Chapter 6

DETECTING SOCIAL ENGINEERING

Michael Hoeschele and Marcus Rogers

Abstract This paper focuses on detecting social engineering attacks perpetrated over phone lines. Current methods for dealing with social engineering attacks rely on security policies and employee training, which fail because the root of the problem, people, are still involved. Our solution relies on computer systems to analyze phone conversations in real time and determine if the caller is deceiving the receiver. The technologies employed in the proposed Social Engineering Defense Architecture (SEDA) are in the proof-of-concept phase but are, nevertheless, tractable. An important byproduct of this work is the generation of real-time signatures, which can be used in forensic investigations.

Keywords: Social engineering, defense architecture, attack signatures, forensics

1. Introduction

Much of the research literature on social engineering focuses on its use in perpetrating computer-related crime as opposed to detecting or preventing social engineering attacks. One exception is the work by Rogers and Berti [15] that discusses how to prevent social engineering attacks. Most of their suggestions are related to security policies and training personnel to be aware of what social attacks may look like; both of these are dependent on the human element. Dolan [4] makes similar suggestions, claiming that successful social engineering attacks rely on the employees of an organization. He emphasizes that to contain such attacks, employees must be well-trained and familiar with common social engineering techniques.

Social engineering is a clear threat. In a 1997 article, Bort [3] noted that: "Of the 384 respondents who confessed to being attacked over the last year, social engineering was to blame in 15 percent of the cases – the second largest cause." Other more recent publications, such as the

annual FBI/CSI Computer Crime and Security Surveys [5, 12, 14], do not consider social engineering attacks *per se*. But the survey results and accompanying discussion allude to social engineering being an ever present threat.

According to Bort [3], no hardware or software can defend information systems against a human being telling a convincing lie. While this may still be true, natural language processing researchers focusing on the problem of lie detection have made some progress. In particular, Raskin and co-workers [13] have proved that the problem is tractable and have created a working model.

This paper considers the problem of detecting social engineering attacks perpetrated over phone lines. The solution relies on computer systems to analyze phone conversations in real time and determine if the caller is deceiving the receiver. While the technologies employed in the proposed Social Engineering Defense Architecture (SEDA) are still in the proof-of-concept phase, they are, nevertheless, tractable [13]. In any case, the attack signatures generated as a consequence of this work can be used in forensic investigations of social engineering attacks.

2. Problem Statement

Social engineering attacks are a threat to all organizations. However, forensic investigations principally focus on attacks perpetrated using computer systems. Little, if any, work has considered the forensic analysis of social engineering attacks. Signatures have not been identified for social engineering attacks nor have systems been developed to log activity associated with such attacks. This means that even if it can be determined that a social engineering attack has occurred, it is very unlikely that the crime can be traced back to the perpetrator, let alone be prosecuted. Social engineering attack signatures are somewhat elusive because of the nature of the attacks and their avenue, which is most commonly the telephone system. Other avenues for attack, such as face-to-face conversations, are quite rare because of the risk incurred by attackers. This work considers the problem detecting of social engineering attacks perpetrated over phone lines. The proposed Social Engineering Defense Architecture (SEDA) can generate attack signatures in real time. These signatures allow a logging facility to also serve as a defense mechanism as in a typical network intrusion detection system.

3. Social Engineering

Dolan [4] defines social engineering as "using relationships with people to attain a goal." For the purposes of this work, the term social engineering is discussed with reference to attackers who attempt to illegally compromise an organization's assets. It should be noted that the types of organizations under attack are not limited to faceless multinational corporations. Educational institutions, banks and even the corner video store are equally at risk.

3.1 Methods

Social engineers, like computer hackers, take advantage of system weaknesses. As Dolan [4] states:

> *"Social engineers use tactics to leverage trust, helpfulness, easily attainable information, knowledge of internal processes, authority, technology and any combination there of. They often use several small attacks to put them in the position to reach their final goal. Social engineering is all about taking advantage of others to gather information and infiltrate an attack. The information gained in a phone book may lead to a phone call. The information gained in the phone call may lead to another phone call. A social engineer builds on each tidbit of information he or she gains to eventually stage a final, deadly attack. A successful social engineering attempt could result in great financial loss for the target company. A motivated attacker will be willing to gain information in any way possible."*

Social engineering is successful because people, in general, have a desire to help others and gain satisfaction from it [4]. An expert social engineer has the ability to establish trust and usually masquerades as someone the victim would trust.

