CyberSecurity Programme

You receive a Certificate of Achievement when you complete the course.

What you will learn

Over the next 12 months, you will learn about:

- Social engineering attacks
- How to keep your mobile phone secure
- Types of malware
- Malicious bots and denial-of-service attacks
- Network security and data control

Protect your device

This week, you will:

- Meet the team here at the SHG Academy
- Meet and introduce yourself to the other learners on this course
- Learn about the three pillars of cybersecurity
- Find out different ways in which attackers may try to gain access to your devices and accounts
- Consider the Computer Misuse Act and its efficacy in legislating against cybersecurity crimes

What are you most looking forward to learning about? Let us know in the comments section.

All of the content based on the latest pedagogical research and SHG Trainer feedback. It provides an innovative progression framework where computing content (concepts, knowledge, skills, and objectives) have been organised into interconnected networks called learning graphs.

Using the course content

Cybersecurity & data

In this step, you'll learn about the kinds of data that need to be protected and why this data is valuable to others. You will also be introduced to some core concepts of cybersecurity, and its aims and objectives.

What is data?

In this section, the term **data** refers to the information that we create and use when interacting with online services like apps and websites, or the information that we store on electronic devices. We might generate this data intentionally, or we might not know that it exists. All of this data has the potential to be valuable and needs to be protected.

Examples of data include personal information, content, information about user behaviour, and data about other people.

Personal information is any information about individuals, for example bank details or forms of ID such as driver's licences. We often share our personal information with others to prove who we are, or to access services.

Content includes social media posts and private messages that we generate when interacting with websites, and the content that other people create that we consume and share. It is also the content that we create for private networks: our work files and personal documents can all be considered data, even if they are not shared publicly.

Information about **user behaviour** is information about the way in which we produce data, for example, our browsing history, the time of day we perform a search, and how long we spend browsing.

Data about other people is any information that we hold about another person. This could include information about our relationships with other people, or any of the types of data described previously.

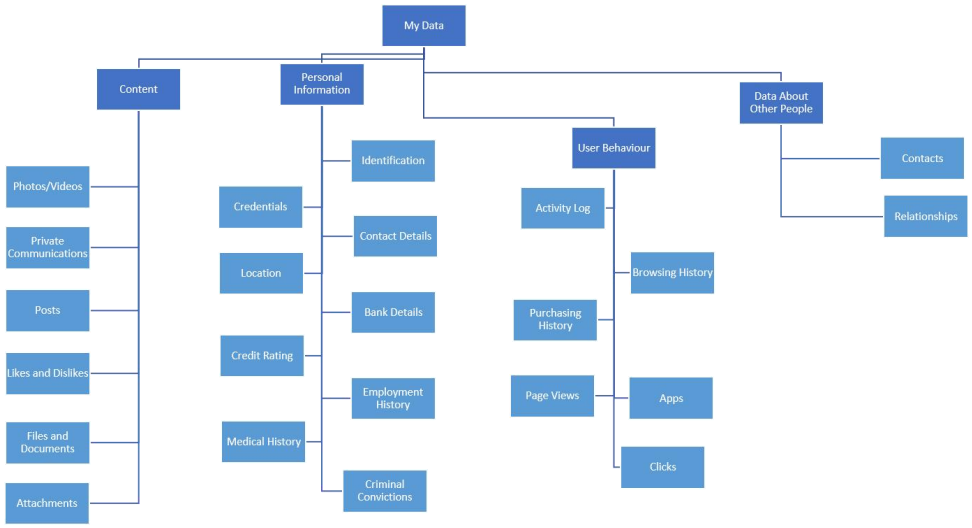

How much data do you think you have generated today? It is estimated that we collectively generate 2.5 quintillion bytes of data each day, which is the equivalent of nearly 7 billion days worth of YouTube videos.

The data economy

Much of our data is valuable to us, but it is also valuable to other people and to companies. For example, user behaviour and content is valuable to companies because it helps them to optimise their service for you, it increases their ability to sell to you, and it makes their product more efficient.

For example, a clothing company may be interested in acquiring your browsing history as it might tell them what sorts of products you are trying to purchase, enabling them to tailor an advert to your interests. A telemarketer might wish to buy your phone number so that they can contact you about their products.

This data has its own market value and is sold between companies. There are specific companies called **data brokers** that collect and sell customer data.

In addition to the data economy, there is an underground trade of stolen data. Data such as login details for website accounts, bank account details, and contact details is sold in batches and could be used by attackers to hack your accounts, steal your money, or impersonate you. Even data that may seem harmless, such as the name of your first pet or your favourite number, has market value, because it can be used to guess your passwords.

Organisations have a responsibility to protect your data from malicious attackers. They use cybersecurity processes to do this. You can also use cybersecurity practices to protect your own data.

Cybersecurity

Cybersecurity is the practice of protecting data from threats such as theft and corruption. These threats include attacks as well as accidental breaches.

Cybersecurity is often context-specific, so different individuals and different sorts of organisations will have different requirements. They will also have different resources available to implement their cybersecurity policy.

However, all cybersecurity approaches should utilise the **three pillars of cybersecurity**. These are **people**, **processes** and **technology**. This is because good cybersecurity practice requires a knowledgeable and motivated person who is able to enact good practice, a clear process that defines what is expected of them, and the right technology so that they can follow this process.

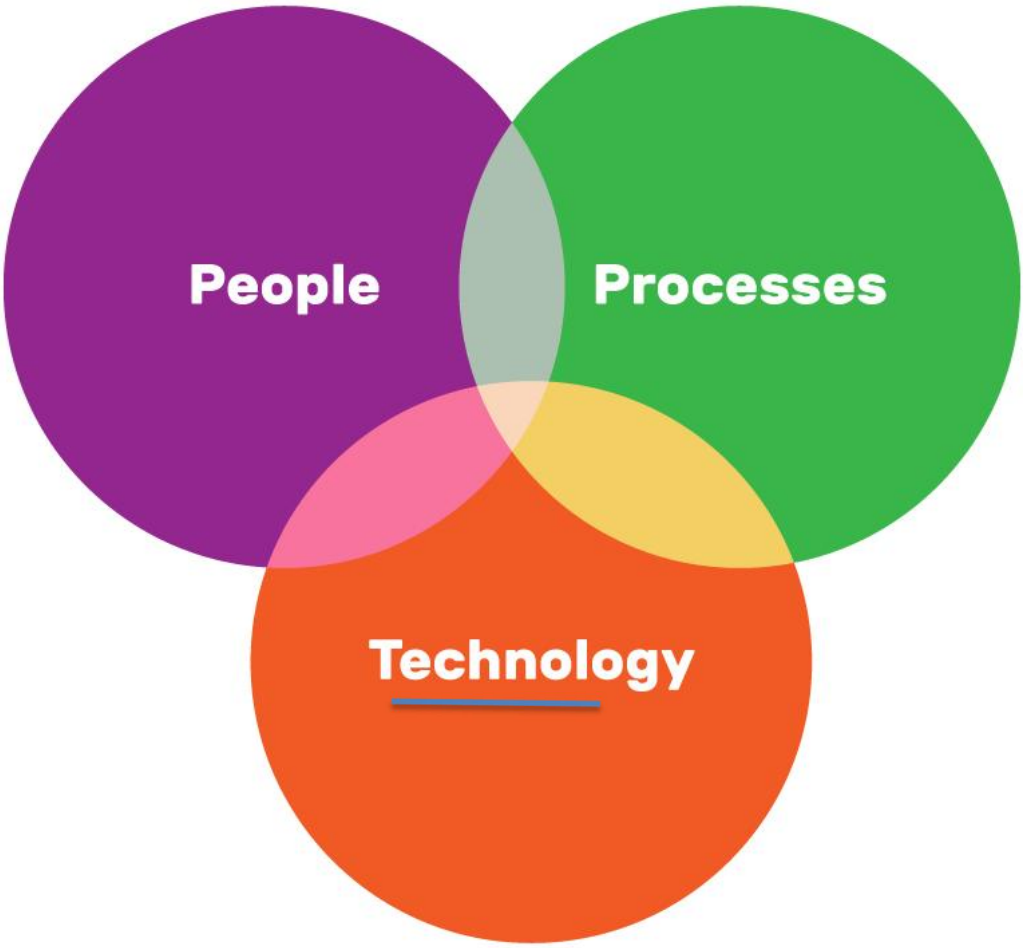

These pillars are all interlinked, and every cybersecurity approach requires all three. In the table below, a description is provided of the role of each pillar in protecting data on your phone.

Pillar	**Role**
People	A careful owner who keeps possession of their phone and turns on password protection.

Pillar	Role
Process	An effective password protocol with a strong password.
Technology	A phone that is password-protected, and free from other vulnerabilities.

Questions

- What role do you think people play in an organisation's password policy?
- What about the role of process or technology?

Share your answers in the comments

Next steps

Without good cybersecurity, we could become the victim of social engineering, malware, and ransomware attacks, among others. In the next steps, you will learn about some of these attacks.

Social engineering

Automated social engineering

In the previous step, you saw the value of your data. Now, you will learn about social engineering attacks, in which attackers try to steal your data. In this step, you will be introduced to phishing, pharming, and name generator attacks.

What is social engineering?

Social engineering is the name given to the type of attack that deceives victims into sharing valuable personal data.

There are many different types of social engineering attack. In this step, you will learn about three kinds:

- Phishing attacks
- Pharming attacks
- Name generator attacks

Phishing attacks

A **phishing attack** is an attack in which the victim receives an email disguised to look like it has come from a reputable source, in order to trick them into giving up valuable data.

The email will either ask for the information directly, or provide a link to another website where the information can be inputted. This attack may also come via phone call or text message.

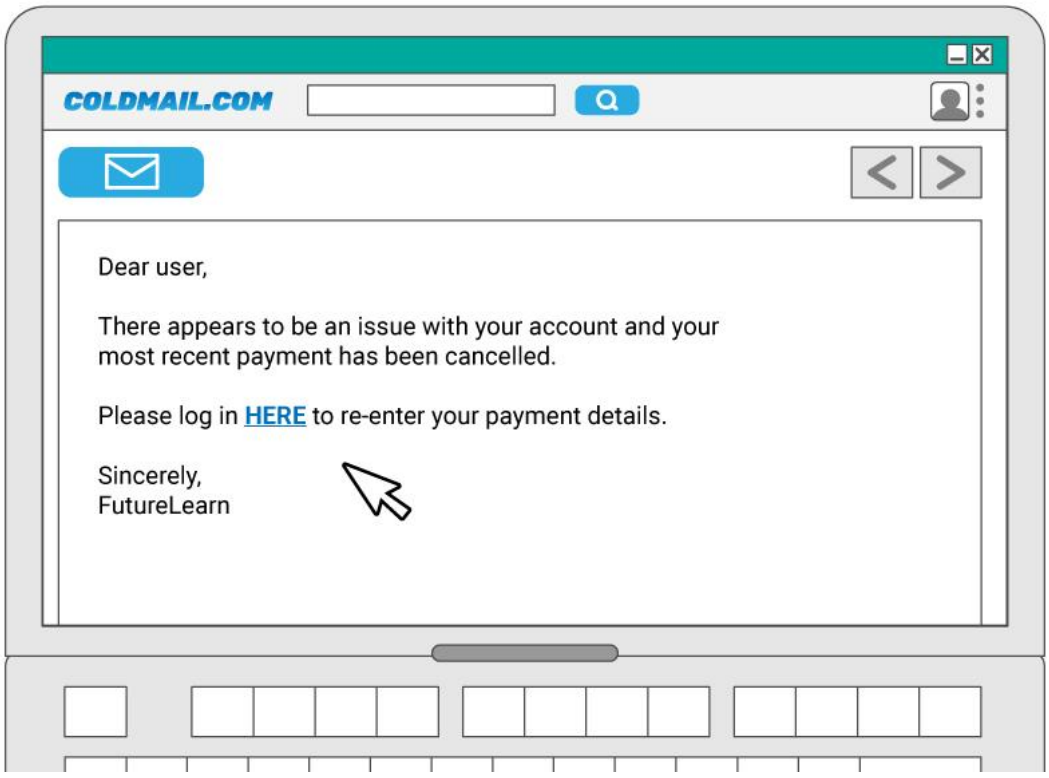

Phishing emails can be recognised in a number of ways. Key indicators to look out for include:

- Any unexpected email with a request for information

- Sender email addresses that contain spelling errors, lots of random numbers and letters, and/or domain names that you don't recognise
- Suspicious hyperlinks:
    - Text that appears to be hyperlinked but does not contain a link
    - Text that is hyperlinked to a web address that contains spelling errors and/or lots of random numbers and letters
    - Text that is hyperlinked to a domain name that you don't recognise and/or isn't connected to the email sender
- Generic emails that don't address you by name or contain any personal information that you would expect the sender to know

Some phishing attacks are more sophisticated and target specific individuals or groups of people, for example, by pretending to be from a company that the person has an account with. This is called **spear phishing**.

To avoid phishing attacks, you should not fill out forms or click on links in emails that you are not expecting.

Pharming attacks

A **pharming attack** is an attack in which malware redirects the victim to a malicious version of a website. The malware could infect the victim's computer or the DNS server (the database that allows your browser to find the website you're visiting. Then, when the victim enters a web address into their browser, they visit a website controlled by the attacker, rather than the legitimate website. The attacker can then collect any data that the victim inputs into the website. Links in phishing emails may also redirect victims to pharming websites.

As with phishing attacks, pharming attacks can be identified from aspects of the website that seem out of place or incorrect. For example, any of the following could indicate a pharming attack:

- Spelling errors or incorrect logos
- Broken or missing links
- A notification from your browser warning you that the webpage is insecure
- The lock symbol that your browser uses to confirm that a webpage is secure is missing

http://www.socialhealthgrowth.org

https://www.socialhealthgrowth.org

If you suspect that a website is malicious, you should close your browser and run up-to-date antivirus software on your computer, then reload the page to see if it has changed.

Name generator attacks

A **name generator attack** is an attack in which the victim is asked in an app or social media post to combine a few pieces of information or complete a short quiz to produce a name.

Attackers do this to find out key pieces of information that can help them to answer the security questions that protect people's accounts.

To protect yourself from name generator attacks, you should avoid providing apps with the following pieces of information or posting this information publicly on social media sites:

- Your mother's maiden name
- Names of current or previous pets
- Previous or current addresses
- Your age or birthdate
- Your lucky number
- Any of your favourite things (such as your favourite place or author)
- Any information that you know you have used to create a password or set up a security question

Next step

In the next step, you will learn about two types of social engineering attack that require the attacker to interact with the victim more personally.

Questions

- What are social engineering attacks used for?
- Why do you think social engineering attacks are effective?
- Of the three types of social engineering attack discussed, which do you think is the most likely to be successful?

Interactive social engineering

In the last step, you learned about social engineering attacks that can be conducted by attackers who create one program and send it to lots of victims. In this step, you will find out about two types of attack that require the attacker to interact

with the victim more personally. These attacks are blagging and shouldering.

Blagging

Blagging (also known as **pretexting**) is an attack in which the perpetrator invents a scenario in order to convince the victim to give them data or money. This attack often requires the attacker to maintain a conversation with the victim until they are persuaded to give up whatever the attacker has asked for.

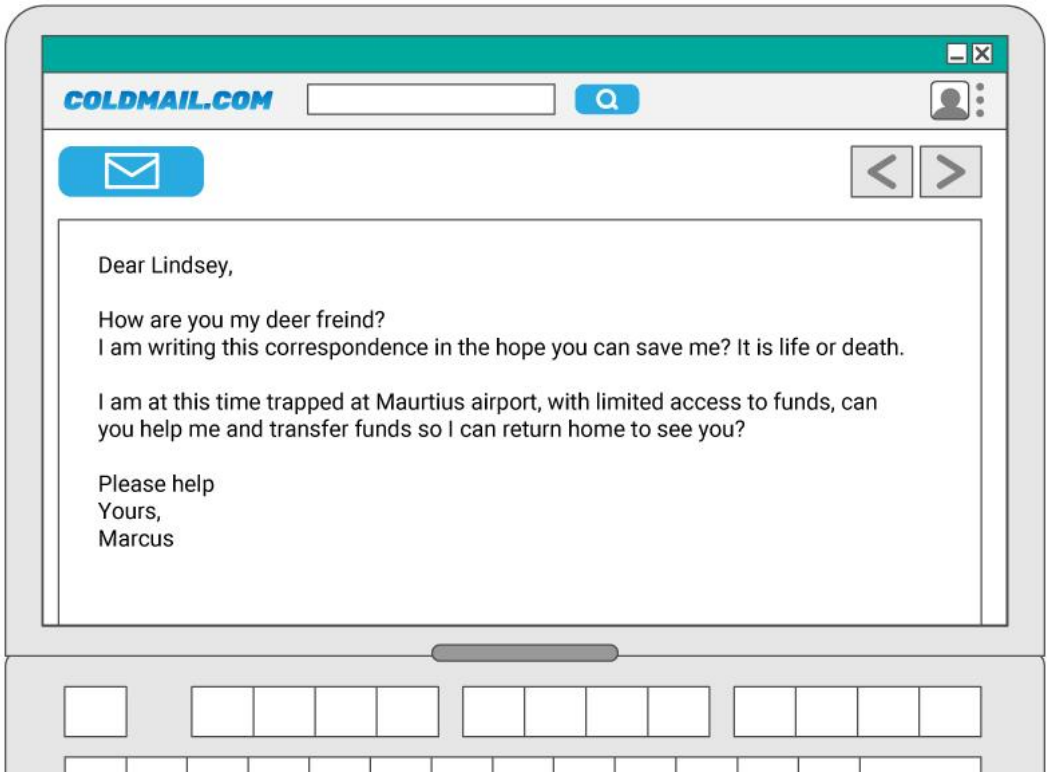

For example, the victim might receive an email from an attacker pretending to be a friend trapped in a foreign country in desperate need of a money transfer that will be repaid with interest as soon as they are safe.

