National Cyber Security Centre
a part of GCHQ

CPNI
Centre for the Protection of National Infrastructure

Phishing attacks:
Defending your organisation

How to defend your organisation from email phishing attacks.

Contents

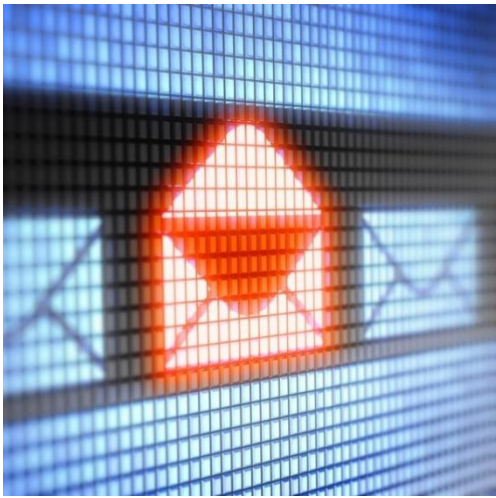

Introduction

'Phishing Attacks: Defending Your Organisation' contains advice on how organisations can defend themselves against malicious emails that use social engineering techniques.

It outlines a multi-layered approach that can improve your resilience against phishing, whilst minimising disruption to user productivity. The mitigations suggested are also useful against other types of cyber attack, and will help your organisation become more resilient overall.

This guidance is aimed at technology, operations or security staff responsible for designing and implementing defences within for medium to large organisations. This includes staff responsible for phishing training.

Staff within smaller organisations will also find this guidance useful, but should refer to the NCSC's Small Business Guide[1] beforehand.

This guidance concludes with a real-world example that illustrates how a multi-layered approach prevented a phishing attack from damaging a major financial-sector organisation.

[1] https://www.ncsc.gov.uk/guidance/avoiding-phishing-attacks

What is phishing?

Phishing describes a type of social engineering where attackers influence users to do 'the wrong thing', such as disclosing information or clicking a bad link. Phishing can be conducted via a text message, social media, or by phone, but these days most people use the term 'phishing' to describe attacks that arrive by email. Email is an ideal delivery method for phishing attacks as it can reach users directly and hide amongst the huge number of benign emails that busy users receive.

Phishing emails can hit an organisation of any size and type. Aside from the theft of information, attacks can install malware (such as ransomware), sabotage your systems, or steal money through fraud. You might get caught up in a mass campaign (where the attacker is just looking to collect some new passwords or make some easy money), or it could be the first step in a targeted attack against your company, where the aim could be something much more specific, like the theft of sensitive data. In a targeted campaign the attacker may use information about your employees or company to make their messages even more persuasive and realistic. This is usually referred to as spear phishing.

Why phishing works

Phishing works because it exploits people's social instincts, such as being helpful and efficient. Phishing attacks can be particularly powerful because these instincts also make us good at our jobs, and shouldn't be discouraged.

The mitigations included in this guidance require a combination of technological, process, and people-based approaches. They must be considered as a whole for your defences to be really effective. For example, if you want to encourage people to report suspicious emails, then you need to back that up with a technical means of doing so, and a process behind it that will provide timely feedback on the email they submitted. Only then will the user obtain any value from reporting, and the mitigation be effective.

Phishing defences: a multi-layered approach

Typical defences against phishing are reliant on users' abilities to detect phishing emails, and the NCSC has discussed the limitations of doing this[2]. However, by widening your defences, you can improve your resilience against phishing without disrupting the productivity of your users. You'll also have multiple opportunities to detect a phishing attack, and then stop it before it causes harm to your organisation. Accepting the fact that some will get through will help you plan for the day when an attack is successful, and minimise the damage caused.

This guidance splits the mitigations into four layers on which you can build your defences:

1. Make it difficult for attackers to reach your users
2. Help users identify and report suspected phishing emails
3. Protect your organisation from the effects of undetected phishing emails
4. Respond quickly to incidents

Some of the suggested mitigations may not be feasible within the context of your organisation. If you can't implement all of them, try to address at least some of the mitigations from within each of the layers. As a result, you'll be in a much better place to defend against phishing attacks.

