Advanced Social Engineering Attacks[☆]

Katharina Krombholz, Heidelinde Hobel, Markus Huber, Edgar Weippl

SBA Research, Favoritenstraße 16, AT-1040 Vienna, Austria

Abstract

Social engineering has emerged as a serious threat in virtual communities and is an effective means to attack information systems. The services used by today's knowledge workers prepare the ground for sophisticated social engineering attacks. The growing trend towards *BYOD* (bring your own device) policies and the use of online communication and collaboration tools in private and business environments aggravate the problem. In globally acting companies, teams are no longer geographically co-located, but staffed just-in-time. The decrease in personal interaction combined with a plethora of tools used for communication (e-mail, IM, Skype, Dropbox, LinkedIn, Lync, etc.) create new attack vectors for social engineering attacks. Recent attacks on companies such as the New York Times and RSA have shown that targeted spear-phishing attacks are an effective, evolutionary step of social engineering attacks. Combined with zero-day-exploits, they become a dangerous weapon that is often used by advanced persistent threats. This paper provides a taxonomy of well-known social engineering attacks as well as a comprehensive overview of advanced social engineering attacks on the knowledge worker.

Keywords: security, privacy, social engineering, attack scenarios, knowledge worker, bring your own device

1. Introduction

The Internet has become the largest communication and information exchange medium. In our everyday life, communication has become distributed over a variety of online communication channels. In addition to e-mail and IM communication, Web 2.0 services such as Twitter, Facebook, and other social networking sites have become a part of our daily routine in private and business communication. Companies expect their employees to be highly mobile and flexible concerning their workspace [10] and there is an increasing trend towards expecting employees and knowledge workers to use their own devices for work, both in the office and elsewhere. This increase in flexibility and, conversely, reduction in face-to-face communication and shared office space means that increasing amounts of data need to be made available to co-workers through online channels. The development of decentralized data access and cloud services has brought about a paradigm shift in file sharing as well as communication, which today is mostly conducted over a third party, be it a social network or any other type of platform. In this world of ubiquitous communication, people freely publish information in online communication and collaboration tools, such as cloud services and social networks, with very little thought of security and privacy. They share highly sensitive documents and information in cloud services with other virtual users around the globe. Most of the time, users consider their interaction partners as trusted, even though the only identification is an e-mail address or a virtual profile. In recent years, security vulnerabilities in online communication and data sharing channels have often been misused to leak sensitive information. Such vulnerabilities can be fixed and the security of the channels can be strengthened. However, even security-enhancing methods are powerless when users are manipulated by social engineers. The term *knowledge worker* was coined by Peter Drucker more than 50 years ago and still describes the basic characteristics of a worker whose main capital is knowledge [17]. The most powerful tool an attacker can use to access this knowledge is *Social Engineering*: manipulating a person into giving information to the social engineer. It is superior to most other forms of hacking in that it can breach even the most secure systems, as the users themselves are the most vulnerable part of the system. Research has shown that social engineering is easy to automate in many cases and can therefore be performed on a large scale. Social engineering has become an emerging threat in virtual communities. Multinational corporations and news agencies have fallen victim to sophisticated targeted attacks on their information systems. Google's internal system was compromised in 2009 [2], the RSA security token system was broken in 2011 [1], Facebook was compromised in 2013 [4], as was the New York Times [40]. Many *PayPal* costumers have received phishing e-mails [45] and many have given the attackers private information such as credit card numbers. These recent attacks on high-value assets are commonly referred to as

[☆]This paper is an extended version of the conference paper [31]

Advanced Persistent Threats (APTs). APTs often rely on a common initial attack vector: social engineering such as spear-phishing and water-holing. The awareness for software security issues and privacy-enhancing methods has increased as serious incidents have been reported in the media. For example, the awareness for social engineering attacks over e-mail, which is without doubt the most frequently used communication channel on the Internet and is flooded by scammers and social engineers every day, has increased among users. However, the awareness for social engineering in cloud services and social networks is still comparatively low.

The main contributions of this article are the following:

- We discuss social engineering with regards to knowledge workers.
- We provide a taxonomy of social engineering attacks.
- We give an overview of current attack vectors for social engineering attacks.
- We discuss real-world incidents of successful social engineering attacks.

The goal of this paper is to provide a comprehensive and complete overview of social engineering attacks on the knowledge worker, to monitor the state of the art of research in this field, and to provide a comprehensive taxonomy to categorize social engineering attacks and measure their impact. Our paper significantly extends the state of the art by including novel, non-traditional attacks such as APTs. Our taxonomy extends and combines already existing work in this field, e.g., by Ivaturi et al. [28] and Foozy et al. [36]. Furthermore, our taxonomy systemizes operators, channels, types and attack vectors as well. The remainder of this paper is structured as follows: Section 2 contains a brief introduction to social engineering. In Section 3, we provide a detailed classification of social engineering attacks. In Section 4, we describe advanced social engineering attacks in online social networks, cloud services and mobile applications. Before concluding our work in Section 6, we discuss recent real-world social engineering attacks in Section 5.

2. Background

This section discusses the state of the art of social engineering and computer-supported collaborative work (*CSCW*). Attacks are divided into four different categories: physical, technical, social and socio-technical approaches.

2.1. Social Engineering (SE)

Social engineering is the art of getting users to compromise information systems. Instead of technical attacks on systems, social engineers target humans with access to information, manipulating them into divulging confidential information or even into carrying out their malicious attacks through influence and persuasion. Technical protection measures are usually ineffective against this kind of attack. In addition to that, people generally believe that they are good at detecting such attacks. Research, however, indicates that people perform poorly on detecting lies and deception [42, 33]. The infamous attacks of Kevin Mitnick [35] showed how devastating sophisticated social engineering attacks are for the information security of both companies and governmental organizations. When social engineering is discussed in the information and computer security field, it is usually by way of examples and stories (such as Mitnick's). However, at a more fundamental level, important findings have been made in social psychology on the principles of persuasion. Particularly the work of Cialdini [16], an expert in the field of persuasion, is frequently cited in contributions to social engineering research. Although Cialdini's examples focus on persuasion in marketing, the fundamental principles are crucial for anyone seeking to understand how deception works.

