The Cycle of Deception - A Model of Social Engineering Attacks, Defences and Victims

M. Nohlberg[1]and S. Kowalski[2]

[1] School of Humanities and Informatics, University of Skövde, Skövde, Sweden
[2]SecLab, Department of Computer and Systems Sciences, Stockholm University/Royal Institute of Technology, Stockholm, Sweden
e-mail: marcus@nohlberg.com

Abstract

In this paper we propose a model for describing deceptive crimes in general and social engineering in particular. Our research approach was naïve inductivist and the methods used were literature study and interviews with the lead investigator in a grooming case, as we see many similarities between the techniques used in grooming, and those used in social engineering. From this we create cycles describing attacker, defender, and the victim and merge them into a model describing the cycle of deception. The model is then extended into a possible deception sphere. The resulting models can be used to educate about social engineering, to create automated social engineering attacks, to facilitate better incident reporting, and to understand the impact and economical aspects of defenses.

Keywords

Social engineering, fraud, deception, security models, computer crime

1. Background

Social engineering is a term used for techniques to con, or trick, victims into giving the attacker sensitive information or to get them to perform actions that the attacker wants them to do. A good definition is given by Kevin Mitnick in an interview by Tanneeru (2005): *"Social engineering is using manipulation, influence and deception to get a person, a trusted insider within an organization, to comply with a request, and the request is usually to release information or to perform some sort of action item that benefits that attacker. It could be something as simple as talking over the telephone to something as complex as getting a target to visit a Web site, which exploits a technical flaw and allows the hacker to take over the computer."* Since many users do not believe that anyone would ever trick or con them, because they are not "rich and famous", and that hackers "cannot do much damage anyway" (Brostoff et al, 2002) these attack techniques are often quite successful. This is further complicated by the fact that most users do not understand how security works, and therefore construct their own, often incorrect, models (Adams & Sasse, 1999). There are a lot of studies on the "gullibility" of users, both academic and non-academic. One example is a study performed by Treasury Department inspectors, where one third of the Internal Revenue Service (IRS) employees gave away their

login and password to auditors who called pretending to be computer technicians (Dalrymple, 2005). This and several other studies demonstrate a high degree of susceptibility to social engineering attacks.

Social engineering as a term has been used for quite a while in the security sector. We have found references dating back to 1995 (Winkler & Dealt, 1995), and there was of course the boom of notoriety of social engineering in connection with the case of Kevin Mitnick (Mitnick, 2002). His warrants, and later imprisonment, lead to a great deal of publicity for the technique. One of the more influential contributions from Mitnick in an academic setting is the social engineering attack cycle, SEAC, as seen in figure 1 below (Mitnick, 2002). This descriptive cycle is frequently used both by security professionals and academics when describing social engineering attacks.

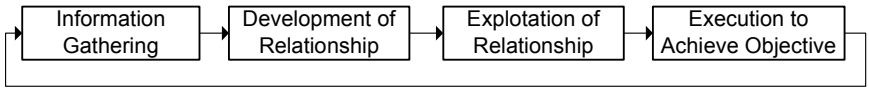

Figure 1: The Social Engineering Attack Cycle (From Mitnick, 2002)

In this paper we propose a new conceptual model of the social engineering attack cycle, which includes descriptions of both defenders and victims. This new cycle addresses the predominant problems with the existing model. We find that the existing model is overly simplistic, while at the same time being obscure. It is quite common that those using the model put too great an emphasis on describing a step-by-step approach while giving little support for the iterative reality of most attacks. SEAC does not provide any suggestions for proper protection, making the model of limited use. Our aim is to provide a new model of the social engineering attack cycle that can be used as both a teaching aid and a framework for developing a holistic protection strategy.

The paper is structured with an introductory section on the research area and problem followed with a discussion on the method used. Then the model is presented and the paper ends with a brief discussion concerning the strengths, weaknesses and possible use of the model.

