Mastering Defensive Security

Effective techniques to secure your Windows, Linux, IoT, and cloud infrastructure

Cesar Bravo

Foreword by Darren Kitchen – Founder, Hak5

Mastering Defensive Security

Effective techniques to secure your Windows, Linux, IoT, and cloud infrastructure

Cesar Bravo

Packt

BIRMINGHAM—MUMBAI

Mastering Defensive Security

Group Product Manager: Vijin Boricha
Publishing Product Manager: Shrilekha Malpani
Senior Editor: Arun Nadar
Content Development Editor: Yasir Ali Khan
Technical Editor: Nithik Cheruvakodan
Project Coordinator: Shagun Saini
Proofreader: Safis Editing
Indexer: Manju Arasan
Production Designer: Jyoti Chauhan

First published: October 2021

Production reference: 1211021

Published by Packt Publishing Ltd.
Livery Place
35 Livery Street
Birmingham
B3 2PB, UK.

ISBN 978-1-80020-816-2

www.packt.com

Contributors

About the author

Cesar Bravo is a researcher and inventor who has more than 100 inventions related to cybersecurity that are being patented in the US, Germany, China, and Japan. Those inventions include cybersecurity hardware, secure IoT systems and devices, and even cybersecurity systems for autonomous cars.

He loves to share knowledge and he has been working with several universities to teach cybersecurity at all levels, from introductory courses for non-IT people up to a master's degree in cybersecurity (for which he has also served as a thesis director).

In recent years, Cesar has become a recognized speaker (including delivering a TEDx talk), giving international presentations about cybersecurity and innovation in the UK, Germany, Mexico, the US, and Spain.

First, I want to thank all my students, who always encourage me with their questions and comments to become a better professional.

To my peer masters in cybersecurity, who took the challenge to learn about new topics and explore a new universe of possibilities, I am super grateful and proud of all of you.

To the cybersecurity community, who invest countless hours to stay up to date with new threats to make the world a better and more secure place to live, for you that live and work in the shadow of your desk, let me say that YOU are the real heroes!

And to my family and friends, who have always supported and encouraged me to become the best version of myself, to all of you, THANK YOU!

About the reviewers

Smith Gonsalves is the director and principal consultant of CyberSmithSECURE, a boutique consulting firm that specializes in providing cybersecurity services to MNCs worldwide. He has been known and recognized in the industry as one of India's youngest cyber evangelists and information security professionals of the time. His key area of work is in the instrumentation of orchestrating cyber capabilities for safeguarding high-end enterprises and institutions. Smith is a Cert-In Certified Auditor and has completed industry-nominated certifications including CISA, OSCP, CEH, CHFI, and TOGAF during his 7+ years of experience.

Yasser Ali is a cybersecurity consultant and red teamer at **Dubai Electricity & Water Authority** (**DEWA**).

Yasser has an extensive background in consultancy and advisory services. His experience in vulnerability research, pentesting, and reviewing standards and best practices has made Yasser a highly sought-after expert for enterprises.

Yasser's passion is mostly spent on the development of red teaming labs and offensive training where cybersecurity professionals sharpen their skills and learn new tradecraft-emulating **techniques, tactics, and procedures** (**TTPs**) used by adversaries.

Yasser was showcased in the BBC documentary movie *How Hackers Steal Your ID*. He is a specialized trainer and is regularly invited to participate in global information security conferences and discussion panels.

I wish to thank Shagun, Ali Mehdi, and the Packt team for their time and for allowing me the opportunity to review this book.

Big thanks to all security researchers and InfoSec communities such as HackerOne, Hackers Academy, and Malcrove. Without their contribution, innovation, and willingness to break the rules but not the law and to help one another, cybersecurity wouldn't be what it is today.

Lastly, a special heartfelt thanks to my caring and loving parents and siblings for always supporting me.

4
Patching Layer 8

"Cybersecurity measures are frequently focused on threats from outside an organization rather than threats posed by untrustworthy individuals inside an organization. However, insider threats are responsible of many millions losses in critical infrastructure nowadays."

– Ricardo Gazoli – Head of IT executive

Users are, by far, the most vulnerable factor in cybersecurity. In fact, a recent study revealed that more than 50 percent of attacks are caused by insiders either by accident (inadvertent users) or intentionally (malicious insiders).

One common mistake is to prepare cybersecurity specialists to deal with technical challenges such as servers and networks, and not prepare them to address all the risks related to the human factor (inadvertent users and malicious insiders). In fact, many people agree that managing the users is far more complex than dealing with systems because, in the end, you cannot just simply patch them!

Therefore, managing users is an art; in this chapter, I am going to show you all the different attack vectors aimed at the user, but also how you can master a plurality of techniques, methods, and tools to prevent those kinds of attacks.

In this chapter, we are going to cover the following main topics:

- Understanding layer 8 – the insider threat
- Mastering the art of social engineering

- Defending against social engineering techniques
- Defending against social engineering attacks (patching layer 8)

Understanding layer 8 – the insider threat

As you probably know, users are also called **layer 8** (as a joke) because they are on top of the 7-layer OSI model.

Another, more *professional*, way to call users is **insiders**. These insiders are a serious threat because they are already inside the network; therefore, many of our defensive systems and mechanisms (which are used to prevent users from accessing our network) will not apply to them.

Now, we will cover the different types of users that you need to consider when creating your cybersecurity strategy.

The inadvertent user

Based on a study from the Ponemon Institute, around 24 percent of data breaches are caused by *innocent* human errors. We call them innocent errors because they are normally user mistakes in which there is no user intention to cause harm to the data or the systems.

Many people believe that these kinds of incidents are rare or cause minimal impact. However, as you can see in the following diagram, a study from the Ponemon Institute, in 2020, shows a very different panorama:

Figure 4.1 – The cost of insider threats

I am going to summarize the most common mistakes or errors caused by inadvertent users, as follows:

- The use of weak passwords
- The repetition of passwords across systems
- The use of the same password for personal systems
- A lack of understanding of cybersecurity policies

- The misuse or abuse of privileged accounts
- Unattended devices
- The mishandling of data
- The installation of unauthorized software
- The inadvertent disruption of systems
- Careless internet browsing
- The use of free or open Wi-Fi
- "Click before think" (that is, in email attachments or links)
- The inadvertent disclosure of sensitive information

As mentioned earlier, these are *innocent* mistakes with no intention from the user to *harm* the company. However, there is another type of threat in which users are motivated to perform an attack, and they are known as **malicious insiders**, which we will be discussing next.

The malicious insider

First, let's try to gain an understanding of what types of motivation might cause an employee to turn into a malicious insider:

1. Offers from external attackers to provide data or perform actions in exchange for money
2. A lack of cybersecurity regulations and corporate sanctions
3. A lack of controls
4. The concentration of power
5. Bad management
6. Poor performance appraisals
7. Disagreements with corporate policies, strategies, and coworkers
8. Layoffs

The following figure shows the difference in terms of motivation between a malicious insider and an inadvertent user:

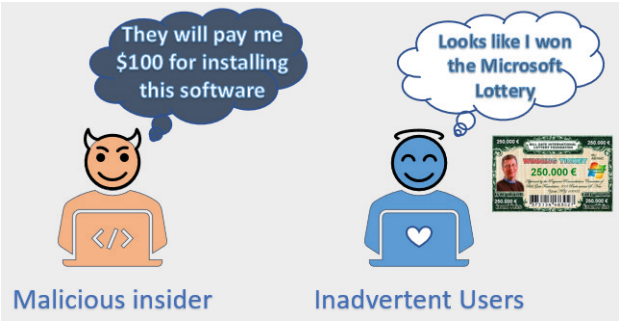

Figure 4.2 – Types of insider threats

As you can see in *Figure 4.2*, understanding those motivations will help you to work with management to create strategies to avoid users turning into *malicious insiders*. Additionally, implementing a training and education strategy will be your best ally to prevent mistakes from *inadvertent users*.

How do you spot a malicious insider?