2.2. Types of Social Engineering Attacks

Social engineering attacks are multifaceted and include physical, social and technical aspects, which are used in different stages of the actual attack. This subsection aims to explain the different approaches attackers use.

2.2.1. Physical approaches

As the name implies, physical approaches are those where the attacker performs some form of physical action in order to gather information on a future victim. This can range from personal information (such as social security number, date of birth) to valid credentials for a computer system. An often-used method is *dumpster diving* [20], i.e., searching through an organization's trash. A dumpster can be a valuable source of information for attackers, who may find personal data about employees, manuals, memos and even print-outs of sensitive information, such as user credentials. If an attacker can gain access to a targeted organization's offices - e.g., in open-plan workspaces - they may find information such as passwords written on Post-it notes. Less sophisticated physical attacks involve *theft* or *extortion* to obtain information.

2.2.2. Social approaches

The most important aspect of successful social engineering attacks are social approaches. Hereby attackers rely on socio-psychological techniques such as Cialdini's principles of persuasion to manipulate their victims. Examples of persuasion methods include the use of (purported) authority. One common social vector that is not explicitly addressed by Cialdini is curiosity, which is, e.g., used in spear-phishing and baiting attacks. In order to increase the chances of success of such attacks, the perpetrators often try to develop a relationship with their future victims.

According to [20], the most prevalent type of social attacks is performed by phone.

2.2.3. Reverse social engineering

Instead of contacting a potential victim directly, an attacker can attempt to make them believe that he/she is a trustworthy entity. The goal is to make potential victims approach him, e.g., to ask for help. This indirect approach is known as *"reverse social engineering"* [20, 35] and consists of three major parts: sabotage, advertising and assisting [38]. The first step in this is sabotaging the company's computer system. This can range anywhere from disconnecting someone from the company's network to sophisticated manipulation of the victim's software applications. The attackers then advertise that they can fix the problem. When the victim asks for help, the social engineer will resolve the problem they created earlier while, e.g., asking the victim for their password ("so I can fix the problem") or telling them to install certain software.

2.2.4. Technical approaches

Technical attacks are mainly carried out over the Internet. Granger [20] notes that the Internet is especially interesting for social engineers to harvest passwords, as users often use the same (simple) passwords for different accounts. Most people are also not aware that they are freely providing attackers (or anyone who will search for it) with plenty of personal information. Attackers often use search engines to gather personal information about future victims. There are also tools that can gather and aggregate information from different Web resources. One of the most popular tools of this kind is Maltego[1]. Social networking sites are becoming valuable sources of information as well (see Section 4 for more details).

2.2.5. Socio-technical approaches

Successful social engineering attacks often combine several or all of the different approaches discussed above. However, socio-technical approaches have created the most powerful weapons of social engineers. One example is the so-called *baiting* attack: Attackers leave malware-infected storage media in a location where it is likely to be found by future victims. Such *"road apples"* could, e.g., be a USB drive containing a Trojan horse [48]. Attackers additionally exploit the curiosity of people by adding tempting labels to these road apples (storage media), such as "confidential" or "staff lay-off 2014". Another common combination of technical and social approaches is phishing. Phishing is usually done via e-mail or instant messaging and is aimed at a large user group in a rather indiscriminate way, similar to spam. Social engineering, in contrast, is typically directed at individuals or small groups of people. Scammers hope that by sending messages to a vast number of users, they will fool enough people to make their phishing attack profitable. Herley and Florencio [22] argue that classical phishing is not lucrative, which might explain why phishing attacks are moving towards more sophisticated "spear-phishing" attacks. Spear-phishing attacks are highly targeted messages carried out after initial data-mining. Jagatic et al. [29] used social networking sites to mine data on students and to then send them a message that looked like it had been sent by one of their friends. By using such "social data", the authors were able to increase the success rate of phishing from 16 to 72 percent. Hence, spear-phishing is considered a combination of technological approaches and social engineering.

2.3. Computer-supported collaboration

Businesses and employees use a wide range of technologies to facilitate, automate and improve daily tasks. We also see collaborative business structures emerging: Computer-supported collaboration tools for file sharing or collaborative workspaces, internal or external communication, blogs, wikis, etc., help connect staff within the company and to customers, allow widespread and instant information exchange about the entire business domain, and establish a constant communication channel to the customers and partners of the company.

Considering the wide range of different communication channels created by these computer-supported collaboration tools, social engineering attacks have a huge attack potential. However, in the business context, we differentiate between office communication and external communication. This enables us to make predictions about a victim's ability to detect a social engineering attack.

2.3.1. Office communication

Modern communication tools have changed communication flows among staff members enormously, making the high-speed exchange of information possible. There are sophisticated technologies that protect the security of data transfer. However, the majority of these countermeasures cover technical attacks, while social engineering attacks remain unconsidered. In enterprise environments, face-to-face communication is often replaced by e-mails or instant messages, generating a novel attack surface for social engineers. Obviously, social engineering attacks coming from internal accounts or e-mails with forged internal addresses are more likely to slip through the defenses of a potential victim. For instance, Parsons et al. [39] conducted a role-play scenario experiment in which 117 participants were tested on their ability to distinguish between phishing e-mails and benign e-mails. Their results indicated that people with a higher awareness level are able to identify significantly more phishing e-mails. Valuable personal information gained through social engineering attacks could have direct consequences, such as the exploit of a bank account,

[1] Maltego is an open source intelligence and forensics application. It allows the mining and gathering of information as well as the representation of this information in a meaningful way. `http://www.paterva.com/maltego/`

or indirect consequences, such as reputation loss [50]; it could also be used to improve the effectiveness of further social engineering attacks. Overall, we face multifarious social engineering attacks - once an attack is successful, the external adversary can use the information to become an insider and perform even more successful social engineering attacks.

