

Two methodologies for physical penetration testing using social engineering

Trajce Dimkov, Wolter Pieters, Pieter Hartel
Distributed and Embedded Security Group
University of Twente, The Netherlands
{trajce.dimkov, wolter.pieters, pieter.hartel}@utwente.nl

Abstract—During a penetration test on the physical security of an organization, if social engineering is used, the penetration tester directly interacts with the employees. These interactions are usually based on deception and if not done properly can upset the employees, violate their privacy or damage their trust towards the organization, leading to law suits and loss of productivity of the organization. This paper proposes two methodologies for performing a physical penetration test where the goal is to gain an asset using social engineering. These methodologies aim to reduce the impact of the penetration test on the employees. The methodologies are validated by a set of penetration tests we did in a period of two years.

Keywords: penetration testing, physical security, methodology, social engineering, research ethics

I. INTRODUCTION

The strength and alignment of an organization's security mechanisms can be accessed through auditing and penetration testing. During an audit, the auditor maps the physical and IT security mechanisms to a standard such as ISO17799 [1], and determines to what extent the security of the organization complies with the standard. When the organization aligns all its security mechanisms to the standard, it usually performs a penetration test. In penetration testing, a penetration tester takes an adversarial role and starts attacking the organization until she achieves a predetermined goal [2, 3]. A successful attack during the test shows the organization that there is misalignment between the security mechanisms.

A penetration test can assess both the IT security and the physical security in the organization. If the penetration tester assesses the IT security, the goal is to obtain marked data located deep in the organizations network. Similarly, if the assessment is on the physical security, then the goal is to obtain a specific asset, such as a laptop or a document. During the execution of a penetration test, if social engineering is allowed, the tester interacts with the employees as means to achieve the goal.

Recently, physical penetration testing gains in importance because of the direct impact of physical theft

to the digital security of the organization. Marshall et al. [4] show that in 2008, 25% of the data loss incidents involved a computer theft, affecting more than 6.2 million people. With the proliferation of laptops and handhold devices this trend will continue, making physical penetration tests where the goal is to obtain an asset increasingly gain value.

A physical test executed in combination with social engineering provides an overview of both the physical security mechanisms and the level of security awareness of the employees. Measuring the resilience of an employee against social engineering in physical penetration test is harder than in digital penetration test. In digital penetration tests the resilience of an employee is measured indirectly, by making phone queries or sending fake mail that lure the employee to disclose secret information. Finn and Jakobsson [5] show that such indirect interaction between the penetration tester and the employee reduces the ethical impact to a minimal level while Soghoian [6] shows how these interactions can be performed while staying within legal boundaries.

However, in physical penetration tests, the tester enters the facility of the organization and directly interacts with the employees. The absence of any digital medium in the communication with the employees makes the interaction between the penetration tester and the employee personal. When entering the premises, the tester either deceives the employee, trying to obtain more information about the goal, or urges the employee to help her, by letting the tester inside a secure area or giving the tester a credential.

There are two main consequences from personal interaction between the tester and the employee. First, the employee might be stressed or not treated with respect. Second, when helping the penetration tester to enter a secure location, the employee loses the trust from the people who reside in the secure location. For example, employees might stop trusting the secretary when they find out she let an intruder inside their office. The penetration test should consider and maintain these trust relationships.

To avoid ethical and legal implications, organizations may avoid physical penetration testing with social engineering. Avoiding this type of penetration test leaves the

This research is supported by the Sentinels program of the Technology Foundation STW, applied science division of NWO and the technology programme of the Ministry of Economic Affairs under projects number TIT.7628.

organization unaware of attacks where the thief uses the help and information from employees to steal an asset.

This paper tackles the problem how to perform a physical penetration test using social engineering in the most respectful manner, while still getting results that can lead to improving the security of the organization.

The contribution of this paper is two methodologies for physical penetration tests using social engineering where the goal is to gain possession of a physical asset from the premises of the organization. The methodologies are validated by performing 14 live penetration tests in the last two years, where students tried to gain possession of marked laptops placed in buildings of two universities in Netherlands.

The rest of the paper is structured as follows. In section 2 we present related work and in section 3 we set the requirements for the methodologies. Sections 4 and 5 outline the methodologies, section 6 provides evaluation of the structure of the methodologies and section 7 concludes the paper.

II. RELATED WORK

In the computer science literature, there are isolated reports of physical penetration tests using social engineering [7, 8]. However, these approaches focus completely on the actions of the penetration tester and do not consider the impact of the test on the employees.

There are a few methodologies for digital penetration testing. The Open-Source Security Testing Methodology Manual (OSSTMM) [9] provides an extensive list of *what* needs to be checked during a physical penetration test. However, the methodology does not state *how* the testing should be carried out. OSSTMM also does not consider direct interaction between the penetration tester and the employees. Barret [10] provides audit-based methodology for social engineering using direct interaction between the penetration tester and an employee. Since this is an audit-based methodology, the goal is to test *all* employees. Our methodologies are goal-based and focus on the security of a specific physical asset. Employees are considered as an additional mechanism which can be circumvented to achieve the goal, instead of being the goal. Türpe and Eichler [11] focus on safety precautions while testing production systems. Since a test can harm the production system, it can cause unforeseeable damages to the organization. In our work the penetration test of the premises of an organization is seen as a test of a production system and few of the mechanisms suggested by Türpe and Eichler are visible in our methodologies. Our approach focuses on the order of executing the mechanisms and their effect on the environment.

