

Social engineering is one of the most prolific and effective means of gaining access to secure systems and obtaining sensitive information, yet requires minimal technical knowledge. Attacks vary from bulk phishing emails with little sophistication through to highly targeted, multi-layered attacks which use a range of social engineering techniques. Social engineering works by manipulating normal human behavioural traits and as such there are only limited technical solutions to guard against it. As a result, the best defence is to educate users on the techniques used by social engineers, and raising awareness as to how both humans and computer systems can be manipulated to create a false level of trust. This can be complemented by an organisational attitude towards security that promotes the sharing of concerns, enforces information security rules and supports users for adhering to them. Even so, a determined attacker with sufficient skill, resources and ultimately, luck, will be able to retrieve the information they are seeking. For this reason, organisations and individuals should have measures in place to respond to, and recover from, a successful attack.

Introduction

In cyber-security, social engineering refers to the manipulation of individuals in order to induce them to carry out specific actions or to divulge information that can be of use to an attacker. Social engineering in itself does not necessarily require a large amount of technical knowledge in order to be successful. Instead, social engineering preys on common aspects of human psychology such as curiosity, courtesy, gullibility, greed, thoughtlessness, shyness and apathy.¹

Social engineering techniques are commonly used to deliver malicious software (malware²) but in some cases only form part of an attack, as an enabler to gain additional information, commit fraud or obtain access to secure systems. Social engineering techniques range from indiscriminate wide scale attacks, which are crude and can normally be easily identified, through to sophisticated multi-layered tailored attacks which can be almost indistinguishable from genuine interactions.

Social engineers are creative, and their tactics can be expected to evolve to take advantage of new technologies and situations. This paper outlines some of the most common and effective forms of social engineering.

Wide scale attacks

Phishing

The most prolific form of social engineering is phishing, accounting for an estimated 77% of all social-based attacks with over 37 million users reporting phishing attacks in 2013.³ Phishing is the fraudulent attempt to steal personal or sensitive information by masquerading as a well-known or trusted contact. Whilst email phishing is the most common, phishing attacks can also be conducted via phone calls, text messages and fax, as well as other methods of communication, including social media.

A large amount of wide scale email phishing attacks remain unsophisticated and will be recognised by most (although not all) computer users as illegitimate. However, email phishing is becoming increasingly sophisticated and attackers will use a variety of techniques to either make the email appear legitimate or to lure the victim into acting before thinking. Attackers may disguise the address the email is sent from so that it appears to be from a well-known organisation and common ones include banks, utility companies, couriers, recruitment agencies and government. Better designed phishing emails will actually appear to be very similar imitations of legitimate emails from these organisations (see example 1). Another common technique is to make use of major news events by posing as having new information on the event, or asking the recipient to take action (donate money, sign a petition, etc.) in relation to the event.

Despite increasing competency in wide scale campaigns, there are still indicators that frequently appear in phishing emails:

- Messages are unsolicited (i.e. the victim did nothing to initiate the action)
- Messages are vague, not addressed to the target by name and beyond purporting to be from a known organisation, contain little other specific or accurate information to build trust
- May be from an organisation with which the target has no dealing with

¹ How hackers exploit 'the seven deadly sins', *BBC News* <http://www.bbc.co.uk/news/technology-20717773>

² See CERT-UK's 'An introduction to malware': <https://www.cert.gov.uk/wp-content/uploads/2014/08/An-introduction-to-malware.pdf>

³ The Social Engineering Infographic <http://www.social-engineer.org/social-engineering/social-engineering-infographic/>

- Contain poor spelling and grammar, typos or use odd phrases; whilst this is becoming less common as attackers are becoming more proficient, mistakes are still made
- Are too good to be true or make unrealistic threats, often with a sense of urgency
- Are sent from an email address that, whilst perhaps similar, does not match ones used officially by an organisation
- Contain incorrect or poor versions of an organisation's logo, and may contain web links to sites that, whilst perhaps similar, are not ones used by that organisation

Phishing emails often ask the user to follow a link to a website or open an attachment. Some may ask the user to reply to the email, after which they will be engaged in an exchange of messages to elicit confidential information. When asked to click on a link, it may be designed so that the text the victim clicks on appears to be for a known website, but the link takes them to a completely different website (a technique known as obfuscation). At the website, the victim will then be asked to enter confidential information or may unknowingly download a file which will subsequently infect their machine with malware. Likewise, any attachment on a phishing email is likely to contain malware.

Example 1 – A phishing emailⁱⁱⁱ Although it appears to be from O2, closer inspection, by right clicking or hovering over the name, shows the email address has been spoofed. For example, 'user123@o2-mail.com'. An official O2 email would come from '@o2.co.uk'.

The subject title is O2 (zero-two) not O2

It is addressed generically, not to the customer by name

Hovering over the link here will show that it will not take the user to O2's website, but to a completely unrelated website

A comma is used instead of a decimal point

By the time this email was sent, O2 has discontinued their 'Lucy' virtual assistant

This message was sent with High Importance.

