Active Cyber Defence - One Year On

Dr. Ian Levy, Technical Director
UK National Cyber Security Centre

5th February, 2018

Executive Summary

In November 2016, the government launched the new National Cyber Security Strategy. This strategy had many effects, including setting up the National Cyber Security Centre (NCSC) as part of GCHQ and giving it a mandate to pursue the radical action required to better protect the UK's interests in cyberspace. A key strand in this new approach is the NCSC's Active Cyber Defence (ACD) programme, which aspires to protect the majority of people in the UK from the majority of the harm, caused by the majority of the attacks, for the majority of the time. It is intended to tackle the high-volume commodity attacks that affect people's everyday lives, rather than the highly sophisticated and targeted attacks, which we deal with in other ways.

The NCSC was founded on principles of evidence and transparency and one purpose of this paper is to make good on that commitment by cataloguing some of the work we have done in 2017 under the ACD programme, publishing data that has been generated by the work, trying to draw outcomes from it and making all of that available for scrutiny and challenge.

More importantly, it is a call to action for UK public sector organisations, UK industry and our international partners to implement these or similar measures so that collectively we make cyber crime less profitable and more risky globally. We believe that the data in this paper, and the conclusions that can be drawn from it, show the effect of our technical interventions and demonstrate that the basic defences we have been calling for can be deployed at scale.

The ACD programme consists of a number of interventions or real services - each free at the point of use - that perform a particular security service for public sector organisations. For each of the services in general use today, we give headline security outcomes below.

- Takedown Service
  This service works by requesting that hosting providers remove malicious content that is pretending to be related to UK government and also certain types of malicious content hosted in the UK.
  - In 2017, we removed 18,067 unique phishing sites across 2,929 attack groups that pretended to be a UK government brand, wherever in the world they were hosted. As a consequence, we have reduced the median availability of a UK government-related phishing site from 42 hours to 10 hours. That means that these sites are available for much less time to do harm to UK citizens. 65.8% of those are down in 24 hours, up from 39% before we started takedowns.
  - In 2017, we removed 121,479 unique phishing sites across 20,763 attack groups physically hosted in the UK, regardless of who it was pretending to be. As a consequence, we have reduced the median availability of a phishing site physically hosted in the UK from 26 hours to 3 hours, again giving them much less time to do harm. 76.8% of those were down in 24 hours, up from 47.3% before we started takedowns.
  - In 2017, we worked with 1,719 compromised sites in the UK that were being used to host 5,111 attacks, intended to compromise the people that visited them. As a consequence, we have reduced the median availability of these compromises from 525 hours to 39 hours.
  - Over the year 2017, the month-by-month volume of each of these has fallen, suggesting that criminals are using the UK government brand less and hosting fewer of their malicious sites in UK infrastructure.
  - In 2017, we notified email providers about 3,243 Advance Fee Fraud attacks, pretending to be related to UK government.
  - In 2017, we have stopped several thousand mail servers being used to impersonate government domains and sending malware to people, in the expectation that the

  government link makes them more realistic. We have also removed a number of deceptive domains that were registered with the sole intention of deceiving people.
  - While the volume of global phishing we can see has gone up significantly (nearly 50%) over the last 18 months, the share hosted in the UK has reduced from 5.5% to 2.9%.

- DMARC
  DMARC helps email domain owners to control how their email is processed, making it harder for criminals to spoof messages to appear as though they come from a trusted address. Organisations that deploy DMARC properly can ensure that their addresses are not successfully used by criminals as part of their campaigns. We have committed to helping the public sector lead in deploying DMARC.
  - We are prioritising 5,322 government domains for adoption in the first instance. Those with SPF policies (a prerequisite of full DMARC) has risen from 26.85% at the start of 2017 to 38.56%. DMARC adoption is up from 5.58% at the start of 2017 to 18.3%. We expect that to accelerate over the coming months as we demonstrate the benefit, as outlined in this paper.
  - At the end of 2017, we have 555 (about 10%) government domains reporting to our Mail Check service.
  - We have seen the number of messages spoofed from an *@gov.uk* address (for example, taxrefund@gov.uk) fall consistently over 2017, suggesting that criminals are moving away from using them as fewer and fewer of them are delivered to end users.
  - Across the 555 public sector email domains reporting to Mail Check, we are seeing an average of 44.1 million messages a month which fail verification, with a peak of 78.8 million in June. Of those, an average of 4.5 million are not delivered to the end users. The peak in June saw 30.3 million spoofed messages not delivered to end users. Deploying DMARC can be hard for a complex enterprise and it is often the case that a few months of monitoring is needed - when spoofed messages are reported but still delivered - before setting the policy to reject spoofed emails.
  - The coming months will see a push to have more public sector bodies to set their domain policies so that spoofed emails are rejected by receivers. Our Mail Check platform is critical to providing the data to help that happen.

- Web Check
  Web Check performs some simple tests on public sector websites to find security issues. It provides clear and friendly reporting to the service owners, along with advice on how to fix the problems.
  - Between 18th April 2017 and 31st December 2017, Web Check performed 1,033,250 individual scans running 7,181,464 individual tests.
  - In that period, we scanned 7,791 unique URLs across 6,910 unique domains ingesting a total of 7,748 unique pages.
  - In that period, we produced 4,108 advisories for customers, covering a total of 6,218 different issues.
  - In that period, we found 2,178 issues relating to certificate management, 1 relating to HTTP implementation, 184 relating to out of date content management systems, 1,629 relating to TLS implementation, 76 relating to out of date server software and 40 other issues.
  - Most issues were fixed by the service owner within 2 days of being reported.

- Public Sector DNS
  The Public Sector DNS service provides protective DNS services to public sector bodies that subscribe to it. It blocks access to known bad domains, where the block lists are

derived from a combination of commercial, open source and NCSC threat feeds. It also performs analytics on the resolution data to find other security issues. The intent of the service is not just to block bad things, but to notify system owners so they can perform remediation.

  - At it peak in December 2017, the public sector DNS services was responding to 1.23 billion requests a week.
  - During that peak week, 273,329 requests were blocked, of which 5,768 were unique.
  - During 2017, over 3 terabytes of DNS data has been analysed for security threats.
  - During 2017, 134,825 unique DNS queries were blocked.
  - Nearly all organisations have benefited from the blocking of DNS queries and, on average, 1 in 6 organisations joining the service have some security issue identified that requires further remediation.
  - The domain generation algorithm detection analytic has found traffic linked to malware in 9 organisations. The malware families involved were Wannacry, BadRabbit, Ramnit and Conficker. Traffic for these was handled appropriately.
  - The domain generation algorithm detection analytic has also found unknown DGA-like activity in some customers. The investigation into this continues.

- Signalling and Routing
  We describe the work we are doing to make both source and destination address spoofing in IP space much harder and the consequent impact this could have on using UK infrastructure as part of a DDoS attack and traffic hijacking. We also describe the work on securing SS7 which intends to make abuse of UK mobile networks harder and some early (but successful) experiments into tackling SMS spoofing.
- We also describe our integration platform, the Threat-o-Matic, that links all the Active Cyber Defence measures and the early experiments we have done with others to prove event sharing and the benefits it could bring.

We committed to providing evidence that the measures we are taking are having a positive effect on the security of the UK. The work we have done already seems to show that these measures have a security benefit, but we need to incentivise others to do similar things in order to scale the benefits to protect the UK from the harm caused by commodity cyber attack in a measurable way. This paper is intended to provide (at least the start of) an evidence base to incentivise others - both industry and other governments worldwide - to take similar measures to help reduce the harm caused by commodity cyber attack globally.

We do not claim that what is presented here is sufficient or optimal, but it is a set of measures that provide objective benefit in a measurable way. We hope that this is the start of a change in how cyber security is handled at national scale in the UK and shows the benefit of the government's approach as codified in the national strategy.

Contents

List of Figures

List of Tables

1 Introduction

In 2015, the National Security Strategy confirmed that cyber remained a top threat to the UK's economic and national security. As a result, in November 2016, the Chancellor of the Exchequer launched the UK's new five year National Cyber Security Strategy, detailed here [1]. It set out ambitious policies to defend our people, deter our adversaries and develop our skills and capabilities, with a vision that by 2021 the UK is secure and resilient to cyber threats, prosperous and confident in the digital world.

This strategy was intended to be a significant departure from the UK's previous national strategy in this area and, by extension, significantly different to any other national strategy around cyber security. The creation of the National Cyber Security Centre (NCSC) recognised the need to simplify the landscape to make the UK's response more efficient and effective. It also demonstrated the commitment set out in the national strategy for the government to be more interventionist in the security of the UK, but recognising that we cannot do this alone. It is this partnership ethos that is at the heart of the NCSC and the Active Cyber Defence programme.

The NCSC has repeatedly stated that we want cyber security to become a true science, with public policy decisions (and enterprise and personal security decisions) based on high quality data and analysis, rather than hyperbole and fear. One year on, we've published this paper to start that process of transparency.

This paper and the work behind it doesn't represent the totality of the work the NCSC is doing to protect the UK. There are many other strands of work that we are driving to protect the UK. We reduce harm every day by helping the victims of cyber attack mitigate and recover through our incident management function. We provide actionable and pragmatic advice to end users and enterprises to help them be more secure, for example our world-leading password guidance. We track, demotivate and disrupt sophisticated adversaries using both our unique national assets as part of GCHQ and by working with our industry partners around the world. Much of this other work is difficult to talk about in public in great detail. The Active Cyber Defence programme is intended to mitigate the high volume cyber attacks we see which cause significant harm and also be easy to talk about in public in detail to help demystify cyber security.

We welcome feedback on the programme and our analysis by email at acd-feedback@ncsc.gov.uk.

1.1 The Active Cyber Defence programme

When the Chancellor launched the national strategy, the NCSC published a high level description of the Active Cyber Defence programme here [2]. The intent for this programme is to protect the majority of people in the UK from the majority of the harm caused by the majority of the cyber attacks the majority of the time. The advice that the cyber security community has given to the public has, in some cases, been useless - for example around 'trusting' emails. The ACD programme intends to just stop the majority of the attacks ever reaching end users, and mitigating the harm caused by those that remain, making their lives easier and hopefully making the more sophisticated attacks more obvious to users. We also expect that the programme will have an effect on some of the misaligned incentives in the cyber security arena. We want to make sure that attacks and adversaries are talked about honestly, reserving the descriptor 'advanced' for those that really are and being honest about the more commodity attacks.

The ACD programme is not intended to be perfect and it's not intended to deal with highly targeted attacks undertaken by the most sophisticated actors. It is intended to make the UK an unattractive target to cyber criminals and some nation states by increasing their risk and

[1] https://www.gov.uk/government/publications/national-cyber-security-strategy-2016-to-2021

[2] https://www.ncsc.gov.uk/blog-post/active-cyber-defence-tackling-cyber-attacks-uk

reducing their return on investment. It is not intended to imply retaliation ('hack back') by victims or militarisation of the internet - in this case 'active' means getting off our backside and doing something, rather than any of the more esoteric definitions. It is intended to automate protection at national scale for a good proportion of the commodity attacks we see, leaving the skilled network defenders across the UK to deal with the more sophisticated attacks that we cannot currently protect against automatically.

The ACD programme isn't static. Some of the services will cause attackers to change their behaviour and we will need to adapt. We are continually having new ideas about how to better protect the UK, but we do not have the monopoly on good ideas and welcome ideas from others on what else we should try. Those potential services will need to be looked at, tested, prototyped and possibly put into production as the programme moves forward. Not every intervention will work, but we'll be honest about that as well as our successes.

1.2 The adversary

There is much hyperbole about the capabilities of cyber actors. Certainly, some nation states invest huge sums of money and significant highly skilled resources in their cyber programmes and use those for various things that are detrimental to the interests of the UK. However, the vast majority of people in the UK will not be directly harmed by these actors. They are much more likely to fall victim to cyber crime, whether directly by being targeted or indirectly through one of their service providers being compromised. In order to help explain what we are trying to achieve, we have published a description of the cyber crime ecosystem here [3]. This ecosystem is complex and there are many parties involved and - for a significant part of this problem - profit margins are small. It is easy to see from figure 1 that there are many places that an intervention could cause an effect on the attackers' profits and therefore return on investment and therefore motivation.

You don't need to eradicate cyber crime and it would be unrealistic to think we could. But we do want to make it as hard as possible and that means making it as unprofitable and risky as we can for cyber criminals to act in the UK. That's not to say that UK people won't continue to be part of general campaigns, but those campaigns should be less likely to target UK brands and infrastructure, and therefore less likely to be effective against UK citizens. What follows is an overview of the things we've done in the first year of the programme to try to achieve this.

1.3 A word about data

The measures we'll talk about in this paper have generated lots of data, of various scales and quality. We haven't yet had the time to conduct thorough analyses of these datasets to the quality we would want - in the opinion of the author, it is more important to put something out quickly to show our intent to come good on our promise of transparency. It would be very easy to draw conclusions from some of the graphs and data we present here. What we have learned is that it's difficult to draw concrete conclusions - especially about causality - from our current analysis of the data. There are also some anomalies in the data that we don't understand yet. We've tried our best to be clear about our confidence in our conclusions in this paper. People will almost certainly disagree with some of the conclusions we draw here. That's probably a good thing as it starts to engender an evidence-based discussion about what cyber security policy should look like going forward. By all means get in touch to discuss if you think we've missed something or you disagree with our analysis. We'll be doing a more thorough analysis over the coming months which we'll publish as it becomes available.

[3] https://www.ncsc.gov.uk/news/ncsc-publishes-new-report-criminal-online-activity

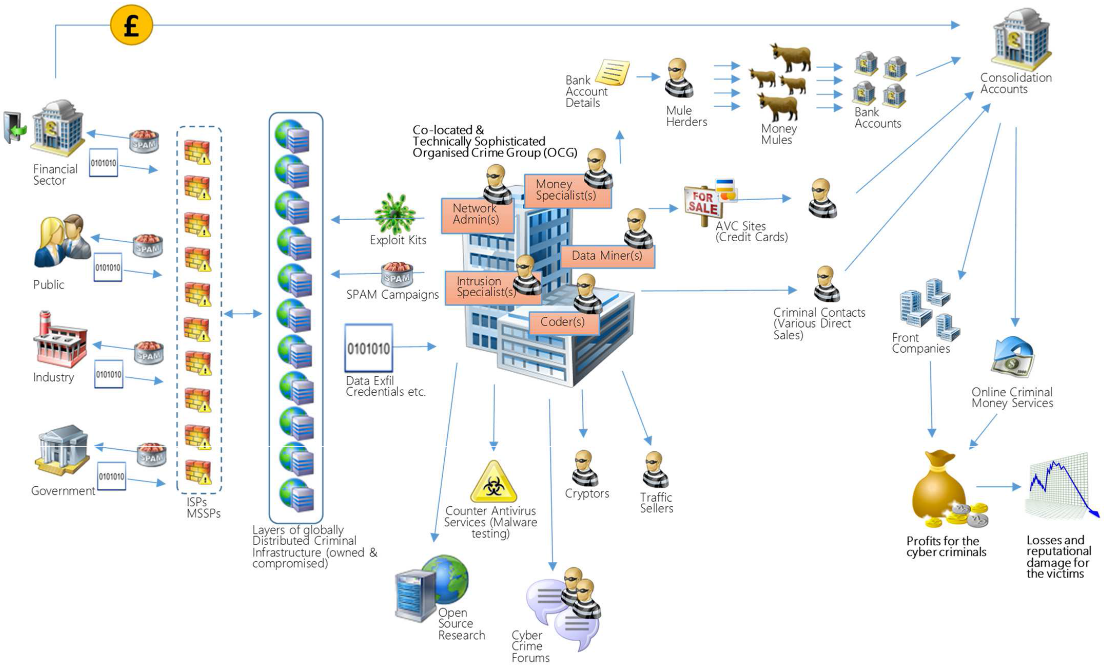
£
Financial Sector
Public
Industry
Government
SPAM
0101010
ISPs
MSSPs
Layers of globally Distributed Criminal Infrastructure (owned & compromised)
Exploit Kits
SPAM Campaigns
0101010
Data Exfil Credentials etc.
Co-located & Technically Sophisticated Organised Crime Group (OCG)
Network Admin(s)
Money Specialist(s)
Data Miner(s)
Intrusion Specialist(s)
Coder(s)
Bank Account Details
Mule Herders
Money Mules
Bank Accounts
Consolidation Accounts
FOR SALE
AVC Sites (Credit Cards)
Criminal Contacts (Various Direct Sales)
Front Companies
Online Criminal Money Services
Profits for the cyber criminals
Losses and reputational damage for the victims
Counter Antivirus Services (Malware testing)
Cryptors
Traffic Sellers
Open Source Research
Cyber Crime Forums

Figure 1: The Cybercrime Ecosystem

One of the things we should be very clear about up front is the potential for selection bias in the data. Bluntly, we can only see what we can see and our measurements are skewed by that. We think that the sources we're using are, on average, of the right character for us to be able to make the statements we do, but there is definitely bias in some of the data. For example, we don't know exactly what proportion of global phishing that is actually sent is in our phishing feed. So, while we can be confident about the effect on the data we can see, we don't have 100% confidence that we're seeing everything. However, using companies who have a long track record in the field mitigates that risk to some degree. We'd like to do work over the next year to better characterise our biases and ensure we can give more confidence in the conclusions we draw.

2 Takedowns

We have known for a while that HMG-related brands are often used as a lure for the various types of commodity cyber attack and we thought that we could do something about this, reducing the usefulness of those brands to criminals so they were not first choice. We also wondered if we could ‘clean up’ the UK’s IP space, reducing the amount of malicious content hosted in the UK. We know that malicious content all over the world affects people in the UK, but could the UK be a proving ground for some more systemic mitigations?

The process of taking down sites by asking hosting providers to do something, and further mitigating harm by adding sites to the various safe browsing lists so that modern browsers won’t access those sites, is well understood, but really based on trust. We could have built a capability from scratch and started to build that trust with each hosting provider, but there are companies who are expert in this already. We are currently working with Netcraft, a small UK company, with a strong track record of performing this sort of work for corporate brands. For clarity, the hosting company is told about a site that is impersonating another brand for the purposes of causing harm (rather than a protest site, for example). If the hosting company agrees that content is contrary to its terms and conditions, it removes the content.

2.1 Government-related abuse

2.1.1 Phishing

Let’s start with the simple case. We get spam and phishing feeds, look for phishing sites that are pretending to be a UK government brand and ask the hosting provider nicely to take the site down. This isn’t rocket science - lots of big brands do this sort of thing. While some government departments do their own brand protection, most don’t and it is simpler and cheaper for us to do this centrally. No-one but HMG is going to protect HMG brands.

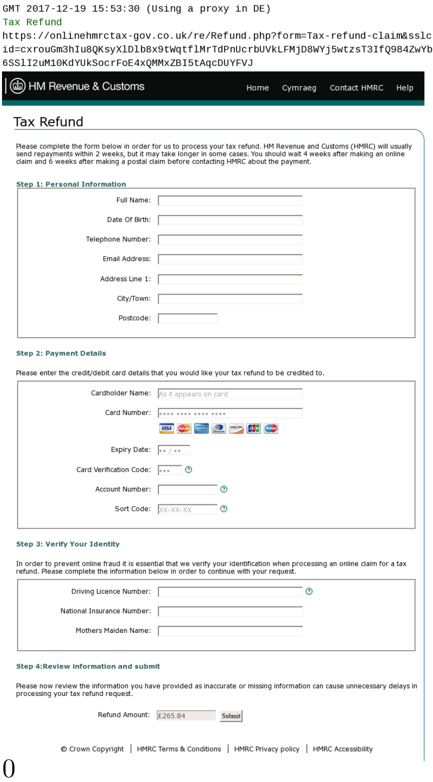

```
GMT 2017-12-19 15:53:30 (Using a proxy in DE)
Tax Refund
https://onlinehmrctax-gov.co.uk/re/Refund.php?form=Tax-refund-claim&sslchannel=true&session
id=cxrouGm3hIu8QKsyXlDlb8x9tWqtflMrTdPnUcrbUVkLFMjD8WYj5wtzsT3IfQ984ZwYbY96TuSzgcdo9bIfHcOa
6SSlI2uM10KdYUkSocrFoE4xQMMxZBI5tAqcDUYFVJ
```

Tax Refund

Please complete the form below in order for us to process your tax refund. HM Revenue and Customs (HMRC) will usually send repayments within 2 weeks, but it may take longer in some cases. You should wait 4 weeks after making an online claim and 6 weeks after making a postal claim before contacting HMRC about the payment.

Step 1: Personal Information

Full Name:
Date Of Birth:
Telephone Number:
Email Address:
Address Line 1:
City/Town:
Postcode:

Step 2: Payment Details

Please enter the credit/debit card details that you would like your tax refund to be credited to.

Cardholder Name: As it appears on card
Card Number: •••• •••• •••• ••••
Expiry Date: •• / ••
Card Verification Code: •••
Account Number:
Sort Code: XX-XX-XX

Step 3: Verify Your Identity

In order to prevent online fraud it is essential that we verify your identification when processing an online claim for a tax refund. Please complete the information below in order to continue with your request.

Driving Licence Number:
National Insurance Number:
Mothers Maiden Name:

Step 4:Review information and submit

Please now review the information you have provided as inaccurate or missing information can cause unnecessary delays in processing your tax refund request.

Refund Amount: £265.84 Submit

© Crown Copyright | HMRC Terms & Conditions | HMRC Privacy policy | HMRC Accessibility

Figure 2: An example of a phishing site impersonating HMRC

For the purposes of this, let’s define phishing to be impersonation of a UK government brand (things like HMRC or Student Loans Company) for the purposes of duping someone into disclosing personal details or passwords. Figure 2 is a screenshot of a typical phishing site for context. You can see that it asks for information that really wouldn’t be needed by the real tax system, but is obviously very useful to criminals. Also notice the domain name that’s been used *onlinehmrctax − gov.co.uk*. That’s intended to deceive the user into thinking this is a real HMRC site. Not all phishing sites use domains like this and many are hosted in areas of legitimate sites that have been compromised by the criminal.

Phishing sites are also automatically added to a number of industry safe browsing lists that are consumed by the major browsers and so even if the hosting provider doesn’t respond, or it takes long time for the site to be removed, users of modern browsers with the default security settings are protected anyway[4].

We define the *availability* of an attack as the total amount of time the phishing site is available from when the Netcraft service first becomes aware of the attack through to when it is finally taken down. This accounts for the times when an attack is reinstated by the criminal after first being taken down by the provider, which can happen multiple times in some cases. It is also often the case that a single attack can involve multiple spoof sites, hosted on the same server. If there are many phishing URLs in a single attack, they can easily skew statistics through the responsiveness or otherwise of the hosting provider. Given a group of attacks are all hosted on the same 'server', we group these together taking the longest time any one of them is available as the availability for that group. This is intended to better reflect the realities of hosting provider takedowns.

Over the last calendar year, we've taken down 18,067 HMG-related phishing sites, distributed as per figure 3. For comparison, in the previous 6 months [5], the volume was 19,443 sites, also shown on the chart. It's clear that we have performed fewer HMG-related phishing takedowns in 2017 and the trend is generally downward. Given how the service is driven, it's reasonable to assume that it sees a relatively constant percentage of the global phishing and so this strongly suggests that there has been less HMG-related phishing this year than last. However, there is insufficient data and analysis available to prove this is caused directly by our intervention. However, it is very likely (in the opinion of the author) that this work has had a direct impact on the viability of criminal phishing targeting HMG brands, making them less lucrative and therefore less likely to be used.

