful for both of us.

The help I have received from my supervisor, Patrick Bours, has been invaluable. His input has helped me get through this project, and encouraged me when needed. Every time I have had questions, he has been quick to reply, even at the strangest hours.

I would also like to thank my colleagues that have looked at my work and helped with technical testing, making sure everything has been sane. A particular thank you goes to Jørn Lemika, who has helped me with SQL, HTML and CSS. Gorm Tvedt in Alta Municipality has been my contact there, helping me to deliver the emails in their system.

A special mention also to the other colleagues of the IT departments of both Alta and Hammerfest municipalities. Thank you all for handling the phone storms I have caused with the project.

Finally, I also need to thank Åge Olai Johnsen, a former colleague and fellow student that I could have meaningful conversations with about everything relating to my studies.

Contents

Figures

Tables

Code Listings

Acronyms

DKIM DomainKeys Identified Mail. 5

DMARC Domain-based Message Authentication Reporting and Conformance. 5

DNS Domain Name System. xix

GDPR General Data Protection Regulation. 6

MX Mail Exhanger. 30

OSINT Open Source Intelligence. xix

Glossary

Domain-based Message Authentication Reporting and Conformance DMARC is an email authentication protocol that helps domain owners protect their domain from being misused by others, so-called spoofing.. xvii, 5

DomainKeys Identified Mail DKIM is an email security standard for email using signatures in DNS to securely identify an email sending entity.. xvii, 5

General Data Protection Regulation GDPR is an EU regulation applied 25 May 2018, harmonizing data privacy laws across Europe, making the regulation far more strict in most areas.. xvii, 6

Mail Exhanger A server that handles incoming email from the internet. xvii, 30

Open Source Intelligence A method of collecting information from public sources, usually in a way that the target cannot discover. xvii

phish An email designed to lure a user to disclose personal information. xix, 2

phishing The process of using phish to lure a user to disclose personal information. 1

phishing site A web site used in conjunction with the phish to retrieve personal information. 15

spearphish a phish designed for a particular (often very small) set of users. xix

spearphishing The process of using spearphish to lure a user to disclose personal information. 11

whaling a phish specifically designed for a particular high value user, often with extensive OSINT research in advance. 7

Chapter 1

Introduction

1.1 Topic covered by the project

Different forms of scams have predated electronic communication by far, the motive usually being either personal gain or the benefit of an organization, or even some strange thrill from destroying other people's values. With electronic communication, email in particular, the effort needed to reach out to many potential targets has become substantially lessened. A single email can be sent to tens of thousands of recipients in seconds, and all of these recipients are potential victims of the scammer. Because of this, only a small percentage of victims need to "bite" on these campaigns for them to be worthwhile for the scammer. This is also the reason for the term "phishing", as the scammers or attackers are fishing for a victim to bite onto the scam. There are of course many different kinds of such phishing attacks, depending of the attackers' motivation, and they all have their distinctions and similarities.

From a systems administrator point of view, the goal is to prevent their users from biting, whether it being clicking on malicious links, transferring money into illicit accounts or leaking logon credentials or other personal information to the attacker in one form or another. Technical countermeasures do exist, in the form of spam filters, web filters, next generation firewalls, anti virus software, etc. However, the attacks evolve and circumvent the countermeasures continuously, therefore some of these phishing emails will always get through to the end users. When these technical barriers are bypassed, the only way to prevent an attack from being successful, is by educating the end users enough to recognize the attack and ignore the bait.

1.2 Keywords

Unsolicited electronic mail, Social Engineering attacks, Phishing

1.3 Problem description

Because users need to be able to distinguish phish from legitimate email, and business owners or employers want to spend as little time or resources as possible on activities that do not produce revenue, there is a need to find the most effective ways to train or educate the users. Having a broad base of users, the most efficient method may not be apparent.

1.4 Justification, motivation and benefits

In the late summer and early fall of 2020, a massive phishing campaign hit many Norwegian organizations, including local government [1]. The emails were mainly formulated as urging the users to open infected attachments, claiming that they were invoices needing attention. Some organizations got lucky, since the attachments were stripped away by antivirus software, even though the emails themselves got through, others ended up in the media as their systems got infected by malware. Earlier the same year, several other campaigns containing links to a malicious web site intended for stealing logon credentials were identified by Hammerfest municipality. One of these campaigns was stopped quickly, and the link was monitored in the firewall. Within 30 minutes, about 10 users had clicked the link, clearly indicating that there is some potential for improvement.
Most systems administrators will be able to identify some telltale signs of a phishing campaign and discard the fraudulent emails rather quickly. This, however, is a skill that is learnt by experience, and that many end users do not possess. The purpose of this research is therefore to identify effective ways to transfer this skill onto end users, and try to measure their effectiveness.

1.5 Research Questions

- **Can we mitigate the effects of phishing emails through group specific education?**
  This main question is asked in order to potentially identify differences in receptibility of education methods in various user groups. By grouping the users and providing each group with different education, we may find such differences.
  - **Are some groups more vulnerable to phishing than others?**
    Even though this question has been asked multiple times before, it has not been conclusively answered. By broadening the selection of users and running a study with many subjects, we may find new information.
  - **Is there a significant difference in the efficiency of various education methods when enlightening users about phishing emails?**
    This question has also been asked a number of times, but is not answered conclusively either. By asking this question separately in different user

groups, we may find information that can help us answer the main question.

1.6 Planned contributions

Finding the best education methods may not be as simple as finding an overall most effective one. This project tries to identify what methods are most effective for users who have different ways they use email. At the same time, the planned study is much larger than most others, with a much broader participant base. Some of the education methods have not yet been compared either.

Chapter 2

Related Work

Fighting unsolicited email in general and phish in particular is such an important part of administering an email system today that there have been numerous studies already trying to identify the best way to fight these problems. In fact, statistics [2] show that the volume of unsolicited email is about 50% of all email today, and has actually been as high as 70%. Although just speculation, this decrease over the last few years coincides with the rising use of technical countermeasures like DomainKeys Identified Mail (DKIM) [3] and Domain-based Message Authentication Reporting and Conformance (DMARC) [4]. Most of the studies agree that while technical countermeasures in the form of traditional signature based filters and sandbox systems are an essential part of protecting the users against these threats, it is not enough. Surely, some of the attacks get past the barriers in front of the user, even with modern machine learning algorithms [5], and the important last part of the protection of the users is the education or training of the users themselves.

An important aspect here is also the amount of time it takes from a phishing email arrives, to a user clicks on the links in them. The authors of [6] estimate that half of the persons that do click on the links will do so within 2 hours. After 8 hours, over 90% of the users that click the links will have done so. If this was a new phish that technical countermeasures have not seen yet, there is a very small time frame where these technical countermeasures would have to be updated with new information before they can protect users. Even with extreme update rates, there is no spam filter that can keep up with this, and underlines the need for users to be trained to catch these attempts themselves, simply because the technical systems cannot catch them all.

Of course, there are still ongoing efforts to use technical countermeasures as well. In [7], the authors describe a method where machine learning is used to scan existing email in a user's mail box, and provide a filter or warning system that, in the test cases, were 100% successful. This system detected small variations of legitimate emails, triggering on subtle differences that may not otherwise have been detected. This may be a very good approach, at least for the test cases in this study, however very few automatic systems are completely fail-safe.

Other studies also show promising results from machine learning, like [8], where the authors compare different machine learning techniques. The best algorithm provided a 97.7% accuracy in detecting phishing URLs.

The problem of phishing seems to be ever growing. It has been discussed for decades, and is still a more-or-less loosing battle today. In [9], which is a recent study, the authors were trying to discover the difference in susceptibility between different age groups. They were surprised to discover that both younger and older adults were equally bad at recognizing spam and unsafe emails, both because earlier studies they referred to had shown a difference between the age groups and the low accuracy (66%-72%) in the the participants' ability to recognize harmful emails.

Testing susceptibility for phishing attacks has a few different approaches in itself. Particularly in earlier tests, the subjects were shown different emails or web sites, and were asked to distinguish the phish from the legitimate emails and sites [10]. Later, more studies use the PhishGuru [11] or similar methods, where subjects are served specially crafted phishing emails from their administrators or security teams have been more common. In this way, the crafted phish is provided in between the normal day-to-day activity, and since it is much more close to an actual phishing campaign and captures who actually bites, probably gives a much more accurate overview of the current situation.

In 2018, a fairly extensive literature study [12] was done. Here, the author reviewed a large amount of studies that focused on information security awareness and information security awareness training. Particularly in light of the General Data Protection Regulation (GDPR) that was applied in May that year, the need for systems to have information security built in by design has been important. This study emphasises the need for users to become aware of how they behave, in order to actually comply their business' information security policies. It also points out the fact that the most common threats are insider threats, at least when unintentional breaches of protocol is counted. In order to get the users' attention and focus on these incidents, the author recommends the use of game based learning in particular.

Some studies try to focus not only on *teaching* users the skills they need to heighten information security. For example, the authors of [13] propose a combination of different techniques for awareness training, which they call awareness training 3.0. In this concept, they use both theoretical, market-oriented and an emotional-based approach to make the raised awareness last longer. The focus on games and social communication to trigger an emotional response is central to this idea.

Somewhat in the same area is the concept of Value Based Compliance, as described in [14]. The thought is to use the underlying values that the information security is meant to protect to create the policies that need to be in place. This creates an awareness that helps the user understand what the policy is meant to achieve, instead of it being a top-down approach the user cannot relate to. The same paper describes methods for showing management what causes the users to act the way that they do. This will help create systems and policies that are easier or more convenient to follow.

The news articles demonstrating that phishing still is a big part of the threat picture are many. For example, Skien municipality in Norway was hit with a so-called OWA-phish attack in early 2021 [15]. Norway's largest ISP, Telenor, has tried to identify the biggest online threats of 2021 [16], and explains that phishing is still at the top of the list for the average user. Other reports from for example PwC, name phishing as one of the top five cyber threats of 2020 [17], claiming that phishing is something we should look out for also in 2021.

The fact that the world is facing a pandemic does not make the problem of phishing any less prominent. On the contrary, security firm CheckPoint points out that the activity is clearly on the rise [18], actually utilizing people's fears or need for information to create more malicious content.

Even the central government is not spared when it comes to these problems. During 2020, the Norwegian parliament was under attack [19] [20]. Although little information has been disclosed about the attack, we know that there has been a breach in several email accounts. This, in combination with the fact that the Storting is using the same email systems that has been targeted in the aforementioned OWA-phish attacks, the likelihood of phishing having a central role in the initial parts of the attack is very large. Further, the Norwegian Ministry of Foreign Affairs did also have a breach in 2021 [21]. How the breach was first made is not disclosed, but the data retrieved was used to create a phishing campaign against a list of applicants that should have been confidential. It is fairly obvious from all of these cases that phishing targets all levels of society, and can potentially cause serious issues.

Even as this thesis is being finalized, news articles [22] are written about how information and access from serious, world-wide security breaches are being used for phishing campaigns to gain further access and create further security breaches through malicious software.

2.1 Can we mitigate the effects of phishing emails through group specific education?

There is a variety of different forms of phishing. These range from random wide targeted (if targeted at all) attacks, where the general population (usually of a service) receives a very general phish, to spear phishing, which target a much more specific group, and finally to whaling [23], where an individual is specifically targeted. For each "step up the ladder", the amount of research and effort the attacker has to do is increased. Likewise, the payoff is also significantly higher with each step. For example, whaling usually targets high-level executives of a company, reaching for sensitive information or direct monetary gain, in order to maximize profit, whereas regular phishing hits everyone with the same or very similar attack, usually with an intent to capture login credential or personal information that can be sold off. Spear phishing is somewhere in the middle, utilizing some in-

formation already known in order to specifically target a smaller group of people. The study I am conducting in this Master's thesis will fall into the category spear phishing, despite hitting all users with the same attack, because I use knowledge from the inside of the organizations, both to formulate my emails and creating a user list. The article is mainly concerned with whaling, as the title suggests, and names several cases where whaling has cost companies several million US dollars, illustrating clearly why there is an industry built around these methods.
Everyone is a target - At some point, an argument could be made to say that small businesses (or individuals for that matter) were not targets. However, in recent years, everything is revolving around the internet. This means everyone has a presence online, and obviously can be a potential target. Lately, even the insurance industry has been showing an interest [24].
The broader picture of finding differences in methods of educating users for distinct groups of users seem to be harder to find. While we can find examples of both studies trying to identify which groups are most susceptible to bite on phish [25] and studies that try to identify the best education methods [10], [6] overall, most of these seem to either be too narrow in the selection of participants or in the selection of education methods to be able to answer this question. In order to answer the question properly, there would be a need to look closer at several other groupings of the participants of the studies, not just gender or age. Perhaps even more interesting in this regard is the possibility to study any differences in education levels or what kind of work situation they normally have.

2.1.1 Are some groups more vulnerable to phishing than others?

Some effort has been put into identifying what is the most vulnerable groups of users, particularly age groups and gender differences have been compared. The research is not conclusive in this area, as some studies show a higher vulnerability to fall for phishing attacks with the younger users (ages 18-26) and female users [25], while others see no differences [26], [27]. The explanation for the differences found in [25] is attributed to less education and technical training, however this might not be applicable any longer, since the study is already 10 years old.
There does not seem to be many recent studies of the demographics of users that fall for phishing attacks, and those that try to identify differences usually are narrow in some dimension, for example by focusing on a uniform set of users, like in schools, universities or a single company. That way makes it possible to group the users by age and gender, but further diversity like education level and type of work is often not possible. In order to facilitate this, a study could benefit from having subjects from more diverse companies or even across multiple companies that span more than one industry. The different ways people handle email in their perhaps hugely different workdays can amount to large differences in how they respond to phish. For example, a person washing floors obviously has a completely different relation to email than a medical doctor, a teacher or a government official. Since they all use email in a different way, that may have more of an impact

on their susceptibility to phishing attacks than other groupings. Closely related, there may also be differences in education levels. This may of course be harder to isolate as a factor, since the education level and work situation in many cases depend on each other.

Last year, another Master's thesis [28] was delivered with a similar topic as this one, based on Gjøvik municipality, also in Norway. In his thesis, the author does a similar experiment as I have done, but with a broader angle on whether or not users in general could be educated to withstand the temptation to click on phishing links. His findings are somewhat surprising, as an increasing amount of users clicked the phishing links in the different tests, apart from the final one. He does have an interesting explanation for this, as the different tests were triggering different emotions in the users, creating more appeal on each subsequent test. This makes it harder to see any significant effect of training at all. The conclusion drawn is emphasising that voluntary one-time training does not provide satisfactory results, but also points out that in general, all studies agree that user education is useful and effective. Therefore, there seem to be other variables involved in this study that affect the results.

One very interesting point, though, is the clear tendency that users in management or administrative positions seem to be most likely to bite. This agrees with my previous personal experience, since most ransomware attacks, which is where we usually get our only indications of compromise, have been hitting departments higher in the hierarchy.

Since the study in [28] also was conducted in a Norwegian municipality, there are a number of relevant observations and caveats that I need to be aware of in my own study. In particular, the decision about keeping the participant list static is one I have also made, for the same reasons. There is a need to see an effect on the participants throughout the study, so allowing more participants in the study after it has been started does not make sense. One could argue that removing a participant (and all data created by that participant) from the study if a user for some reason does not work in the organization any longer, but from my experience, this is difficult in practice. I will elaborate on this in Chapter 5. There would also need to be a connection between the user that bite on the phish and the bite itself, not a mere counter.

2.1.2 Is there a significant difference in the efficiency of various education methods when enlightening users about phishing emails?

Of course there have been studies also focusing on the different forms of training as well, partly answering the question raised here. Since there are many approaches to both exploiting training opportunities and timing and how to perform the training itself, different studies tend to compare some of these mechanisms, but often focus on only traditional education or training [10] [29] or focus on different approaches to opportunistic education [11], but not necessarily both. This leaves some room to examine and compare these approaches more closely in this context. The method of immediate feedback to the user is discussed widely, but its role as a perfect method appear to be somewhat disputed [30] [31].