Much of the information necessary to perpetrate social engineering attacks is publicly available. Reverse phone look-up directories, such as www.reversephonedirectory.com, are freely available on the Internet. Once a phone number and address are obtained, other useful information can be obtained effortlessly. The web pages of many organizations hold considerable information, e.g., organizational charts and biographical sketches, that can be used in social engineering attacks.

One common social engineering technique is to call the main switchboard of an organization and ask to be transferred to an employee. The receiver of the transfer call does not typically have a phone number that is posted publicly, so he assumes that the call is from an insider. This can prove to be a strong enough credential to allow more internal numbers to leak out, furthering the social engineer's cache of information. Arthurs [2] identifies other examples of social engineering attacks:

- **IT Support:** A social engineer claiming to be from the company's IT support group phones a user and explains that he is locating faults in the company network. He has narrowed the fault to the user's department but he needs a user ID and password from a department employee to identify the problem. Unless the user has been properly educated in security practices, he will very likely give the "trouble-shooter" the requested information.
- **Manager:** A social engineer, using a perceived position of authority, phones the help desk demanding to know why he cannot login with his password. He intimidates the help desk into giving him a new password by emphasizing that he has only a limited time to retrieve information for a report to the company vice president. He may also threaten to report the help desk employee to his supervisor.
- **Trusted Third Party:** A social engineer phones the help desk claiming to be the vice-president's executive assistant. She says that the vice-president has authorized her to collect the information. If the help desk employee balks, she threatens to inform the employee's supervisor.

It can be seen from these examples that the majority of social engineering attacks are committed over the telephone and rely on the fact that the receiver of the call takes the caller's word about his/her identity. Typically, there is no authentication other than answering questions pertaining to information that only an employee would know.

3.2 Motives

The Hackers Manifesto [10] explains why hackers desire to break into secure systems. The primary drivers are the quest for knowledge and the challenge. While harm may not be the intent, it is clear that considerable damage can be caused.

Much more dangerous are social engineering attacks that are intended to compromise an organization's assets. Examples are a recently fired employee seeking vengeance and a seasoned social engineer attempting to steal corporate information. These individuals have a better chance of succeeding and causing damage as they have greater motivation to succeed and potentially more resources at their disposal [4].

3.3 Targets

The targets of social engineering attacks range from personal information to intellectual property. However, it is important to note that

these are only the end targets of the attack, and many small pieces of information must be obtained before the final target can be reached. The information includes organization policies, protocols, hierarchy, phone books and server names. Much of this information is available on an organization's website or is obtained by making phone calls to employees. The real challenge is protecting such seemingly trivial information without interfering with day-to-day operations [15].

4. Current Solutions

Security policies and employee training are the two main approaches for preventing social engineering attacks. However, these approaches are fundamentally flawed as they rely on humans to patch security holes. The flaw is that human trust is the vulnerability that is exploited by social engineers; the notion of trust is deeply embedded in Western culture and is difficult to overcome [15]. A good social engineer can be very convincing that he or she needs the requested data, and that the individual who receives the request is hurting the organization by not helping. Employee training may stop novice attackers, but a seasoned social engineer will likely not be considered a risk during a conversation, so the training will never be triggered. This point is evident in the social engineering attacks described in the previous section (see [15] for additional details).

4.1 Security Policies

Security policies alone cannot prevent break-ins. Often, a security policy is effective only in the sense that after the policy is broken, it is easy to display the policy and show how it was violated. However, a security policy that classifies data into different levels of sensitivity can be quite effective as it requires a social engineer to obtain higher credentials to gain access to sensitive information. The end result is more work, which would deter most casual social engineers. However, an experienced, persistent social engineer will likely keep working until he has all the credentials needed to access the information.

Allen [1] advocates security policies because they provide clear direction on what is expected of employees in an organization. Equally important is limiting data leakage by reducing the amount of specific data that is available. This has the effect of making social engineering attacks arduous and time consuming.