The top ten brands abused as part of these attacks are shown in Table 1. In the first year of the service, we've taken a very wide view of what we mean by 'HMG-related' to help us understand the problem. It's obvious from the table that the vast majority of HMG-related phishing attacks continue to use the HMRC brand. That's unsurprising given that most adults have a relationship with them and everyone would welcome a tax refund.

Brand	No of phishing URL	No of attack groups	Median group Availabiity (hours)
HM Revenue & Customs	16,064	2,466	10
Gov.uk	1,541	241	15
TV Licensing	172	93	5
DVLA	107	53	11
Government Gateway	46	22	6
Crown Prosecution Service	43	26	15
A UK University	23	9	0.7
Student Loans Company	19	11	17
Student Finance Direct	13	3	3
British Broadcasting Corporation	8	7	35

Table 1: Top ten HMG-related brands used for phishing

The campaign related to the UK university is one of 13 that targeted UK universities. From what we can see, all of these were attempting to steal credentials for users' email accounts.

[4] You really should all be using a modern browser with the security features turned on.

[5] We only started the takedown service in June 2016.

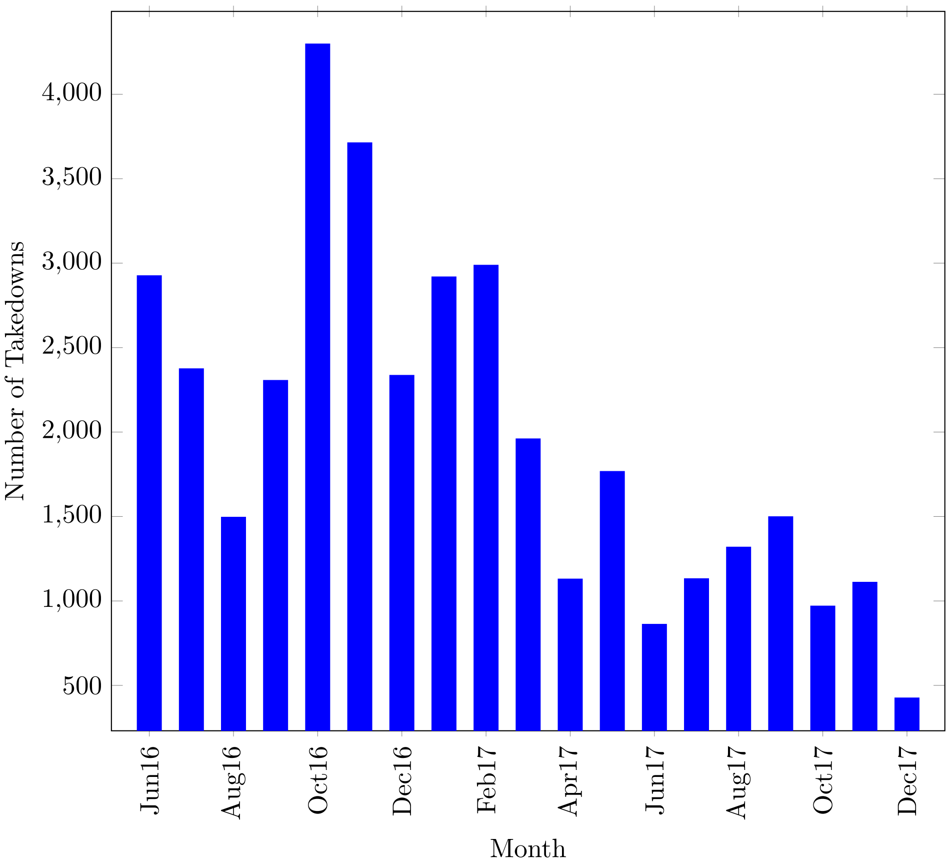

Figure 3: Volume of HMG-related phishing URLs removed

Ignoring the addition of these malicious sites to the various safe browsing lists [6], an important question is how quickly these sites are being removed from the internet.

In order to understand the effect that the service has on phishing sites, we ran the service for three months without contacting any hosting providers, in order to baseline how these sites are taken down without our intervention. The statistics, based on attack groups, for those three months compared with the statistics for 2017 where we were sending real takedown requests are shown in table 2. The datasets are plotted in figures 4a and 4b. The vertical line at the left of the graphs denotes 24 hours, and the vertical blue bars are 'chunks' of takedowns done within defined time periods. You can easily see that there's more blue bars on the left of the 24 hours line in the 2017 chart compared to the baseline one, showing that more of the attacks are being removed more quickly.

You can see from these numbers and the associated histograms that the majority of the phishing URLs seen by the service are being removed much more quickly in 2017 than they were before this work started. The median availability has dropped from 42.6 hours to 10 hours, and the percentage of sites unavailable within 24 hours has increased from 39.0% to 65.8%. While the majority of HMG-related phishing sites are removed quickly, as these statistics show, there is

[6] This is useful because not all browsers support safe browsing lists, not everyone has them turned on and many mobile browsers do malicious site protection differently due to bandwidth limitations.

Measure	Baseline Value	2017 Value
Mean	281.1 hours	104.8 hours
Median	42.6 hours	10.0 hours
Skewness	3.2	7.2
25th percentile	6.4 hours	1.7 hours
75th percentile	250.8 hours	42.9 hours
Sites down in 4 hours	21.4%	39.5%
Sites down in 24 hours	39.0%	65.8%

Table 2: Statistics of availability of HMG-related phishing sites, before and after intervention

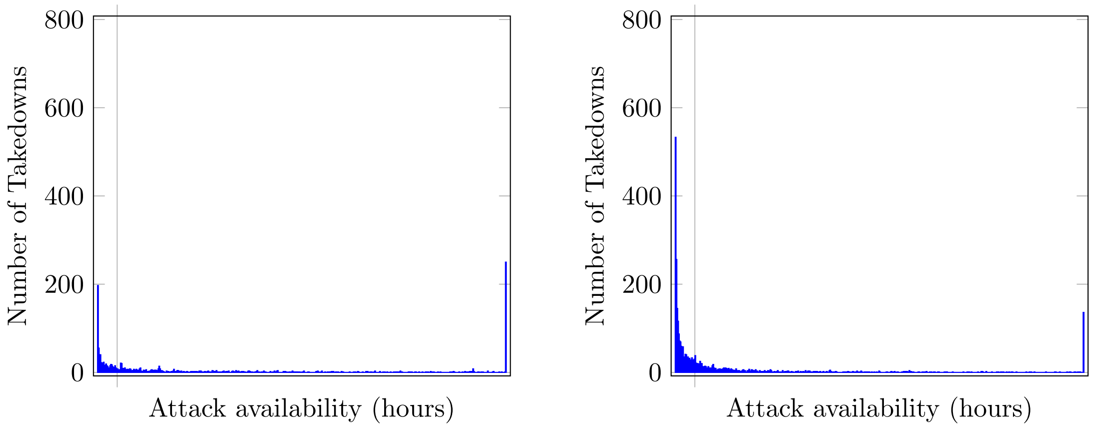

(a) Histogram of baseline availability for HMG-related phishing

(b) Histogram of availability in 2017 for HMG-related phishing

Figure 4: Comparison of availability of HMG-related phishing sites between baseline and 2017.(500 bins,$x_{max} = 500$, 24 hours noted by vertical line)

a recalcitrant tail in the distribution. You can see on the right of the histograms in figure 4 that there are a number of takedown requests that are taking a long time to be actioned - the last bucket contains anything that takes more than 500 hours to be removed. Figure 5 concentrates on the first 72 hours. You can see in the baseline graph that, before we started this work, most of the sites would take longer than 72 hours to be taken down (the far right vertical bar). In 2017, most of the sites are down in the first few hours. We're going to see if there's commonality among those long-lasting sites (including those that don't get removed at all), for example if they're all hosted on a specific set of hosting companies or in specific Autonomous Systems. Depending on what that shows, we may be able to take some other action to further mitigate the harm.

We've probably made HMG-related phishing sites go away a lot more quickly, but in order to really understand the protective effect of this work, we need to understand the percentage of people who clicked the link before the takedown occurred, who could have been harmed, and those who clicked the link after the takedown happened and so were protected. In other words, are the takedowns fast enough? Finding out when people click links has proved very, very difficult. We talked to ISPs because maybe their DNS would tell us, but it turns out not to be the case for a variety of reasons specific to this scenario. We can't consistently get data from the browser vendors, so we're a bit stuck. We haven't given up yet, but if anyone has a great idea about how to get this understanding, or someone wants to do the work for us, please get in

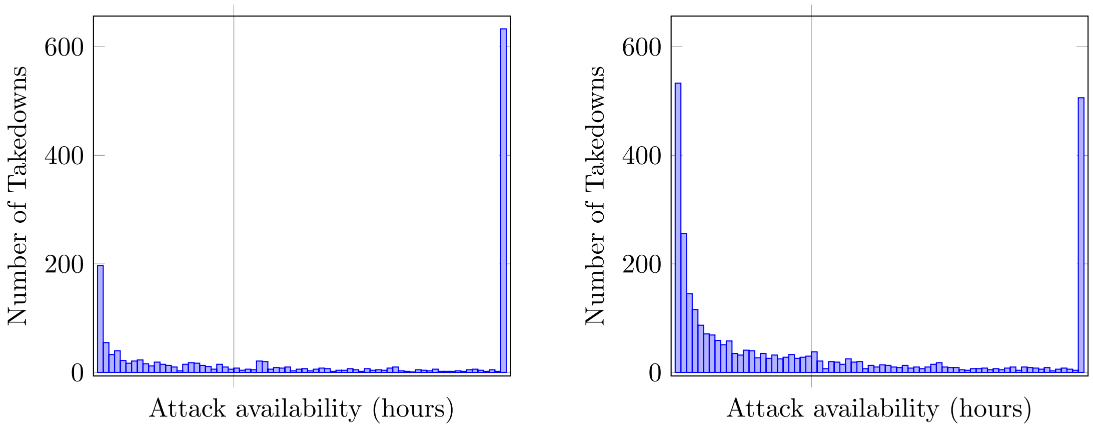

(a) Histogram of baseline availability for HMG-related phishing

(b) Histogram of availability in 2017 for HMG-related phishing

Figure 5: Comparison of availability of HMG-related phishing sites between baseline and 2017.(72 bins,$x_{max} = 72$, 24 hours noted by vertical line)

touch.

Table 3 shows the ten top companies that are used to host HMG branded phishing. In total, there were 587 different hosting companies used in 2017 to host HMG branded phishing attacks with these 10 accounting for 36.5% of all attacks. That's not necessarily surprising though - they are among the bigger hosting providers in the world and most offer a free tier of hosting.

Hosting Company	Percentage of hosting
Endurance International Group	9.3%
GoDaddy	6.9%
Amazon	4.3%
OVH	3.7%
United Internet	3.3%
Hetzner Online	2.1%
Dynamic Network Services	1.9%
Host Europe Group	1.8%
Cloudflare	1.6%
LiquidWeb	1.6%

Table 3: Top ten hosters of HMG-related phishing sites in 2017

2.1.2 Government sending malware

The feeds of spam and malicious emails we use to drive the phishing takedown work showed something else. Early in 2017, we noticed that criminals had started to send malware as attachments using 'from' addresses that were spoofed from both real and non-existent email addresses under gov.uk. In June, we contracted Netcraft to start to mitigate these attacks. Mails which purport to come from HMG addresses and contain a suspicious attachment are automatically processed and useful characteristics extracted. The most interesting characteristics are the mail server that's been used to send the attack emails and the infrastructure the malware in the attachment tries to talk to in order to get its instructions from the criminals.

For the command and control infrastructure, we treat them the same as phishing URLs and notify the hosting provider that they're hosting bad stuff. The hosting provider will, generally speaking, remove the malicious entities when they're told about it although, as with phishing URLs, different providers react very differently. The mail servers that are used to send these malicious emails are generally not infrastructure that is owned and operated by the criminals, but more often mail servers that are badly configured and are being abused by the criminals. In either case, we inform the provider hosting the mail server and ask them to mitigate the problem - either fix a badly configured mail server or remove a malicious one.

By doing these two actions, we can reduce the harm of the campaign that we've just seen, through removing the command and control infrastructure for the malware, and reduce the likelihood of the same mail infrastructure being abused again to send more malicious mail. This service has only been running since June 2017, so the data we have is relatively limited. However, the volume of mail server notifications sent is shown in Figure 6 and the malware infrastructure takedowns, relating to around 8,100 unique malicious binaries, in figure 7. The number of infrastructure takedowns in June is possibly anomalous due to a backlog; the August and September mail server numbers are real. We do not currently understand the criminal behaviour that drives these numbers, but as the service runs for more time, we should understand more about exactly what this tells us about how criminals operate in this space.

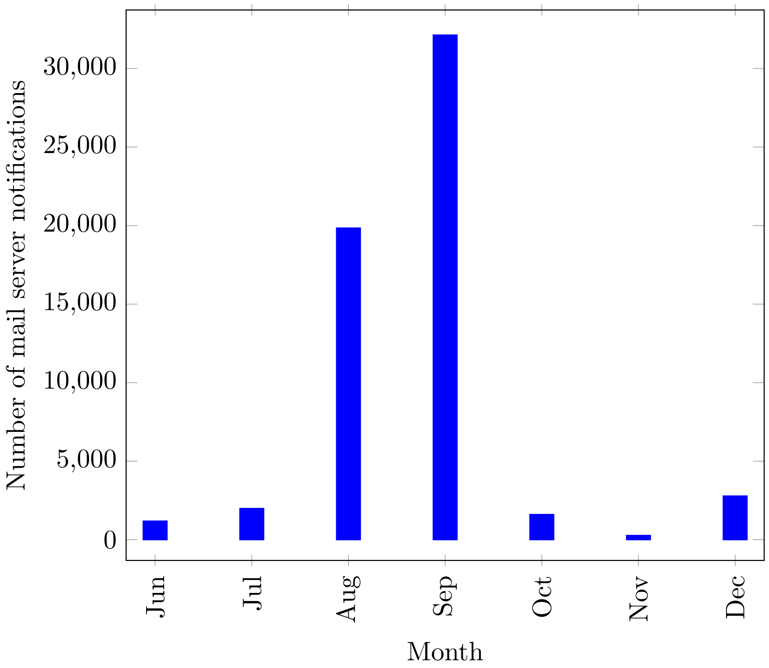

Figure 6: Volume of mail servers (unique IP address) notified as sending malware using an HMG from address

As an example, on 23rd August we saw a very large malware campaign, where the mails were either forging gov.uk addresses or contained text that suggested to the reader that the message was related to government. In the Netcraft mail feed, the campaign totalled around 109,000 messages, 552 distinct malicious attachments and 133 distinct forged gov.uk subdomains used as

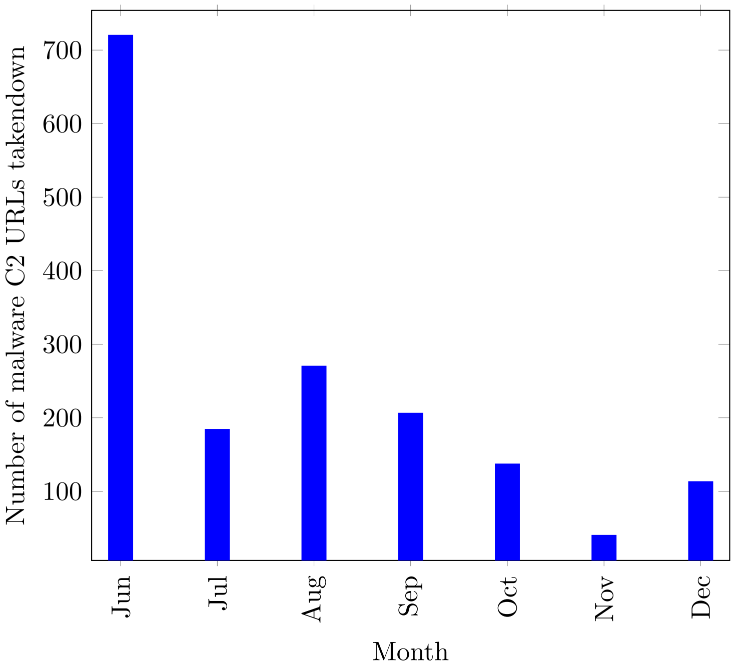

Figure 7: Volume of malware C2 URLs taken down

from addresses. The campaign was sent from around 20,000 distinct mail servers.

The malware was distributed as RAR archives, with names similar to Fax175343452255e8b.rar, attached to mails purportedly from a fax-to-email gateway service. The attached archive contains a Windows Script file with a filename similar to Fax5155564o9d964cf.js. When this script is extracted from the archive and run, it downloads an additional malicious executable payload from a remote server used to compromise a victim's machine. The malware itself is a new variant of the Locky ransomware known as Lukitus. However, despite the scale of the mail server infrastructure used, the variety of spoofed addresses used to give legitimacy to these mails and the large variety of malware used, the entire campaign was reliant on only 16 command and control URLs for download and control of the ransomware. These fell to takedown requests very quickly indeed, reducing the harm caused by the campaign. A campaign of similar scale was seen in September. It will be interesting to see what the future holds for criminals sending malware using gov.uk addresses. As our DMARC work continues (see section 3.2) this will have an effect on the ability of criminals to spoof gov.uk addresses. It's worth saying again that the big, international mail platforms do a relatively good job of not delivering this sort of email to people. However, not everyone uses those (relatively few) platforms and so this work provides harm reduction benefit to those who are not otherwise protected. Defence in depth is a good thing.

2.1.3 Advance fee fraud

Advance fee fraud is the generic name for the '419 scam' where someone pretends to want to give you a lot of money but need you to pay a relatively small processing fee up front. Below is an extract (text removed marked by $< snip >$) of a real advance fee fraud sent on 9th December 2017 with a reply-to address of $bankofengland.uk1@yahoo.com$.

FROM THE DESK OF CENTRAL BANK OF UNITED KINGDOM
FOREIGN SECRETARY OUTLINE THE GOVERNMENT
VISION FOR UK FOREIGN POLICY.
Bank of England Threadneedle Street London, EC2R 8AH

<snip>

Welcome To Bank of England Office London.
Batch Number: 573661545-UK/2017
Ref.Number: BOE 14-M-246-05

Payment Transfer File:/UK11007845/17.
Payment amount: ($2,500,000.00USD) TWO MILLION FIVE HUNDRED THOUSAND UNITED STATES DOLLARS.

BANK OF ENGLAND OFFICIAL PAYMENT NOTIFICATION.

Dear Winner.

The Foreign Exchange Transfer Department Bank of England (BOE)has decided to bring to your attention, that you were listed as a beneficiary in the recent schedule for payment of outstanding debts incurred by the BRITISH GOVERNMENT Pending since year 2016 to 2017 According to your file record with your email address, Your payment is categorized as: Contract type: Lottery unpaid contract funds/ Undelivered Lottery fund/

<snip>

Therefore, we are writing this email to inform you that ($2,500,000.00USD) will be release to you in your name, as it was committed for (BOE) Governor that Beneficiary will have to pay Fund Release Order Fee Charges ($153USD) only.

<snip>

Please you are advised to fill the form below and send it immediately to our transfer department for verification through email below for prompt collection Contact. Manager Name Joseph Pounch in charges of foreign exchange Transfer department
Contact E-mail:bankofengland.uk1@yahoo.com

Fill the Form Below and RETURN BACK IMMEDIATELY FOR YOUR PAYMENT:
1.Full Names:
2.Residential Address:
3.Mobile Number:
4.Occupation:
5.Sex:
6.Age:
7.Religion:

```
8.Nationality:
9.Name of your State:
10.Country:
11.Marital Status:
12.E-mail id:
13.Bank Name:
14.Account Number:
15.Account Holders Name:
16.Bank Branch:
17.Bank Address:
18.Ifsc Code:
19.Swift Code:
20.Next Of Kin
21.Mother Name:
22.Father Name:
23.Winning Amount:
24.Sacn Copy Of Your Id Proof:
25.Working Office Address:


NOTE: Fund release Order charges ($153USD) cannot be paid directly to Bank
of England, due to the fact that we do not Bank Publicly others Bank will
be assigned to receive Payment from you.
<snip>
DR MARK CARNEY.
GOVERNOR.
BANK OF ENGLAND
```

And yes, people really do fall for these scams.

Advance fee fraud is normally purported entirely on email and there are no phishing sites to take down. In this case, we notify the email provider of the fraud and ask them to delete the mail account. This stops future fraud taking place as a result of an existing campaign, but doesn't do anything for the people who've already responded to the fraudster. We're looking at whether there's a process we could create whereby law enforcement could get notification of people who've engaged with the criminal and proactively help them minimise the harm. That may or may not be possible, given all the different constraints in play.

In the year 2017, we have seen 3,243 advance fee fraud attacks across the HMG-related brands we're looking for [7], distributed as figure 8. Obviously, advance fee fraud campaigns aren't anything like as prevalent as phishing and malware email campaigns. The brands targeted are obvious [8]and the top three are shown in table 4. We cannot currently discern a pattern to the timing or targeting of these advance fee fraud attacks.

The distribution of email hosting providers is also interesting. There have been 120 different email hosting providers used by criminals for advance fee fraud in 2017. 61 of them only hosted one attack each during the year. The top 5 hosting providers are shown in table 5. As you'd expect, the big free email providers dominate this list, but they're also pretty good at taking down these fraudulent accounts when asked.

It's also clear that not all big email providers have good abuse notification and processing.

[7] We won't list them all for obvious reasons.

[8] They are where the money is.

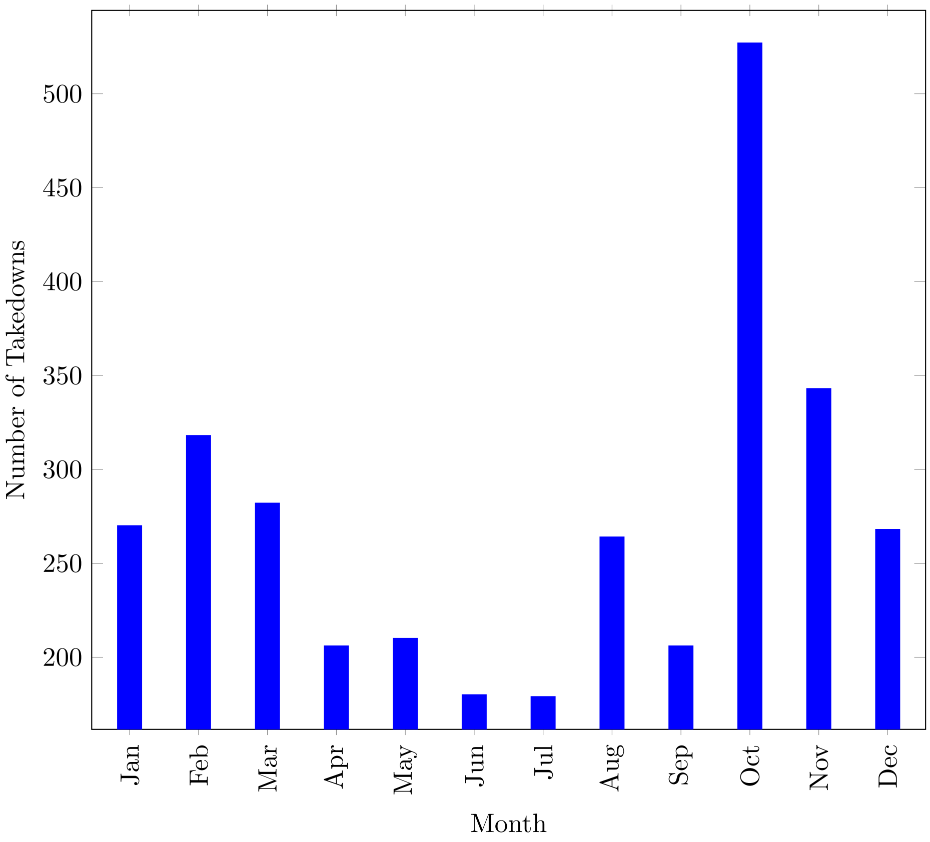

Figure 8: Volume of HMG-related advance fee fraud attacks

Brand	No. of AFF attacks
National Lottery	1,172
Financial Conduct Authority	743
Bank of England	73

Table 4: Top three HMG-related brands used for Advance Fee Fraud

Hosting Provider	No. of attacks hosted
Google	846
Yahoo	551
Verizon	371
Microsoft Corporation	367
1&1	320

Table 5: Top five providers hosting advance fee fraud email accounts

One large provider, which only hosted 8 attacks from February to November, only took action on 5th December, removing all eight malicious accounts. It's not clear what we should do about this, apart from calling out the companies who consistently fail to take fraud and security seriously.