Here is a list of *behaviors* or *actions* that can help you to identify a malicious insider before is it too late:

- The download of big amounts of data (or a dump of databases)
- After-hours access to systems and information
- Escalation of privileges
- The download of sensitive information without a business need
- The creation of accounts without following established processes and controls
- The increased upload of data to unknown external addresses
- Repeated access requests to sensitive systems or data
- Requests for exceptions of a given cybersecurity policy
- The increased usage of external storage devices
- Abnormal attachments on emails (by size or by the number of files)
- Evidence or signs of the execution of hacking tools
- Unexpected or increased connection of personal devices to the corporate network

If you don't think that you will face a malicious insider, then think twice. A study published by *inc.com* shows that almost one in five employees will be willing to sell their password to an external attacker, and as you can see in *Figure 4.3*, they will do so for a very low price:

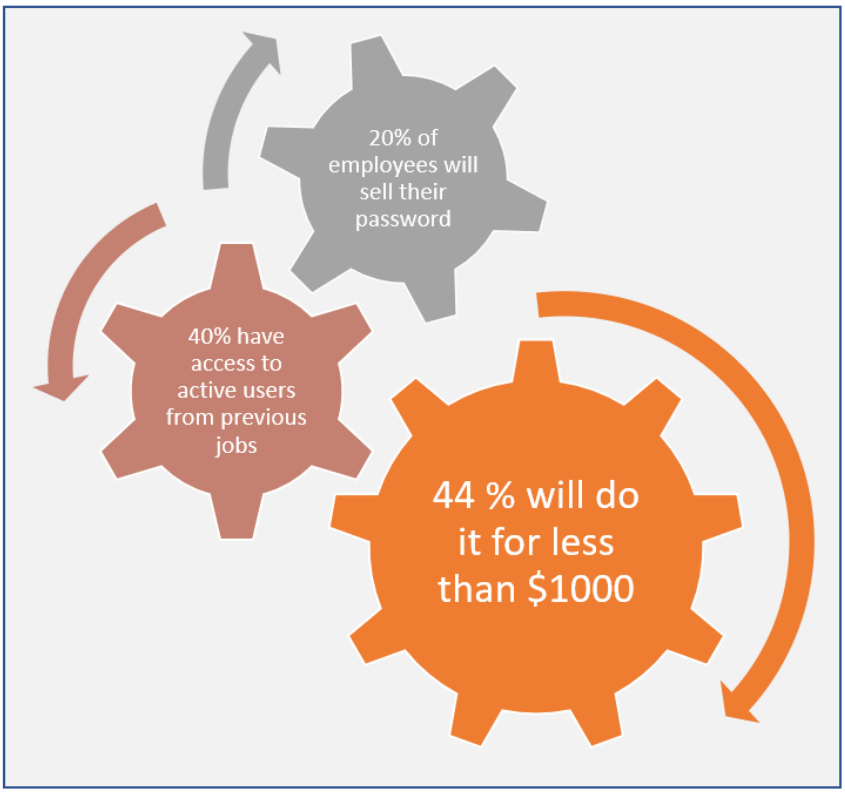

Figure 4.3 – The value of a corporate password

Now that you see that this is a serious threat, let's consider some actions that you can perform in order to reduce the probability and impact of risks associated with those malicious insiders.

Protecting your infrastructure against malicious insiders

Let's take a look at the tools, systems, and strategies that you can implement to protect against this threat.

The segregation of duties

This is one of the core activities that you *must* do as part of your defensive security strategy, and it is based on two main actions or activities:

- The first one is about the identification of *the most critical tasks on your infrastructure*. Here you need to ask yourself: what are the human actions that (if performed by a malicious insider) will cause a *considerable impact on the systems and data?*
- Second, once you identify those activities, you need to set controls to *ensure that a single person cannot perform those tasks by themselves*.

> **The importance of segregating duties**
>
> Researchers agree that the biggest hack to a social media platform (for example, the Twitter hack of 2020) could be prevented if a segregation of duties was put into place.

Now, let's take a look at some examples of how you can leverage and implement this strategy.

An example of the segregation of duties

As you can see in *Figure 4.2*, allowing a system administrator to create privileged accounts allows a malicious insider the possibility to perform a dangerous attack. Instead, you *must* put some systems and processes in place to elaborate a *flow for the creation of new users*, which requires the involvement of several groups, *reducing the probability of the attack*.

In the following example, you can see a flow in which the system administrator will have to create a ticket with the request. The request is then sent for approval, and once it is approved, it will be sent to the **Identity Management Access** (**IAM**) team to fulfill the request.

Notice that filtering all the communications through the helpdesk (in both directions) is a great way to prevent direct communications between the malicious insider and the person in charge of creating the accounts, which greatly enhances the security of this method:

Figure 4.4 – The segregation of duties

Another great example is related to backups because a malicious insider might know that deleting some files will not do any harm. This is because they can be retrieved from the backups, and in those cases, they will target the backups to prevent any restoration attempt.

To prevent this dangerous scenario, you can use the segregation of duties to ensure that a single user won't be able to delete those backups because a flow is in place to perform that action (*supported by a policy, a process, and enforced by the system*).

The use of mailboxes

When the segregation of duties is implemented, a malicious insider might try to persuade or convince another person to do some actions to help with the attack.

To avoid this, you can use **mailboxes** for the communication of highly sensitive teams, such as approvers, the helpdesk, the IAM team, and more. This will avoid exposing the identity of the people on those jobs, preventing any direct attempt to persuade or blackmail them:

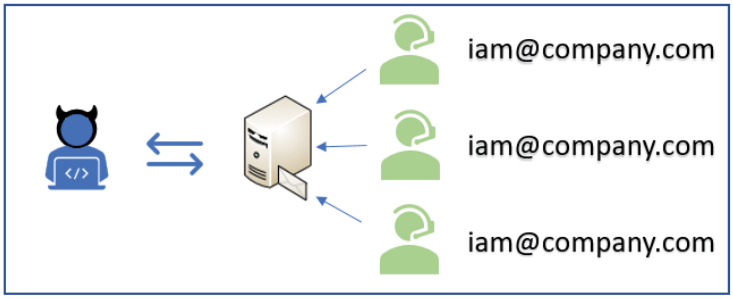

Figure 4.5 – Using mailboxes

As you can see in *Figure 4.5*, even if there is a direct channel of communication with support teams (which is normal in small companies), malicious insiders won't be able to identify who is the person is working on the request.

Job rotation

Another good practice is job rotation for IT support personnel. This consists of the creation of a policy that requires IT personnel to switch roles from time to time. This requires the implementation of *cross-training*, *mentoring*, and *skill-development programs*, which is also *motivational* for IT personnel.

This simple policy gives you some extra advantages in your defensive security strategy, including the following:

- **Reduces the risk of downtime**: You must ensure that you have people trained to avoid risks due to a lack of skills regarding a given technology. For example, "*Oh, we will have to wait until next week because Maria is out, and she is the only one who knows about DB2."*
- **Reduces the risk of fraud**: When a person stays "fixed" in the same role, they might be able to cover their own tracks (in the case of any illegal activity). However, if you keep rotating them, there is a reasonable probability that the new person will discover some "anomalies" that could uncover that illegal activity.
- **Decrease the impact of the attack**: By reducing the time for which a person is doing the same role, it will also reduce the time a person will have to perform an illegal activity; therefore, the impact of such an attack (on your data and systems) will be lower.

Mandatory vacations

This is based on the same principle as job rotations, and it serves as a way to *detect and stop frauds*. The way this works is very simple: first, it is known that insiders who commit fraud tend to be paranoid of being discovered. Therefore, they avoid taking vacations to prevent someone from finding their malicious activity.

