Social engineering: the art of deception

Matthieu Paques

In a typical penetration test (hacker test) attempts are made to gain unauthorized access to systems or data by exploiting technical vulnerabilities. The "weakest link" in the information security chain is often overlooked in these tests: *users*. It appears that this "link" has increasingly become the target of attackers. The media have reported a large number of incidents involving this type of attack ([security.nl]). This is reason enough to also put this "link" to the test within the scope of an audit or security test. This act of "hacking people" is called "social engineering". This article describes how social engineering tests are performed, provides some real-life examples, and discusses what measures can be taken against such attacks.

M.B. Paques
is a manager at
KPMG IT Advisory.
paques.matthieu@kpmg.nl

What is a social engineering test?

KPMG IT Advisory has performed social engineering assignments for a large number of clients. The purpose of such tests is twofold:

- identify the risks to the organization being evaluated
- make employees aware of these risks (training)

During the tests, attempts are made to manipulate employees so that unauthorized access to confidential information is obtained. These attempts vary from a simple "phone call test" in which employees are tricked into disclosing

passwords or a so-called phishing attack (in which the attacker uses forged emails and/or websites), to a physical attack where a client's premises are entered by a tester undercover using counterfeited access badges (or sometimes disguised like a pizza delivery person or fireman) to gather confidential information from the inside. The findings are usually quite remarkable. To name just a few, unauthorized access has been gained to safes in banks, heavily secured government areas and large data centers. In several of these cases, the assignment also included a penetration test. In these combined tests, also known as *red teaming* (Figure 1), the team first has to gain unauthorized physical access to the building and then has to hack internal systems and eventually leave with confidential information without being caught.

The main difference between a penetration test where you can attempt to access systems multiple times and a social engineering test is that in the latter the tester usually has only one chance of success. There are no "try-outs", it must be successful the very first time. The tester has to be prepared for unforeseen situations and must have a made-up story (the *pretext*) ready in case his presence is being questioned. If his story is not credible, there is a risk of being taken away in handcuffs. The employees of the organization for which the test is performed are generally not informed in advance about the test. Often, only a few executives are aware of the test and even they do not know exactly when the test will be carried out. Security staff is not meant to be put on the alert and take extra precautions. This approach makes it possible to obtain a realistic impression of the risks. As a result of this approach security personnel may take drastic measures if the tester is unmasked as an "intruder" (especially when he has a stack of confidential documents in his possession).

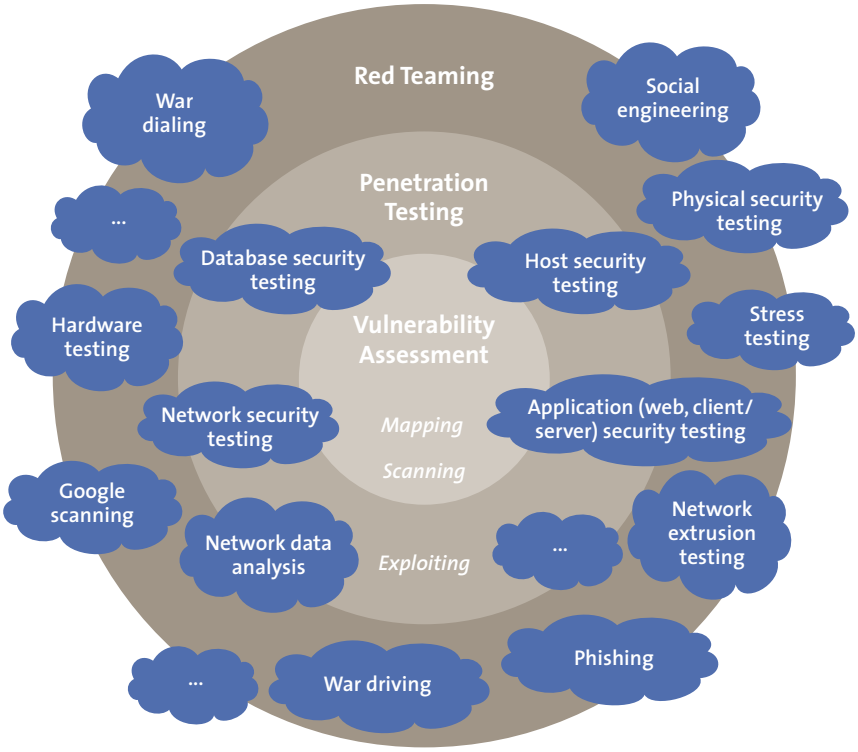

Figure 1. **Red teaming is a test approach where different attack techniques are combined to simulate an actual attack.**