Advance fee fraud is, generally speaking, well mitigated by the big email platform providers who generally direct these messages to the users' spam or junk folders. However, not all email providers or mail clients perform this analysis and so, while advance fee fraud is relatively limited in actual impact, this remains a useful mitigation for UK users.

2.1.4 Deceptive domains

We've seen an increase in people registering deceptive domains to try to make their criminal campaigns more effective. These domain names are intended to look like the real brand being used as the hook for the campaign and are often quite complex to further confuse attentive recipients who try to check the site address. The modus operandi at the moment appears to be for criminals to register the domains and host benign content on them for a few weeks. Only after that do they switch to hosting malicious content and then send out their email campaigns trying to drive people to the site. We've asked Netcraft to look through daily zone transfers (which they can get as a commercial company) for domains that look like they're trying to impersonate a government brand. If they find one, the site is visited to determine if it's malicious. If it is, a takedown is started immediately. If it isn't, the site is monitored frequently to see if the content has changed. If it has, a maliciousness determination is made again. As soon as malicious content is detected, we start a takedown request with the hosting provider.

Some examples of the domain names that have been used as part of phishing campaigns are shown in table 6.

nationalcrime-agency.com	onlinehmrc-gov.uk
hmrclogin.online	hm-revenue-customs-tax-refund-secure.se
hm-revenue-gov.uk	hm-revenue-customs-tax-refund-secure.top
dvla-refunded.com	hm-revenueandcustoms-gov.com
hmrc-gov-rebate.com	hmrc.today
hmrc.info.tm	rebates-hmrc-account.1gb.ru
hmrcapplication-gov.uk	hmrc-online-refund.com
hmrc-govuk-refund.su	rebates-dvla-tax.ml
pointshomeoffice-gov-uk-sms.com	hmrcgovrefunds.co.uk
tax-returngov.uk	refunds-dvla.co.uk
fceo-gov.uk	

Table 6: Examples of deceptive domains relating to HMG brands

We are also seeing a small number of SSL certificates associated with some of the deceptive domains. This is particularly pernicious as the cyber security community has trained the public to treat an SSL certificate (giving rise to the 'secure padlock' in the browser) as a designator of trust, even though it isn't. Average users will not reliably distinguish Domain Validation, Organisational Validation and Extended Validation certificates - which all try to say very different things about the site involved - and so this criminal method is likely to continue. The certificates we have seen appear to fall into two categories. Firstly, free SSL certificates for domains like the ones shown in table 6 from providers like 'Let's Encrypt!'. The second appear to be real sites that have been compromised and had deceptive subdomains added, of the form $hmrevenue - gov - taxrefundapplicationprocess.realdomain.co.uk$. All the instances we have seen in this second category so far have been certificates issued by the cPanel AutoSSL feature, making it likely that vulnerable cPanel installations are being targeted to some degree. These deceptive SSL certificates are currently low in number, but we will continue to monitor the situation.

2.2 Malicious hosting in UK IP space

We are also working to reduce the volume of malicious things hosted in the UK. More specifically, we're trying to reduce the longevity and incidence of phishing and web-injects hosted in IP space delegated to a UK-owned Autonomous System. The idea is that if we can show a real effect on the ability to host malicious content[9] in the UK then we can ask other governments why they're not doing the same sort of thing. Through this, we hope to be able to build the UK's reputation as a responsible hoster and, hopefully, reduce the ease by which malicious content can be hosted.

2.2.1 Phishing

We have asked Netcraft to follow exactly the same process as they do for phishing for HMG-related brands to find any phishing that they see that resolves to hosting in the UK. When they have such a resolution, they follow the exact same takedown process as detailed previously. For some brands, it is possible that the brand owner is also performing a similar process for their brand that will intersect with this work when UK hosting is involved. We don't think that is an issue. The worst case is that we assign to our work effects that the brand owner may have taken independently. The data that follows suggest that is not a significant concern.

Figure 9 shows the volume of phishing sites removed from UK IP space in 2017 and in the latter half of 2016 (when the service was also running). This shows that the number of phishing URL takedowns has risen over the period of the service - the monthly number of phishing URLs removed is higher in 2017 than it was in 2016. As stated previously, it's probably reasonable to assume that the service sees a relatively constant proportion of the total phishing that exists, which would then suggest that this trend implies that attacks hosted in UK IP space have increased. However, there is an alternative explanation for the trend.

We mentioned grouping of attacks in section 2.1.1. Consider phishing sites hosted at `www.yourbankhonest.com/currentaccount/secureupdate.php` and `www.your-bank-honest.com/savingsaccounts/transfer.php`. If notification of these occurs roughly contemporaneously, then it's reasonable to conclude that they are part of the same campaign and we group them for the purposes of statistics. Similarly, if we saw
`secure.update.hmrc.online.gov.uk.iansfreehosting.com` and `refund.dvla.online.gov.uk.iansfreehosting.com` and they resolved to the same IP address (knowing they're subdomains of the hosting provider `iansfreehosting.com`) then it's reasonable to assume these are also part of the same campaign and we group these for the purposes of statistics.

If we take that into account and graph the volume of attack *groups* over the last year, as in figure 10, we can immediately see the volume is much lower than in 2016. This implies that there are less phishing campaigns being hosted in the UK, but when they are the campaign seems to use many more related URLs.

Figure 11 shows the average number of URLs per attack group by month over the last year, comparing those hosted in the UK with all others. While there are trends in this data, there is insufficient data or context currently to draw any firm conclusions as to whether attackers behave differently when hosting their phishing sites in the UK.

Similar to the HMG-related phishing statistics, we show the pre-service baseline availability statistics and statistics for 2017 in table 7. It is obvious from these statistics that UK hosted phishing sites are much less long lived than before we started. There is an obvious harm reduction consequence for this.

[9] 'Malicious content' in the cyber security sense, not in the sense of porn, protest sites, objectionable content etc.

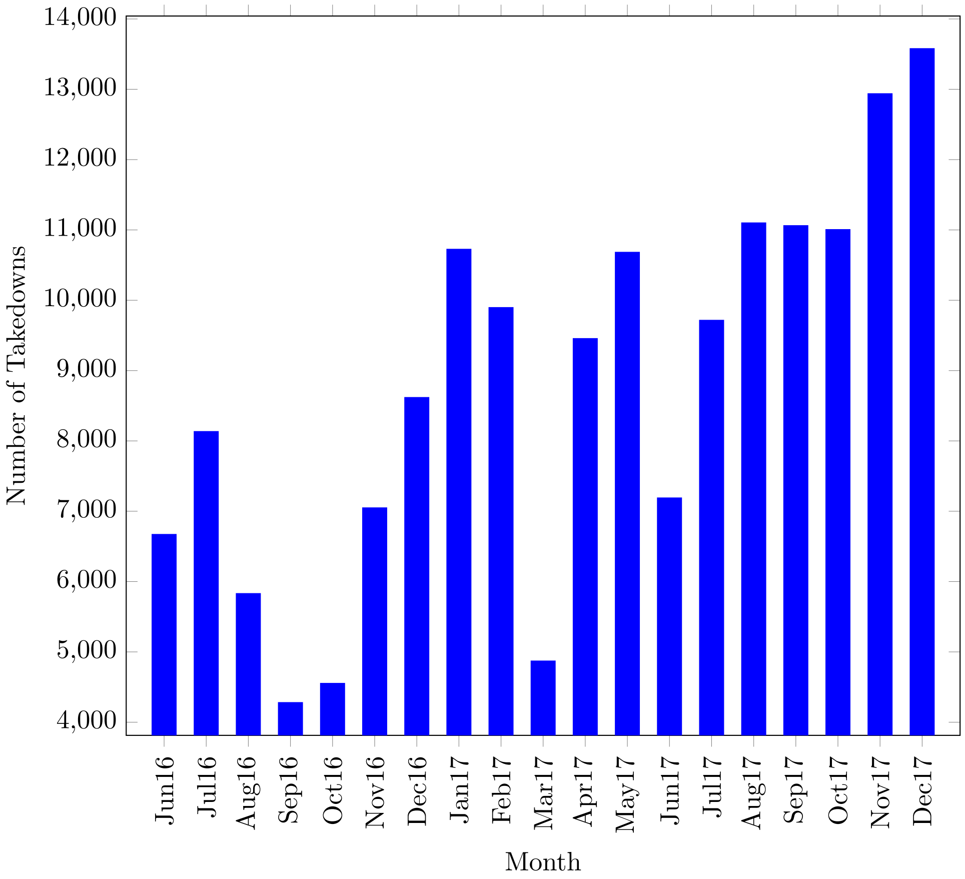

Figure 9: Volume of phishing URLs removed from UK IP Space

Measure	Baseline Value	2017 Value
Mean	254.8 hours	66.9 hours
Median	26.3 hours	3.5 hours
Skewness	6.9	10.2
25th percentile	3.2 hours	0.5 hours
75th percentile	173.6 hours	23.0 hours
Sites down in 4 hours	25.9%	56.7%
Sites down in 24 hours	47.3%	76.8%

Table 7: Statistics for baseline and 2017 availability of phishing sites in UK IP space

2.2.2 Web-injects

We also look for sites in the UK that are hosting web-injects. Normally these are real sites that have been compromised by the attacker due to some vulnerability and the web-inject is there to compromise visitors to that site. Obviously, you can't just take that site down as it's a real one - you've got to get it fixed instead which obviously takes longer. The volume of web-injects removed from UK IP space is shown in figure 12. It's worthy of note that, as far as we can tell,

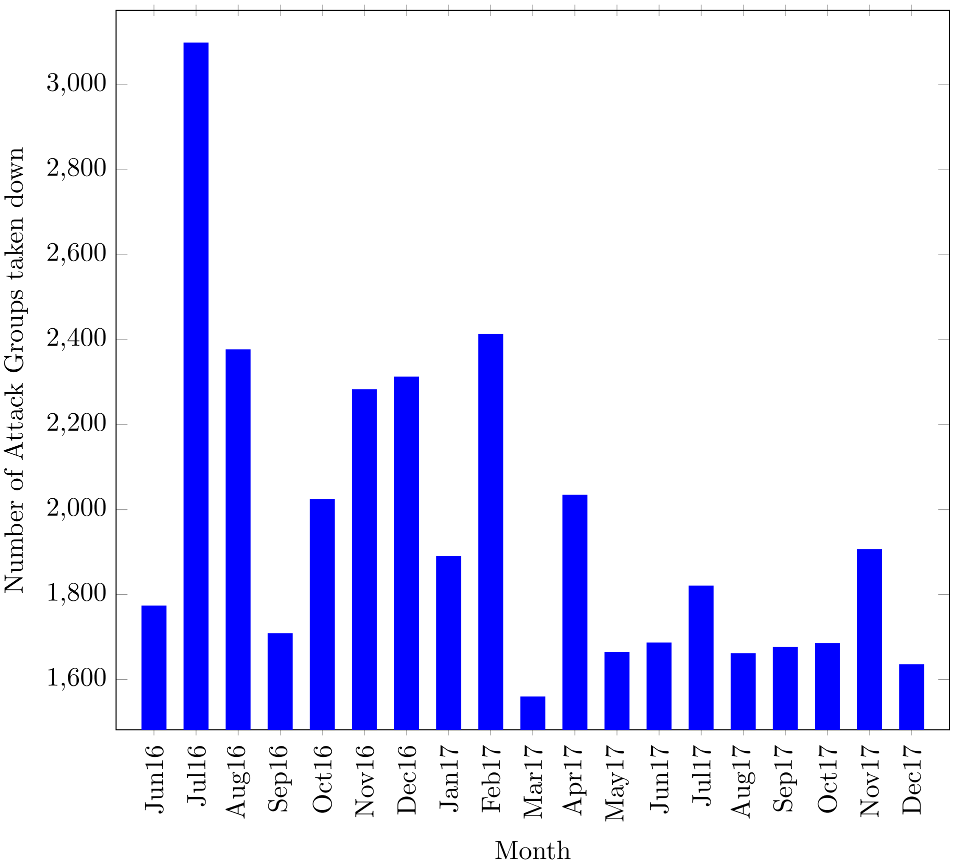

Figure 10: Volume of grouped phishing attacks removed from UK IP Space

the legitimate owners of these sites are not aware of the compromise and there's no concerted effort to fix them, apart from the work detailed here.

From this graph, it's pretty obvious that the volume of web-injects we've seen in the UK has dropped sharply since the service started. That could mean that sites hosted in the UK are 'safer' than those elsewhere, since criminals are seeing reduced payback for compromising UK sites (since they're around for less time) and may be compromising sites in other hosting areas. We don't have the data to make that statement definitively yet, but we'll try to get it over the coming months.

And for completeness, table 8 shows the statistics for the availability of web-inject compromises in UK IP space. Again, the availability of these sites is reduced from the baseline, suggesting that the service is having a real effect of the longevity of these compromises and therefore the harm they can cause.

2.3 Conclusion

The data we can garner from the takedown service strongly suggests that the activities are having a real harm reduction effect and are helping to protect users. HMG-related brands are generally going to target UK individuals and the potential harm caused by these campaigns has been reduced - disproportionately protecting UK users. The more general work to remove malicious

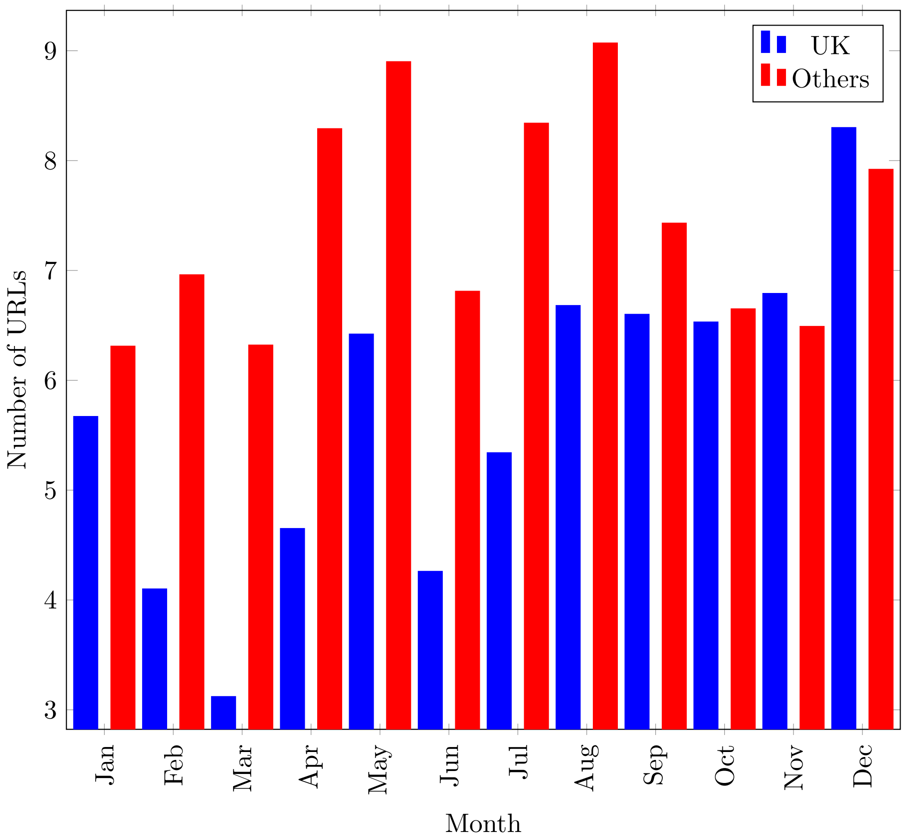

Figure 11: Average phishing URLs per attack group

Measure	Baseline Value	2017 Value
Mean	807.2 hours	513.4 hours
Median	525.1 hours	39.1 hours
Skewness	3.9	3.2
25th percentile	125.8 hours	3.55 hours
75th percentile	1,084.9 hours	340.3 hours
Sites down in 4 hours	3.4%	17.3%
Sites down in 24 hours	9.9%	33.8%

Table 8: Statistics for baseline and 2017 availability of web-inject sites in UK IP space

content in UK IP space shows that concerted effort can affect the ability of these campaigns to cause harm.

According to Netcraft's data, phishing has risen sharply in the last 18 months, since we started the service. We define a measure of *unique phishing attack* as the existence of a phishing target (e.g. HMRC, Student Loans) on a given IP address in a given month. In this case, 10 HMRC-related phishing attacks on the same IP address in June would be counted once, but if the same IP address hosted both an HMRC and a Student Loans phishing attack during June, that would be two unique phishing attacks. Note that this way of counting is different to the others

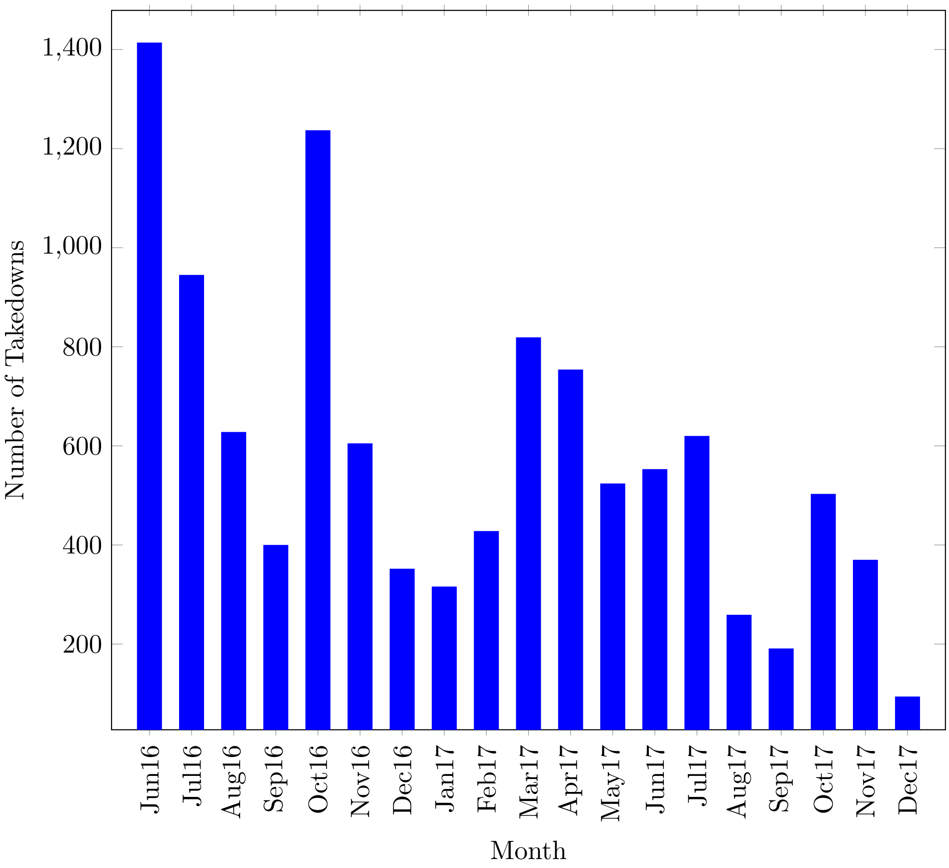

Figure 12: Volume of web-inject URLs removed from UK IP Space

in this section. In June 2016, they identified a total of 65,000 unique phishing attacks hosted on 21,000 unique IP addresses. In November 2017, this number had risen to 131,000 phishing attacks hosted on 30,000 unique IP addresses. If we look at the UK's share of this hosting, we were responsible for hosting 5.3% of the attacks in June 2016, but only 3.1% of the hosting by November 2017. Figure 13 shows the UK's share of phishing hosting since we started actively taking down sites in UK delegated IP space.

Obviously, phishing and web-inject attacks are not confined to the UK's IP space and most campaigns of these types are hosted elsewhere. There needs to be concerted international effort to have a real effect on the security of users. We return to this in section 8.

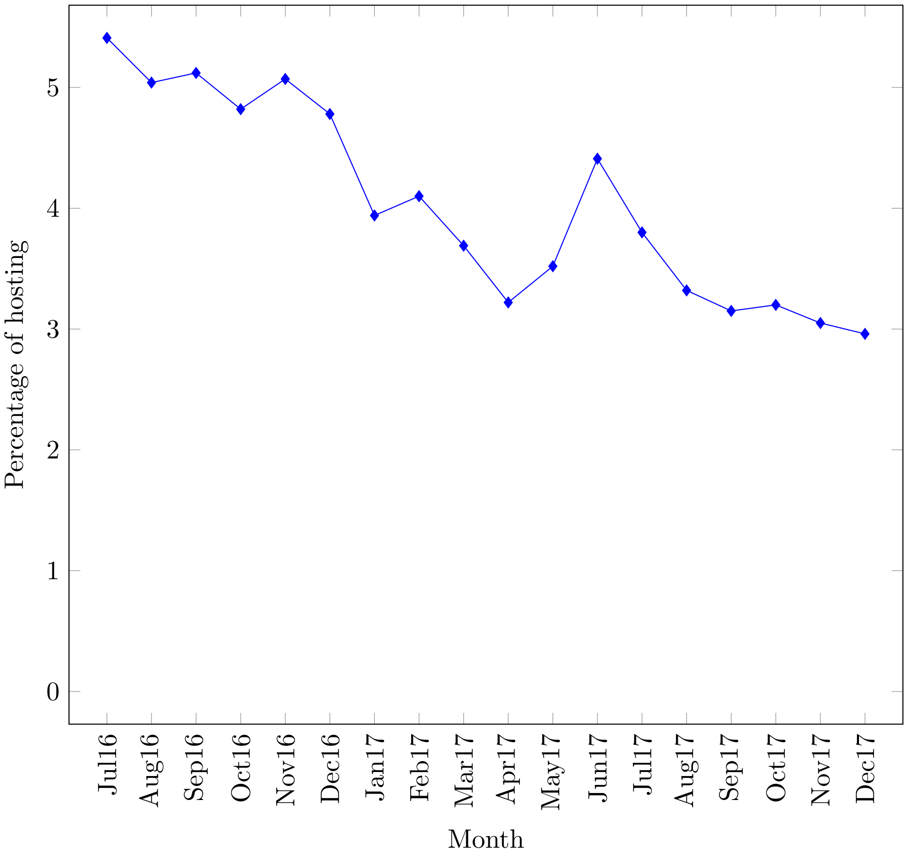

Figure 13: Percentage of global phishing hosted in UK IP space

3 DMARC

3.1 Domain discovery

We need to have a good handle on the domains used by the public sector to deliver services so that we can protect them better. There's well over 3,000 domains in the official list, recorded in the gov.uk DNS zone file. However, we know that the public sector hasn't limited itself to names under gov.uk and there's no central registry of names that have been used by government over the years.