2. Method

The general approach in this study was of the naïve inductivist (Kowalski, 1994) as illustrated in figure 2 below.

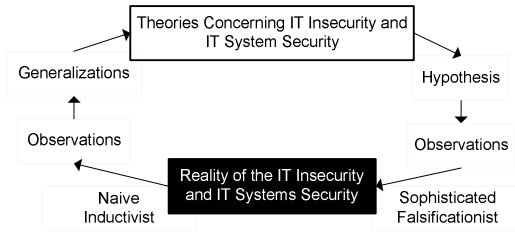

Figure 2: Naïve Inductivist and Sophisticated Falsificationist approaches to studying IT systems Security and IT insecurity (Kowalski, 1994, pp. 4)

We started out observing the problem with the need for an improved model of social engineering attacks and the assumption that this could be created using the body of knowledge existing but also by studying actual crimes. The observations were made on criminals using grooming, and a new model was created using generalizations from the data gathered. The two methods used for observations in this study are literature study and interviews. The literature study has been ongoing for several years, and has covered a large part of the written materials on the subject. This has given an understanding for the concept and what it consists of, but also on the current misconceptions and areas where improvements are needed. In order to understand more about how attackers actually use social engineering, we set out on trying to find public, well documented, cases of social engineering. This proved to be difficult. Either the attacks were poorly documented or purely anecdotal. This is a common problem when it comes to security research; it is hard to get access to well documented data about successful attacks, as most organizations tend to keep those secret, if they even know about the attack. In order to address this problem we looked outside of traditional social engineering attacks for other crimes using similar patterns and techniques. We found that the attack patterns used in grooming matched social engineering. Grooming is the term for when an adult, most often a man, tries to convince a child or a minor to sexual acts (O'Connell 2003). The advantage of studying grooming is that once grooming is reported to the police, a thorough, and often public, investigation is made. In our study we chose to focus on the most infamous Swedish groomer, who during a period of several years developed an ever increasingly efficient method of grooming using a combination of advanced manipulative techniques and technology (The Local, 2005). In order to study this case we studied the public records from the trial, but also did telephone interviews with the lead investigator in the case. The first was an open interview aiming to get background of the case, and the second interview was semi-structured, as described by May (2001), and developed with the SBC-model (Kowalski, 1994) as a foundation for the questions asked. The SBC-model is a model of information security that covers both technical and social aspects of security. The questions were designed in order to get specific knowledge while still being open enough as to make it possible to ask complimentary questions and to have a dialogue with the subject. The subject was informed about the usual ethical issues before the interview and also allowed to read through the transcripts and give comments to ensure correctness, upon which a couple of minor misunderstandings were corrected.

From the data gathered by the interviews, the steps in the attacks were identified and described. By analyzing these steps and using knowledge gained from the literature study, a model for the attack cycle was developed, followed by the victim cycle. They were complemented by a defense cycle adapted from Kowalski (2002). These three cycles were then merged into a complete model, which was later discussed with security professionals and especially well received by students when used as a teaching aid.

3. Results

When studying the grooming attacks, it was apparent that the attacks had eight steps. The steps have been somewhat simplified in the description below, but the general gist is maintained.

1. Create a pretext, in this case a fictional female model.
2. Contact the victims, or set up web pages so the victims contacted the attacker believing that he was the female model.
3. Get interest, gather information and get compliance from the victims.
4. Move the victim to unfamiliar location physically, use information gathered earlier.
5. Perform crime, in this case sexual abuse or rape.
6. Contact the victims again afterwards, get the victim to agree that what happened was OK.
7. Try to recruit the victim into finding new victims.
8. Relive the crime, from stored data/pictures, and convince that no crime happened if confronted for instance by parents or the victims.

Some of the steps in the description above are specific for grooming, but five generalizable steps in the attack cycle are apparent and are used for the attack cycle.

3.1. The attack cycle

The attack cycle is about the behavior of the attacker, and the actions he or she will take in an attack. In figure 3 below, the stages of the attack cycle are described.