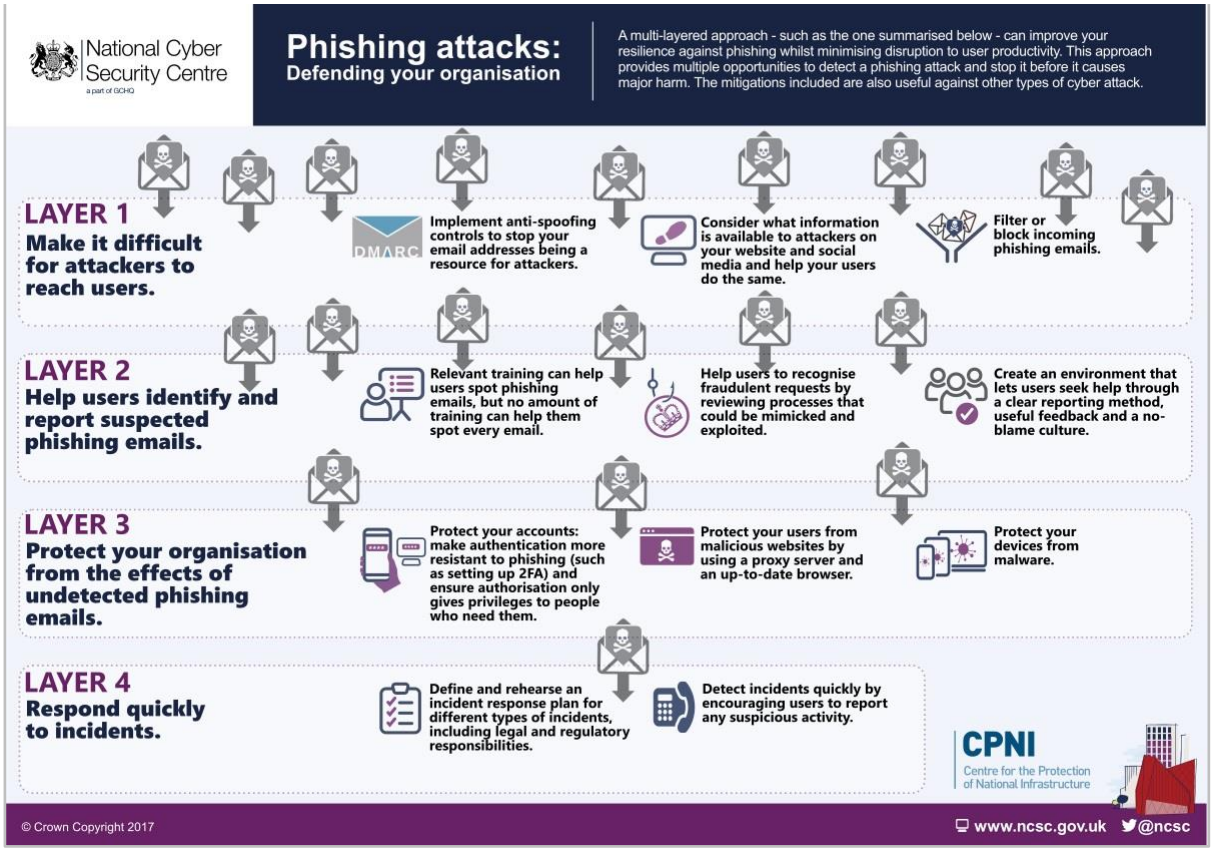

Figure 1: Summary of multi-layered approach to phishing defences

[2] https://www.ncsc.gov.uk/blog-post/im-gonna-stop-you-little-phishie

Layer 1: Make it difficult for attackers to reach your users

This section describes the defences that can make it difficult for attackers to even reach your end users.

Don't let your email addresses be a resource for attackers

Attackers 'spoof' trusted emails, making their emails look like they were sent by reputable organisations (such as yours). These spoofed emails can be used to attack your customers, or people within your organisation.

How do I do this?

Make it harder for email from your domains to be spoofed by employing the anti-spoofing controls: DMARC, SPF and DKIM[3], and encourage your contacts to do the same.

More information on implementing anti-spoofing controls like DMARC can be found in the NCSC guidance on protecting emails[4].

Reduce the information available to attackers

Attackers use publicly available information about your organisation and users to make their phishing (and particularly spear phishing) messages more convincing. This is often gleaned from your website and social media accounts (information known as a 'digital footprint').

How do I do this?

