Phishing attacks: defending your organisation

How to defend your organisation from email phishing attacks.

Introduction

This guidance suggests mitigations to improve your organisation's resilience against phishing attacks, whilst minimising disruption to user productivity. The defences suggested in this guidance 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 for medium to large organisations. This includes staff responsible for phishing training.

> **Note:**
> Staff within smaller organisations will also find this guidance useful, but should refer to the NCSC's Small Business Guide beforehand.

What is phishing?

Phishing is when attackers send scam emails (or text messages) that contain links to malicious websites. The websites may contain malware (such as ransomware) which can sabotage systems and organisations. Or they might be designed to trick users into revealing sensitive information (such as passwords), or transferring money.

Phishing emails can hit an organisation of any size and type. You might get caught up in a mass campaign (where emails are sent indiscriminately to millions of inboxes), or it could be the first step in a targeted attack against your company, or a specific employee. In these targeted campaigns, the attacker uses information about your employees or company to make their messages even more persuasive and realistic. This is usually referred to as **spear phishing**.

The mitigations described in this guidance are mostly focused on preventing the impact of phishing attacks within your organisation, but if you implement these measures, you will be helping to protect the whole of the UK. Setting up DMARC, for example, stops phishers from spoofing **your** domain (that is, making their emails look like they come from your organisation). There are numerous benefits in doing this:

1. Your own company's genuine emails are more likely to reach the recipients' inboxes, rather than getting filtered out as spam.
2. From a reputational aspect, no organisation wants their name becoming synonymous with scams and fraud.
3. The more organisations set up DMARC, the harder it is for the phishers to succeed.

Why you need a multi-layered approach

Phishing mitigations often place too much emphasis on users being able to spot phishing emails. As we explain below, this approach risks wasting both time and money without improving security. Instead, you should widen your defences to include technical measures, with user education being just one aspect of your approach. A layered approach means you'll have multiple opportunities to detect a phishing attack, and then stop it before it causes harm. Some phishing attacks will always get through, so you should plan for incidents which means you can minimise the damage they cause.

The mitigations below require a combination of **technological**, **process**, and **people-based** approaches. They *all* must be considered for your defences to be really effective. More specifically, the 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 messages
3. Protect your organisation from the effects of undetected phishing emails
4. Respond quickly to incidents

If you can't implement all of the mitigations, try to address at least some of the mitigations **from within each of the layers**.

The problems with phishing simulations

No training package, including phishing simulations, can teach users to spot *every* phishing attempt. Asking users to examine, in depth, every email they receive will not leave enough hours in the day for work tasks. It's an unrealistic and counter-productive goal because responding to emails and clicking links is an integral part of work.

Phishing simulations can also create legal risk. Since no one can be expected to spot all phishing emails, punishing people for clicking on emails you've sent starts to resemble entrapment. For this reason, you should always check with your HR department before undertaking any phishing simulations (the NPSA has a set of free resources to help you design training.)

More practically, blaming users for clicking on links doesn't work. People click for a range of reasons. These could be personality traits or situational (for example, if a person is busy and stressed). Threatening someone with punishment doesn't change these factors.

Phishing simulations also erode trust between employees and security. Employees who are afraid for their jobs will not report mistakes. Employees should instead create a positive cyber security culture so employees feel comfortable reporting phishing incidents, and in this sense, they can be a valuable early warning system.

So why are phishing simulations so popular? One reason is that they allegedly provide clear, quantitative metrics that demonstrate how progress (in an area you care about) is being made. However, metrics express an organisation's values, and if you appear to value the absence of reports of problems, you incentivise people to keep quiet about issues. You should consider how you can formulate your security metrics to also include successes. For example, as well as measuring how many people clicked on a phishing email, focus on how many people reported it.