Additionally, it will be easy to identify deltas in the activities between a new admin and the previous admin (who is now on vacation) that might lead to the discovery of malicious activities:

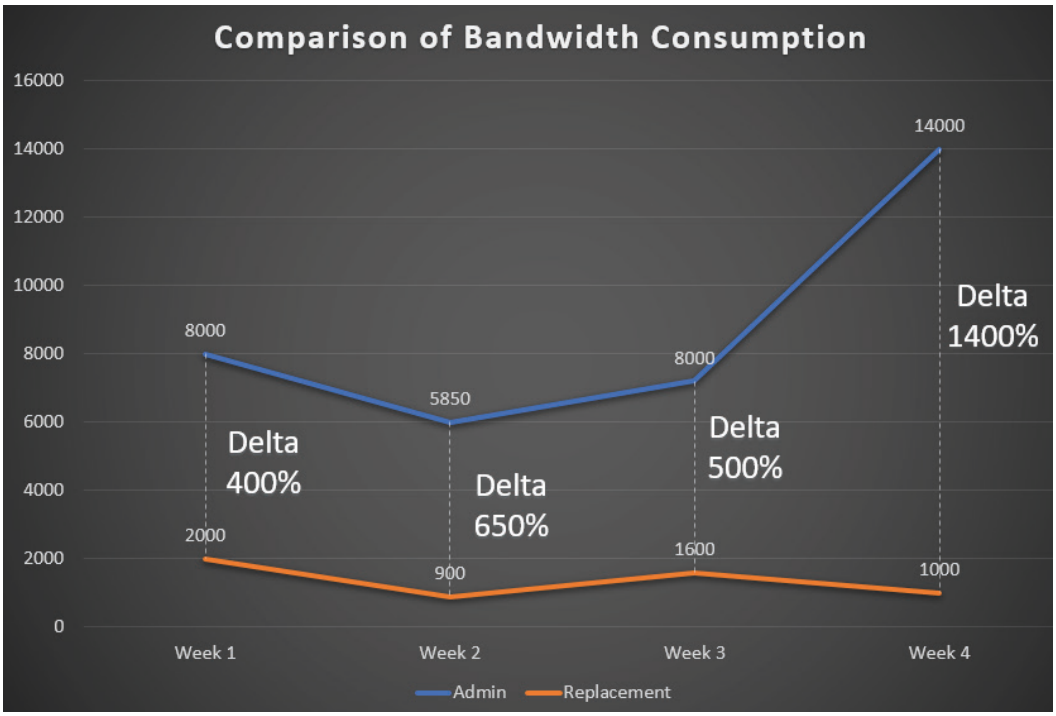

Figure 4.6 – The detection of malicious activity with mandatory vacations

Figure 4.6 shows a scenario in which a privileged user used to download more than 4 GB of data per week (which was considered to be *normal*) until the person was forced to take vacations and the use of bandwidth decreased by more than 400 percent, which confirmed that the user was engaging in unauthorized use of corporate bandwidth.

The analysis and correlation of logs

Unprocessed data might not trigger any alarms, but as demonstrated in the previous example, when it is properly correlated, it could show very interesting information.

Logs are a gold mine; however, you need to dig relatively deep to uncover useful information. One of the most basic ways to gather that information is by correlating data between users and systems to identify outliers.

Additionally, when doing the analysis, you need to determine which are the events that are higher or lower than the average and those are the ones that you should investigate further.

Also, there are many systems that automate the analysis of logs. So, instead of giving you some brands and names, I am going to show you the type of tools that will help you to achieve this, so you can search and find the solution that better suits your organization. Additionally, I would suggest you look for alternatives that leverage machine learning algorithms in order to improve detection and reduce false positives.

The systems are as follows:

- Behavior analytics systems
- Threat intelligence
- Anomaly detection
- Predictive alerts

However, I want you to know that not having those systems is not an excuse to waste your valuable data. In fact, I remember a very interesting scenario in which, by analyzing several logs, we found a system administrator that was illegally using corporate assets to *"mine" bitcoins*.

How did we find it?

By simply checking the logs, we discovered that several systems and non-production servers were normally turned on from 10 pm until 4 am with the sole purpose of mining bitcoins. Additionally, those logs contained the level of detail required to *identify the users* involved but also the evidence required to *pursue the associated penalties and sanctions against them*.

Alerts

Another great way to identify a malicious insider is by setting up monitors to give you alerts when a cybersecurity system is disabled by the user.

This is especially useful in those companies that give *administrative rights to all the employees* because they think they can bypass the security mechanisms (such as disabling the antivirus software or a firewall); however, what they don't know is that you are already one step ahead.

> **Important note**
>
> There are several ways in which to prevent a user from disabling some security features; however, not all companies or IT departments have the tools, knowledge, or interest in doing so, and that is why it is important for you to learn how to deal with these scenarios.

Now, let's take a look at one example of a very common practice in IT departments, which is a really bad practice for security.

Shared credentials

By default, the best practice says that *shared credentials should NOT be allowed in your infrastructure*. However, in case you do have them, you need to set up additional controls in place such as **Multifactor Authentication** (**MFA**), **Role-Based Access Control** (**RBAC**), and **Privileged Access Management** (**PAM**).

PAM works by *locking shared credentials* into a repository that can only be accessed by authenticated employee accounts (enabling accountability). Once the credential is used by the system administrator, the credential is *reset* for the next employee. Although PAM solves the challenge of shared accounts, it is very expensive to implement:

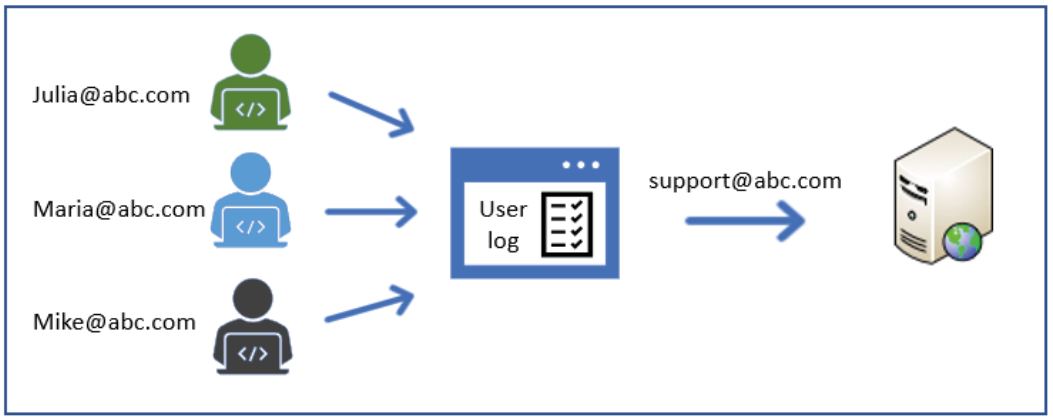

Figure 4.7 – An example of a PAM system

The preceding diagram is an example of a PAM system and shows how each user logs into a centralized system (to enable accountability), and then from there, *inject* the shared account into the server.

Audits

This topic was already covered in detail in *Chapter 3*, *Comprehending Policies, Procedures, Compliance, and Audits*; however, I want to highlight that *audits are one of the most effective ways to detect malicious insiders*, so make sure your infrastructure is audited regularly (either internally or externally).

Cybersecurity policies

As discussed in *Chapter 3*, *Comprehending Policies, Procedures, Compliance, and Audits*, policies need to be well defined and communicated. Additionally, those policies *must include the associated sanctions in case there are any violations to them*. Those sanctions are a great mechanism to dissuade malicious insiders, and that is why it is so important to make sure that all employees are aware of your cybersecurity policies.

We already talked about two types of insiders: the *inadvertent* users and the *malicious* insiders. However, there is another attack vector. In this attack, *an outsider will use the power of influence and psychological manipulation to convince or persuade an employee to perform a set of actions aimed to disrupt the systems or gather/modify sensitive data*. This technique is known as **Social Engineering**, which we will be discussing next.

Mastering the art of social engineering

Social engineering is one of the most fascinating topics in security. In fact, many experts define social engineering as an *art*: an art that requires a lot of social skills that enables the attackers to *gain access* to the victim's mind to gather personal information or even persuade the victim to perform certain actions that will benefit the attacker.

This is like hacking into the human brain to read the user's data or inject instructions for the human to perform.

As I mentioned earlier, this is a very exciting and important topic, so I will try to summarize it as much as I can.

> **Important note**
>
> As a professional in defensive security, you *must* master this topic because the better you understand how this works, the better you can defend against it.

Now, let's take a look at the attacks that are aimed to trick the user. Additionally, keep in mind that while not all authors agree on classifying these attacks as types of social engineering, the truth is that these attacks share the same concepts and strategies as social engineering attacks.

The social engineering cycle

There are many techniques that attackers can use to perform a social engineering attack, but they need to be orchestrated, as follows, to improve the efficiency of the attack:

1. **Information gathering**: First, the attacker will gather as much information as possible about the target individual or organization. The more the attacker knows about the organization, the higher the chances to succeed. For example, an attacker would be very interested in knowing the organizational structure, the processes, and procedures as inputs for the next steps.
2. **Building trust**: Here, the attacker will use the data gathered plus a combination of social techniques to gain trust. In more elaborate attacks, the attacker will need to gain trust from a plurality of individuals to bypass additional security layers before reaching the real target or victim.