From: O2 Billing
To: er.co.uk
Cc:
Subject: Your O2 bill is ready #1035346

Dear Customer

Your O2 bill for 28/05/14 is now ready. You can [look at your bill here](#).

In total, your bill for this month comes to £372,85. We'll request this amount from your chosen account on, or just after, the date in your bill.

To see your bill, you'll need the username and password you were given when you joined O2. If you've forgotten them, we can give you a [reminder](#).

Is your bill more than you were expecting?
If so, here's a few reasons why this might be:

- You could have gone over the minutes, texts or data that's in your allowance.
- You could have called or sent texts to numbers that can't be taken from your allowance such as International, 0800, 0845 numbers or directory enquiries.
- You have used your phone for calls, text or data whilst abroad.

To view any charges outside your allowance [click here](#)

If you have any questions, just [ask Lucy](#). She's our online virtual agent. You can also find out more about what's included in your bill with an [online demonstration](#).

Best regards

O2

This email is sent from Telefónica UK Limited. Registered office:
260 Bath Road, Slough, Berkshire, SL1 4DX. Registered number: 7270332.
Please do not reply.

More sophisticated phishing campaigns may even extend to taking victims to a close replica of a legitimate website that is designed to trick them into entering username, password or other confidential information.

Baiting

Another form of wide scale attack is baiting through the use of online adverts and websites. As with phishing, these will usually have offers that are too good to be true or with an urgent warning. This includes some websites that allow the user to download or stream videos (i.e. movies or TV shows), or pop-ups that purport to have detected a problem with the victim's system which clicking on the pop-up will solve. Following the links provided in the bait, a user may then be tricked into giving away

⁴ O2 Phishing alert mid-2014 <http://news.o2.co.uk/2014/05/29/phishing-alert-may-2014/>

personal information, or their machine may automatically download malware. These attacks can be crude, but others are sophisticated and persistent (see example 2).

Example 2 – Shylock

Shylock is a sophisticated, hard to detect, adaptable piece of malware that enables criminals to steal victims' credentials and support financial cybercrime. Whilst UK-led law enforcement activity in mid-2014 successfully disrupted and reduced the prevalence of Shylock, it still continues to be used by criminals. The UK was a major target for Shylock, and at its height around 61% of all infected websites were UK based – the majority of these being from the retail sector.^{iv}

Shylock is distributed via a variety of methods, including malicious code embedded into online adverts which subsequently appear on legitimate websites and the direct targeting and compromise of popular websites. Some compromised websites may display a 'missing plugins' message with a button to install the missing plugins. When clicked this downloads and installs Shylock onto the victim's system.

Once running, Shylock, can send any data entered into a computer to the attacker, including website credentials and other sensitive information. It can even create a false 'chat' window on a banking website, enabling the attacker to interact with the intended victim in order to persuade them to give up additional sensitive information – effectively a second layer of social engineering.

Another mass form of baiting is the use of 'free' Wi-Fi hotspots, although this requires some technical knowledge. The attacker creates a Wi-Fi hotspot that is clearly labelled as 'free', typically in public areas such as coffee shops, airports and hotel rooms. Whilst they may provide a victim with an internet connection, any data sent over this connection can be intercepted by the attacker, through what is known as a 'man-in-the-middle' attack.⁶ The ability to intercept the victim's data can extend even to secure connections to services such as online banking. The attacker may also be able to remotely install malware on to the victim's system, allowing a range of further exploits to be carried out.

Focusing the attack

Spear phishing

Spear phishing is used by more sophisticated attackers who will limit the target audience and increase the precision of their messages, increasing the appeal of the message and apparent legitimacy. A spear phishing attack may target individuals within a particular business sector, who work in the same company, in the same department, or who share some other common attribute. A spear phishing email may even target just one specific individual if they are seen to be of sufficient value to the attacker. Whilst this decreases the number of potential victims, it is also likely to result in a higher proportion falling for their attack. Some spear phishing attacks can still be crude, and still remain easy to spot as they contain some of the indicators listed above. Others can appear legitimate and are extremely difficult to identify as malicious.

A competent attacker will research their target(s) in order to maximise their chances of success. They will try to find out information about the organisation, including organisational charts, contact details

⁵ BAE Shylock Whitepaper <http://info.baesystemsdetica.com/rs/baesystems/images/ShylockWhitepaper.pdf>

⁶ This is similar to another technique called 'Evil Twin' where an attacker creates a Wi-Fi network with the same name as a known public network (e.g. 'BTOpenzone' or 'Starbucks'), that a smartphone, tablet or computer will automatically connect to. As with baiting, this now enables the attacker to intercept all data sent by the victim. Unlike baiting however, this can happen without the user's knowledge if their device is set to automatically connect to known public networks.

and combine this with knowledge obtained from their victim's social media profiles and other publicly available information. Rather than a generic greeting, a recipient is likely to be addressed by name and the message will probably include other personalised details. An attacker is likely to use the identity of a third party that is to be known or of interest to the intended victim(s), such as a supplier, to leverage existing trust relationships. Similarly, the attacker may try to replicate the third party's email address and use their research to assume the identity of someone who is employed by the third party, potentially someone who they believe their victim(s) know. They may even attempt to gain access to a third party's email account (see example 3).