The ingredients of a successful attack

There are two decisive factors that determine the success of a social engineering attack: *information* and *timing*. Thorough preparation is crucial. In such a test as much *information* as possible is assembled about the target prior to the actual attack. About 90% of the time is spent to research and make preparations for the actual *hit*. Information is gathered not only about the organization in scope (e.g., via the corporate homepage, Google Maps, search engines, newsgroups or job vacancy websites,) but also about the organization's employees, their hobbies, address and contact info (Facebook, Hyves, LinkedIn etc are very useful). After this step the tester usually makes several telephone calls to the company's general telephone number and the phone numbers of employees found through publicly available sources. Large organizations often use series of telephone numbers. The known numbers in a series allow for other numbers in the series to be determined and called. When an employee answers, they are told that it must be a wrong number as it is Mister X who is needed (Mister X being a name that was found in the earlier research, for example on LinkedIn). The correct number of Mister X and their name, job title, and department are then verified. Information obtained in this manner can then be used to extract further information. All information that is gained is potentially interesting. For a test on a highly secure data center we went there months before the test and photographed the building from all sides with a camera with a 500mm lens to determine the location of all the cameras and entrances, to observe how employees were dressed, what time they went home, etc. Information like this is used for elaborating a detailed attack scenario. We determine who we will impersonate, what time we must arrive (to walk in with the rush of the daily crowd), the best clothing to wear, and what route to take once we are "in" to avoid as many risks as possible (e.g., cameras and security guards).

The *timing* of an attack is also very important. Often, the help of an employee is required to get past a gate, fence, reception or other secured entrance. The exact moment that a suitable employee is present may be a matter of seconds. With a good story, improvisation skills for unanticipated situations, the ability to make contact easily and sometimes nerves of steel an attacker might even be able to penetrate the most secure environments.

During an attack it is useful to know what people to approach and who to avoid. For example, secretaries often know a lot about what is happening in a company. Their knowledge can be of tremendous value. However, because they know a lot about what is happening in the company, a good story that is well supported is a prerequisite when you approach them. Complete improvisation may be like a game of Russian roulette and result in a premature and undesired end of the test. Case study 1 describes a test case in which the individuals approached were specifically selected to make the chance of success as high as possible.

Case study 1

A colleague and I carried out an advanced phishing attack on one of our clients. My colleague placed himself at the entrance of the client office building and selectively asked employees who entered whether they wanted to take part in a survey about the upcoming Christmas activity. We focused our "selection" of employees on the younger female employees to minimize the risk of accidentally speaking with managers or IT staff. (They would know whether such a survey existed and thus figure out quite quickly that there was an attack underway.) Beforehand, we had examined the LinkedIn and Facebook profiles of key people in the organization so we could recognize and avoid these "risky people".

Participants would be included in a raffle for an iPod Touch. The employees who wanted to participate were given a sealed envelope containing a letter explaining the activity and a link to our forged web page with the survey that we set up beforehand. After logging in with their credentials, the employees were presented with ten questions about their ideas for the perfect Christmas activity. They could also supplement these with their own suggestions. After submitting their responses, they were thanked for their participation. Of course, we were not at all interested in the employees' "party ideas", but just in their login details. I had taken position around the corner to keep an eye through the window to see whether anything suspect happened inside. If it became necessary, I could warn my colleague via our two-way radio transceiver and inform him that it was time to take to his heels. At the same time I watched my smartphone that provided "real-time" updates on the number of users that logged in on the web page. In a matter of minutes several users had entered their passwords on our web page already. After about 35 minutes, we both left the location in different directions. We estimated that this was the minimum amount of time it would take to be detected. In the discussion with the client afterwards we discovered that only a few minutes passed after we left until two alarmed people came outside to demand an explanation.

Employee training

An important aspect of a social engineering test is *to make the employees aware of the risks*. Nevertheless, the attack scenarios should be selected in such a way that the impact experienced by employees is kept to the absolute minimum required. Therefore, we do not give the client any details (insofar as possible) about which employees played a role in the tests (for example, which employees did provide their password). Details are anonimized as much as possible. The least sensible thing a client can do (and, of course, highly undesirable, but not inconceivable) is taking disciplinary action against these employees. The outcome of such an action is that employees who do become "victims" of a *real* social engineering attack may not report it in fear of reprisals and the organization does not become aware of the attack until it is too late and it has to deal with the consequences.