The traditional training is usually some combination of instructor or class room training, computer based training (which again can be devided into instructional videos, cartoons or other computer aided instructions, or even gamification, where the training is made into a game) or written information.

The more opportunistic education methods take advantage of the user actually clicking the link in the emails. In 2010, [11] made an evaluation of different online training programs, and added their own implementation of an anti-phishing education system called "PhishGuru". This system sends emails from systems administrators or other people that want to train the users, and much like real phishing emails, the users get coerced into clicking a link, where the PhishGuru system provides an "intervention", telling the user that they should be more careful with these links. In this way, the system utilizes a "teachable moment", where the user is supposed to be extra susceptible to acquire knowledge. The study shows this to be a very effective way of educating users. The study did test the users also after some time had passed, and was able to show that the users retained the knowledge longer after applying this method.

The same study also used a game called "Anti-phishing Phil" to educate users. This is an example of the gamification method, which transforms the training into a fun activity, the idea being that knowledge is easier to acquire when the person is having fun. The game is simple, where a fish called Phil is presented with worms he wants to eat. Each worm is presenting a URL that the user needs to identify as legitimate or bad, and they need to steer Phil toward the legitimate URL so he can eat it. This of course focuses on the user being able to identify suspicious URLs. However, this may not be so simple today, as we are much more likely to be using different mobile devices nowadays. On these devices, as there is no mouse pointer to hover over the link to see where it leads, the actual URL may be very hard to determine if you do not know what you are looking for. On a traditional computer, most email clients will show this URL.

In [11], these these methods are tested against the security notices that most admins use to raise awareness of phishing attacks, and concludes that these security notices are ineffective, and that both embedded training in PhishGuru and gamification in "Anti-phishing Phil" are effective at training and retaining knowledge.

The authors of [10] looks at different ways of making users identify potential problems in URLs. It used three different formats of education, 45 minutes of instructor based training, an Android app that the user could use at their own pace, or textual information. In this study, the instructor based training showed the best results, both in regard to how well the participants received the knowledge, how confident it made them of their performance, and how well they liked the training. Text-based training was the most time-efficient of the three. What is very interesting is in the discussion section, where the authors points out that all tests are done in a short time, exposing them to many more phishing emails in a short time period than what would normally be expected, possibly making it easier to keep the users in an anti-phishing mind set. The fact that the study is done in a school can also cause a bit of skew in the results as the students are already used to instructor based teaching.

Another way of identifying a phishing attack, is by looking at the web sites these links point to. There are quite a few methods of attacking users that only click on the link in an email, by infecting their computers with viruses or other malware. Phishing attacks are usually only after stealing personal information, however, and the easiest way is to lure the user into providing this information themselves, making this a viable option also for educating users, since some are bound to click the links anyway. In [32], the focus is on this approach, and while the authors of this study do not test different training techniques, they try to identify what indicators users are looking for and what indicators are usually missed, making it an interesting starting point for creating a training program, showing where there is room for improvement.

Up until this point, we have assumed that a phish is a phish, and that they all look the same. However, there are many forms of phish, from the general poorly Google-translated mails that get sprayed around the entire Internet to the carefully crafted spearphishing attempts often called CEO Frauds, where the attacker often uses signatures or even writing styles of insiders of the company they are attacking. This may become a factor when determining what kind of training is needed, as well as help assess the training methods. In [33], the authors have made an effort to systematically assign a difficulty level to phishing emails in order to help CISOs or administrators determine how good their training is. To achieve this, they introduced a measurement called "Phish Scale". This system can be helpful when designing a phishing campaign for the purposes of studies like these.

Chapter 3

Method

3.1 Overview

In order to answer the research questions, I have done a literature study, before a study is conducted, composed by a phishing campaign, then some form of education or training, before another phishing campaign.

After the experiment is finished, the different training methods are compared within the different worker groups, to see whether the different training methods give different results within each group.

Finally, after running the second phishing campaign and having collected the results, all users are informed of the experiment, informing them of their rights in regard to personal information collected and thanking them for their participation.

Two municipalities in Norway have been selected for this study, both because of their size - Municipalities are often the largest employers in their area - and because of their diversity in jobs - The municipality handles everything from road sweeping and kindergartens to medical practices and city planning. Alta and Hammerfest are the largest municipalities in the former county of Finnmark, in Northern Norway. Since the researcher is an employee of Hammerfest municipality, easy access to the user database made this a natural choice. Alta was invited to join after the initial idea was formed, as the two municipalities also cooperates in different areas, making the threshold to invite Alta municipality lower. There should, however, be no conflict of interest in this study, as there is no consequence or prejudice involved in this particular study. There may be differences in the results for Alta and Hammerfest, particularly due to the infrastructure used in the study being located in Hammerfest, and that the injection of email into each organization is not the done the same way. The delivery will be uniform within each municipality, however, and comparing the municipalities is not within the scope of this study.

The study will be tracking user behaviour by collecting the users' username/email address. To comply with rules and regulations, the project needed acceptance from the Norwegian Centre for Research Data, NSD. Being a deception research project, this was a lengthy process, finally getting approval 25.03.2021. This, in turn, has made the schedule for the project tight, and may account for some shortcomings, discussed later.

3.2 User Groups

The different types of workers are selected based on their different ways of handling computers in general, partly also their pay grade and education levels. An outdoor worker, like a gardener or road maintenance worker, will obviously have a different approach to everyday computer usage than a case worker handling incoming various applications from residents. Departments that are similar in this regard are grouped together to form larger groups.

The grouping is determined by their placement in the user catalogs of the municipalities, and their group memberships in the catalog. This may be somewhat inaccurate, which will be discussed later, but should be sufficient for this study.

The following grouping is used in this study:

- Management
- Kindergarten
- School
- Treatment
- Care
- Physical
- Economy and HR
- Culture
- Technical Case Worker
- Other Case Worker
- Other

Some of these are self apparent, like Kindergarten, School, Economy and HR and Other. The others may need some extra information. First, Management is a group of people somehow tagged with having a management position. Since a person cannot be a part of different groups, these persons will not appear elsewhere, even if they technically belong to other groups.

Treatment is extruded from the healthcare departments, as the persons actually providing treatment is often a little different from the other care giving professions, both in education and how they work. The persons in this group are mainly from medical practices and physical therapists offices.

Care is mainly composed of care institutions for the elderly and others, mostly nurses and similar professions.

Physical is the group of people doing more physical labour, for example firefighters, gardeners, mechanics, road maintenance, etc.

Culture includes libraries, culture institutions, cultural schools and similar.

The Case workers have been divided because there are mainly two different types - The technical case workers, that are involved in building applications and area planning and projects, while the others handle other case types.

3.3 First phish

Establishing a baseline is done through a specially designed phishing campaign. The first phish needs to be sufficiently difficult to recognize in order to get enough responses for statistics. One of the education methods I want to provide is immediate feedback (See Section 3.4 for all methods), therefore the numbers also need to be large enough to be able to draw users from the ones that bite.
A natural start for a hard-to-discover phish, is mimicking an email from an often seen sender that may send to the entire organization. One such email is a status or warning from the IT department. Since many Norwegian municipalities have been hit with so-called OWA phishing (Stolen credentials are used to log on to an internal account via Outlook Web Access and sending phish from a legitimate account, in order to steal more credentials) lately, the theme of this first campaign was chosen to be in the form of an IT employee, urging the users to click on a link to check whether their password was stolen.
This method is somewhat easier in smaller organizations, since people may know many of the IT people by name, if not necessarily in person. To make it possible to distinguish this email from legitimate emails at all, a few indicators where crafted particularly for this purpose.

- Only the given name of the "sender" was used, "Gorm" for Alta, "Audun" for Hammerfest. Surnames could not be found anywhere.
- The greeting of the email was "Hallo!", which is normally only used in oral communication.
- The signature of the email only used a first name, and used the abbreviation "IT-avd.", instead of the full name of the department along with a work title, which we usually do.
- The sender address was made to look like, but not actually be, one of the official addresses of the "senders". Both municipalities use domains that end in `.kommune.no`, but in order to do this study, I had registered the domain `kommmune.no` (with 3 m's), which was available.
- Each mail was sent to an individual recipient, not to mailing lists, which are usually used.

These indicators should be enough for the users to be able to tell the difference, while still leaving enough room for people to do mistakes. We are not interested in either impossible to distinguish emails or completely obvious ones.
The link itself was for a site in the aforementioned `kommmune.no` domain, placed internally in Hammerfest municipality's infrastructure, but with a public IP address, so it was reachable over the internet, as well as from Alta. Every mention of Hammerfest was also replaced with Alta for the emails that went to Alta. Otherwise, the emails were identical.
The phishing site was then a copy of the actual ADFS logon page in Hammerfest,

and the webmail login page from Alta. In order to not compromise the users' passwords, the pages were modified to not include that field of the logon form when submitted.

The logic on the phishing pages registered both who clicked the links and who entered information with a time stamp.

3.4 Education

Each of the user groups was then divided into five different education groups;

- Text based education
- Web based education
- Web based education with immediate feedback
- Video based education
- Game based education

Text based education is, as the name suggests, an informational text sent in the form of an email, with a list of the indicators mentioned earlier.

Web based education is a web page with images from the phishing emails, where the indicators are circled and explained.

The immediate feedback group needed to be drawn as users bit on the phishing campaign, in order for it to work. Since there was no possible way of predicting who would actually bite, every fifth person that took the bait from each of the user groups was served this method. They were redirected to a page telling them that they shouldn't have clicked the link, nor left their username and password, before being directed to the web education page.

The video education is a 2.5 minute video where I explained the same indicators as the earlier methods, while showing them on-screen.

The game is a simple game where the user is presented with 10 different images of emails, and get to choose whether the image displayed is a legitimate email or a phishing email. The user also has to indicate how certain he/she is of his/her answer, which in turn affects the score. The score for a correct answer is directly connected to the confidence, with 0 points for 0 confidence, and 100 for full, in steps of 10. If the answer is wrong, the user looses points for every step the confidence is above 30. After the final question, the user is presented with the score, how many have reached the maximum score, and finally their rank among all players, based on their best score. There is nothing stopping them from trying again to improve their scores.

These education methods were chosen for their simplicity and the time consumption in a busy day. The cost would therefore be more or less equal, measured in time used for each employee, and it would be kept to within about five minutes, which should be reasonable and possible to find room for in most situations.

3.5 Second Phish

The final test to determine the effect of the training also needed to be difficult, but not impossible, for the user to expose as a phish. Since both the names, domains and logon sites used in the first test now would arise suspicion, the next phish needed to be carefully designed.

Information and links from the IT departments were out of the question because of the rising suspicion. Links coming from other departments internally in the organization would expose the project for the people working closely with those departments, making that approach less useful. The "source" of the campaign therefore needed to be external, but at the same time have significance for all employees.

Because of the large amount of real phishing campaigns during the last year, users seem to be generally more careful when it comes to the larger companies operating in Norway. There was also the point of reaching every user, and even though one might hit broad with one of the major telecom companies, there are still many people that can reveal the campaign simply because they do not have any connection to the company. In the end, the choice fell on using some local businesses that give discounts to employees. That way, the campaign would be relevant for all users.

In the last minute, the local business here in Hammerfest withdrew their support. However, the business in Alta was willing to let me use their name for both municipalities. This was of course a setback, but the name and reputation of the Alta based business is well known also in Hammerfest, making it a good substitute.

The second phish will therefore be an email asking the user to register to receive a 50% discount on his/her next purchase with Kokkejævel, which is known as a popular local food store. Technically, the store is called Hoftepluss, with the owner calling himself Kokkejævel, but this confusion in names is one of the indicators I want to provide to the users, so they may discover that it is a phish if they pay attention. Again, a domain similar to (one of) the legitimate one(s) is registered, in this case I used `kokkejavel.no`, as opposed to the legitimate `kokkejævel.no`. The phishing web site used the layout from `hoftepluss.no`, which served my purpose better.

Of course, the user is in this way not asked to enter their current password, but by registering on a fraudulent site, the scammer has a potential usable username/-password combination, since most people reuse their passwords.

There are a set of indicators also in this second phish that the user could become aware of;

- The sender address was `hoftepluss@kokkejavel.no`.
  - The domain is not associated with the business it is claiming to be from.
  - The business names are mashed up in a way that is unnatural - They should be the other way around.
- The email was apparently sent to a distribution group with all employees.

  - The sending of email to these kinds of groups is restricted to a few select employees, and never allowed from outside the organization.
  - The distribution groups do not exist. The names are made up, and cannot be found in the address lists.
- There are no logos or graphics in the email.
- The email is urging the user to hurry.
- The link also points to the domain `kokkejavel.no`, which is not associated with the business it claims to be.
- The layout of the web page is taken from `hofteplus.no`, not `kokkejævel.no`, which is the domain name I am mimicking.

That way, there is a possibility for both users biting, and for users to discover that this is another phish.

Chapter 4

Results

4.1 Initial Results of First Campaign

While not directly relevant for the research questions, there are a few interesting observations that were made already at this point.

Within ten minutes of sending the first emails, the support phones of both IT departments where ringing constantly with both users warning us that there was an ongoing attack and worried users that wondered if this was a legitimate email. It has not been possible to measure the volumes of these inquiries, since most phone calls are not registered, and there are a number of different people answering these lines. The volume was, however, so large that I was asked to inform the department in advance and be present at the office myself to help for the next round. In addition, I was contacted in virtually every possible way by people in my other circles. This included both Facebook Messenger, SMS, Microsoft Teams Chat, Microsoft Teams Calls, as well as direct phone calls both to my office and personal mobile phone.

Again, these inquiries were ranging from people warning me that I had been "hacked" to asking whether or not the email was real to people realizing that this actually was a test.

In Alta, Gorm has reported the exact same reaction from their users, indicating that this is not an isolated incident.

Table 4.1 shows the distribution of bites and information entered for each of the worker groups in each of the municipalities.

As we can see from the table, all groups have at least one bite where the user has entered their information. It is also clear that in this first phishing campaign, the users in Hammerfest are far more happy to both click the link and enter their login information than the users in Alta. This makes the numbers less comparable, so in order to distinguish what groups are more susceptible to phishing, I have split the table into separate figures, shown in Figure 4.1.

While some groups do appear high on the list in one municipality, and low in the other, there are a few groups that stand out as more universal. For example, the management groups are among the most eager to click and leave information in

Table 4.1: Distribution of bites and information entered by each of the worker groups

Worker Group	Total mails sent	Clicks	Information Entered
Hammerfest Management	142	73 (51.4%)	59 (41.5%)
Hammerfest Kindergarten	186	24 (12.9%)	16 (8.6%)
Hammerfest School	341	77 (22.6%)	62 (18.2%)
Hammerfest Treatment	94	19 (20.2%)	12 (12.8%)
Hammerfest Care	917	116 (12.6%)	80 (8.7%)
Hammerfest Physical	157	22 (14.0%)	16 (10.2%)
Hammerfest Economy and HR	60	13 (21.7%)	11 (18.3%)
Hammerfest Culture	28	8 (28.6%)	5 (17.9%)
Hammerfest Technical Case Worker	24	3 (12.5%)	3 (12.5%)
Hammerfest Other Case Worker	29	13 (44.8%)	11 (37.9%)
Hammerfest Other	65	1 (1.5%)	1 (1.5%)
Alta Management	85	21 (24.7%)	12 (14.1%)
Alta Kindergarten	190	13 (6.8%)	9 (4.7%)
Alta School	554	81 (14.6%)	45 (8.1%)
Alta Treatment	266	31 (11.7%)	17 (6.4%)
Alta Care	1165	91 (7.8%)	45 (3.9%)
Alta Physical	126	12 (9.5%)	7 (5.6%)
Alta Economy and HR	34	3 (8.8%)	1 (2.9%)
Alta Culture	41	12 (29.3%)	6 (14.6%)
Alta Technical Case Worker	84	12 (14.3%)	10 (11.9%)
Alta Other Case Worker	59	7 (11.9%)	2 (3.4%)
Alta Other	26	3 (11.5%)	2 (7.8%)
Total	4673	656 (14.0%)	433 (9.3%)

both municipalities, as are the culture and school groups. In the other end of the scale, we find the care and kindergarten groups, which seem less prone to bite.