To try to better understand the public sector's real internet footprint, we've built a self-service Domain Discovery tool, currently in alpha. This uses public information plus knowledge of how government generally works to try to build a catalogue of domains that are likely registered and used by government. We've also made it easy to export any new domains found in the catalogue to other tools, such as Web Check. At the time of writing, the tool has discovered 29,198 domains that we have the evidence to link to government. Noting that the tool remains in alpha, we have reasonable confidence in that number. It's worth noting that the number includes multi-level subdomains under *gov.uk*, but there are plenty under suffixes other than *gov.uk*. Some of those are in active use but a good number of them aren't. Obviously, we need to ensure that those that aren't being actively curated are taken under control, any services that are still running shut down and the domains protected so they can't be abused. It's not clear yet whether they should be removed or not, especially if they have a government brand associated with them. Once we understand the situation better, we'll work out exactly what to do. However, our initial findings do suggest that responsible domain owners bear a long term cost - in terms of curating the domain after they have finished with it, assuming it has currency with the public.

3.2 DMARC

We talked in section 2.1.1 about how attackers spoof email addresses to make victims more likely to open their phishing emails. The way internet email works is a lot like the way the real post works - there's an envelope and the content. The envelope is used to decide how to route the message and includes a 'to' address and a 'from' address, much like you'd write on the front and back of a real envelope. The internet uses that information to get the message to the intended recipient. Just like in the real world, there's nothing at all binding the addresses on the envelope to the addresses written on the letter inside. Just like in the real world, you can write a completely fake 'from' address on the back of the envelope and the message will still be delivered.

There's an internet standard[10] called Domain-Based Message Authentication, Reporting and Conformance, commonly known as DMARC, that attempts to address some of these issues by making sure that domain owners can have more control over who can use their email addresses as 'from' addresses. DMARC relies on two other standards : Sender Policy Framework (SPF)[11] and DomainKeys Identified Mail (DKIM)[12]. DMARC isn't perfect - far from it - but we believe that getting DMARC employed at scale will significantly raise the cost of performing attacks for a significant group of attackers. As we said right at the start of this paper, cyberattack is a business and reducing the return on investment - or increasing risk or cost - can have a significant effect. We believe that getting all Government domains using DMARC will help de-risk implementations for others by ensuring that we have actually implemented the technology

[10] RFC7489.

[11] RFC7208

[12] RFC6376

at scale and worked through any business impacts it could cause. It will also help us decide what sort of tools are useful to implementors and build some of those where appropriate.

DMARC, along with SPF and DKIM, allows domain owners to make a few different sorts of statement to potential recipients about their email - or, more accurately, to the destination mail servers that hold the recipients' inbox. These statements are made by publishing DNS records associated with the domain and are, basically :

- Email from my domain will only come from servers with the following IP addresses (SPF)
- Email from my domain will be signed cryptographically by my mail infrastructure using the following key (DKIM)
- When you get email apparently from me which fails the first two tests, enact this policy (DMARC)
- When you get email apparently from me which fails the first two tests, tell me at this email address (DMARC)

The 'policy' referenced can be *p=none*, which asks the receiver to take no special action, *p=quarantine*, which asks the receiver to put the message in the user's spam mailbox or *p=reject* which asks the receiver to not bother delivering the email to the end user. This works because, in general, attackers won't be able to use the mail servers listed in the SPF record and won't have access to the secret cryptographic key for the domain used by DKIM. So, they can't send from the right IP ranges and can't sign the messages properly, so the authentication should fail. DMARC then lets the domain owner decide what to do with those failures and also provides data back to the domain owner to help them understand how their brand is being abused by attackers. Our stated intention is to get all of public sector using DMARC (and therefore SPF and DKIM) on their email domains with a policy of *p=reject*.

3.3 Implementation

3.3.1 Mail Check

Mail Check is our platform for assessing email security posture. It currently includes collecting, processing and analysing DMARC reports for the public sector, but will be extended over the coming months to test for other email security features required by the Government's Email Security Standard[13], such as use of TLS.

There are two parts to implementing DMARC - 'setting the DNS records that expose organisational policy' and 'processing and analysing the resulting reports generated by receivers'. The second of these is by far the more difficult to achieve and a centralised service is the most obvious way to manage this for public sector. Mail Check is under constant development, but our intention is to build a system that allows individual domain owners to perform their own analytics on the failures related to their domains and for NCSC to perform analysis across the totality of data for the public sector.

Mail Check receives the email reports sent by receivers in accordance with a DMARC policy, processes them to extract relevant information and entities and stores them in a set of databases. There is a dashboard to help domain owners and us understand the macro effects, but also an analytic platform to help individual analysts answer more subtle questions.

The system is built entirely in public cloud infrastructure and so scales very well with the number of customers and number of DMARC reports processed. We have committed to keeping

[13] https://www.gov.uk/guidance/set-up-government-email-services-securely

Mail Check open source, and we would welcome contributions. Mail Check is available from the NCSC GitHub repo[14].

3.3.2 gov.uk

The *gov.uk* domain itself shouldn't ever be used to send email. That is, no-one should ever send email from an address like *ian@gov.uk*, it's always *ian@department.gov.uk*. However, it can be very convincing when used as a from address by attackers. Consider getting email from *taxrefund@gov.uk*. That's probably a lot more convincing than *taxrefund.uk2018@gmail.com*. In late 2016, we set the following SPF record relating to *gov.uk*:

```
gov.uk. 2252 IN TXT "v=spf1 -all"
```

This tells anyone receiving an email from an address *@gov.uk* that there are no valid mail servers that could possibly send that email. We also set the following DMARC record :

```
_dmarc.gov.uk. 2252 IN TXT "v=DMARC1;p=reject;sp=none;adkim=s;aspf=s;fo=1;
rua=mailto:dmarc-rua@dmarc.service.gov.uk;
ruf=mailto:dmarc-ruf@dmarc.service.gov.uk"
```

This tells anyone receiving an email from an address *@gov.uk* where the SPF check fails (which they all will, since we say there are no valid servers for an email address of this form) that the email should not be delivered to the intended recipient and that we should be told about the spoof at our reporting points. There are two sorts of report generated by receivers, according to the internet standard. Forensic reports provide full details of the spoofed message, including the message body and attachments. Aggregate reports provide details of the number of spoofed messages targeting a particular domain, along with the IP addresses of the mail servers that sent the spoofed messages and the action taken by the receiver (known as the *disposition*). Many mail receivers don't send forensic reports, probably because of privacy concerns around misconfigurations[15]. What follows is based only on the processing of aggregate reports.

Because of our DMARC statements, we should end up with no email from *@gov.uk* ever being delivered to end users. By processing the reports , we should get to know more about how attackers are abusing email addresses under *@gov.uk* and seeing whether we have had an effect on attacker behaviour. For emails apparently sent from *@gov.uk* over 2017, figure 14 shows the volume of email that was reported as not delivered to end users - the intent of our DMARC policy.

It is probably safe to assume that a number of these spoof emails are being used to attack people through malware by abusing the government brands. In section 2.1.2 we provided data for the number of mail servers involved in sending malware using government domains. For the purposes of this section, we will assume that the number of mail servers involved is related in a simple way to the number of messages involved.

Figure 15 plots the volume of messages reported as spoofing *@gov.uk* addresses and the volume of mail servers notified as sending malware from any government address. Obviously, data from only six months does not provide for a strong statistical base and so we must be very careful drawing conclusions from this dataset, but it looks as though the two things may be related. Operation of the services over the next year will provide more data from which we may be able to draw more concrete conclusions.

[14] github.com/ukncsc.

[15] Such misconfigurations in the domain owner system configuration could cause real messages to be sent to the DMARC processor.

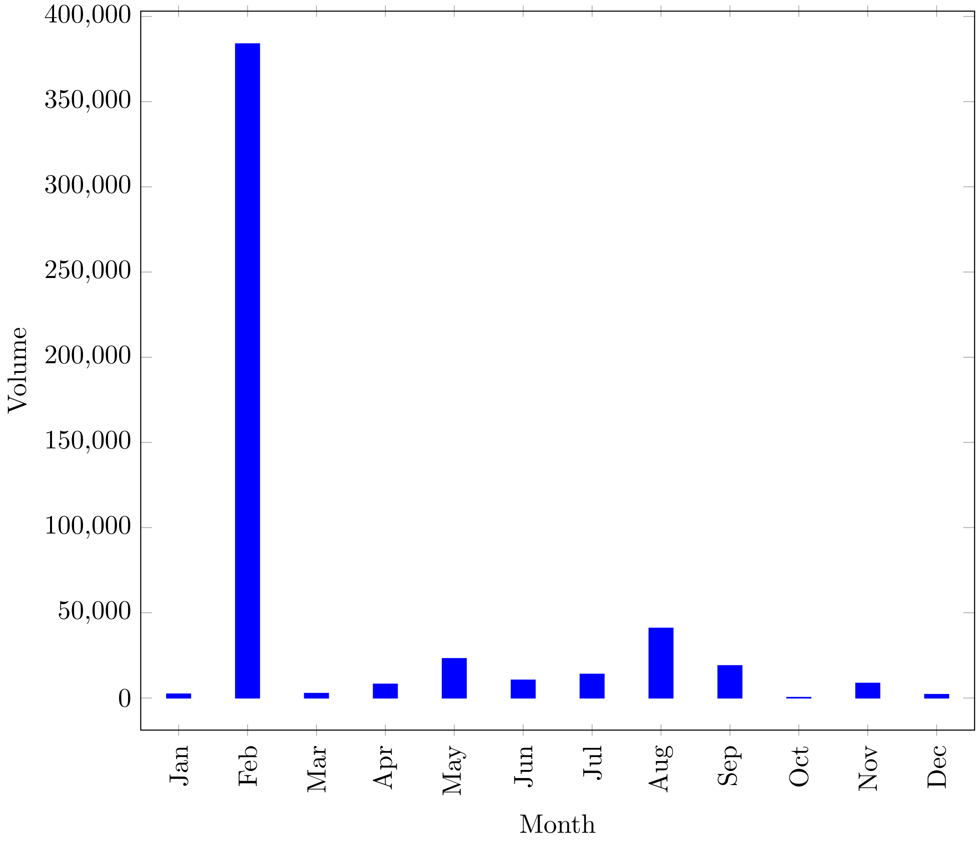

Figure 14: Volume of spoofed email using *@gov.uk* rejected by the receiver

Figure 16 shows the volume of spoofed *@gov.uk* email that receivers delivered in some way (which should not happen). At the time of writing, we have absolutely no clue what is going on, but are reaching out to the relevant receivers to try to find out.

For the year 2017, we had 167 different $(receiver, disposition)$ pairs reported to us with 159 unique receiver names (for spoofs of *@gov.uk*). For clarity, a receiver that only rejects spoofs from *gov.uk* will be counted once. A receiver that rejects some messages but delivers others will be counted two or three times. That means that 8 receivers process messages spoofed against *gov.uk* inconsistently. While this applies only to a relatively small number of receivers and messages this needs more investigation as to the underlying cause.

It is interesting to look at the top receivers who report rejections to us for *@gov.uk* spoof mails. We assume that all major receivers provide aggregate reports, although we know there are some providers who do not. The top three receivers who report rejections of *@gov.uk* messages is shown in table 9 and figure 17. As expected, the majority of the rejection messages come from large platform providers, such as Microsoft and Yahoo!. The inclusion of receiving mail services which are not primarily in the English language, for example qq.com (Chinese) and mail.ru (Russian)[16], suggests that the campaigns using *@gov.uk* addresses are not solely targeting UK users, although this likely speaks mainly to the quality of the email address list used by the attackers. However, the relatively small number of messages now using *@gov.uk* addresses as

[16] For clarity, this says only that people likely in Russia and China are potential victims of these campaigns.

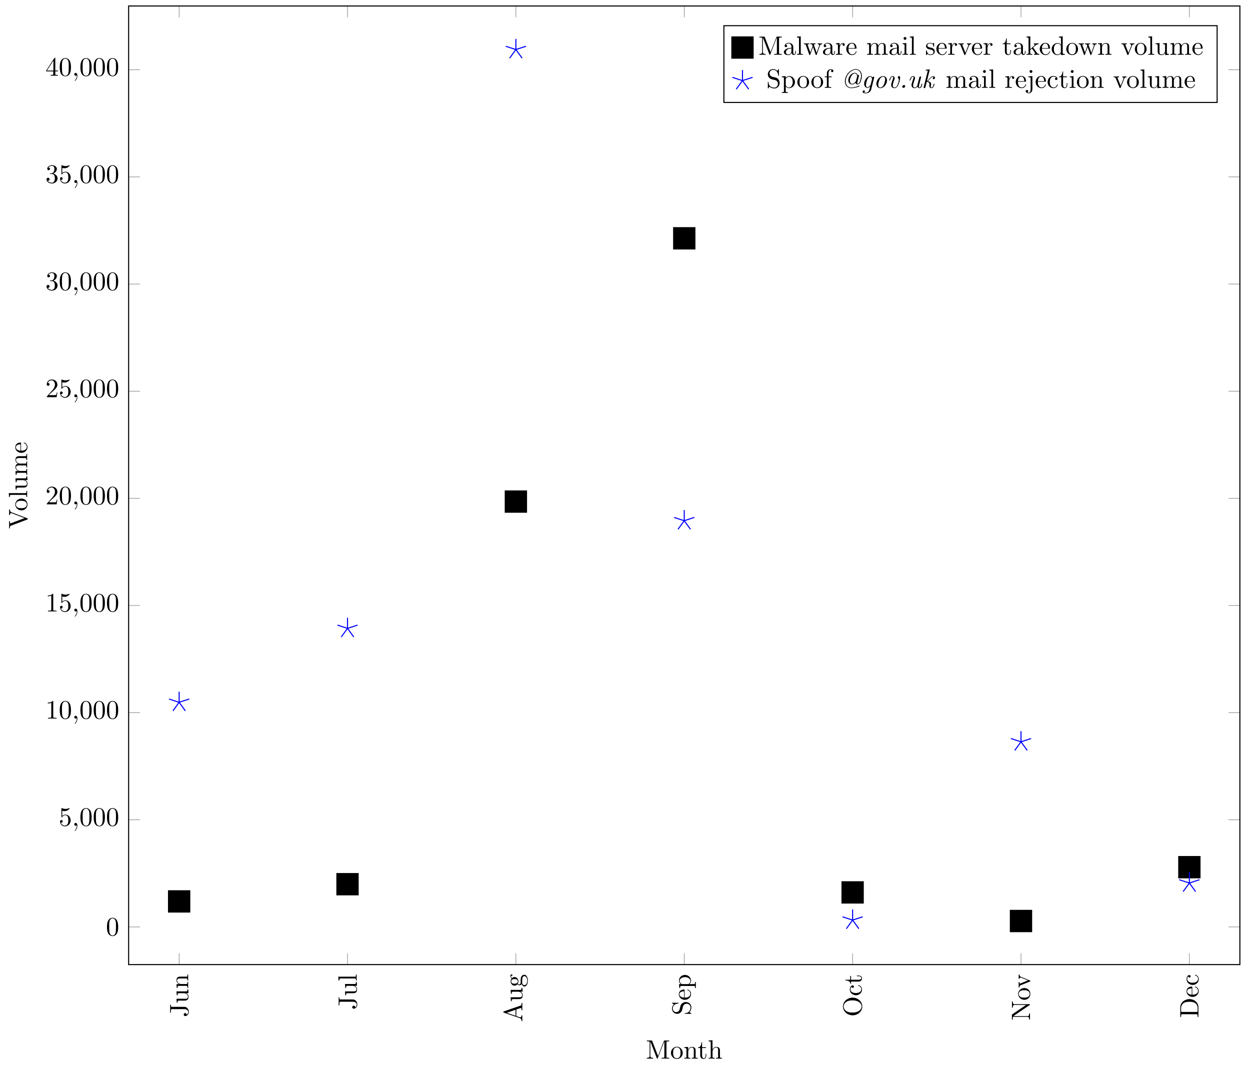

Figure 15: Volume of mail servers sending malware and *@gov.uk* mail rejected due to DMARC

their spoof probably means that we should not draw hard inferences from these data.

3.3.3 All registered public sector

The public sector standard for email includes a requirement for DMARC, among other things like proper support of hop-by-hop TLS encryption. We are prioritising 5,322 government domains for adoption in the first instance. These are mainly first and second level domains under the suffixes *llyw.cymru, gov.scot, gov.uk, gov.wales, police.uk and nhs.uk.* Those with SPF policies (a prerequisite of full DMARC) has risen from 26.85% at the start of 2017 to 38.56%. DMARC adoption is up from 5.58% at the start of 2017 to 18.3%. As of the time of writing, we have 555 public sector organisations reporting to our DMARC platform, representing approximately 10% of our target. Many of these have published policy of $p = none$, which is the first step in any sensible implementation, and we will be working with these domain owners to help them move

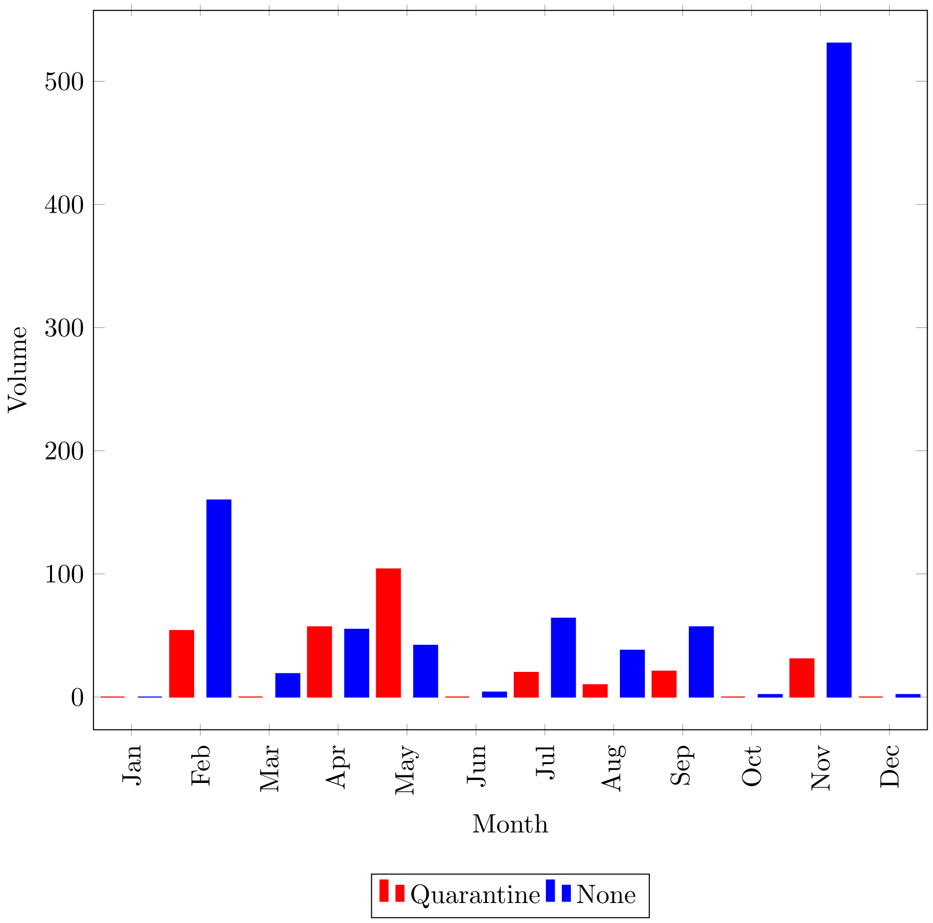

Figure 16: Volume of spoofed email using *@gov.uk* reported as not rejected by the receiver

to $p = reject$.

The volume of messages that fail either SPF or DKIM and so are reported to our platform are shown in figure 18. It is worthy of note that we are not saying that these are uniquely spoofed emails. When organisations start to use DMARC with a policy of $p = none$ it is entirely possible that there are misconfigurations or unknown parts of infrastructure that would cause some valid messages to be processed as a DMARC failure and reported. This is much less likely at $p = quarantine$ and $p = reject$, because any such misconfigurations would likely have been ironed out as the owner moved away from $p = none$. Therefore, it would not be reasonable to say that we are stopping an average of 44 million spoofed emails a month sent in the name of UK public sector. It's certainly reasonable to say that we are stopping an average of 4.5 million spoofed emails a month sent in the name of gov.uk (the total of those with disposition 'quarantine' and 'reject'), and allowing analysis of another 40 million or so. While we will do much more analysis on this dataset, it is not immediately clear what solid inference around actor behaviour can be drawn from this dataset until the number of public sector domains with a disposition other than $p = none$ increases significantly.

Figure 19 and table 10 show the three *gov.uk* domains that are reported with the highest monthly message rejection volumes. We have chosen only those with domains with $p = reject$

Month	No. Of Reported Rejects	Reporting Receiver
Jan	1,745	Yahoo! Inc
	243	qq.com
	183	Microsoft Corp
Feb	314,280	Microsoft Corp
	35,646	AOL
	30,941	Yahoo! Inc
Mar	2,360	qq.com
	206	Yahoo! Inc
	103	Microsoft Corp
Apr	3,758	qq.com
	2,456	Microsoft Corp
	1,236	Yahoo! Inc
May	14,537	Microsoft Corp
	4,834	Yahoo! Inc
	2,072	qq.com
Jun	4,402	Microsoft Corp
	2,544	Yahoo! Inc
	1,972	qq.com
Jul	10,075	Microsoft Corp
	2,234	Yahoo! Inc
	1,403	AOL
Aug	19,037	Microsoft Corp
	12,921	Yahoo! Inc
	6,794	AOL
Sep	17,682	AOL
	285	mail.ru
	271	Yahoo! Inc
Oct	6,845	Yahoo! Inc
	802	comcast.net
	561	emailsrvr.com
Nov	13,528	Yahoo! Inc
	13,488	AOL
	8,839	mail.ru
Dec	1,259	Yahoo! Inc
	773	AOL
	71	FastMail Pty Ltd

Table 9: Three receivers who report the most rejection of @gov.uk messages per month

since they are unlikely to generate large scale misconfiguration-related DMARC failures. However, it appears that the rejections for the **.kent.gov.uk* and **.hillingdon.gov.uk* domains are due to misconfigurations which are almost certainly benign. The Kent misconfigurations were fixed earlier in the year. The rest of the rejections are almost certainly attack-related. Remember that these are organisations that have put in place measures to ensure that their domains are much harder to spoof and so this list shows public sector organisations that are doing the right thing.

It is interesting that HMRC - consistently the most abused government brand[17] - does not

[17] For reasons previously discussed.