A particularly common type of blagging attack is an **online dating scam**. In an online dating scam, the attacker might meet the victim through a dating app or chat room and begin an online relationship with them. Then, they might ask the victim for money or gifts, or if they find out that the victim is married, threaten blackmail. Action

Fraud estimates that around $27 million was lost to this kind of scam in 2014/15.

Shouldering

Shouldering (also known as **shoulder surfing**) is an attack designed to steal a victim's password, or other sensitive data. It involves the attacker watching the victim provide sensitive information, for example, over their shoulder. This type of attack might be familiar as it is often used to find out someone's PIN at a cash machine.

Questions

- How would you avoid becoming the victim of a blagging attack?
- How would you protect your password or PIN from a shouldering attack?

Share your answers in the comments

Suggested classroom exercises

In order to design defences against different types of attack, security professionals often pretend to be attackers. This allows them to identify weaknesses in cybersecurity systems, and better understand the mechanics of the attacks.

You may have heard of white hat hackers — people who hack into websites and security systems, not to steal data, but to expose security flaws so that they can be fixed. This approach is taken with other types of attack as well, and can be used to help your students avoid blagging and shouldering attacks.

Students could do role-play activities in order to better understand how these attacks work:

- When teaching about blagging, you could ask your students to write a script from the perspective of the attacker. To do this, they will have to contemplate what language an attacker might use and why.
- When teaching about shouldering, you could ask your students to get into pairs of one attacker and one victim, and see if the victim can write a secret word or phrase on a piece of paper without the attacker reading it.

Passwords

What do we protect with our passwords?

Most of our online accounts are protected by a username and password combination. These passwords protect the data that we store in our accounts, whether that is our bank details, our purchase history, or our home address.

Exercise

How many accounts do you have that are password-protected? How many different passwords do you have?

Read the rest of this step, and comment on how secure you think your passwords are

Passwords do not form a perfect defence. There are many ways in which attackers can find out your passwords and then use them to enter your accounts.

How are passwords hacked?

You will learn about three different types of password attack: brute force attacks, theft of individual passwords, and theft of batches of passwords.

Brute force attacks

In a **brute force attack**, an attacker guesses passwords until they find the correct one. This might involve guessing a combination of characters, or creating a list of passwords beginning with the most common, as in the more specialised **dictionary attack**. The dictionary that attackers use contains passwords centred around real words and combinations of real words.

Theft of individual passwords

An attacker could **steal a victim's password**, for example, by using the **social engineering** techniques discussed previously, or by infecting the victim's device with a form of **malware** that records their activity, including the letters that they type. You will learn more about malware next week.

Some websites take more precautions to protect your accounts than others. You might have multiple strong passwords for an online banking account, but you might not take the same precautions when setting up a social media account. Attackers know this, and will target weaker accounts to help them to guess the passwords for more secure accounts.

Theft of batches of passwords

An attacker could hack a website and **steal batches of passwords**. This can give them access to lots of accounts at the same time. To learn more about how websites store passwords securely, refer to the attached PDF.

How do you make a strong password?

Passwords should be memorable for the individual, but difficult for an attacker to guess. As you have seen, password attacks often rely on victims using common combinations of characters and similar passwords across multiple accounts. Therefore, all of your passwords need to be different and unpredictable.

Avoid using personal details and dictionary words

You should avoid using any personal details, like your pet's name or your favourite sports team, as a basis for your password. To protect yourself from a brute force attack, you should avoid dictionary words altogether, even if you're substituting some letters for numbers or symbols — if "password" is in the attacker's dictionary, so is "p@ssw0rd".

Increase the length and complexity of your password

You should also increase the length of your password and add in more types of character. The more types of character you include and the longer your password is, the more guesses the attacker has to make.

Use a strong password generator

Rather than finding a strong password, it is better to design a strong password generator that you can use to easily create lots of memorable passwords that appear random. Here are three methods of generating passwords:

- Create a phrase from random words — you can still defend against a dictionary attack if you combine words in an unpredictable way. Choosing words at random is the easiest way to do this. For example, this website helps you to choose words at random with dice. Once you have chosen the words, you should add numbers and symbols into the password.

- Use a memorable phrase as the basis of your password, instead of using words. For example, you could turn the phrase ‘Socialhealthgrowth’ is the number one online learning platform’ into the password ‘S0cialHeaLTHgr0wth’. You can tailor this phrase to the purpose of your account to make it more memorable. For example, you could use a phrase about shopping to make a password for an eBay account.
- If you have a visual memory, create a grid of characters (arranged randomly) and choose your password by drawing a pattern. Then, you would just need to learn the pattern, not the actual password.

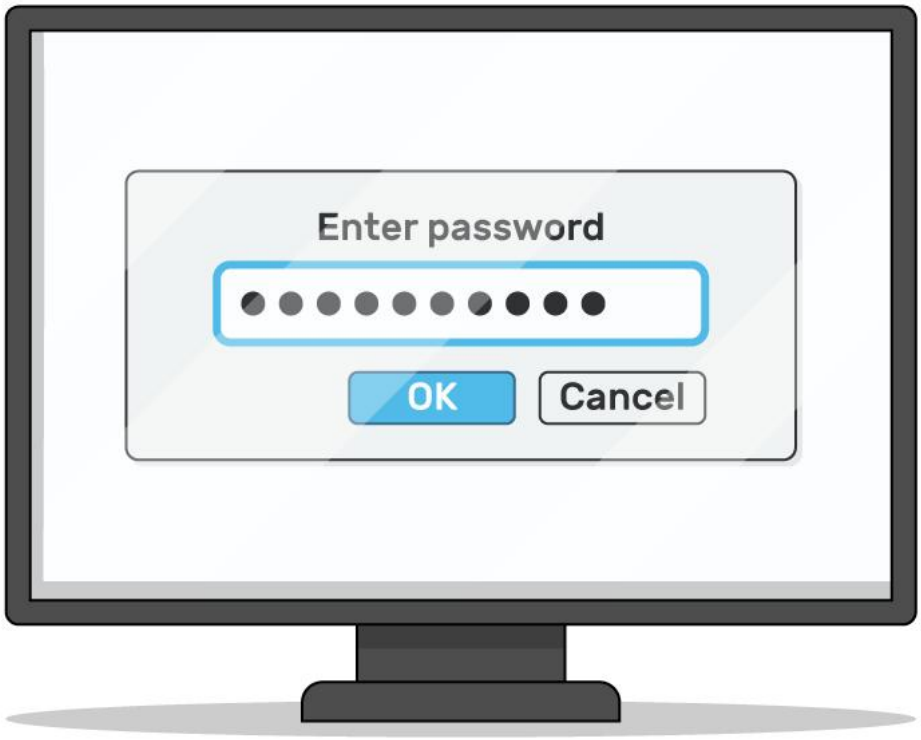

Next step

In the next step, you will learn how to keep your phone secure.

Questions

- What kind of information might an attacker use to guess your password?
- Why is a longer password more secure?

- How might you avoid an attacker stealing your password through a phishing attack?

Phone security

Many of us carry around extremely sensitive data in our pockets: our phones give us access to our bank accounts, to our work files, and to our apps, among other things. All of this data is only as secure as our phones and so, in this step, you will learn about phone security, and how to ensure that your phone is secure.

Password protection

Most phones are protected by a passcode, which is usually a series of digits or a pattern. As with passwords, the longer and more random the passcode is, the harder it is to guess. As passcodes are typically four to six digits long, they are easier to guess than a password. However, many phones are programmed to lock down or wipe their contents if the wrong passcode is inputted too many times. This fail-safe should discourage an attacker from trying to guess the combination.

However, phones can give physical clues to hackers which can give away the passcode. This is because when the user touches their phone screen, they leave grease from their fingers behind. This can point out the places on the screen that they touch often, and even hint at the order in which the keys have been touched. Therefore, it is important to wipe your phone screen after use.

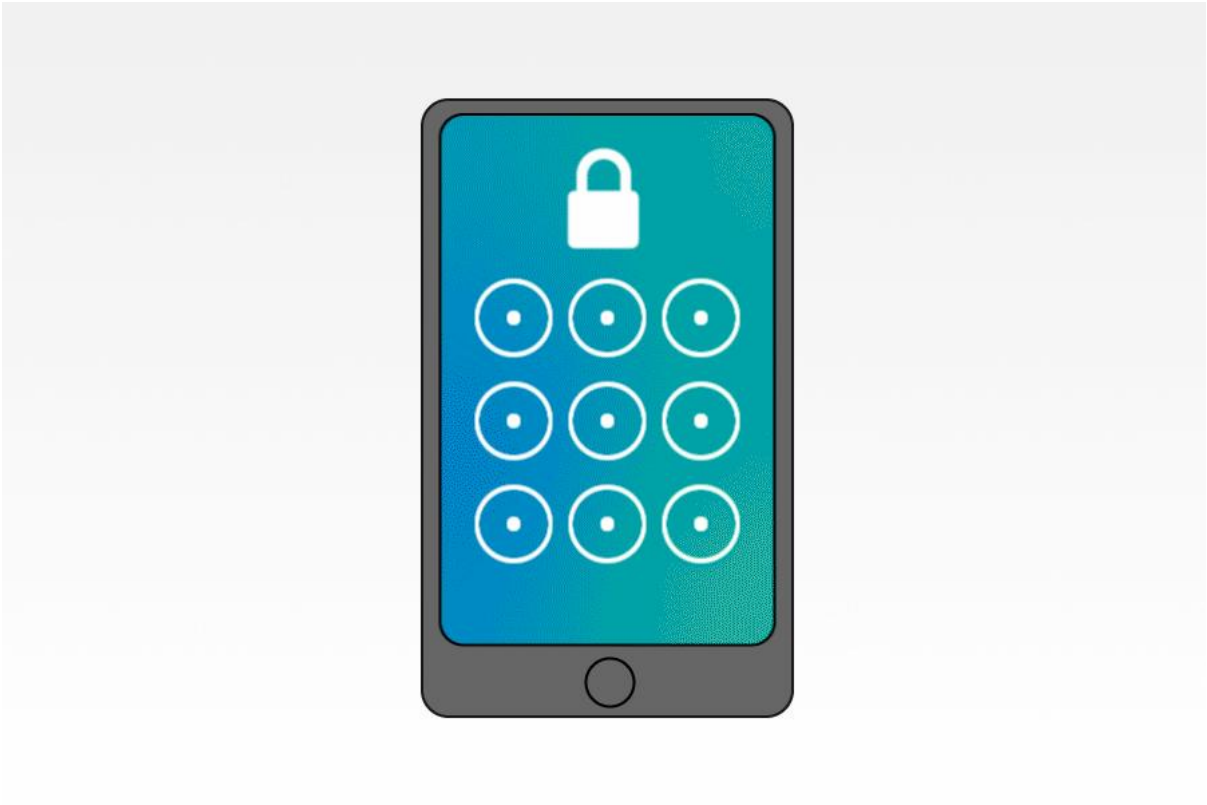

Biometrics systems

A new development in phone security is the use of **biometrics systems**. Biometrics systems use unique biological characteristics as authenticators. Some examples of biological characteristics used as biometric data include:

- Fingerprints
- Iris/retina patterns
- Facial features
- Voice patterns

Biometric data has several advantages over passwords. For example, users can’t forget it, and they take it everywhere with them. Because biometric data is unique, users don’t need to be able to create something secure.

However, biometrics systems are more expensive to build than password-based systems, and it is possible to forge biometric data. For example, fingerprints for mobile phone biometrics systems can be recreated using everyday materials such as sticky tape and lipstick, or a glue gun and tinfoil.

This is an additional video, hosted on YouTube.

Remote access

Attackers may also try to break into your phone remotely. Remote access attacks work without the user needing to connect with a hacker's phone or give them access permissions. Phones are vulnerable when they communicate with other networks, and devices broadcast information when Bluetooth is turned on or the device is allowed to search for wireless networks.

The information that a device broadcasts might help an attacker to tailor hacking software to the device, which would improve their chances of successfully carrying out an attack. Hackers can also use a device's connection to gain access to the data stored on the

device, or even gain control of the device and use it to make payments or phone calls.

To protect your phone from remote access attacks, you should keep Bluetooth and wireless network connectivity switched off until you are ready to connect to a trusted and secure network. It is also important to keep your operating system (OS) security updated. In addition, apps that require two-factor authentication (2FA) are better protected from remote attacks, because if an attacker gains access to a device, they will still need an additional password or other form of authentication to use the apps. You will find out more about OS security updates and 2FA later.

Permissions

So far, you have learned about how to protect your phone from external attackers, but the data is also threatened from within. When a user downloads an app, it asks for permission to access the data stored on their phone, including their contacts, camera, photos, location data, and more. Apps can even be given permission to turn on the user's microphone and record their conversations.

App providers might use this data to optimise the app's performance, or they may be trying to harvest and sell the data as part of the data economy. Malicious apps could also be used to collect data to steal the user's identity and/or hack their accounts.

App stores like Apple and Google take steps to check apps for malware, but they cannot protect phones from every threat. Therefore, when an app requests access to your data, it is important to double-check what data it wants to access, and think critically about whether you need to share that data with them.

Exercise

Which of the following do you think the game Candy Crush requests access to and why?

- Name
- Contact details

- Contact list
- Social media profiles
- GPS data
- Camera and microphone

You might find the website Terms of Service; Didn't Read helpful to complete this exercise, or to check the permissions that you have granted to other apps.

Share your thinking in the comments section.

Next step

We rely on an up-to-date OS to protect our phones, and you will learn more about this in Week 2. In the next step, you'll learn more about a form of data protection called two-factor authentication.

Questions

- Why might an attacker choose to carry out a remote access attack instead of one of the other types of attack discussed in this step?
- How might you best protect yourself from an attack conducted using a malicious app?
- Is your phone more secure when protected by a passcode or by a form of biometric data?

Two-factor authentication (2FA)

In the previous steps, you learned about flaws in the protocols designed to protect data, accounts, and systems. Two-factor authentication has often been proposed as the solution to many of these flaws. You will find out more about two-factor authentication in this step.

What is two-factor authentication?

Two-factor authentication (or **2FA**) is a method of verifying a user, in which the user has to present two forms of proof of their identity. Usually, this involves asking for any two of the following:

- Something the user has
- Something the user knows
- Something the user is

For example, the user might have to submit a password (something the user knows) and a code sent via text message or generated on a token (something the user has). Alternatively, the user might have to scan their fingerprint (something the user is) and answer a security question (something the user knows). An authentication method is only considered 2FA if the user needs to provide information from two categories (for example, having to provide two passwords would not count as 2FA).

How secure is 2FA?

2FA systems are generally more secure than 1FA (one-factor authentication) systems because hacking them requires more effort — instead of stealing one type of credential, an attacker would need to steal two. In addition, the weaknesses of one form of authentication can be mitigated by another form.

- Credentials that the user “knows” could be guessed, or stolen if they are stored somewhere like a website’s database, so they are susceptible to remote attacks.
- Credentials that the user “has” would need to be physically stolen by an attacker, which prevents an attacker from acting remotely.
- Credentials that the user “is” would need to be forged, which can be more difficult. However, if an attacker manages to forge a biometric key, that key is compromised forever, whereas passwords and physical keys can be replaced.

