Social Engineering Threats and Countermeasures In An Overly Connected World

By Shane MacDougall, shane@tacticalintelligence.org

Social engineering (hereafter referred to as SE), is rapidly becoming one of the hot topics in information security, which is curious since it has been an oft-used attack vector for centuries. Its emergence into the world of espionage and warfare has been well established since Sun Tzu wrote The Art of War, and most likely considerably before that.

Today's media often conflates phishing with SE, but this is not an accurate use of the term. SE involves actual social interaction with a target – a person speaking or interacting with another person, be it over the phone or in person. A good litmus test is if it can be automated (such as phishing attacks) it's not social engineering.

For an SE attacker to have maximum efficacy, (s)he must utilize many different elements, from OSINT harvesting and dossier building, to efficient application of the "social handshake" (the presentation of the pretext to the target). Efficacy lift can be gained via use of push polling (presentation of information gleaned by OSINT or observation), psychological gambits such as the use of humor or sympathy, trust transference (using an established social handshake to gain trust of the peer of a target), through the use of authority-based pre-texts, and many other methods.

The social handshake is the most critical part of a successful SE attack. It occurs during the initial interaction between the attacker and the victim, and is somewhat akin to the TCP 3-way handshake (although I don't wish to delve into an esoteric technical comparison of social interaction to packet exchange).

It is essentially SYN (presentation of attacker and his pre-text), SYN-ACK (the target validating the attacker and the pre-text), and ACK (acceptance of the attacker's pretext by the target). As in the TCP 3-way handshake, if any part of the interaction fails, the attack (connection) will fail. While a target can terminate the "connection" at any time, the author's experience is that once the social handshake is completed, the attack will usually succeed (as long as the attacker sticks to a well-defined and targeted pre-text).

Establishing this initial connection is not nearly as difficult as one might think. In fact, past history has shown that intelligence agents have an over 80% success rate with establishing connections over the phones with targets with extremely basic introductions, such as asking "can I have a few minutes of your time?"(1) As long as the caller has a proper response to the initial pushback (if, indeed, any is received), the success rate is high. What is not necessarily high is the acquisition rate of the flags (the information or end result the attacker seeks). That is determined by the strength of the pre-text and the ability to continually groom the target throughout the call (methods such as humor, complements, sympathy, et al can be used).

The ideal end goal of the attacker is to create the PST (persistent social threat), an ongoing relationship with a target that lasts well beyond the initial interaction, and can be repeatedly tapped in the future. Such PST's can be incredibly devastating to an enterprise, since they are not (usually) detectable via firewall logs or blocked by anti-virus solutions. If the company becomes aware of a PST resulting in a leak, the duped target is not likely to come forward and reveal themselves as being the weak link, thus the attacks are an ideal attack vector.

An effective social engineer will also exploit common psychological weaknesses. For example, it has been shown that humans are inefficient at detecting deception, on average only detecting deception correctly in 54% of interactions (2). This is because most people rely on inaccurate methods such as eye contact and body languages that are not reflected in reality (3). As such, an effective attacker will maintain eye contact when telling critical lies, or avert eye contact when they want their target to think they are lying.

Why would one want a target to think they were lying? An example could be if the attacker is posing as employee, talking with the target about experiences in a positive manner when the mark knows that not to be the case (happy with pay/management/work environment/promotion process while most employees within the company are not – something easily discerned via resources such as GlassDoor) – i.e. trying to maintain an outwardly upbeat view of the company while "betraying" their true feelings. This helps to establish a trust relationship with the target, even though they think the attacker is lying. While this may seem counterintuitive, it has worked many times in the field for the author.

Creation of a thorough dossier or OSINT compilation is also essential to creating an effective PST. Items that often fall through the cracks in the Maltego type information gathering model include photo sharing sites, which can yield tremendous information that can be utilized in an attack. The author has successfully leveraged a single photo of a target's work cubicle to discern marital status, family structure, ownership of pets, college affiliation, and hobbies. Another photo set yielded employee names, badge layout, internal corporate event information, cultural backgrounds, and physical plant layout.

These data points were easily leveraged to establish credibility with targets, and were successfully used in the Defcon 19 SECTF competition.

Non-traditional OSINT activities can include lobby-surfing, which can yield information such as names of client (gleaned from lobby "welcome visitor" signs), vendor/client/employee names (gleaned from the sign-in sheet), employee names, positions and home addresses (harvested from magazine labels), and a general assessment of overall security posture (is CCTV present? PTZ/fixed? Are proximity cards or biometrics in use? Guards? Is piggybacking allowed?). All these data points can prove to be extremely helpful to an attacker, and can be gathered in a lowly intrusive/kinetic method. It is important to note that SE and OSINT efforts are often used as a complementary attack vector, not as the end attack method, although the latter is often devastating enough that other attack vectors are rendered useless.