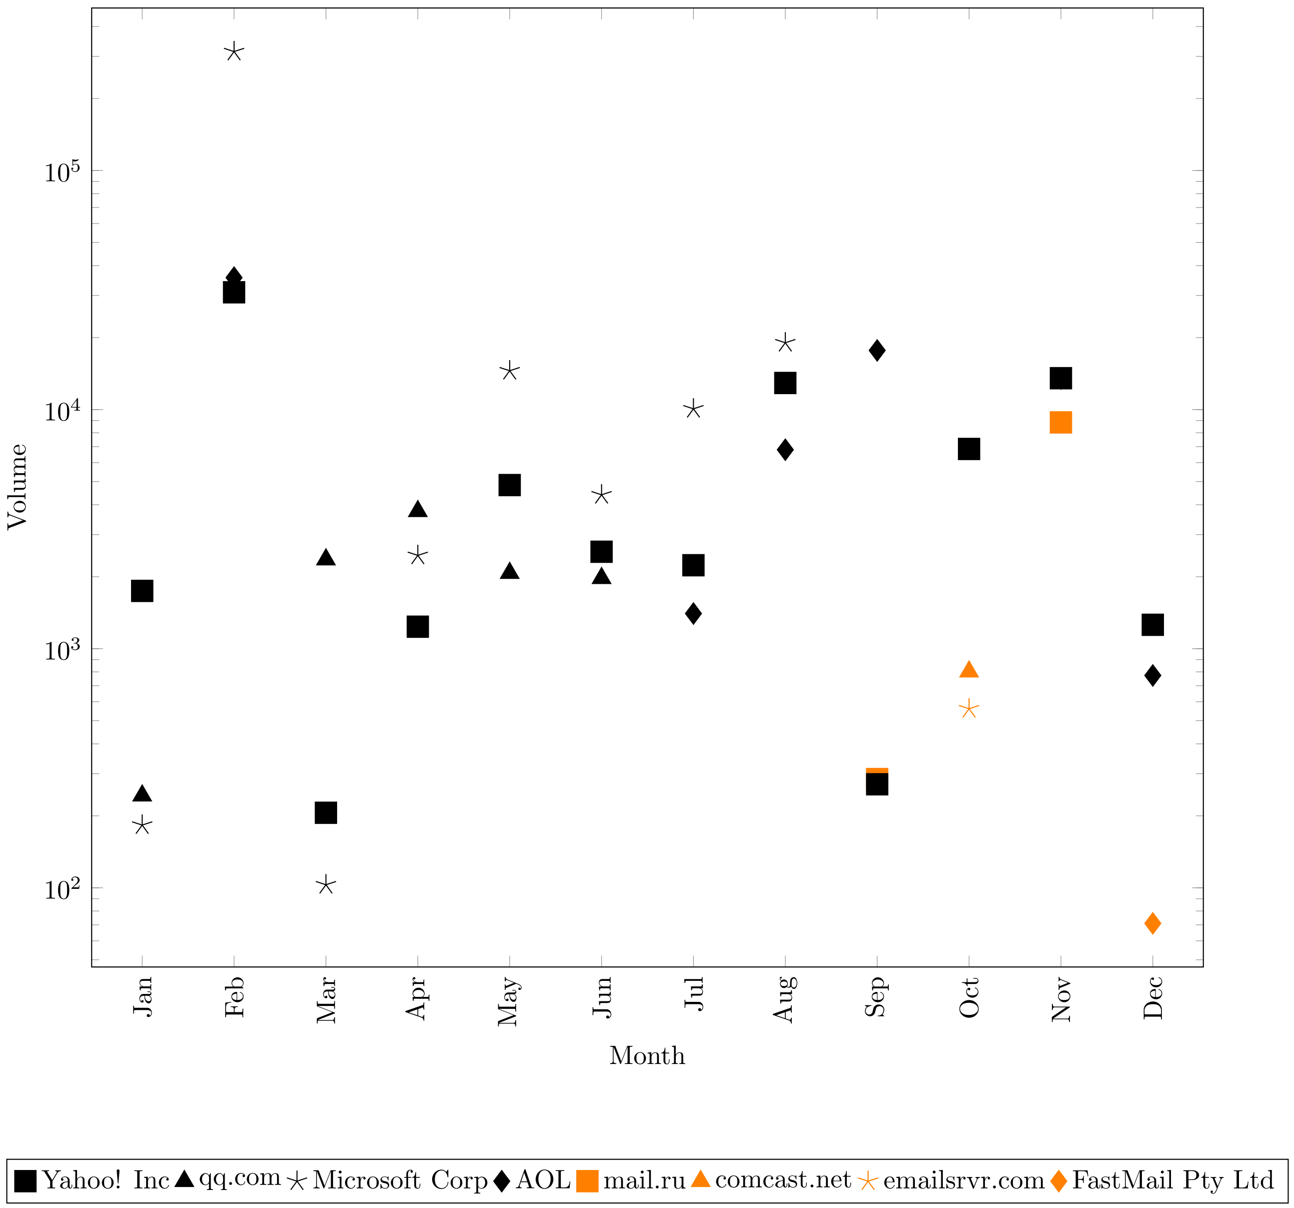

Figure 17: Three receivers who report the most rejection of @gov.uk messages per month (note logarithmic scale)

appear in the top three every month. It is feasible that HMRC is prevalent in January and February because this is the UK Self Assessment Income Tax peak[18] and the most convincing time for refund notifications to be sent. April is the end of the financial year for many companies and so the May peak may be targeting corporation tax. Without forensic reporting, it is impossible to know for sure. HMRC were the first major government department to properly implement DMARC. It is easy to suggest that HMRC's absence from the top 3 spoofed domains most

[18] UK taxpayers must submit by 31st January or face penalties.

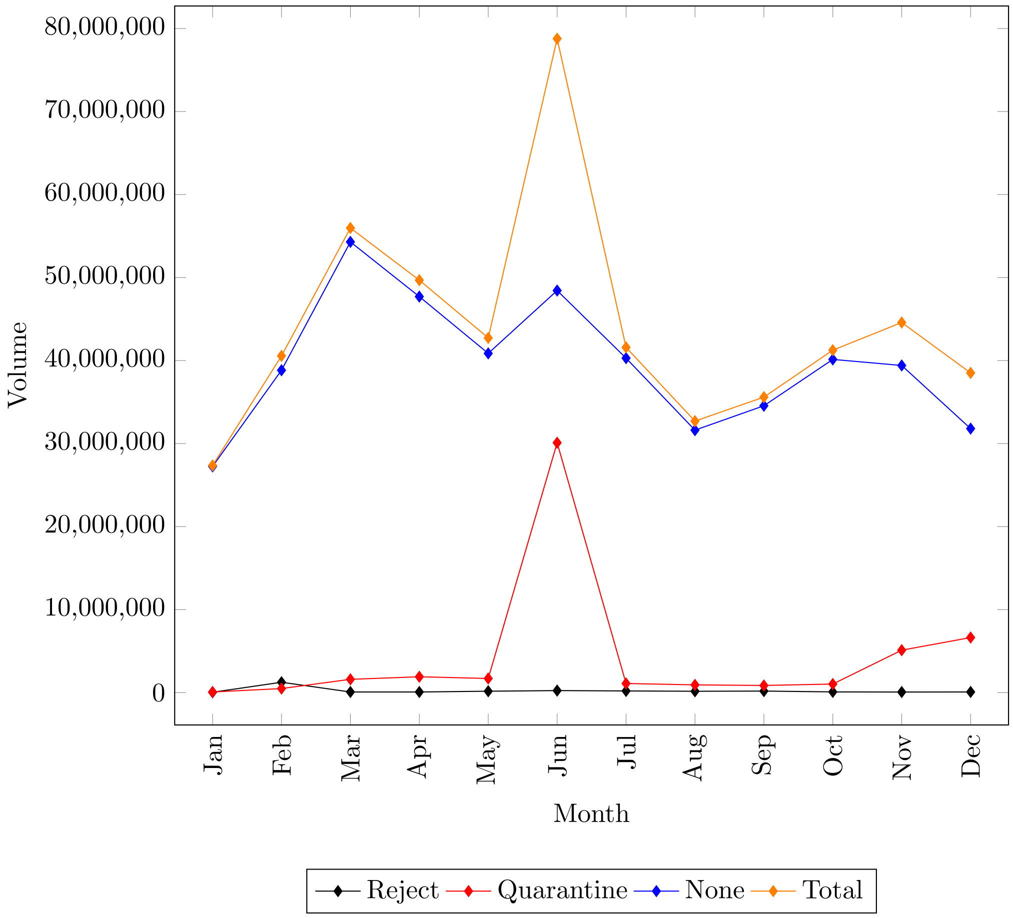

Figure 18: Volume of reported mail using any registered public sector domain, split by receiver disposition

months is a direct result of their early adoption of DMARC, especially when taken together with the use of the HMRC brand in other forms of commodity cyberattack. While this is probably true, we do not have the data to prove it conclusively and it is unlikely that we can get data from before HMRC's implementation that is sufficiently similar in collection characteristics to allow comparison. We will, however, be able to track the effect of applying DMARC to other HMG brands much better and will report on that in due course.

Recall in figure 13 there is a peak in UK hosting of phishing in June 2017. There is a similar peak in June in the DMARC reporting. It would be easy to be convinced that the peaks are related - that someone tried to use UK public sector domains to launch a campaign physically hosted in the UK. Given the analysis we have performed so far, we cannot make that link.

We expect more detailed analysis of the dataset we currently have and future datasets to be very informative. We will publish our analyses.

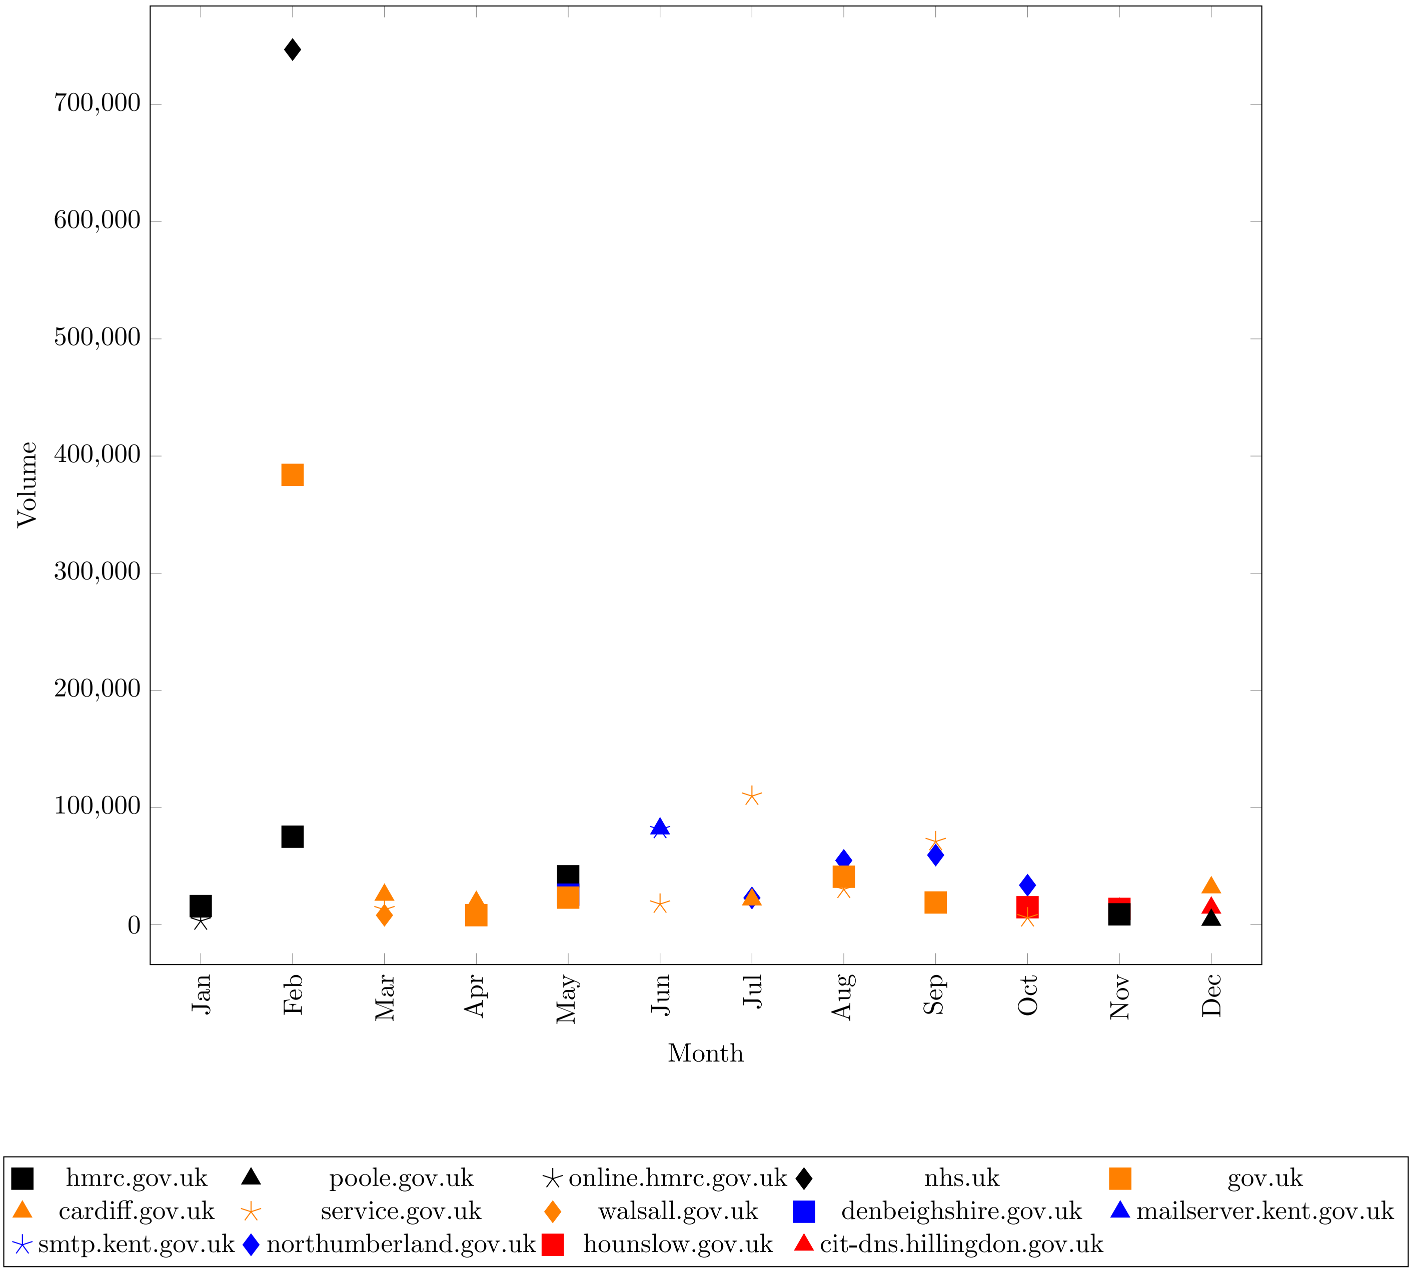

Figure 19: Volume of rejected mail from top 3 registered public sector domains by month

3.4 Future work

3.4.1 Adoption

The biggest challenge is getting adoption for DMARC. We believe that the excellent, groundbreaking work undertaken by HMRC - which has very stringent deliverability requirements and a very complex set of enterprise requirements and networks - initially helped us show that

Month	No. Of Reported Rejects	Spoofed Domain
Jan	15,907	hmrc.gov.uk
	11,076	poole.gov.uk
	3,154	online.hmrc.gov.uk
Feb	746,957	nhs.uk
	383,927	gov.uk
	75,088	hmrc.gov.uk
Mar	25,424	cardiff.gov.uk
	12,789	service.gov.uk
	8,047	walsall.gov.uk
Apr	19,308	cardiff.gov.uk
	11,009	walsall.gov.uk
	8,113	gov.uk
May	41,332	hmrc.gov.uk
	25,627	denbighshire.gov.uk
	23,115	gov.uk
Jun	82,031	mailserver.kent.gov.uk
	81,168	smtp.kent.gov.uk
	17,620	service.gov.uk
Jul	109,739	service.gov.uk
	22,844	northumberland.gov.uk
	21,400	cardiff.gov.uk
Aug	54,833	northumberland.gov.uk
	40,943	gov.uk
	30,221	service.gov.uk
Sep	70,979	service.gov.uk
	59,405	northumberland.gov.uk
	18,956	gov.uk
Oct	33,685	northumberland.gov.uk
	14,896	hounslow.gov.uk
	5,926	service.gov.uk
Nov	13,528	cit-dns.hillingdon.gov.uk
	13,488	hounslow.gov.uk
	8,839	hmrc.gov.uk
Dec	31,728	cardiff.gov.uk
	14,424	cit-dns.hillingdon.gov.uk
	4,041	poole.gov.uk

Table 10: Three registered domains with highest reported rejection volume by month

complex DMARC deployment was possible. Recently, Government made a concerted effort to adopt DMARC across the core email domains, which we've equated to 5,322 domains. As stated previously, at the start of 2017, only 26.85% of those domains had any sort of SPF record. Now it is 38.56%. At the start of 2017, only 5.58% of those domains had any DMARC record, which has risen to 18.3% over the year. That doesn't mean that all those SPF and DMARC records are as we would want, but it's a good start. There are 555 public sector domains reporting to Mail Check.

The new release of Mail Check that allows individual departments to do analysis on their own

DMARC records will help further adoption across public sector. There are also some *incentivising* options if we need them. Once we have public sector adoption at a respectable value, we will be looking at brands and sectors that have high public trust with the UK public and start to understand where we need to push businesses and entire sectors to implement DMARC at scale to better protect the public. We are also talking to other governments who would like to protect their email domains from being spoofed. We will continue to offer the benefit of our experience of implementing DMARC at a country level and the source code to our tools freely.

3.4.2 Move to *p=reject*

As adoption continues, we believe that we will naturally see the evolution of DMARC policies from $p = none$ through $p = quarantine$ to $p = reject$. We know that there have been instances of internet infrastructure in the UK which, through poor design or implementation, actively broke the ability to accurately evaluate SPF, DKIM or DMARC policies. These cause most issues when an organisation has a policy other than $p = none$ as it can interfere with the delivery of valid email to large swathes of customers. Where we have become aware of those issues, we have worked with the owner of the infrastructure to try to fix the issue. It is almost certain that there will be more of these infrastructure issues that will need resolution. We hope that infrastructure owners will proactively fix their implementations and we will obviously provide help if it is required. We are also willing to intervene if particular infrastructure owners are intransigent in fixing their networks.

3.4.3 Synthetic DMARC

Towards the end of 2017, we saw an increase in the number of attacks using non-existent subdomains of *gov.uk* to mount their attacks - *taxrefund.gov.uk* looks legitimate, but really isn't. Since there are no valid domain name records for these non-existent domains, it's not possible to create DMARC and SPF records for them. With the help of the IETF DMARC community, we are currently testing synthesising SPF and DMARC records for non-existent domains [19] [20]. While we know the current implementation isn't perfect, it does appear to be working reasonably well. At the time of writing, the experiment has been running for a couple of weeks.

We can confirm that we can synthesise the appropriate SPF and DMARC records in DNS for non-existent domains and we have reports in Mail Check that show attackers are using this sort of domain as part of their campaigns. The data we have also shows that there are further (non DMARC-related) misconfigurations on some public sector systems that we were not previously aware of. We will work to better understand the potential impact of these misconfigurations and work with the affected organisations.

From what we have seen so far, we should be able to implement the effects of DMARC on non-existent domains under public sector suffixes, further restricting how attackers can abuse public sector brands as part of their attacks. We are currently testing this with a policy of $p = none$ and, once we have enough data, will do the analysis to ensure that moving to $p = reject$ will not break anything. Regardless, this is giving us further insight into how attackers work.

[19] We've nicknamed this 'synthetic DMARC' partly because it describes the solution but mainly because it sounds kind-of-cool.

[20] We are aware of the potential for a bad implementation of this to be used as part of an amplification attack.

3.4.4 The Public Suffix List

The Public Suffix List is a community resource maintained by the Mozilla Foundation[21]. It is intended to list the domains under which internet users can directly register names and is used in a number of ways to help prevent abuse on the internet. So, *.com, .co.uk, service.gov.uk* are all on the Public Suffix List. So is *gov.uk* itself.

The DMARC standard defines the concept of an *organisational domain*[22] which is used to try to work out what organisation an email has been received from in order to retrieve the correct set of DMARC-related policies from the DNS. The standard defines an algorithm that expects to define an organisational domain as a primary subdomain of an entry on the public suffix list. When email is spoofed from a *@gov.uk* address, the policy should obviously be retrieved from the DNS records for *gov.uk* and *_dmarc.gov.uk*. However, since *gov.uk* is also on the public suffix list, it is unclear from the standard as to what the organisational domain should be set to, since there are no labels below the name on the public suffix list. Empirically, many implementations of DMARC processors must allow for policies to be retrieved from entities that are on the public suffix list as we are receiving reports. However, more work needs to be done to understand how the main DMARC processing implementations act in this case, especially if other governments are to be encouraged to do something similar as most government-level domains are also on the public suffix list.

3.4.5 Understand how receivers process better

The potential issue around protecting domains on the Public Suffix List is one example of indeterminate receiver behaviour. We have seen other odd behaviour in receivers, for example the multiple disposition reports for contemporaneous spoofed emails described in section 3.3.2. To quote one of our analysts, ‘There’s plenty of weird in the data.’ and we do not understand some of the artefacts we are seeing. Some of that will be due to quirks in how indiviual receivers process spoofs and organisational policies, which they may not publish. Some may be due to interactions between various standards that we - and others - may not aware of. We will be talking to individual receivers over the coming months to get more details about their policies and intentions and how they play out in the data we are collecting.

3.4.6 Data analysis

As more public sector bodies register for the Mail Check service and move towards a DMARC policy if $p = reject$, our data volume and fidelity will improve. More advanced analytics can be performed on that dataset and, as we learn more about the vagaries of the various email receivers and how they have interpreted the standards, we should be able to draw concrete conclusions about how email is being abused by criminals and further steps to make it more difficult.

3.5 Conclusion

Our DMARC implementation continues to evolve and as more domains are moved to $p = quarantine$ and $p = reject$ the fidelity of the data we receive will increase. The data we are collecting shows that government domains are actively being spoofed for use in attacks and those public sector bodies that have set a policy of $p = reject$ are actively protecting both their brand and potential targets of cybercrime. The spoofed government domains that we can see being rejected by receivers (and therefore almost certainly being used by attackers) do not seem to follow

[21] See *www.publicsuffix.org*.

[22] RFC7489, section 3.2.

any specific pattern or obey any specific logic. As more public sector bodies move to $p = reject$ we will get more specific data which may help us better understand attacker mentality.

We have discovered a number of complications - both technical and driven by adversary behaviour - around implementing DMARC at a national level. We will continue to work on solutions to these issues and work with receivers to better understand how we can collectively make the ecosystem more predictable and less risky for domain owners.

Our experience with the trend of reports for the *@gov.uk* spoofs (and we can be sure they are real spoofs) suggests that they are being used much less than at the start of the work, implying that criminal behaviour is changing in the face of our defences. If the number of misconfigurations is low in the current population of all domains reporting to Mail Check, the data we are collecting suggests that tens of millions spoof emails are being sent using government-related domains each month. As we onboard more public sector domains and move their policies to $p = reject$ we will start to both understand the scale of this abuse and also stop the mails from being delivered to end users, reducing the harm these attacks can cause and increasing the cost of attacks.

However, in order to draw more high confidence conclusions, we need more data and more complex analysis. Once we have that analysis, we expect to be able to identify more concrete solutions to make it more difficult for criminals to abuse email at scale.