2.3.2. External communication

As with intra-office communication, there is a trend towards the use of e-mail services, cloud, blogs, etc., for external communication, creating the same challenges as in internal communication. However, as the organizational border becomes increasingly blurred, it is difficult to decide which information may be published or passed on to an external communication partner. For instance, marketing blogs are useful for advertising purposes, but also carry the risk of unwanted information leakage. Another example is the release of information, e.g., about staff members, on LinkedIn, where a potential adversary can find out how many people are employed over a number of years and infer the economic status of the respective company from this data [8]. The strongest potential risk of external communication lies in the broad range of possible communication channels. Furthermore, new trends increase the number of channels, such as Bring Your Own Device (BYOD) [34] and the idea of "technology gets personal", which is used by Thomson [52] to explain the impact of using mobile devices to work with corporate information in insecure environments, such as cafés or public transport systems. He refers to mobile technology as the "window into the enterprises". Of course, security systems are installed on most of these devices; however, these systems offer no protection from social engineering attacks.

3. Social Engineering Taxonomy

In this section, we propose a taxonomy for the classification of social engineering attacks. Figure 1 illustrates the structure of our taxonomy and the attack scenarios, which we describe in detail in this section.

To classify social engineering attacks, we first introduce three main categories: **Channel**, **Operator**, and **Type**.

Attacks can be performed via the following channels:

- **E-mail** is the most common channel for phishing and reverse social engineering attacks.
- **Instant messaging** applications are gaining popularity among social engineers as tools for phishing and reverse social engineering attacks. They can also be used easily for identity theft to exploit a trustworthy relationship.
- **Telephone, Voice over IP** are common attack channels for social engineers to make their victim deliver sensitive information.
- **Social networks** offer a variety of opportunities for social engineering attacks. Given their potential to create fake identities and their complex information-sharing model, they make it easy for attackers to hide their identity and harvest sensitive information.
- **Cloud** services can be used to gain situational awareness of a collaboration scenario. Attackers may place a file or software in a shared directory to make the victim hand information over.
- **Website**s are most commonly used to perform waterholing attacks. Furthermore, they can be used in combination with e-mails to perform phishing attacks (e.g., sending an e-mail to a potential customer of a bank that contains a link to a malicious website that looks just like the bank's original website).

We also classify the attack by operator. The originator (operator) of a social engineering attack can be:

- **Human**: If the attack is conducted directly by a person. The number of targets is limited due to the lower capacity compared to an attack conducted by software.
- **Software**: Certain types of attacks can be automated with software. Examples include the Social Engineering Toolkit (SET), which can be used to craft spear-phishing e-mails [53]. A number of authors have discussed automated social engineering based on online social networks, such as Boshmaf et al. [13], Huber et al. [25] and Krombholz et al. [32]. The main advantage of automated attacks is that the number of possible targets that can be reached within a short period of time is considerably higher than with purely human attacks.

Furthermore, we categorize social engineering attacks into four types, namely:

- **Physical** as described in Section 2.2.1
- **Technical** as described in Section 2.2.4
- **Social** as described in Section 2.2.2
- **Socio-technical** as described in Section 2.2.5

Concerning social engineering, we determine the following attack scenarios: Attackers perform social engineering attacks over a variety of different channels. They are mostly conducted by humans as well as by software and furthermore categorized as physical, technical, social or socio-technical. The boundaries of the individual types of attack are highly expandable and have, in most cases, not yet been technically exhausted.

- **Phishing** is the attempt to acquire sensitive information or to make somebody act in a desired way

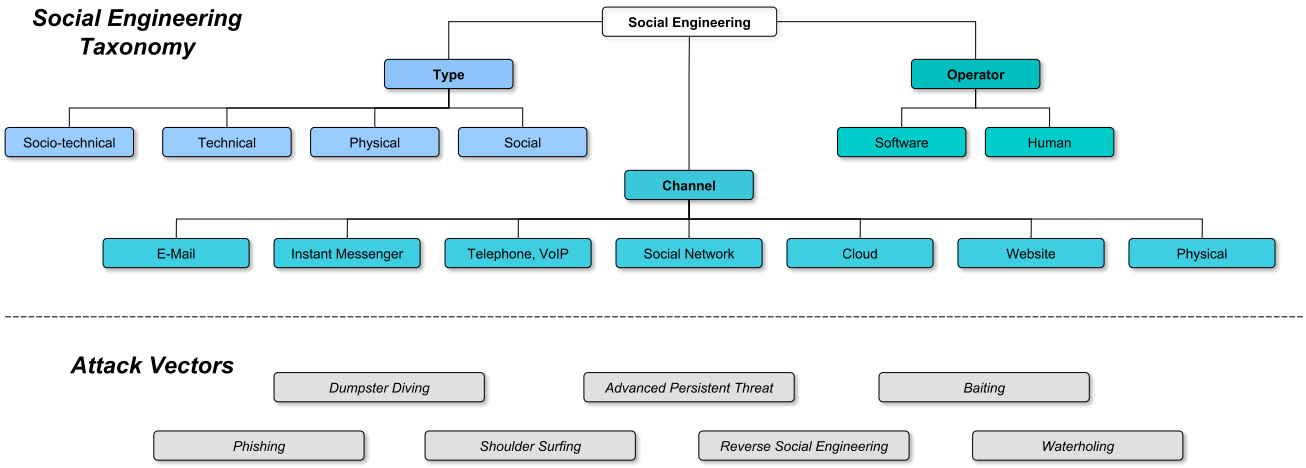

Figure 1: Overview of our classification of attack characteristics and attack scenarios.

by masquerading as a trustworthy entity in an electronic communication medium. They are usually targeted at large groups of people. Phishing attacks can be performed over almost any channel, from physical presence of the attacker to websites, social networks or even cloud services. Attacks targeted at specific individuals or companies are referred to as *spear-phishing*. Spear-phishing requires the attacker to first gather information on the intended victims, but the success rate is higher than in conventional phishing. If a phishing attack is aimed at high-profile targets in enterprises, the attack is referred to as *whaling*.