While a 2FA system is stronger than a 1FA system, it is not perfectly secure, as 2FA can be implemented poorly. Here are two examples of poor implementation:

- If both forms of authentication are on the same device, then an attacker could hack that device and therefore both credentials. For example, if a user needs to enter a combination of a password and voice pattern to access their phone and an attacker was able to hack the user's phone to record their keystrokes and any recordings made by their phone's microphone, then the attacker could steal both the user's password and voice pattern.
- If a service provider accepts 1FA as well as 2FA, then attackers could gain access once they have obtained the first credential. For example, some email providers accept either 2FA or 1FA, as not all customers are prepared to use 2FA. Attackers could therefore access the user's account just by guessing their password.

It might be harder for an attacker to access an account protected by 2FA, but it is not impossible, as long as each form of authentication has vulnerabilities. For example, it might seem secure to have a code sent to your mobile phone, but it is possible for attackers to intercept the message sent to you by hacking the telecommunications network itself, and thereby gain access to your account. However, because of the added layer of complexity, attacks on 2FA systems are often more specialised, so it is harder for attackers to target lots of victims at once.

Next step

In the next step, you will find out about USB flash drives as a tool for storage, security, and sabotage.

Questions

- Which combination of the three forms of authentication used in 2FA do you think is the most secure and why?

- Even when 2FA is given as an option, users do not always implement it. Why do you think that is?
- Which of the 3 pillars of security, that we have discussed previously, apply to 2FA?

USB flash drives

USB flash drives, also known as USB sticks, are devices that are often used to back up files or transport them from one computer to another. But they can also be used to protect or attack computers.

Storing passwords on USB flash drives

As you learned in a previous step, the longer and more random a password is, the stronger it is. However, we don't all use very long, randomly generated passwords, because they are too difficult to remember. This kind of password would be easy for a computer to remember, and could be kept on a USB drive to make it portable. However, if the USB drive gets lost or stolen, the user would no longer be able to access their device, and someone else could use the USB drive to access the user's data.

Malicious USB flash drives

USB flash drives are often used to transport files, so they connect to lots of different devices. If the user plugs their USB drive into an infected computer, the malware could spread to their USB drive, and their USB drive would then be a host to the malware. Their USB drive would then infect every computer it connects with subsequently. The problem is made worse in that we often instruct computers to allow USB drives to run automatically, which enables the malware to spread.

Attackers may deliberately target victims with infected USB flash drives. For example, they might infect a USB drive and send it to their victim under the guise of a freebie or promotion. Alternatively, they might leave it lying around in a public space in the hope that someone will find it and plug it into their device (either to use it or to find its owner).

Next step

In the next step, you will learn about the legislation surrounding stealing data or hacking into a device.

Questions

- How could you minimise the chances of an infected USB flash drive installing malware on your computer? What ideas do you have for teaching this to your students?
- In the previous steps, you learned about several ways of restricting access to your data (passwords, biometrics, and now, USB flash drives). Which do you think is the most secure

and why? Is every form of authentication appropriate for every type of device?

- In the previous steps, you learned about how to keep data secure. Can you think of any problems that data security might cause for the police investigating criminal activity?

The legal perspective

In this step, you will learn about the legislation designed to protect our devices and our data, and discuss its efficacy.

The Computer Misuse Act and the Data Protection Act are two pieces of legislation which define criminal offences involving computers, and the rights we have to protect our data. You will find out more about both acts, but first, here is an exercise to see how much you already know.

Exercise

Which of the following is illegal under the Computer Misuse Act?

- You access your friend's phone to get their contacts for a surprise birthday party
- Your friend gives you their Netflix password so that you can watch a film
- For a school project, you write a script that can hack someone's laptop through their wireless network

Once you've read through the rest of this step, reflect on your answers. Have you changed your mind?

The Computer Misuse Act

The Computer Misuse Act (1990) and its amendments were created so that unauthorised access to computers and crimes committed using a computer could be prosecuted. The act is based on three principles and makes the following actions illegal:

- Unauthorised access to digital/computer material. This means a person asking a computer to perform any function with the intent to access anything on the computer when they do not have permission, and know that they do not have permission.
- Unauthorised access to digital/computer material with intent to commit or facilitate the commission of further offences. This means a person gaining access to a computer without permission in order to commit another crime or enable someone else to commit a crime.
- Unauthorised acts with intent to impair, or with recklessness as to impairing, the operation of a computer. This means a person intentionally impairing the operation of any computer or program, or intentionally preventing access to any data or program on any computer. This includes creating or supplying materials that could be used to carry out this offence.

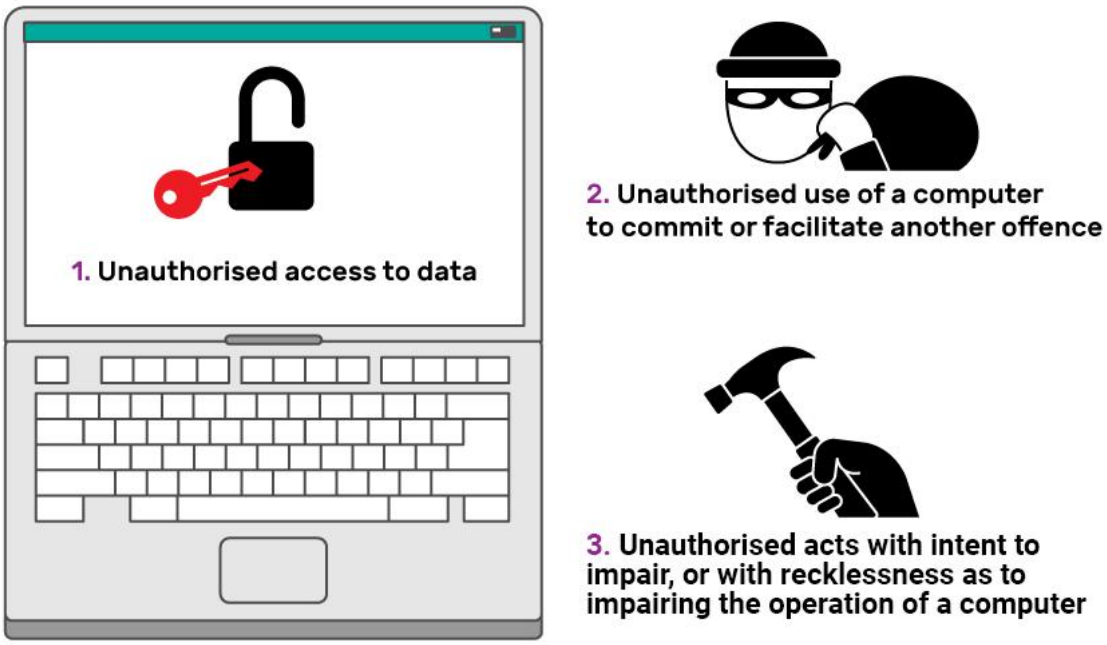

This legislation is commonly used to prosecute crimes such as hacking and ransomware attacks.

The Computer Misuse Act also has three levels of punishment:

- Unauthorised access to digital/computer material is punishable by up to 2 years in prison and a $5,000 fine
- Unauthorised access to digital/computer material with intent to commit or facilitate the commission of further offences is punishable by up to 10 years in prison and an unlimited fine determined by the damage caused and the severity of the crime
- Unauthorised acts with intent to impair the operation of a computer is also punishable by a prison sentence of up to 10 years and an unlimited fine, but if the act puts life at risk or endangers national security, then the sentence may be extended to life imprisonment

Questions

- Do you think the police should have to prove intent in order to enforce these laws?
- Do you think the punishments are justified?
- In 2015, law enforcement authorities were granted immunity from the legislation concerning cybercrimes. Do you think this is legitimate, or do you think it has the potential for abuse?

Share your answers in the comments

The Data Protection Act (Singapore PDPA) Do follow the attachment

The Data Protection Act (2018) is the legislation that dictates how data can and should be protected. It describes the rights of the individuals whose data is being held, and the responsibilities of those holding the data.

This legislation is an updated version of the Data Protection Act (1998). It is similar to the EU General Data Protection Regulation (GDPR) which currently governs UK data protection policy, even though the UK has left the European Union so far this law has remained in place.

The Data Protection Act has six principles that dictate how anyone who uses data must protect it. They must ensure that the data is:

- Used fairly, openly, and in accordance with the law
- Used for a specific and stated reason
- Used only in a way that is necessary and sufficient for the purpose for which it was collected
- Accurate and up-to-date
- Only kept for as long as it is needed
- Protected against loss, damage, and unauthorised access

The Data Protection Act gives individuals rights to access and manage their data. Individuals have the right to:

- Find out how their data is being used (by an organisation)
- Access the data that an organisation has about them
- Update their data
- Have their data deleted
- Stop an organisation from processing their data
- Transfer their data to a different organisation

In some circumstances, individuals also have the right to object to their data being

processed.

Questions

- Do you know how to access the data held by an organisation about you?
- Do you feel that you have the rights laid out in the Data Protection Act?
- Do you believe that companies in the UK abide by the six principles of data protection?

Congratulations on completing 3rd and 4th month of the course. This week, you've learned about the main idea of cybersecurity, and looked at social engineering attacks that attempt to exploit technology users. You've also considered ways of securing a device, including two-factor authentication, and thought about what situations are covered by the Computer Misuse Act and the Data Protection Act.

Next week, you'll look at malicious software, bots, and denial-of-service attacks.

What did you find particularly interesting or surprising this week? Leave a comment and let us know!

Welcome to 5th Month

Welcome to 5th month of *Introduction to Cybersecurity for Mother and Father*. Last week, you looked at ways in which hackers could gain access to your device through social engineering, as well as some of the ways of protecting your device. Now, you'll find out more about malicious code and how attackers can use it.

This week, you will learn about:

- The different types of malware
- Ways to prevent attacks using antivirus software
- How to try to avoid leaving vulnerabilities in software
- How attackers can use compromised machines in other attacks

Malware

In this step, you will learn about some different examples of malware (malicious software). More specifically, you will learn about worms, viruses, Trojans, ransomware, and spyware.

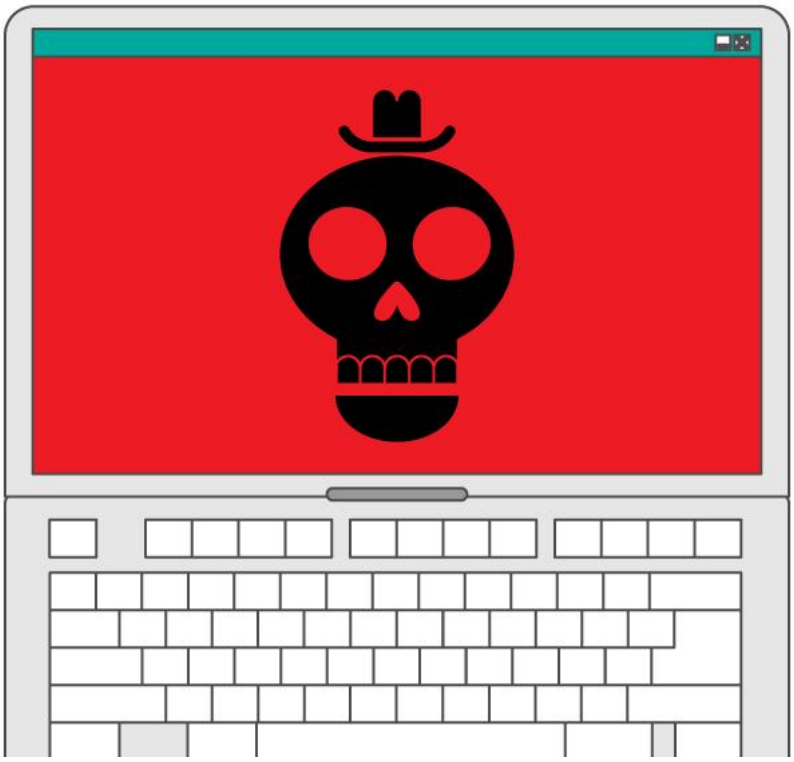

Worms

Worms are a form of malware that replicate themselves in order to spread to other computers. They can be programmed to damage software, but are often designed only to spread and infect as many devices as possible.

A famous example is the **Mydoom** worm, which was first launched in 2004 and became the fastest-spreading computer worm of its time. It was transmitted via email and infected hosts, causing them to carry out a distributed denial-of-service (DDoS) attack against several tech companies (you will find out more about DDoS attacks later). It is estimated that Mydoom infected between 16 and 25% of all emails in 2004 and caused approximately $38 billion worth of damage.

Worms can be detected and removed by antivirus software. You will learn how later.

Viruses

Viruses are another form of malware that operate differently to worms. Worms come in the form of their own file or program, whereas viruses infect a file or program that already exists. In order for a virus to operate, the device user has to open the file or run the program. They can be harder to detect because they hide themselves in their host file. Viruses are less common than worms.

The **Anna Kournikova virus** was a famous example of a virus. The virus infected victims' Outlook so that it sent an email to all of the contacts in their address book. The virus was disguised as a photograph of Anna Kournikova to trick victims into opening it.

Trojans

Trojans are pieces of malware disguised as legitimate software. They are named after the famous wooden horse in Greek mythology that snuck Greek soldiers into Troy. Unlike worms, Trojans aren't self-replicating, and unlike viruses, Trojans don't infect other files. However, like viruses, they can only operate if the victim runs the program.

The **Zeus Trojan** often infects computers through phishing and pharming scams (as you learned about in Week 1). It takes control of the infected computers and turns them into **bots**. You will learn more about bots later. It can use these bots to conduct more criminal activity. For example, it can obtain the victim's banking passwords by recording their key strokes. Because the Zeus Trojan uses malware to obtain sensitive information from its victim's devices, it is also known as **spyware**.

Ransomware

Ransomware is a form of worm called a **cryptoworm**. Cryptoworms are worms that use cryptography to cause damage to the devices they infect. In particular, ransomware encrypts files on a victim's device and then demands payment from the victim in exchange for the decryption key.

A recent example of ransomware is **WannaCry**. It is thought to have had hundreds of thousands of victims. In 2017, it infected computers within the NHS infrastructure. Over just one week, it caused an estimated 19,000 appointments to be cancelled because staff were unable to access patient records or equipment. The attack is estimated to have cost the NHS tens of millions of pounds, even though the ransom was not paid.

Many believe that the attack was enabled by out-of-date software. You will learn more about how to protect devices from malware in later steps. With ransomware, if your files are backed up on another device, then you are in a better position to refuse to pay a ransom, because it doesn't matter if you don't ever get access to your encrypted files.

Adware

Adware can be a worm, virus, or Trojan. It infects a computer and causes it to download or display malicious adverts or pop-ups when the victim is online. If you are online and see a large number of adverts, or you are receiving pop-ups when you're offline, then your computer might have been infected by malware.

Next step

In the next step, you will learn about how malware accesses our computers.

Questions

- Which of the attacks, worms, viruses, Trojans, and ransomware do you think would be the easiest to detect?
- How would you protect your device from a Trojan?
- Why do you think viruses are less common than worms?

Vectors of attack

In the last step, you learned about different kinds of malware. In this step, you will learn about how malware infects our computers. This is known as a malware's *attack vector*.

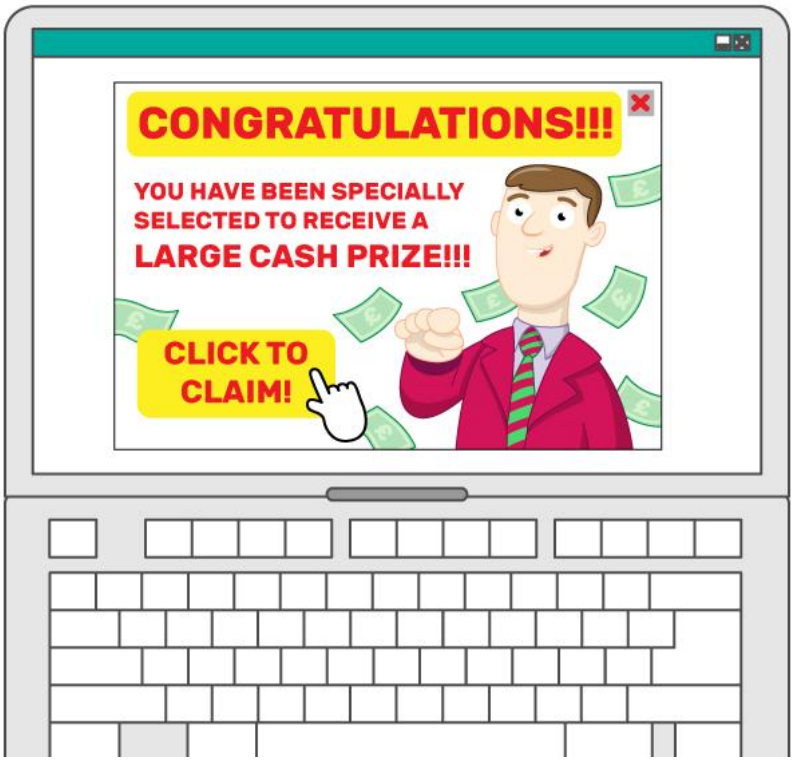

In this step, you will learn about two types of vector — **passive** and **active**.