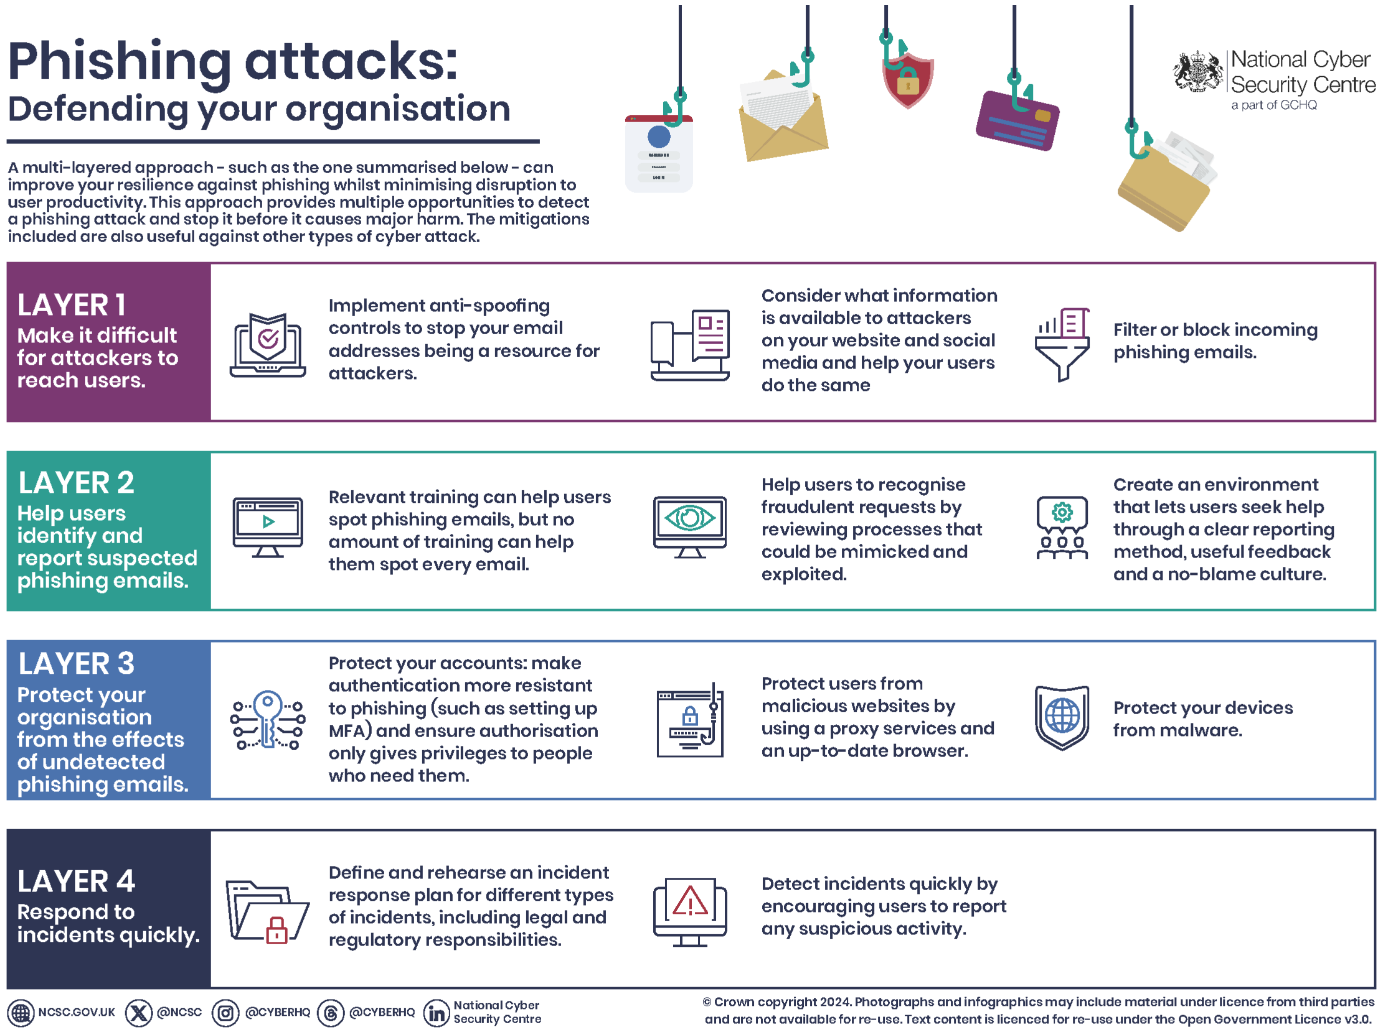

Phishing attacks:
Defending your organisation

National Cyber Security Centre
a part of GCHQ

A multi-layered approach - such as the one summarised below - can improve your resilience against phishing whilst minimising disruption to user productivity. This approach provides multiple opportunities to detect a phishing attack and stop it before it causes major harm. The mitigations included are also useful against other types of cyber attack.