   As you can see in the following screenshot, the attacker might also leverage some technical knowledge to gain the trust of the victim. For example, the attacker might impersonate an IT person by telling the user that their computer was reported as infected by a virus and ask the user to check whether the `svchost` Windows process is presented on the Task Manager:

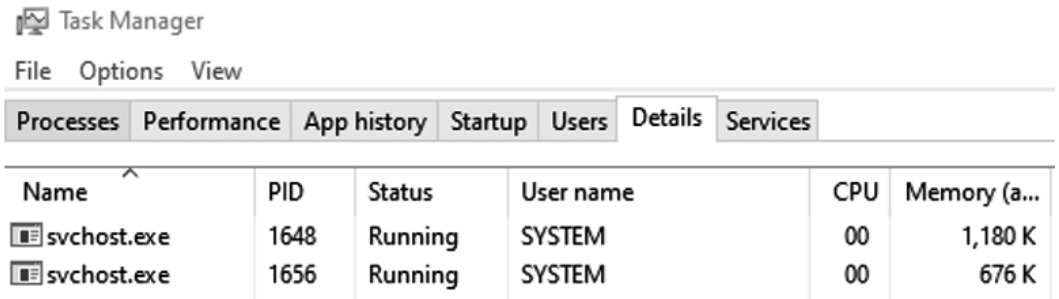

Figure 4.8 – SVCHOST running on Windows

   Of course, the attacker knows that such a process will be *always* there, and so when the user finds it, it will be a way to legitimize the attacker, gaining full trust and opening the door for the next step.
3. **Influencing the victim**: Having gained the trust of the user, the attacker can manipulate the victim to either provide some confidential information (such as usernames and passwords) or to perform some actions (such as resetting a password, opening a terminal, and opening a web page).

4. **Executing the attack**: By now, the attacker might have valid credentials, full remote control of the computer, and many other paths in which to execute their final attack (such as deleting, modifying, or copying confidential data, accessing a given system, and more).
5. **Erasing tracks**: Once the attack is completed, the attacker may want to *cover their tracks* to avoid detection and prosecution, but also to retain access to the systems and data for a longer period of time:

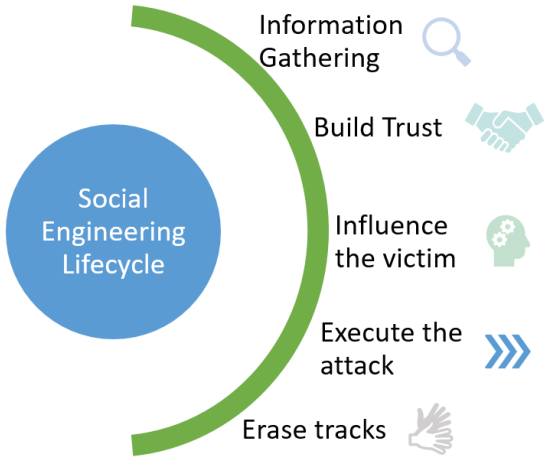

Figure 4.9 – Social Engineering Lifecycle

Now, let's take a quick look at some of the techniques used by attackers to successfully launch a social engineering attack.

Social engineering techniques

Here are few techniques that an attacker can use to launch a social engineering attack:

- **Impersonation**: One of the most common techniques used by attackers is to present themselves as someone else; for example, as someone with authority, someone with power, or someone representing a reputable company or group.

  This is normally used to gain trust from the victim in order to either gain information or to make the victim execute a given action.

  Some of the most common impersonations are impersonating an IT person, a government representative, a bank employee, or a reputable business.

- **Fear**: Attackers can use fear to persuade the user to comply with a given action. For example, imagine an email that says the following:

  "Your computer is infected, click here to scan before the computer is blocked and blacklisted from the corporate network."

- **Reciprocity**: The attacker will do something that appears to be beneficial for the victim. That way, the victim will be prone to comply with the attacker's request (either to provide some information or perform some action) to return the favor.

- **Exploiting user greed**: This will exploit a fundamental human weakness, for example:

  "You won a cruise vacation, click here to claim your prize!"

- **Exploiting user curiosity**: In this scenario, an attacker might drop some malicious USB drives near the target, hoping that an employee will pick them up and plug them in. Attackers could put a label on the USB such as *"My pictures"* or *"Confidential"* to increase the level of curiosity and, therefore, increase the effectiveness of the attack.

  As a fun fact, most sources consider that *Stuxnet* (the virus that damaged Iran's nuclear program) was spread by infected USB drives.

- **Social validation**: Another tactic is to use social validation to push the victim. For example, an attacker could tell you that *"This was already tested by other systems administrators,"* to give you some sense of assurance that the request is safe because it has already been carried out by others.

- **Technical validation**: Attackers can use technical jargon to confuse the victim. Normally, this is used in conjunction with other techniques such as a sense of urgency and fear. *Figure 4.8* is a great example of the application of this technique.

- **Authority figures**: The attacker might impersonate an authority figure to make you comply with a given request. In some cases, the attacker might not impersonate the person but say they are acting on behalf of the authority figure instead. For example, *"If you do not install this software, it will be escalated to Mr. Satori ."* Notice in this example they called the person by the name (Mr. Satori) instead of by the title (CEO), which is a technique used by the attackers.

- **Scarcity**: Here, the attacker will make the victim think that if the action is not performed quickly, the user (victim) could lose a potential reward. For example, imagine a supposed email from IT that says the following:

  "We have 20 new MacBooks to upgrade old computers; click here to fill in the form. Remember there are only 20 laptops and they will be provided to the first 20 that complete the form (first come, first served)."

- **Sense of urgency**: This classic scam could say the following:

  "If you don't reset your password in the next 30 minutes, your computer will be blocked from the network."

  "Your computer is infected, click here NOW, before your information is stolen."

Figure 4.10 describes the entire flow of a social engineering attack and the tactics used by attackers:

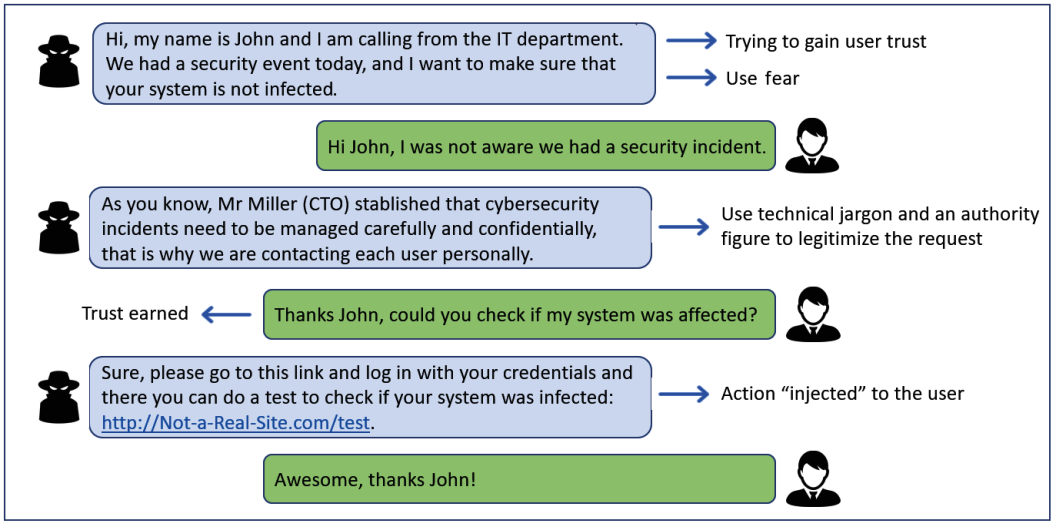

Figure 4.10 – An example of a social engineering attack

At this point, you know the flow of the attacks and the techniques used by the attackers to gain the user's trust and perform the attacks.

Now, it's time for us to take a look at the most common types of social engineering attacks where those techniques are used.

Types of social engineering attacks

Here, I am going to summarize the most common attacks that are based on social engineering techniques.

Phishing

As you probably already know, this concept is very simple. The attacker sends a fake email trying to impersonate a reputable person or company. To increase their chances of success, the attacker will first try to convince the victim that the email is legitimate (by using company logos or impersonating an email account or domain) and then request the user to take some action, normally to access a link or open an attached PDF.

Let's view some examples next.