- Passive vectors do not require the victim to engage with the malware for it to run
- Active vectors require the victim to engage with the malware for it to run

Passive vectors can be difficult to detect because the victim might not realise that they have downloaded any malware until they receive an alert from their antivirus software or, even worse, suffer the consequences of the malware.

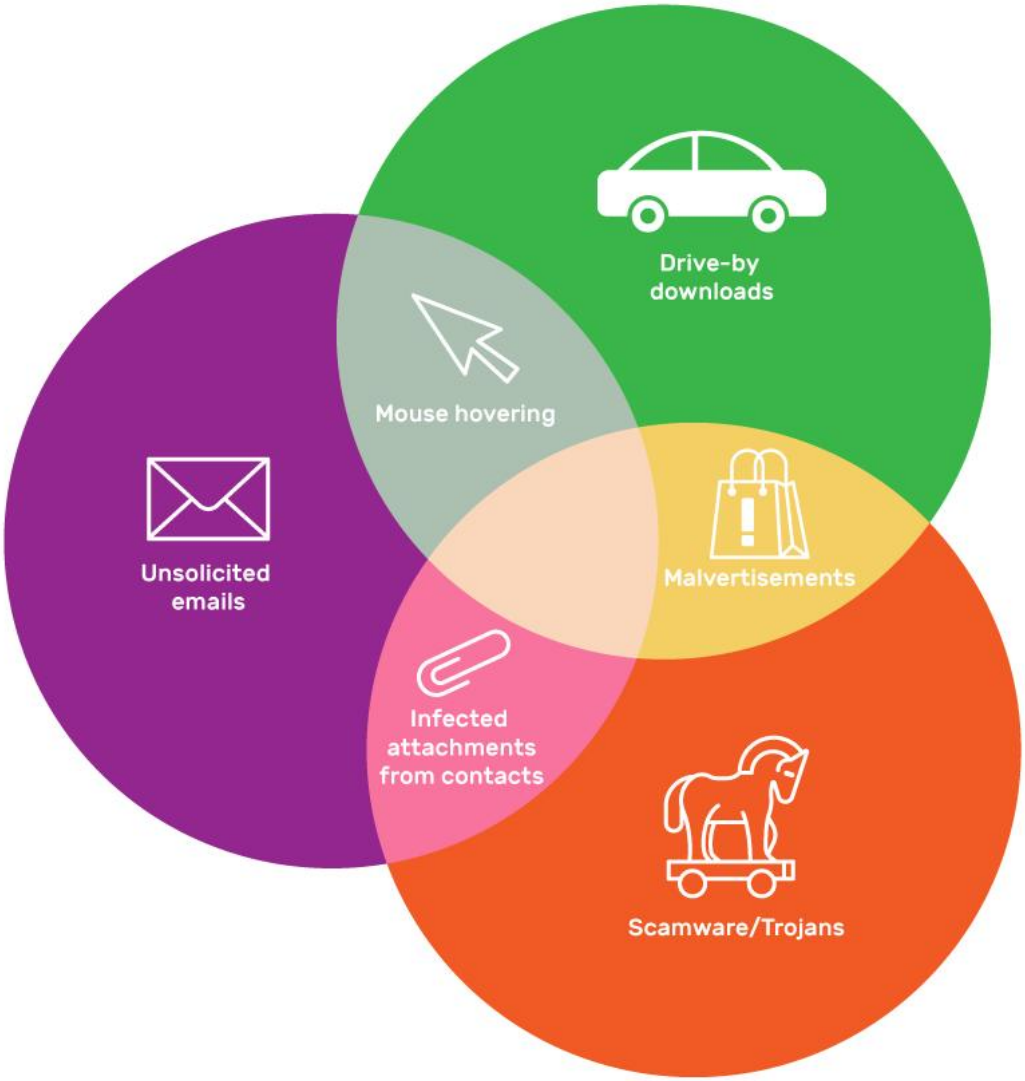

Unsolicited emails

A common method of infecting computers is attaching malware to emails (as in the Anna Kournikova attack). Users can prevent this by avoiding **unsolicited emails** and not opening files or clicking on links contained in them.

Self-propagation

Worms (and some viruses) can **self-propagate**:

- A worm/virus may make the infected device send copies of the worm/virus to the victim's contacts via email.

- A worm/virus may be hidden on a device or in a file. If an infected device is connected to another device, or an infected file is downloaded onto another device, then that second device will also be infected. Therefore, victims of worms and viruses may unknowingly play a part in the infection of their devices and the spreading of the malware.

Trojans/scamware

The victim might run malware (Trojans) because they think that it is legitimate software. A common form of Trojan is called **scamware**, which is malware designed to look like antivirus software. This malware tells the user that it has detected threats and offers to clean their device. If the user accepts this offer, then the malware can infect their computer.

To defend against Trojans:

- Scan files or programs before you run them (even if they are from a friend)
- Only download and run software from trusted sites
- Do not open or run a file or software package if you don't remember downloading it

Drive-by downloads

Some malware can be downloaded without the user clicking on a download link. In a **drive-by download**, malware is downloaded onto the victim's computer when they visit a malicious website, even if they don't click on any of the links or download any software. This attack vector exploits vulnerabilities in the victim's browser to gain access to their computer.

Users can try to protect themselves from drive-by downloads by avoiding websites that might be hosts to malware, and keeping their browsers up-to-date.

Malvertisements

Malvertisements are advertisements that contain malware:

- Some malvertisements install malware onto the user's computer if the user clicks on them
- Some more sophisticated versions of malvertisements begin the attack if the user just visits the website hosting them

Malvertisements can be uploaded to legitimate websites if the attacker has attacked the website or the company that supplies their adverts. It is therefore more difficult to avoid malvertisements than to avoid drive-by downloads.

Mouse hovering

In 1st month, you learned about phishing emails and how they can be identified by hovering your mouse over a hyperlink to see if the redirect address is suspicious. **Mouse hovering** exploits that behaviour and uses it to enable the download of malware onto the victim's computer. This attack has been seen in infected PowerPoint files, which are sent by email and contain hyperlinks that release malware when the victim hovers their mouse over them.

The attack can only take place if the victim has downloaded a suspicious file onto their computer. The attack is also enabled by security flaws in the software used to run the file, so it is important to keep the software on your computer (such as Microsoft Office) up-to-date.

Next step

To avoid becoming a victim of malware through many of these attack vectors, you need to keep your antivirus software up-to-date. You will find out about antivirus software and how it works in the next step.

Questions

- Which of the attack vectors described do you think are passive, and which are active?
- Because of the way that some malware spreads, defending yourself relies on both you and your contacts having good

computer security. How can you help your friends and colleagues avoid infecting their devices as well as yours?

- Which of the attacks mentioned so far do you think is the hardest to avoid?

2.4

You've completed 3rd month

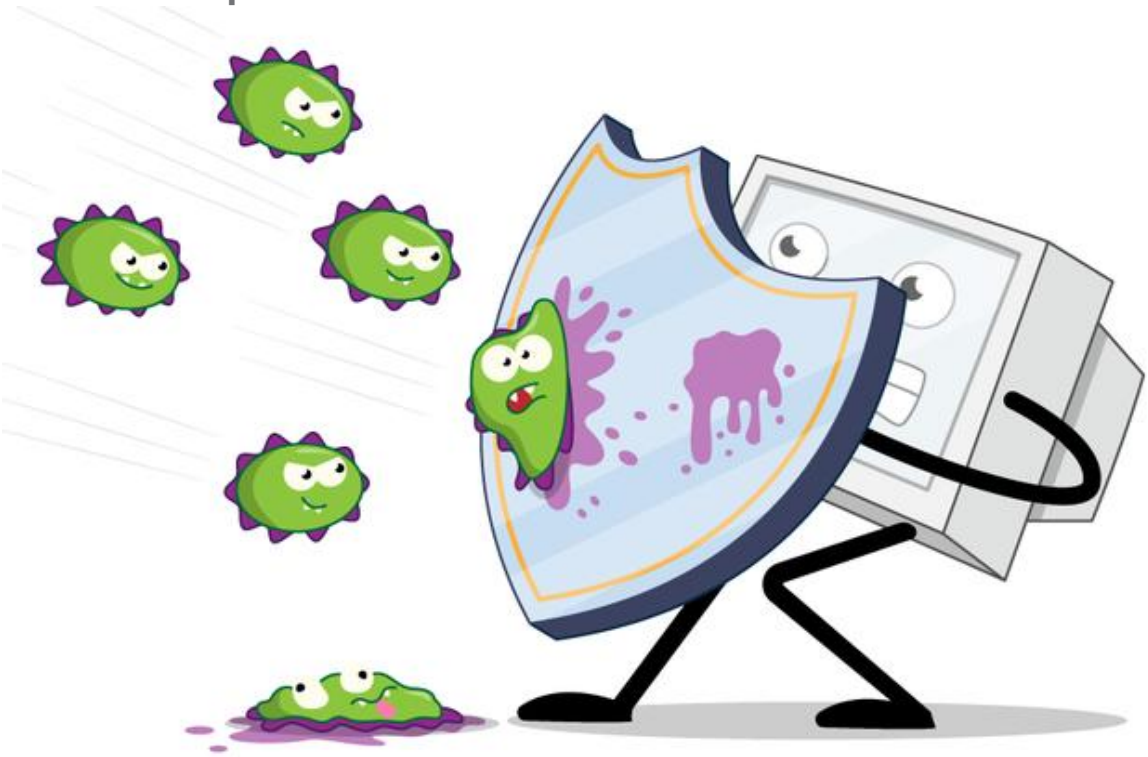

Antivirus software

In the last step, you learned about how malware can access computers. In this step, you will learn about a major defence against malware — antivirus software. More specifically, you will learn what it is and how it detects and removes malware.

What is antivirus software?

Antivirus software, also referred to as **anti-malware software**, is a type of software designed to identify and remove malware from your computer. It can scan a computer for suspicious files and activity, and it can scan specific files or programs, attachments, and

downloads. Some programs can also give updates on a computer's performance.

Most antivirus software can be set up to scan a computer regularly, but it is a good idea to scan for malware if you notice a reduction in your computer's performance, for example, if it is running slower than usual or is unable to run particular programs, or if it is showing pop-ups when you're offline.

How does it detect malware?

Antivirus software uses lots of different approaches to detect malware. The first is a **dictionary approach**, which involves comparing files on your computer with a list of known **malware signatures** in order to find matches. A malware signature is a unique string of code in the malware that identifies it.

Attackers know that this is how antivirus software works, so they adapt their malware by slightly altering the code that runs it, in order to make it undetectable. Antivirus software therefore searches for similarities between the code in a suspicious file and the known malware in the dictionary, instead of making direct comparisons.

Antivirus software can only be effective if it has encountered the malware or a variant of it before. Therefore, it is important to keep your antivirus software up-to-date, so that it can learn about new strains of malware.

To identify new types of malware, antivirus software also takes a **heuristic approach**. This involves monitoring files for suspicious activity (for instance, if a program asks to change settings in your OS). The software might even run suspicious files or programs in a quarantined setting to see how they behave, without endangering the computer.

How does it remove malware?

When antivirus programs identify malware, they generally present three options: **clean**, **quarantine**, or **delete**. The most appropriate

approach to take is usually determined by the type of malware and the type of file or program that has been infected.

You can **clean** the file/program if you still need the file/program that has been infected. If you were to delete it, then you could lose the file, or if the malware has infected a program in your OS, your computer's ability to function could be impaired.

You can **delete** the file/program if the malware is in the form of a worm or a Trojan, because these types of malware are contained in a separate file/program (as you learned earlier).

You can **quarantine** the file/program if you are unsure, and don't want to risk deleting an important file. As the term suggests, this isolates the malware so that it can't infect any other files or programs. This allows you to check that your computer can run without the file/program, before it is deleted. It also allows you to keep malware until your antivirus software has the tools to destroy it.

Next step

In this step, you learned how antivirus software defends against malware. In the next step, you will learn about preventing attacks.

Questions

- How regularly should you run your antivirus software?
- If a friend sends you an email with an attachment that your antivirus software flags as suspicious, would you clean, quarantine, or delete it?
- Do you have any recommendations for antivirus software?

How to avoid problems

Attackers are constantly trying to find new vulnerabilities to exploit. Here, you will learn how software developers build secure software and respond to new attacks.

A continuous cycle of improvement

When designing software, developers take known types of attacks into consideration and build specific protections against them into their software. However, developers can't predict every possible attack — especially new types. Therefore, protecting software requires a **continuous cycle of improvement** in which new protections are added as frequently as possible. Whenever new attacks are discovered, software can be updated to be protected against them — but only if the software was designed to be adapted.

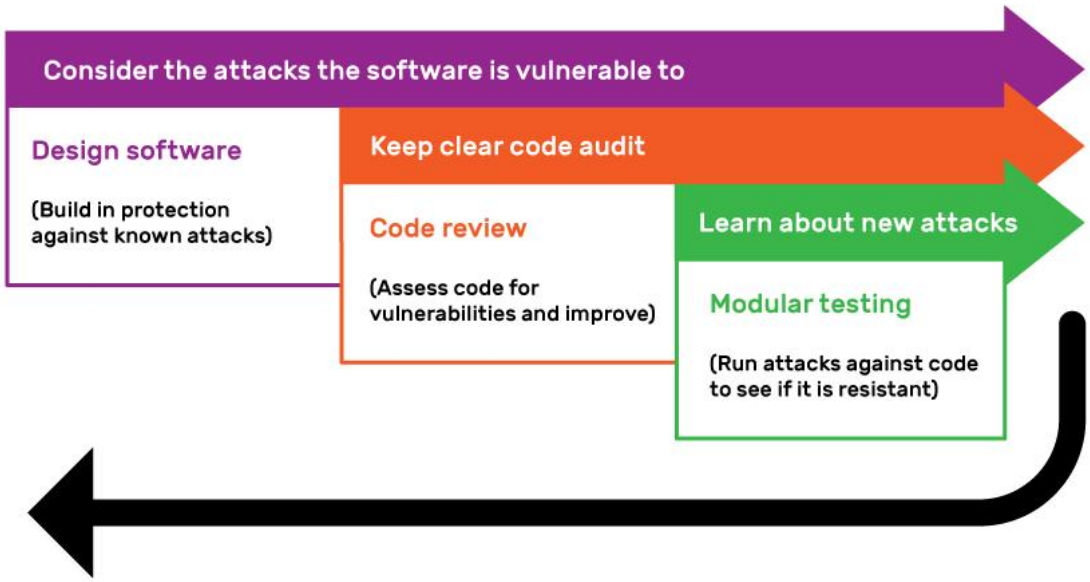

The cycle includes:

- Code review — assessing the code to identify vulnerabilities that hadn't been noticed previously
- Modular testing — testing the code in pieces to make it easier to identify and isolate problems
- Code audit — keeping a record of changes in order to understand the changes and easily fix any problems created by the changes

14 Oct 2024

Zero-hour threats

New threats to a system or piece of software are called **zero-hour threats** (also known as **0-hour threats**, **zero-hour attacks**, or **zero-day vulnerabilities**), because they are carried out before the developers are aware of their existence (i.e. in the zeroth hour).

Design
Often, attackers design zero-hour threats to cause as much damage as possible in the first 24 hours, because they know that developers will soon become aware of the attack and fix the vulnerability that it exploits.

Detection
By definition, zero-hour threats are new. This means:

- They are often harder to detect than other threats. For example, antivirus software cannot identify zero-hour viruses, because they are not in its dictionary. These threats are usually only discovered if and when victims notice the damage that they have caused.
- There are no known ways to prevent zero-hour threats or mitigate their effect. The time it takes for developers to understand a threat and work out how to protect against it is called the **vulnerability window**.

Dozens of zero-hour threats are released every day, so protecting against them is a constant process.

Patching

Zero-hour threats are fixed through **patches** (also called **bug fixes**). These are updates to software code that try to improve performance and usability, or fix a vulnerability.

- Patches that improve performance are released regularly
- Security patches are usually released on an ad hoc basis in response to new threats

When a zero-hour threat is discovered, developers work out how to prevent it, either by working out how to change the code that has

created the vulnerability, or warning the software when the attack is taking place so that it can be stopped.