Successful pretext generation is often driven by the end data points that the attacker desires to gather. For example, in the Defcon SECTF competition, some of the more esoteric data points that were to be gathered included where the target got their office supplies from, who performed the onsite cleaning, who provided food service, and dumpster disposal. To grab these items from a target required a pre-text that would fit in alongside the gathering of more obvious data points such as OS version, anti-virus type, current patch level, etc.

To do this effectively, the author pitched the target with a pre-text of doing information gathering to see whether or not the company would be eligible for preferential placement for a potential government contract, under the guise of minority or veteran owned businesses. "If we can tie ourselves to a

minority owned business that will give us a real edge in winning this. Do you know if we get our office supplies from a minority owned business? What about something like dumpster disposal? Who does that? Any chance they are minority or veteran owned?"

The pretext worked perfectly in both SECTF 19 and 20, since the questions were not posed as "who does our dumpster disposal" which could arouse suspicion as plain information gathering, but rather were being asked for another reason – gaining the company a competitive advantage. This manner of elicitation has long been used by both intelligence agencies and corporate information gathering specialists with great success.

Another very effective method of trust building is the utilization of what I call "delayed validation." The attacker contacts the target and informs them that they will be coming on-site in the near future and that the call is to introduce himself and the purpose of the visit. This method is very effective since the attacker at first does not attempt to elicit any information, just to set the date. After some small talk, the attacker asks for hotel recommendations for the area, tourist attractions, and other minutiae. As this is happening, the target is subconsciously at ease, since they think they will have lead-time to establish any inconsistencies, and as such they are at a lowered state of suspicion. Additionally, the target is lulled into a sense of safety since the attacker is not attempting to gather any information, rather establish logistics. It is only later in the conversation that the attacker begins to ask some questions, usually under the guise of getting some prep work out of the way while they have the target on the phone. What happens in most instances is that the lowered level of suspicion is maintained, since the most critical part of the social handshake comes at the beginning of the call. Since the target has already

accepted the initial risk, and eventually established the social connection, they usually fail to reset their "awareness parameters" that they have in place during the initial interaction.

On-site SE attacks pose a completely new set of issues for the attacker. While over the phone interactions allow them the freedom to simply hang up with little threat of being physically detained, the attackers are somewhat hamstrung by the inability to visually gauge reactions from the target, as well as a lessened level of situational awareness.

However, the simple act of presenting oneself to the target has a surprising effect: there is an added legitimacy factor added to the victim's mind. Surely nobody would be bold enough to come onto the premises to try to fool their way inside. The common perception that attacks like this only take place in spy movies and TV shows like "Burn Notice" biases most targets, and is usually reinforced by a "I would know a scammer if I met one / it would never happen to me" attitude. Unfortunately for most targets, the history of espionage and computer security is littered with people who suffered from all-too-common cases of hubris.

On-site attacks that are especially effective are those that involve the attacker presenting himself at a facility, while referencing a worker who is unavailable as being his reason for visiting ("I was hired by Ted Johnson to work on the tape silos"). If the reference target utilizes social media heavily, the attacker may even be able to refer to where the target currently is. Sites such as FourSquare and Twitter are remarkably effective reference points, since they can give attackers real-time reconnaissance information. Out of office

emails and voicemail messages are also a common source for this information.

The attacker can make sure the reference target is not reachable by phone via one of a few methods. War-faxing the reference target's cell phone has worked very well for the author on multiple occasions. After a few minutes of answering the phone, only to be greeted with the annoying beep of a fax machine usually results in the phone being turned off in short order. Another method that works well is to ensure the reference target is physically separated from his phone. This is achievable in a few ways, one of which is the fake job interview. The target is offered a "too good to turn down" job opportunity, and while they are off at the "interview", the attack takes place. When the receptionist or guardian of the site is unable to reach the reference target, they will either turn the attacker away, or allow them access. The latter is much more likely if a proper dossier has been collected, since the attacker can drop many reference points that infer they must be in the company's circle of trust.

If the receptionist decides not to allow the attacker in, this can often be mitigated by calmly asking for them to call a taxi to take the attacker back to the airport. "I'll have to reschedule and fly back in sometime next month." In the author's experience, this has an exceptionally high success rate when it is delivered in a non-aggressive, resigned manner. The receptionist realizes that a significant delay (and possible cost) may be incurred, so they err on the side of self-preservation, or in what they mistakenly believe to be in the best interests of the company.

With all the subterfuge and tools available to attackers, it may seem that all is lost for the corporate security group when it comes to defending against social engineering attacks. Nothing could be farther from the truth. There are many things that companies can do to mitigate these threats, but they involve much more than basic awareness training.

