GUIDANCE

Phishing attacks: defending your organisation

How to defend your organisation from email phishing attacks.

Introduction to Phishing

Phishing attacks: defending your organisation provides a multi-layered set of 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.
- Staff within smaller organisations will also find this guidance useful, but should refer to the NCSC's Small Business Guide 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.

Note: 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.

What is phishing?

Phishing is when attackers attempt to trick users into doing 'the wrong thing', such as clicking a bad link that will download malware, or direct them to a dodgy website.

Phishing can be conducted via a text message, social media, or by phone, but the term 'phishing' is mainly used to describe attacks that arrive by email. Phishing emails can reach millions of users directly, and hide amongst the huge number of benign emails that busy users receive. Attacks can install malware (such as ransomware), sabotage systems, or steal intellectual property and money.

Phishing emails can hit an organisation of any size and type. 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**.

Every organisation can play a part

The mitigations described here are mostly focused on preventing the impact of phishing attacks within your organisation, but they include some measures that will help protect the whole of the UK. For example, setting up DMARC 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 wider community will also benefit if your contacts (such as suppliers, partners and customers) are encouraged to register their details with DMARC. This can give you much greater assurance that the email asking for information (or money) actually comes from where you think.

The NCSC are encouraging organisations to lead by example and set up DMARC, and then start asking their contacts to do the same. It's in everyone's interest to promote widespread adoption, as the more organisations that take part, the harder it is for the phishers to succeed.

Phishing defences: why you need a multi-layered approach

Typical defences against phishing often rely exclusively on users being able to spot phishing emails. This approach will only have limited success. Instead, you should widen your defences to include more technical measures. This will improve your resilience against phishing attacks without disrupting the productivity of your users. You'll have multiple opportunities to detect a phishing attack, and then stop it before it causes harm. You also acknowledge that some attacks will get through, as this will help you plan for incidents, 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**. The mitigations within each layer are summarised in the following infographic.