4 Web Check

The public sector has a lot of internet properties and websites and they are not all secured to a common high standard. There is an intent described in the Active Cyber Defence blog to have all government web properties scanned for vulnerabilities automatically. While this is not technically that difficult, getting all of the public sector comfortable with us automatically testing their sites without notice will take some time[23]. We also need to help public sector organisations get better at managing vulnerabilities that are reported to them, regardless of whether the source of those reports is an automated tool such as Web Check or a security researcher responsibly disclosing something they have found. Our vulnerability disclosure pilot detailed here[24] is the work we are undertaking to help build a process for all of the public sector to follow which will minimise the difficulty in reporting a vulnerability in a public sector system and maximise the chance of it being fixed quickly and competently.

This is mainly a business change problem, rather than a complex technical one. The current Web Check capabilities are simple and intended to be low risk. We need to ensure that public sector organisations are confident that we will not break their service so Web Check scans currently need to be initiated by the service owner. As we help the public sector build confidence in the technology and their ability to respond to reports, we will expand the capabilities of Web Check and - importantly - move towards scanning all government services, whether under *gov.uk* or not.

4.1 Capabilities

Web Check's capabilities are implemented in modules that perform specific checks. The modules currently in use are :

- CMS
  - Attempts to determine if a content management system being used to serve the website and checks its version.
  - Supports Drupal, Jadu, Joomla, Umbraco, and WordPress.
  - Reports whether the CMS is no longer supported, supported but not on the current major version, or on the current major but not minor version.
  - Fetches the URL, parses as HTML, and checks a small number of well-known paths if necessary.
  - Needs to be kept up to date with version numbers.
- Domain
  - Checks whether the URL contains a domain reserved for use by the public sector (eg foo.gov.uk, foo.ac.uk, not foo.co.uk, foo.com).
  - Only uses DNS.
- HTTP
  - Checks whether the website redirects from a reserved domain to a public domain (eg foo.gov.uk → foo.com)
  - Checks whether the website redirects from HTTPS to HTTP.

[23] We do note, of course, that adversaries do this often and without regard for the stability of the target site.

[24] https://www.ncsc.gov.uk/blog-post/vulnerability-co-ordination-pilot

  - Checks whether the website returns a valid HTTP status code.
  - Checks for HTTP/2 support.
  - Checks a variety of HTTP security headers, but does not currently report findings for them.
  - Makes a small number of HTTP requests.

- Server Version
  - Attempts to determine the server software being used to serve the website and checks its version.
  - Supports Apache, Nginx, and OpenSSL.
  - Reports whether the server is no longer supported.
  - Fetches the URL and inspects the Server HTTP header.
  - Needs to be kept up to date with version numbers.

- TLS
  - Checks whether SSL/TLS is supported and performs a full SSL/TLS vulnerability assessment.
  - Supports SSLv2, SSLv3, TLSv1.0, TLSv1.1, TLSv1.2, TLSv1.3 draft 23
  - Identifies all known SSL/TLS vulnerabilities.
  - Connects to the host over SSL/TLS a medium number of times to assess vulnerabilities (around 30).

- x509:
  - Checks the website's certificate chains for issues (including expiry)
  - Checks all certificate chains for issues, such as too many or too few certificates being sent by the server, whether the domain matches etc.
  - Checks for any certificates blacklisted by browsers (e.g. Symantec, DigiNotar).
  - Connects to the host over TLS three times to fetch each certificate type.

- WannaCry:
  - Checks for artefacts of the WannaCry infection in the server.
  - Checks whether ports 137, 138, 139, 445, or 3389 are open for TCP or UDP.
  - This module should probably be renamed - having those ports open is bad whatever the cause.

- WordPress:
  - Checks whether the website is being hosted using WordPress and checks whether the version is supported.
  - Needs to be kept up to date with version numbers.
  - This module has been deprecated by the CMS module, but is included for completeness.

This is a relatively simple set of capabilities, but the tests present a very low risk to the services under test, while providing real security benefit. Web Check's capabilities will be expanded over the coming months.

4.2 Results

4.2.1 Scan results

The current implementation of Web Check has been operational since 18^{th} April, 2017. Between 18^{th} April, 2017 and 31^{st} December, 2017 Web Check performed 1,033,250 scans involving 7,181,464 individual checks. During that period, Web Check scanned 7,791 unique URLs across 6,910 unique domains, ingesting a total of 7,748 unique pages. Across those tests, Web Check has generated 4,108 advisories to customers, covering 6,218 different issues that needed resolution. We do not currently account for the cases where an issue is reported, fixed and then reoccurs, so the actual number of issues may be slightly higher, but anecdotally the rate of recurrence is low. The numbers of issues that are reported by each module are shown in table 11 and are grouped by unique URLs (rather than unique domains). An advisory is a notification to customers (e.g. *You need to address certificate issues*) while issues are the individual weaknesses that must be resolved (e.g. *Certificate close to expiry* and *Certificate does not match domain name*). We report by advisory in order to not artificially inflate the problems found.

Web Check Module	Number of reported issues
x509	2,178
Domain	0
HTTP	1
CMS	184
TLS	1,629
Server Version	76
WannaCry	36
WordPress	4

Table 11: Number of advisories reported by Web Check modules 18th April-31st December 2017

It is obvious from these results that certificate management is an issue for the public sector websites we have scanned in 2017. Some of these errors are potentially benign, for example a certificate that is nearing expiry. Providing the service owner acts quickly when receiving this notification, no security issue actually occurs. We do, however, have examples of service owners - or their contracted IT companies - failing to renew a certificate in sufficient time to avoid security issues.

However, Web Check has identified serious security issues where certificate management is the underlying cause. When Web Check tests a site, it attempts to negotiate an SSL or TLS connection with the site. Where it can successfully negotiate an encryption session, but cannot build a valid certificate chain for whatever reason, we record an error of *bad certificate*. Where this is the case, we make a further check. If we are unable to build a valid certificate chain and the certificate subject does not match the URL's domain, we check to see if the certificate subject matches another domain which resolves to the same IP address. If it does, we record a secondary error of *bad provider* as this is almost certainly due to misconfiguration by the underlying hosting provider. In the period we can report on, Web Check identified 1,137 unique URLs with *bad certificate* errors. Of these, 824 - 72% - we also subject to *bad provider* errors. More investigation is needed to understand the distribution of this error across hosting providers.

It is well understood that not keeping software up to date is a primary cause of cyber security issues. If we take the results of the *CMS* and *Server Version* modules as a broad indicator of how well patched these web properties are, then we see that around 3.3% of URLs tested appear to have patching issues. While an ideal value here would be zero, it is likely impractical. However, we do not know what a *good* value would be. Patching of internet-facing services is particularly critical as it is easier for adversaries to find and exploit known vulnerabilities in out of date software when they can communicate with the service easily, as is the case with a website. The trend of this value over the coming months will be interesting.

The population of Web Check users is taken from across the public sector. However, local authorities dominate the customer set as a sector (some 38% of all customers) and, on average, are smaller and can be less well resourced. It is more likely that such departments may need help in identifying issues that are not directly service affecting, simply due to resource levels. It may be that Web Check provides a significant benefit to these departments, despite being relatively simple in the nature of the current testing.

4.2.2 Issue resolution

Identifying issues is obviously only part of the problem. If issues are not resolved quickly, then there is little point in performing the scans in the first place. Figure 20 shows the resolution time in days for issues reported by Web Check, while figure 21 shows the cumulative number of issues resolved over time. While the majority of issues are fixed quickly - the vast majority within two days - there remain a number of issues outstanding after an extended period of time. We are continuing to investigate what we can do to help in these cases.

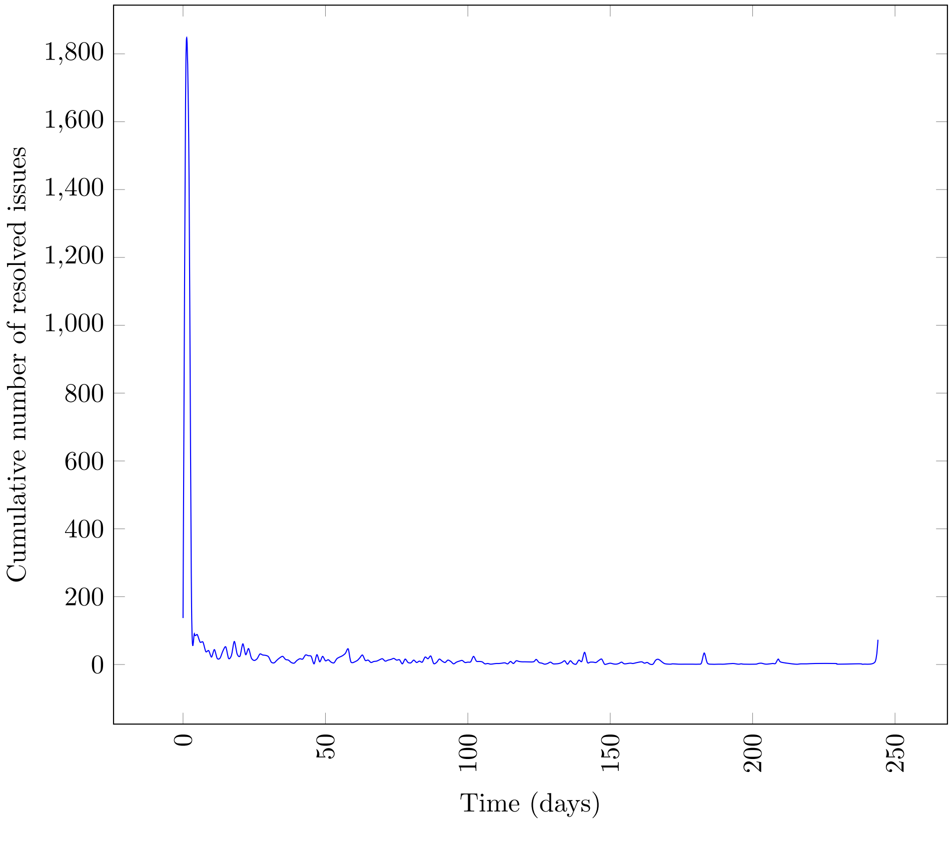

Figure 20: Resolution time for Web Check issues

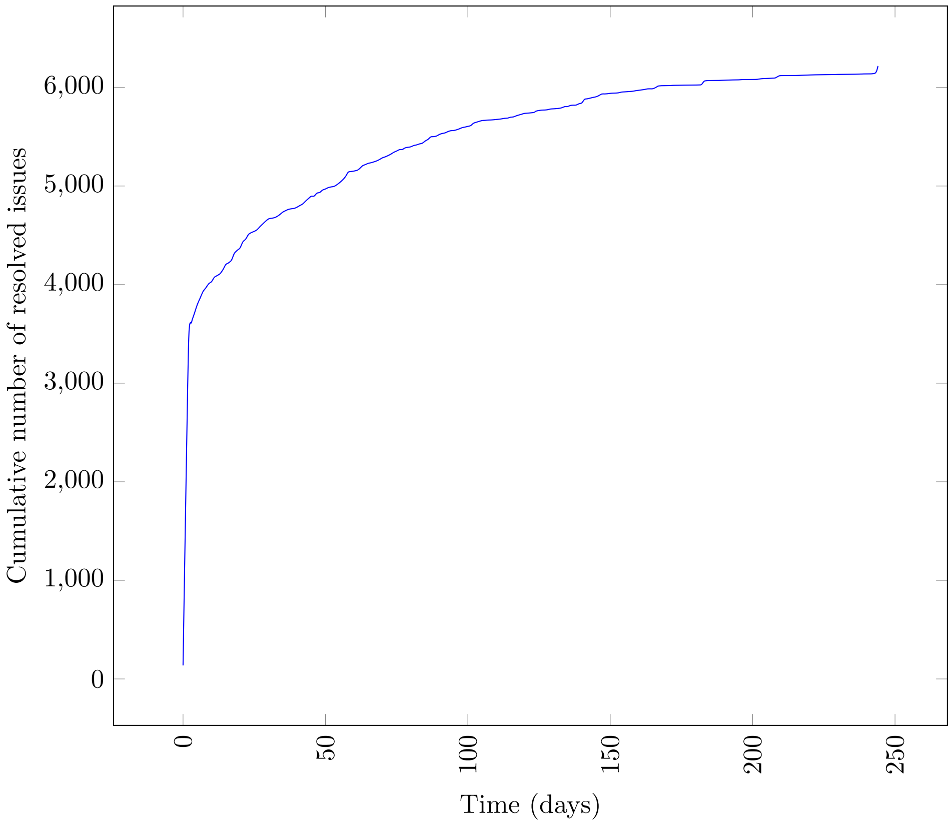

Figure 21: Cumulative number of Web Check issues resolved

4.3 Conclusion

Web Check is still early in its development and adoption by the public sector. However, it has already proven to be a useful tool in discovering relatively simple security issues at scale across public sector organisations. There is business change needed in public sector organisations to ensure that they are capable of receiving and responding to vulnerability reports and this will be an ongoing requirement. Web Check has shown that simple tests, at scale, can have a measurable positive effect on the the security of the web sites involved. We will continue to evolve Web Check with more complex tests and wider application to all public sector websites, without the need for service owner action.

5 Public sector DNS

5.1 Introduction to DNS

The Domain Name System (DNS) is critical to the operation of the internet and all the services that run on top of it. At its most basic level, it's like the phonebook for the internet, translating names that humans understand, for example *www.ncsc.gov.uk* into IP addresses that end user devices, servers and the internet infrastructure understand, for example *54.230.14.55.* But DNS is also used for lots of other things, like finding out how to deliver email to an organisation and various other infrastructure-related services. From a cyber security point of view, almost all cyber attacks use DNS at some point in their lifecycle. Whether it's the phishing link or compromised website, or the attacker's command and control system for their malware (and lots in between), DNS is almost always used at some point.

To understand how DNS can be used in a protective manner, it's important to understand how it works, at least at a high level. DNS is based on a hierarchical structure and is a *recursive* system, in that the clients walk up the tree until they get an answer (or everyone decides there isn't an answer).

If we look at where harm manifests as a result of a cyber attack, it's the client machine (or server if it's an infrastructure attack, but the principle ends up being the same). Consider the simple example where someone has opened a malicious attachment and the script in there is trying to call out to download the real malware. Events would proceed something like this :

1. User opens attachment and malicious script runs
2. Malicious script tries to connect to *evilmalwaredownload.co.uk*
3. User's machine doesn't know the IP address of *evilmalwareownload.co.uk* and so asks the departmental DNS server if it knows of it
4. The departmental DNS server doesn't know anything about *evilmalwaredownload.co.uk* and so asks the upstream Internet Service Provider's DNS if they know about it
5. The ISP's DNS doesn't know anything about it, so asks the Authoritative DNS for the Top Level Domain (TLD) involved. In this case, it's *.uk* so the question goes to Nominet
6. Nominet's Authoritative Server replies to the ISP's DNS service '*You need to go and ask the authoritative server dns.registrar.co.uk what the address is.*'
7. The ISP's DNS server asks *dns.registrar.co.uk* for the address of *evilmalwaredownload.co.uk*
8. *dns.registrar.co.uk*'s server replies to the ISP's DNS service '*The IP address for evilmalwaredownload.co.uk is 1.2.3.4. You can remember that value for an hour before asking again.*'
9. The ISP's DNS service then responds to the department's DNS service '*The IP address for evilmalwaredownload.co.uk is 1.2.3.4. You can remember that value for an hour before asking again.*'
10. The department's DNS service then responds to the user's machine '*The IP address for evilmalwaredownload.co.uk is 1.2.3.4. You can remember that value for an hour before asking again.*'
11. The user's machine goes off to the IP address 1.2.3.4 and retrieves whatever the script told it to.

In this example, each DNS server is called a *recursive server* apart from the *.uk* server and the registrar's server *dns.registrar.co.uk* which are called *authoritative servers*. If any of the servers in this hierarchy knew up front that *evilmalwaredownload.co.uk* was likely to cause harm to someone visiting it, then they could lie in their response and say *'Hey, I can't find that. Sorry.'*. In most cases, the script on the user's machine would give up and that potential harm would be mitigated. Our public sector DNS service is a service that effectively replaces the ISP's recursive DNS service in this example and performs that blocking function, with some other analysis detailed later. The cyber security functionality of blocking malicious domains is achieved through the use of standard Response Policy Zones and some more complex back end analytics. 'Departments' in this example would be all public sector organisations (eventually) and users are employees of those organisations[25]. It's worth saying that the DNS service at that level sees everything that's trying to get out of the department's network to the internet, regardless of whether it was initiated by a user's action, a server doing some processing or anything else. It sees all lookup attempts, good and bad, for all connections the department's IT is attempting to make. When the service sees a resolution request for a harmful domain, we only know that something in the customer department requested it. From the DNS service, we can't work out which machine is making the request. That needs the co-operation of the department. Basically, this service defines the department's view of the internet. If we can get all public sector organisations using the same DNS service, we can curate that service to block access to malicious domains and mitigate harm. By doing some data analysis on the queries, we can also discover interesting things.

Obviously, if every public sector organisation is using the same DNS service, it becomes absolutely critical to the operation of those organisations. NCSC could have built a DNS service, but we have no experience of building and running such a thing, so why would anyone trust us to do so? Instead, we chose to work with the UK Registry, Nominet, who already run the critical services that keep *.uk* on the internet and have an established track record of running national scale DNS services. Nominet's function includes running the UK's authoritative DNS service, which will be useful later.

5.2 The public sector DNS

The public sector DNS service actually consists of two services, one on the internet and one on the Public Services Network (PSN), a private network[26] used only by the public sector. There's an introduction to the service on GOV.UK here [27]. Nominet have built a highly available DNS infrastructure that is scaled to be able to serve the entire public sector, in total around 1,000 organisations. While the internet-facing service is (obviously) on a public IP address, we're filtering access to the service so that only registered and approved organisations can use it. That's partly for scaling reasons and partly because we can be more aggressive in protecting public sector organisations than if we had a more general population. There are also limitations in our licenses for the commercial threat feeds. As explained in the original ACD blog post, no-one is daft enough to suggest that NCSC should be running a national DNS service for the country. There is more on how we scale out in section 8.

The service provides analytics of various sorts, including using Nominet's world leading *turing* platform, and also provides reporting directly to NCSC and organisations that are using the service. We'll be expanding those as we learn more and better understand what our customers

[25] Strictly speaking, people who use the IT systems provided by those organisations, but you get the point.

[26] Actually a set of interconnected private MPLS clouds. More details are at https://www.gov.uk/government/groups/public-sector-network.

[27] https://www.gov.uk/guidance/introducing-the-uk-public-sector-dns

need to know to help them mitigate issues.

It's worth saying that NCSC will only ever look to block cyber security-related malicious content. If a department wants to block their users looking at gambling sites, they need to sort that out themselves. Some departments already use DNS to implement their policies and others may wish to do so and we are looking at how we could merge these services.

One benefit that we didn't think about up front, but which is pretty obvious in retrospect, is that we can understand what technologies are in use in which organisations to a large degree. For example, if we see DNS queries for *update.ianssoftware.com* it's likely that there are machines in that department running something from *Ian's Software Inc.* We're looking to do analytics on this data to help us target communications and action when vulnerabilities are announced.

5.2.1 Protection using DNS

The first way the service provides protection is to block access to known bad domains. If we think the domain is malicious or likely to cause harm, we give a different answer to the department - either saying it doesn't exist at all or replying with an address under our control. The latter can be useful for some types of threat. Our particular implementation has some limitations in that action is taken for all users of the service - we can't do specials for particular departments. That means any decision whether to block has to be weighed against the potential impact on all users.

The service consumes a number of threat feeds, some open source, some commercial, some from NCSC and uses those to decide whether to block a particular query. We analysed a number of non-NCSC feeds before selecting the ones we have chosen. These were chosen to minimise overlap between feeds and to get some assurance of quality of the data. The majority of the feeds we're using at the moment are malware-related, rather than phishing, mainly for volume reasons. We want to sort out the business change piece that needs to go with this service using alerts that are relatively limited in number. If an organisation comes onto the service and immediately gets thousands of alerts, that's not a great user experience. So, we're starting with threats that are likely to cause harm that aren't mitigated in some other way to start with. We'll move onto other threat types as the service (and its customers) settle down.

The other way the service provides protection is by mining the resolution data after the fact and looking for weird things that aren't *a priori* bad. Domain Generation Algorithms (DGA) are a way for malware authors to make it much more difficult for defenders to take down their command and control servers. The malware would have a way of calculating (say) 5,000 random-looking domains to use each day and the criminal would pick a small number to use on a particular day. Taking these down is hard because you either have to pre-register all the domains, which is expensive, or invoke a formal takedown process for each active domain which often takes longer than the time the malware uses it. However, blocking known DGAs is relatively simple as you know the method the malware will use to calculate the domain names on a given day. So, you can pre-calculate all the domain names and spot them if they're in use. Seeing DGA names is a good indicator that there's active malware somewhere in an organisation and using your DNS to block them can mitigate the harm until the malware is removed.

But what about unknown domain generation algorithms? One of the analytics run on the names requested is Markov analysis. This particular statistical test indicates whether a set of names requested by an organisation are likely to be part of a DGA, even if we don't know exactly what it's related to. The same analytic will also pick up DNS tunnelling, which is a mechanism malware authors sometimes use to get data out of a well protected network.

There's a plethora of other analytics we want to build into the service over time. Some will be very simple. For example, consider the effect of blocking resolution of a name the first time it is ever requested. A name[28] being resolved for the first time across all public sector will be an interesting event and probably warrants investigation. Others will be much more complex and they'll all evolve as both our understanding of the networks in our customers and the adversaries evolve. For example, the *turing* platform can alert to a change in query behaviour which might indicate a malware infection beginning to spread through a network.

5.2.2 The feeds

Before deciding which feeds to adopt into the protective DNS service, we spent time characterising a large number of feeds and understanding their relative strengths and weaknesses. In the end, we ranked feeds according to uniqueness, descriptiveness of block reason and the apparent quality of the data and process for curating it and then bought as many as we could afford. We then did some experiments where we logged what could be blocked but still resolved correctly. This showed that naively blocking on the basis of a single threat feed can be dangerous. We have worked with our threat feed providers to really understand what they mean by particular confidence levels or maliciousness scores so that we can give confidence to our customers about what we block and why. We've also made sure that we have global whitelists to ensure that popular services and particular infrastructure domains, like software update mechanisms and major content delivery networks, don't get blocked.

5.2.3 Adoption and service statistics

The internet-facing DNS resolver launched properly in May 2017. As of the end of December 2017, we have 121 unique organisations actively using the service as their primary resolver. Depending on how you count, that's probably 10%-15% of the total we'd like. The service is now fully live and so we're hoping that adoption rates will increase in 2018. We do note that some System Integrators (who run some departmental IT) are trying to charge ridiculous amounts of money for what is essentially a single configuration change. If that turns out to be a systemic issue or one that isn't simply resolved by departments, we'll be looking to do something to fix that (and it probably won't be friendly).

Since we are continually on-boarding customers to the service and those customers are of very different sizes and have very different DNS usage profiles, trends over the course of the year are not useful. Instead, we provide data in table 12 for the last few weeks of 2017, with the obvious drop off of traffic - and associated blocking activity - at the end of the year. Of the total queries received, only 0.04% are IPv6 and almost all of this is test traffic. The *Total Blocks* column details the total number of queries that were actively blocked because we believe the query implied malicious activity of some sort. The *Unique Blocks* is the number of unique names that caused the totality of the blocks. The *Total Blocks* is useful as a measure of ongoing malware-related activity in a network, however *Unique Blocks* probably gives a more realistic view of the security benefits. This is because one block of a particular nature could cause a disproportionate number of blocking events - for example very chatty malware (which calls out often), a very good deceptive domain (where users keep refreshing) or a false positive. Additionally, blocked responses may generate a higher than normal level of requests as the time to live (TTL) of a blocked record is set by our DNS resolver to 5 seconds. In most cases, records that we do not block have a TTL of minutes or hours. As such, clients may more frequently request blocked records as each record times out of their cache more quickly.