- **Dumpster diving** is the practice of sifting through the trash of private individuals or companies to find discarded items that include sensitive information that can be used to compromise a system or a specific user account.
- **Shoulder surfing** refers to using direct observation techniques to get information, such as looking over someone's shoulder at their screen or keyboard.
- **Reverse social engineering** is an attack where usually trust is established between the attacker and the victim. The attackers create a situation in which the victim requires help and then present themselves as someone the victim will consider someone who can both solve their problem and is allowed to receive privileged information. Of course, the attackers try to choose an individual who they believe has information that will help them.
- **Waterholing** describes a targeted attack where the attackers compromise a website that is likely to be of interest to the chosen victim. The attackers then wait at the waterhole for their victim.
- **Advanced Persistent Threat** refers to long-term, mostly Internet-based espionage attacks conducted by an attacker who has the capabilities and intent to comprise a system persistently.
- **Baiting** is an attack during which a malware-infected storage medium is left in a location where it is likely to be found by the targeted victims.

Table 1 outlines the relationship between our proposed social engineering taxonomy and current attack scenarios. We classified current social engineering attack scenarios based on our taxonomy. We can, for example, observe that a number of social engineering attacks exclusively rely on a physical attack channel, such as shoulder surfing, dumpster diving and baiting. To protect against this class of attacks, physical security needs to be improved. The table furthermore highlights that the majority of today's social engineering attacks rely on a combination of social and technical methods. Hence, to effectively protect against socio-technical attacks, user awareness for social engineering attacks needs to be improved and their devices protected on a technical level.

4. State-of-the-Art Attacks

This section provides an overview of state-of-the-art social engineering attacks. These attacks often use personal information from online social networks or other cloud services and can be performed in an automated fashion.

4.1. Online Social Networks (OSNs)

While the more traditional forms of social engineering use information collected through dumpster diving or phone calls, OSNs contain a wealth of personal information that can be misused as an initial source for social engineering attacks. Huber et al. were among the first researchers to argue that OSNs enable automated social engineering (ASE) attacks [23] because information harvested from OSNs is easy to process. The authors showed

Table 1: Classification of social engineering attacks according to our taxonomy.

		Phishing	Shoulder Surfing	Dumpster Diving	Reverse Social Engineering	Waterholing	Advanced Persistent Threat	Baiting
Channel	E-mail	✓			✓		✓	
	Instant Messenger	✓			✓			
	Telephone, VoIP	✓			✓			
	Social Network	✓			✓			
	Cloud	✓						
	Website	✓				✓	✓	
	Physical	✓	✓	✓	✓			✓
Operator	Human	✓	✓	✓	✓			✓
	Software	✓		✓	✓	✓	✓	
Type	Physical		✓	✓				✓
	Technical					✓	✓	
	Social				✓			
	Socio-technical	✓			✓	✓	✓	✓

that information on employees of a given target company can be collected in an automated fashion and potentially misused for automated social engineering. Reverse social engineering describes a particular social engineering technique where an attacker lures the victim into initiating the conversion as described in 2.2.3. Irani et al. [27] argue that OSNs enable reverse social engineering attacks and describe three potential attack vectors. The authors evaluated their proposed attack vectors on three different OSNs: recommendation-based reverse social engineering on Facebook, demographic-based reverse social engineering on Badoo and visitor-tracking-based reverse social engineering on Friendster. Their results show that reverse social engineering attacks are feasible in practice and can be automated by exploiting the features of current online social networks. While social spam is usually sent via an OSN's primary communication channel, attackers who harvest information can also send traditional e-mail messages to deliver spam because users provide their e-mail addresses on their profiles. If spam is delivered via traditional e-mail instead of OSN platforms, these malicious messages cannot be detected by the OSN's provider. Balduzzi et al. [9] showed that OSNs can be misused for automated user profiling, to validate large sets of e-mail addresses and to collect additional personal information corresponding to these sets.

Social phishing and context-aware spam

Phishing is a widely-spread threat on the Internet and consists of an attacker attempting to lure victims into entering sensitive information like passwords or credit card numbers into a faked website that is controlled by the attacker. It has been shown that social phishing [29], where "social" information specific to the victim is used, can be extremely effective compared to regular phishing. Jagatic et al. [29] found that when phishing e-mails impersonated a target's friend, the success rate increased from 16% to 72%. The social graph is, therefore, not only of value for the social network operator, but also for attackers. This is the case especially if it contains additional information like a valid e-mail address or recent communication between the victim and a friend whom the attacker can impersonate. With automated data extraction from social networks, a vast amount of further usable data becomes available to spammers. Prior conversations within the social network, such as private messages, comments or wall posts, could be used to determine the language normally used for message exchange between the victim and his friends, as a phishing target might find it very suspicious to receive a message in English from a friend with whom they normally communicate in French. Context-aware spam misuses personal information extracted from OSNs to increase the appearance of authenticity of traditional spam messages. Brown et al. [14] identified three context-aware spam attacks: relationship-based attacks, unshared-attribute attacks, and shared-attribute attacks. Relationship-based attacks solely exploit relationship information, making this the spam equivalent of social phishing. The two other attacks exploit additional information from social networks, information that is either shared or not shared between the spam target and the spoofed friend. An example of an unshared attack are birthday cards that seem to originate from the target's friend. Shared attributes, e.g., photos in which both the spam target and her spoofed friend are tagged, can be exploited for context-aware spam. Huber et al. [24, 26] found that the missing support for communication security can be exploited to automatically extract personal information from online social networks. Moreover, the authors showed that the extracted information could be misused to target a large number of users with context-aware spam.