In the crime science community, Cornish [12] provides mechanisms how to structure multiple crimes into universal crime scripts and reasons about mechanisms how to prevent them. We adopt a similar reporting format to

present the results from a penetration test. However, instead of using the crime scripts to structure the multiple attacks, we use the scripts to isolate security mechanisms that continuously fail or succeed in stopping an attack.

In social research, the Bellman report [13] defines the ethical guidelines for protection of humans in testing. The first guideline in the report states that all participants should be treated with respect during the test. Finn [14] provides four justifications that need to be satisfied to use deception in research. We use the same justifications to show that our methodology is ethically sound.

III. REQUIREMENTS

A penetration test should satisfy five requirements to be useful for the organization. First, the penetration test needs to be realistic, since it simulates an attack performed by a real adversary. Second, during the test all employees need to be treated with respect [13]. The employees should not be stressed, feel uncomfortable nor be at risk during the penetration test, because they might get disappointed from the organization, become disgruntled or even start legal action. Finally, the penetration test should be repeatable, reliable and reportable [10]. We call these the R* requirements:

Realistic - employees should act normally, as they would in everyday life.

Respectful - the test is done ethically, by respecting the employees and the mutual trust between employees.

Reliable - the penetration test does not cause productive loss of employees.

Repeatable - the same test can be performed several times and the if the environment does not change, the results should be the same.

Reportable - all actions during the test should be logged and the outcome of the test should be in a form that permits a meaningful and actionable document of findings and recommendations.

These are conflicting requirements. For example:

In a realistic test penetration test, it might be necessary to deceive an employee, which is not respectful. The employees should not be aware of the test taking place. The moment the employee is informed about the penetration test, she will stop reacting normally, making the test not realistic. Deceiving employees is clearly not respectful.

In a realistic test, arbitrary employees might be social engineered to achieve the goal, which is unreliable. Freedom of approaching arbitrary people gives opportunity to the penetration tester to spontaneously adapt her strategy during the execution of the test. However, the penetration tester cannot foresee the outcome of these interactions, making the test unreliable.

In a reportable test, all actions of the penetration tester need to be logged, which is unrealistic. The tester should provide evidence of all actions which contribute to achieving the

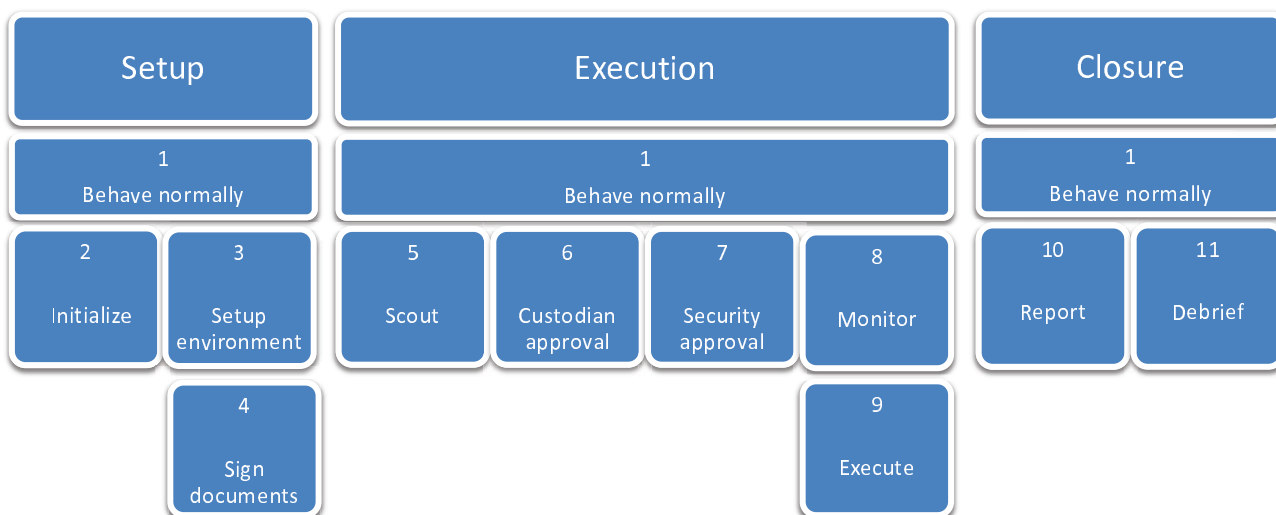


Figure 1. Sequence of events in the first methodology. Each box represents an event which happens in sequence or parallel with other events. For example, event 3 happens after event 2 and in parallel with events 1 and 4.

goal, through recording her actions with logging equipment such as microphones or cameras. This information is used both to repeat the test and to prove the tester was not harassing the employees. Logging equipment conflicts with two requirements. First, the burden of carrying additional devices reduces the set of available attacks to the tester, making the test less realistic. Second, the employee does not know she is recorded without her knowledge, which is not ethical and in some cases illegal [15].