4.2 Training

The education of the users in itself should not provide any significant results, since all training, except gamification, is only sent to the user without requiring feedback. However, I have collected information on whether the user accessed the web based education, the video or the game. The text based education was just plain text in an email, meaning I could not track the user when reading. The amount of users receiving the training was disappointing, as Table 4.2 shows, with only a few groups above 50%. The numbers were even lower before sending a reminder to all users not having received their training.

This also has the effect of opening up another possibility. The users not having received their training can be seen as a sixth training group, in addition to the earlier defined ones.

Also, there was once again a fair amount of feedback from users wondering if this information was also a scam, naturally being sceptical towards any links coming from myself because of the first phishing attempt. This was despite having informed the users several times from other accounts that the information was coming, how it was coming, and why it was sent using the same mechanism as the phish. Again, the feedback received was going through several channels, including the support lines, making it impossible to measure.

Table 4.2: Number of users registered as having received education by worker group

Worker group	Total mails sent	Training received	Training not received	Percentage received
Hammerfest Management	142	61	81	57.0%
Hammerfest Kindergarten	186	64	122	34.4%
Hammerfest School	341	139	202	40.8%
Hammerfest Treatment	94	39	55	41.5%
Hammerfest Care	917	311	606	33.9%
Hammerfest Physical	157	55	102	35.0%
Hammerfest Economy and HR	60	32	28	53.3%
Hammerfest Culture	28	14	14	50%
Hammerfest Technical Case Worker	24	10	14	41.7%
Hammerfest Other Case Worker	29	18	11	62.1%
Hammerfest Other	65	18	47	27.7%
Alta Management	85	37	48	43.5%
Alta Kindergarten	190	54	136	28.4%
Alta School	554	182	372	32.9%
Alta Treatment	266	81	185	30.5%
Alta Care	1165	334	831	28.7%
Alta Physical	126	40	86	31.7%
Alta Economy and HR	34	16	18	47.1%
Alta Culture	41	16	25	39.0%
Alta Technical Case Worker	84	33	51	39.3%
Alta Other Case Worker	59	26	33	44.1%
Alta Other	26	10	16	38.5%

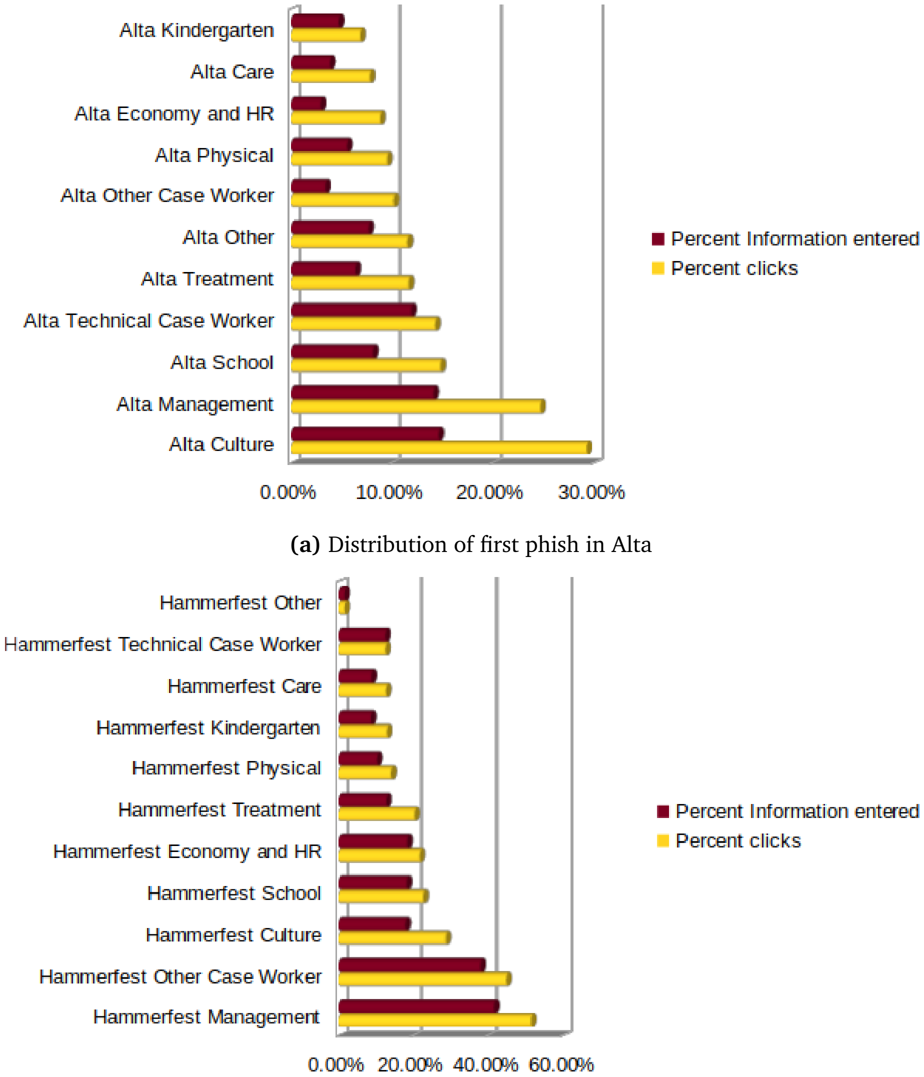

(a) Distribution of first phish in Alta

(b) Distribution of first phish in Hammerfest

Figure 4.1: The distribution of first phish bites by municipality

4.2.1 Gamification

The only education method involving response from the users was the gamification, where users got scores that went into a scoreboard. As stated in the previous section, response was slow as in the other education methods, until the reminder was sent out. This reminder also told the users they could try the game several times, and informed them that there were only two top scorers at the time. The scoreboard has reached 137 entries, 8 of these having a maximum score.
Getting a full score is not likely, given the form of the game, as the confidence the user has on each question is influencing the score. One likely explanation for having users reach a full score, is that they have been trying multiple times, and hopefully learned something from their experience.

4.3 Second Phish

The second phish was delivered to the users in the afternoon of May 18, 2021. In order to have time for analyzing the data, they were collected at noon, May 21, 2021. The issues concerning this short time span will be discussed in Chapter 5. In short, the results from the first phish have been limited to the bites from a similar time period.
At the time of data collection, there were 344 bites registered, whereof 145 users had entered some information in the form, potentially leaving usable credentials. Compared to the 621 bites and 410 information enters of the three first work days after the first phish, at least we are seeing users being significantly more sceptical to links in the second attempt. Digging deeper into the numbers, I have extracted some special observations (Table 4.3) from the data, counting the number of users that did not fail any tests, the users that did bite on all phish and, perhaps most interestingly, how many bit on only one of the phish. In particular, the number of users that only clicked the link and entered their information *after* I had put focus to the phishing problem. Even users tagged with "training received", which indicated that they either had clicked the training link they received, or that they received textual education, which couldn't be tracked, have a surprising amount of fails of the second test. Although the numbers in this study did not rise with each attempt, like in [28], this may in fact be another indication that we are seeing some of the same effects in this study.
In Figure 4.2, we can see that the distribution of bites is a bit different from in the first phish. Most noteworthy is the fact that both management groups have moved to the other end of the list, while both school groups and the Hammerfest culture group keep clicking and entering information. Also, the Alta culture group has gone from being the "worst" group in Alta with regards to the first phish, to not having a single user entering information on the second phish. This may indicate a connection between the different types or approaches of the phish and the group that is most susceptible to that particular type.

Table 4.3: Special observations made after experiment concluded

Observation	Count
User did not bite on any phish	3778
User bit on only first phish	551
User bit on only second phish	240
User bit on only second phish (training received bit set)	111
User entered information only on first phish	401
User entered information only on second phish	117
User entered information only on second phish (training received bit set)	55
User entered information on both phish	28

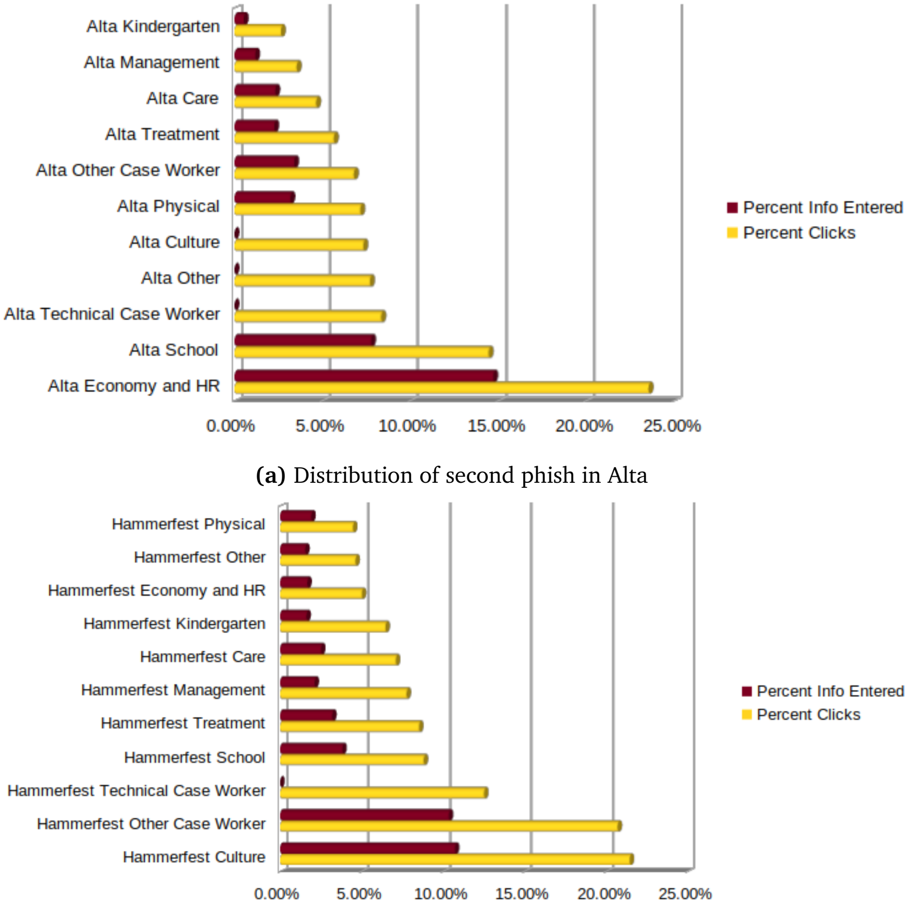

(a) Distribution of second phish in Alta

(b) Distribution of second phish in Hammerfest

Figure 4.2: The distribution of second phish bites by municipality

4.3.1 Education methods

As mentioned in Section 4.2, the number of users receiving their selected education was disappointing. In some cases, there were groups that did not receive some education methods at all. Since the groups in Alta and Hammerfest should contain the same type of people, and I by the time of sending out the second phish had overcome technical difficulties in sending mail to Alta users, there should be no reason to keep them separate. In addition, some groups with similar functions have been aggregated in order to have reasonably sized education-per-worker-group groups.
The numbers of total users in each new group, the number of clicks for each group, and the number of users who entered their information are listed in Table 4.4, Table 4.5 and Table 4.6, respectively. After calculating the percentages of clicks and information enters from the total in each group (Table 4.7 and Table 4.8), we end up with the distribution shown in Figure 4.3 and Figure 4.4. While the 'No education' type seems to be doing best, this is due to the fact that this education method also includes every user that did not respond in any way. This includes every person that for some reason hasn't been checking their mail, whether it being they are on vacation, have quit, or even that the initial list of users included users that do not exist. Of course, it also includes those users who (believe they) know everything about detecting phishing attempts or believed the education itself was a phish. It is included to illustrate where the rest of the users in the experiment is located in the data set. Most of the same argument is true for the text based education as well, since this has not been trackable.

Table 4.4: Total number of users in each worker group, by type of education

Group name	Text	Immediate	Web	Video	Game	None
Management/economy/HR	71	9	25	16	27	148
Kindergarten/School	313	13	41	25	47	832
Care	515	11	38	40	41	1437
Treatment/Case workers	138	8	20	22	20	348
Physical/Culture/Other	106	4	19	13	18	296

Table 4.5: Number of clicks from users in each worker group by type of education

Group name	Text	Immediate	Web	Video	Game	None
Management/economy/HR	7	3	1	1	4	9
Kindergarten/School	34	3	7	6	11	66
Care	35	1	5	8	9	61
Treatment/Case workers	14	2	4	5	4	14
Physical/Culture/Other	8	0	1	3	5	13

Table 4.6: Number of users who entered information in worker group by type of education

Group name	Text	Immediate	Web	Video	Game	None
Management/economy/HR	1	2	0	0	1	6
Kindergarten/School	16	0	3	2	5	34
Care	17	0	1	2	5	25
Treatment/Case workers	5	1	1	0	3	4
Physical/Culture/Other	3	0	1	0	1	6

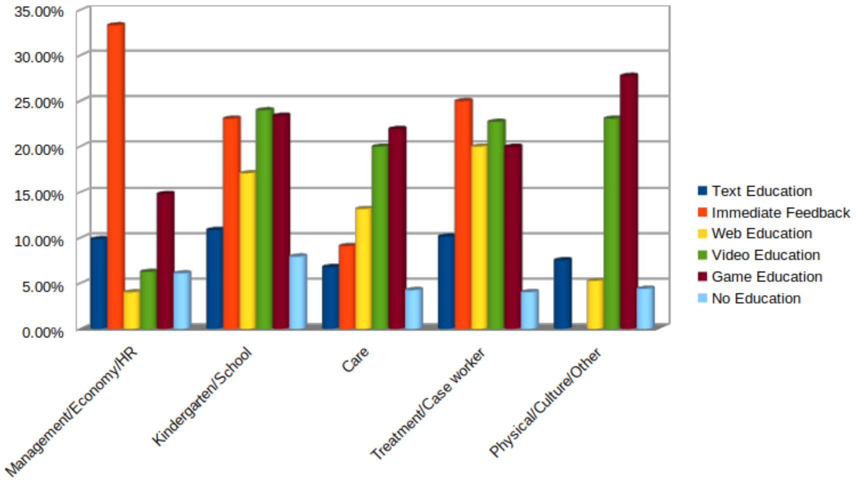

Figure 4.3: Distribution of clicks in each group by education type

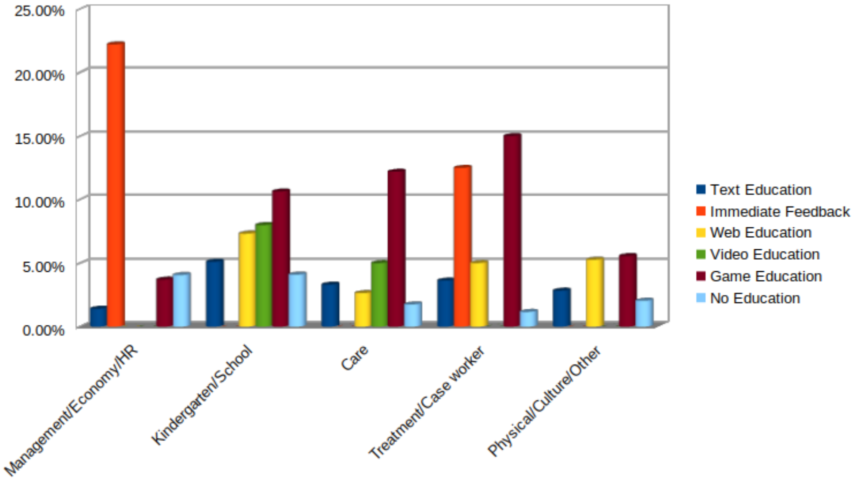

Figure 4.4: Distribution of information entered in each group by education type

Table 4.7: Percentage of clicks from users in each worker group/education group combination

Group name	Text	Immediate	Web	Video	Game	None
Management/economy/HR	9.86%	33.33%	4.00%	6.25%	14.81%	6.08%
Kindergarten/School	10.86%	23.08%	17.07%	24.00%	23.40%	7.93%
Treatment/Case workers	10.14%	25.00%	20.00%	22.73%	20.00%	4.02%

Table 4.8: Percentage of users who entered information in each worker group/education group combination

Group name	Text	Immediate	Web	Video	Game	None
Management/economy/HR	1.41%	22.22%	0.00%	0.00%	3.70%	4.05%
Kindergarten/School	5.11%	0.00%	7.32%	8.00%	10.64%	4.09%
Care	3.30%	0.00%	2.63%	5.00%	12.20%	1.74%
Treatment/Case workers	3.62%	12.50%	5.00%	0.00%	15.00%	1.15%
Physical/Culture/Other	2.83%	0.00%	5.26%	0.00%	5.56%	2.03%

Looking at the remaining education methods, we can see that there are a few education methods that stand out as apparently better than the others for each group with regard to clicking the links. For instance, Management/Economy and HR seem to favour web or video based education, kindergarten and school users score slightly better with web based education, treatment and case workers have only slight variations, while both care and physical, culture and other groups seem to be best off with immediate feedback and web based education.