Fake profiles

At the time of writing, the only requirement for the creation of a social networking account is a valid e-mail address, which makes it rather easy for attackers to create fake accounts. A study by Sophos published in 2007 with randomly chosen Facebook users showed that approximately 41% of social networking users accepted friendship requests from a fake profile [46]. Ryan and Mauch [5] further showed that fake profiles can be misused to infiltrate social networks: they set up a profile for a fictional American cyber threat analyst, called "Robin Sage", and were able to gain access to sensitive information in the military

and information security community. Bilge et al. [11] outlined two sophisticated fake profile attacks that could be used to infiltrate the trusted circles of social networking users: profile cloning attacks, where attackers clone existing user profiles and attempt to "reinvite" their friends, and cross-profile cloning attacks, where attackers create a cloned profile on an online social network where the target user does not yet have a profile and then contact the targets' friends. If a user, for example, has a Facebook account but no LinkedIn account, an attacker could clone the Facebook profile to create a LinkedIn profile and then contact the target's Facebook friends who are also on LinkedIn. Bilge et al. showed that their attacks can be fully automated and are feasible in practice. If an attacker is able to create fake accounts on a large scale, Sybil attacks on OSNs are possible. OSN providers therefore use various protection mechanisms to limit the creation of large amounts of fake accounts [49]. Boshmaf et al. [13] however found that OSNs can be infiltrated on a large scale. They evaluated how vulnerable OSNs are to a large-scale infiltration by socialbots - computer programs that control OSN accounts and mimic real users. The authors created a *Socialbot* Network (SbN): a group of adaptive socialbots that are orchestrated in a command-and-control fashion on Facebook. The authors used 102 fake profiles to send friendship requests to 5,053 randomly selected Facebook users. 19.3% of these users accepted the friendship requests. Next, the SbN tried to infiltrate the circle of friends of the users who had accepted their fake friendship requests. Within 8 weeks, the SbN was able to further infiltrate the network and gain access to personal information. A recent survey by Alvisi et al. [7] provides an overview of Sybil defenses for online social networks and proposes community detection algorithms.

4.2. Cloud services

Cloud services provide a new channel through which social engineers can conduct attacks on the knowledge worker. Knowledge workers frequently collaborate with others who do not work at the same location. Sharing information on a cloud service has therefore become popular. In this scenario, an attacker exploits this situation and uses the cloud as a channel for the social engineering attack. Recent publications described a variety of possible attacks in the cloud, e.g., an attacker placing a malicious file into another user's cloud as described by Gruschka et al. [21] and then using social engineering to make them execute the malicious file. A malicious piece of software can also be used to extract personal information from the victim's account, which is then used to perform more targeted attacks. Mulazzani et al. [37] provide countermeasures to reduce the risk by preventing the attacker from placing malicious files on Dropbox, one of the currently most commonly used cloud services. The level of trust between users of a shared directory or file is not always as high as desired. Social engineers can exploit this fact by using a fake identity or a compromised user account to invite the victim to share specific information with the attacker in the cloud. According to Roberts et al. [43], one of the biggest weaknesses of cloud services is that the users - companies and individual users - lose control over their data when they store and access it remotely. On traditional servers that are owned by a company itself, it can restrict access and define customized access policies. In cloud services, the responsibility for that is shifted to a third party. Therefore, if a cloud service is to be used for the exchange of sensitive information, a certain level of trust must be established not only between collaborating users, but also between the cloud hosting company and the user. The most commonly observed attacks on cloud services are spear-phishing and APTs.

4.3. Mobile applications

The increased use of mobile applications in both business and private contexts makes them an increasingly popular channel for social engineering attacks. In business communication, mobile messaging and e-mail applications are of high interest to social engineers. *BYOD* policies established by companies often include the use of mobile phones and tablets. More and more employees use their smartphones to check their company e-mails or to read documents that are stored in the cloud. However, many smartphone users use highly vulnerable smartphone applications that can be misused to conduct social engineering attacks. Schrittwieser et al. [44] presented two different attack scenarios that can serve as a starting point for such an attack. In their work [44], they demonstrated how sender ID spoofing can be done on popular mobile messaging applications such as *WhatsApp* [6]. A social engineer can use this to send a message to a victim while pretending to be one of his friends. The authors also highlighted how vulnerabilities can be exploited to hijack user accounts, which can then be used to perform social engineering. Considering that many smartphone applications are highly vulnerable and can leak sensitive information, we can conclude that such mobile devices offer a variety of attack vectors for social engineering and other attacks on user privacy. Moreover, some smartphone applications request permissions to access sensitive data on the user's device. If an attacker were to create such an application, they would obtain the information and could use it as a starting point for a social engineering attack. Chin et al. [15] discussed how inter-application information exchange can be sniffed on smartphones and then be misused to violate application policies and permissions. In some cases, such as described by Potharaju et al. [41], the attacker simply plagiarizes a popular smartphone application and deploys it in order to perform an attack.

5. Real-world Examples

In this section, we describe how targeted attacks against the knowledge worker are performed in real-world scenarios. Two methods were prevalently used in recent social

engineering attacks, namely spear-phishing and waterholing attacks. We discuss these two methods in detail and in the context of recent real-world attacks.

5.1. Spear-Phishing

Many companies deploy highly sophisticated end-point security controls to protect their networks. Nevertheless, targeted attacks such as spear-phishing are an increasing threat for knowledge workers because of their targeted precision. In practice, the first step within an attack scenario is that the attacker seeks publicly available information on the company's Internet site and public profiles on social networks to obtain precise information on the targeted victim. Then the attacker constructs an e-mail using the gathered information to gain the victim's trust. In general, such e-mails are only sent to a carefully selected small group of people. In most cases, they contain attachments with malicious software to provide a remote control tool to the attacker. Zero-day exploits are a good way of installing a backdoor via an existing vulnerability. The remote control functionality is then used to harvest sensitive information and to get into internal company networks. In this section, we discuss three real-world spear-phishing attacks and their impact on knowledge workers.

RSA, 2011. As described in [1], RSA suffered from an attack by an advanced persistent threat. A small number of employees received an e-mail with "2011 Recruitment Plan" in the subject line. Even though most of them found this e-mail in their junk mail folder, the e-mail was prepared well enough to convince the receiver of its trustworthiness. A number of employees thus directly opened the e-mail from the junk mail folder. The e-mail had a spreadsheet attached. According to [1], the spreadsheet contained a zero-day exploit that installed a backdoor via an Adobe Flash vulnerability. The attacker chose to use a Poison Ivy[2] variant to obtain remote control of the target's device which initially was not detected. After the initial social engineering phase was completed, the attackers compromised further machines in the local network. The attackers then successfully compromised a number of strategic accounts and were able to steal sensitive information on RSA's SecurID system. Eventually RSA had to replace millions of SecurID tokens [51] due to this successful social engineering attack.