A good follow-up to a social engineering test is to present the results back to all employees so that the test can be a learning experience and they are better prepared against a real attack. We experience that in practice, most untrained employees are susceptible to a social engineering attack and employees can be misled at every level in the organization.

Psychological tricks

For each test the attack scenario is completely different because it is tailored to the client's specific circumstances. Nonetheless, some fundamental psychological principles or "tricks" are regularly used:

- *Making a personal connection:* mentioning a common problem or interest is typical. Social media can be a valuable source of information. Indicating that you have worked for the same company or play the same sport builds trust. You can also say you have a friend or acquaintance in common. After the connection is made, it is harder for the "victim" to refuse a request.
- *Time pressure:* create a situation where the "victim" does not have enough time to make a proper decision because circumstances are described in such a way that a quick decision must be made. The Windows operating system often shows the name of the last user that logged in (but not the password). Sitting at a user's (locked) PC, you can usually block that user's account by entering the wrong password five times. After blocking the account, you can call the help desk and say that you must give an important presentation *within five minutes* and need to get into your blocked account. Due to the time pressure the help desk employee (after checking that the account

Figure 2. **Security badge costing a few dollars that a social engineer can use to "exude" authority.**

is actually blocked) may issue a temporary password and give it over the telephone. Now, you have access to the system.

- Referring to a *senior person in the organization* (authority). This trick often works very effectively combined with the "time pressure" element. Indicate that the "victim" is hindering the actions of a high ranking person in the organization and that the victim must immediately assist with the request. A variation of this is using clothing and accessories that "exude" authority (see also Figure 2). Wearing a suit and tie makes it sometimes much easier to get into a building without being questioned than wearing jeans and a T-shirt. I once entered a bank in a soaking construction worker's jacket announcing that there was a leak on the floor above. I said something like: "I just want to take a quick look to see if any water is coming through the ceiling." The staff were happy that they had been warned in time and without asking questions allowed me access to the restricted areas in the building that should only be accessible by bank staff.
- *Asking for help:* for example, ask someone to print a file from a USB memory stick that is infected with malware that infects the pc of the victim as soon as the file on the stick is accessed, or borrow an access badge because "you left yours on your desk". A request made by a man (the tester) to a woman (the victim) and vice versa is usually fulfilled easier than when the gender is the same.
- Using *recognizable* items related to the organization that is being evaluated. Employees may believe they are dealing with a co-worker because you have an access badge (possibly forged), similar style of clothing, business cards, jargon, knowledge of work methods or names of information systems or colleagues (name dropping). All are less likely to prompt critical questions. If the name on the (fake) badge also has a LinkedIn or Facebook profile that refers to the company being evaluated, even the most suspicious people may be convinced that they are dealing with a co-worker.
- Another method is to *request one employee to give information to another employee* (for example, communicate with an internal department to have them forward a "wrongly addressed email"). Using these internal reference points increases credibility. Another example is recording the hold music that companies use when callers are put on hold. You can call an employee, then say after a few minutes: "Wait a minute please, I have to get the other line". You then put the "victim" on hold and play the hold music that you previously recorded making the victim unconsciously think: "Hey, that's our music, he must work for our company".
- Indicating that *all colleagues of the "victim" have acted the same way* so that it makes the request seem completely normal. People are inclined to believe something is correct when others have made the same choice. A variation of this is the gradual escalation of requests (for information). If someone has already fulfilled a number of requests (for example, they looked up trivial information) it is then more difficult to refuse a request for confidential information.
- Creating the need to *return a favor*. Giving people something creates an emotional obligation where they feel they owe you something back. This makes it easier than usual to get someone to fulfill a request. When you have done something for someone (even when they did not ask for it), it becomes more difficult for that person to refuse a request.
- Creating the impression that the actual request already is a *concession*. When all that is needed is five minutes inside, it can be useful to request a tour on the premises. If this is refused, insist that it will only take five minutes to have a quick look around.
- Offering something that leads to *a personal benefit*. For example, send a phishing email with a code to receive a personal Christmas packet.
- Creating *unexpected situations* so that employees (especially security guards) are no longer able to follow their usual routine. We once dressed up as "Sinterklaas" (a traditional Winter holiday figure celebrated in the Netherlands) and his helper and have even penetrated a high security data center in this manner (Figure 3). The data center was at a secluded location and surrounded by high fences with barbed wire, dozens of cameras and an earthen wall that hid the building from view. We called security on the phone a week in advance and pretended to be from the HR department. We told them that we were calling about the Sinterklaas activities at the different locations. To get onto the premises, we first had to get through a checkpoint where a security guard behind bullet-proof glass consulted his colleagues inside the building when we showed up. Somewhat to