Everyone dreams of free money, and attackers know that. So, to leverage that human desire, an attacker will impersonate a company that wants to transfer your money but claim that they were unable to do it because the number was incorrect. Then, to be able to *gain* that money, you just need to open a *harmless* PDF that, of course, will be packed with all kinds of viruses, from a keylogger to deadly ransomware.

In *Figure 4.11*, we have highlighted the common aspects to help you to identify these types of attacks:

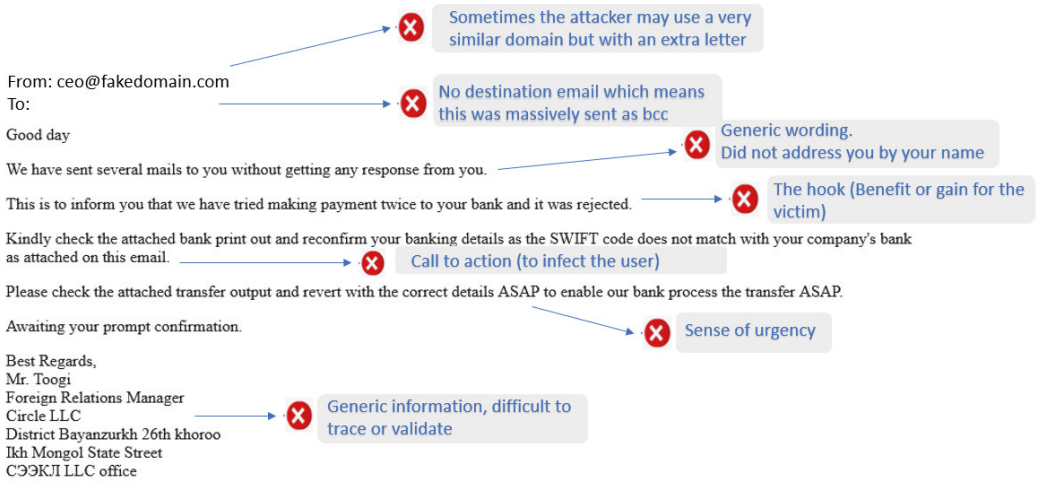

Figure 4.11 – A "free money" phishing example

Everyone shops online these days, and the probability that you are expecting a package is very high, so attackers leverage that and indiscriminately send these types of phishing emails, hoping that anyone that is expecting a real delivery falls into the trap.

As shown in *Figure 4.12*, the attacker will impersonate a well-known delivery company, but there are still several items that you can use to confirm that it is a phishing email.

The first one (which might look very obvious) is that the address is from Hotmail.

The second one is that your email is not listed in the **To** field.

The third one is the use of a generic salutation. However, the most important one is regarding the link.

Here, you can hover your mouse over the link to see where it is pointing. In this example, it is very obvious that the link is not pointing to the real DHL domain:

Figure 4.12 – A package delivery phishing example

As you know, attack vectors normally evolve to bypass defensive mechanisms, so, let's take a look at some social engineering variants that have evolved from phishing attacks.

SMishing

This attack shares the same attributes as a phishing attack, with the only difference that this is distributed by **Short Message Service** (**SMS**).

While this might sound like a small change, this is one of the most dangerous attack vectors because users do not associate old SMS with viruses, so they tend to *trust* these messages and fall into the trap. If your company has a **Bring-Your-Own-Device** (**BYOD**) policy, then you must do the following:

Educate the users on this type of threat.

Disable hyperlink capability in SMS.

Now, let's take a look at *Figure 4.13* to show you how to easily spot these threats:

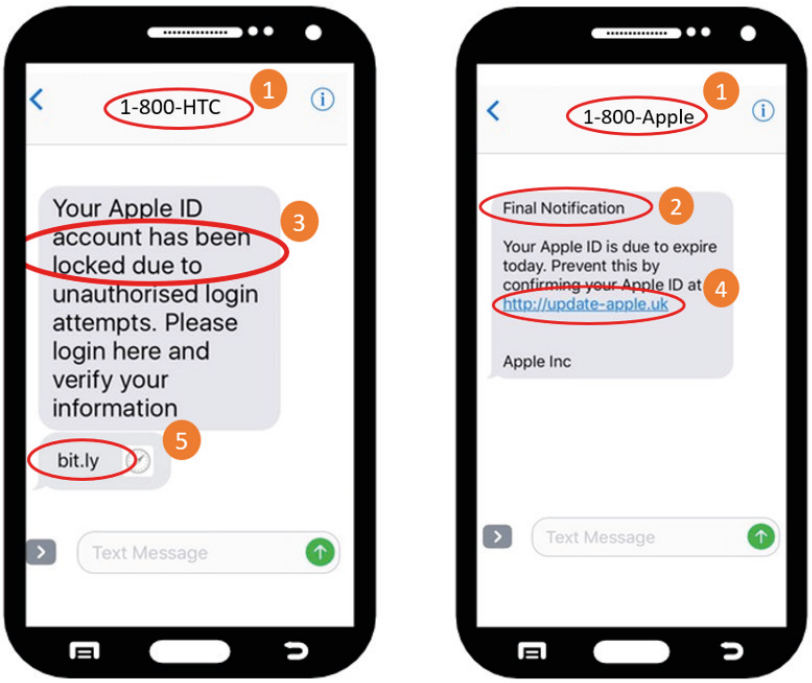

Figure 4.13 – A SMishing example

1. Most of the time, attackers will use a random number that is easy to spot and block. However, in more elaborate attacks, the attacker might use some tools (such as **BurnerApp** and **SpoofCard**) to spoof the caller ID and impersonate a more credible number.
2. Again, another common factor here is the *sense of urgency* to perform the requested action.
3. Another common factor is the *use of fear* to make the user believe that something bad will happen if the requested action is not performed.

4. As mentioned earlier, here, the main attack vector is *sending the victim to a malicious website* that will either steal your credentials, infect your device with malware, or both.
5. In one possible scenario, the attacker will purchase a domain that looks like the original; however, in other situations, the attacker will use a *link shortener* to mask the name of the site.

Spear phishing

This is a targeted attack in which the attacker starts by conducting in-depth research on the victim and the company. Then, the attacker will use all of that knowledge to create a *customized phishing attack*.

Normally, these types of attacks are targeted at high-value targets such as managers, financial personnel, or system administrators (because of the value of their administrative credentials).

Now, let's analyze a real example of spear phishing (as shown in *Figure 4.14*):

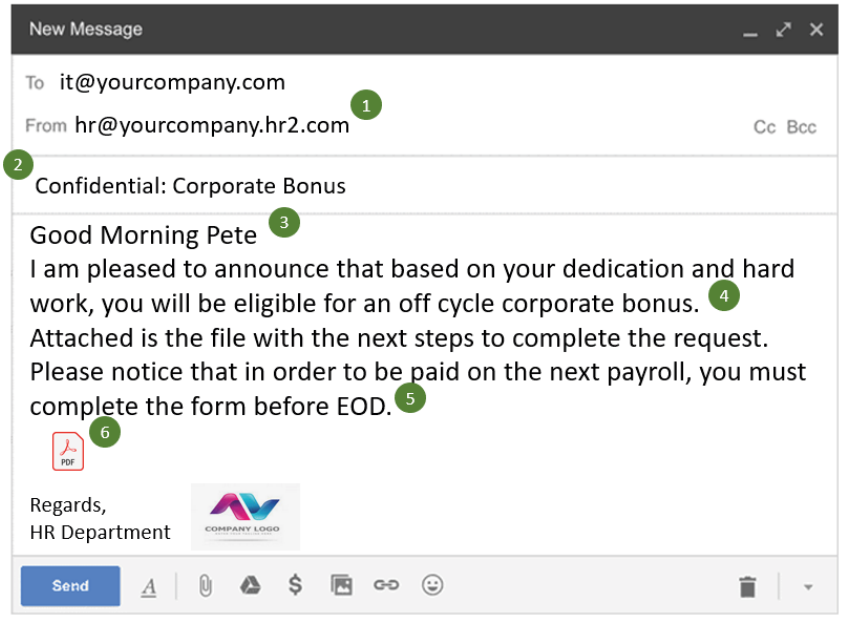

Figure 4.14 – An example of spear phishing

From the preceding screenshot, we can decipher the following:

1. Most of the time, an attacker will *use a very similar domain to trick their victim*. As you can see, in this example, in more elaborate attacks, the site might look legitimate. However, if you look carefully, you will see that the name of the company is just a subdomain of the attacker's domain.
2. The attacker uses a *catchy subject* but also with some sense of urgency (to prevent the attack from being discovered).
3. The email will be properly directed to the victim by *using the real name and title* (in some situations, the attacker will even use nicknames to reduce suspicion by sounding familiar).
4. A hook will be used to catch the victim's attention and persuade them to open the malicious file.
5. A *sense of urgency* is used, which motivates the user to execute the requested action (for example, to open the attachment) without further verification with management or any other employee who could identify this as a potential attack.
6. Again, an apparently *innocent PDF* will be the gateway used by the attacker to finally execute the attack (for example, installing a particular ransomware, opening a backdoor, or installing a keylogger).