[28] A name that's not on an obvious whitelist.

Date	No. of Customers	Total Queries	Total blocks	unique blocks
28 Oct - 3 Nov	95	931,362,079	86,973	6,280
4 Nov - 10 Nov	99	990,383,153	216,487	5,810
11 Nov - 17 Nov	105	1,071,878,213	281,993	5,562
18 Nov - 24 Nov	106	1,129,390,734	397,575	6,087
25th Nov - 1st Dec	115	1,131,271,855	378,004	5,889
2 Dec - 8 Dec	117	1,233,632,389	273,239	5,768
9 Dec - 15 Dec	120	1,136,316,409	494,027	4,499
16 Dec - 22 Dec	121	1,120,468,992	314,601	1,858
23 Dec - 29 Dec	121	732,559,201	124,188	600

Table 12: Statistics from the Public Sector DNS Resolver

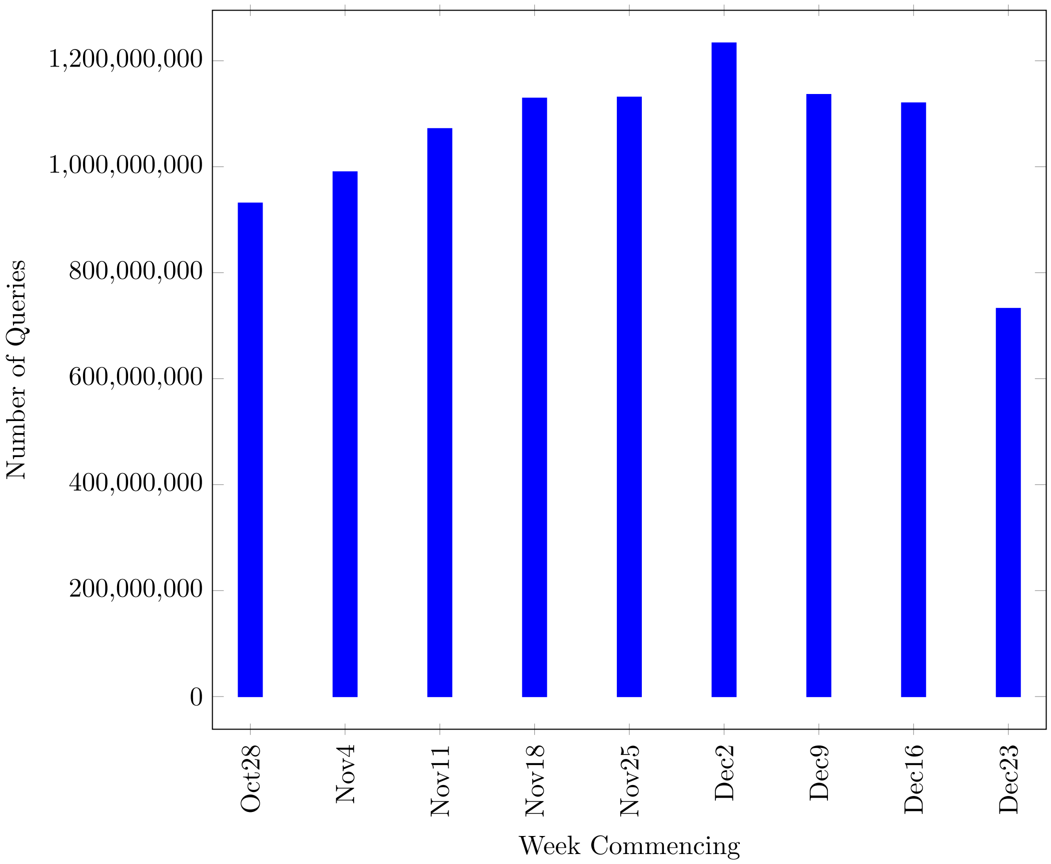

Figure 22: Volume of queries received by public sector DNS service

5.3 Security outcomes

5.3.1 Misconfigurations

One of the first things we saw when on-boarding customers onto the public sector DNS was locally scoped queries being sent to our resolver. These locally scoped queries are of form *3.2.1.10.in-*

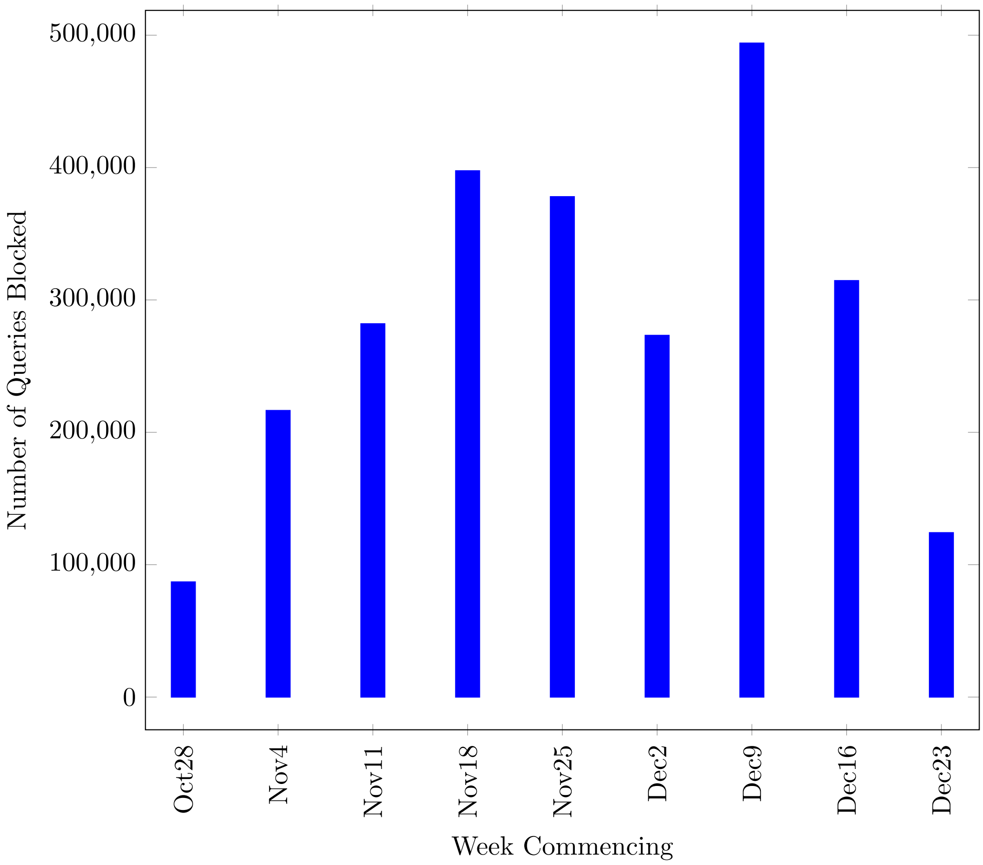

Figure 23: Total resolutions blocked by public sector DNS service

addr.arpa,which would be from a machine trying to resolve something about the internal only address *10.1.2.3*. This *10.x.x.x* range is one of the ranges reserved in RFC1918 that should never be routed to the internet. These queries potentially leak intelligence about the internal configuration of a network that could be useful to an attacker in planning their work. Given the number of organisations that leak these queries when they onboard to the DNS service, we believe this could be a very common misconfiguration of enterprise networks. We suggest that enterprise network administrators check their DNS logs to ensure that their systems are not leaking locally-scoped queries.

5.3.2 Direct blocks

Currently, we are blocking malware-related domains. These could be malware download, command and control or other sites that are involved in the distribution and control of malware. We have not yet started blocking phishing domains (and similar) but intend to do so as we and our customers become more confident in the service and that we are not over blocking, a risk with all filtering systems. Even so, much of the harm will be mitigated because of the blocking of any malware-related domains that the phishing attacks may need.

Our feed of bad domains comes from three commercial feeds (Commercial 1, Commercial 2 and Commercial 3), two non-commercial feeds (OS1 and OS2) and two feeds from NCSC (NCSC1 and NCSC2). Volumes of blocks attributed to the individual feeds is shown in tables 13 and 14. Feed NCSC1 was reset in mid-December for ease of curation going forward. Feed NCSC2 is manually added to by NCSC analysts. Blocks are attributed to every feed in which they appear, so adding the per-feed totals in tables 13 and 14 will not provide the number of blocks detailed in table 12. However, there is relatively little overlap in the various feeds, as can be seen by comparing the unique names blocked in totality and the total across all feeds.

It is important to note that the service is not intended just to block, although that harm mitigation is a primary benefit. It is also intended to help public sector organisations fix problems. At the moment, the service records very limited information when identifying potential problems. Recording and optimising blocks is important, but future intent is to ensure that automated incident alerting and resolution is better addressed by the service itself (rather than the NCSC staff).

Date	Feed Name	Total Blocks	Unique Names
Nov4-Nov10	Commercial 1	203,204	1,519
	Commercial 2	21,645	453
	Commercial 3	17,205	157
	OS1	18,608	1,339
	OS2	2	2
	NCSC1	38,783	2,723
Nov11-Nov18	Commercial 1	249,677	1,502
	Commercial 2	33,548	1,159
	Commercial 3	22,880	154
	OS1	18,446	1,314
	OS2	3	3
	NCSC1	38,252	2,828
Nov19-Nov26	Commercial 1	350,578	1,476
	Commercial 2	37,702	1,371
	Commercial 3	34,370	158
	OS1	16,134	1,340
	OS2	836	6
	NCSC1	48,260	3,225
Nov 27-Dec 2	Commercial 1	315,596	1,035
	Commercial 2	57,314	1,439
	Commercial 3	25,617	194
	OS1	17,636	1,342
	OS2	13	5
	NCSC1	47,141	3,282

Table 13: Blocks per week by threat feed (November)

Date	Feed Name	Total Blocks	Unique Names
Dec 2-Dec 8	Commercial 1	221,650	1,196
	Commercial 2	19,610	1,397
	Commercial 3	17,656	201
	OS1	17,761	1.337
	OS2	19	6
	NCSC1	38,076	3,175
Dec9-Dec15	Commercial 1	660,409	1,071
	Commercial 2	29,165	1,459
	Commercial 3	7,199	107
	OS1	15,919	1,354
	OS2	3	3
	NCSC1	32,918	2,080
	NCSC2	642	1
Dec16-22	Commercial 1	281,566	831
	Commercial 2	24,074	821
	Commercial 3	7,067	118
	OS1	11,649	944
	OS2	38	3
	NCSC1	260	3
	NCSC2	1,281	1
Dec23-29	Commercial 1	104,203	497
	Commercial 2	5,964	50
	Commercial 3	15,523	56
	OS1	8	2
	OS2	7	2
	NCSC1	84	1
	NCSC2	1,123	1

Table 14: Blocks per week by threat feed (December)

5.3.3 DGA detection

A DGA detection implies that malware is already running in the department's network and is trying to call out to its criminal owner. A set of queries that would match the DGA detector is shown below.

- lwciztdl.biz
- ffxuxfkxaj.net
- uvdesfvs.info
- kkojupb.cc
- jnsuzak.cc
- lmhdatzftv.org
- ztzbjkqrfqn.org
- bovmzfhx.net
- usygtpuvjvv.com
- ziusdawj.cc

It is obvious from inspection that these are not likely to be typed into a browser by a user and are machine generated. We use a variety of open source DGA feeds and also our commercial feeds to classify the malware family that is likely related to a set of queries that match the DGA classifier. In the case above, this set of queries is related to the Conficker malware, which was first seen in November 2008. In the period of 1st November 2017 to 31st December 2017, we have observed DGA and command and control traffic that suggests customers of the DNS service are infected with Wannacry, BadRabbit, Ramnit and Conficker. The ability of customers to find the individual infected machines within their networks varies greatly. This is something we must address going forward, providing more detailed help to those that require it, either directly or through our industry partners.

We have also discovered DGA-like resolution requests that do not match any known DGA. These are currently under investigation but it is possible that these are DNS tunnelling artefacts. DNS tunnelling is a technique that uses the DNS communication channel as a way of exfiltrating data from an otherwise well-secured network. It is also used by some tools to subvert network monitoring and filtering products. Whether this turns out to be malware, user-installed software to get around corporate network monitoring or something else, it is worthy of investigation. For example, in some cases, we have seen anti-virus software use DNS to look up hashes of unknown files against a cloud-based database. These queries have all the properties of a data exfiltration event, but obviously should not be blocked.

5.4 Conclusion

Even with the small number of customers and only part of the overall feed of malicious domains, the public sector DNS has already proved its value. In the last two months of 2017 alone[29], it has blocked over 2.5 million malicious resolution requests driven by direct blocks from our feeds. It is possible that some of these are not directly related to real compromise, for example

[29] This time period obviously includes the Christmas holidays where we see a significant slowdown in traffic.

security staff doing research or security products resolving malicious domains as part of their normal operation, but that is likely to be a relatively small proportion of the blocks. Most organisations using the service have benefitted from the blocking of requests for bad domains. On average, approximately 1 in 6 of the organisations using the service have had some security issue discovered that required investigation and remediation. As the number of customers scales up and the threat feeds are expanded to cover other threat types, we expect the protective effect to be even more pronounced. Over the coming months, we will be more proactively analysing the blocks in order to see what further information we can pass back to customers of the DNS service to help them understand their cyber security posture.

In the last six months of the year, the domain generation algorithm analytic has discovered that ten customers have had confirmed malware infections and that at least two have suspicious-looking DNS traffic leaving their networks that needs investigation.

Nominet also runs the authoritative server for the *.uk* Top Level Domain. The analytics that can be performed on the authoritative resolution data can provide a different sort of intelligence. For example, analysis of the geography of requests for a given UK name may give pre-warning of a DDoS attack against the service referred to by that name. Over the coming months, we intend to further explore how analytics on the various DNS data we and our partners can access can help better secure the UK as part of the Active Cyber Defence ecosystem.

6 Signalling and routing

6.1 IP routing and BGP

There are two major security problems caused by the way traffic is routed on the internet. The first is the triviality of using compromised machines to generate traffic from spoofed IP addresses as part of a scaled DDOS. The second is arbitrary rerouting of internet traffic at large scale - normally through an adversary announcing a very low cost route to a particular destination. Think of these as source address spoofing and destination address spoofing.

6.2 Destination address spoofing

BGP is the protocol used by internet service providers and carriers to describe how traffic should flow around the internet. From a security point of view, the majority of implementations are pretty terrible. Networks (in BGP speak, *Autonomous Systems*) advertise routes and associate a cost with those routes. Autonomous Systems *peer* in order to exchange this routing information and be able to route traffic for each other. Autonomous Systems may announce something like 'I can get you to the IP subnet 25.25.0.0/16 for a cost of 3.'. Other Autonomous Systems will make similar announcements and a router will make a least cost decision to decide how to send packets to that IP subnet. Routers build *paths* based on these routing decisions.

There have been lots of attempts to create a 'secure BGP' but they all seem to fail. We believe that's because the attempts so far have been very complicated, requiring a public key infrastructure that interoperates between international ISPs or other expensive infrastructure. The real issues with BGP appear to be that there's little incentive for ISPs to change as there's no direct commercial benefit for fixing the problem and that your peering partners can be your attackers, so the concept of trust is difficult in this space. Furthermore, changes to the BGP ecosystem must be done carefully - bad BGP configuration can drop entire networks or countries off the internet. However, we believe we have a viable approach.

We believe that better management of BGP peering relationships will significantly help in ensuring traffic is not rerouted trivially. We will establish an open community of UK ISPs whose use of BGP will be process-driven and predictable. This will both improve the robustness of the UK's internet and make false, disruptive announcements easier for everyone to detect. Our initial design for this includes:

- Better grooming of BGP announcements and path updates onto specific physical interfaces, which will better represent the peering arrangements;
- Applying some simple and sensible constraints to paths and
- Implementing Unicast Reverse Path Forwarding (uRPF) at scale.

uRPF is a way of ensuring that traffic can enter a network only if it comes down a path that makes sense to the receiving network. There are likely other technical requirements that will need to be addressed as we begin to change the BGP ecosystem in the UK, but our intention is to better use existing standards, rather than try to design something completely new.

To really tackle the BGP hijack problem, we first need to be able to measure what's going on. We're working with BT to build a community BGP monitoring platform that will be free for all UK ISPs to use so that we can collectively run some analytics on the collected routing path data. We've written analytics for all sorts of things that could happen in the world of BGP and we're expecting all the big BGP talkers in the UK to use this platform to help the community understand what's really going on - both in terms of attacks and how their peering relationships actually work. The Network Security Information Exchange (NSIE), inherited from CPNI[30], is a fantastic group where really collaborative security work gets done by the telecoms sector. They've set up a BGP working group to take this forward. We've collectively got a reasonably good idea what we're going to do, but we need the monitoring platform first. Once it's working, we'll start talking about what it's finding.

BGP is one of those things that you can't just do to yourself and expect everything to be OK. To get the right sort of protection, you need to get other big Internet Exchanges and big telco and transit providers to do something similar. Our intention is to publish the standard we end up with and the data to show the effect it has and then strongly encourage others to implement. There may be other ways of incentivising people, as discussed at the end of this document.

6.2.1 Source address spoofing

DDOS attacks generally work by creating a massive volume of traffic and pointing it at a service. A good proportion of them make it harder for victims to block the bad traffic by spoofing the source IP address with a single compromised machine pretending to be a large number of machines. If ISPs filtered traffic coming into their networks from the edge - something called ingress filtering - then we could make it harder to spoof source addresses. Luckily, there's an internet standard called BCP38 that explains how to do just that. If implemented correctly on UK networks it would make it much harder to generate lots of traffic from UK networks that can't be easily filtered. To be totally clear, this would make it harder for UK based infrastructure and machines to be used to DDOS others. It has no direct effect on whether UK services can be subject to a DDOS attack - that takes the scaling that we talk about later. It is worth noting that there are many different types of DDOS and this work will not affect them all. However, it should reduce the volume and allow service owners to invest in a more service specific way.

The great news is that most of the big UK ISPs have told us they've now implemented BCP38. The even better news is that the Centre for Applied Internet Data Analysis (CAIDA[31]) runs a programme to generate data about which ISPs allow spoofing. You can see the latest reports on the Spoofer project website here[32]. The latest reports on UK ISPs are listed here[33]. At the time of writing, the UK looks pretty good according to that page - there's only a small number of ISP netblocks that allow spoofing.

However, the CAIDA data is limited by the clients that are run by users and ISPs to generate the data. At the moment, it looks like the project tested about 1,700 /24 IPv4 blocks in the UK in 2017, or about 435,000 IPv4 addresses. The RIPE database suggests the UK has about 100 million IPv4 addresses, but Ofcom say it's about 70.5 million addresses, with 84% reachable[34]. So, according to Ofcom, the maximum number of reachable IPv4 addresses we could test is about 59 million. So, at best, CAIDA are testing about 0.74% of the UK IP space each year. It would be useful to get more complete coverage of UK networks so we've got more confidence in the quality of the data we're working from. To that end, we're going to ask ISPs to run the CAIDA *Spoofer* on their own networks to ensure that spoofing is harder. Running it frequently will make sure that ISPs notice if other changes in the network affect the ability of their networks to be used for source address spoofing. Any readers of this paper who are willing to help us generate better data on the spoofability of UK IP space can download and run the *Spoofer* tool from the CAIDA website.

[30] The Centre for Protection of National Infrastructure, part of MI5. The cyber function moved into the NCSC.

[31] www.caida.org

[32] https://www.caida.org/projects/spoofer

[33] https://spoofer.caida.org/recent_tests.php?country_include=gbr&no_block=1

[34] We're looking into the discrepancy.

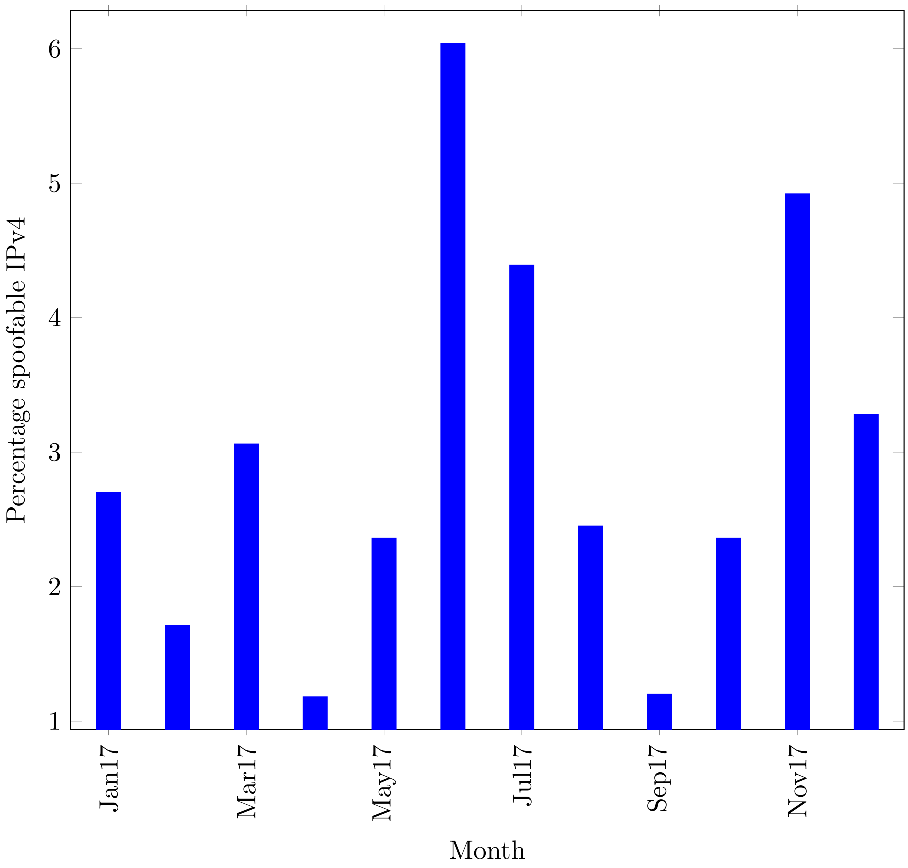

Figure 24: Percentage of spoofable IPv4 prefixes in the UK over 2017

With thanks to the CAIDA Spoofer Project[35], figure 24 shows the trend of spoofability of UK IP space tested by the project over 2017. Obviously, this is a very small sample of all UK UP space, but suggests a broadly static distribution over the year, implying that there isn't much proactive work on these netblocks going on to make this better at the moment. Obviously, more data will make this more useful, and we will need to take into account any effect of increased deployment of the *Spoofer* tool on the statistics, but this should allow us to make statements about our interventions over the coming year and their effects with some certainty.

As described in section 1.2, criminals compromise machines for a number of reasons, but one of those reasons is to use the machine as part of a distributed denial of service attack, which they can charge other criminals for. If UK networks are configured to make it harder to generate traffic from spoofed sources, the utility compromised machines in the UK have as traffic generators goes down and so the value of those machines to criminals goes down. We expect that making machines in the UK less able to participate in a scaled DDoS will demotivate attackers from targeting machines in the UK.