Patch releases

Software developers release the patches to users, who can download them to update their software. Anyone who doesn't download a security patch is still vulnerable to the respective attack.

Some attackers send out false patches in order to trick people into downloading more malware. Because of this, developers have to make sure that users can tell authentic patch from false ones.

Because users are unwilling to constantly update their software, developers work to release as few patches as possible. They do this in two ways:

- If developers anticipate that a threat will be redesigned and will therefore require multiple patches, they may wait to release all of the patches as one update (if the threat in question is not significant)
- If a regular performance patch is due soon, developers may include the security patch in the update

Bug bounties

Software vulnerabilities are sometimes found by the developers who wrote the code, but they can't be expected to find every problem. Therefore, software companies will often offer payment (referred to as **bounties**) to anyone who can find vulnerabilities that they have missed.

Outcomes

When a member of the public (often a researcher, programmer, or professional bug hunter) comes forward with a vulnerability, generally, the company pays the bounty and fixes the vulnerability.

However, sometimes the company does not fix the vulnerability. In these instances, whoever found the vulnerability may go public with their findings in order to warn other users, potentially forfeiting their bounty by doing so.

Next step

In the next step, you will look at SQL injections, a particular example of an attack which can be adapted by attackers, making it difficult to pre-empt.

Questions

- Is the vulnerability window a fixed time period for all types of threats and for all victims? If not, why not?
- How might you tell the difference between a legitimate security patch and a false one?
- What problems might arise when a bug hunter publicises a vulnerability before the developers have patched it?

SQL injections and input sanitisation

In this step, you will learn about SQL injections, and using input sanitisation to defend against them.

SQL injections

SQL is a programming language used to communicate with databases.

When a user signs in to a website that requires a username and password, SQL is used to send a request to the database containing their username and password. If the combination is found, the user's account is returned, and if it isn't found, an error is returned (for example, "Wrong username or password").

An **SQL injection** is an attack which exploits how the requests to the database are formed. Instead of submitting a username and password, the attacker can submit two strings that trick the database into giving up its information. The characters in the strings that allow this attack are called **bad characters**.

The request to the database might look something like this:

```
SELECT row FROM database WHERE Username = "username" and Password =
"password"
```

When the user enters their username and password, these values replace ‘username’ and ‘password’ in the request. For example, if your username is ‘GoTfan’ and your password is ‘password1’, the request becomes:

```
SELECT row FROM database WHERE Username = "GoTfan" and Password =
"password1"
```

This request says “go through each row in the database and if both statements

1. Username = “GoTfan”
2. Password = “password1”

are true, return the account”.

If an attacker was able to input a string such that both statements are always true, it would trick the database into returning all of the accounts.

To do this, the attacker could submit

```
"x or ""x=x"
```

for both the username and password. If you input these values into the request, you get:

```
SELECT row FROM database WHERE Username = ""x or ""x=x"" and Password = ""x
or ""x=x""
```

This request says “go through each row in the database and if both statements

1. (Username = x) OR (x=x)

2. (Password = x) OR (x=x)

are true, return the account".

The statement `x=x` is always true, which means that even if the username and password are not both equal to `x`, both statements will return as true, so the database will return the account.

Protecting against SQL injections

When designing a database, you could design the queries that request data from the database so that the input to the form is not added directly to the query. Instead, you could search for and remove bad characters. This process is called **input sanitisation**.

Most often, any `'` or `"` characters are removed, but they are not the only characters that can be used to manipulate the request. To defend against an SQL injection, you need to know all of the possible bad characters. If you were to find a new version of the attack, you would need to release another security patch to defend against it.

If you sanitised the attacker's input to remove any `'` or `"` characters, you would get:

```
x or x=x
```

When this is inputted, the query becomes:

```
SELECT row FROM database WHERE Username = "x or x=x" and Password = "x or
x=x"
```

This request now says "go through each row in the database and if both statements

1. Username = "x or x=x"
2. Password = "x or x=x"

are true, return the account". This query is requesting an account where the username and password are both "x or x=x".

Alternatively, you could work out which characters cannot be used to carry out an SQL attack, and only allow users to create usernames and passwords from these characters. You could then reject any submission which contained anything other than good characters.

You could also hash passwords, as you learned about last week. If the passwords in the database are hashed, then the attacker's input will be altered automatically.

If the attacker used the same input as before, the request would become:

```
SELECT row FROM database WHERE Username = ""x or ""x=x"" and Password =
"e6a1826aebc1012b1444d9933684e9b5"
```

Even though the first statement is always true, the second isn't. This makes the attack harder (but not impossible, because the attacker could, in theory, produce a hash output that manipulates the second statement to always be true).

Next step

In the next step, you will learn about why people might not download security patches or keep their software up-to-date.

Questions

- Think back to the cycle of improvement described earlier. At what stage of the cycle would it be most advantageous to discover that your system is vulnerable to something like an SQL injection?
- There are many other forms of SQL injection. Can you find an example to share in the comments?
- Why are databases vulnerable to SQL injections? (Hint: think about how the website communicates with the database,

which stages of the process the attacker is modifying, and why the attacks are successful)

2.7

You've completed 3rd month

Software updates

So far this week, you have learned about many different threats to our computers and the software they host. You have also looked at the different ways to defend against these threats. These defences almost always require keeping software (antivirus or otherwise) up-to-date, so you will now find out more about the human element in these defences.

Is your software up-to-date?

You likely receive lots of offers to update different kinds of software, such as your antivirus software, your computer's OS, your phone's OS, and your apps.

- Do you keep your software up-to-date?

- Are there updates that you are more likely to install than others? For example, are you more likely to run an update on your computer than on your phone?
- What reasons might you or other users have for not keeping software up-to-date?

Share your answers in the comments.

Suggested classroom exercise

Software programs need a way of updating. They might employ a system that automatically downloads software updates when a device connects to the internet, or one that asks the user if they want to download updates every month. To understand the difficulties of convincing users to update their software safely, you could ask your students to design a software update system. Ask them to think about how frequently they need to release updates, what the updates should look like, and how much control users should have in accepting, rejecting, or scheduling updates.

What if there are no updates?

Users can only update their software if a developer is still trying to improve it and defend it against new vulnerabilities. But, this is not always the case. Earlier this the week, you learned about the WannaCry ransomware attack in 2017. Many of the attacks by WannaCry were on computers running Windows 7 or Windows XP. Support for Windows XP had ended before the attack, meaning that Windows had stopped releasing updates for it.

Computer OSes may be supported for a decade or more, but the OSes on our phones are usually only supported for a few years, perhaps three at most. Given that the average smartphone lasts for two to three years, it would be a waste of time and money for developers to support software for longer. But do developers stop supporting phones' OSes because phones are replaced so frequently, or do we replace our phones so frequently because the software on them is no longer supported? Not everyone is able to buy a new phone every three years, and given how much data we

store on our phones, the short window of OS support for phones could make a lot of customers vulnerable to attack.

Next step

In the next step, you will learn about a crucial tool in the internet's infrastructure: the bot.

Questions

- Why do you think developers stop releasing upgrades for software? Do you think this is fair?
- Do you think it is the responsibility of the software provider or the user to ensure that software is up-to-date?

What are bots, and why can they be a problem?

A lot of the attacks that you have learned about so far in the course are only effective if the attacker can repeat the same action many times. Instead of the attacker repeating this action manually, they often automate the process by using bots. In this step, you will learn what a bot is, why they are useful, and how they can be misused.

Bots

Bots are automated programs that perform tasks repeatedly. Ideally, these tasks are simple, repetitive, and performed much more quickly by bots than humans. **Internet bots** (which are also referred to as **bots**, and are the kind of bots that you will learn about in this step) perform these tasks over the internet.

Bots are a crucial part of the internet's infrastructure and perform lots of useful tasks. For example, bots identify and index new websites for search engines so that they can be included in search results. Given that an estimated 4 million blog posts are created

every day, the task is too big for humans to manage, so a bot that can process the information much faster is needed.

If a task is too large for one bot, a **botnet** might be used instead. A botnet is a network of computers which are all programmed to perform the same, or a similar, repetitive task. The bots in the net(work) can communicate with each other to effectively share the workload.

Malicious bots

However, not all bots are good. Attackers use bots to increase the scale of their attacks and to reduce their overhead costs (such as computing power and storage).

For example, bots can be used in SQL injections. An attacker might not know which websites are vulnerable to which SQL injections, so they would have to try lots of different inputs in lots of different websites. If they can program a bot to perform the same action, the process will be much faster, allowing the attacker to find more vulnerabilities. Furthermore, if the process is automated, the attacker can do something else while the program is running.

Attackers also use botnets. These are particularly useful to attackers if they want to send repeated requests to a website which rate limits the number of requests it receives from one IP address (i.e. websites that seek to prevent one individual trying to perform the same operation too many times). If an attacker controls a botnet, they can send the request from lots of different computers with different IP addresses, which fools the website into thinking that the requests are all coming from different people. This setup allows the perpetrator to carry out attacks such as DDoS attacks (you will learn about this in the next step).

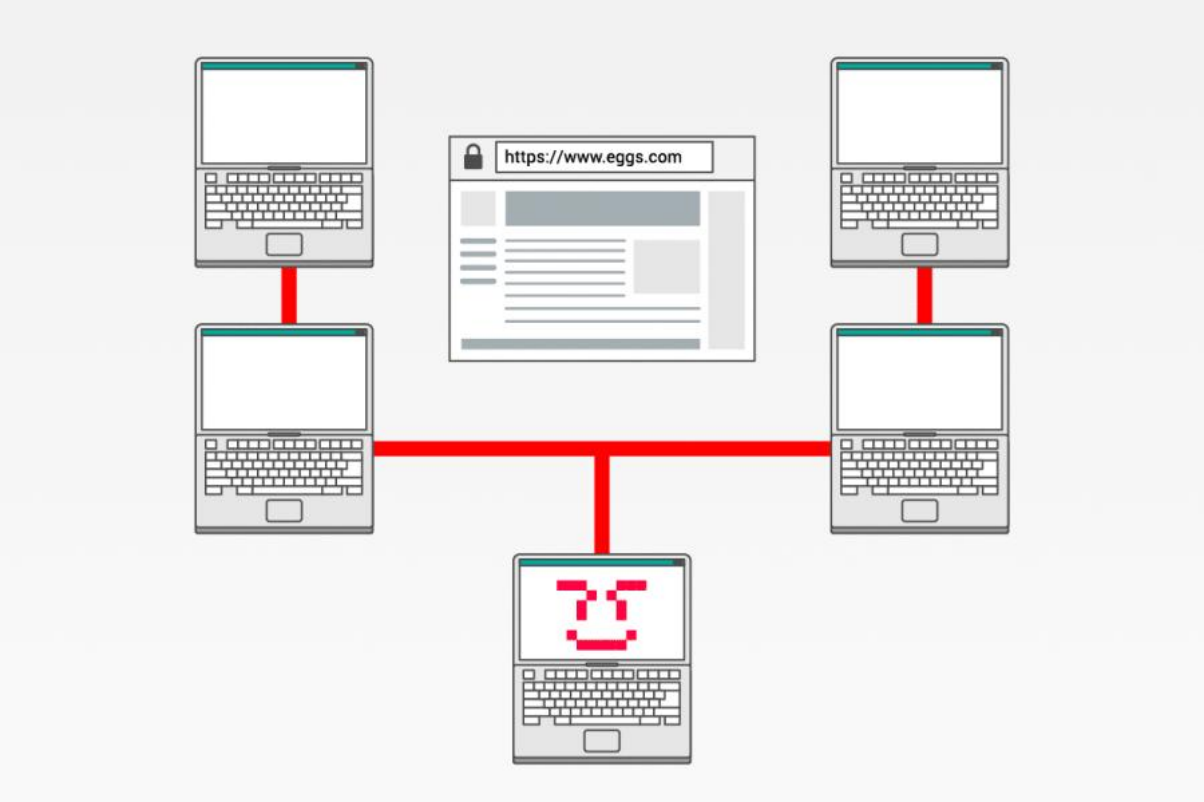

Attackers often create botnets by infecting other people's computers with malware, as you learned about earlier. This means that their victims are also paying for the energy needed to conduct the attack. When computers are infected with malware and start running more slowly, the computer may be operating as a bot in a botnet, so part of its processing power is being diverted to the task set by the attacker.

Some statistics

It is estimated that just over half of all activity on the internet is conducted by bots, rather than by humans. Furthermore, more bot activity is performed by malicious bots than by regular ones. This means that an alarmingly high proportion of internet activity is malicious; in fact, about one in every three visitors to a website is a malicious bot. These statistics come from a report by the security company Imperva — you can read more here.

Next step

In the next step, you will find out more about distributed denial-of-service attacks.

Questions

- Besides indexing new websites, what other internet processes do you think are performed by bots?
- You have learned about a few examples of malicious bots in this step. Can you think of any others? Are there any that you have encountered?
- How might websites protect themselves from botnets?

DoS and DDoS attacks

In this step, you will learn about two attacks often carried out by bots, denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks.

Denial-of-service attacks

A **denial-of-service attack**, or **DoS attack**, is any attack that aims to prevent access to a service for legitimate users. That service might be a website, an email account, a network, or a device. The attack can target any potential users of the service, or one user in particular. For example, a DoS attack could target one person's device to prevent them from accessing the internet, or it could target a website to deny access to all of its visitors.

Attackers can use DoS attacks to make companies lose business, or hold companies to ransom by threatening attack. They might also use DoS attacks to distract their victim from other types of attacks, for example, as a cover to break into a server and steal sensitive data. Sometimes this form of attack has political motivations, for example, the hacker collective Anonymous uses DoS and DDoS attacks to take down government and corporate websites that they disagree with.

There are lots of different ways of conducting a DoS attack, but broadly, they fall into two types:

- Sending illegitimate data (**teardrop attack**)

- Flooding the victim with data (**flooding attack**)

In a **teardrop attack**, the attacker sends data to the victim that the victim doesn't know how to process. It spends so long or so many resources trying to interpret the data that the service slows down or stops. For example, the attacker might send large data packets, broken down into fragments to be reassembled by the victim. The attacker might change how the packet is broken down so that the victim doesn't know how to reassemble it.

In a **flooding attack**, the attacker floods the victim with so many messages that it overwhelms them. The service slows down or stops for legitimate users, because it cannot handle so many simultaneous demands.

DoS attacks are difficult to defend against. One technique to defend against flooding is to **rate limit** users, which means only allowing individuals to send a certain number of requests per minute. However, the distributed denial-of-service attack helps attackers to get round this defence.

Distributed denial-of-service attacks

In a **distributed denial-of-service** (or **DDoS**) attack, the attacker carries out a DoS attack using several computers. These computers are often infected bots, which we discussed in the previous step.

Controlling lots of computers at the same time allows an attacker to send a greater number of messages, which increases the chances of their DoS attack being effective. Also, the bots that the attacker controls could be located anywhere in the world and would all have separate IP addresses. This means that protections like rate limiting won't stop the attack.

In a standard DoS attack, if the victim can identify the attacker, they might be able to block their messages. However, when the attacker is made up of lots of different computers, the victim might not be able to tell the difference between the bots and the legitimate users. Sometimes websites just receive a high quantity of traffic because lots of people want to use their service, and it can be extremely

difficult to tell when this is happening and when a DDoS attack is taking place. In addition, even if the victim is able to identify a few bots, they can't stop the attack unless they can identify all of them.

Next step

In the next step, you will learn about Sybil attacks, another kind of attack.

Questions

- Can you find any examples of successful DoS or DDoS attacks?
- How might a victim prevent a DoS attack in the form of illegitimate data being sent to them?
- Can you think of some creative ways to teach your students about DoS attacks?

You've completed 4th Month

Sybil attacks

In this step, you will learn about another attack that can be increased in scale, called a Sybil attack. It is an attack that can be used to cheat e-voting systems and social media websites.

What is a Sybil attack

A **Sybil attack** is when an attacker creates multiple fake accounts. When the attacker is in control of multiple accounts, they are able to gain special privileges in systems designed for one account per person. For example, on a social media site where having a lot of friends or followers increases your influence, people buy fake accounts to artificially boost their influence.

These fake accounts can be controlled by bots, and they can be programmed to share particular kinds of content. For example, after the last US presidential election, it was found that thousands of

Facebook accounts were actually run by bots programmed to share pro-Trump and anti-Clinton fake news stories. Some bots, such as chatbots, are even able to simulate a conversation.

Exercise (also suitable for use in the classroom)

Building your own chatbot is a fun way to develop your (or your students') coding skills whilst also learning about how chatbots work and how to spot them on social media sites.