Seemingly, the immediate feedback and video education is very effective with several groups for preventing users from entering information. However, as the numbers of Table 4.6 did show, there simply is so little data that if one single user made another decision, the results could be completely altered. The only two groups that have some data to compare, are the kindergarten and school group and the care group. Here, we can see that immediate feedback has a perfect score, while web and video education score better than game education.

Chapter 5

Discussion

Although care has been taken to control other variables in the project, there is always the chance that there is uncontrollable conditions involved.

5.1 Group Definitions

The groups of workers are based on the different typical departments in a Norwegian municipality. Because of the different ways of organizing the departments, there may of course be other ways to define these groups. My connection with Hammerfest municipality made the divisions here natural, and mostly also fits with Alta. The groups are of varying sizes, as is to be expected.

The lines between each group may still be unclear in some cases. Particularly when it comes to persons working in different departments, determining where to put that person becomes a bit more difficult. While Alta uses a hierarchical structure for their users, Hammerfest users are organized in groups, but otherwise flat. This flat structure made deciding which department was a user's main workplace difficult. The hierarchical structure is easier to make a first decision, but might not always give the correct placement either.

Sending the same phish twice to any person would reveal to the user that this is not legitimate, so any user is represented only once. Since the database collecting the data and providing the sending scripts with the email addresses is a standard SQL server, making sure the email address was unique also provided a possibility.

All groups are populated in sequence. Management is defined as anyone having having a management role, personnel responsibility or similar, so these users will turn up in more of the groups. However, they will already be in the database in the Management group, so they will not be added a second time.

Further down the list of groups, we may encounter people having multiple roles, but as with the Management group, if the user is found in the database already, it will not be added a second time. While the sequence of the groups may be a topic for discussion, the order has been consistent between the organizations in this study.

5.2 Other errors in worker lists

Extracting user lists based on workplaces or departments is always a difficult task, regardless of method. In my experience, even determining what people actually are on the payroll is a horrendous task, and even worse if you cannot involve the HR department. The IT department are sometimes the last to know when people are hired, and rarely get any notice when someone quits. This, in turn, will affect the lists I have had to work with, and subsequently, there may be users in the list that do not work in the organization any longer.

This is of course also an issue after the list has been extracted and entered into the infrastructure used for these experiments. Because the user list is hard to determine in the first place, there is no clear path to updating such a list in the middle of the experiment. Due to this fact, I know there are people in the list that are no longer working in the municipalities. They have not been removed from the system in an attempt not to skew any numbers, and let all groups be treated equally by not altering the list after the experiment started.

5.3 Differences between the municipalities

5.3.1 Technical difficulties

A lot of effort has been put into making the experiment equal for all groups across both municipalities. However, there have been technical difficulties especially when it comes to delivering emails to Alta. Since the infrastructure has been located in Hammerfest, and the email security systems here are well known to the researcher, delivering emails has been a fairly easy task, without much interference. In Alta, however, the email systems did not act very friendly when confronted with the amount of email coming in through their Mail Exhanger (MX) from our experiment infrastructure. This resulted in the first phish being dumped into the email systems in Hammerfest within about half an hour, while the same phish to Alta was delivered over a longer period of time, about a day. There may also have been emails that were lost in transit, but those that were known to be lost, were resent as soon as it was known. This may account for some of the lower hit counts in Alta in the first phishing campaign, as people got time to discuss what they had seen with their coworkers before they received their own phishing email. This was of course not possible to discover before the first phish, but was later taken into account, so that for each Alta-email, the script would pause for 15 seconds. Since the script sent the emails in alphabetical order, Alta and Hammerfest recipients were shuffled, and the 15 second delays also affected Hammerfest users, so the rate was similar across both municipalities. Another problem occurred when using this method, as the delay was so long that the MX everything went through reset the connection after about 75% of the list. This usually occurred at night, and was easily restarted the following morning. Also, since the list was alphabetically sorted, the remainder should be fairly evenly distributed between both municipalities and groups.

5.3.2 Other differences

There are some subtle differences between the way Hammerfest and Alta send their emails. This experiment has been done using my personal style of writing, and in that way probably also been easier to determine to be fake in Alta than in Hammerfest. In particular, the first phish was made to look like it was sent from a named person in the IT department. Although I did confer with the person that should be impersonated in Alta, and some of the indicators of a scam was supposed to be present, they may have been stronger in Alta, simply because the users there are used to another style of writing and quite a different signature at the bottom.

This, and what was discussed in the previous section, may account for the rather large difference between the bite rates of the first phish between the municipalities. In the second phish, the technical differences were eliminated, and the phish was imitating an external sender for both localities. However, since the possible skew in difficulty of the phish is uniform across all the groups of the municipality, this should not pose serious issues.

5.4 First phish

The large amount of users that reported this to the IT departments could of course indicate that our users are very good at discovering fraudulent emails, but we have no indication of this being the case. It is far more likely that we simply had too many indicators available for the users to discover the nature of the email. Another possible reason is our use of well-known names within the organizations, making it easier to spot the difference between the phish and our normal style of communication than anticipated. As the previous section showed, this effect would be even larger with the Alta users.

5.5 Second phish

5.5.1 Aggregating groups

When looking at the data after the second phish, it quickly became apparent that the individual groups had too little data to make sense in an analysis by themselves. Therefore, I have combined some groups where users are more or less similar. Economy and HR do use computers in many ways similar to those in management positions, and are all valuable targets for an attacker. Kindergartens and schools probably use computers quite differently, but they still are a kind of educators, which indicates they can be combined. Both kinds of case workers and the treatment users use computers similarly, as they have their specialized applications for their functions, and generally use computers often. Also, the education level is fairly high in most cases. Physical, culture and others are generally users that do not use computers as frequently as others. There is of course a wide range of different kinds of computer users in each of the groups, but this should be a fair division of the groups to be able to find any potential differences.

5.5.2 Short time between second phish and collection of data

Due to the time pressure to get the research done before the deadline, the data collection from the second phish needed to be done only a few days after sending the email. Also, the owner of the business whose name I used in this campaign was very clear that I needed to expressly explain the situation. In order to compensate for this short amount of time, all comparisons against the first phishing campaign is done with data from only the first three work days after the first phish was sent. Previous studies, like [6] have shown that if a user is going to click a link in a phishing email, they are going to do so within a very short amount of time. This can also be seen in this study, as the amount of users clicking the link in the first phish for over one month after the initial three work days included in the data is less than 10 percent (60 users of 656 total clicks, 7 of those *after* they had received training) This, in addition to the adjustment of the data for comparison in the first phish, should be enough for the data from the second phish to be adequate for the study.

5.6 Interpreting the results

5.6.1 First phish

The distribution of users biting is quite different in the two municipalities, with only some worker groups showing similar trends for the first phish. This indicates that the results may be more individual than anticipated, or that other factors play a more significant role. Seeing that the management group is eager to both click the links and leaving their username and password on the phishing site is consistent with what was found in [28]. This comes as no surprise to me personally, as this is also my experience from over 18 years as an IT administrator in the public sector.

Care and kindergarten groups are a bit different, as they use computers differently. In kindergarten, there are a few users that have the role of educational leader. These persons use the computer in preparing activities for the children and have contact with parents. However, most of the employees in the kindergartens do not have this role, and do not generally use computers in their work. All employees have email, and will have received the phish, but not all users will check their work email on a regular basis. The same is true for the care users, as many of these users only use computers to read or write the reports on their shifts. Using an email client is not necessarily something that they do every day. Also, working shifts mean they often have longer periods off. This makes it hard to determine exactly why these groups do not bite as often as others.

School users, on the other hand, are employees that use email more frequently, as most have day-to-day communication with colleagues and parents. My knowledge of how culture workers use email in their everyday situation is limited, but librarians and museum personnel will almost certainly use email for communication frequently, and I can imagine the same being the case for culture schools, cinemas and the like.

This could indicate that the difference in results for the groups is not actually related to the work that people do directly, but is tighter related to the frequency that the person is checking their email.

5.6.2 Second phish

In this part of the study, it became apparent that we had many users that responded to the second phish without responding to the first. This was completely unexpected, particularly with users that had received their education as well. There is a chance that people by now was expecting such campaigns from trusted sources, and were curious as to what this could be, but that would probably not account for all of these. Using other methods to lure the users into clicking the link than the original technical urgency and fear of losing their passwords into the wrong hands, may play a role here. In the second phish, I played on the user's wish to save money personally and make a good deal. This may trigger completely different responses in people, and as such, may also trigger a response in a different set of people.

The idea was that this would not necessarily be an issue, as everyone received one of the predetermined, randomly assigned education methods. These methods were also specially designed to convey the same information (with one exception), just using different media. Since every user also received the same phish, the education methods was meant to be the variable that would show differences in the responses afterwards. I initially thought it strange that the gamification education method seemed to do worse than the others. This was the one method I personally had assumed would do better than the rest. However, this method was the one that did not explain in detail what to look for to discern a phish from legitimate email. The same information was there, but the user was not told how to find it, and may therefore be the reason why it scored badly.

As mentioned in Chapter 4, particularly the data set of users leaving information on the phishing site is limited, which creates large deviance if a single user had decided differently. Since the potential for harm is still very much present by only clicking malign links, I have decided to focus on this part of the data set. The same problem really applies to the data of the immediate feedback method, and it may be just curiosity from single users that make extreme responses in the statistics, like in the management, economy and HR graph, where this method is by far doing the worst.

Chapter 6

Conclusion and future work

6.1 Conclusion

We started off with the research questions below,

- Can we mitigate the effects of phishing emails through group specific education?
    - Are some groups more vulnerable to phishing than others?
    - Is there a significant difference in the efficiency of various education methods when enlightening users about phishing emails?

It would be natural to conclude the two subquestions before reaching a final conclusion for the main research question.

6.1.1 Are some groups more vulnerable to phishing than others?

It appears reasonably clear that this question is answered quite thoroughly, as we do see a clear tendency in several of the groups. In particular, the management, school and culture groups appear more vulnerable than the others, while the care and kindergarten groups seem less vulnerable. However, the reason why these groups behave differently is not entirely clear. The reason may, in fact, be as simple as the frequency at which the user is logged into an email client.

6.1.2 Is there a significant difference in the efficiency of various education methods when enlightening users about phishing emails?

As we have seen, the different groups do seem to respond differently to the different education methods, although they seem to respond mostly different to one or two single methods, while the other methods are not as much affected. In this case, we saw that management, economy and HR users seemed to do better with web or video education, kindergarten and school users scored slightly better with web education, while care and physical, culture and other users did best with immediate feedback and web education. Treatment and case worker users did not show any particular improvement with any of the educations. This effect is sufficient to say there is a difference, but to determine what can be done to exploit this has yet to be discovered.

6.1.3 Can we mitigate the effects of phishing emails through group specific education?

From this study, there do seem to be some differences that perhaps can be exploited. The differences, however, point in the direction that only the web based education tends to do better than the others in different degrees, while the remaining methods are doing equally bad. There is some evidence that video education might work for management, and that immediate feedback can work for care users and physical, culture and other users.

But can group specific education help mitigate the effects of phishing emails? That is not likely, at least not in the form tried in this study. There is still a possibility that we could train users based on spearphish, or even whaling in extreme cases, targeted specifically for each group. The effort needed to create such education unfortunately also increases rapidly, and there is a good chance this is simply not a viable option. The chance that generic education, if provided frequently enough, will have a much better effect in users than targeting specific users with specific education. At least, creating group specific education will not be cost effective in general, but can be tried if such user groups already stand out and are requesting a change in method. Also, generic education could use a mix of different methods, in order to get positive effects from every method. Finally, making the education mandatory would at the very least make sure that awareness is raised about phishing. As the numbers of this study shows, there are a lot of users that have not seen any of the education methods.

6.2 Future work

There are some areas that have weaknesses in this study. In particular, if one is to pursue this path further, more phishing attempts need to be made, in order to test a wider array of triggers in the users. Here, I have mostly gone for fear of losing data (login credentials) and fear of missing out and personal gain (discount that need a response). There are several other emotions that can trigger a response from a user.

This of course needs a lot of time, as users will get wary if these attempts come too close together.

Also, the education needs to be mandatory. Without mandatory education, the data sets will have a risk of not being large enough to give meaningful results.

What really should be further looked into, is whether or not the differences between the user groups is simply a matter of frequency of email use, or that management and other high value targets actually are more vulnerable for other reasons. Since these are the highest valuable targets, they are also the ones that need the most protection, and if there is a deeper reason for them to be vulnerable, that reason should be discovered and used to help that group be less vulnerable.

Having a study look into the password health of users can be an interesting thought. In my profession as an IT administrator, we often come across (very) insecure passwords, almost to the point of ridiculousness. Some general studies have been done in this field, but I have yet to see how this might look in local government or other wide ranged user environments. Both investigating the general strength of passwords or authentication methods and the reuse of these passwords across systems could be interesting.

As I have shown here, there are still many users that do bite on phish. From my experience, there are also a number of users that do bite on the more obvious, at least for a trained eye, phish. This means there is some damage done. Trying to determine the amount of damage or what the cost of such damage translates into, could be an interesting angle for future work.

Bibliography

[1] S. B. C. VG. (2020). 'Over 10.000 kommunalt ansatte utsatt for alvorlig e-post-angrep,' [Online]. Available: https://www.vg.no/nyheter/innenriks/i/8m3Ayr/over-10000-kommunalt-ansatte-utsatt-for-alvorlig-e-post-angrep (visited on 30/05/2021).

[2] Statista.com. (2020). 'Global spam volume as percentage of total e-mail traffic from january 2014 to september 2020, by month,' [Online]. Available: https://www.statista.com/statistics/420391/spam-email-traffic-share/ (visited on 01/06/2021).

[3] Builtwith.com. (2021). 'Dkim usage statistics,' [Online]. Available: https://trends.builtwith.com/mx/DKIM (visited on 01/06/2021).

[4] Builtwith.com. (2021). 'Dmarc usage statistics,' [Online]. Available: https://trends.builtwith.com/mx/DMARC (visited on 01/06/2021).

[5] T. Nagunwa, 'Complementing blacklists: An enhanced technique to learn detection of zero-hour phishing urls,' eng, *International journal of cybersecurity and digital forensics*, vol. 4, no. 4, pp. 508–520, 2015, ISSN: 2305-0012.

[6] P. Kumaraguru, J. Cranshaw, A. Acquisti, L. Cranor, J. Hong, M. Blair and T. Pham, 'School of phish: A real-world evaluation of anti-phishing training,' eng, ser. SOUPS '09, ACM, 2009, pp. 1–12, ISBN: 1605587362.

[7] P. H. N. Rajput, *Phish muzzle: This fish won't bite*, eng, 2017.

[8] S. Kumar, A. Faizan, A. Viinikainen and T. Hamalainen, 'Mlspd - machine learning based spam and phishing detection,' eng, in *Computational Data and Social Networks*, ser. Lecture Notes in Computer Science, vol. 11280, Cham: Springer International Publishing, 2018, pp. 510–522, ISBN: 9783030046477.

[9] D. M. Sarno, J. E. Lewis, C. J. Bohil and M. B. Neider, 'Which phish is on the hook? phishing vulnerability for older versus younger adults,' eng, *Human factors*, vol. 62, no. 5, pp. 704–717, 2020, ISSN: 0018-7208.