That said, most awareness programs are sorely lacking when it comes to training regarding social engineering. A truly effective program takes more than a 30 minute discussion of the threats, it requires the users to actually learn some of the tools of the trade. One technique that the author has found to be exceptionally effective, is to complement a robust social engineering awareness lecture with having all employees take a few hours to create a dossier on a randomly assigned co-worker from social media sites. The offer of a cash prize for the best dossier also seems to make this process even more effective. Not only does the information gathering process make all employees aware of the dangers of social media, it makes them think like attackers, and helps modify their own online behavior. The author has also observed that these efforts often lead to the employees in turn educating friends and family members to the threats, which helps secure the internet community, if even a few users at a time.

Awareness programs also need to move away from the traditional "we won't ask for passwords over the phone" sort of mindset. A skilled SE will never ask for a password, since that instantly raises flags. Instead they will spend long times building a trust relationship, garnering information from the target that they can leverage in many other ways. If they need data from a system, they will target a user with access to that system and get the information from that user, or gain information by compromising the user's machine.

Still, every company should have a list of “hot button” items that all employees should have close by, that trigger training to kick in if asked. A formal reporting process needs to be set up, and any attempted social engineering attacks logged and investigated. I know of several major companies that maintain “suspicious phone call” databases that are monitored both by InfoSec and the physical security groups. It is important that those two groups share information on a regular basis; an ongoing social engineering attack may very well be followed up by attempts to gain physical access, and vice versa.

Companies need to perform ongoing social media monitoring, and not in the traditional sense. Currently social media monitoring is usually driven by corporate communications, and is mainly focused on negative comments by customers. The people responsible for these programs need to be trained to identify sensitive information that is being leaked on the web in addition to doing corporate public image monitoring. The author has seen cases of companies focusing on negative employee comments in Glassdoor forums, while being completely oblivious to potentially sensitive technical information being leaked in the same forum. For social media monitoring to be most effective, it should be a cross-functional task, with InfoSec involved, at least on a semi-regular basis.

Employee use of social media should also be monitored aggressively to identify those who use sites such as FourSquare or Twitter to post their location/activity, especially during business hours. It should be noted that some federal and state laws require specific rules about this sort of monitoring, and it might require this be performed by an external third party. This can also be made part of any regulatory penetration test process.

The creation of a daily or weekly internal password is another effective defense mechanism. Anyone who is calling in to a company and asking for information, or is identifying themselves as an employee should be asked the password. If they cannot produce the password, the call should be terminated and the attempt logged and investigated. An additional layer of defense is to create a callback process for all calls from the field from people the employee is not familiar with. While this adds an additional delay to the business process, it allows the company to log a contact number, and gives the employee a chance to verify the caller is who they say they are.

It is also important that employees be taught it is okay to say "no." The traditional adage that "the customer is always right" has led to call center employees and customer service personnel being overly eager to ensure customers are happy, and as a result, are hesitant to risk annoying a client. This attitude must change; employees need to begin to adopt a sense of distrust in every single interaction. That is not to say that they need to act in an anti-social manner, but they need to be aware that there are people out there whose sole goal is to take advantage of them. One can be polite while remaining wary.

With this in mind, it is critical that employers let their personnel know that they will never be punished for erring on the side of caution. An angry customer can usually be assuaged with proper messaging and outreach; a breached company is much harder to fix.

Additionally, employees need to be taught situational awareness, to identify instances where they may be being

"played." If a caller is overly nice, and is making the employee laugh and enjoying themselves, they should step back and ask themselves – "why am I enjoying this conversation so much?" (4)

Finally, one of the most critical components of a successful program is to constantly be testing employees. Routinely try to social engineer information from random employees, and offer a financial incentive to those who detect or rebuff the attempts. A quarterly award of $500 to an employee who successfully detects or reports a potential social engineering attack is a surprisingly effective enticement. The award need not be given to everyone who reports, however they should be put in contention to qualify for it.

The corporate world is only now beginning to understand the enormous risks that social engineering poses. Rigorous training and monitoring can mitigate these risks, however it cannot be a half-hearted effort, nor can it be a one-time thing. It must involve all employees, cross departmental boundaries, and have management buy-in. It also must be part of an overall robust information security program that is driven by a dedicated resource. Properly implemented, the author feels it is possible to finally have a patch for "human stupidity."

_____

(1) "Confidential: Business Secrets - Getting Theirs, Keeping Yours, 2nd Ed.", John Nolan, Yardley-Chambers,June 1999

(2) "Accuracy of Deception Judgements", Charles F. Bond Jr. (Texas Christian University) & Bella M. DePaulo (University California Santa Barabra), Personality and Social Psychology Review, 2006

(3) “What Every Body Is Saying: An Ex-FBI Agent’s Guide to Speed Reading People”, Joe Navarro and Marvin Karlins, William Morrow Paperbacks, 2008

(4) “It’s Not All About Me: The Top Ten Techniques For Building Quick Rapport With Anyone”, Robin Dreeke, People Formula, 2011