Orchestrating a penetration test is striking the best balance between the conflicting requirements. If the balance is not achieved, the test might either not fully assess the security of the organization or might harm the employees.

IV. FIRST METHODOLOGY

We propose two methodologies for conducting a penetration test using social engineering. In both methodologies the goal of the test is to gain possession of a target asset.

In the first methodology, the owner of the asset is aware that the test takes place. This makes the methodology suitable for tests where the owner is out of scope of the test, such as in tests assessing the security of the laptop belonging to the CEO or the security of equipment in storage areas.

In the following sections, first we define the actors in the first methodology. Then, we introduce all events that take place during the setup, execution and after the penetration test. Finally, we validate the methodology by conducting three penetration tests and present some insights from the experience.




Figure 2. Actors in the first methodology

A. Actors

The penetration test consists of four different authors.

Security officer - an employee responsible for the security of the organization. The security officer orchestrates the penetration test.

Custodian - an employee who owns the assets, sets up and monitors the penetration test.

Penetration tester - an employee or a contractor trying to gain possession of the asset without being caught.

Employee - person in the organization who has none of the roles above.

The actors and the relations between them are shown in Figure 2. The majority of actors treat each other with respect. No respect relation between two actors means either the actors do not interact during the penetration test (for example between the tester and the custodian) or do not have a working relationship (between the penetration tester and the employee). In this methodology, the tester deceives the employee during the penetration test, presented in the figure with a red dashed line.

1. Social engineer might pass from an employee.
2. Enter the building early in the morning.
3. Social engineer the cleaning lady to access the office.
4. Cut any protection on the laptop using a bolt cutter.
5. Leave the building during office hours.

Figure 3. Example of an attack scenario

B. Setup

Figure 1 provides the sequence of events that take place during the setup, execution and closure of the penetration test. During all three stages of the penetration test, employees should behave normally (1 in Figure 1).

As in other penetration testing methodologies, before the start of the test, the security officer sets the scope, the rules of engagement and the goal (2 in Figure 1). The *goal* is gaining physical possession of a marked asset. The *scope* of the testing provides the penetration tester with a set of locations she is allowed to enter, as well as business processes in the organization she can abuse, such as processes for issuing a new password, or processes for adding/removing an employee. The *rules of engagement* restrict the penetration tester to the tools and means she is allowed to use to reach the target. These rules, for example, define if the tester is allowed to force doors, to break windows or to use social engineering.

The custodian first signs an informed consent and then sets up the environment, by marking an asset in her possession and installing monitoring equipment. The asset should not be critical for the daily tasks of the custodian. Thus, when the penetration tester gains possession of the asset, the productivity of the custodian using the asset will not be affected. The custodian leaves the asset in her office or an area without people (storage area, closet). If the custodian shares an office with other employees, the monitoring equipment should be positioned in such a way that it records only the asset and not the nearby employees. The custodian knows when the test takes place, and has sufficient time to remove/obscure all sensitive and private assets in her room and around the marked asset (3 in Figure 1).

Meanwhile, the penetration tester needs to sign the rules of engagement (4 in Figure 1). The rules of engagement should include the following:

- 1) Each attack scenario needs to be approved by the custodian and the security officer.
- 2) The penetration tester should not harm nor intimidate any employee.
- 3) The penetration tester should not damage any property.

C. Execution

The execution stage of the penetration test depends mainly on the skills and experience of the penetration tester. The security officer should choose a trustworthy

penetration tester and monitor her actions during the execution stage.

When the penetration test starts, the tester first scouts the area and proposes a set of attack scenarios (5 in Figure 1). An example of an attack scenario is presented in Figure 3. The proposed attack scenarios need to be approved first by the custodian (6 in Figure 1) and then by the security officer (7 in Figure 1). The custodian is directly involved in the test and can correctly judge the effect of the scenario on her daily tasks and the tasks of her colleagues. The security officer needs to approve the scenarios because she is aware of the general security of the organization and can better predict the far-reaching consequences of the actions of the tester.

After approval from the custodian and the security officer, the tester starts with the execution of the attack scenarios (8 in Figure 1). The custodian and the security officer remotely monitor the execution (9 in Figure 1) through CCTV and the monitoring equipment installed by the custodian.

The penetration tester needs to install wearable monitoring equipment to log her actions. The logs serve three purposes. First, they ensure that if an employee is treated with disrespect there is objective evidence. Second, the logs prove that the penetration tester has followed the attack scenarios, and finally, the logs provide information how the mechanisms were circumvented, helping the organization repeat the scenario if needed.

D. Closure

After the end of the test, the penetration tester prepares a report containing a list of attack traces. Each attack trace contains information of successful or unsuccessful attack (10 in Figure 1). Based on the report, the security officer debriefs both the custodians and the deceived employees (11 in Figure 1).

Reporting. The attack traces are structured in a report that emphasizes the weak and the strong security mechanisms encountered during the penetration test. In the crime science community, there are 25 techniques for situational crime prevention [16]. For different domains there are extensive lists of security mechanisms to enforce the 25 techniques (for example, [17]). The combination of the attack traces together with the situational crime prevention techniques gives an overview of the circumvented mechanisms [18].