Figure 3: The attack circle

Goal & Plan: The attacker must have a purpose with the attack, a goal, and a plan how to reach it. Here traditional criminological knowledge comes into play. The four classic traits that the attacker must possess are method, motive, opportunity, and

means (Pfleeger & Pfleeger, 2003). In order to be able to perform an attack, the perpetrator must know about what kind of attacks are possible, the method. Some methods are obvious and require no great cunning or planning, while others require certain skill or knowledge. There are three basic ways to acquire this knowledge. The method might be known beforehand, it might be searched for with the specific intention to use for attacks, or it might be found serendipitously. The perpetrator might discover an attack that works well on the first try, or might find a book or text describing attacks without having any intention to perform attacks beforehand. It is notable here that Sunderland's Differential Association Theory (DeMelo, 2007) states that once a potential perpetrator learns the methods required, he or she can easily pick up the required motive from just about anyone. So by learning the methods required it is probable that the perpetrator will also pick up the motives needed. The criminal culture, as discussed by Ferrell (1995), can be seen as the major factor determining crime. In fact, one of the flaws of traditional criminological reasoning is that the contemporary culture is sometimes neglected in the consideration of criminological analysis. The criminal subculture spans more than simply proximity, something that is ubiquitous in a connected world, it also concerns motives, drives, rationalizations and attitudes as well as certain appearances, group specific language and self presentation, style (Ferrell, 1995). *Map & Bond:* This is where the attacker tries to get information needed for the attack. This can be done by using traditional social engineering techniques such as dumpster diving or desktop hacking, or by searching the web for data and by studying other open sources of information. It can also, however, be when the attacker befriends the victim or someone with usable knowledge, and use manipulative techniques to get them to divulge the information the attacker needs, or to "prepare" them for the next step. In order to create a deceptive relationship the attacker uses influence techniques, for instance authority, scarcity, liking & similarity, reciprocation, commitment & consistency, social proof, and involvement (Cialdini, 1993). The influence techniques then exploits certain social psychological weaknesses, as suggested by the taxonomy put forth by Jordan & Goudey (2005). In other words the attacker manipulates the victim into trusting the attacker. *Execute:* The execute-step is where the attacker does something that is clearly illegal or not allowed, for instance when the target is asked to submit his log-in information, or when the nefarious e-mails are sent. *Recruit & Cloak:* Cloak is the actions performed after the execution, actions performed in order to hide the illegal activities. It can be to continue with the "friendship" to normalize the actions, moves to make the victim seem untrustworthy, or more advanced techniques to hide the crime. In some cases the victim can be recruited to either work for the attacker or as an ambassador/reference for the attacker. *Evolve/Regress:* This is where the attacker learns from the process and creates an internal justification for what has happened. There are basically two choices for the attacker here. Either the attack evolves, moving into another phase of the attack if the process has been successful up to this step. The other choice if the results to this point have been unsuccessful is to regresses, which can either be to stop the attack or to move to a more basic level of attack in order to be successful again.

3.2. The defense cycle

The defense cycle describes the general options available to the defender. The defender might in some cases be the same person as the victim or it might be security professionals in an organization or similar. This section is based on the work of Kowalski (2002), from which the terms and definitions have been taken and the flow is identified.

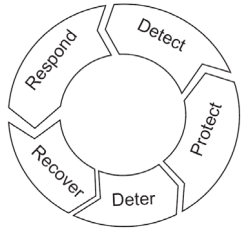

Figure 4: The defense cycle (Adapted from Kowalski, 2002)

The description of defenses given by Kowalski (2002) has been adapted to a circle to match this model, as shown in figure 4 above. Below are a couple of examples of implementations, but they can of course consist of many other measures. The description is based on what the defender has to do to be successful in providing defenses. By having a good, public policy or a rumor of reporting incidents to the police, you can *deter* an attacker. By making little sensitive data available, and educating employees about the risks and methods of attackers bonding with them as well as providing a strong policy on how to act, you *protect* the organization. By running a surveillance of the network communication, you can see when sensitive data are being sent, or when sensitive data are accessed, and by having well-educated employees that know when they are asked illicit questions, you *detect* an attack. By making it easy and without social or professional stigmata to report social engineering incidents, and by making the employees aware of how they can be manipulated into acting on the behalf of attackers, you are able to *respond* to an ongoing attack. By knowing the value of your data, having attacks reported and a well-designed policy, you can *recover* from the attack and learn from it. Hopefully you can be able to find the attacker to prevent him from evolving and attacking you, or others, in the future.