Summary of multi-layered approach to phishing defences

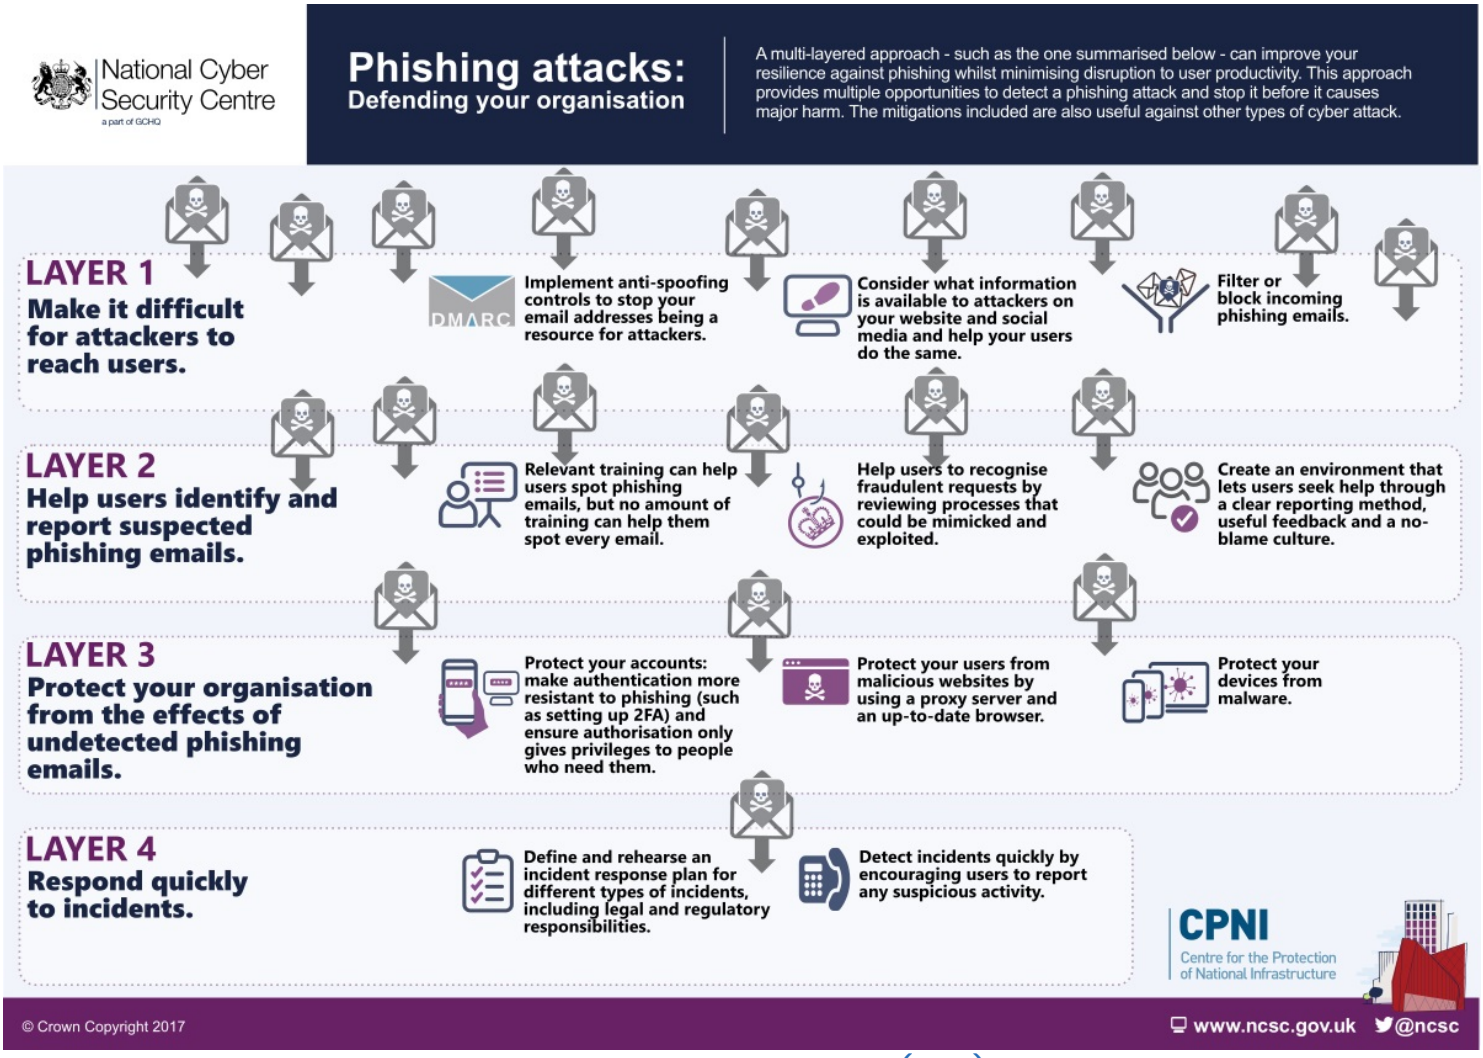

Download the phishing attacks infographic below (pdf)

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 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?

- 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. CPNI's Digital Footprint Campaign contains a range of useful materials (including posters and booklets).
- Be aware of what your partners, contractors and suppliers give away about your organisation online.

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. Your filtering/blocking service might be a cloud-based email provider's built-in service, or a bespoke service for your own email server.

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).
- 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, and to respond to your business's changing needs and ways of working.
- 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

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 defence. Your users cannot compensate for cyber security weaknesses elsewhere. Responding to emails and clicking on links is a huge part of the modern workplace, so it's unrealistic to expect users to remain vigilant all the time.

Spotting phishing emails is hard, and spear phishing is even harder to detect. Even experts from the NCSC struggle. The advice given in many training packages, based on standard warnings and signs, will help your users spot some phishing emails, but they cannot teach everyone to spot all phishing emails.

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. Never punish users who are struggling to recognise phishing emails; it's a bad idea for many reasons. Users who fear reprisals will not will not report mistakes promptly, if at all.
- Training should encourage your users' willingness to report future incidents, and re-assure them that it is OK to ask for further support when something looks suspicious. This message needs buy-in across all departments including HR, support and senior management.
- 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 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 provides a range of materials to deliver security messages on this topic.
- Using phishing simulations will not make your organisation more secure. Some companies have user training that gets the participants to craft their own phishing email, giving them a much richer view of the techniques used. Others are experimenting with workshops, quizzes and gamification, making a friendly competition between peers (rather than an 'us vs them' situation with security).

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 how to review and improve them so phishing attacks are easier to spot.

In addition, think about how the emails you send to suppliers and customers will be received. Can your recipients easily distinguish your genuine email from a phishing attack? After all, their users (like yours) cannot be expected to look for and recognise every sign of phishing. Don't assume providing personal information will verify your identity; stolen or researched information is used by phishers to make their emails more convincing.

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. 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 attempts (including ones that that are 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 they think a phishing attempt 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?
- 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 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 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. 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) 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, see our security 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.
- Read the NCSC's End user device (EUD) security guidance.

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, 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 Two Factor Authentication (2FA), which is also called 'Two Step Verification' 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 login 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 the NCSC'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 Incident Management section of 10 Steps to Cyber Security.
- 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 have created Exercise In A Box, an online tool which helps you to find out how resilient you are to cyber attacks, and where you can practise your response in a safe environment.

Case study: how multi-layered phishing mitigations defended against Dridex malware

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, or in the case of relying on cleaning up quickly afterwards, be very costly and prohibitively time consuming.

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.