[10] S. Stockhardt, B. Reinheimer, M. Volkamer, P. Mayer, A. Kunz, P. Rack and D. Lehmann, 'Teaching phishing-security: Which way is best?' eng, ser. IFIP Advances in Information and Communication Technology, vol. 471, Cham: Springer International Publishing, 2016, pp. 135–149, ISBN: 3319336290.

[11] P. Kumaraguru, S. Sheng, A. Acquisti, L. Cranor and J. Hong, 'Teaching johnny not to fall for phish,' eng, *ACM Transactions on Internet Technology (TOIT)*, vol. 10, no. 2, pp. 1–31, 2010, ISSN: 1533-5399.

[12] M. Scholl, 'Awareness in information security,' eng, *Journal of systemics, cybernetics and informatics*, vol. 16, no. 4, pp. 80–89, 2018, ISSN: 1690-4524.

[13] M. Scholl, F. Fuhrmann and D. Pokoyski, 'Information security awareness 3.0 for job beginners,' Oct. 2016.

[14] E. Kolkowska, F. Karlsson and K. Hedström, 'Towards analysing the rationale of information security non-compliance: Devising a value-based compliance analysis method,' *The Journal of Strategic Information Systems*, vol. 26, no. 1, pp. 39–57, 2017, ISSN: 0963-8687. DOI: `https://doi.org/10.1016/j.jsis.2016.08.005`. [Online]. Available: `https://www.sciencedirect.com/science/article/pii/S0963868716301639`.

[15] A. L. Varden. (2021). 'Utsatt for dataangrep - nå advarer kommunen,' [Online]. Available: `https://www.varden.no/nyheter/kommunen-utsatt-for-dataangrep/` (visited on 31/05/2021).

[16] Telenor. (2021). 'Dette er de største truslene på nett i 2021,' [Online]. Available: `https://www.telenor.no/privat/artikler/internett/dette-er-de-storste-truslene-2021/` (visited on 31/05/2021).

[17] PwC. (2021). 'Fem cyber-trender du må se opp for i 2021,' [Online]. Available: `https://www.pwc.no/no/publikasjoner/fem-cyber-trender-du-ma-se-opp-for.html` (visited on 31/05/2021).

[18] CheckPoint. (2020). 'Update: Coronavirus-themed domains 50% more likely to be malicious than other domains,' [Online]. Available: `https://blog.checkpoint.com/2020/03/05/update-coronavirus-themed-domains-50-more-likely-to-be-malicious-than-other-domains/` (visited on 31/05/2021).

[19] P. K. e. a. NRK. (2020). 'Stortinget utsatt for et omfattende it-angrep,' [Online]. Available: `https://www.nrk.no/norge/stortinget-utsatt-for-et-omfattende-it-angrep-1.15143406` (visited on 31/05/2021).

[20] Stortinget. (2020). 'It-angrep mot stortinget,' [Online]. Available: `https://www.stortinget.no/no/Hva-skjer-pa-Stortinget/Nyhetsarkiv/Pressemeldingsarkiv/2019-2020/it-angrep-mot-stortinget/` (visited on 31/05/2021).

[21] M. S. Aftenposten. (2021). 'Hackerangrep mot søkere til uds aspirantkurs. kan misbrukes av fremmede makter, mener professor.,' [Online]. Available: `https://www.aftenposten.no/norge/i/41mX3e/hackerangrep-mot-soekere-til-uds-aspirantkurs-kan-misbrukes-av-fremmed` (visited on 31/05/2021).

[22] E. K. Digi. (2021). 'Microsoft advarer: Solarwinds-hackerne er på ferde igjen med ny kampanje,' [Online]. Available: `https://www.digi.no/artikler/microsoft-advarer-solarwinds-hackerne-er-pa-ferde-igjen-med-ny-kampanje/510561` (visited on 01/06/2021).

[23] D. Pienta, J. B. Thatcher and A. Johnston, 'Protecting a whale in a sea of phish,' eng, *Journal of information technology*, vol. 35, no. 3, pp. 214–231, 2020, ISSN: 0268-3962.

[24] M. Carr, 'How social engineering fueled the cyber-attack business,' eng, *Property & Casualty 360*, 2017, ISSN: 2331-5326.

[25] S. Sheng, M. Holbrook, P. Kumaraguru, L. Cranor and J. Downs, 'Who falls for phish?: A demographic analysis of phishing susceptibility and effectiveness of interventions,' eng, in *Proceedings of the SIGCHI Conference on human factors in computing systems*, ser. CHI '10, ACM, 2010, pp. 373–382, ISBN: 1605589292.

[26] O. A. Zielinska, R. Tembe, K. W. Hong, X. Ge, E. Murphy-Hill and C. B. Mayhorn, 'One phish, two phish, how to avoid the internet phish: Analysis of training strategies to detect phishing emails,' eng, 1, vol. 58, Los Angeles, CA: SAGE Publications, 2014, pp. 1466–1470.

[27] R. Dhamija, J. Tygar and M. Hearst, 'Why phishing works,' eng, in *Proceedings of the SIGCHI Conference on human factors in computing systems*, ser. CHI '06, ACM, 2006, pp. 581–590, ISBN: 9781595933720.

[28] B. B. Bilet, 'Security awareness training as a countermeasure to phishing attacks,' M.S. thesis, Norwegian University of Science and Technology, Gjøvik, Norway, Dec. 2020.

[29] C. B. Mayhorn and P. G. Nyeste, 'Training users to counteract phishing,' eng, *Work (Reading, Mass.)*, vol. 41 Suppl 1, pp. 3549–3552, 2012, ISSN: 1051-9815.

[30] J. R. Anderson, L. M. Reder and H. A. Simon, 'Situated learning and education,' eng, *Educational researcher*, vol. 25, no. 4, pp. 5–11, 2016, ISSN: 1935-102X.

[31] S. A. Mathan and K. R. Koedinger, 'Fostering the intelligent novice: Learning from errors with metacognitive tutoring,' eng, *Educational psychologist*, vol. 40, no. 4, pp. 257–265, 2005, ISSN: 0046-1520.

[32] M. Alsharnouby, F. Alaca and S. Chiasson, 'Why phishing still works: User strategies for combating phishing attacks,' eng, *International journal of human-computer studies*, vol. 82, pp. 69–82, 2015, ISSN: 1071-5819.

[33] M. Steves, K. Greene and M. Theofanos, 'Categorizing human phishing difficulty: A phish scale,' eng, *Journal of cybersecurity (Oxford)*, vol. 6, no. 1, 2020, ISSN: 2057-2085.

Appendix A

Database

The database used is a simple MySQL database with a few tables that hold all information. The initialization script Code listing A.1 is almost self-explanatory. The table phish is the main table, the table extra is used to pick up entries where the data entered deviates from what is given in the links (for example if a user borrows another user's link to enter their own information), the table scoreboard is the scoreboard for the game education, while the table hitspercategory was used simply as a counter during the first phish. The last table was unnecessary, because the same information could be read easily from the other tables.

Code listing A.1: Database initialization script

```
USE phish;
--DROP TABLE phish;
--DROP TABLE extra;
--DROP TABLE hitspercategory;
--DROP TABLE scoreboard;

CREATE TABLE phish (username VARCHAR(255) NOT NULL PRIMARY KEY,
                    training INTEGER,
                    worker_group INTEGER,
                    first_mail_sent TIMESTAMP,
                    first_bite TIMESTAMP,
                    first_info_entered TIMESTAMP,
                    first_pass_length INTEGER,
                    first_client_ip VARCHAR(15),
                    first_numhits INTEGER,
                    first_numenters INTEGER,
                    training_received TIMESTAMP,
                    second_mail_sent TIMESTAMP,
                    second_bite TIMESTAMP,
                    second_info_entered TIMESTAMP,
                    second_pass_length INTEGER,
                    second_client_ip VARCHAR(15),
                    second_numhits INTEGER,
                    second_numenters INTEGER
                    );

CREATE TABLE extra (username VARCHAR(255) NOT NULL PRIMARY KEY,
                    orgusername VARCHAR(255),
                    training INTEGER DEFAULT 99,
                    worker_group INTEGER DEFAULT 99,
                    first_info_entered TIMESTAMP DEFAULT CURRENT_TIMESTAMP,
                    first_pass_length INTEGER,
                    first_client_ip VARCHAR(15),
                    first_numenters INTEGER,
                    training_received TIMESTAMP,
                    second_mail_sent TIMESTAMP,
                    second_bite TIMESTAMP,
                    second_info_entered TIMESTAMP,
                    second_pass_length INTEGER,
                    second_client_ip VARCHAR(15),
                    second_numhits INTEGER,
                    second_numenters INTEGER
                    );

CREATE TABLE hitspercategory (category INTEGER NOT NULL PRIMARY KEY,
                              hits INTEGER
                              );

CREATE TABLE scoreboard (username VARCHAR(255) NOT NULL PRIMARY KEY,
                         score INTEGER
                         );
```

Appendix B

Script for sending emails

The first phish for Alta recipients was sent using the php script in Code listing B.1. All emails have been sent using alterations of this script.

Code listing B.1: PHP script sending first phish to Alta recipients

```
<?php
require_once "Mail.php";
require_once "Mail/mime.php";


// Variables for connecting to the database.
$sqlserver = "localhost";
$sqldb = "phish";
$sqluser = "*****************";
$sqlpass = "*****************";

//Initializing the connection
try {
  $sqlConn = new PDO("mysql:host=$sqlserver;dbname=$sqldb", $sqluser, $sqlpass);
  $sqlConn->setAttribute(PDO::ATTR_ERRMODE, PDO::ERRMODE_EXCEPTION);

  $sql = 'SELECT␣username,worker_group␣FROM␣phish;';
  $q = $sqlConn->query($sql);
  $q->setFetchMode(PDO::FETCH_ASSOC);

}
catch (PDOException $e)
{
  echo "<p>Connection␣failed:␣" . $e->getMessage();
}

$recipients = $q->fetchAll(PDO::FETCH_ASSOC);
while ($row = $q->fetch()) {
  $recipients[] = $row;
}

$from    = 'Gorm␣<gorm@alta.kommmune.no>';
$subject = 'Stjålne␣passord';
$date    = date('r', time());

$crlf = "\r\n";
$mimeparams['text_encoding']="8bit";
```

```
$mimeparams[’text_charset’]="UTF-8";
$mimeparams[’html_charset’]="UTF-8";
$mimeparams[’head_charset’]="UTF-8";

$params   = array(
  ’host’=>’hkmvask.hammerfest.kommune.no’,
  ’auth’=>FALSE,
  ’username’=>’gorm@alta.kommmune.no’,
  ’localhost’ => ’hammerfest.kommmune.no’,
  ’persist’=>TRUE
);

$headers = array(
  ’From’         => $from,
  ’Return-Path’  => $from,
  ’Subject’      => $subject,
  ’Date’         => $date,
  ’Content-Type’ => ’text/html;␣charset="utf-8"’,
);



$mailfactory = Mail::factory(’smtp’,$params);

$sendcount=0;

foreach ($recipients as $recipient) {
  $sendcount++;
  $headers[’To’] = $recipient[’username’];

  if ($recipient[’worker_group’] > 20) {
    $bodyhtml = ’<html><body>’;
    $bodyhtml .= ’Hallo!<br>’;
    $bodyhtml .= ’<br>’;
    $bodyhtml .= ’Vi␣har␣hatt␣flere␣tilfeller␣av␣stj&aring;lne␣passord␣i␣det␣siste.’;
    $bodyhtml .= ’␣␣For␣&aring;␣finne␣omfanget␣av␣dette␣problemet,␣’;
    $bodyhtml .= "<a␣href=’https://alta.kommmune.no/passjekk.php?oun=";
    $bodyhtml .= $recipient[’username’] . "’>logg␣p&aring;␣her</a>␣for␣&aring;␣se␣";
    $bodyhtml .= "om␣du␣er␣p&aring;virket.<br>";
    $bodyhtml .= ’<br>’;
    $bodyhtml .= ’Hilsen<br>’;
    $bodyhtml .= ’Gorm,<br>’;
    $bodyhtml .= ’IT-avd.<br>’;
    $bodyhtml .= ’</body></html>’;
    $bodytxt = ’Hallo!\n’;
    $bodytxt .= ’\n’;
    $bodytxt .= ’Vi␣har␣hatt␣flere␣tilfeller␣av␣stjålne␣passord␣i␣det␣siste.␣␣’;
    $bodytxt .= ’For␣å␣finne␣omfanget␣av␣dette␣problemet,␣’;
    $bodytxt .= "<a␣href=’https://alta.kommmune.no/passjekk.php?oun=";
    $bodytxt .= $recipient[’username’] . "’>logg␣på␣her</a>␣for␣å␣se␣om␣du␣er␣";
    $bodytxt .= ’påvirket.\n’;
    $bodytxt .= ’\n’;
    $bodytxt .= ’Hilsen\n’;
    $bodytxt .= ’Gorm,\n’;
    $bodytxt .= ’IT-avd.\n’;
  }

  $mime = new Mail_mime($crlf);
  $mime->setTXTBody($bodytxt);
  $mime->setHTMLBody($bodyhtml);
  $body = $mime->get($mimeparams);
  $headers = $mime->headers($headers);

  $mail = $mailfactory->send($recipient['username'],$headers,$body);

  if    (PEAR::isError($mail)) {
        echo date("␣Y-m-d␣H:i:s␣␣->␣␣");
        echo "email␣to␣{$recipient['username']}␣failed,␣details␣:␣";
        echo $mail->getMessage() . "\n\n";
  }
  else  {
    echo date("␣Y-m-d␣H:i:s␣␣->␣␣");
    echo "email␣sent␣to␣{$recipient['username']}␣!\n\n";
    $sql = "UPDATE␣phish␣SET␣first_mail_sent=CURRENT_TIMESTAMP␣WHERE␣username=?";
    $sqlConn->prepare($sql)->execute([$recipient['username']]);
  }

  if ($sendcount%100 == 0) sleep(10); // This was too fast for Alta.
                                      // Needed to sleep(15) for each Alta-mail.
}
?>
```

Appendix C

First Phish

The first phish was designed to mimic how the IT department can send emails, and used different links and names for each municipality. Figure C.1 shows the email sent to Hammerfest users, Figure C.2 shows the email sent to Alta users.

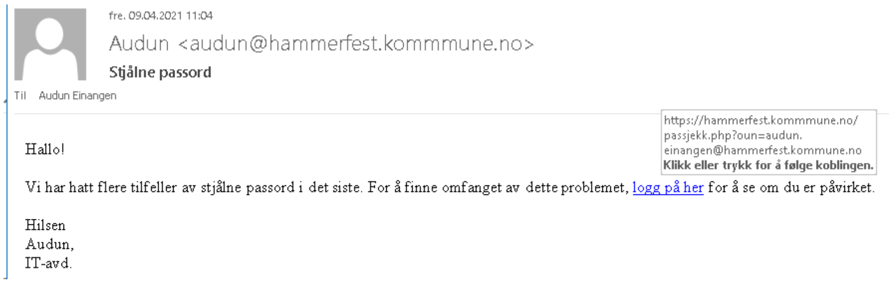

fre. 09.04.2021 11:04

Audun <audun@hammerfest.kommmune.no>

Stjålne passord

Til Audun Einangen

https://hammerfest.kommmune.no/passjekk.php?oun=audun.einangen@hammerfest.kommune.no
Klikk eller trykk for å følge koblingen.

Hallo!

Vi har hatt flere tilfeller av stjålne passord i det siste. For å finne omfanget av dette problemet, logg på her for å se om du er påvirket.

Hilsen
Audun,
IT-avd.