3.3. The victim cycle

The victim cycle is focused on the behavior of the individual victim, the person, in the attack. A common mistake when analyzing crime is to focus too much on the attacker, and to forget the victim. In fact, many crimes can be more readily prevented by focusing on the victim rather than the attacker. The flow is described in figure 5 below.

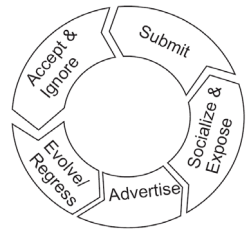

Figure 5: The victim cycle

By having something of value and making it known, either knowingly or unknowingly, the victim *advertises* its suitability as a target. By *socializing* with the criminal the victim sets itself up for deception, and by *exposing* valuables they make them accessible to the attacker. When the actual crime is executed, the victim *submits* to it, for instance by giving out the secret information. After the crime has been executed, the victim can chose to *accept* it, for instance by believing that it was not so "serious", or by simply *ignoring* it, either knowingly or by actually not knowing about the crime. By learning from the crime, the victim can *evolve* into someone who is harder to victimize in the future, but it is also possible that the victim can *regress*, turning into someone who accepts the role as a victim and is an easier pray in the future.

3.4. Adding the element of control

The attack circle could be perceived as the attackers' strive to reach the target in the center, and the way of reaching the center is by increasing control. Once enough control is gained through the process, the risk is reduced to a level acceptable by the attacker, and the attack can be performed. The level of acceptable risk is individual to each attacker. This is illustrated in figure 6 below.

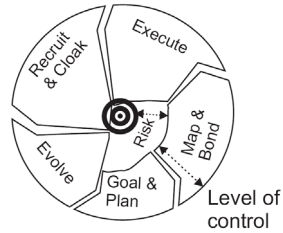

Figure 6: The attack cycle with added control

This aspect of control is also true for the victim and the defender, where they both strive to have a higher level of control than that of the attacker.

3.5. The cycle of a social engineering attack

When merging the three different cycles and adding a target in the center, a more holistic view of the prerequisites of a social engineering attack appears as seen in figure 7 below. One of our theories is that in order to have a "successful" social

engineering attack, all the steps in all the cycles have to fall in place. The attacker needs to succeed with the first three steps in the attack in order to be successful and the fourth and fifth in order to be able to continue attacking in the future. This is based on the reasoning that if the attacker is unable to provide a plan and a method for the attack, he will most likely fail. If he cannot learn about the potential victim or perform the attack, he will fail. If the attacker is unable to hide the attack, he will most likely get caught, and if the internal rationalization of the attacker does not judge the attack as a "good" experience, he will most likely not continue.

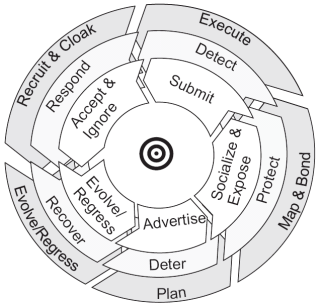

Figure 7: The cycle of deception

The same is true of the defender. If any one of the steps in the defense cycle is good enough to stop the attacker, then the attack will obviously fail or let the attacker be caught. In contrast, if there is no single part of the cycle that stops the attacker, then the attack will not fail due to the activities of the defender. Looking at the victim cycle we assume that the victim must submit in each of the sections in the model in order for the attacker to succeed. There are perhaps exceptions to this assumption, but based on sound reasoning it should be mostly true. This gives a possibility to consider the economic impact of this model. Choosing to invest in sections that are among the first three will stop the crime from happening. Investing in the fourth and fifth might stop the crimes from happening in the future. When considering purchasing defenses or educating the users it is relevant to consider where in the cycle of deception that particular investment should be placed. By knowing this it is possible to see if the investment will have the intended consequences.