Applications of a Sybil attack

Sybil attacks can have a wide variety of applications. Any platform on which users make decisions, for example e-voting platforms, and make conclusions based on majority votes, is vulnerable to Sybil attacks. If one user can create multiple fake accounts, they can increase their influence and swing votes in their favour.

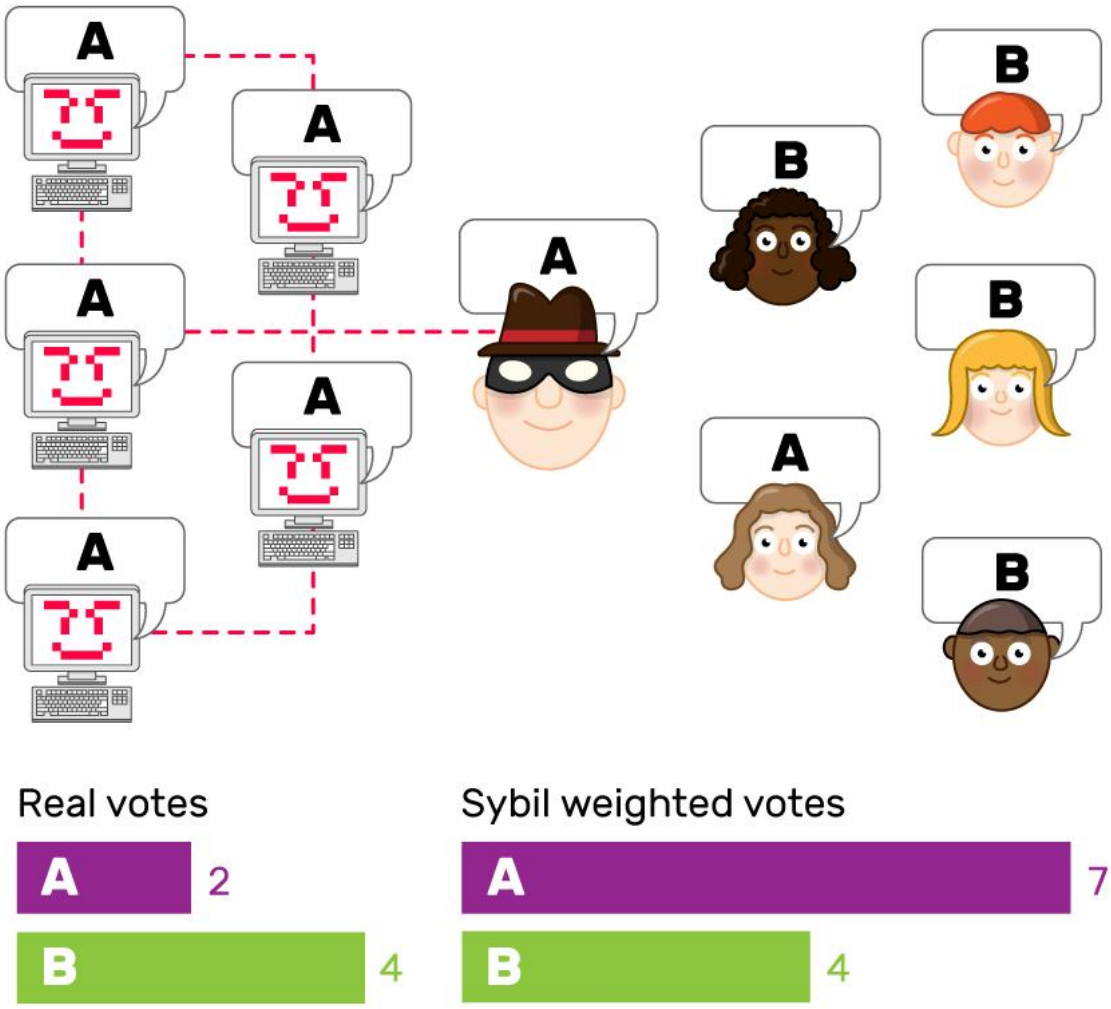

A common use of a Sybil attack is on shopping sites like eBay and Amazon. Vendors can create fake buyer accounts and leave positive reviews on their own products, tricking other buyers into believing that they are reputable. This is also seen on social media sites, where users create fake accounts to positively rate their own posts and comments, causing them to be seen by more users.

Preventing Sybil attacks

In order to prevent Sybil attacks, platforms might introduce features like reputation systems so that fake accounts have less influence than real users. For example, if a user posts a lot of content and engages with lots of other users, their votes (or reviews or content,

etc.) might be weighted more highly compared to users who interact with the platform less.

Unfortunately, when Sybil accounts are controlled by bots, they can produce lots of content, so this mitigation technique is less effective. In these instances, platforms rely on their users to report fake accounts so that they can be removed.

Users may also be required to pay to use a platform. The payment could be a literal payment, or could involve performing a task that costs the users time or other resources. For example, in order to prevent cheating, many cryptocurrencies require users to expend large amounts of computing power performing tasks such as brute forcing hash function inputs (you explored hash functions in the downloadable pdf in Step 1.7). The intention is to make it prohibitively costly to carry out a Sybil attack.

However, this technique does not apply when users are not willing to pay to use a service. This is often the case for social media sites, for example. When developing techniques to mitigate Sybil attacks, it can be difficult to make it inconvenient or expensive for attackers, but not for regular users.

Furthermore, even if a platform takes precautions to prevent attackers from creating new fake accounts, they might not be able to prevent attackers from hacking legitimate accounts and using these in a Sybil attack instead.

Next step

In the next step, you will look at other methods of mitigating Sybil attacks, and malicious bots more generally.

Questions

- What other attacks can you think of that use Sybil accounts?
- Have you ever encountered a fake account? How could you tell that it was not run by a real person?

- Can you think of any other measures that might prevent an attacker from creating lots of fake accounts?

2.11

You've completed 4th month

How to mitigate against bots

Mitigating bots

As you have seen, bots can cause a lot of damage to the internet and to individuals. You are now going to learn about how their actions might be mitigated.

CAPTCHAs

Bots are programs that perform repetitive tasks, sometimes whilst carrying out an attack. Often, they are disguised to look like legitimate users. Therefore, to stop attacks that utilise bots, users are often asked to confirm that they are human. This is often done using a **Completely Automated Public Turing test to tell Computers and Humans Apart**, also known as a **CAPTCHA**.

You have undoubtedly done many of these tests before. They ask simple questions that take a few seconds to answer if you are human, but are designed to be impossible for computers to answer. CAPTCHAs might ask you interpret some text or identify specific objects in images, or sometimes just tick a box.

CAPTCHAs are also used to collect data to help to improve machine learning tools. For example, when you answer a CAPTCHA that requires you to spot cars in pictures, the data might be used to help to train number plate recognition technology, and when you read a word, your answer can be used to help to digitise texts.

Email addresses and phone numbers

Users are also often asked for authentication. For example, websites sometimes ask for an email address or phone number when you create an account with them. This not only gives the website a way of contacting you about your account, but slows down bots who create lots of accounts (for example to conduct Sybil attacks), because they have to create a unique email address or get a new phone number for each new account.

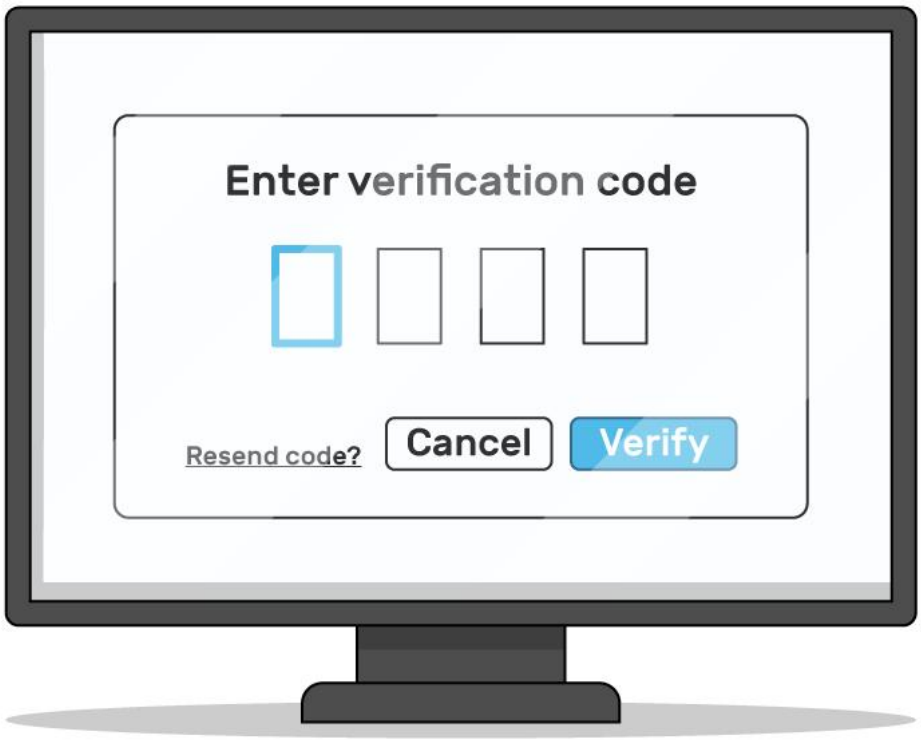

Websites might put restrictions on the type of email address or phone number, or they might send you an account activation message to verify that you own the address/phone number.

While these measures might slow down some attackers, they are not always effective, because attackers can purchase email addresses and phone numbers online for a few dollars per identity.

Suggested exercises for the classroom

To help your students to understand how CAPTCHAs work, you could ask them to design their own. To do this, they will have to think about how CAPTCHAs tell the difference between humans and robots — what sorts of tasks can humans do easily that robots can't? Don't limit your students to simple picture CAPTCHAs, get them to be creative!

Questions

- Why do you think CAPTCHAs are able to stop bots?

- What kind of attacks would a CAPTCHA prevent? Are these different to the attacks that authentication requirements are designed to prevent?
- Can you think of any other tools that mitigate against bots?

End of 4th Month

Congratulations on reaching the end of 4th month. This month, you've learned about how attackers use malware and bots in their attacks, as well as some ways to protect against these attacks.

Next week, you'll look at ways of protecting networks from attacks from both inside and outside those networks.

Was there anything you learned this week that surprised you? Let us know in the comments!

Welcome to 5th Month and 6th Month

Welcome to the final month of *Introduction to Cybersecurity Programme*. You have looked at vulnerabilities in software and how attackers could use them. Now, you'll look in more detail at networks, and how to protect them from attacks.

This week, you will learn about:

- Physically securing your network
- Using a firewall to block unwanted traffic
- Managing the network by limiting who can connect to it, and the resources that they can access

You'll then finish the course by completing an assignment based on what you've learned over the past 4 months.

You've completed 4th month

Time to reflect

It's important to take some time out to reflect on what you've learnt so far.

People who take time to reflect on their learning are much more likely to make use of it later. Not only that, but by sharing your reflections with us and our team at the Social Health Growth you will help to make our courses even better.

It will help you connect the dots between all the new knowledge you've taken in, and get your synapses firing to find new ways to apply your learning.

All the information collected is anonymous, and will be stored and handled according to Form Assembly's privacy policy/terms and conditions of SHG.

You've completed 4th month

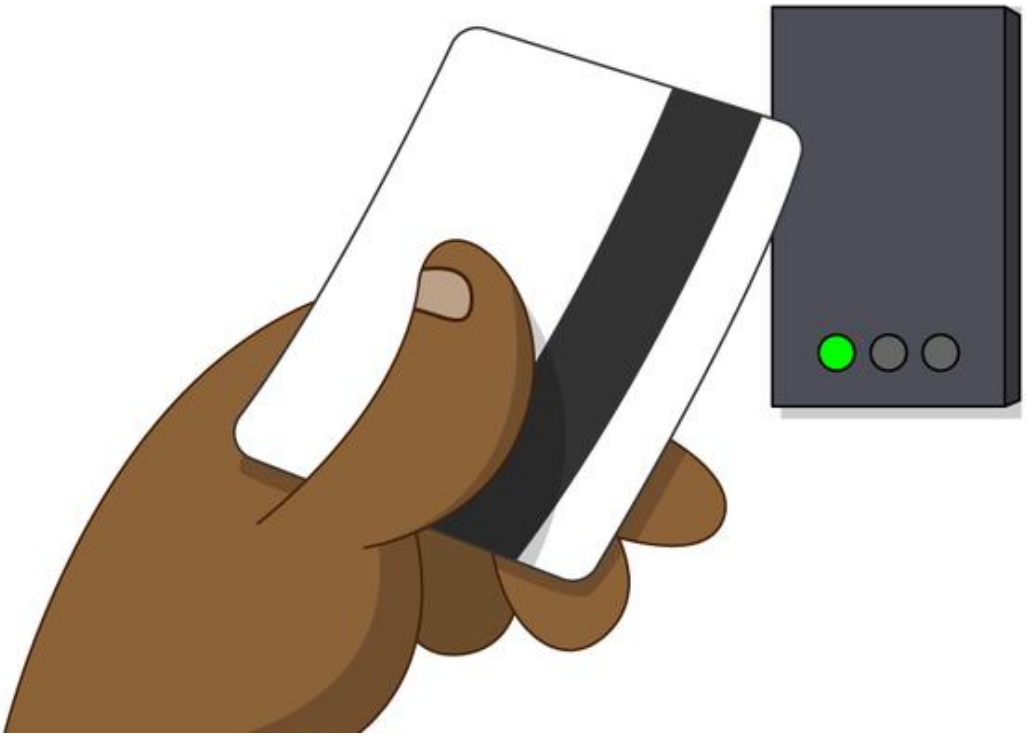

Physical security

So far in this course, you have learned about different ways to protect software and data from remote attackers. But what about the hardware that this information is stored on? In this step, we will look at how to physically protect our data systems.

Intruders

You may have developed a highly sophisticated network which can defend against attackers hacking your system, but if an intruder can easily walk onto your premises and access a computer, then your security may be redundant. Ensuring the physical security of your network may require a different set of solutions to securing your network, but many of the same principles apply.

Access to the premises

You need a method of granting and denying access to your premises which is effective and proportionate. Many systems use **key cards**, which are pieces of plastic programmed to unlock digitally secured locks, for instance, on doors. This technology can be highly specialised. For example, different members of the same organisation may have different access authorisations, so their key cards can be programmed to reflect their individual access authorisations.

Key cards are not perfect. They can be lost or stolen quite easily, and could therefore be used by the wrong person. This can be mitigated by requiring a key code to be used as well as the card (this is a form of two-factor authentication) or by turning key cards into photo ID cards. Security staff can then inspect card users to make sure that they are the legitimate owners.

As with device security, an alternative to key cards (or passwords) is **biometrics systems**. Retinal, fingerprint, and facial scanners, and voice recognition software authenticate users by verifying biological features. It is harder to lose or steal this biometric data, but biometrics systems are not infallible and can be tricked.

A further issue with door access is **tailgaiting**. Tailgaiting is when someone who does not have access permissions follows someone who does through a door or gate. Many organisations have policies to prevent tailgaiting, but these can be difficult to enforce. For instance, organisations could require employees to unlock doors to enter or leave a room or building. If, for example, only employees who have used their key card to open a door into a room are able to open the door to leave, then employees are incentivised to unlock doors themselves, instead of following their co-workers. This makes tailgaiting by attackers easier to identify.

Access to the network

Even when organisations have secured their premises against unauthorised access, they still need to protect their computers and other devices with passwords and/or the other forms of security discussed in Week 1. Passwords are only useful if the devices they protect are kept locked, so organisations may put a policy in place

to deter employees from leaving their computer unlocked and unattended.

Attackers might not need to access an employee's account to carry out an attack. If they have access to a device, they may be able to infect the network with malware through a USB flash drive. That is why some organisations disable the USB ports on their computers. In addition, if the objective is to destroy data rather than steal it, getting access to and corrupting a server might be enough to take a service offline or permanently delete important files. Therefore, it is important to keep backups in a separate location.

Next step

In the next step, you will learn about a tool for protecting networks called a firewall.

Questions

- Where necessary, many organisations allow visitors temporary access to a system. How might an organisation give visitors access to the system without granting them the same privileges as employees?
- Are you aware of the security policies concerning access to your school premises? How well do you think they are enforced?
- Are external threats the only threats an organisation needs to be wary of?

Firewalls

In this step, you will learn about a tool for protecting a network from both internal and external threats: the firewall.

What is a firewall?

Firewalls are tools that protect networks by deciding how and what information can enter and exit the network. They can be used to

protect large networks and individual computers from malware and data theft.