- Understand the impact of information shared on your organisation's website and social media pages. What do visitors to your website need to know, and what detail is unnecessary (but could be useful for attackers)?
- Be aware of what your partners, contractors and suppliers give away about your organisation online.
- Help your staff understand how sharing their personal information can affect them and your organisation, and develop this into a clear digital footprint policy for all users. This is not about expecting people to remove all traces of themselves from the Internet. Instead support them as they manage their digital footprint, shaping their profile so that it works for them and the organisation.
- CPNI's Digital Footprint Campaign contains a range of useful materials (including posters and booklets) to help organisations work with employees to minimise online security risks.

Filter or block incoming phishing emails

Filtering or blocking a phishing email before it reaches your users not only reduces the probability of a phishing incident; it also reduces the amount of time users need to spend checking and reporting emails.

[3] https://www.gov.uk/government/publications/email-security-standards/domain-based-message-authentication-reporting-and-conformance-dmarc

[4] https://www.ncsc.gov.uk/guidance/email-security-and-anti-spoofing

How do I do this?

- Check all incoming email for spam, phishing and malware. Suspected phishing emails should be filtered or blocked before they reach your user. Ideally this should be done on the server, but it can also be done on end user devices (ie in the mail client). Your filtering/blocking service might be a cloud-based email provider's built-in service, or a bespoke service for your own email server.
- For inbound email, anti-spoofing policies of the sender's domain should be honoured. If the sender has a DMARC policy in place with a policy of quarantine or reject, then you should do as requested if validation checks fail.
- If you use a cloud-based email provider, ensure that their filtering/blocking service is sufficient for your needs, and that it is switched on by default for all your users. If you host your own email server, ensure that a proven filtering/blocking service is in place. This can be implemented locally and/or purchased as a cloud-based service. Again, ensure that it is switched on by default for all your users.
- Filtering services usually send email to spam/junk folders, while blocking services ensures that they never reach your user. The rules determining blocking or filtering need to be fine-tuned for your organisation's needs. If you filter all suspicious emails to spam/junk folders, users will have to manage a large number of emails, adding to their workload and leaving open the possibility of a click. However, if you block all suspicious emails, some legitimate emails could get lost. You may have to change the rules over time to ensure the best compromise.
- Filtering email on end user devices can offer an additional layer of defence against malicious emails. However, this should not compensate for ineffective server-based measures, that could block a large number of incoming phishing emails entirely.
- Email can be filtered or blocked using a variety of techniques including: IP addresses, domain names, email address white/black list, public spam and open relay black lists, attachment types, and malware detection.

Layer 2: Help users identify and report suspected phishing emails

his section outlines how to help your staff spot phishing emails, and how to improve your reporting culture.

Carefully consider your approach to phishing training

Training your users - particularly in the form of phishing simulations - is the layer that is often over-emphasised in phishing defence. Your users can provide a valuable contribution to your organisation's defences, but they cannot compensate for weaknesses elsewhere[5]. This is why it is important to take a holistic approach with appropriate technical mitigations, and changes to the wider security culture of the organisation.

How do I do this?

- Make it clear that phishing messages can be difficult to spot, and you do not expect people to be able to identify them 100% of the time. Instead foster a mindset where it is OK to ask for further guidance or support when something feels suspicious, unexpected or unusual.
- Never punish users who are struggling to recognise phishing emails. Training should aim to improve your users' confidence and willingness to report future incidents.
- Ensure that your users understand the nature of the threat posed by phishing. Where possible use real examples[6] and case studies to make the threat tangible, without overwhelming people.
- Help your users spot the common features of phishing messages, such as urgency or authority cues that pressure the user to act. CPNI's Don't Take the Bait! Campaign[7] provides a range of materials to deliver security messages on this topic.
- There are many approaches you could consider for phishing training. Your users may find hands-on approaches such as quizzes, or workshops where they craft their own phishing messages, more engaging and informative.
- Some areas of your organisation may be more vulnerable to phishing. Customer-facing departments may receive high volumes of unsolicited emails, whereas staff authorised to access sensitive information, manage financial assets, or administer IT systems will be of greater interest to an attacker (and may be the target of a sophisticated spear phishing campaign). Ensure these staff are aware of the risks, and offer them additional support.
- Think carefully before you consider using phishing simulations. They can help you gain an understanding of susceptibility to specific types of phishing messages (or a clearer picture of vulnerable areas in your organisation), but 'phishing your users' may have unintended consequences. For example, it could impact on productivity by creating uncertainty about whether to respond to normal emails or users feeling 'tricked' by your organisation.
- Liaise with HR to ensure your simulations comply with your organisation's HR policies.
- CPNI, in partnership with the University of Bath, have produced a guide (PDF) for what to consider when running phishing simulations[8].