New York Times, 2013. The New York Times was hit by a similar attack as RSA. Chinese hackers performed a 4-month targeted attack, infiltrating The New York Times computer systems and harvesting the employees' user credentials [40]. Reports suggest that there is evidence of political motives behind the attack. The attackers broke into e-mail accounts, tried to cloak the source of traffic to the The New York Times and to route traffic caused by the attacks through university computers located in the United States. Again, the initial attack vector had been a spear-phishing attack which sent fake FedEx notifications. The New York Times hired computer security experts to analyze the attack and prevent a persistent threat. They found that some of the methods used to break into the company's infrastructure were associated with the Chinese military. Additionally, the malware that was installed to gain access to the computers within the company's network followed the pattern of previously reported Chinese attacks. Perlroth [40] also reported that the same university computers in the United States had been previously used by hackers from the Chinese military. With this attack, the hackers stole the passwords of all employees at The New York Times and were thus able to access the personal devices of 53 people. However, according to The New York Times, no customer data was stolen. The characteristics of the attack clearly indicate a political motive for this APT. China's Ministry of National Defense stated that hacking is clearly prohibited under Chinese law and denied being the originator of the attack. According to Perlroth [40] China is using such attacks to control its public image in the West and therefore authorized attackers to injure organizations that might damage the reputations of Chinese authorities. They furthermore reported that hackers assigned by Chinese authorities had stolen sensitive information from more than 30 Western journalists.

The Red October Cyber-espionage Network. Kaspersky Lab [3] recently released a new research report on spear-phishing attacks against diplomatic, governmental and research organizations. The majority of the organizations targeted were located in former USSR Republics in Eastern Europe and Central Asia. The attack was launched in 2007 and remained active until the beginning of 2013. Sensitive data was not only stolen from research institutions but also from nuclear and energy groups as well as aerospace organizations. Similar to the attacks against RSA and The New York Times, the attackers sent a spear-phishing e-mail to a carefully selected group of people. For example, the attackers advertised cheap diplomatic cars in spear-phishing messages, which included custom malware. According to Kaspersky [3], the malware architecture consisted of malicious extensions, info-stealing modules and a backdoor Trojan that exploited Microsoft Office security vulnerabilities. The attackers also exfiltrated an enormous amount of sensitive data from the infiltrated networks. Stolen credentials were arranged in lists and then used to guess passwords of additional systems. The Red October APT remained active for almost six years. The in-depth analysis revealed artifacts within the executables of the malware that indicate that the attackers were located in a Russian-speaking country.

5.2. Waterholing

Recently, waterholing attacks have been the major vector in attacks on multi-national corporations alongside

[2] `http://www.poisonivy-rat.com/`

spear-phishing. Instead of directly targeting employees with customized phishing messages, the attackers target websites that are likely to be visited by their victims. They infect specific websites with malware and expect that some of their target companies' employees will visit them.

Apple, Facebook, Twitter. In 2013, a zero-day vulnerability in Java was exploited to specifically infiltrate corporate networks [55] via waterholing attacks. Apple and Facebook fell victim to this. The attackers first compromised the development site "iPhoneDevSDK", which is a popular forum for iOS application developers. The website was modified to exploit the Java vulnerability on devices which visited the website. A number of Apple's employees visited the infected website and their devices were compromised. The infected machines were connected to Apple's cooperate network and the attackers were thus able to infiltrate Apple's internal networks [54]. Facebook's internal network was infiltrated by the same waterholing campaign [56]. In both reported cases, it is assumed that no confidential data was stolen. A similar breach was reported earlier; however, it was not confirmed that the main cause was related to Java's zero-day vulnerability. In this attack, the attackers exploited Twitter's network and were able to compromise about $250,000$ user accounts, stealing user names, e-mail addresses, session tokens and encrypted passwords [30]. At first, it was believed that the attack was conducted by Chinese attackers instructed by the country's government or military, but others disagree and believe that one of Twitter's employees visited the same infected website as Apple's and Facebook's employees [55].

Reports on real-world examples focus to a large extent on the initial attack vector of social engineering and do not cover the technical aspects of these attacks. The first systematic analysis of targeted attacks at a large scale was conducted by Le Blond et al. [12]. In addition to the social aspects of targeted attacks, Le Blond et al.'s analysis discussed technical aspects based on malicious emails received by an NGO over a four-year period. The authors found that malicious office documents are the most popular attack vectors, followed by malicious archives. The analysis also showed that attackers tend to use well-known instead of zero-day vulnerabilities. The real-world examples of recent attacks performed by APTs underline that social engineering remains the most successful strategy to compromise large-scale organizations. The examples also highlight that the predominant channels of social engineering are e-mail communication and websites. Suggestions to counter social engineering attacks focus mainly on security policies and staff training [35, 18]. Gragg [19] points out that any education on social engineering must include psychology and persuasion in order to understand and counter attacks. Srikwan [47] recommends cartoons to teach users about social engineering and phishing. In the light of our discussed examples, user education might indeed help to counter spear-phishing attacks. Waterholing attacks, however, are hard to counter even with additional user awareness training and security policies. One possible approach to counter waterholing attacks could be to identify the most popular websites visited by employees to conduct an additional monitoring of these websites.

6. Conclusions

In this paper, we described common attack scenarios for modern social engineering attacks on knowledge workers. *BYOD*-policies and distributed collaboration as well as communication over third-party channels offers a variety of new attack vectors for advanced social engineering attacks. We believe that a detailed understanding of the attack vectors is required to develop efficient countermeasures and protect knowledge workers from social engineering attacks. To facilitate this, we introduced a comprehensive taxonomy of attacks, classifying them by attack channel, operator, different types of social engineering and specific attack scenarios. We discussed real-world examples and advanced attack vectors used in popular communication channels and the specific issues of computer-supported collaboration of knowledge workers in the business environment such as cloud services, social networks and mobile devices as part of *BYOD* policies. In this paper, we not only discussed complex advanced attack scenarios, but also provided a comprehensive classification that can serve as a basis for the development of countermeasures and further interdisciplinary research in the field.