Figure C.1: First Phish email for Hammerfest

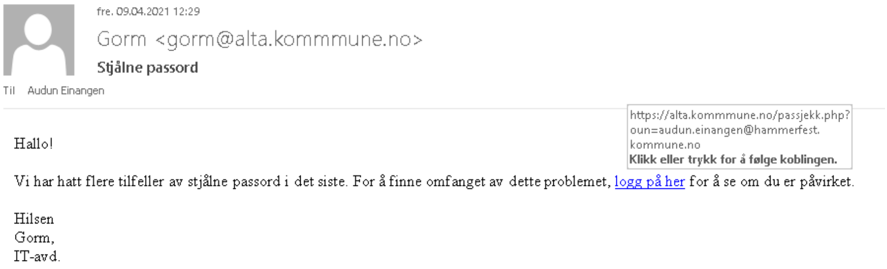

fre. 09.04.2021 12:29

Gorm <gorm@alta.kommmune.no>

Stjålne passord

Til Audun Einangen

https://alta.kommmune.no/passjekk.php?oun=audun.einangen@hammerfest.kommune.no
Klikk eller trykk for å følge koblingen.

Hallo!

Vi har hatt flere tilfeller av stjålne passord i det siste. For å finne omfanget av dette problemet, logg på her for å se om du er påvirket.

Hilsen
Gorm,
IT-avd.

Figure C.2: First Phish email for Alta

The link sent the users to different sites according to where the user belongs. Hammerfest users saw the site in Figure C.3. Alta users saw the site in Figure C.4. The username field was prepopulated by the link.

Unless the user was selected for the immediate feedback education, the user was presented with one of two responses, determined by whether the password entered was shorter than nine characters (Figure C.5), or ten or more characters

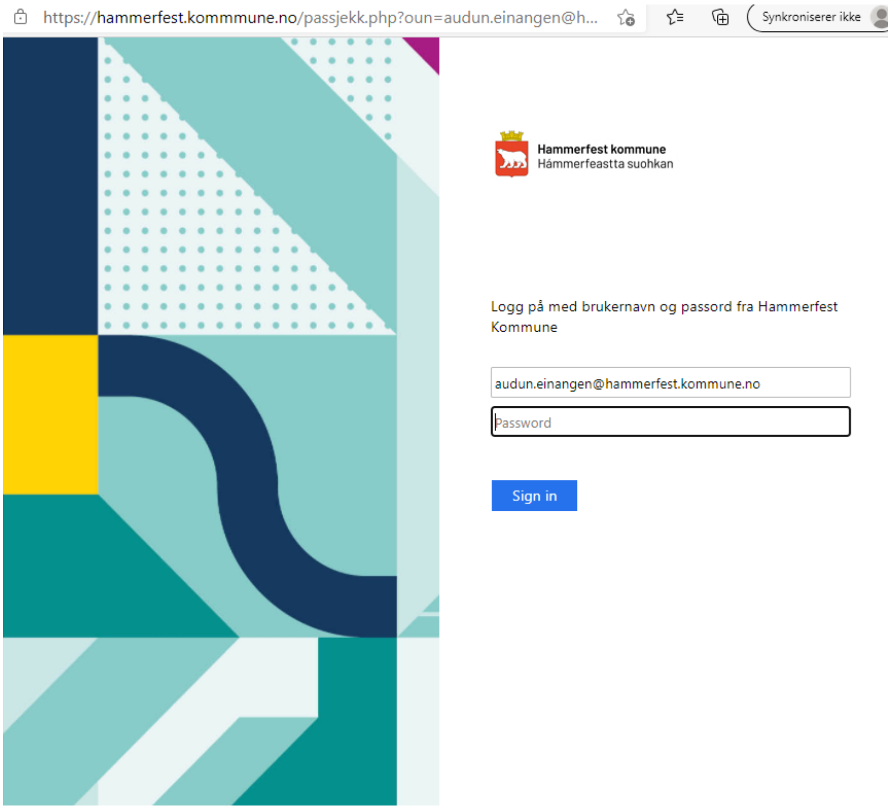

https://hammerfest.kommmune.no/passjekk.php?oun=audun.einangen@h...

Hammerfest kommune
Hámmerfeastta suohkan

Logg på med brukernavn og passord fra Hammerfest Kommune

audun.einangen@hammerfest.kommune.no

Password

Figure C.3: First phish site for Hammerfest

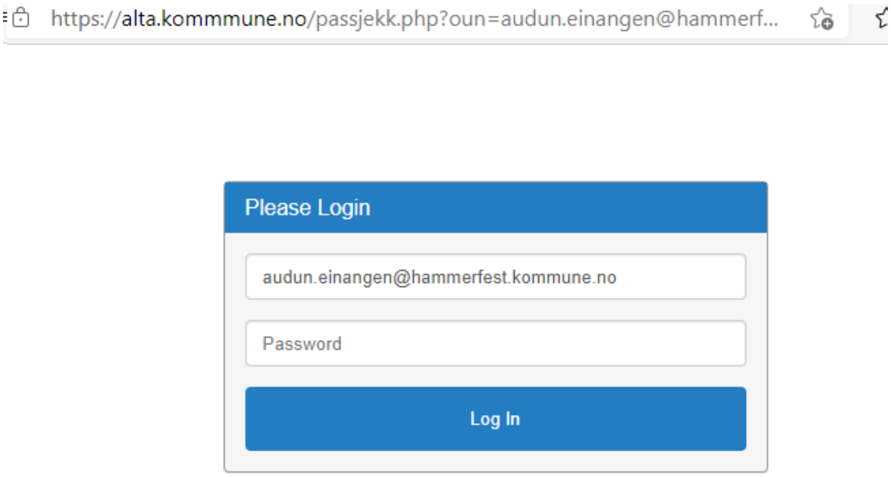

https://alta.kommmune.no/passjekk.php?oun=audun.einangen@hammerf...

Please Login

audun.einangen@hammerfest.kommune.no

Password

Figure C.4: First phish site for Alta

long (Figure C.6).

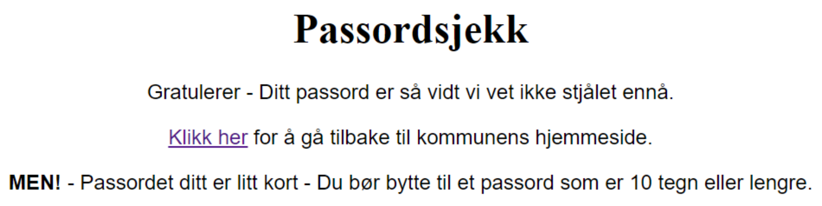

Passordsjekk

Gratulerer - Ditt passord er så vidt vi vet ikke stjålet ennå.

Klikk her for å gå tilbake til kommunens hjemmeside.

MEN! - Passordet ditt er litt kort - Du bør bytte til et passord som er 10 tegn eller lengre.

Figure C.5: Information entered response when password is short

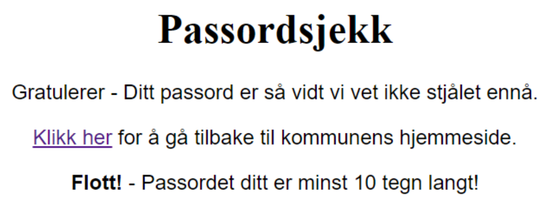

Passordsjekk

Gratulerer - Ditt passord er så vidt vi vet ikke stjålet ennå.

Klikk her for å gå tilbake til kommunens hjemmeside.

Flott! - Passordet ditt er minst 10 tegn langt!

Figure C.6: Information entered response when password is long

Appendix D

Assigning Education

Some users received their training immediately - Every fifth user that entered information in each worker group was assigned the immediate feedback education method. The mechanism was similar to what is shown in Code listing D.1, which shows how the rest of the education was assigned. Because the immediate feedback education was numbered "1" in the database, the script reassigns the number "1" to "4" in the script, that way I populate the training field equally with the numbers "0","2","3" and "4", while keeping the users with this field set to "1" untouched.

Code listing D.1: PHP script assigning education methods

```
<?php

// Variables for connecting to the database.
$sqlserver = "localhost";
$sqldb = "phish";
$sqluser = "phish";
$sqlpass = "MISEB";

//Initializing the connection
try {
  $sqlConn = new PDO("mysql:host=$sqlserver;dbname=$sqldb", $sqluser, $sqlpass);
  $sqlConn->setAttribute(PDO::ATTR_ERRMODE, PDO::ERRMODE_EXCEPTION);
}
catch (PDOException $e)
{
  echo "<p>Connection␣failed:␣" . $e->getMessage();
}

$sql_fetchallusers = $sqlConn->prepare("SELECT␣*␣FROM␣phish;");
$sql_fetchallusers->execute();

$allusers = $sql_fetchallusers->fetchAll(PDO::FETCH_ASSOC);

$counter[1]=0;
$counter[2]=0;
$counter[3]=0;
$counter[4]=0;
$counter[5]=0;
$counter[6]=0;
```

```
$counter[7]=0;
$counter[8]=0;
$counter[9]=0;
$counter[10]=0;
$counter[11]=0;
$counter[51]=0;
$counter[52]=0;
$counter[53]=0;
$counter[54]=0;
$counter[55]=0;
$counter[56]=0;
$counter[57]=0;
$counter[58]=0;
$counter[59]=0;
$counter[60]=0;
$counter[61]=0;


foreach ($allusers as $user) {
    $category = (int)$user[’worker_group’];
    if ((int)$user["training"] != 1) {
        $training = $counter[$category] % 4;
        if ($training == 1) $training = 4;
        $counter[$category]++;
        $sql_updatetraining = "UPDATE␣phish␣SET␣training␣=␣" . $training;
        $sql_updatetraining .= "␣WHERE␣username␣=␣’" . $user[’username’] . "’";
        $updatequery = $sqlConn->prepare($sql_updatetraining);
        $updatequery->execute();
        echo $sql_updatetraining . "\n";
    }
    else {
        echo $user[’username’] . "␣har␣allerede␣training␣";
        echo $user[’training’] . "...\n";
    }
}
?>
```

Appendix E

Education

Based on the field training in the database, each user got sent an email, either with text information, shown in Figure E.1, or links to the web education site (email shown in Figure E.2), video education site or the game education site. The emails for the last two were very similar to the one for web education.
The web education site is shown in Figure E.3, Figure E.4, Figure E.5 and Figure E.6. The immediate feedback education showed Figure E.7 before linking to the web education page. The game education site front page is shown in Figure E.8. When the user clicks on start, he/she is given questions like in Figure E.9. The video page is simply a video where I show the images in the web education and explain the same information as is written in the web education verbally.

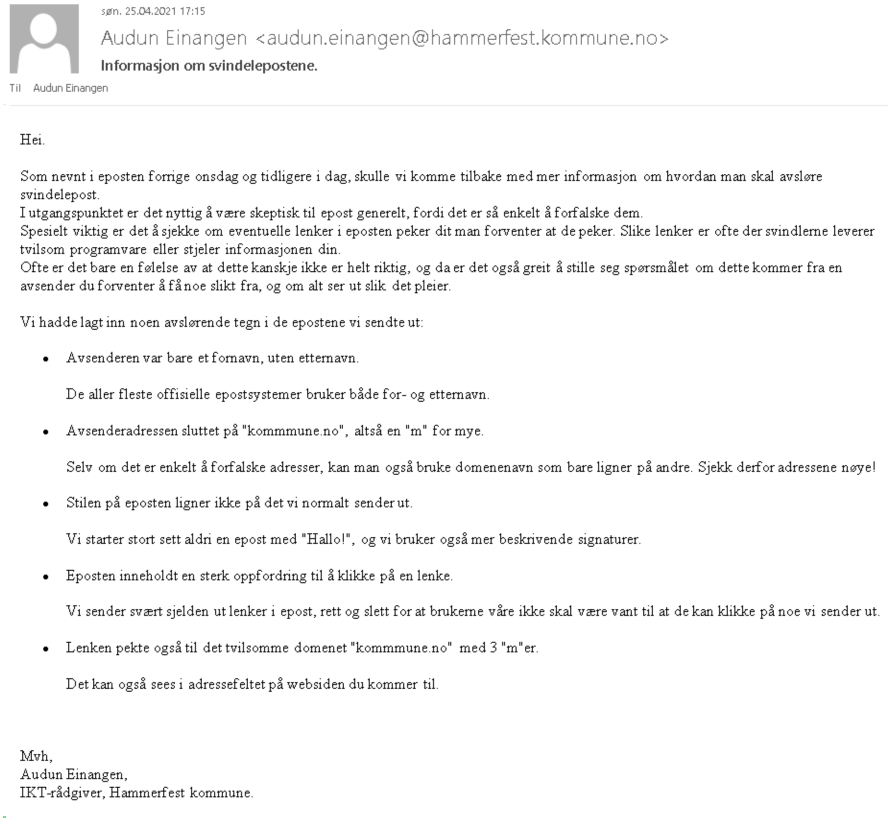

søn. 25.04.2021 17:15

Audun Einangen <audun.einangen@hammerfest.kommune.no>

Informasjon om svindelepostene.

Til Audun Einangen

Hei.

Som nevnt i eposten forrige onsdag og tidligere i dag, skulle vi komme tilbake med mer informasjon om hvordan man skal avsløre svindelepost.
I utgangspunktet er det nyttig å være skeptisk til epost generelt, fordi det er så enkelt å forfalske dem.
Spesielt viktig er det å sjekke om eventuelle lenker i eposten peker dit man forventer at de peker. Slike lenker er ofte der svindlerne leverer tvilsom programvare eller stjeler informasjonen din.
Ofte er det bare en følelse av at dette kanskje ikke er helt riktig, og da er det også greit å stille seg spørsmålet om dette kommer fra en avsender du forventer å få noe slikt fra, og om alt ser ut slik det pleier.

Vi hadde lagt inn noen avslørende tegn i de epostene vi sendte ut:

- Avsenderen var bare et fornavn, uten etternavn.

  De aller fleste offisielle epostsystemer bruker både for- og etternavn.

- Avsenderadressen sluttet på "kommmune.no", altså en "m" for mye.

  Selv om det er enkelt å forfalske adresser, kan man også bruke domenenavn som bare ligner på andre. Sjekk derfor adressene nøye!

- Stilen på eposten ligner ikke på det vi normalt sender ut.

  Vi starter stort sett aldri en epost med "Hallo!", og vi bruker også mer beskrivende signaturer.

- Eposten inneholdt en sterk oppfordring til å klikke på en lenke.

  Vi sender svært sjelden ut lenker i epost, rett og slett for at brukerne våre ikke skal være vant til at de kan klikke på noe vi sender ut.

- Lenken pekte også til det tvilsomme domenet "kommmune.no" med 3 "m"er.

  Det kan også sees i adressefeltet på websiden du kommer til.

Mvh,
Audun Einangen,
IKT-rådgiver, Hammerfest kommune.

Figure E.1: Text education email

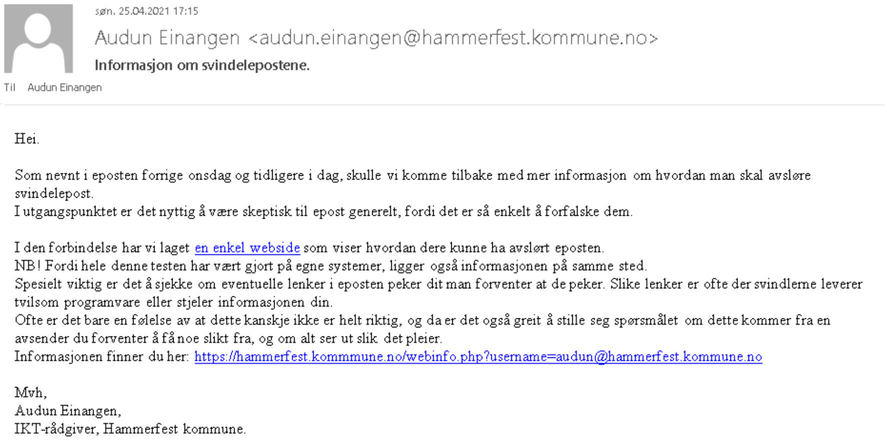

søn. 25.04.2021 17:15

Audun Einangen <audun.einangen@hammerfest.kommune.no>

Informasjon om svindelepostene.

Til Audun Einangen

Hei.

Som nevnt i eposten forrige onsdag og tidligere i dag, skulle vi komme tilbake med mer informasjon om hvordan man skal avsløre svindelepost.
I utgangspunktet er det nyttig å være skeptisk til epost generelt, fordi det er så enkelt å forfalske dem.