[5] https://www.ncsc.gov.uk/blog-post/im-gonna-stop-you-little-phishie

[6] https://www.ncsc.gov.uk/blog-post/serious-side-pranking

[7] https://www.cpni.gov.uk/dont-take-bait

[8] https://www.cpni.gov.uk/system/files/documents/51/d7/phishing_simulations_guide.pdf

Make it easier for your users to recognise fraudulent requests

Attackers can exploit processes to trick users into handing over information (including passwords), or making unauthorised payments. Consider which processes could be mimicked by attackers, and the importance of reviewing and improving them so phishing attacks are easier to spot (while still enabling your organisation to function).

How do I do this?

- Ensure that everyone involved is familiar with your processes, so that they are equipped to recognise unusual requests.
- Make processes more resistant to phishing by ensuring that all important email requests are verified using a second type of communication (SMS/website/phone/post/in-person). Other examples of changing processes include using a different login method, or sharing files though an access-controlled cloud account, rather than sending files as attachments.
- Think about how your outgoing communications appear to suppliers and customers. For example, do you send unsolicited emails asking for money or passwords? Will your emails get mistaken for phishing emails, or leave people vulnerable to an attack that's been designed to look like an email from you?
- Consider telling your suppliers or customers of what they should look out for (such as 'we will never ask for your password', or 'our bank details will not change at any point').

Create an environment which empowers users to seek help

Building a good reporting culture enables your users to ask for help, and gives you vital information about what types of phishing attacks are being targeted at your organisation. Both of these can help you improve your defences. You can also learn what type of emails are getting mistaken for phishing, and what impact this might be having on your organisation.

How do I do this?

- Have an effective process for users to report when they think phishing attempts may have made it past your organisation's technical defences. Is the process clear, simple and convenient to use? Do users have confidence that reports will be acted on?
- Provide feedback on what action has been taken, and make it clear that their contributions make a difference. Feedback will be more effective if it is quick and specific.
- Think about how you can use informal communication channels (through colleagues, teams, or internal message boards) to create an environment where it is easy for users to 'ask out loud' for support and guidance when they may be faced with a phishing attempt.
- Avoid creating a punishment or blame-oriented culture around phishing. It is important that users feel supported to come forward even when they have 'clicked' and later believe that something may be suspicious.

Layer 3: Protect your organisation from the effects of undetected phishing emails

Since it's not possible to stop all attacks, this section outlines how to minimise the impact of undetected phishing emails.

Protect your devices from malware

Malware is often hidden in emails or fake websites that a user is directed to. Well configured devices and good end point defences can stop malware installing, even if the email is clicked.

How do I do this?

- Prevent attackers from using known vulnerabilities by only using supported software and devices. Make sure that software and devices are always kept up to date with the latest versions from software developers, hardware suppliers and vendors.
- Prevent users accidentally installing malware from a phishing email, by limiting administrator accounts to those who need those privileges. People with administrator accounts should not use these accounts to check email or browse the web.
- There are many other defences against malware and you will need to consider your security needs and ways of working to ensure a good approach. Some defences are specific to particular threats (such as disabling macros[9]) and some may not be appropriate for all devices (anti-malware software may be preinstalled on some devices and not needed on others, see our blog on using antivirus on mobile devices[10].) For more information, see the NCSC end user device guidance.
- The impact of malware on your wider system will depend on how your system has been set up. For more information on this see our security design principles[11].

Protect your users from malicious websites

Links to malicious websites are often a key part of a phishing email. However, if the link is unable to open the website, then the attack cannot continue.

How do I do this?

- Most modern, up-to-date browsers will block known phishing and malware sites. Note that is not always the case on mobile devices.
- Organisations should run a proxy service, either in house or in the cloud, to block any attempt to reach websites which have been identified as hosting malware or phishing campaigns.
- Public sector organisations should use the Public Sector DNS service[12], which will prevent users resolving domains known to be malicious.
- The NCSC's Active Cyber Defence program[13] is working to block and takedown malicious websites.

[9] https://www.ncsc.gov.uk/guidance/macro-security-microsoft-office

[10] https://www.ncsc.gov.uk/blog-post/av-or-not-av

[11] https://www.ncsc.gov.uk/guidance/design-principles-reducing-impact-compromise

[12] https://www.ncsc.gov.uk/information/uk-public-sector-dns-service

[13] https://www.ncsc.gov.uk/active-cyber-defence

Protect your accounts with effective authentication and authorisation

Passwords are a key target for attackers, particularly if they are for accounts with privileges such as access to sensitive information, handling financial assets, or administering IT systems. You should make your login process to all accounts more resistant to phishing, and limit the number of accounts with privileged access to the absolute minimum.