Example 3 – A spear phishing email^v

Subject: employee making negative comments about you and the company

From: <name>@<compromised company's domain>

I noticed that a user named FinanceBull82 (claiming to be an employee) in an investment discussion forum posted some negative comments about the company in general (executive compensation mainly) and you in specific (overpaid and incompetent). He gave detailed instances of his disagreements, and in doing so, may have unwittingly divulged confidential company information regarding pending transactions.

I am a longtime client and I do not think that this will bode well for future business. The post generated quite a few replies, most of them agreeing with the negative statements. While I understand that the employee has the right to his opinion, perhaps he should have vented his frustrations through the appropriate channels before making his post. The link to the post is located here (it is the second one in the thread):

<http://forum.<domain>/redirect.php?url=http://<domain>%2fforum%2fequities%2f375823902%2farticle.php%par>

Could you please talk to him?

Thank you for the assistance,

<name>

The example to the left has been attributed to a group called 'Fin4' by FireEye, a cyber-security company. Fin4 were identified as targeting individuals within companies who had access to information about market catalysts (i.e. events that would cause changes in stock prices). In this example, Fin4 successfully compromised the email of an individual at a public company (possibly through the use of social engineering), and then used the compromised account to send a message that would play on a chief executive's concerns: damage to reputation and disclosure of confidential information. Prompted by this, the victim would click on the link in the email, which would result in the download of malware.

Watering hole attacks

Watering hole attacks, similar to baiting, use trusted websites to infect victim's computers. They are typically more sophisticated than most other social engineering techniques as they also require some technical knowledge. A watering hole attack works by compromise a trusted third party website to deliver malicious code against the intended victim's computer. As with other targeted social engineering attacks, the attacker will research their intended victim(s) and identify one or more trusted websites that they are likely to access. This may be a supplier's website, an industry journal, think tank or some other website that the attacker has identified as of interest to the intended victim. Having identified a suitable website (or websites), the attacker will seek out vulnerabilities within the server that hosts the website, and having found one, insert code that will enable malware to be downloaded, sometimes with little or no interaction from the victim (known as a 'drive-by' attack).

Attacking on multiple fronts

A determined attacker may adopt a multi-layered approach along with additional techniques to increase their target's trust, or confusion, in order to maximise the chance of success. Whilst somewhat indiscriminate, an attacker could begin dialling random numbers within an organisation claiming to be IT support (potentially using a real name from the IT support department gleaned from social media) until they eventually find a victim that does have an IT issue. In their attempt to solve the problem, they will trick the user into giving them login, password or other information that will be useful in compromising their computer. Alternatively, the attacker may pretend to be an executive, urgently demanding to be sent an important (and sensitive) document to their personal email address as they cannot access their work account. In both cases, the victim is put under pressure to do something they should know they should not do: they do not want to question someone who knows more than them (IT support), or who is senior to them (the executive), and refusal to comply could get them in trouble. Some attackers may be even more creative (see example 4).

Example 4 – A sophisticated social engineering attack

As part of a vulnerability assessment for an organisation, an assessor carried out some information gathering and found the locations of servers, IP addresses, email addresses, phone numbers, physical addresses, mail servers, employee names, titles and much more. Through Facebook, he was also able to get other personal details about the CEO, such as his favourite restaurant, sports team and that he was involved in cancer fundraising. Using this information, he called the CEO and posed as a fundraiser from a cancer charity the CEO had dealt with in the past. He informed him they were running a raffle with prizes including tickets to a game played by his favourite sports team and tickets to his favourite restaurant. The CEO was interested and agreed to let the assessor send him a PDF with more information on the fundraising and the raffle. The assessor even managed to get the CEO to tell him which version of Adobe reader he was running. Soon after he sent the PDF, the CEO opened it, installing malware that enabled access to his machine.^{vi}

Such sophisticated attacks are usually reserved for targets who will have access to valuable information, such as chief executives; this type of spear phishing is known as whaling.

Physical baiting

An attacker may also use hardware to bait a target or group of targets. The nature of this type of social engineering means that it is typically only used by more sophisticated attackers against a particular sector, organisation or individual. A common example of baiting is to leave a form of digital media (e.g. a USB flash drive, CD, DVD) unattended, perhaps labelled with something alluring to, and in a location frequented by, the intended victim (like a car park). The intent is that they will pick it up and then use it on a personal or work computer, whereupon that computer is infected with malware. Another form of physical baiting can be at conferences or other events, where the attacker is in a position to hand out free USB drives as gifts, or provide further information on digital media, which is secretly loaded with malware.