3.6. The spherical view of deception

The earlier descriptions presented here are simplified and omit the fact that most attacks span several cycles and include several smaller attacks that are parts of the larger attacks. Therefore there is a need for a third dimension in the description. This added information is useful for instance when trying to create a piece of software using this model, but also in order to illustrate a more complete image of the attack.

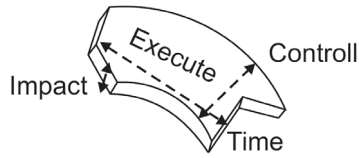

Figure 8: A 3D view of the Execution Phase

What is added in figure 8 above is the element of control, the time, and the impact. The impact is how noticeable the attack is for the victim and controlling organization. The goal for the attacker is to keep the impact as low as possible, while still being able to achieve a high level of control in a short time. The separate cycles in the whole attack each belong to one of the general attack cycles. For instance, there might be several smaller attacks performed in order to facilitate the first step (goal & plan) of a larger attack. Using this imagery and extending it into a more holistic description, we get the spherical view of attacks, presented in figure 9 below.

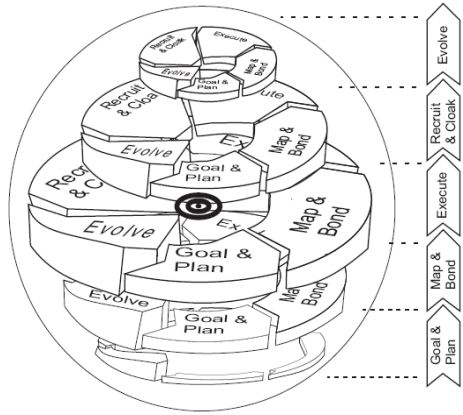

Figure 9: The spherical view of attacks

This model adds a third dimension to the description. The target is reached in the centre of the sphere, in the execution phase.

4. Discussion

This proposed model has several uses. It is can be used when educating about social engineering. As the model covers activities by all involved – victims, attackers, and the protecting organization – it gives a holistic and understandable view. It also facilitates a deeper understanding of the criminal process, and perhaps most importantly, it gives an easy way to understand how to develop a protection strategy. This model is a good tool in order to prioritize and understand the implications of spending in a certain area while neglecting other areas.

From an academic point of view, this model presents an excellent starting point for security researchers trying to position themselves. When looking at this model from a computer science perspective, it can be used to facilitate the implementation of an

automated social engineering AI-bot. By using this model, the attack flow of the AI-bot, as well as the target points, that is, where the AI-bot has gained enough information to move on, can be implemented. This allows for the creation of sophisticated AI-bots which in turn can be used to train users to avoid falling for real attacks, as well as be used for social engineering penetration testing. Security professionals can use the model when studying and visualizing the readiness of an organization. The model can also be used to investigate and understand attacks made against the organization, as well as making it easier for the potential victims to report their experiences, as they can be guided in the process by the steps involved in this model. This can provide improved incidence reports, something that is quite important in fighting social engineering.