7. Acknowledgements

This research was funded by the Austrian Science Fund (FWF): P 26289-N23 and COMET K1, FFG - Austrian Research Promotion Agency.

References

[1] Anatomy of an attack. available online: `http://blogs.rsa.com/anatomy-of-an-attack/`, last accessed on 2013-07-17.

[2] Google hack attack was ultra sophisticated. available online: `http://www.wired.com/threatlevel/2010/01/operation-aurora/`, last accessed on 2013-07-17.

[3] Kaspersky lab identifies operation "red october," an advanced cyber-espionage campaign targeting diplomatic and government institutions worldwide. available online: `http://www.kaspersky.com/about/news/virus/2013/Kaspersky_Lab_Identifies_Operation_Red_October_an_Advanced_Cyber_Espionage_Campaign_Targeting_Diplomatic_and_Government_Institutions_Worldwide`, last accessed on 2014-01-10.

[4] Microsoft hacked: Joins apple, facebook, twitter – InformationWeek. available online: `http://www.informationweek.com/security/\Attackacks/microsoft-hacked-joins-apple-facebook-tw/240149323`, last accessed on 2013-07-10.

[5] The robin sage experiment: Fake profile fools security pros. available at `http://www.networkworld.com/news/2010/070810-the-robin-sage-experiment-fake.html?t51hb`, last accessed on: 2013-07-14.

[6] Whatsapp. available online: `http://www.whatsapp.com/`, last accessed on 2013-07-18.

[7] L. Alvisi, A. Clement, A. Epasto, S. Lattanzi, and A. Panconesi. Sok: The evolution of sybil defense via social networks. *IEEE Symposium on Security and Privacy*, 2013.

[8] G. Bader, A. Anjomshoaa, and A. Tjoa. Privacy aspects of mashup architecture. In *Social Computing (SocialCom), 2010 IEEE Second International Conference on*, pages 1141–1146, 2010.

[9] M. Balduzzi, C. Platzer, T. Holz, E. Kirda, D. Balzarotti, and C. Kruegel. Abusing social networks for automated user profiling. In *Recent Advances in Intrusion Detection*, pages 422–441. Springer, 2010.

[10] R. Ballagas, M. Rohs, J. G. Sheridan, and J. Borchers. Byod: Bring your own device. In *In Proceedings of the Workshop on Ubiquitous Display Environments, Ubicomp*, 2004.

[11] L. Bilge, T. Strufe, D. Balzarotti, and E. Kirda. All your contacts are belong to us: automated identity theft attacks on social networks. In *Proceedings of the 18th international conference on World wide web*, pages 551–560. ACM, 2009.

[12] S. L. Blond, A. Uritesc, C. Gilbert, Z. L. Chua, P. Saxena, and E. Kirda. A look at targeted attacks through the lense of an ngo. In *23rd USENIX Security Symposium (USENIX Security 14)*, San Diego, CA, Aug. 2014. USENIX Association.

[13] Y. Boshmaf, I. Muslukhov, K. Beznosov, and M. Ripeanu. The socialbot network: when bots socialize for fame and money. In *Proceedings of the 27th Annual Computer Security Applications Conference*, pages 93–102. ACM, 2011.

[14] G. Brown, T. Howe, M. Ihbe, A. Prakash, and K. Borders. Social networks and context-aware spam. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, CSCW '08, pages 403–412, New York, NY, USA, 2008. ACM.

[15] E. Chin, A. P. Felt, K. Greenwood, and D. Wagner. Analyzing inter-application communication in android. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*, MobiSys '11, pages 239–252, New York, NY, USA, 2011. ACM.

[16] R. Cialdini. *Influence: science and practice*. Allyn and Bacon, 2001.

[17] P. F. Drucker. *Landmarks of tomorrow: a report on the new "post-modern" world*. Harper, New York, 1st edition, 1959.

[18] Gartner Inc. Protect Against Social Engineering Attacks. *Gartner Security webletter*, 1(1), Feb. 2002. [Retrieved 2008-11-13].

[19] D. Gragg. A Multi-Level Defense Against Social Engineering. *SANS Reading Room, March*, 13, 2003.

[20] S. Granger. Social Engineering Fundamentals, Part I: Hacker Tactics. *SecurityFocus*, 2001.

[21] N. Gruschka and M. Jensen. Attack surfaces: A taxonomy for attacks on cloud services. In *IEEE CLOUD*, pages 276–279, 2010.

[22] C. Herley and D. Florencio. Phishing as a Tragedy of the Commons. *NSPW 2008, Lake Tahoe, CA*, 2008.

[23] M. Huber, S. Kowalski, M. Nohlberg, and S. Tjoa. Towards automating social engineering using social networking sites. In *Computational Science and Engineering, 2009. CSE'09. International Conference on*, volume 3, pages 117–124. IEEE, 2009.

[24] M. Huber, M. Mulazzani, M. Leithner, S. Schrittwieser, G. Wondracek, and E. Weippl. Social snapshots: digital forensics for online social networks. In *Proceedings of the 27th Annual Computer Security Applications Conference*, 2011.

[25] M. Huber, M. Mulazzani, S. Schrittwieser, and E. Weippl. Cheap and automated socio-technical attacks based on social networking sites. In *3rd Workshop on Artificial Intelligence and Security (AISec'10)*, 10 2010.

[26] M. Huber, M. Mulazzani, E. Weippl, G. Kitzler, and S. Goluch. Friend-in-the-middle attacks: Exploiting social networking sites for spam. *IEEE Internet Computing*, 15(3):28–34, 2011.

[27] D. Irani, M. Balduzzi, D. Balzarotti, E. Kirda, and C. Pu. Reverse social engineering attacks in online social networks. *Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 55–74, 2011.

[28] K. Ivaturi and L. Janczewski. A taxonomy for social engineering attacks. 2011.

[29] T. Jagatic, N. Johnson, M. Jakobsson, and F. Menczer. Social phishing. *Communications of the ACM*, 50(10):94–100, 2007.

[30] R. King. Twitter: More than 250K user accounts have been compromised. online, 2013. available at: http://www.zdnet.com/twitter-more-than-250k-user-accounts-have-been-compromised-7000010711/, last accessed on 2014-01-21.