LAYER 1 Make it difficult for attackers to reach users.	Implement anti-spoofing controls to stop your email addresses being a resource for attackers.	Consider what information is available to attackers on your website and social media and help your users do the same	Filter or block incoming phishing emails.
LAYER 2 Help users identify and report suspected phishing emails.	Relevant training can help users spot phishing emails, but no amount of training can help them spot every email.	Help users to recognise fraudulent requests by reviewing processes that could be mimicked and exploited.	Create an environment that lets users seek help through a clear reporting method, useful feedback and a no-blame culture.
LAYER 3 Protect your organisation from the effects of undetected phishing emails.	Protect your accounts: make authentication more resistant to phishing (such as setting up MFA) and ensure authorisation only gives privileges to people who need them.	Protect users from malicious websites by using a proxy services and an up-to-date browser.	Protect your devices from malware.
LAYER 4 Respond to incidents quickly.	Define and rehearse an incident response plan for different types of incidents, including legal and regulatory responsibilities.	Detect incidents quickly by encouraging users to report any suspicious activity.	

NCSC.GOV.UK @NCSC @CYBERHQ @CYBERHQ National Cyber Security Centre

Infographic: Summary of layered defences (download available below)

Four layers of mitigation

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, and encourage your contacts to do the same.

Reduce the information available to attackers

Attackers use information freely available on your website and social media accounts (known as your 'digital footprint') to make spear-phishing messages more convincing.

How do I do this?

- Consider what visitors to your website *need* to know, and what detail is unnecessary (but could be useful for attackers)? This is particularly important for high-profile members of your organisation, as this information could be used to craft personalised whaling attacks (a type of spear phishing that targets a 'big phish', such as a board member who has access to valuable assets).
- 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. The NPSA's Digital Footprint Campaign contains a range of useful materials that can help with this.
- Be aware of what your partners, contractors and suppliers give away about your organisation online.

Filter or block incoming phishing emails

Emails should be filtered/blocked for spam, phishing and malware before they reach your users. Ideally this should be done on the server, but it can also be done on devices (ie in the mail client). 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 will need to be fine-tuned for your organisation's needs.

How do I do this?

- 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.
- 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 clicking on a bad link. 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, and to respond to your business's changing needs and ways of working.
- Filtering email on 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 allow/deny list, public spam and open relay deny lists, attachment types, and malware detection.

Layer 2: Help users identify and report suspected phishing emails

This 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 defences. However, spotting all phishing emails is hard, and spear phishing attacks are even harder to detect. The advice given in many training packages (based on standard warnings and signs) will help your users spot some phishing emails, but they can't teach everyone to spot all phishing emails.

How do I do this?

- Ensure that your users understand the nature of the threat posed by phishing, especially those departments that may be more vulnerable to it. 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 more vulnerable staff are aware of the risks, and offer them additional support.
- Help your users identify the common features of phishing messages. The NCSC has produced guidance on how to spot scam messages.
- Don't reprimand users who are struggling to recognise phishing emails. Users who fear reprisals will not report mistakes promptly, if at all. Training should re-assure users that they won't get in trouble if they report phishing incidents. This message needs buy-in across all departments including HR, support and senior management.
- Rather than using simulations, some companies ask participants to craft their own phishing emails, giving them a much richer view of the influence techniques used. A friendly competition between peers can avoid unhelpful 'us vs them' scenarios that phishing simulations may engender, where staff may feel they are being tested by the security team.

Make it easier for your users to recognise fraudulent requests

Attackers can exploit 'ways of working' to trick users into handing over information (including passwords), or making unauthorised payments.

Consider which processes could be mimicked by attackers, and how to review and improve them so phishing attacks are easier to spot.

How do I do this?

- Ensure staff are familiar with the normal ways of working for key tasks (such as how payments are made), so they're better 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 (such as SMS message, a phone call, logging into an account, or confirmation by post or 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. Can your recipients easily distinguish your genuine email from a phishing attack? Is the recipient expecting an email, and will they recognise your email address? Do they have any way of knowing if links are genuine?
- Consider telling your suppliers or customers what to look out for (such as 'we will never ask for your password', or 'our bank details will not change at any point'). This gives the recipient another chance to detect a phish.