One of the potential problems with this model is that it is developed from a study of grooming, which is not the same crime as social engineering. After comparing what is known about social engineering and our case of grooming it is, however, apparent that these two attacks share many of the same methods, manipulative techniques, and the same flow of the attack, even if the end goals and motivations are quite different. In fact, our model might be valid and have the same merits for other crimes and nefarious acts purposefully carried out by rational perpetrators. The model has a slight focus in examples given on traditional social engineering, with an attacker having interaction with a victim, but it should be valid also for more technical attacks such as phishing. The major difference is that those attacks often have a shorter timeframe, so the parts of the cycle might be moved through quicker. Whether an attacker calls or uses a spear-phishing attack makes no major difference for the cycle, however. The attacks that are not covered by this model are those that are based almost purely on random successes from exposure to large numbers of users. For instance, the most basic form of phishing, or malware, where the attack is a generic and non-specific message used against a group of non-tailored victims, obviously has very little “Map & Bound”, apart from the message sent out to deliver the attack. The model could be argued to cover these attacks too, but we feel that the greatest use of the model comes when applying it to crimes with a more specific intent than random frauds.

In the future this proposed model could be studied further, and analyzed in more types of attacks and crimes than only social engineering to validate a broader use for it. There is a possibility that this model can help us understand and prevent more crimes than social engineering.

References

Adams A. and Sasse M. (1999), *Users are not the Enemy: Why users compromise computer security mechanisms and how to take remedial measures,* Commun. ACM 42.

Brostoff S., Sasse A. and Weirich D. (2002), Transforming the ”weakest link”: A Human-computer Interaction Approach to Usable and Effective Security, *BT Technology Journal* 19(3), 122-131.

Cialdini, R. (1993), *Influence: the psychology of persuasion*. New York, Quill, ISBN: 0688128165.

Dalrymple, M. (2005), *Auditors Find IRS Workers Prone to Hackers*. [Online]. AP. Available from: http://sfgate.com/cgi-bin/article.cgi?file=/news/archive/2005/03/16/national/w162055S07.DTL, (Accessed 6 Mar 2006)

DeMelo, D. (2007), *Sutherland's Differential Association*. [Online]. Available from: http://home.comcast.net/~ddemelo/crime/differ.html, (Accessed 29 Jan 2007)

Ferrell, J. (1995), Culture, Crime, and Cultural Criminology. *Journal of Criminal Justice and Popular Culture*, 3(2) (1995) 25-42.

Jordan, J. and Goudey, H. (2005), The signs, signifiers and semiotics of the successful semantic attack. Presented at the *14th Annual EICAR Conference*, St.Juliens/Valletta, Malta, 2005.

Kowalski, S. (1994), *IT Insecurity: A Multi-disciplinary Inquiry*. Diss. University of Stockholm. Report series No. 94-040, Stockholm.

Kowalski, S. (2002), Value Based Risk Assessment: The Key to a Successful Security Target for the Telecommunication Industry, 3rd International Common Criteria Conference (ICCC) Ottawa, 2002.

May, T. (2001), *Social Research: Issues, Methods and Process*. Buckingham: Open University Press, ISBN: 0335206123.

Mitnick, K. (2002), *The Art of deception.* Indianapolis:Wiley Publishing, Inc., ISBN: 076454280X.

O'Connell, R. (2003), *A typology of child cybersexploitation and online grooming practices*. [Online]. University of Central Lancashire: Preston. Available at: http://www.uclan.ac.uk/host/cru/docs/cru010.pdf, (Accessed 11 Dec 2007)

Pfleeger, C. and Pfleeger, S. H. (2003), *Security in Computing* (3rd ed). Upper Saddle River: Prentice Hall, ISBN: 0130355488.

Tanneeru, M. (2005) A convicted hacker debunks some myths. [Online]. CNN. Available at: http://edition.cnn.com/2005/TECH/internet/10/07/kevin.mitnick.cnna/index.html (Accessed 28 Mar 2008)

The Local (2005), Man sexually abused "at least sixty girls". [Online]. Available at: http://www.thelocal.se/2756/20051228, (Accessed 1 Dec 2007)

Winkler, I. and Dealt, B. (1995), *Information Security Technology? ...Don't rely on it A case Study in Social Engineering.* [Online]. Proceedings of the Fifth USENIX UNIX Security Symposium. Available at: http://www.usenix.org/publications/library/proceedings/security95/full_papers/winkler.ps, (Accessed 12 Apr 2004)