Figure 4 shows an example of a generic script instantiated with an attack trace. First we define the generic script, which encompasses the stages of all attacks. In the example, they are: enter the building, enter the office, identify and get the laptop, and exit the building. For each step in a trace, we identify both the mechanisms (if any) that were circumvented and mechanisms that stopped an attack. For failed attacks, the table shows which mechanisms were circumvented up to the failed

Generic Script	Attack trace	Circumvented mechanisms	Successful mechanisms	Recommendations
Prepare for the attack	Buy a bolt cutter and hide it in a bag. Scout the building and the office during working hours. Obtain an after working hours access card.	Access control of the building entrances during working hours. Credential sharing policy.		Keep entrance doors to the building locked at all time. Provide an awareness training concerning credential sharing.
Enter the building	Enter the building at 7:30 AM, before working hours. Hide the face from CCTV at the entrance using a hat.	CCTV pre-theft surveillance.		Increase the awareness of the security guards during non-working hours.
Enter the office	Wait for the cleaning lady. Pretend you are an employee who forgot the office key and ask the cleaning lady to open the office for you.	Challenge unknown people to provide ID. Credential sharing policy.		Reward employees for discovering intruders.
Identify and get the laptop	Search for the specific laptop. Get the bolt cutter from the bag and cut the Kensington lock. Put the laptop and the bolt cutter in the bag.	A.1 Kensington lock.		Get stronger Kensington locks. Use alternative mechanism for protecting the laptop.
Leave the building with the laptop	Leave the building at 8:00, when external doors automatically unlock for employees.	CCTV surveillance. Access control of the building entrances during working hours.		The motion detection of the CCTV cameras needs to be more sensitive .

Figure 4. Reporting the results.

action, and the mechanism that successfully stopped the attempt.

After structuring all attack scenarios in this format, for each generic step we get a list of the mechanisms which were commonly circumvented, and the mechanisms which stopped the most attacks.

Debriefing the employees and the custodian. The interaction between the penetration tester and the employee might be a stressful event, in which the penetration tester makes the employee bend or violate a policy of the organization. After finding they were deceived by the same organization they work for, the employees might get disappointed or disgruntled. At the end of the test the security officer fully debriefs the custodian and the employees. The debriefing should be done carefully, to maintain or restore the trust between custodian and the employees who helped the tester to gain the asset.

E. Validation

We orchestrated three penetration tests using the first methodology. As penetration testers, we used three teams of master students who were instructed to steal three laptops from the custodian (the first author). We locked the laptops with Kensington locks and hid the keys in an office desk. To monitor the laptops, we installed motion detection web cameras which streamed live feeds to an Internet server. Since the custodian shares the office with four other colleagues, the cameras were positioned in such a way to preserve the privacy of the colleagues. We told the colleagues we are doing an experiment, but we did not reveal the nature nor the

goal of the experiment.

Since we knew about the penetration test, we did not allow the students to gain possession of the laptops in our presence. During the experiment, we carried on the normal work, thus the students were forced to carry on the attacks after working hours or during the lunch break.

The three teams first scouted the building and wrote a list of attack scenarios they want to execute. At the beginning the students tried to gain the laptops through the secretary and the colleagues from the office. However, none of them agreed to hand in a laptop without permission from the custodian. In desperation, one of the teams asked the security guard of the building. The strategy was to directly go to the security guard and provide a fake e-mail stating the students are authorized to pick up the laptop. The guard believed the story and opened the door for the students. The rest of the teams followed the same strategy and social engineered the same guard. At the end of the test, all three teams successfully obtained the target laptops.

Later, all teams wrote the successful and unsuccessful attempts in the format shown in Figure 4. After the penetration test, we individually debriefed the security officer, the security guard, the secretary and the colleagues.

F. Lessons learned from the penetration tests

The analysis of the actions that took place before, during and after the penetration test, lead us to the following observations:

1) *The attack scenarios should be flexible*: Although the students provided scenarios prior to all attacks, in all cases they were forced to deviate from them, because the target employee was either not present or was not behaving as expected. Attack scenarios assure the custodian and the security officer that the actions of the penetration tester are in the scope of the test, but at the same time they need to be considered more as a framework than as a specific sequence of steps.

2) *The methodology does not respect the trust relationship between the custodian and the employees*: After the penetration test, the custodian knows which employees were deceived, and the trust relationship between them is shaken. For example, if the secretary lets the penetration tester inside the office of the custodian, the custodian might not be able to trust her again. We concluded that in some instances it would be better to hide from the custodian who contributed to the asset loss. We address this problem in the second methodology.

3) *During the penetration test, separating the custodian from the employees is hard*: Whenever the students approached a colleague from the office, the first reaction of the colleague was to call the custodian and ask for guidance. This led to uncomfortable situations where we were forced to shut down our phones and ignore e-mails while outside the office.