[31] K. Krombholz, H. Hobel, M. Huber, and E. Weippl. Social engineering attacks on the knowledge worker. In *Proceedings of the 6th International Conference on Security of Information and Networks*, SIN '13, pages 28–35, New York, NY, USA, 2013. ACM.

[32] K. Krombholz, D. Merkl, and E. Weippl. Fake identities in social media: A case study on the sustainability of the facebook business model. *JoSSR*, 4(2):175–212, 2012.

[33] K. Marett, D. Biros, and M. Knode. Self-efficacy, Training Effectiveness, and Deception Detection: A Longitudinal Study of Lie Detection Training. *lecture notes in computer science*, 3073:187–200, 2004.

[34] K. Miller, J. Voas, and G. Hurlburt. Byod: Security and privacy considerations. *IT Professional*, 14(5):53–55, 2012.

[35] K. Mitnick and W. Simon. *The Art of Deception: Controlling the Human Element of Security*. Wiley, 2002.

[36] F. Mohd Foozy, R. Ahmad, M. Abdollah, R. Yusof, and M. Mas' ud. Generic taxonomy of social engineering attack. 2011.

[37] M. Mulazzani, S. Schrittwieser, M. Leithner, M. Huber, and E. Weippl. Dark clouds on the horizon: using cloud storage as attack vector and online slack space. In *Proceedings of the 20th USENIX conference on Security*, SEC'11, pages 5–5, Berkeley, CA, USA, 2011. USENIX Association.

[38] R. Nelson. Methods of Hacking: Social Engineering. online, 2008. available at: http://www.isr.umd.edu/gemstone/infosec/ver2/ papers/socialeng.html, last accessed on 2013-07-04.

[39] K. Parsons, A. McCormac, M. Pattinson, M. Butavicius, and C. Jerram. Phishing for the truth: A scenario-based experiment of users' behavioural response to emails. In L. Janczewski, H. Wolfe, and S. Shenoi, editors, *Security and Privacy Protection in Information Processing Systems*, volume 405 of *IFIP Advances in Information and Communication Technology*, pages 366–378. Springer Berlin Heidelberg, 2013.

[40] N. Perlroth. Chinese hackers infiltrate new york times computers, Jan. 2013. available at `https://www.nytimes.com/2013/01/31/technology/chinese-hackers-infiltrate-new-york-times-computers.html`, last accessed on: 2013-07-01.

[41] R. Potharaju, A. Newell, C. Nita-Rotaru, and X. Zhang. Plagiarizing smartphone applications: attack strategies and defense techniques. In *Proceedings of the 4th international conference on Engineering Secure Software and Systems*, ESSoS'12, pages 106–120, Berlin, Heidelberg, 2012. Springer-Verlag.

[42] T. Qin and J. Burgoon. An Investigation of Heuristics of Human Judgment in Detecting Deception and Potential Implications in Countering Social Engineering. *Intelligence and Security Informatics, 2007 IEEE*, pages 152–159, 2007.

[43] J. C. Roberts, II and W. Al-Hamdani. Who can you trust in the cloud? a review of security issues within cloud computing. In *Proceedings of the 2011 Information Security Curriculum Development Conference*, InfoSecCD '11, pages 15–19, New York, NY, USA, 2011. ACM.

[44] S. Schrittwieser, P. Fruehwirt, P. Kieseberg, M. Leithner, M. Mulazzani, M. Huber, and E. Weippl. Guess Who Is Texting You? Evaluating the Security of Smartphone Messaging Applications. In *Network and Distributed System Security Symposium (NDSS 2012)*, 2 2012.

[45] SocialEngineer. What is phishing - paypal phishing examples. available online: http://www.social-engineer.org/wiki/archives/Phishing/Phishing-PayPal.html, last accessed on 2013-07-04.

[46] Sophos. Sophos facebook id probe shows 41% of users

happy to reveal all to potential identity thieves, 2007. available online: http://www.sophos.com/en-us/press-office/press-releases/2007/08/facebook.aspx, last accessed on 2013-07-13.

[47] S. Srikwan. Using Cartoons to Teach Internet Security. *Cryptologia*, 32(2):137–154, 2008.

[48] S. Stasiukonis. Social Engineering, the USB Way. 2006. available at http://www.darkreading.com/security/perimeter/showArticle.jhtml?articleID=208803634, last accessed on: 2013-07-02.

[49] T. Stein, E. Chen, and K. Mangla. Facebook immune system. In *Proceedings of the 4th Workshop on Social Network Systems*, SNS '11, pages 8:1–8:8, New York, NY, USA, 2011. ACM.

[50] L. Tam, M. Glassman, and M. Vandenwauver. The psychology of password management: a tradeoff between security and convenience. *Behav. Inf. Technol.*, 29(3):233–244, May 2010.

[51] The Wall Street Journal. Security tokens take hit, 2011. Available at `http://online.wsj.com/news/articles/SB10001424052702304906004576369990616694366`, last accessed: 01/12/2013.

[52] H. Thompson. The human element of information security. *Security Privacy, IEEE*, 11(1):32–35, 2013.

[53] TrustedSec. Social-engineer toolkit, 2013. available online at: `https://www.trustedsec.com/downloads/social-engineer-toolkit/`, last accessed 03/12/2013.

[54] Z. Whittaker. Apple hacked by same group that attacked Facebook. online, 2013. available at: http://www.zdnet.com/apple-hacked-by-same-group-that-attacked-facebook-7000011509/, last accessed on 2014-01-21.

[55] Z. Whittaker. Facebook, Apple hacks could affect anyone: Here's what you can do. online, 2013. available at: http://www.zdnet.com/facebook-apple-hacks-could-affect-anyone-heres-what-you-can-do-7000011520/, last accessed on 2014-01-21.

[56] Z. Whittaker. Facebook hit by 'sophisticated attack'; Java zero-day exploit to blame. online, 2013. available at: http://www.zdnet.com/facebook-hit-by-sophisticated-attack-java-zero-day-exploit-to-blame-7000011390/, last accessed on 2014-01-21.