4.2 Employee Training

Employee training is currently the most effective deterrent to social engineering attacks. The training ranges from annual multi-day seminars to constant reminders via posters and mailings. The idea is that if employees know how social engineers execute attacks and gain trust, they will be able to detect attacks as they occur and take steps to defeat them. Employees are also encouraged not to release certain information over the phone, e.g., passwords and ID numbers. But the problem is that an expert social engineer never shows any signs of being an attacker; often, the social engineer appears to be a very conscientious employee. Therefore, it is not realistic to expect employees to be the primary defense against social engineering attacks. It is, however, logical to make them aware of social engineering attacks. This also helps in gaining employee acceptance of policies and systems designed to defend against social engineering attacks [15].

4.3 Evaluation of Current Solutions

Unfortunately, it is impossible with the available data to determine the effectiveness of current methods for dealing with social engineering attacks. It could be reasoned that the general lack of data shows the inadequacy of current methods of detection. For example, if network intrusion detection systems existed, how would one measure the number of attacks? It is clear that social engineering poses serious security threats, but no metrics exist for measuring its impact. Therefore, in addition to enhancing existing solutions and developing new prevention and detection techniques, research efforts must focus on measuring the impact of social engineering attacks.

5. Proposed Solution

This section describes the Social Engineering Defense Architecture (SEDA). SEDA is intended to detect social engineering attacks perpetrated over telephones and to generate logs for forensic investigations. The focus on the telephone medium is crucial as most social engineering attacks are carried out over the phone [4, 6].

SEDA is designed to detect attacks based on intent and deception instead of the attack target. Detecting a social engineering attack based on its target is difficult because social engineers typically first pursue targets with seemingly very little importance as discussed in Section 3. However, this trivial information is then used to obtain more sensitive

and well-guarded information. By detecting lying and deception, SEDA will help prevent social engineering attacks in their early and late stages.

5.1 Attack Detection

The primary purpose of SEDA is to make recipients of phone calls within an organization aware of callers who are attempting to deceive them or obtain unauthorized information. The "muscle" of the system is a text-independent voice signature authentication system. Markowitz defines text-independent verification as "[accepting] any spoken input, making it possible to design unobtrusive, even invisible, verification applications that examine the ongoing speech of an individual" [9]. According to Markowitz, the ability of text-independent technology to operate unobtrusively in the background makes it attractive for customer-related applications, because customers need not pause for security checks before moving on to their primary business objectives. The result is a system of authentication that hinders workflow marginally, if at all.

The voice signatures collected by SEDA will be linked to a database of personal information that includes the employee name, corporate association, job title, and all the phone numbers used to place calls. The types and amount of information gathered would depend on the needs of the organization employing SEDA.

The success of SEDA's strategy lies in the fact that social engineers often masquerade as employees to gain trust. The system would prevent social engineers from claiming to be employees, even if they have all the information to pass as one in a phone conversation. It would also complicate matters for a social engineer who keeps changing his name. The first time the attacker calls, the name he uses is associated with his voice signature; this would require him to modify his voice if he calls again under another name. While this is a way to defeat SEDA, most attackers would be deterred. In any case, attackers who modify their voices would still have to deal with SEDA's other attack detection systems.

The SEDA design also incorporates a voice-to-text engine, which can convert voice conversations into text in real time. Several prototype voice-to-text systems have been developed (see, e.g., [7, 8]). It is important that the voice-to-text conversion be performed rapidly and accurately. If the generated text cannot be sent for analysis fast enough, the attacker could obtain the requested information before the recipient of the call can be notified that an attack is in progress. Also, the voice-to-text conversion must be robust enough to deal with bad phone connections.