How do I do this?

- Add additional security to your login process by setting up Two Factor Authentication (2FA), which is also called 'Two Step Verification' on some web services. 2FA is supported by many web services, with some offering enterprise solutions in addition to the basic options. Having a second factor means that an attacker cannot access an account using just a stolen password.
- Passwords are often stolen by tricking a user into typing their password into a fake website. Some password managers can recognise real websites and will not autofill on fake websites. Similarly, you could use a single sign-on method (where the device recognises and signs into the real website automatically). Adopting these techniques means that manually entering passwords becomes unusual, and a user can more easily recognise a suspicious request.
- Consider using alternative login mechanisms that require more effort to steal, like biometrics or smartcards.
- The damage an attacker can cause is proportionate to the privileges allocated to the credentials they have stolen. The more your users can do, the more harm an attacker may achieve. Only provide privileged access to people who need it for their roles. Regularly review users that have been provided with privileged accesses to ensure that they are still needed. If this is no longer the case, then privileged accesses should be revoked. Privileged accesses should be the exception not the norm, and only provided for as long as they are needed.
- Remove or suspend accounts that are no longer being used, such as when a member of your organisation leaves or moves to a new role.
- Consider evaluating your existing policies. For example, a password policy should reduce the chance that people will reuse a work password on a personal account (where it may be more vulnerable to phishing). For more information see the NCSC password guidance[14] or the Secure by Default case studies[15].

[14] https://www.ncsc.gov.uk/guidance/password-collection

[15] https://www.ncsc.gov.uk/information/case-studies-secure-default-partnership-programme

Layer 4: Respond quickly to incidents

All organisations will experience security incidents at some point, so make sure you're in a position to detect them quickly, and to respond to them in a planned way.

Detect incidents quickly

Knowing about an incident sooner rather than later allows you to limit the harm it can cause.

How do I do this?

- Users should feel confident reporting incidents without fear of punishment or blame (or they may be reluctant to report future incidents).
- Users should know in advance how they can report. Bear in mind that they may be unable to access normal means of communication if their device is compromised.
- Having a security monitoring capability can pick up on incidents your users are not aware of, although this is not suitable for all organisations as it is very resource intensive. As a starting point you can gain visibility of your systems/networks by collecting logs (for example history of emails received, web addresses accessed and connections to external IP addresses). For those with enough resources, and a strong security need, this can be expanded into reactive monitoring against known threats.
- To collect this information, you can use monitoring tools built into your off-the-shelf services (such as cloud email security panels), build an in-house team, or outsource to a managed security monitoring service. The amount you collect and store will depend on your budget, the volume of logs, and how much you are able to analyse. Cloud storage can prevent storage capacity being a limiting factor.
- Once a monitoring capability has been set up, it needs to be kept up to date to ensure it remains effective.

Have an incident response plan

Once an incident is discovered, you need to know what to do to prevent any further harm as soon as possible.

How do I do this?

- Ensure that your organisation knows what to do in the case of different types of incidents. For example, how will you force a password reset if the password is compromised? Who is responsible for removing malware from a device, and how will they do it? For more information refer to the Incident Management section of 10 Steps To Cyber Security[16].
- Ensure your response plan complies with the legal and regulatory responsibilities of your organisation.
- Incident response plans should be practiced before an incident occurs. At a minimum, ensure that everyone is familiar with their roles and knows who to call for further support.
- To help improve your defences against phishing incidents, you may want answers to certain questions. For example, how, when, and to what extent the incident has affected the organisation? If your organisation is collecting logs as part of your monitoring, these can be used to help you answer these questions.

[16] https://www.ncsc.gov.uk/guidance/10-steps-incident-management

Real-world example of multi-layered phishing mitigations

The following example illustrates a real attack on a company in the financial sector and shows how effective layering of their defences gave them a strong overall defence.