I den forbindelse har vi laget en enkel webside som viser hvordan dere kunne ha avslørt eposten.
NB! Fordi hele denne testen har vært gjort på egne systemer, ligger også informasjonen på samme sted.
Spesielt viktig er det å sjekke om eventuelle lenker i eposten peker dit man forventer at de peker. Slike lenker er ofte der svindlerne leverer tvilsom programvare eller stjeler informasjonen din.
Ofte er det bare en følelse av at dette kanskje ikke er helt riktig, og da er det også greit å stille seg spørsmålet om dette kommer fra en avsender du forventer å få noe slikt fra, og om alt ser ut slik det pleier.
Informasjonen finner du her: https://hammerfest.kommmune.no/webinfo.php?username=audun@hammerfest.kommune.no

Mvh,
Audun Einangen,
IKT-rådgiver, Hammerfest kommune.

Figure E.2: Web education email

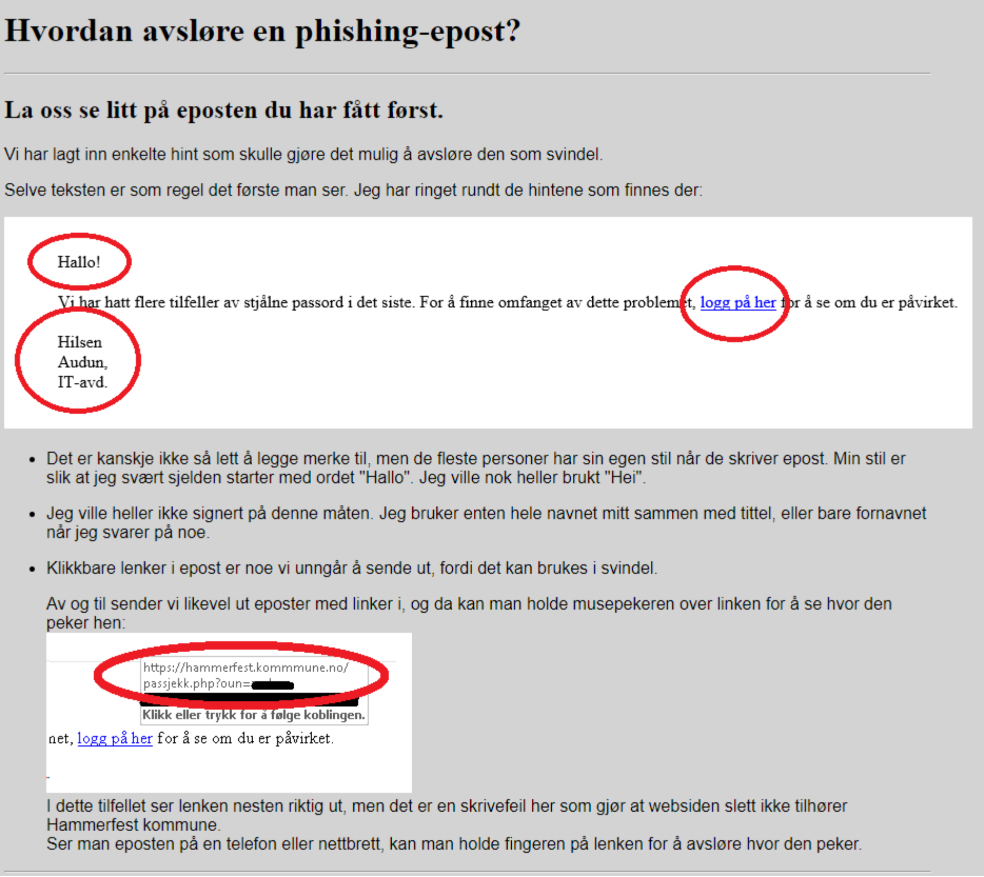

Hvordan avsløre en phishing-epost?

La oss se litt på eposten du har fått først.

Vi har lagt inn enkelte hint som skulle gjøre det mulig å avsløre den som svindel.

Selve teksten er som regel det første man ser. Jeg har ringet rundt de hintene som finnes der:

Hallo!

Vi har hatt flere tilfeller av stjålne passord i det siste. For å finne omfanget av dette problemet, logg på her for å se om du er påvirket.

Hilsen
Audun,
IT-avd.

- Det er kanskje ikke så lett å legge merke til, men de fleste personer har sin egen stil når de skriver epost. Min stil er slik at jeg svært sjelden starter med ordet "Hallo". Jeg ville nok heller brukt "Hei".
- Jeg ville heller ikke signert på denne måten. Jeg bruker enten hele navnet mitt sammen med tittel, eller bare fornavnet når jeg svarer på noe.
- Klikkbare lenker i epost er noe vi unngår å sende ut, fordi det kan brukes i svindel.

  Av og til sender vi likevel ut eposter med linker i, og da kan man holde musepekeren over linken for å se hvor den peker hen:

  https://hammerfest.kommmune.no/passjekk.php?oun=[redacted]
  Klikk eller trykk for å følge koblingen.
  net, logg på her for å se om du er påvirket.

  I dette tilfellet ser lenken nesten riktig ut, men det er en skrivefeil her som gjør at websiden slett ikke tilhører Hammerfest kommune.
  Ser man eposten på en telefon eller nettbrett, kan man holde fingeren på lenken for å avsløre hvor den peker.

Figure E.3: Web education site, part 1

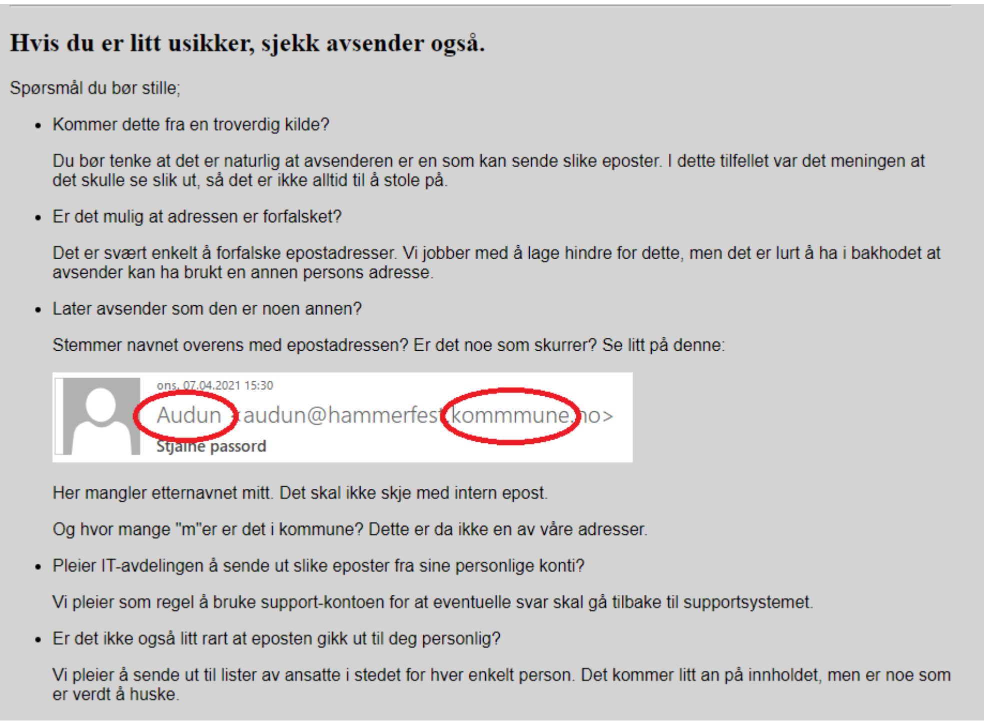

Hvis du er litt usikker, sjekk avsender også.

Spørsmål du bør stille;

- Kommer dette fra en troverdig kilde?

  Du bør tenke at det er naturlig at avsenderen er en som kan sende slike eposter. I dette tilfellet var det meningen at det skulle se slik ut, så det er ikke alltid til å stole på.

- Er det mulig at adressen er forfalsket?

  Det er svært enkelt å forfalske epostadresser. Vi jobber med å lage hindre for dette, men det er lurt å ha i bakhodet at avsender kan ha brukt en annen persons adresse.

- Later avsender som den er noen annen?

  Stemmer navnet overens med epostadressen? Er det noe som skurrer? Se litt på denne:

  ons. 07.04.2021 15:30
  Audun <audun@hammerfest.kommmune.no>
  Stjålne passord

  Her mangler etternavnet mitt. Det skal ikke skje med intern epost.

  Og hvor mange "m"er er det i kommune? Dette er da ikke en av våre adresser.

- Pleier IT-avdelingen å sende ut slike eposter fra sine personlige konti?

  Vi pleier som regel å bruke support-kontoen for at eventuelle svar skal gå tilbake til supportsystemet.

- Er det ikke også litt rart at eposten gikk ut til deg personlig?

  Vi pleier å sende ut til lister av ansatte i stedet for hver enkelt person. Det kommer litt an på innholdet, men er noe som er verdt å huske.

Figure E.4: Web education site, part 2

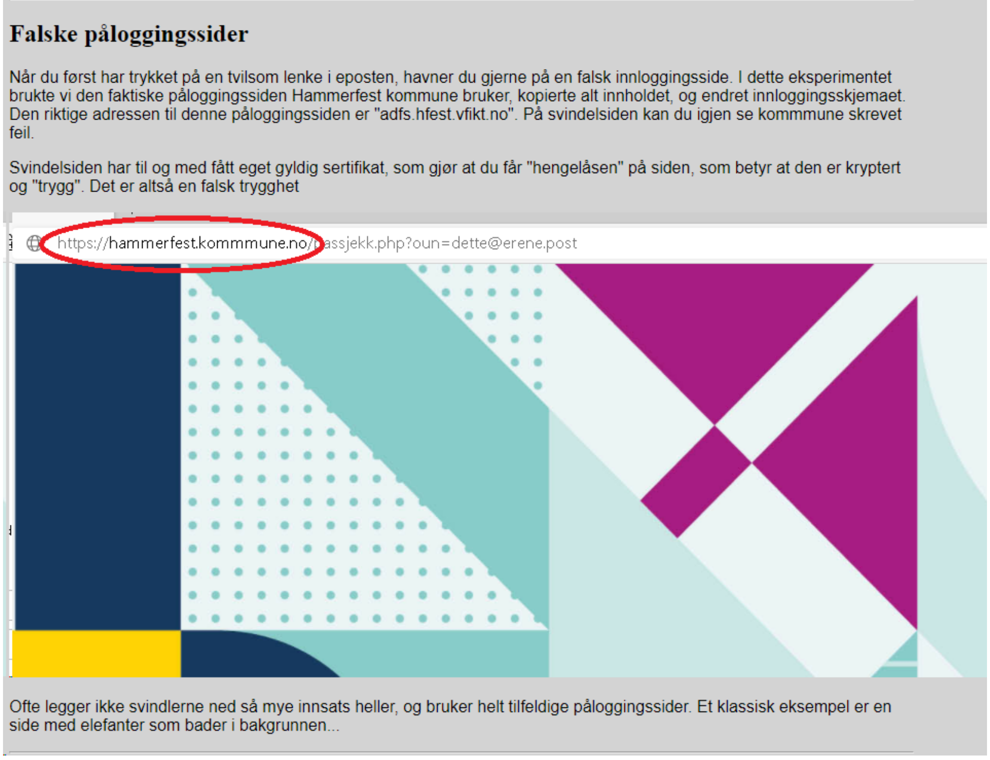

Falske påloggingssider

Når du først har trykket på en tvilsom lenke i eposten, havner du gjerne på en falsk innloggingsside. I dette eksperimentet brukte vi den faktiske påloggingssiden Hammerfest kommune bruker, kopierte alt innholdet, og endret innloggingsskjemaet. Den riktige adressen til denne påloggingssiden er "adfs.hfest.vfikt.no". På svindelsiden kan du igjen se kommmune skrevet feil.

Svindelsiden har til og med fått eget gyldig sertifikat, som gjør at du får "hengelåsen" på siden, som betyr at den er kryptert og "trygg". Det er altså en falsk trygghet

https://hammerfest.kommmune.no/passjekk.php?oun=dette@erene.post

Ofte legger ikke svindlerne ned så mye innsats heller, og bruker helt tilfeldige påloggingssider. Et klassisk eksempel er en side med elefanter som bader i bakgrunnen...

Figure E.5: Web education site, part 3

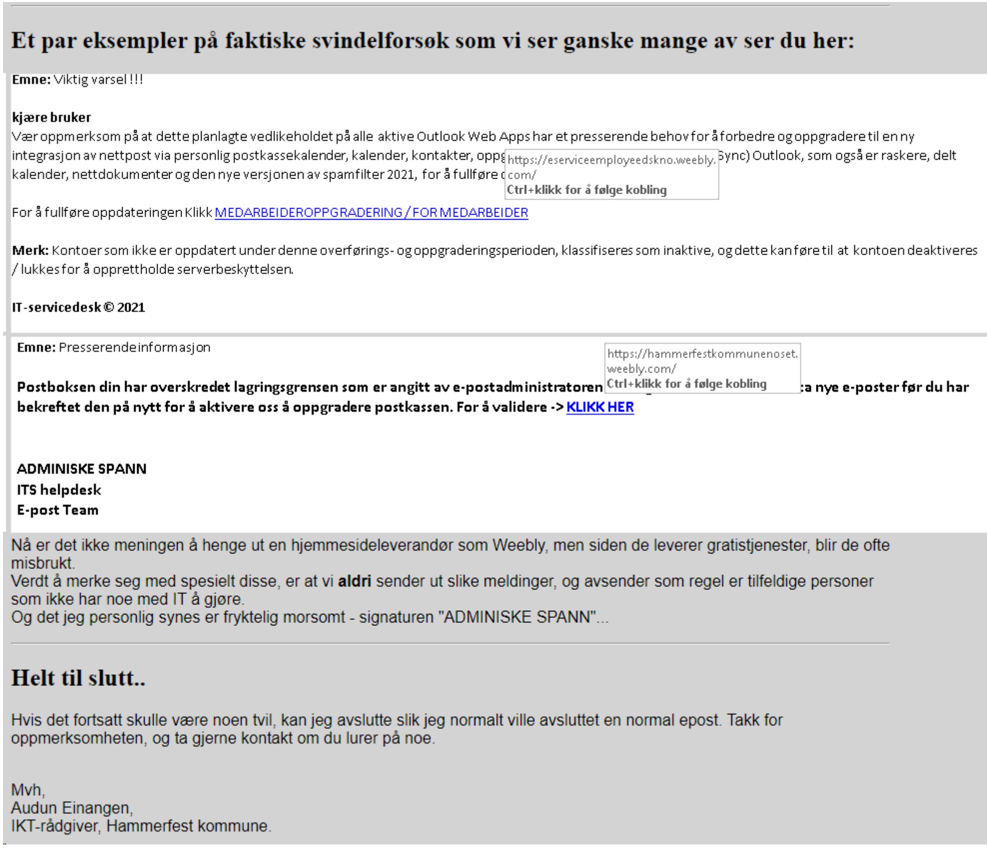

Et par eksempler på faktiske svindelforsøk som vi ser ganske mange av ser du her:

Emne: Viktig varsel !!!

kjære bruker
Vær oppmerksom på at dette planlagte vedlikeholdet på alle aktive Outlook Web Apps har et presserende behov for å forbedre og oppgradere til en ny integrasjon av nettpost via personlig postkassekalender, kalender, kontakter, oppg https://eserviceemployeedskno.weebly.com/ **Ctrl+klikk for å følge kobling** Sync) Outlook, som også er raskere, delt kalender, nettdokumenter og den nye versjonen av spamfilter 2021, for å fullføre d

For å fullføre oppdateringen Klikk MEDARBEIDEROPPGRADERING / FOR MEDARBEIDER

Merk: Kontoer som ikke er oppdatert under denne overførings- og oppgraderingsperioden, klassifiseres som inaktive, og dette kan føre til at kontoen deaktiveres / lukkes for å opprettholde serverbeskyttelsen.