There are two forms of firewall:

- A **network firewall** acts as a barrier to the entire network. It normally comes in the form of a separate piece of hardware which sits between the network and the internet or any external networks.
- A **host firewall** is software which is downloaded onto an individual device and only protects this device, known as the host device.

Network and host firewalls have different strengths and weaknesses and so are used in different situations:

- If your network is made up of one computer, then you only need a host firewall. This technology is cheaper and often easier to use if you are not an IT professional. Host firewalls can also be more specialised to their host and can show more clearly if the host has been compromised.
- If your network is made up of hundreds of computers, then a network firewall offers equal protection for all of the computers. If there was no network firewall in place, and any computer in the network had an out-of-date host firewall, then it could make the whole network vulnerable.

Most companies will use a combination of network and host firewalls for extra protection.

How do they work?

Firewalls can perform lots of different functions and try to protect the network from different attacks in different ways. Here are some examples of the role a firewall can play in defending a network:

- **Filter**: The firewall filters traffic entering and leaving the network. It evaluates incoming data to identify malware and

other threats, but also checks that any outgoing data is authorised to leave the network.

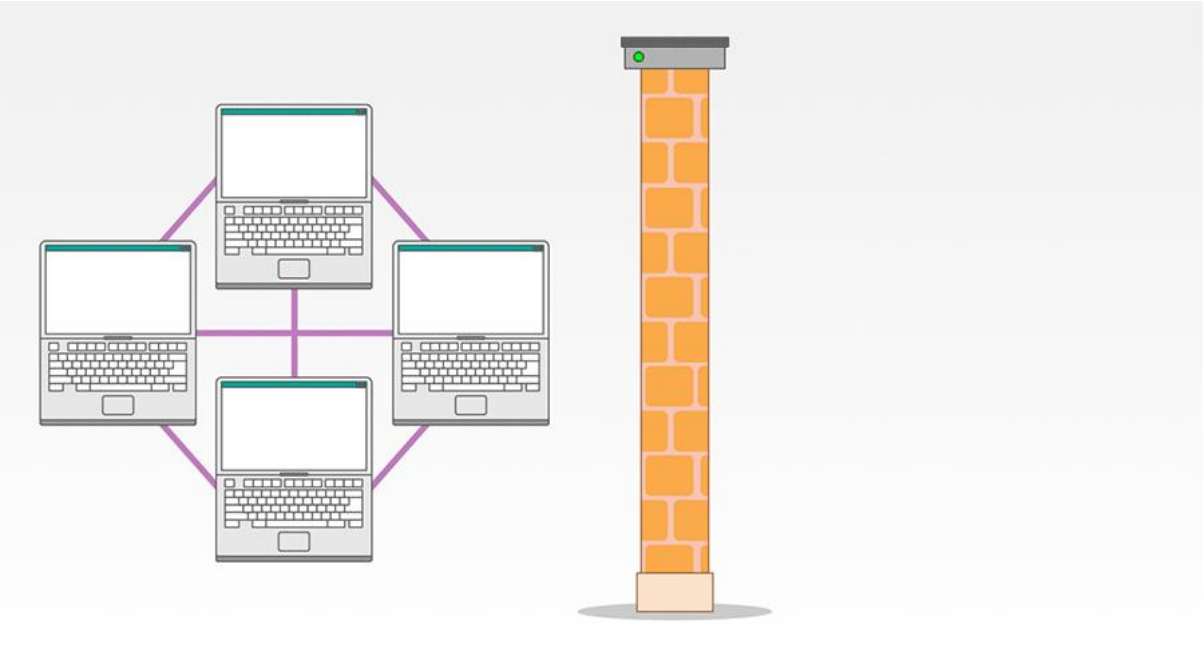

- **Access control**: Firewalls can be used to prevent external devices from accessing the network. You will find out more about this later.
- **Proxy service**: Network firewalls can act as a proxy when communicating with websites. This means that any information requests sent outside the network are sent to the firewall, which passes the requests to the recipients on the user's behalf. When the recipients respond, they respond to the firewall, which can pass the response to the user. This allows the firewall to screen potentially harmful messages, and stops the recipient from gaining access to the network. It can also speed up the user's interactions with the internet — if the firewall stores websites that the user visits regularly, then it can load the information without having to get it from the website itself.

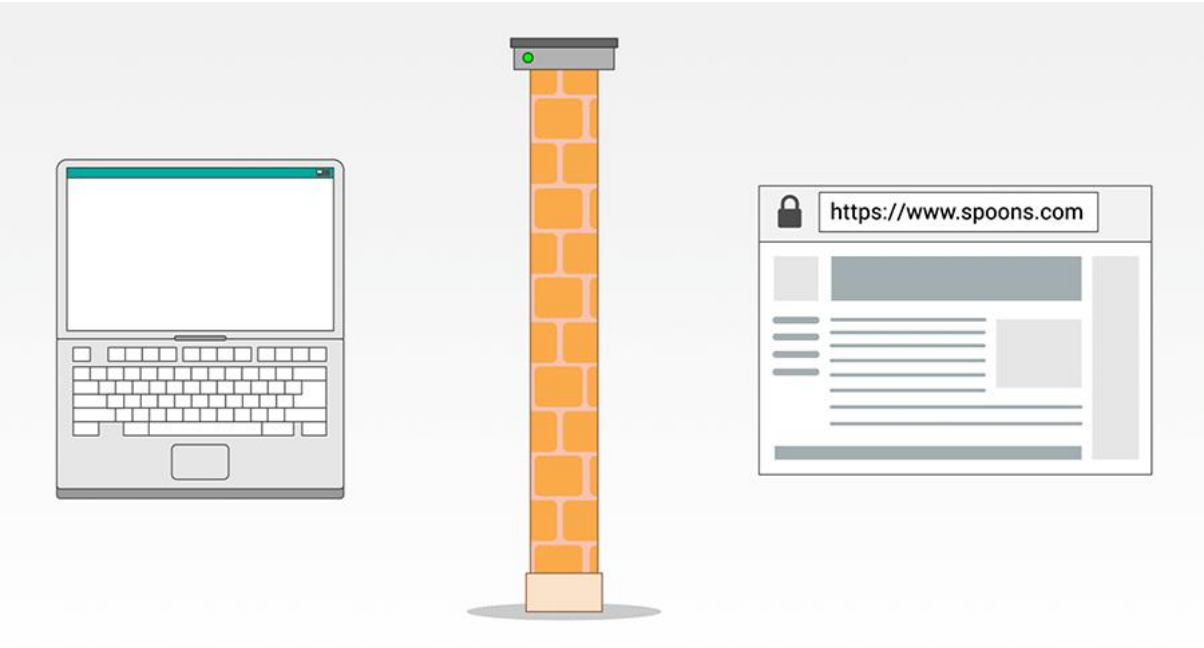

- **Block websites**: The firewall checks outgoing messages as well as incoming messages, so it can block certain requests. This allows organisations to prevent their employees from visiting or using particular websites. This is often done by implementing a **blacklist** of websites that are blocked, or by using a key word search that blocks any websites containing a particular word.

The restrictions that firewalls apply to the network can be specific to a device or an account. For example, the firewall could be set up to allow more senior staff to send information out of the network, but prevent less senior staff from doing so.

How effective are they?

Firewalls can only protect against the threats that they can identify. As you learned last week, threats change all the time, so firewalls must be kept up-to-date in order to be effective.

In addition, many firewalls are not effective because they are not configured properly. If a firewall is not set up to scan an organisation's internal network, then the network is not protected from internal threats.

Furthermore, many firewalls don’t or can’t scan encrypted traffic, so if an organisation is receiving a high number of encrypted files, it may be under threat.

Next step

In the next step, you will learn about how devices can be added to a network securely.

Questions

- How do you think the firewall detects malware trying to access the network?
- Does your school use a firewall to block certain websites? If so, what kind of sites and why? How effective is it?
- If the firewall in your school applied different restrictions to students and staff, what kind of differences would there be and why?

You’ve completed

Data access control

In the last step, you learned about how to manage access to a network. Similar systems are needed to manage access to data held in the network.

Who needs access to what data?

There are lots of different types of data contained in a school network. There might be:

- Teaching materials and files
- Student academic records
- Staff employment records and financial information
- Personal and medical information

Can you think of any others?

Now, consider all of the people who have access to your school network: teaching and administrative staff, students, guests, and potentially even attackers. Not all of the people in the network should have equal access to the data it stores, so a system is needed that can decide which people can access which data in which capacity.

What are the different kinds of access?

There are lots of different ways in which you can access data on a system. The type of access granted is often called a **permission**. Broadly, permissions can be split into four categories:

- **Read** — the ability to see data
- **Write** — the ability to change data
- **Execute** — the ability to run a program
- **Delete** — the ability to remove data

Having these four types of permission allows you to share data in different ways with different people. For example:

- Distinguishing between reading and writing allows you to share data with users who need to see it but should not be able to modify it
- Restricting who can execute files allows new software to be downloaded or updated in a controlled manner
- Defining who does and does not have delete permissions can help to prevent attacks that are aiming to destroy data

Data access control systems

A **data access control system** is a method of deciding who gets what permissions. There are several different approaches to designing a data access control system. The traditional method is for a single entity to assign classification levels to data, and levels of clearance to network users. The user can then only access data

if they have the relevant clearance. This approach is called **mandatory access control**.

Alternatively, the data owners can define the access permissions for their data. This means that if you produce a file, you can choose who can access it. This is called **discretionary access control**.

Permissions can be assigned to individuals in a network or to groups of individuals. This is often called **role-based access control**. The permissions might also be condition-specific, for example, some data may only be accessible at certain times of day to make it harder for attackers to access the data outside of working hours. This is known as **rule-based access control**. Role- and rule-based access control systems can be used in conjunction with a mandatory or discretionary access control system.

To enforce a data access control system, the methods of authentication discussed previously (such as passwords and biometrics systems) can be employed. When you sign into your account on a device, the network can grant you certain permissions to interact with data. As an added layer of security, you may be asked to re-enter your password before performing certain actions, to verify your identity.

Combining data and network access control

Any data and network access control systems used in a network have to be compatible with each other. These two systems can work together to optimise the safety of data on the system.

The access permissions granted to certain users can also be applied to machines. For example, it may only be possible to access very sensitive data on particular computers that have certain software and hardware restrictions.

The data that a user can access in a particular network is determined by what the user as an individual is authorised to access, and what the machine that they are using is authorised to access.

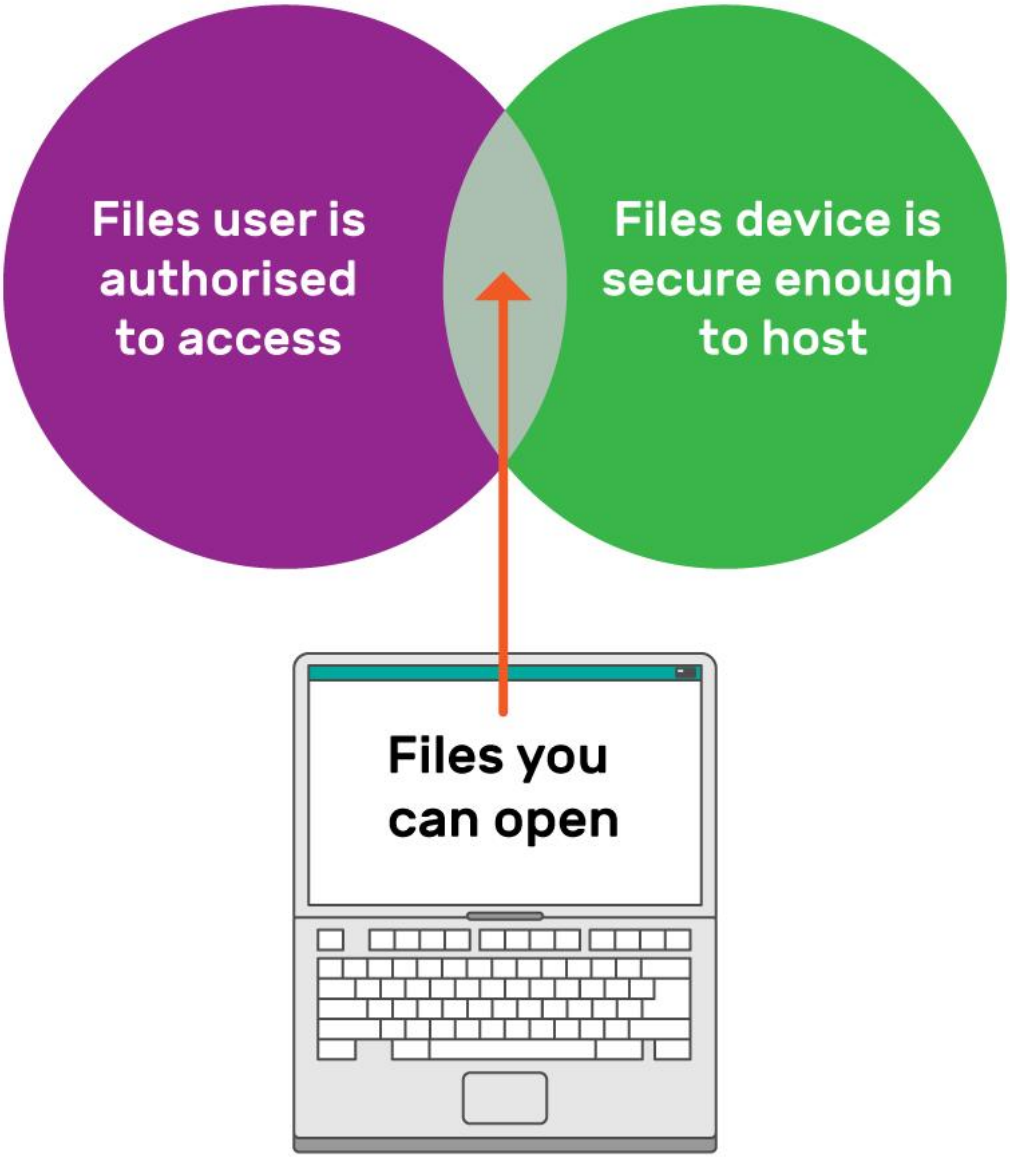

If the two systems are designed without consultation between the parties involved, users might find that they have permission to access certain data, but not the device needed to access the data.

Next step

In the next step, you will learn about internal threats to data and network access control systems.

Questions

- Can you give an example of data in your network to which you have read access but not write access?
- Which data access control system do you think is most appropriate for a school and why?
- Can you think of a scenario in which the data and network access control systems for a network are incompatible?

You've completed

Internal threats

In many of the attacks discussed so far, we have presumed that the attacker was external to the organisation. But this is not always the case. In this step, you will learn about the possibility of internal threats, and how to defend against them.

There are two types of employees to consider: **careless employees** who fail to follow the correct protocols and thereby make it easier for attackers to target the system, and **malicious employees** who intentionally try to harm the system.

Careless employees

Careless employees can enable attacks by failing to follow security protocols correctly. For example, they might leave their computer unlocked and unattended, misplace their access card, or click on a phishing link. The probability of these scenarios occurring can be reduced through staff training and instituting suitable punishments for poor security practices. However, accidents do happen, so systems need to be designed to mitigate human error.

Malicious employees

Sometimes, employees want to attack the system. They may be angry with their employers and trying to seek revenge, they may have been bribed or blackmailed into stealing data, or they may be whistleblowers, attacking their own systems to expose crimes or malpractice, or for political reasons.

Different precautions need to be taken to protect against internal threats rather than external threats. These precautions are often centred around preventing data from leaving the system (as opposed to malware or unauthorised individuals entering the system).

As we saw in a previous step, firewalls inspect traffic as it leaves a network, as well as when it enters it. This can help to identify when employees are trying to transfer sensitive data out of the system.

In some organisations, scanners that detect electronic devices are used to stop employees bringing in or taking away portable storage devices (like phones or USB flash drives). In other organisations, employees are not allowed to take their phones into any room where they have access to customers' data, like customers' names and addresses. This is to protect customers from identity theft or harassment from the employees.

Protocols

Often, the same security protocols can be used to protect a system from both careless and malicious employees. As you saw earlier, networks can restrict access to data on an employee-by-employee basis. Employees may have to undergo additional training and vetting in order to be granted higher levels of clearance, which may help to reduce the risk of carelessness. A data access control system can also help to identify which employee has broken a security protocol.

Segregation of duties can also help to protect data against the actions of employees. When this principle is used, no single employee has the power to both authorise and execute an action. For example, in order to transfer data out of the network, one employee would need to approve the request, and a second would need to transfer the data. This prevents a single employee from falling for a social engineering attack, or a single malicious employee from selling company data to a competitor. The use of the principle is a legal requirement for some actions, particularly those pertaining to sharing sensitive data.