4) *Debriefing proved to be very difficult*: After the test, we fully disclosed the test to all involved employees. Debriefing the security guard who opened the office for the penetration testers three times was the hardest. During the debriefing we focussed on the benefits of the penetration test to the university and their help setting up the test. After the debriefing, we concluded that we caused more stress to the guard during the debriefing than the students had caused during the penetration test.

5) *Forbidding the penetration testers to damage property is too restricting*: For example, the price of a Kensington lock is small. In a real scenario, the thief would cut the lock instead of trying to locate the correct key and social engineer it from an employee.

V. SECOND METHODOLOGY

In the first methodology, the custodian is aware of the penetration test. The knowledge of the penetration test changes her normal behavior and thus influences the results of the test. Since the asset belongs to the custodian, and the asset is in the office of the custodian, in many environments it is desirable to include the custodian's resistance to social engineering as part of the test.

After performing the first series of penetration tests, we revisited and expanded the first methodology. The second methodology can be seen as a refinement of the first methodology, based on the experience from the first set of penetration tests. In the second methodology the custodian is not aware of the test, making the second

methodology suitable for penetration tests where the goal is to check the overall security of an area *including* the level of security awareness of the custodian.




Figure 6. Actors in the second methodology

A. Actors

There are six actors in the second methodology.

Security officer - an employee responsible for the security of the organization.

Coordinator - an employee or contractor responsible for the experiment and the behavior of the penetration tester. The coordinator orchestrates the whole penetration test.

Penetration tester - an employee or contractor who attempts to gain possession of the asset without being caught.

Contact person - an employee who provides logistic support in the organization and a person to be contacted in case of an emergency.

Custodian - an employee at whose office the asset resides. The custodian should behave normally during the whole penetration test (1 in Figure 5).

Employee - person in the organization who has none of the roles above. The employee should also behave normally during the penetration test (2 in Figure 5).

Figure 6 shows the actors and the relations between them. In this methodology, the penetration tester deceives both, the employees and the custodian. Moreover, the contact person also needs to deceive the custodian. These relations are discussed in greater depth in section VI.

B. Setup

At the beginning, similar to the first methodology, the security officer initializes the test by defining the target, scope and the rules of engagement. The security officer at this point assigns a coordinator for the penetration test and provides the coordinator with marked assets and equipment for monitoring the assets (3 in Figure 5). The marked assets should be similar to the asset of interest for which the security is measured. The monitoring equipment should be non-intrusive and its purpose is to

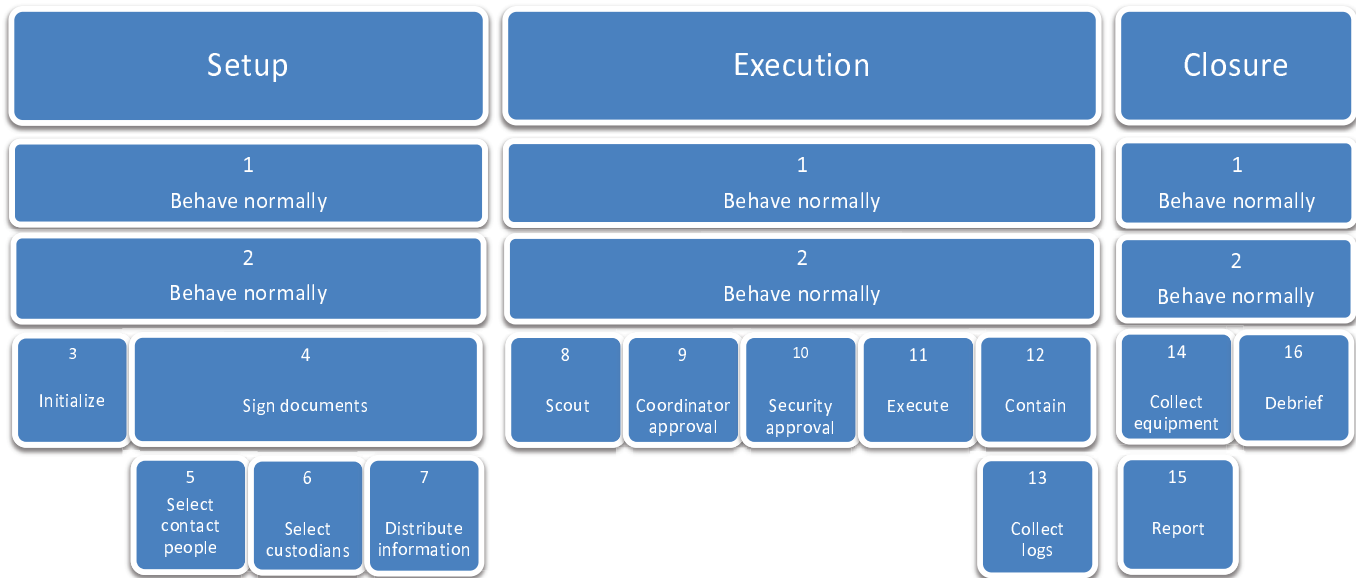


Figure 5. Sequence of events in the second methodology

have an additional information on the activities of the penetration tester.

The rules of engagement are similar to the first methodology. The difference is in the third rule which forbids the penetration tester to damage property:

The penetration tester *can damage* property of the organization up to a certain threshold.