Vishing

Also known as phone elicitation or phone scams, vishing is a type of phishing based on a phone conversation between the attacker and the victim in which the attacker will try to convince the victim to perform a series of actions or to inadvertently disclose some kind of confidential information.

This is one of the most complex attacks from the attacker's point of view because it requires the attacker to master most of the social engineering concepts that we previously covered.

However, the attacker can also augment some of these techniques to ensure compliance from the victim. For example, the attacker could call a help desk and request a reset for a password and ask the agent to provide the password over the phone. If the agent denies the request, the attacker could threaten the agent that they are about to close a multimillion deal, and if the password is not provided over the phone, then the deal will not be signed and the agent will be held responsible. This simple trick confirms why this attack is the preferred mechanism for experienced social engineers.

> **Phishing in numbers**
>
> Phishing is the most common type of social engineering. In fact, *Verizon's Data Breach Investigations Report 2019* showed that more than 30 percent of confirmed data breaches were associated with phishing attacks. On the other hand, *global losses for vishing attacks are estimated at $46 billion*.

As you can see, attackers are good at finding new and clever ways to expand or evolve their attacks, so you must stay up to date to uncover any new and potential variation of phishing that could impact your employees.

Scareware

Scareware is based on deceiving the victim into thinking that the computer was infected by a virus, whereas the reality is that the computer is fine. The aim of the attacker is to then convince the victim to install *antivirus software* to delete those viruses, but the *antivirus software* is fake.

There are two main variations of this attack. The first one is based on a *"free antivirus"* that is, in fact, an actual virus that *"opens the door"* to additional viruses.

The other variation is based on selling you *"antivirus"* software that will *"remove"* some viruses that do not even exist (so, it is essentially fraud).

Normally, this will come in several ways, such as the following:

- A popup from a malicious website
- An add-on for a legitimate site (for example, YouTube)
- A script that will be executed when Windows starts

- A fake antivirus program:

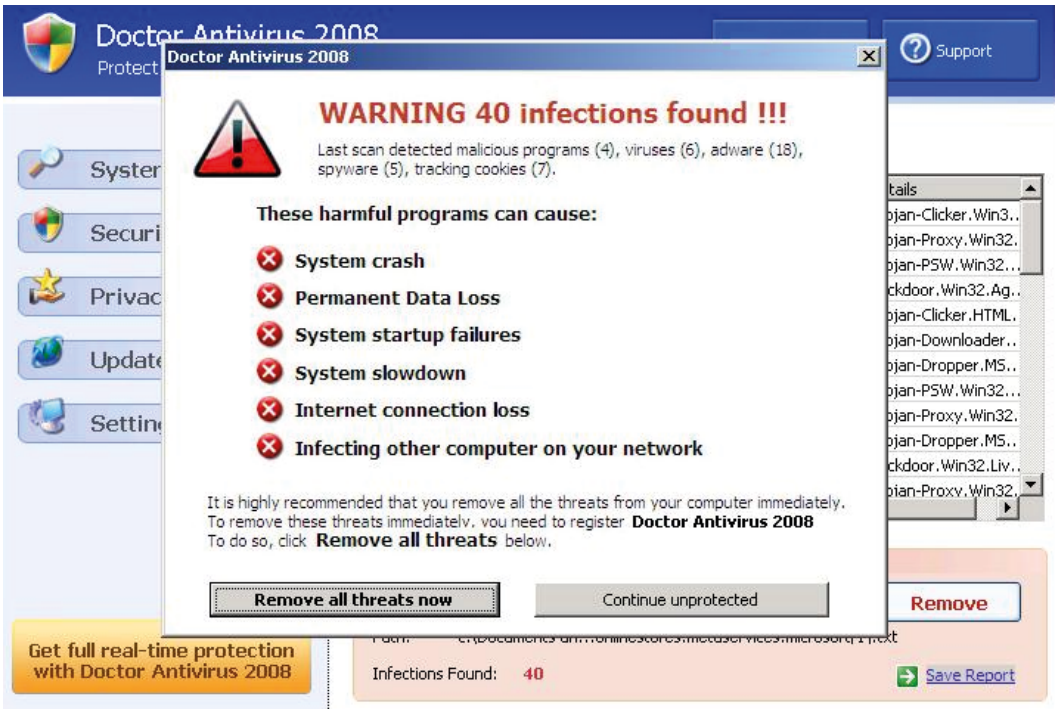

Figure 4.15 – Scareware examples

This type of threat was very popular a few years ago (that is, in the Windows XP era); however, nowadays, they are less common, though it's still dangerous.

However, in recent years, this threat seems to be moving to another platform and they are now targeting smartphone users.

As you can see in *Figure 4.16*, the attack is very similar. To prevent this, the Play Store constantly bans these apps; however, they just keep appearing under another name:

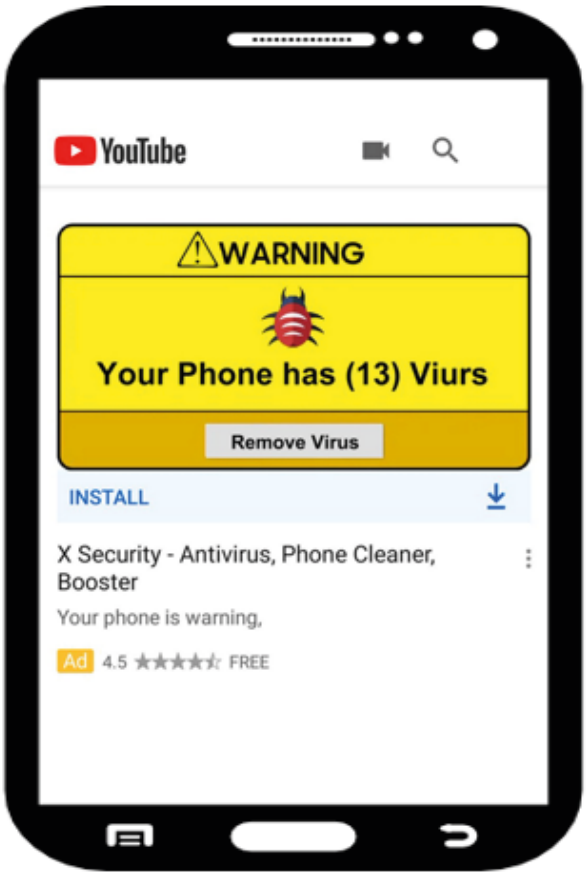

Figure 4.16 – Scareware on smartphones

One way to prevent this threat on corporate workstations is to limit the permissions to install third-party apps (which is a topic we will cover in the *Defending against social engineering attacks* section). This limitation also applies to smartphones and *must* be enforced if the company allows employees to access their systems with their smartphones (BYOD).

A great way to achieve this is by leveraging **Android Business**, which enables companies to create a virtual environment in which they can apply more controls to safeguard their data and to control who has access to their networks, systems, and data.

For more information, please visit their official site at `https://www.android.com/enterprise/`.

Additionally, deploying an **ad blocker** is a great idea. This can be done at three levels:

- Locally on the workstation
- Using the corporate firewall
- Using a DNS

> **Creating your own DNS ad blocker**
> In *Chapter 10*, *Applying IoT Security*, I will show you how you can create your own **ad blocker** DNS for less than $50 using a **Raspberry Pi**.

While Scareware was very popular in previous versions of Windows, this attack is still relevant not only to protect your infrastructure but also to protect people's money.

Baiting

Primarily, this is a technique used by attackers to exploit the victim's curiosity and trick them into a trap. Here, the main goal of the attacker is to make the victim access a fake or bogus web page, open an infected file, provide their credentials on a fake page, or download a trojan.