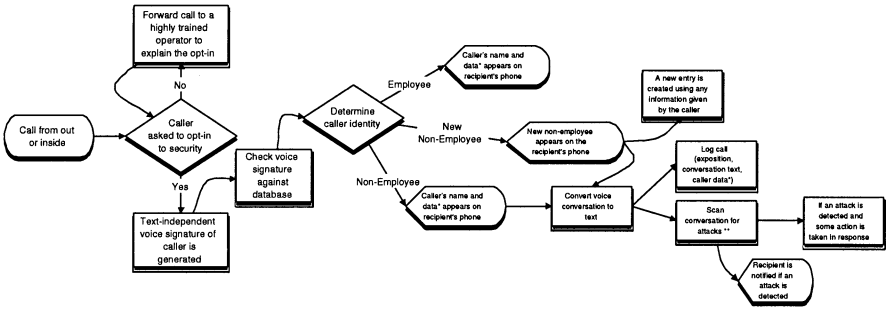

Figure 1. SEDA decision tree.

To support forensic investigations, all conversations originating from outside the company's phone switch should be recorded and the text of the conversations should be linked to the callers' and receivers' voice signatures. Because of the need to record conversations for security purposes, each caller would have to opt-in by calling a specific number when he/she first calls. This removes the expectation of privacy associated with telephone conversations, and ensures that the use of SEDA does not violate wire tap laws [16]. Otherwise, the caller would be transferred to an operator who is trained to resist social engineering attacks; this operator would explain the purpose of the opt-in process to the caller.

A textual conversation analysis tool is the "brain" behind SEDA. Raskin and co-workers [13] have developed a content analysis tool that uses sophisticated natural language processing techniques to determine if a person is lying. While the tool is not yet ready to be incorporated within SEDA, the research shows that the problem of parsing a conversation and determining if someone is lying is tractable. Due to its computational needs, such a tool will have to run on multiple servers to analyze conversations in real time. Nevertheless, it may only be a matter of time before conversation analysis tools are used in SEDA and other applications [13].

A simpler content analysis tool with a narrower scope could monitor conversations for specific strings used in social engineering attacks. These strings are similar to virus signatures. For example, if a caller says, "Please read me your username and password," it is clear that either the caller has malicious intent or he is violating the security policy, neither of which is acceptable. These rules would have to be customized for each organization. Figure 1 presents SEDA's decision tree structure. An expanded view of the attack detection process is shown in Figure 2. It should be noted that no action, aside from notifying the receiver of the call, is taken when an attack is detected. Further research is needed

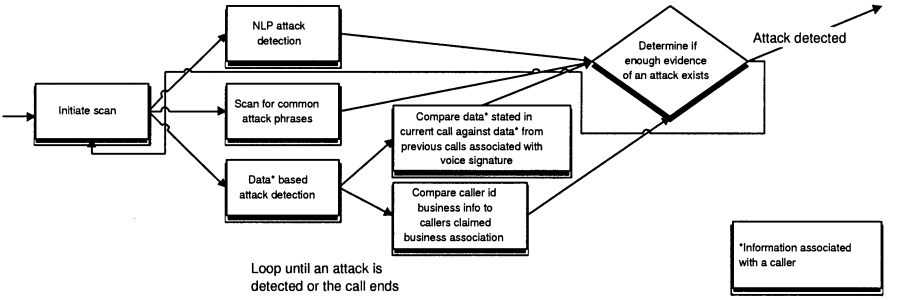

Figure 2. Attack detection process.

to determine the best course of action when a social engineering attack is detected.

5.2 Attack Signature Generation

As with attacks on computer systems, some social engineering attacks will get through no matter what is done to prevent them. In the case of a skilled social engineer breaking into an organization, SEDA provides call logs to perform a forensic analysis of the attack. As noted above, every conversation originating outside the company phone switch is recorded in text format with the voice signatures linked for caller identification. Storing conversations in text format reduces storage needs and permits scanning for clues without having to convert voice into text. Logs of conversations provide forensic investigators with information to trace criminal activity from the final target to the attacker. The logs could also be used in conjunction with other forensic methods to identify attackers.

6. Limitations

One of the major limitations of the proposed solution is its inability to deal with voice modulation. If an attacker were to mask his voice using a modulation device during every call in an attack, the ability to link the calls during a forensic investigation would be greatly decreased. However, voice modulation would have no effect on SEDA's ability to detect deception based on conversation content. Even a resourceful social engineer would not be able to bypass all of the SEDA's levels of protection.

Problems could also arise if SEDA is unable to handle poor telephone connections, e.g., from a cell phone. This situation can be viewed as another form of voice modulation.