The threshold value defines to which extent the penetration tester can use force in circumventing specific mechanisms. The damage is localized only to property of the organization. An additional rule is:

The penetration tester should immediately abort the test, whenever an abortion condition is met.

The abortion conditions are defined by the security officer. The penetration tester should sign these rules of engagement before the start of the execution stage (4 in Figure 5).

The coordinator selects a number of contact people and provides them with the marked assets and the monitoring equipment (5 in Figure 5). Furthermore, the coordinator provides a cover story which explains why the custodian is given the asset. The contact person selects a number of custodians based on the requirements from the security officer (random, specific roles, specific characteristics) and distributes the marked assets and the monitoring equipment to the custodians. After giving the monitoring equipment, the contact person should get a signed informed consent from the custodians (6 in Figure 5). If the asset can store data, the document must clearly state that the custodian should not store any sensitive nor private data in the asset.

Before the penetration test starts, the coordinator dis-

tributes a list of penetration testers to the security officer, and a list of asset locations to the penetration tester (7 in Figure 5).

C. Execution

The first steps of the execution stage are similar to the previous methodology. The penetration tester scouts the area and proposes attack scenarios (8 in Figure 5). The coordinator and later the security officer should agree with these scenarios before the tester starts executing them (9 and 10 in Figure 5). After approval from both actors, the tester starts executing the attack scenarios. If a penetration tester is caught or a termination condition is reached, the penetration tester immediately informs the contact person. Thus, if the custodian stored sensitive data in the asset, the data is not exposed.

When the penetration tester gains possession of the target asset, the tester informs the contact person and the coordinator and returns the asset to the contact person (11 in Figure 5). The contact person collects the monitoring equipment and informs the security officer (12 in Figure 5). If the tester gains possession of the asset without the knowledge of the custodian, the contact person needs to reach the custodian before the custodian reaches the office and explain to the custodian that the test is terminated. The security officer obtains surveillance videos from the CCTV and access logs and gives them to the coordinator (13 in Figure 5).

D. Closure

After the execution stage, the penetration tester writes a report of all attempts, both failed and successful, in the form of attack traces and gives them to the coordinator (14 in Figure 5). The coordinator has two tasks. First, she

collects the marked assets and monitoring equipment from the contact person (15 in Figure 5) and returns them to the security officer. Second, the coordinator debriefs the security officer and the custodians and provides the custodian a form of reward for helping in the assessment (16 in Figure 5).

Not all employees that were social engineered should be debriefed. Employees who were treated with respect and to whom the penetration tester did not cause discomfort during the interaction should not be debriefed, because the debriefing can cause more stress than the interaction with the penetration tester. The decision which employees need to be debriefed lies with the security officer, and is based on the logs from the penetration tester and the monitoring equipment.

Debriefing the custodians. All custodians should be debriefed, because they sign an informed consent at the beginning of the test. However, to preserve the trust between the custodian and the employees, the custodian should not know which employee contributed to the attack.

Three elements should be considered before the debriefing. First, the custodians were deceived by the organization they work for (more specifically, by the contact person). Second, in case of direct interaction, their privacy might be violated by the logging equipment from the tester. Third, they might be stressed from the penetration test either directly, through interaction with the penetration tester, or indirectly, by finding their asset is gone before the contact person reaches them.

The debriefing should focus on the contribution of the custodian in finding the security vulnerabilities in the organization, and the custodian should be rewarded for the participation.

E. Validation

We orchestrated eleven penetration tests with the second methodology. As targets we used eleven marked laptops that just came out on the market. We selected three PhD students (friends of the author) as contact persons, who chose eleven friends as custodians. The custodians were ten PhD students and one professor spread across two universities who worked in nine different buildings.

We gave the laptops to the custodians for two weeks usage. The cover story was that the universities are interested in the usability of the new generation of laptops. Thus, we need to monitor the behavior of the custodian and measure their satisfaction level. The contact persons explained that they cannot tell the custodians exactly which behavior we measure, since it might change the results of the experiment. We instructed the custodians to lock the laptops with Kensington lock and to keep them in the office. To reduce the risk of data leakage and loss of productivity, we asked the custodians not to store any private or work data on the laptops. Later, thirty-three

master students in computer security, divided in eleven teams of three, tried to gain possession of the laptops by using physical and social means.

The favorite approach of the students was to directly confront the custodian and ask for the laptop. The students took roles as service desk employees, students that urgently need a laptop for a few hours or claimed they were sent by the coordinator. The students used mobile phones and pocket video cameras to record the conversation with the employees. In one case they took a professional camera and a cameraman, and told the custodian the recording is part of a study to measure the service quality of the service desk.

The resistance of the employees varied. In five cases, the employees gave the laptop easily after being showed fake email and being promised they will get the laptop back in a few hours. In two cases the custodian wanted a confirmation from a supervisor or the coordinator. The students succeeded in the attempt because the custodian called a number provided by the students. Needles to say, the number was of another team member disguised as the coordinator. In one case a colleague of the custodian got suspicious and sent an email to the campus security. Since only the main security officer knew about the penetration test, in few hours the security guards were all alerted and started searching for suspicious students.