Some examples of baiting are presented in the following diagram:

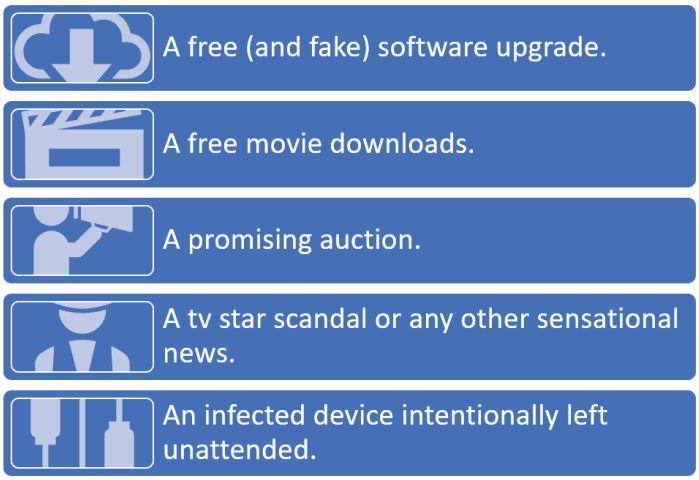

Figure 4.17 – Common baiting examples

There is also a subtype, called **clickbait**, that is mainly focused on presenting very interesting fake news with the hope that the user will click on it. Most of the time, clickbait is used to generate traffic or earn money with ads; however, there are also cases when they are used to infect a system with malware.

Shoulder surfing

This might sound very basic, but a lot of information is leaked using this simple method. Essentially, it involves looking over the shoulder of the victim to gather sensitive information, such as usernames, passwords, and more.

This is normally done by outsiders, so having a strong physical security system is key to prevent this type of attack.

Another recommendation for employees who constantly travel is the use of *privacy screens* that prevent others from reading your screen. Additionally, the use of *password vaults* also reduces this risk, because passwords don't need to be typed and, therefore, there is no risk of disclosure.

Tailgating

Now that we have mentioned physical security, it is time to talk about tailgating.

This is one of the most common methods used by attackers to gain physical access to a restricted location.

Here, the attacker will leverage the human characteristic of being *polite or friendly* to keep the door open for the person behind you.

Attackers are very creative and, many a time, they will carry a big pizza box or a couple of cups of delicious coffee as an excuse for not using a badge to access the building and hoping that a "good person" (that is, an inadvertent user) keeps the door open for them.

Besides the application of user training, the best way to fight this type of threat is by using additional verification mechanisms such as cameras to detect outsiders. In fact, cameras can now be used to detect outsiders using other mechanisms beyond face recognition. These include movement patterns, user counts (if two people enter but the system just reads one badge), analytics (based on the detection of unusual paths), and more.

Dumpster diving

One of the most famous hackers, *Kevin Mitnick*, who, in fact, was the first *hacker* to appear on the FBI's most-wanted list, made this tactic famous (this is explained very well in his books).

He commented that he was able to obtain a lot of information by simply searching in the company's trash, looking for unshredded documents with confidential information. In some cases, the attacker might be lucky to find sensitive information such as user credentials; however, in other cases, the attacker will use this to gather important information about the company, which can be successfully used to perform other attacks (such as impersonation).

To avoid this attack, you *must create a data classification and management policy* that clearly defines the following:

- The different types of documents (such as sensitive, confidential, public, and more)
- The appropriate way to dispose of each document type
- The appropriate way to dispose of physical documents (for example, notes, books, sticky notes, and more)

These policies are easier to enforce when users are at the office. However, with an increasing number of users now working from home, you must apply additional mechanisms to ensure these policies are being followed and that your users have the appropriate tools to carry them out. For example, provide a shredder machine to users with sensitive information or restrict the printing of sensitive documents at home.

Quid pro quo

This is a very interesting attack in which the attacker provides some benefit to the victim for free.

A classic example is when an attacker calls the employees of a given company impersonating an IT person doing a callback that is related to an open ticket. As you can see in the following example, the victims are very likely to *"take advantage"* of the call to get something fixed, while in reality, it is the attacker who will take advantage of them:

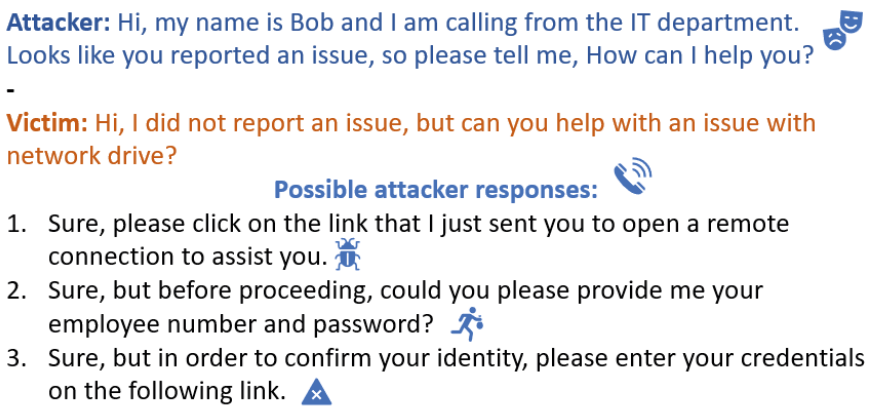

Figure 4.18 – A quid pro quo attack

The best way to prevent this attack is by implementing (and communicating) a policy about IT support, which states the following:

- IT will never call you from an external or blocked number (if possible, assign a friendly number for all IT calls such as 114).
- IT personnel will NEVER ask for your password.
- Never give your password either by phone, email, or text: **NEVER**.

Another good idea is to establish a *two-way validation callback mechanism*. This means that if a user gets a call from IT, then the user will have to call them back (using the official help desk number). This callback will serve as a secondary verification method for the IT person and the employee.

Social media ransom

This is one of the latest attack that is taking place. Here, the attacker will apply multiple techniques to get access to the social networks of your company (that is, Facebook, Instagram, and WhatsApp). The attacker will make sure that they change all your mechanisms to restore your account easily, so while you can contact the social media firm to get back the access, your attacker will have access to it until the issue has been resolved, which could be hours or even days. Attackers know that many companies will not take the risk of leaving their social accounts under the control of the attackers (because of the damage to the brand, customers and followers), so here is where the attackers demand some payment (normally bitcoins) to give the account back to your company.

Here are some tips on to prevent this dangerous attack:

- Always use MFA.
- Use strong passwords. These types of accounts should be managed with a password vault, so why be shy? Use the maximum length, use special characters, and make it immune to dictionary or brute-force attacks.
- Make sure each password is unique.
- Change the password frequently (at least every 3 months). This will be managed by the password manager, so the effort required is just two clicks, four times a year.
- Keep the access to these accounts to the very minimum number of people in order to reduce risks.

Additionally, *make sure that those people managing these accounts* (such as social media managers) *are well trained in cybersecurity* (to prevent these types of attacks).

Extorsion

In this scenario, the attacker will try to convince the victim that their computer or smartphone was hacked and that some sort of private or compromising information will be released if the attacker's demands are not fulfilled in less than 10 hours (of course, they will use the sense of urgency tactic).

As you can see in *Figure 4.19*, one method is based on telling the victim that their computer was hacked, and to prove it, the attacker will paste a password from the victim inside the email:

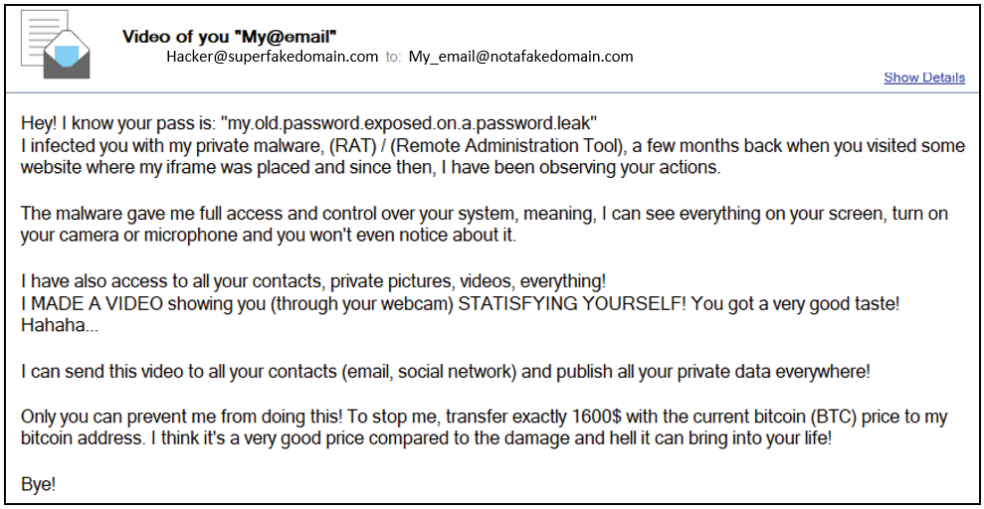

Figure 4.19 – An extorsion email

If the attacker knows the password, does this mean that they have really hacked the victim?