[35] This material is based on research sponsored by the Department of Homeland Security (DHS) Science and Technology Directorate, Homeland Security Advanced Research Projects Agency, Cyber Security Division (DHS S&T/HSARPA/CSD) BAA HSHQDC-14-R-B0005, and the Government of the United Kingdom of Great Britain and Northern Ireland via contract number D15PC00188.

6.3 SS7

Signalling System No 7 (SS7) is the protocol by which international telecoms networks talk to each other in order to route calls, send SMS and allow users to roam between countries. It was originally created in 1975 and has undergone little fundamental change since then. SS7 has no real security built in and given how the telecoms sector has evolved it is now trivial to exploit SS7 weaknesses. Exploiting those weaknesses can allow an attacker to geolocate a user's phone, reroute SMS messages and voice calls so that they can be intercepted, get networks to release encryption keys and other nefarious actions. It is impractical to expect a change in the standard for SS7, but we believe we can better protect users of UK networks from these sorts of attacks while simultaneously ensuring that the next generation telecoms signalling protocol (DIAMETER) is better secured.

Guidance on the secure use of international telecommunications standards is published by the GSMA. The NSIE previously worked with the GSMA on baseline security for SS7 which resulted in published guidance to operators (GSMA document FS.21) that details some simple filtering that a telecoms operator can undertake to better protect their users. We are working with UK operators to implement the GSMA guidance, along with two extra measures which are mentioned elsewhere in standards documents, namely enforce SMS home routing and don't accept your own global title. We are in the process of setting up some independent testing to see what security is really like on the SS7 terminations in the UK - are the measures we have asked for actually in place and, if so, do they actually have the protective security effect we expect? We've got to be careful in the testing because SS7 really is critical to the operation of the networks and so there is risk. Working with the UK mobile operators, we intend to publish the data once the testing is done, although obviously we won't immediately publish if the testing exposes vulnerabilities in the UK operators.

6.4 SMS

As part of the work with the telecoms community, we have started some work that should enable better control of how SMS messages are handled. In particular, it should be possible to make it so that only certain entities could send SMSs with a specific Transport Path Originating Address (TPOA) often the alpha tag used as the 'from address' for an SMS. Our friends at HMRC have already run an experiment with one of the mobile networks and some of the big SMS aggregator companies. They're initial results are very encouraging, seeing a 90% drop in customer complaints around the spoofing of specific, HMRC-related alpha tags on SMS. We're now looking to see how we'd scale and automate that protection into a proper service that can be used by businesses. The idea would be to have a canonical list of protected TPOAs published and ensure that they are really protected. Advice then becomes 'If you get an SMS from HMRC, it's OK; any other address and it's not' which is probably better than the advice today which boils down to 'guess whether this SMS is real'. We are also looking to link the SMS-based phishing work with the takedown work described in section 2 so that even if SMSs are delivered, we can help reduce the harm they cause. While we can probably get some pretty good protection pretty quickly, making this protection really robust probably requires the SS7 hardening to have been done, in particular the home routing changes. This would allow both aggregator level protection and operator level protection.

6.5 Conclusion

We have little hard data to publish at the moment on these projects for this report, but we should be able to generate some in the future. Since we announced the ACD programme, we've been told multiple times by multiple people that fixing BGP, SS7 and trying to secure SMS is impossible. We believe we've done enough work to be confident that we can achieve much of what we've promised, provided the telecoms sector continues to work collaboratively with us.

Our work on securing SS7 in the UK is a strategic piece of work. The testing we intend to do this year will give us some objective data on the actual security of the UK's SS7 terminations. It will almost certainly discover issues that no-one is currently aware of and we will have to collectively work to fix them. This is the first step on the journey to fix the 40 year old problem we've been living with.

HMRC's work on providing a protected SMS capability is groundbreaking and proves that - with the right technical environment and commercial model - we can make the adversaries' job much harder while making the public's life much better.

The work the NSIE has done with us on making BGP and routing better is coming to a conclusion and we should have a standard that has been designed with input from ISPs and so it should be commercially viable to implement. The ISP community monitoring platform will help us - as a community - really understand the effect of the interventions we want to make and help us make the case for broader change.

7 Threat-o-Matic

The *Threat-o-Matic* is the central hub that links the various Active Cyber Defence capabilities together, making them more than the sum of their parts. The idea of the Threat-o-Matic is to enable individual services to communicate to enable processing or action to be taken in multiple places to aid analysis and defence in depth. The cartoonish name is deliberate - it's part of demystifying cyber security and being honest about what we're doing. The Threat-o-Matic isn't just about sharing cyber threat intelligence - there are any number of existing ways of doing that already. It is a platform for service data with an emphasis on automated analysis, feedback tracking and workflow automation, but one of the things that can be pushed around the system is threat intelligence.

Consider the diagram in figure 25 which was part of the initial Active Cyber Defence blog published on the NCSC website here[36].

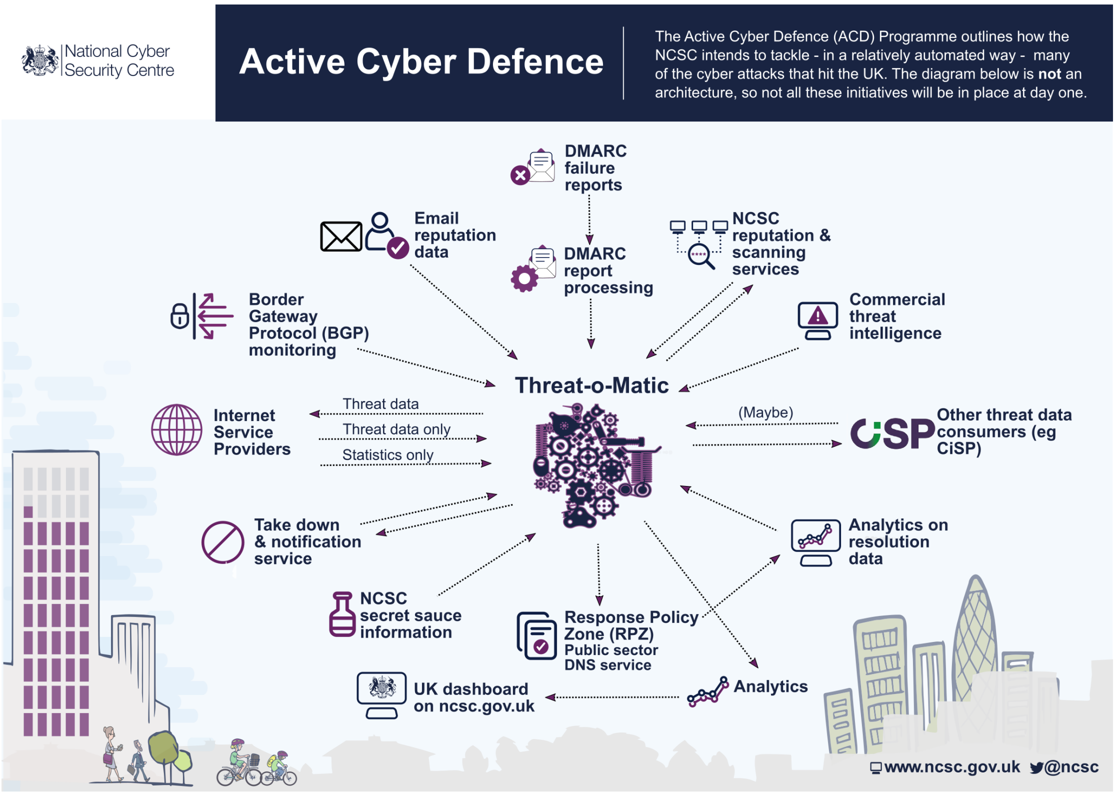
National Cyber Security Centre
Active Cyber Defence
The Active Cyber Defence (ACD) Programme outlines how the NCSC intends to tackle - in a relatively automated way - many of the cyber attacks that hit the UK. The diagram below is **not** an architecture, so not all these initiatives will be in place at day one.
DMARC failure reports
Email reputation data
DMARC report processing
NCSC reputation & scanning services
Border Gateway Protocol (BGP) monitoring
Commercial threat intelligence
Threat-o-Matic
Internet Service Providers
Threat data
Threat data only
Statistics only
(Maybe)
CiSP
Other threat data consumers (eg CiSP)
Take down & notification service
Analytics on resolution data
NCSC secret sauce information
Response Policy Zone (RPZ) Public sector DNS service
UK dashboard on ncsc.gov.uk
Analytics
www.ncsc.gov.uk
@ncsc

Figure 25: The Initial ACD Ecosystem

Consider the public sector DNS service blocking an organisation going to a malicious domain. The service would pass that fact to the Threat-o-Matic detailing the domain that was blocked and why, along with some other metadata - but importantly not the identity of the organisation that caused the block. In this example, we'll assume that the link out to the ISPs and the takedown service get told of this event. The notification to ISPs would arrive automatically and they would decide, automatically, whether to take action and block that domain for their

[36] https://www.ncsc.gov.uk/blog-post/active-cyber-defence-tackling-cyber-attacks-uk

customers. They'd likely want to tell us whether the data was actually any use or not. If the event said that ISPs should block *www.bbc.co.uk* you could imagine a very negative score. If blocking the domain really protected users, then positive feedback would follow. The takedown service would follow the same pattern - if the takedown was successful then good feedback on the event would follow, if the hosting provider rejected the takedown request because it wasn't a malicious site, then negative feedback would follow.

Another potential flow is the DNS service analytics finding that a DDoS against a UK name is likely to happen soon. That event could be passed to the BGP monitoring platform to determine which ISPs are serving the IPs hosting that site and allow them to take proactive measures to limit the impact. We have not yet done any experiments in this area, but the theory seems sound.

It is impossible to predict the full panoply of services that will form the ACD ecosystem over the coming years, nor the new and inventive ways both attackers and defenders will employ. Therefore, the Threat-o-Matic needs to be extensible and self-governing with a mechanism for automatically limiting the impact of a bad source or useless defensive technology. The service-to-service feedback principle is absolutely key to this.

7.1 Threat-o-Matic design

Unfortunately, the development of the Threat-o-Matic is not as far advanced as we would have hoped, but is a priority for the next year. However, we have a design for the system and minimum viable proposition implementation.

The design is effectively an event bus, with traditional publishers and consumers, an event stream manager and enrichers. Each event carries with it its provenance, a serialised view of the history of this event. Consumers rate publishers, with positive cyber security outcomes providing a positive score. This will allow us to automatically generate insight into the value of individual relationships between entities in the ecosystem. We may add publishers rating consumers in the future.

Events are really JSON objects encoding the path the event has taken as it traversed the Threat-o-Matic ecosystem with the payload encoded in STIX2 or raw JSON[37]. The use of open standards ensures that we have a library of tools - both commercial and open source - that will interoperate with this service, but also ensures that exchanging data with other, similar systems should be easy. The current prototype uses the AWS Simple Notification Service as the event bus backbone, but that's certainly not a dependency. The initial test scenario was automated processing of DMARC reports. The prototype took DMARC forensic reports and put them on the event bus as a *email.dmarc.report* event. The email phishing detector publisher - a machine learning classifier to determine whether a DMARC report is related to real phishing - consumes that event and, assuming it is related to phishing publishes a *email.phishing* event. This is consumed by the email phishing takedown publisher which extracts the relevant data and encodes them in a *email.phishing.takedown.request* event and also uses the Netcraft API to start a takedown. The Netcraft API is polled to monitor the status of the takedown and, when it is successful, an *email.phishing.takedown.success* event is published. This is consumed by a feedback publisher that correlates requests and successes and publishes a *feedback.email.phishing.takedown* event, which is consumed by the system scoreboard. All events are archived for later analysis and multiple entities can consume each event, so each event can cause multiple actions in the real world.

Towards the end of 2017, a piece of joint work between Nominet, BT and NCSC built a prototype adapter to connect detection events from the Public Sector DNS service to the BT

[37] The platform is agnostic to actual payload of events and we've decided on these for now.

MISP platform, described in section 8. This remains a prototype but is actively pushing detection events from the Public Sector DNS to the BT MISP platform and from there out to other connected ISPs. More details are in section 8.

7.2 Next steps

The Threat-o-Matic is currently an early discovery platform, but early experiments have been successful and have proved the initial high level design. The next year will see user needs capture and a move of the service from a discovery prototype to a beta quality service with a number of services implemented. It is likely that initial services will concentrate on processing DMARC reports of various types for action (rather than analysis which is what Mail Check does) and linking malicious detection events to the ISP community.

8 Scaling the effect

The Active Cyber Defence programme is intended to provide enhanced protection for the whole of the UK and, in time, more widely, not just the public sector. However, most of the work exposed in this paper is around protecting the public sector. We have a mantra in the NCSC of 'eat your own dogfood', that is we will implement anything we ask others to do ourselves first in order to prove the benefit. This paper is mainly about proving the various initiatives using government as a guinea pig. This section explains how we will seek to take those learnings and scale the positive security effects outside of the public sector in the UK and eventually outside the UK. Each of the measures we trial will need different interventions to try to scale them. There are different technical, privacy, legal and market considerations for each of these initiatives which will need to be addressed.

8.1 BGP and SS7

The BGP and SS7 work both require UK telecoms operators to do the majority of the work. We will use the power of government procurement to incentivise long term change once the standards are in place, for example government could say that it will only procure services from operators that meet specific standards. As we start to prove the benefit of the changes to the implementations, we will be working with other governments, while our telecoms operators work with their foreign partners, in order to push wider adoption of the new standards. We hope that the internet exchange community will also help with implementation. We need to have concrete data on the protective effect and proven implementations to make the case for others to follow.

8.2 Takedowns

There are two parts to our current takedown work and they will scale differently with different effects.

The work to remove malicious content from UK delegated IP space potentially affects all brands and their customers. However, as shown in section 2.2.1, the UK is responsible for relatively small amounts of malicious content used in commodity attacks. So, in order to scale the protective effect, we need other countries to take more responsibility for the malicious content hosted in their delegated IP space. We do not believe that every country should do exactly as we have in the UK, but we are calling on all countries to do more to detect and remove malicious content[38] in their delegated IP space. We are talking to cyber security agencies in other countries about how they can help make hosting cyber attacks more difficult globally. We can, to some degree, infer which countries host the most malicious cyber content from the feed that Netcraft builds. Observing how global hosting changes over the coming year will be interesting.

The second part is treating HMG as a large scale public brand and making it more difficult for attackers to abuse the trust the public has in it. We believe the work has shown that this is a useful thing to do and objectively reduces the harm taken by the public. HMG brands are not the only ones that are useful to attackers, though. There are many brands that could have the same market penetration and monetary potential for the attackers as the primary HMG brands. In order to scale the protective effect for the public, we will be calling on all major brands and key industry sectors to take control of how their brand is used and protected online. This will start with private, bilateral discussions, but again the data is available to allow us to determine where attackers find the most value.

[38] Strictly in a cyber security sense.

8.3 DMARC

Our intention is to have all of public sector using DMARC on internet domains as quickly as possible. This will allow us to understand potential blockers to implementation, the benefits of implementing DMARC and build tooling that will help system owners long term. Once that is done, we will be pushing brands that are well known and entire industry sectors to follow suit. We are already in discussion with other governments about what they can do to better protect their online presence. Our MailCheck tool will always be freely available and open source to help implementers.

There needs to a stick as well, though. DMARC policies are, by their very nature, public. They are published by organisations in publicly accessible DNS records so that other internet infrastructure can locate them. That means that it's relatively simple to determine whether an organisation correctly uses DMARC and so it would be relatively simple for us to measure adoption by sector - for example, determining which retail banks properly protect their email domains from spoofing. As adoption continues, we may end up publishing that data and explaining to the public how adoption of DMARC could be seen as a proxy for how much a company cares about its customers.

8.4 Web Check

Web based vulnerability scanners already exist. Many of them do much more intrusive scanning than Web Check does. However, Web Check wasn't built to be the best technical web vulnerability scanner. It was built as a tool to help organisations that currently have no knowledge of what scanning means or what to do with the results to get better at handling things like this. Only then will its capability be scaled up. Technically, it would be simple to allow Web Check to be run on any domain. However, establishing ownership then becomes critical and would be much more difficult than the mechanism we use today for the public sector. There are many policy questions about how Web Check should be scaled out to parts of the UK outside of the public sector. For example, should we allow charities to use Web Check? Should it be all charities or just those with fewer than 100 people? Should NCSC run the service or allow (say) the Charity Commission to white label the service? What would this do to the market for more intrusive web vulnerability scanners?

These, and other, policy questions will need to be answered before we decide whether and how to scale Web Check.

8.5 Public sector DNS

The Public Sector DNS is providing protection for an increasing proportion of the UK public sector, but the main thrust is to get adoption up significantly from the current approximately 10%. The next year will see the adoption of the service on the PSN at scale, plus ongoing trials with the NHS. We also need to dedicate effort to optimising the service itself, in particular the automated incident reporting function. In order to scale that protection to other parts of the public sector that aren't already covered, we are working with various providers and organisations to talk to them about setting up similar systems to protect their users, if there are valid reasons why our centralised system can't be used. Those conversations will continue and we will end up with a best-fit approach for any solution.

We do not see the 'free at the point of use' public sector DNS service being used by industry in general. Big enterprises should be taking their own security under their own control, as most of them do. There is a policy question around whether we should provide a service to Critical National Infrastructure companies and whether that should be an operational service like the

DNS or whether they can consume our threat information in their own systems. The latter is more likely. How we either provide or engender to be provided appropriate services for small businesses, charities and other sectors remains an open policy question.

The really interesting thing we said when we launched this programme was that we will work with ISPs to protect the UK public better, by default. This has attracted complaints of the 'Great British Firewall' but we hope that the explanation of what we're doing for government, coupled with how we intend to actually make this happen, will allay those fears. In section 5.2.1 we briefly discussed the benefit of aggregating threat feeds. We want to create a capability that will allow the UK public (at least initially) to be protected using an aggregation of multiple threat sources. To enable that, BT have set up an instance of the MISP threat sharing platform[39] which is free for all ISPs to use. Our joint intention is to enable UK ISPs (and those more widely should they choose to join) to share detection events between the parties. By sharing detection events, rather than whole swathes of threat information, we can limit the ability of the platform to be abused[40], while ensuring that it provides real security benefit to the end users. Our experience of commodity attacks is that they are very rarely targeted to a particular ISP's customers and so it is reasonable to assume that attacks will be randomly distributed across the ISPs and therefore early detection is equally likely in each ISP, assuming comparable technical implementations and threat intelligence sources. This mechanism also allows us to leverage the power of the disparate threat feeds that the individual ISPs consume for collective benefit, without sharing the threat feeds themselves, which would be likely be very costly due to licensing.

We have worked with Nominet and BT to build an adapter between the Threat-o-Matic and the BT MISP platform in order to exchange events. Both MISP and the Threat-o-Matic use STIX2 as the payload format and so this should be a relatively simple adapter. While the current implementation is only an alpha, we are successfully exchanging detection events between the platforms. Over the last couple of months, over 200,000 malicious domains have been pushed onto the MISP platform by BT alone, covering more than 45 threats. These have been consumed by the three ISPs currently connected to the MISP platform. While the events sent from NCSC to BT are automatically used to protect customers across the ISPs using the platform, we have more work to do on the NCSC side of this link to enable automated protection for Public Sector DNS users from the events provided by the MISP platform. This is part of the Threat-o-Matic build issues previously identified. There is now a dedicated subgroup of the NSIE to take forward this integration between ISPs.

8.6 The UK ISPs

The UK internet service providers have a key and unique role to play in helping the UK be more secure at scale. The ISP Association[41] is a key partner for us in scaling protection by default for the UK. In 2017, we wrote to ISPA and its members, asking them to consider doing the following set of measures :

- Ensure that DMARC is processed properly and that their infrastructure does not break any of the prerequisite protocols, such as SPF and DKIM.
- Ensure that it is easy for their customers to properly deploy DMARC on their customer domains.
- Implement our SS7 filtering minimum standard and the current BGP suggestions and commit to working with us on future BGP enhancements.

[39] See www.misp-project.org

[40] Abused as in 'used for something other than cyber security protection'.

[41] www.ispa.org.uk

- Properly filter management protocols (for example TR-069 and SNMP) and other potentially dangerous protocols (for example telnet and UPNP) from the internet to residential broadband by default (with an opt out for users)
- By default, protect their users from known cyber attacks (with an opt out for users). We are not specifying how this should be done or what threat feeds are used, but offer ours either directly or through the MISP platform.

We expect ISPA and its members to respond formally early in 2018. A positive commitment by this powerful group would be a significant step in fundamentally changing the security of the UK.

8.7 The UK Registry

Nominet, the UK Registry, is responsible for the *.uk* namespace. They have a similar objective to the NCSC to help make the UK safer online by making the *.uk* namespace one of the most secure in the world. They have been operating their free Domain Health service to help registrars combat cyber crime on domains managed by them since December 2016. Nominet collects, collates and distributes security information from a number of suppliers which highlights *.uk* domains implicated in abuse. A ranking and scoring system is used to nudge registrars into dealing with the issues. Nominet have seen a 10% improvement in the average scores given by the service since it began.

Future initiatives are being planned by Nominet such as detecting and mitigating phishing domains at the point of registration. It is recognised that phishing domains have a short shelf-life and can be difficult to take down once active. The intent is to make it harder for a criminal to even register these harmful domains and make the *.uk* namespace even more trusted. They are also looking to use data science techniques to discover domains exhibiting malicious behaviour that may previously have been undetected.

9 Conclusion

In November 2016, we posited an ecosystem of relatively simple services working together to objectively and measurably make the UK safer from cyber attack. In particular, we suggested that we could begin to tackle the commodity attacks that directly affect the majority of people in the UK, focusing on affecting the return on investment these attacks attract. We wish to protect the majority of the people from the majority of the harm from the majority of the attacks the majority of the time. We are not expecting our interventions to be perfect or to defend against every single type of cyber attack. However, we believe that government *actively doing something*, providing real services and generating real data and analysis has to be a first step in demystifying cyber security and beginning to tackle the impacts of cyber attack at scale.

This paper has documented the work we have done in 2017 for four of the key services in our Active Cyber Defence programme, and some of the underpinning ecosystem work to support them. For each one we have shown the effect they have had already in the UK, using the real, unvarnished data to show the effects we have seen. Each service is having a positive effect already and we will scale these services up, through more public sector adoption, and out, through the initiatives in section 8, over the coming months. As the scaling happens, the data generated will allow for more complex analysis and inferences to be drawn from the data with more confidence. We commit to continuing to publishing the code for the services, where it is appropriate, the evolution of the services as we continue to learn and the data and outcomes we are generating from them.

While the Active Cyber Defence programme is only one year old, we believe that it has already demonstrated the value of the new approach adopted by the government in the National Cyber Security Strategy. However, we are not the only organisation with good ideas and we are not the only country connected to the internet. We would welcome partnerships with people and organisations who wish to contribute to the ACD service ecosystem, analysis of the data or contributing data or infrastructure to help us make better inferences. We also call on other government cyber security agencies around the world to consider the effects of the initial ACD programme and whether they could implement similar services.

We hope that the ACD programme in general and this paper and future ones like it help to better inform people, enterprises and governments about the possibilities of national scale protection through the widespread adoption of relatively simple services that are driven by evidence and data, rather than hyperbole and fear. We believe that evidence-based cyber security policy is a possibility. If you're interested in being a part of it, look on the NCSC website for opportunities to work with us either as an employee, as part of the Industry 100 scheme or even in other ways.