However, in two cases the students were not able to social engineer the custodian directly and were forced to look for alternative approaches. For example, in one of the cases the students entered the building before working hours. At this time the cleaning lady cleans the offices, and under the assumption it is their office let the students inside. After entering the office, the students cut the Kensington lock and left the building before the custodian arrived.

We debriefed only the custodians through a group presentation, where we explained the penetration test and its goal. During all penetration tests, there were no incidents and none of the employees complained about stressful situations.

F. Lessons learned from the validation

1) *We should specify in advance which information can be used by the penetration tester:* For example, the penetration tester should not use knowledge about the cover story used by the contact person. During the validation, six penetration testers used knowledge of the cover story to convince the custodian to hand in the laptop. Thus, these tests were less realistic.

2) *Panic situations need to be taken into consideration in the termination conditions:* Several times the custodian or an employee got suspicious and raised an alarm. Since only the security officer knew about the experiment, and the other security personal was excluded, news of people stealing laptops spread in a matter of hours.

3) *The penetration test cannot be repeated many times:* If a custodian participated in the penetration test once, she knows what will happen. The same holds for the employees she told about the experiments and the employees that were socially engineered.

4) *The custodian does not behave the same toward its own asset and the acquired asset, making the test less realistic:* During the penetration test, we noticed the security awareness of the custodians is low towards the marked asset than to assets she owns for a longer period of time. We assume the custodian does not consider the asset as her belonging and is more willing to give it away. If the marked asset is given a few months before the penetration test starts, we believe the custodian will care of the asset as if the asset belongs to her.

VI. EVALUATION

There are two main elements that define to which extent the penetration test satisfies the requirements: the rules of engagement and the pre-approval of each attack scenario. Less restrictive rules of engagement and approving more invasive attack scenarios make the penetration test more realistic, but make the test less reliable and respectful to the employees. The evaluation below assumes these elements are tuned to the risk appetite of the organization and focuses only on the structure of the methodologies.

Reliable: In the first methodology, the penetration tester gains possession of a non-critical asset which the custodian is prepared to lose. Thus, the result of the penetration test will not affect the productivity of the custodian. In the second methodology, the productivity of the custodian can be affected since the custodian does not know the asset will be stolen. The informed consent she signs, explicitly states not to use the marked asset for daily tasks nor store sensitive information on the asset. The productivity of other employees is not affected, since the penetration tester does not gain possession of any of their belongings without their approval.

Repeatable: The repeatability of any penetration test using social engineering is questionable, since human behavior is unpredictable. Checking if a penetration test is repeatable would require a larger set of tests on a single participant, and a larger number of participants in the test.

Reportable: The approach used in reporting the results of the penetration test completely covers all information needed to perform the attack in a real-life situation and provides an overview of what should be improved to thwart such attempts. The logs from the tester and the monitoring equipment installed by the custodians provide detailed information on all actions taken by the penetration tester, giving a clear overview of how the mechanisms are circumvented.

Respectful: Both methodologies should respect all the employees and the trust relationships between them.

	Methodology I	Methodology II
Reliable	+++	++
Repeatable	-	-
Reportable	+++	+++
Respectful: actors	++	+
Respectful: trust relations	-	++
Realistic	+	+++

Figure 7. Evaluation of both methodologies

In physical penetration testing, the social engineering element is much more intense than in digital penetration testing because the interaction between the penetration tester and the employee is direct, without using any digital medium. Baumrind [19] considers deception of subjects in testing as unethical. The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, clearly states in their first rule of ethical principles: "Respect for persons" [13].

However, in some tests deception cannot be avoided, and Finn [14] defines four justifications that need to be met to make deception acceptable: (1) The assessment cannot be performed without the use of deception. (2) The knowledge obtained from the assessment has important value. (3) The test involves no more than minimal risk and does not violate the rights and the welfare of the individual. Minimal risk is defined as: "the probability and magnitude of physical or psychological harm that is normally encountered in the daily lives" [20]. (4) Where appropriate, the subjects are provided with relevant information about the assessment after participating in the test. Physical penetration testing using social engineering can never be completely respectful because it is based on deception. However, the deception in both methodologies presented in this paper is justifiable.

The first two justifications are general for penetration testing and its benefits, and have been discussed earlier in the literature (for example, [10]).

The third justification states that the risk induced by the test should be no greater than the risks we face in daily lives. In the first methodology, the only actor at risk is the employee. The penetration tester cannot physically harm the employee because of the rules of engagement, thus only psychological harm is possible. If the employees help the penetration tester voluntarily, the risk of psychological harm is minimal. The logging equipment assures the interaction can be audited in a case of dispute. In the second methodology, an additional actor at risk is the custodian. The only case when the risk is above minimal for the custodian is if the tester gains possession of the laptop without custodian's knowledge. When the custodian finds the laptop missing, her stress level might increase. Therefore it is crucial for the contact person to reach the custodian before custodian gets to her office.