Absolutely not!

Here, the attackers leverage the information from previous data leaks, look for email/password pairs, and use that in the attack.

Therefore, while most of the time, the password provided is an old password, the victim will recognize it as one of their passwords and, therefore, is very likely to fall into this scam.

The best way to prevent this attack is by launching a campaign to explain to your employees **how attackers can get hold of your old passwords**. You can take the following steps:

1. Provide a brief explanation about what a data leak is and provide some examples of recent data leaks in big companies (for example, LinkedIn's data leak, Yahoo's data leak, and more).

2. Ask them to check whether their accounts were compromised in any of those leaks. There are several sites to do this, but not all of them can be trusted. As shown in *Figure 4.19*, one of the most trusted/used sites is `https://haveibeenpwned.com/`. This site will show you in which data leak your email was found, so you can go ahead and secure those accounts.
3. Provide them with a list of actions if their account was found in a known data breach, for example, change your password, ensure that you never use a variation or similar password, use MFA, or delete the account (if not in use):

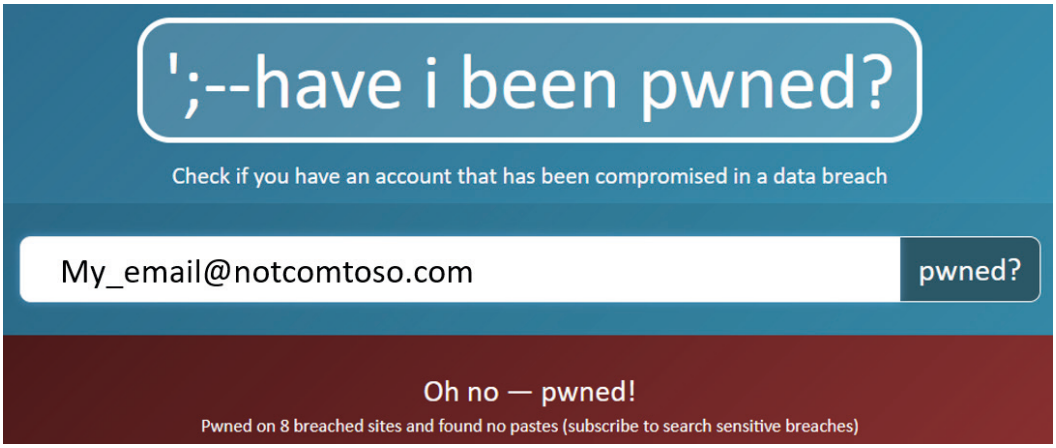

Figure 4.20 – A website to check for breached credentials

By now, you should know the tactics used in social engineering attacks as well as the most common types of attacks.

We have covered some defensive techniques. Now it is time to learn additional **best practices** that apply to **all** of these types of attacks and will help you to reduce the risks related to these threats.

Defending against social engineering attacks (patching layer 8)

"Companies spend millions of dollars on firewalls, encryption, and secure access devices, and it's money wasted because none of these measures address the weakest link in the security chain: the people who use, administer, operate and account for computer systems that contain protected information."

– Kevin Mitnick

Let's learn how to effectively protect your company against these threats.

Creating your training strategy

As you know, *patching* is one of the most important strategies in defensive security, and this strategy can also be applied to people through *education and training*. Therefore, you *MUST* invest time and other resources to make sure you have a **strong training strategy**.

Let's take a look at the key points that you need to consider when creating your own training strategy:

1. *Personalize it* based on your company culture, your threat landscape, and the type of data managed by the company.
2. For smaller companies, you can create a single training session to cover all employees; however, mid-to-large-sized companies and corporations *MUST have different types of customized training*. This training can be segregated based on the type of employee, the organizational level, the data managed, or data access.
3. Define the *delivery method* (for example, live training, webinar, videos, animations, web-based interactive learning, and more).
4. Define the *frequency* of the training.
5. Define the success criteria to *"pass"* the training, for example, by scoring at least 80 percent on the final assessment.
6. Define a *rewarding schema*; for example, providing a digital badge that can be shared on social media.
7. Get buy-in from HR and the senior management to *make the training mandatory*.

> **Tip**
>
> Make the training as interactive as possible, use up-to-date real-life examples, include everything (never assume a topic is too basic to be included), and use the list of attacks that we've just reviewed as a baseline to make sure all major attack vectors are covered.

You need to convince upper management that companies are not **spending** money on cybersecurity education; instead, they are **investing** in securing the most vulnerable cybersecurity factor.

Admin rights

This is a controversial topic because there is no consensus regarding whether giving admin right to all employees it is a good practice or not. However, *from a security standpoint, there is no question that giving administrative rights to all of your employees increases your threat landscape.*

Therefore, you should always push to avoid giving admin right to all users; however, if your company decides to grant admin rights to all employees, then you need to take the following countermeasures:

1. Define a clear policy about software installation.
2. Create a whitelist and blacklist of applications that can be installed.
3. If possible, create a repository to host the whitelisted software (this reduces the risk of a user installing hacked versions of the software).
4. Set alerts if a blacklisted software is installed on a corporate workstation.

Implementing a strong BYOD policy

If you allow employees to use their personal devices for work, then make sure that you have a strong BYOD policy in place.

Additionally, this policy *must* be supported by systems and software to enforce it.

Performing random social engineering campaigns

The best way to evaluate the level of preparedness or exposure that users have against a social engineering attack is by testing them with real-life controlled attacks.

Here is how you can do this:

- **Set up your environment**: Purchase a domain from where you will launch the attacks. Look for similar names as the ones that a real attacker will use, for example, *support-companyname.com*.
- **Test one attack per cycle**: First, you need to define the cycles, such as every 3 months, 6 months, or 1 year. For example, phishing in early 2020, baiting in late 2020, quid pro quo in early 2021, and extorsion in late 2021.
- **Analyze the results**: The goal of this is *not* to chase after your employees and put them on a "wall of shame." Instead, this is about gathering intelligence to determine areas of improvement for upcoming training and education.
- **Set up rewards**: You can set up a rewarding system for those employees that found the *attack* and used the proper channels to report it to the cybersecurity team.

> **Rewards are not always about money**
>
> You can also leverage free perks such as digital badges, a wall of fame, a secure employee of the month (you might want to use a catchier name such as "*Cybersecurity Rockstar*"), a preferential parking spot for a month, or more.

- **Announcements and communications**: You might not want to spoil your assessments by communicating the start of it. However, it is a good idea to send a communication *after* the assessments are finished so that people are aware of these types of initiatives, but also to share some relevant numbers with them (for instance, how many people fell victim to the attack, potential losses, and more).

To avoid disruption of services, we recommend that you roll out these campaigns to a randomly selected group of individuals (depending on the size of the company, this could be between 10 percent to 60 percent of employees). Additionally, you need to ensure that this random selection includes participants from all the organizations across the company (for instance, HR, sales, and IT).

Summary

In this chapter, you learned all about users, including how they can impact your defensive security strategy, their vulnerabilities, and the plurality of attacks aimed at them, but also all the tactics that you can apply to mitigate those risks.

This chapter is extremely important because by securing this attack vector, you will exponentially reduce the scope of the attacks against your infrastructure, systems, and data.

Now, get ready for the next exciting chapter in which we will take a deep dive into more technical stuff. In the following chapter, you will learn about the best penetration testing tools, forensics, networking, and many other technologies that you need to master in order to create the best defensive security strategy.

Further reading

Here is the complete report of *The Cost of Insider Threats: 2020*:

`https://www.ibm.com/security/digital-assets/services/cost-of-insider-threats/#/.`