Create an environment that encourages users to report phishing attempts
Building a culture where users can report phishing emails (including where they've been clicked-on) gives you vital information about what types of phishing attacks are being used. 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 phishing. Is the process clear, simple and quick to use? Quickly provide feedback on what action has been taken, and make it clear that their contributions make a difference.

- 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 'successful' phishing emails

Since it's not possible to stop all attacks, this section outlines how to minimise the impact of phishing emails that reach your users, and are clicked.

Protect your devices from malware
Malware is often hidden in phishing emails, or in websites that they link to. Well-configured devices and good end point defences can stop malware installing, even if the email is clicked. Some defences are specific to particular threats (such as disabling macros) and some may not be appropriate for all devices (anti-malware software may be pre-installed on some devices and not needed on others). Finally, the impact of malware on your wider system will depend on how your system has been set up. For more information, refer to the section on reducing the impact of compromise from our Secure design principles.

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 patches.

- 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.

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 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, which will prevent users resolving domains known to be malicious.

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 multi-factor authentication (MFA), which is also called 'two step verification (2SV)' on some web services. Having a second factor means that an attacker cannot access an account using just a stolen password.
- Consider using password managers, some of which 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 authentication mechanisms (like biometrics or smartcards) that require more effort to steal than passwords.
- The damage an attacker can cause is proportionate to the privileges allocated to the credentials they have stolen. Only provide privileged access to people who need it for their roles. Regularly review these and revoke privileges if no longer 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 reviewing your password policies. Doing so may (for example) reduce the chance likelihood of staff re-using passwords across home and work accounts.

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?

- Ensure users know in advance how they can report incidents. Bear in mind that they may be unable to access normal means of communication if their device is compromised.
- Use a security logging system to pick up on incidents your users are not aware of. 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.

- Smaller organisations that may lack dedicated logging resources may wish to try CISA's Logging Made Easy open source project, which provides a practical way to set up basic end-to-end Windows monitoring of your IT estate.
- 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 NCSC's Incident Management guidance.
- Incident response plans should be practised before an incident occurs. The best way to do this is through exercising. If you're new to this, the NCSC has created Exercise In A Box, a free online tool which helps you to find out how resilient you are to cyber attacks, and lets you practise in a safe environment.

Case study: example of multi-layered phishing mitigations

The following real-world example illustrates how a company in the financial sector used effective **layered** mitigations to defend against phishing attacks. Reliance on any single layer would have missed some of the attacks, and resulted in a costly and time consuming clean-up operation.

The company, which has around 4,000 employees, received 1,800 emails containing 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.