A related problem is a replay attack – recording someone's voice and playing it into the phone to obtain authorization, and then continuing

the conversation. Text-independent voice signatures can be generated over the entire call to deal with such attacks. Many commercial speaker-verification systems look for telltale auditory signals, distortions, exact matches, and other indications that a recording has been used [9]. In fact, some voice signature systems already detect such attacks [9].

7. Future Research

This work suggests several avenues for future research. One open problem is handling internal calls. Our solution treats all calls the same regardless of the location of the caller. This could be problematic when there are a large number of internal calls. A technique for streamlining internal calls would significantly reduce the computational requirements.

As mentioned earlier, research is also needed to determine how to handle social engineering attacks after they are detected. Strategies for dealing with attacks would depend on the security goals of the organization. Nevertheless, guidelines for responding to attacks must be created to assist in the development of countermeasures.

Another area for future research is social engineering forensics. Forentic tools must be developed to parse log files and discover clues in logs generated by SEDA. In addition, policies for conducting forensic investigations using SEDA logs must be created. Essentially, much of the research in digital forensics has to be applied to investigating social engineering attacks.

8. Conclusions

The principal advantage of SEDA is that it takes the human element out of determining a person's identity over the phone. Callers will be identified as employees or outsiders; this alone is crucial to preventing social engineering attacks. The ability to detect deception also means that a social engineer will not be able to appeal to someone's emotions or try to bully him or her into performing an action. Another advantage is that the log files that are generated could support forensic investigations of social engineering attacks, which is an interesting new area of research in digital forensics. The main weakness is that voice modulation makes it possible for one person to call many times under different names and not be tracked. While this is a problem that must be addressed by efforts in voice modulation detection, it does not undermine SEDA. Despite the limitations, SEDA addresses two major problems that are so far unanswered: how to detect social engineering attacks and how to perform forensic analyses of social engineering attacks.

References

[1] M. Allen, The use of social engineering as a means of violating computer systems (www.sans.org/rr/catindex.php?cat_id=51).

[2] W. Arthurs, A proactive defense to social engineering (www.sans.org /rr/catindex.php?cat_id=51).

[3] J. Bort, Liar, Liar, *Client Server Computing*, vol. 4(5), 1997.

[4] A. Dolan, Social engineering (www.sans.org/rr/catindex.php?cat_id =51).

[5] L. Gordon, M. Loeb, W. Lucyshyn and R. Richardson, *2004 CSI/FBI Computer Crime and Security Survey* (www.gocsi.com), 2004.

[6] D. Gragg, A multilevel defense against social engineering (www.sans. org/rr/catindex.php?cat_id=51).

[7] C. Karat, C. Halverson, D. Horn and J. Karat, Patterns of entry and correction in large vocabulary continuous speech recognition systems, Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 568-575, 1999.

[8] J. Lai and J. Vergo, MedSpeak: Report creation with continuous speech recognition, Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 431-438, 1997.

[9] J. Markowitz, Voice biometrics, *Communications of the ACM*, vol. 43(9), pp. 66-73, 2000.

[10] The Mentor, *The Hackers Manifesto* (www.geocities.com/Silicon Valley/Heights/1926/mentor.html).

[11] Nemesysco, The Layer Voice Analysis (LVA) technology (www. nemesysco.com/technology-lvavoiceanalysis.html).

[12] R. Power, *2002 CSI/FBI Computer Crime and Security Survey* (www.gocsi.com), 2002.

[13] V. Raskin, F. Christian and K. Triezenberg, Semantic forensics: An application of ontological semantics to information assurance, *Proceedings of the Forty-Second Annual Meeting of the Association for Computational Linguistics*, Barcelona, Spain, 2004.

[14] R. Richardson, *2003 CSI/FBI Computer Crime and Security Survey* (www.gocsi.com), 2003.

[15] M. Rogers and J. Berti, The forgotten risk, in *Information Security Management Handbook, Volume 3*, H. Tipton and M. Krause (Eds.), CRC Press, New York, pp. 51-63, 2002.

[16] U.S. Department of Justice, 18 U.S.C. 2511 – Interception and disclosure of wire, oral or electronic communications prohibited (www.cybercrime.gov/usc2511.htm).