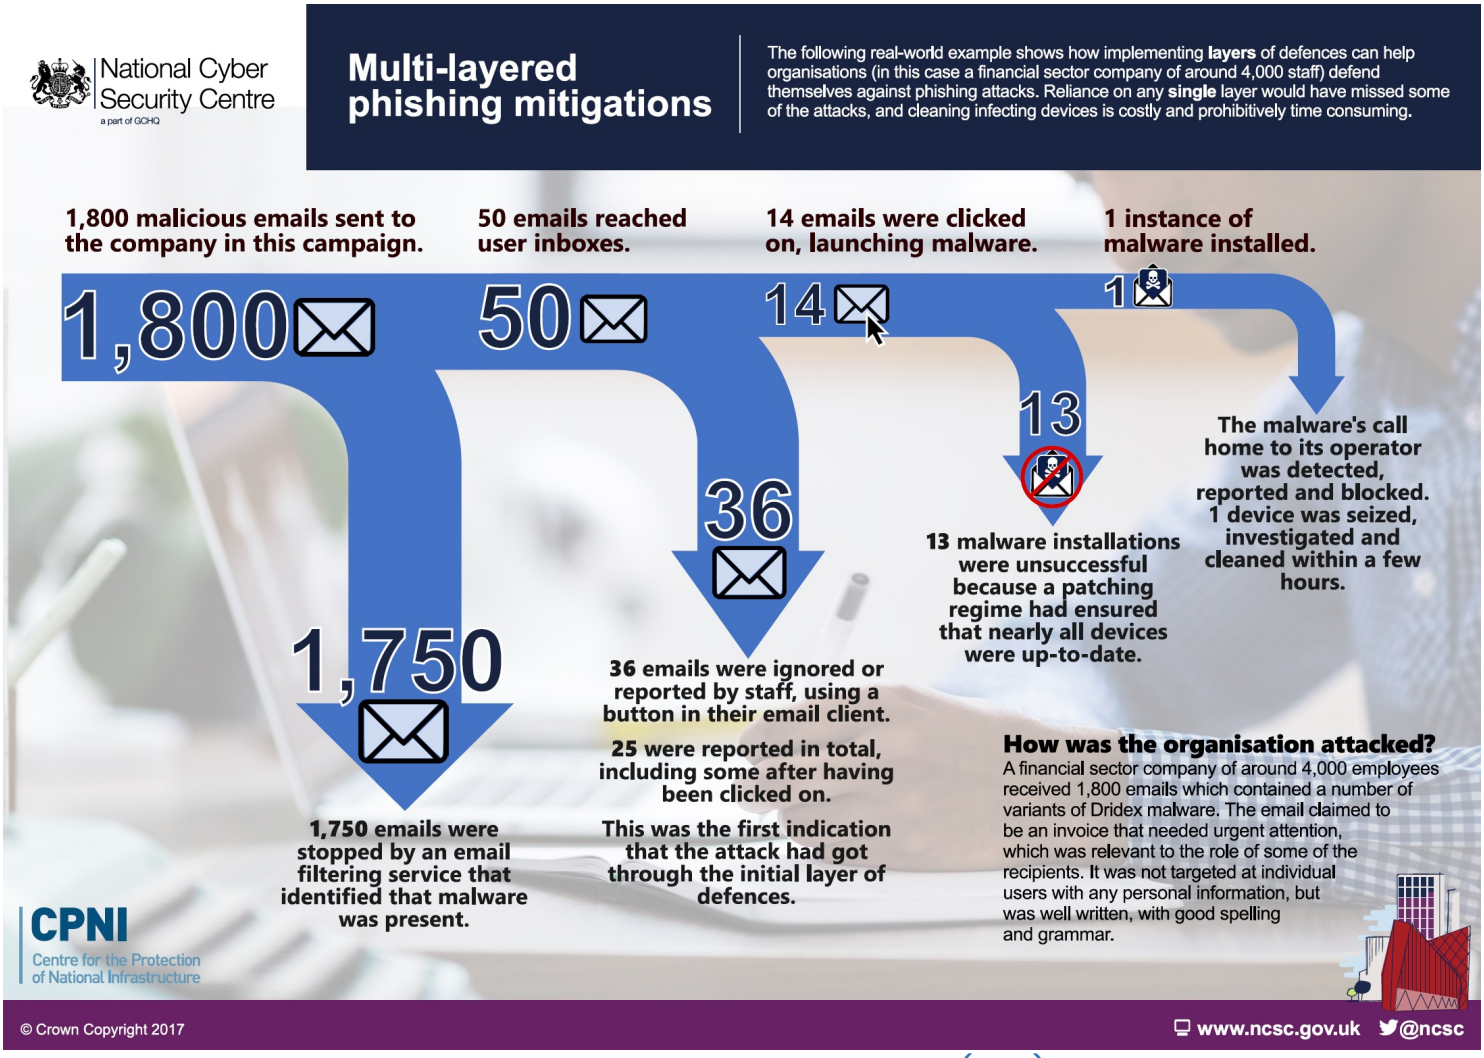

Download the phishing mitigations case study below (pdf)

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.

Video summary: defending your organisation from email phishing attacks

PUBLISHED

5 February 2018

REVIEWED

8 August 2019

VERSION

1.1

WRITTEN FOR

Small & medium sized organisations

Large organisations

Public sector

Cyber security professionals