The fourth justification states that all actors should be debriefed after the exercise. In both methodologies, all actors except the employees are either fully aware of the exercise, or have signed an informed consent and are debriefed after the exercise. Similarly to Fin and Jakobsson [5], we argue that there should be selective debriefing of the employees. The risk for the employees during the penetration test is minimal. Debriefing can make the employee upset and disgruntled and is the only event where the risk is higher than minimal. Thus, an employee should be debriefed only if the security officer constitutes the tester did more than minimal harm.

Besides being respectful towards the participants, the methodology needs to maintain the trust relations between the employees. The first methodology affects the trust between the custodian and the employees and the employees and the organization. This is a consequence of the decision to fully debrief all participants in the test. The second methodology looks at reducing these impacts. First, the custodians are not told who contributed to the attack. Only the coordinator and the security officer have this information, and they are not related to the custodian. Second, the employees are not informed about the penetration test unless it deemed necessary. However, the trust between the custodian and the contact person is shaken. Therefore, the contact person and the custodian should not know each other prior to the test.

In conclusion, the second methodology is less respectful to the custodian than the first methodology, because the custodian is deceived and might get stressed when she finds out the asset is gone. The first methodology does not preserve any trust between the employees, the organization and the custodian. The second methodology preserves the trust bond between the custodian and the employees and between the employees and the organization. However, the trust bond between the custodian and the contact person may be affected.

Realistic: The first methodology allows testing the resilience to social engineering of employees in the organization. Since the custodian knows about the penetration test, she is not directly involved during the execution of the test, making this methodology implementable in limited number of situations.

In the second methodology, neither the custodian nor any of the other employees know about the penetration test, making the test more realistic.

VII. CONCLUSION

We present two methodologies for penetration testing using social engineering. The second methodology was developed after validating the first methodology, thus in many aspects it improves on it. However, the first methodology is more reliable, does not deceive the custodian and fully debriefs all actors in the test.

We present here both methodologies, leaving to the organization to decide which one is more appropriate for its environment.

This paper shows that performing physical penetration tests using social engineering can reduce the impact on employees in the organization, and still provide meaningful and useful information on the security of the organization. We validate the methodologies by performing 14 tests and evaluate them using five requirements.

REFERENCES

- [1] Information technology - code of practice for information security management: International Organization for Standardization (ISO), 2000.
- [2] M. Bishop. About penetration testing. *IEEE Security and Privacy*, 5(6):84–87, 2007.
- [3] D. Geer and J. Harthorne. Penetration testing: a duet. In *Proceedings of 18th Annual Computer Security Applications Conference*, pages 185–195. IEEE, 2002.
- [4] M. Marshall, M. Martindale, R. Leaning, and D. Das. *Data Loss Barometer*. Sep 2008.
- [5] P. Finn and M. Jakobsson. Designing ethical phishing experiments. *Technology and Society Magazine, IEEE*, 26(1):46–58, Spring 2007.
- [6] C. Soghoian. Legal risks for phishing researchers. In *eCrime Researchers Summit, 2008*, pages 1–11, Oct. 2008.
- [7] C. Greenlees. An intruder’s tale-[it security]. *Engineering & Technology*, 4(13):55–57, 2009.
- [8] Wil Allsopp. *Unauthorised Access: Physical Penetration Testing For IT Security Teams*, chapter Planning your physical penetration test, pages 11–28. Wiley, 2009.
- [9] P. Herzog. OSSTMM 2.2–Open Source Security Testing Methodology Manual. *Open source document*, www.isecom.org/osstmm, 2006.
- [10] N. Barrett. Penetration testing and social engineering hacking the weakest link. *Information Security Technical Report*, 8(4):56–64, 2003.
- [11] S. Türpe and J. Eichler. Testing production systems safely: Common precautions in penetration testing. In *Proceedings of Testing: Academic and Industrial Conference (TAIC PART 2009)*. IEEE Computer Society, 2009.
- [12] D. B. Cornish. The procedural analysis of offending and its relevance for situational prevention. In R. V. Clarke, editor, *Crime Prevention Studies*, volume 3, pages 151–196. Criminal Justice Press, Monsey, NY, 1994.
- [13] National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. The Belmont report: Ethical principles and guidelines for the protection of human subjects of research. pages 1–18, 1978.
- [14] P.R. Finn. *Research Ethics: Cases and Materials*, chap-

ter The ethics of deception in research, pages 87–118. Indiana University Press, 1995.

- [15] Hernandez v. Hillside Inc. S147552. Supreme Court of California, 2009.
- [16] D.B. Cornish and R.V. Clarke. Opportunities, precipitators and criminal decisions: A reply to Wortley's critique of situational crime prevention. *Crime Prevention Studies*, 16:41–96, 2003.
- [17] G. Kitteringham. Lost laptops = lost data: Measuring costs, managing threats. Crisp report, ASIS International Foundation, 2008.
- [18] R. Willison and M. Siponen. Overcoming the insider: reducing employee computer crime through situational crime prevention. *Communications of the ACM*, 52(9):133–137, 2009.
- [19] D. Baumrind. Research using intentional deception. Ethical issues revisited. *The American psychologist*, 40(2):165–174, 1985.
- [20] Code of Federal Regulations. Title 45: Public welfare department of health and human services. part 46: Protection of human subjects. pages 1–12. 2005.