Next step

In the next step, you will learn about the potential impact of secure data and network access control systems on employee privacy.

Questions

- If an employee clicks on a phishing link, whose responsibility is it? What should they or their employer do?
- If an employee infects a device on the network with malware, what repercussions do you think there should be?
- How can you tell the difference between a careless employee and a malicious employee?

You've completed

Privacy vs. security

In this step, you will learn about the privacy issues that can arise when designing secure systems.

In the previous few steps, you learned about how systems can be designed to defend against internal and external threats. However, in order for the tools discussed to be implemented, the system needs to store a lot of data about its users. The collection of this data could make users feel uncomfortable, or even put them in danger. You will now look at a few specific scenarios, and consider what an appropriate level of data collection is.

Do you feel that this company is retrieving and using too much data about its employees?

A company is worried that its employees have been visiting insecure websites on their lunch break. The IT security team decides to look through the employees' search histories to see who is at risk of causing a breach. In the process, they notice that one employee has been looking at job openings in other companies, and passes this information on to the employee's line manager.

Do you feel that this company is retrieving and using too little data about its employees?

A company has one wireless network router for their organisation, and every employee is able to use it with the same username and password. An employee connects to the wireless network with their phone and accidentally downloads malware from an insecure website. The malware can access the other computers and devices at the company through the wireless network.

Because the employee had not registered their phone with the company's IT team, it cannot easily be identified. In order to find the source of the malware, the IT team decides to search every employee's phone to see if they are the source.

Next step

In the next step, you will create a learning resource about a particular cybersecurity incident, using the knowledge you've gained in this course. This will be reviewed by a fellow learner, and you'll also be able to review another learner's resource, before reflecting on your own.

Questions

- In the above scenarios, how was the privacy of the employees infringed, and how might this cause problems for them?
- Can you think of any solutions to avoid privacy infringement in the above scenarios?
- Do you think that it is necessary for systems to infringe on their users' privacy to enforce security? If so, what level of privacy infringement do you think is acceptable?

You've completed

Create a learning resource

To consolidate some of your learning on this course, I would like you to create a learning resource. Your resource should focus on the 2020 Twitter verified accounts hack, and should be designed to teach high-school students about social engineering attacks, and the legal implications of these hacks.

In July 2020, hackers were able to gain control of many high-profile Twitter accounts, through what Twitter described as a "coordinated social engineering attack". These accounts were used to post messages claiming that if you sent the cryptocurrency Bitcoin to a particular address, double that amount would be returned to you. Several people were arrested in the USA and the UK in relation to this hack. You can read Twitter's description of the incident here.

To teach your students about this incident, you could create a poster, a presentation, a lesson plan, a class exercise, or something else entirely. You might want to consider some or all of the following few questions:

- How were young, relatively inexperienced hackers able to gain access to the internal tools of such a high-profile company?
- Why would some people believe that by effectively sending money to an unknown address, they would receive more back? How might this relate to other scams that don't involve the internet?
- What methods were used to find and prosecute those responsible for this hack?
- Which accounts were targeted, and why? What else could the hackers have attempted to do with access to those accounts?

You can use the internet to look up more information about this incident, to develop your learning resource.

Guidelines

Make sure that your teaching resource meets the following criteria as well as possible. Your resource should:

1. Cover one of the concepts from this course
2. Be suitable for learners of all skill levels, to assist them in learning about computer security
3. Be deliverable by anyone who has taken this course, without relying on specialist equipment or software

Submitting your resource

Share a link to your resource in the comments section. You could use any cloud-based storage or office solution, such as:

- Dropbox — https://www.dropbox.com
- Google Drive — https://drive.google.com

Reviewing your peers' work

Once you've completed your resource and shared it, I'd like you to look at several of the resources shared by your fellow learners. Pick at least one other resource to review, and **respond to the comment containing the resource with a review of that resource**. Your reviews should contain:

1. A detail from the resource that you particularly **liked**
2. A suggestion for how to **expand the resource**, perhaps to link the resource to some other learning, or to elaborate in more detail about a particular point
3. How the resource met or failed to meet the **guidelines** set out above

Writing well-considered feedback will not only help the author of that resource, but will also help you to consider how you can improve your own resource.

Improving your resource

If you come across anything in another resource that gives you an idea to improve your own, feel free to update your own resource. Make sure to thank the learner who has inspired the change, and

don’t just plagiarise someone else’s work — you should adapt it for your own purposes.

Your actions

For this step, you need to:

1. **Develop your resource and share it**
2. **Examine several other resources that have been shared**
3. **Review at least one other resource**
4. ***(Optional) Use what you have learnt to improve your own resource***

3.10

You’ve completed

What makes a good learning resource about cybersecurity?

Now that you’ve reviewed some of your peers’ work individually, it’s time to reflect on the learning resources in general.

Try to think about aspects of the resources that were specifically relevant to this course.

- Were there any common elements in the resources that worked particularly well?
- Did you spot any common mistakes that you think people should avoid?
- If you had to give advice to someone about creating a learning resource for part of this course, what would you say?

Discuss your thoughts about these resources in the comments section. Please do not link to individual projects in this step.

3.11

You've completed

Penetration testing

Penetration testing is a type of test that helps to identify what kinds of attacks an infrastructure is vulnerable to. It involves intentionally trying to attack the system in order to find its weaknesses and devise ways to defend them. This process is usually conducted through a third party.

Penetration tests can target different parts of the infrastructure and presume different types of attacker. For example, in a **black box test**, the team conducting the test is not given information about the organisation's infrastructure, whereas in a **white box test**, they are given all of the information about the system (for example, what kinds of OSes are in use, where different kinds of data are stored, who has access to which systems, etc.).

An organisation might conduct a penetration test on its internal network to find vulnerabilities in the way in which data is secured and stored, or on its external network, to find leaks or other vulnerabilities in the way in which it connects to the outside world. It might conduct a penetration test on its client-facing infrastructure, for example, by testing its website with an SQL injection.

Penetration tests are not just carried out on the organisation's computers — a penetration tester might send phishing emails to the employees to see if an attack could be facilitated through human error.

A key element of penetration testing is the production of a report, usually in the form of a risk assessment, which allows the organisation to determine which attacks it is vulnerable to, and how cost-effective it would be to take steps to prevent them.

Another continuous cycle of improvement

As you learned last week, providing any form of computer security is a constant and cyclical process. The same is true of penetration testing, which involves multiple steps of research and attack.

Companies often run penetration tests annually, or more regularly if they have introduced new systems, or if they want to check that a vulnerability has been fixed.

A penetration test might be conducted in stages (just as software is often tested module by module, as you learned in 2nd month). These tests are also often performed outside of usual working hours. This is because devastating attacks that take entire systems offline or otherwise disrupt the ability of an organisation to function as normal can be extremely costly. Penetration testing is designed to prevent these kinds of losses, so it would be counterproductive to overwhelm the system with lots of attacks, or to attack the system when it is in use.

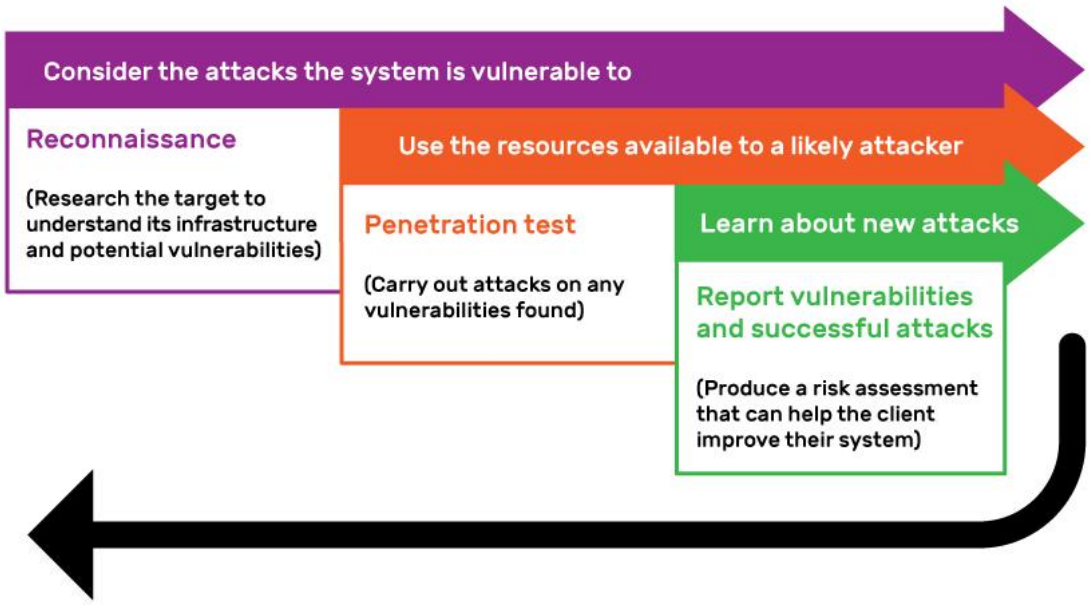

Why organisations use penetration testing

Even though penetration tests cost money, if they help an organisation to prevent more costly attacks in the future, they can save the organisation money overall.

However, this is not the only motivation for organisations to conduct a penetration test. If an organisation handles sensitive data, it may

be required by law to protect the data from theft or corruption. This obligation extends to preventing potential attacks.

In addition, the report produced in a penetration test can be used to demonstrate that an organisation has taken reasonable steps to protect the data that it holds.

Next step

In the next step, you will look at some of the key security principles that have been shared in this course.

Questions

- What kind of attacker is being simulated in a black box test, and what kind of attacker is being simulated in a white box test?
- Why might it be important for an organisation to hire a third party to conduct its penetration test?

Share your answers in the comments

Suggested classroom exercise

To help your students to understand the penetration test process, and to get them thinking about infrastructure vulnerabilities more generally, you could ask them to design their own companies and penetration tests.

Split your class into groups and ask each group to design their own data company (for example, a social media platform or a telemarketing company). In the first stage of the exercise, each group should describe their company's infrastructure (for instance, how many employees it has, who has access to what data, where the data is stored, etc.). They could give presentations to explain their companies, or write reports.

In the second stage, each group should design a penetration test on a different group's company. To do this, they will need to think about the security infrastructure and where the potential vulnerabilities are likely to be, and think of strategies to exploit

these vulnerabilities. For ideas on how to get started, you can direct your students to adverts for real penetration tests. You can find an example here.

End of the course & security principles

Congratulations on completing *An Introduction to Cybersecurity Programme*. You have learned about the different attacks that individuals, devices, and networks are vulnerable to, and how to prevent them. To finish the course, here is a summary of some of the security principles you have learned.

During this course, you have learned about: + Social engineering attacks + Malware + SQL injections + Malicious bots + Physical threats to data, devices, and networks

You have also learned about lots of different ways of protecting your data, website, or network. You've learned about tools like strong passwords, biometrics systems, 2FA, antivirus software, firewalls, and CAPTCHAs, and defence methods like input sanitisation, well-designed security protocols, and data and network access control systems.

Your data has value, so you need to take steps to protect it. Here are three security principles that you have learned during this course to help you to protect your data:

- **Be suspicious**. Look out for phishing emails and pharming websites. Malware can infect your device through the websites you visit, the emails you receive, the files you download, the apps you install, and the networks, devices, and USB flash drives you connect to. If you treat everything as suspicious and take the necessary steps to clean the files you download and screen the devices you connect to, then you are less likely to become a victim of malware.
- **Say 'no' as a default**. Avoid giving data to third parties when you don't need to. For example, don't fill out name generator apps with real data, and only give the minimum permissions to

apps and websites when asked. As you saw earlier, attackers can hack your phone through your Bluetooth connection, so keep this switched off until you need it.

- **Ensure strong defences**. You can't be expected to spot every possible attack, and if you have given your data to someone else, you are not always in a position to prevent the attack. But you can take simple steps that help to protect against attacks and mitigate the attacks that you can't prevent. For example, antivirus software and firewalls can help to identify malware from the devices, websites, and emails that seemed trustworthy. If you use strong, unique passwords, and 2FA where possible, then if one of your accounts does become compromised, the attacker will not have access to your other accounts. Computer security is a **continuous cycle of improvement**, which means that you need to keep your antivirus software up-to-date, and if you become a victim of an attack, change the password that has been compromised.

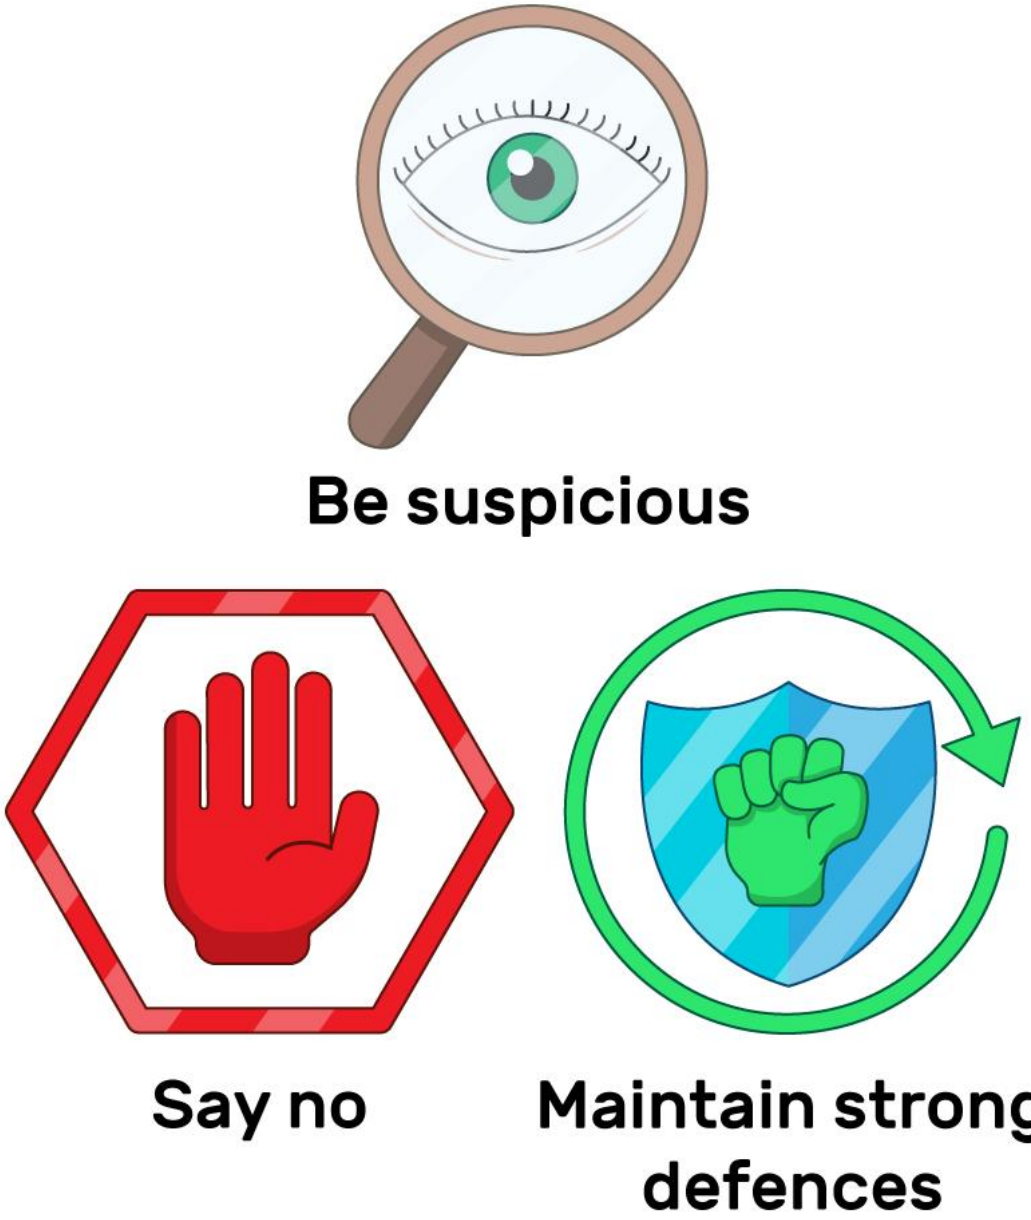

These three principles are easy to build into your everyday routine and can help to protect you and your data from attacks.

Course feedback

Share your thoughts on the course in the comments section. We'd love to hear your feedback. We'd particularly like to hear which parts of the course you enjoyed, which parts you found difficult, or something new that you learned.