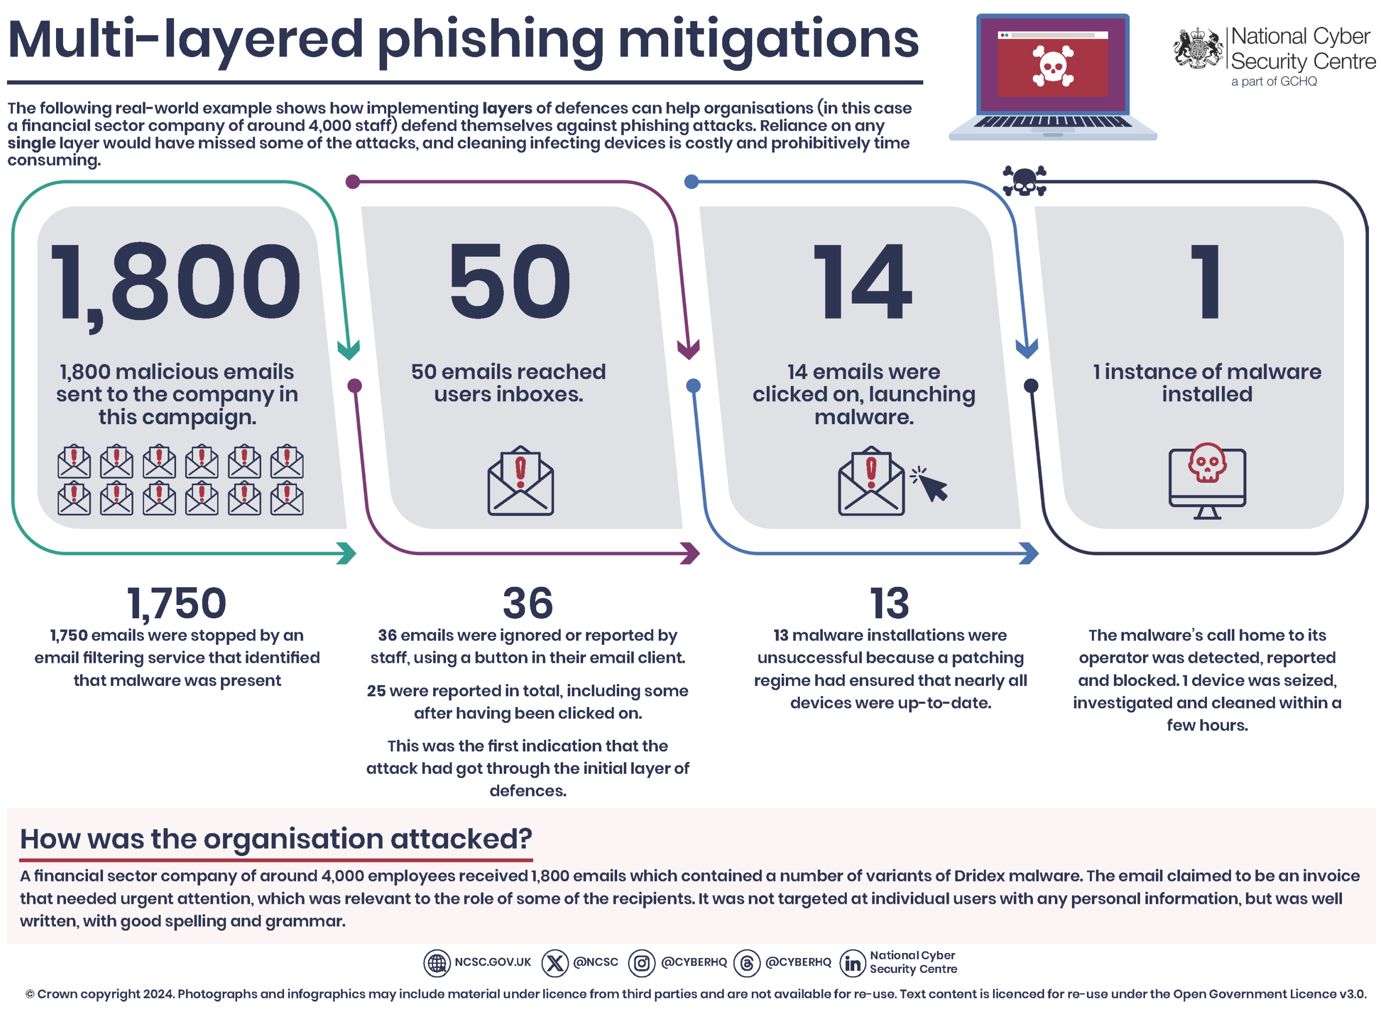

Multi-layered phishing mitigations

National Cyber Security Centre
a part of GCHQ

The following real-world example shows how implementing **layers** of defences can help organisations (in this case a financial sector company of around 4,000 staff) defend themselves against phishing attacks. Reliance on any **single** layer would have missed some of the attacks, and cleaning infecting devices is costly and prohibitively time consuming.

1,800

1,800 malicious emails sent to the company in this campaign.

50

50 emails reached users inboxes.

14

14 emails were clicked on, launching malware.

1

1 instance of malware installed

1,750

1,750 emails were stopped by an email filtering service that identified that malware was present

36

36 emails were ignored or reported by staff, using a button in their email client.

25 were reported in total, including some after having been clicked on.

This was the first indication that the attack had got through the initial layer of defences.

13

13 malware installations were unsuccessful because a patching regime had ensured that nearly all devices were up-to-date.

The malware's call home to its operator was detected, reported and blocked. 1 device was seized, investigated and cleaned within a few hours.

How was the organisation attacked?

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.

NCSC.GOV.UK @NCSC @CYBERHQ @CYBERHQ National Cyber Security Centre

Infographic: Multi-layered phishing mitigations (download available below)

Summary of the phishing attack:

- **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.

PUBLISHED

5 February 2018

REVIEWED

13 February 2024

VERSION

2.0

WRITTEN FOR

Small & medium sized organisations

Large organisations

Public sector

Cyber security professionals