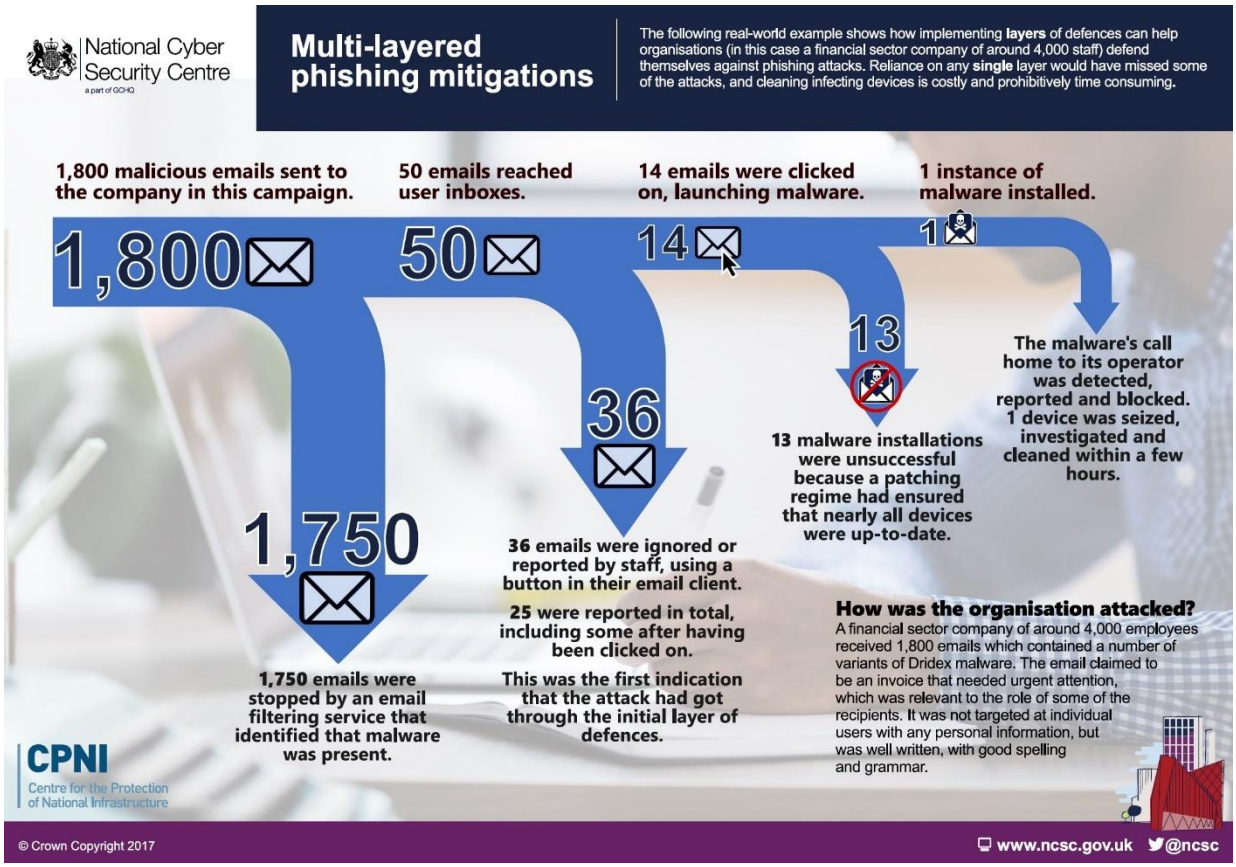

A financial sector company of around 4,000 employees received 1,800 emails which contained a number of variants of Dridex malware. The email claimed to be an invoice that needed urgent attention, which was relevant to the role of some of the recipients. It was not targeted at individual users with any personal information, but was well written, with good spelling and grammar.

- 1,800 emails were sent to the organisation by this campaign
    - 1,750 were stopped by an email filtering service that identified that malware was present.
- This left 50 emails that reached user inboxes.
    - Of these, 36 were either ignored by users, or reported using a button in their email client. 25 were reported in total, including some post click; this was the first indication that the attack had got through the initial layer of defences.
- This left 14 emails that were clicked-on, which launched the malware.
    - 13 instances of the malware failed to launch as intended due to devices being up-to-date.
- 1 instance of malware was installed.
    - The malware's call home to its operator was detected, reported and blocked.
    - 1 device was seized, investigated and cleaned in a few hours.

Reliance on any single layer would have missed some of the attacks, or in the case of relying on cleaning up quickly afterwards, be very costly and prohibitively time consuming.

National Cyber Security Centre
a part of GCHQ

CPNI
Centre for the Protection of National Infrastructure

Phishing attacks:
Defending your organisation

How to defend your organisation from email phishing attacks.