IT-servicedesk © 2021

Emne: Presserende informasjon

https://hammerfestkommunenoset.weebly.com/ **Ctrl+klikk for å følge kobling**

Postboksen din har overskredet lagringsgrensen som er angitt av e-postadministratoren :a nye e-poster før du har bekreftet den på nytt for å aktivere oss å oppgradere postkassen. For å validere -> KLIKK HER

ADMINISKE SPANN
ITS helpdesk
E-post Team

Nå er det ikke meningen å henge ut en hjemmesideleverandør som Weebly, men siden de leverer gratistjenester, blir de ofte misbrukt.
Verdt å merke seg med spesielt disse, er at vi **aldri** sender ut slike meldinger, og avsender som regel er tilfeldige personer som ikke har noe med IT å gjøre.
Og det jeg personlig synes er fryktelig morsomt - signaturen "ADMINISKE SPANN"...

Helt til slutt..

Hvis det fortsatt skulle være noen tvil, kan jeg avslutte slik jeg normalt ville avsluttet en normal epost. Takk for oppmerksomheten, og ta gjerne kontakt om du lurer på noe.

Mvh,
Audun Einangen,
IKT-rådgiver, Hammerfest kommune.

Figure E.6: Web education site, part 4

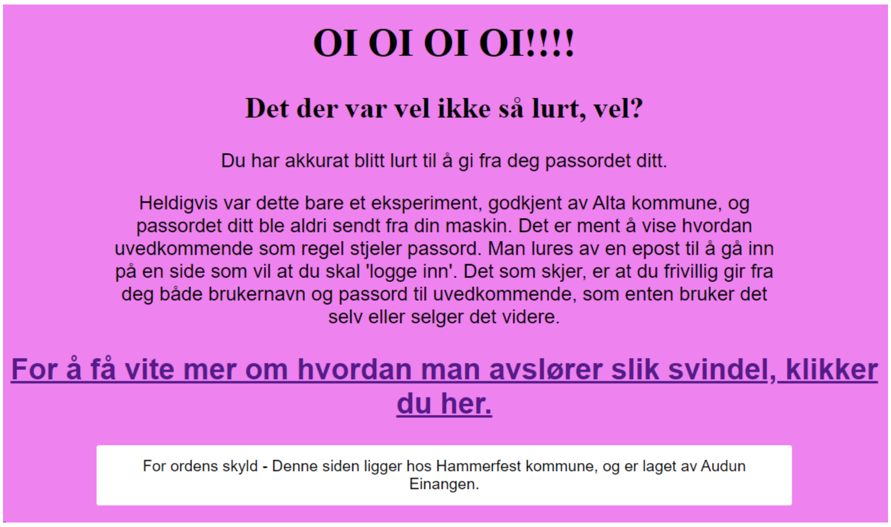

OI OI OI OI!!!!

Det der var vel ikke så lurt, vel?

Du har akkurat blitt lurt til å gi fra deg passordet ditt.

Heldigvis var dette bare et eksperiment, godkjent av Alta kommune, og passordet ditt ble aldri sendt fra din maskin. Det er ment å vise hvordan uvedkommende som regel stjeler passord. Man lures av en epost til å gå inn på en side som vil at du skal 'logge inn'. Det som skjer, er at du frivillig gir fra deg både brukernavn og passord til uvedkommende, som enten bruker det selv eller selger det videre.

For å få vite mer om hvordan man avslører slik svindel, klikker du her.

For ordens skyld - Denne siden ligger hos Hammerfest kommune, og er laget av Audun Einangen.

Figure E.7: Immediate feedback page

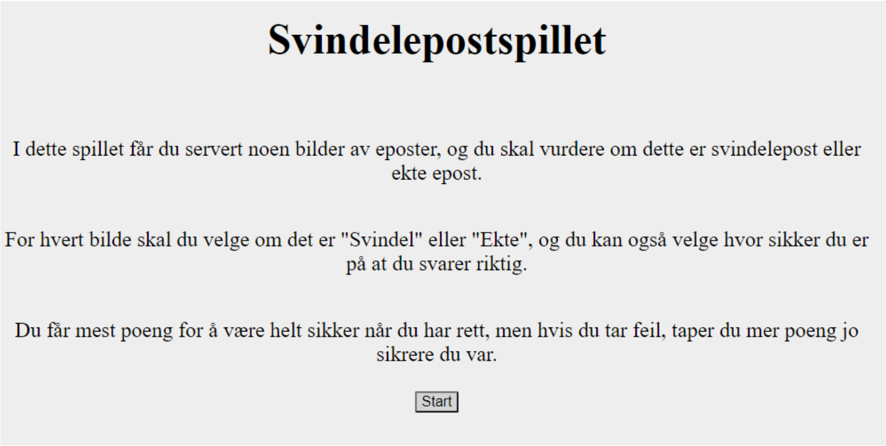

Svindelepostspillet

I dette spillet får du servert noen bilder av eposter, og du skal vurdere om dette er svindelepost eller ekte epost.

For hvert bilde skal du velge om det er "Svindel" eller "Ekte", og du kan også velge hvor sikker du er på at du svarer riktig.

Du får mest poeng for å være helt sikker når du har rett, men hvis du tar feil, taper du mer poeng jo sikrere du var.

Figure E.8: Game education start page

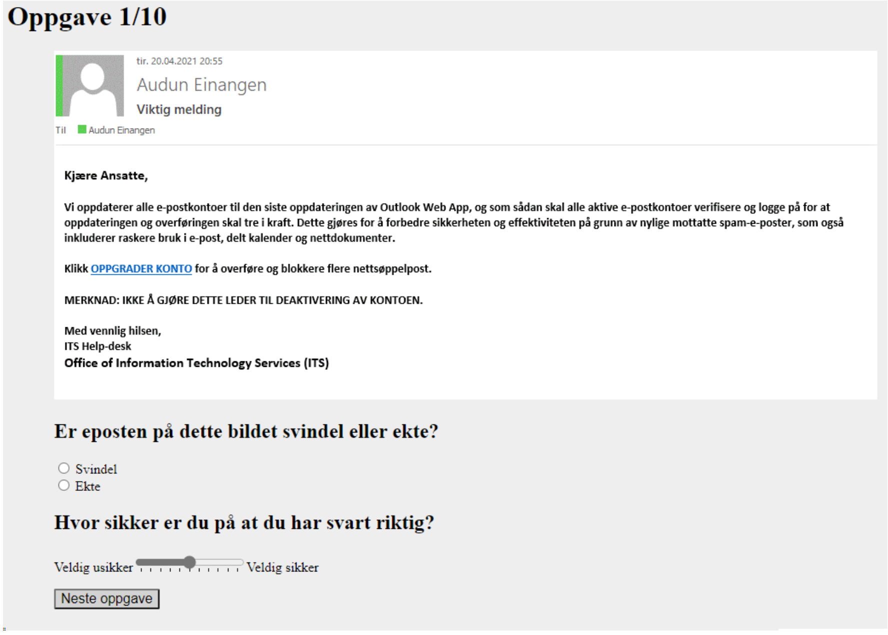

Oppgave 1/10

tir. 20.04.2021 20:55

Audun Einangen

Viktig melding

Til 🟩 Audun Einangen

Kjære Ansatte,

Vi oppdaterer alle e-postkontoer til den siste oppdateringen av Outlook Web App, og som sådan skal alle aktive e-postkontoer verifisere og logge på for at oppdateringen og overføringen skal tre i kraft. Dette gjøres for å forbedre sikkerheten og effektiviteten på grunn av nylige mottatte spam-e-poster, som også inkluderer raskere bruk i e-post, delt kalender og nettdokumenter.

Klikk OPPGRADER KONTO for å overføre og blokkere flere nettsøppelpost.

MERKNAD: IKKE Å GJØRE DETTE LEDER TIL DEAKTIVERING AV KONTOEN.

Med vennlig hilsen,
ITS Help-desk
Office of Information Technology Services (ITS)

Er eposten på dette bildet svindel eller ekte?

○ Svindel
○ Ekte

Hvor sikker er du på at du har svart riktig?

Veldig usikker — Veldig sikker

Figure E.9: Game education question page

Appendix F

Second Phish

The second phish was designed for each municipality individually, but otherwise identical.
The emails are shown in Figure F.1 and Figure F.2, while the respective sites are shown in Figure F.3 and Figure F.4.

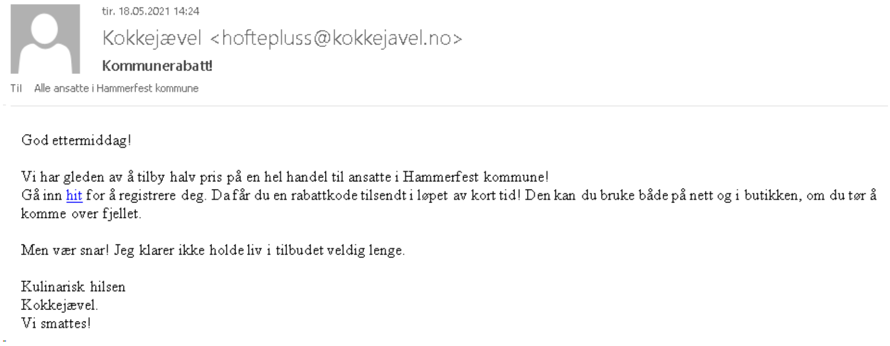

tir. 18.05.2021 14:24

Kokkejævel <hoftepluss@kokkejavel.no>

Kommunerabatt!

Til Alle ansatte i Hammerfest kommune

God ettermiddag!

Vi har gleden av å tilby halv pris på en hel handel til ansatte i Hammerfest kommune!
Gå inn hit for å registrere deg. Da får du en rabattkode tilsendt i løpet av kort tid! Den kan du bruke både på nett og i butikken, om du tør å komme over fjellet.

Men vær snar! Jeg klarer ikke holde liv i tilbudet veldig lenge.

Kulinarisk hilsen
Kokkejævel.
Vi smattes!

Figure F.1: Second phish email - Hammerfest

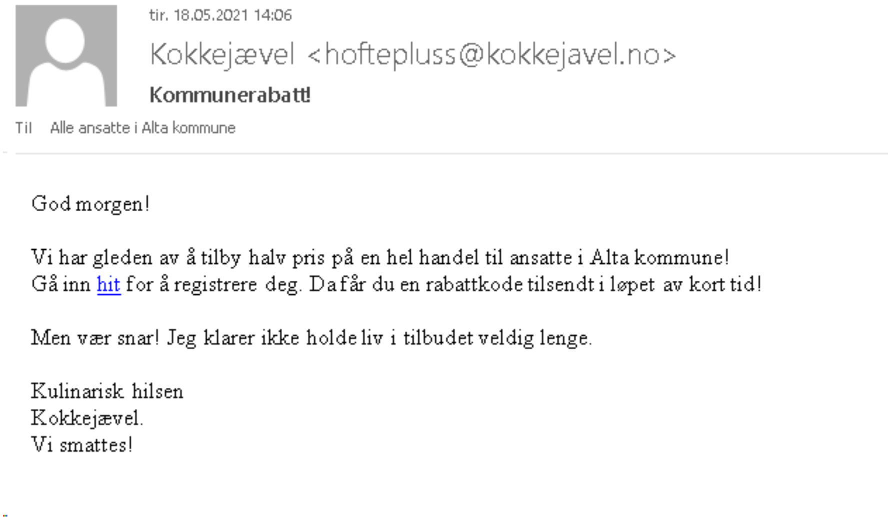

tir. 18.05.2021 14:06

Kokkejævel <hoftepluss@kokkejavel.no>

Kommunerabatt!

Til Alle ansatte i Alta kommune

God morgen!

Vi har gleden av å tilby halv pris på en hel handel til ansatte i Alta kommune!
Gå inn hit for å registrere deg. Da får du en rabattkode tilsendt i løpet av kort tid!

Men vær snar! Jeg klarer ikke holde liv i tilbudet veldig lenge.

Kulinarisk hilsen
Kokkejævel.
Vi smattes!

Figure F.2: Second phish email - Alta

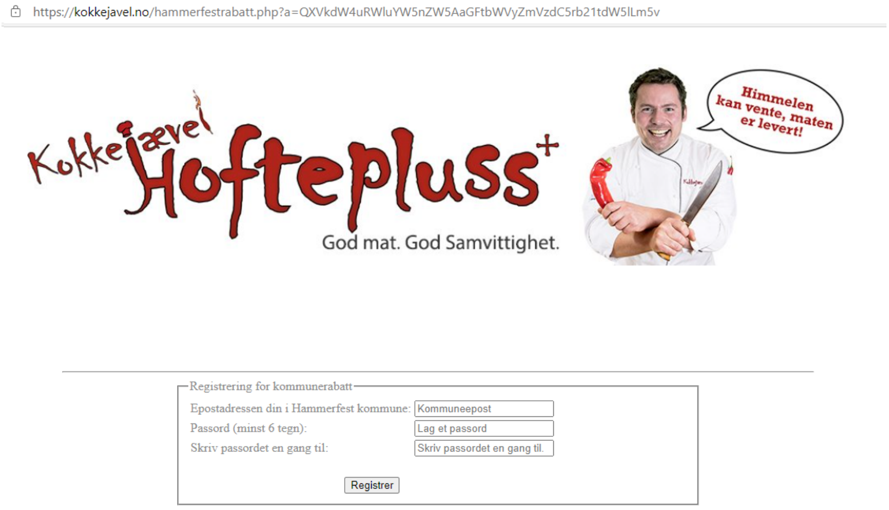

https://kokkejavel.no/hammerfestrabatt.php?a=QXVkdW4uRWluYW5nZW5AaGFtbWVyZmVzdC5rb21tdW5lLm5v

Kokkejævel
Hoftepluss+
God mat. God Samvittighet.
Himmelen kan vente, maten er levert!

Registrering for kommunerabatt

Epostadressen din i Hammerfest kommune: Kommuneepost
Passord (minst 6 tegn): Lag et passord
Skriv passordet en gang til: Skriv passordet en gang til.

Figure F.3: Second phish site - Hammerfest

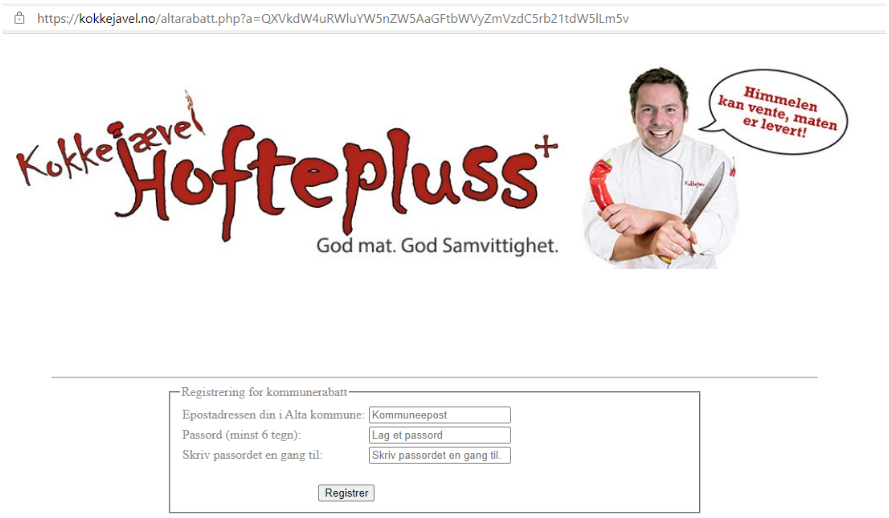

https://kokkejavel.no/altarabatt.php?a=QXVkdW4uRWluYW5nZW5AaGFtbWVyZmVzdC5rb21tdW5lLm5v

Kokkejævel
Hoftepluss+
God mat. God Samvittighet.
Himmelen kan vente, maten er levert!

Registrering for kommunerabatt

Epostadressen din i Alta kommune: Kommuneepost

Passord (minst 6 tegn): Lag et passord

Skriv passordet en gang til: Skriv passordet en gang til.

Figure F.4: Second phish site - Alta

NTNU

Norwegian University of
Science and Technology