Figure 3. **The "Sinterklaas and helper" who managed to penetrate the data center.**

Figure 4. **A button camera that surreptitiously films security sensitive actions such as password keystrokes.**

our surprise, we were allowed to enter the premises and the door was locked again behind us. When we arrived in the data center itself, we walked straight up to a glassed-in security area with five security guards. A quick peak in our heavy bag of "pepernoten" (traditional Sinterklaas cookies) would have sufficed to reveal the recording equipment of the spy camera (Figure 4) and unmask us. "Hello! Well, here we are then!!", we called out, and instead of putting identification into the tray filled it with pepernoten. After bribing one of the guards with a chocolate letter, they allowed us access. We made a tour through the building and we left again with no problems.

- Using *distraction* such as bringing along that attractive female colleague with a short skirt and high heels.

As mentioned before, for each social engineering test specific attack scenarios are elaborated depending on the specific situation of the client. These scenario's often use one or more of the aforementioned techniques . In Case study 2, a *personal connection* was made with the victim, *recognition* was induced by referring to internal departments, a *personal benefit* was offered (not losing data) and a *compromise* was agreed upon (last paragraph). That "help" had been previously given also created the obligation for "*compensation*".

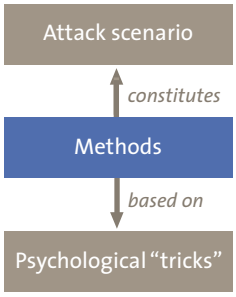

Figure 5. **The relationship between methods, tricks, and attack scenario.**

Case study 2

In a test case where the goal was to gain unauthorized access to a system, I called an employee to report that there was probably a problem with her system as it was causing an enormous amount of traffic on the network. I said that it would eventually crash her system and in the worst case prevent access to existing data. When I asked whether her laptop was very slow lately, I did indeed receive an affirmative answer (of course). After some random tapping on my keyboard, I said that I had found the problem, emphasized how very difficult it was to solve, but that I was working on it. I hung up and called again after half an hour to indicate that the problem was solved. After she had thanked me emphatically, I hung up.

Two days later, I called again and said that, unfortunately, it turned out that the problem was still present and it appeared that changes needed to be made to her laptop. I asked her whether she could bring her laptop along to the local IT department (that I had already called earlier to determine how the process worked and to verify that there actually was a local service point) to give the impression that I actually worked within her company. The employee said she was very busy and it was very bad timing. I said that we could make an exception and that I could try to solve the problem remotely. I said that we, because of security reasons, never asked users for their passwords over the phone, and therefore I asked her to temporarily change her password to "welcome123" so that I could fix the problem remotely. Two minutes later I was able to login to the laptop and I had access to the confidential data that I wanted.

Figure 6. **"Audio bug" with which one can listen in via cell phone calls.**

Figure 7. **Key logger that collects all keystrokes.**

Methods

Some common methods that are used in a social engineering attack are presented below. These methods partly rely on the previously described psychological "tricks". The combination of methods constitutes the attack scenario.

- *Phishing:* this is an attack method using forged email messages or web pages that appear to be legitimate such as those of the employer, but which in reality are controlled by the attacker. These email messages and pages are often aimed at collecting employee data (for example, passwords).
- *Dumpster diving:* searching for valuable information by looking through garbage bins, bins by copiers, or containers outside an organization's premises.
- *Pretexting:* obtaining information under false pretenses (the pretext). For example, calling an employee and pretending you are a colleague.
- *Tailgating:* "hitching" along with an employee through a secured entry gate to get physical access to a secured location.
- *Reverse Social Engineering*: a method in which the "victim" is manipulated so that they ask the social engineer for help. The social engineer creates a problem for the "victim" and then makes himself known as an "expert" who can solve the problem. The social engineer then waits for the "victim" to make a request. Trust is more likely because the "victim" takes the initiative.
- *Shoulder Surfing:* watch when someone enters a password or PIN code. You do not actually have to watch. In several tests we used miniature spy cameras such as a button camera (Figure 4) with which you replace one of the buttons on your jacket. After the entry of a password has been recorded, it can be played back later.
- *Placement of listening devices (bugs), wireless access point or key logger.* Once access is gained to a building, it is often easy to place listening devices. Modern listening equipment is available at low cost. For instance, such a device can dial a previously programmed cell phone number when sound is detected so that the attacker can listen along via the phone (Figure 6). Alternatively, a key logger can be installed (Figure 7). This device can be plugged in between the keyboard and the computer in a few seconds and will then record all keystrokes that are typed in. Current versions of key loggers can then automatically send an email with captured keystrokes to the attacker through a wireless network. Hiding an access point inside a building may also be useful (for example, by hiding it behind a radiator). After it is connected to the network, the attacker can then leave the building. On the outside, say in a car, the attacker then connects to the newly installed access point and

Case study 3

It was just after eight o'clock in the morning when I parked my car a few hundred feet from the building of one of our clients. I had earlier determined that most employees came to work with their car and parked behind the head office in the private parking lot. It seemed best to mimic this habit because walking through the car park would probably draw attention to my presence. In my car mirror, I kept an eye out for employees driving up to the lot. After about ten minutes, a gray car appeared. Once the car passed me, I merged and followed closely behind. Unfortunately, the car drove past the building of today's target and I was forced to circle back to my starting position. The second time, I had more luck and after the employee used his access badge to open the gate I could follow closely behind to get into the private car park behind the building. I waited until the employee left his car and entered through the staff entrance at the rear of the building. I walked to the smoking area near the entrance. I grabbed a new pack of cigarettes out of my pocket and lit one. Fortunately, there were no cameras on this side of the building, so I could just quietly wait until an unsuspecting employee joined this non-smoker who was flaunting a cigarette for the occasion. A woman wanting a smoke appeared after a little while. We talked a little and walked back together – through the door opened with her employee badge – into the building. I was inside! I immediately decided to follow her up the stairwell because it appeared that this client had placed card readers on the doors of each floor.

I followed her to the fourth floor and entered the office, once again she politely opened the door for both of us. Luckily, there was a coffee machine so I could stay there for a while and observe the floor without walking myself into a dead-end part of the building. A little further away, I could see some rooms set up for meetings. I took my coffee with me to a meeting room, removed the cable from the VoIP phone and inserted it into my laptop. While my laptop booted up, I cast a glance at the stack of paper that I had grabbed from the bin near the printer while walking by. It included emails with a lot of addresses of employees in the "To" and "CC" fields. Perfect! These would be the "victims" in my next attack.

After my laptop booted, I performed a port scan on port 80 on nearby IP addresses to look for internal web pages. I also used my web browser to try open a few obvious URL's like "intranet.clientname.com", "intraweb.clientname.com", "search.clientname.com", "directory.clientname.com", and so on. It did not take me long to find an internal web page. I copied the page and adjusted some text and after fifteen minutes I had put together an "employee of the month" voting page that looked exactly like the company web pages including logos and colors. Then, I started a web server on my laptop so that the newly created page could be accessed via the internal network.

A second limited port scan allowed me to identify an internal mail server that had mail relaying enabled (allowing anonymous email to be sent out). At that moment, I had been in the building for at least twenty minutes and had not been questioned by anyone about what I was doing there. Then, I focused again on the "victims". First, I sent an email via the mail server identified that contained the content of an email that I had copied from my spam folder, to some of the addresses in the printed emails. I hoped that this email would trigger an out-of-office message from one of the employees. When I then received just such an email, I copied the signature from it and changed the name and function to fictional ones. I now had a web page and an email message that looked exactly like those used in the organization. Then, I created an email with a reminder for the invitation to vote for the "employee of the month". The message indicated that a random selection of employees could nominate their colleagues for this award. This could be done via an internal web page included in the link at the bottom of the email. Naturally, logging in was required to prevent people from making duplicate votes. The reminder indicated that those who missed the first mail still had the chance to enter their vote up until 12:00 o'clock the same day. I switched to a second window and calmly waited until the password of the first enthusiastic employees appeared in the second window. This took exactly two minutes after sending out the reminder email.

By logging in at the site, the employees, in addition to their password and username, also automatically left behind their IP address. This was all the information that I needed. I started Metasploit (a hacker toolkit) that allowed me to remotely login to the PC of the first survey participant. Meanwhile, I had also found the user in the internal online telephone directory. Unfortunately, it turned out that the first employee worked in the finance department. At this stage, I was really looking for an IT administrator because they often have privileges to access a large number of systems. I decided to dump the local password hashes on the users system. Using the hash of the local administrator account, I tried to authenticate against the system of an arbitrary user on the network. This "trick" has worked at several client sites and was now also successful. Since all (or at least a lot of) desktops where installed from the very same image, the passwords for the local accounts were also identical. At this point, I had been inside for about three quarters of an hour without anyone noticing and I had already taken full control of two systems. Unfortunately, the password hash did not work on the domain controller, so I decided to keep logging into desktop systems until I found a system with a user (or process) that was running with the highest privileges (for example, the IT administrator). After twenty minutes, I found a system where an IT administrator was logged on. The freeware Metasploit tool has a built-in feature allowing you to take over the identity of a user and with it all his privileges. After I took over the identity of IT administrator, I had domain administrator rights and full access to all Windows systems and the data present on the network, including all servers with financial administration and the mailboxes of the board of directors. I made some screenshots and decided that it was time for a second cup of coffee.

then accesses the internal network with little chance of being detected and arrested.

- *Malware:* malicious software that, for example, collects and forwards passwords to the email address of the attacker. Malware can be installed on the systems by, for example, using an infected PDF file ([Paqu01]). The PDF file can be circulated in different ways, for example, by leaving a USB memory stick containing files titled "2011 payroll" or "fraud investigations in 2011" or similar. Ideal places to leave these sticks are in the restrooms or by the coffee machine. When the "victim" opens the PDF the malware is being run in the background automatically.

In Case study 3, some of the above methods are used. This example shows, amongst other things, how information obtained from one attack can be used in another attack to get even more information.

Knowing about possible attack techniques and the weaknesses of the target builds real awareness

Case study 3 shows that it is not always important how many employees are tricked by social engineers. In this particular situation, it was enough for an outsider to deceive only two employees to compromise the entire IT environment.

Countermeasures

Awareness

The keyword in countering social engineering attacks is awareness. More specifically, it is what the targets know about possible attack techniques and their own weaknesses. In one of my assignments, in addition to the usual paper bins alongside printers, the client also placed large enclosed bins for any paper containing confidential information. Nonetheless, the bin for ordinary waste paper provided a huge stack of confidential documents (reports of security incidents, HR information, passwords, and so on). Why? It was probably too much trouble to push the piles of paper through the small slot in the bin for confidential paper and it was just easier to throw it all away in one go.

When clients hear how a trick works at a presentation or training, people often say things like: "you have to be really naive to fall for that, it would never work on me". Our test results shows differently. Therefore, it is useful to perform a test and confront employees with the results within their organization to really raise awareness. It usually shows that people are not so ready for such an attack as they think they are. It is this that leads to real awareness. In addition to promoting awareness, a test is also quite useful in identifying risks.

Guidelines

Alongside awareness, it is essential to draw up guidelines and continue to check compliance with these. Consider drawing up "ten rules for information security". An example is as follows:

1. Never reveal your passwords to others (including IT employees).
2. Do not share internal information with outsiders.
3. Adhere to the clean desk and whiteboard policy.
4. Lock your computer when you leave your workstation.
5. Do not leave any information behind at the printer.
6. Use secure waste bins for confidential information.
7. Verify the identity of the caller when asked for confidential information. (For example, in case of a telephone request, ask the caller to call back on a specific number.)
8. Never save confidential information locally or on a private PC or device (drive, USB stick).
9. Immediately alert the security officer about any suspicious activities.
10. Keep your access badge visible and request colleagues to wear their badge. Any unknown person without a badge should be escorted out of the building and handed over to the reception and/or security.

To ensure that such rules are followed, it is necessary to monitor that employees are actually complying. The outcome of the monitoring (both positive and negative) should be given as feedback to the relevant employees.

Conclusion

After reading this article, you may doubt that the cases described ever happened and that such incidents can succeed in real-life. Unfortunately, the reality is that these and similar attacks occur every day, despite the various security measures. Security personnel, barbed wire fences, access cards, CCTV, alarm systems, and so on, are not enough. Social engineers know how to penetrate into the heart of an organization. Performing a social engineering test can be a good way to identify risks in an organization and raise employee awareness.

References

[Hadg01] Christopher Hadgany, *Social engineering – the art of human hacking*, 2010.

[Mitn01] Kevin D. Mitnick and William L. Simon, *The Art of Deception*, 2002.

[Paqu01] Matthieu Paques, *Hacking with PDF files*, http://www.compact.nl/artikelen/C-2009-4-Paques.htm.

[security.nl]: articles concerning social engineering attacks, http://www.security.nl/tag/social%20engineering.

About the author

M.B. Paques is a manager in the "ICT Security and Control" team for KPMG IT Advisory. He has experience with security testing, social engineering, technical security reviews and the security of new technologies.