



National Cyber
Security Centre
a part of GCHQ



Active Cyber Defence

The Third Year

Contents

Introduction	3
Takedown Service	5
Mail Check	18
Web Check.....	24
Protective DNS	27
Dangling DNS	39
Routing and Signalling	40
Host Based Capability.....	43
Vulnerability Disclosure	44
NCSC Observatory	45
Suspicious Email Reporting Service	47
Exercise in a Box	48
Logging Made Easy	50
MyNCSC.....	51

Introduction

The Active Cyber Defence (ACD) programme's aim 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.'

This third annual report (covering the 2019 calendar year) aims to provide transparency on these efforts, and evidence on their effectiveness. Once again, we invite interested readers to copy, distribute and challenge our analysis, and to share their feedback.

In 2019, much of the effort on ACD was maturing our existing services alongside some more experimental cyber security work. We also started to think about how we bring the services together to make them easier for our customers to use, and more efficient to run. Alongside that, we always have a pipeline of problems and solutions we're investigating.

This report includes sections for each of the services present in the [second year report](#), along with some introductions to newer work from 2019.

To highlight just a small selection of the progress we've made in 2019:

- the Takedown Service continued to result in a significant reduction in 'badness' on the internet, and adapted to the changing behaviour of bad actors
- DMARC adoption continued, we developed a better understanding of how to coach organisations through adoption, and we are continuing to keep track of email security standards
- we developed a prototype capability to detect subdomain hijack vulnerability at scale

To repeat a little from last year's report, interpreting the data provided across the various reports is difficult. It might be the case that an increased statistic (such as number of blocks) indicates success at blocking more badness. Or a reduction in the same statistic might indicate success as the blocking capability has meant attackers are focussing elsewhere. The report contains our best analysis of the data, but if you disagree with our interpretation, please let us know. We've worked to make the numbers within the report comparable year-to-year.

This year's report was developed slightly differently. Each service team wrote their own section, allowing their particular insights, lessons learned, and priorities to shine through. That was supplemented with expert oversight from the NCSC's technical writers, data scientists and data visualisation specialists to ensure accuracy and clarity. Thanks to everyone who put in so much effort during a difficult year.

As always, we need to acknowledge all of the work going on elsewhere in the NCSC, HMG, and our partners internationally and in the private sector who contribute to ACD and the protection of the UK. Cyber security is a team sport, and when we collaborate and share, we can have a bigger impact than doing things individually.

Finally, please do get in touch with any feedback on this paper, the services and approaches it outlines, or other aspects of the ACD approach. We don't have all the answers, so challenge and offers of data or research help are very welcome. You can either send it to ACDenquiries@ncsc.gov.uk or via our social media and other contact channels.

ACD at-a-glance

The ACD programme seeks to stop a range of different attacks ever reaching UK citizens, institutions or businesses. Working in a relatively automated and scalable way, it removes the burden of action from the user and enables attacks to be taken down at scale.

This report covers the following ACD services. For more information, please refer to the Active Cyber Defence website: www.ncsc.gov.uk/section/products-services/active-cyber-defence

Takedown Service (www.ncsc.gov.uk/information/takedown-service)

Finds malicious sites and sends notifications to the host or owner to get them removed from the internet before significant harm can be done. The NCSC centrally manages the service, so departments automatically benefit without having to sign up.

Mail Check (www.ncsc.gov.uk/information/mailcheck)

Helps organisations assess their email security compliance and adopt secure email standards which prevent criminals from spoofing their email domains.

Web Check (www.ncsc.gov.uk/information/web-check)

Helps owners of public sector websites to identify and fix common security issues, making sites in the UK a less attractive target to attackers.

Protective DNS (www.ncsc.gov.uk/information/pdns)

PDNS prevents users from accessing domains or IPs that are known to contain malicious content and stops malware already on a network from calling home.

Dangling DNS

Detecting subdomain hijack vulnerability, at scale, including insight into the services most commonly affected by this vulnerability.

Routing and Signalling

Fixing the underlying infrastructure protocols on which the internet is based: the Border Gateway Protocol (BGP) and the Signalling System No. 7 (SS7).

Host Based Capability (www.ncsc.gov.uk/information/host-based-capability)

Advanced NCSC threat detection capability that can be deployed to detect threats on an organisation's network.

Vulnerability Disclosure (www.ncsc.gov.uk/section/products-services/active-cyber-defence#section_6)

Services based around identifying, reporting and remediating vulnerabilities in government and other key services.

NCSC Observatory

Generating data-driven insights to underpin the NCSC's research and strategy, which includes supporting the other ACD services.

Suspicious Email Reporting Service (www.ncsc.gov.uk/section/products-services/active-cyber-defence#section_8)

Allows the general public to report phishing or suspicious emails they receive in their inboxes. The service analyses the emails for links to malicious sites, and then seeks to remove those sites from the internet to prevent the harm from spreading.

Exercise in a Box (www.ncsc.gov.uk/information/exercise-in-a-box)

A toolkit of realistic scenarios that helps organisations practise and refine their response to cyber security incidents in a safe and private environment.

Logging Made Easy (www.ncsc.gov.uk/information/logging-made-easy)

An open source project that helps organisations to install a basic logging capability on their IT estate enabling routine end-to-end monitoring of Windows systems.

MyNCSC (www.ncsc.gov.uk/information/myncsc)

The NCSC's digital platform that provides a single point of entry to ACD and other NCSC services.

Takedown Service

2019 is the third full year of the NCSC's Takedown Service. Run by Netcraft on behalf of the NCSC, the service finds 'bad stuff' hosted on the internet and seeks to have it removed, the goal being to remove cyber security threats before members of the public (or organisations) fall prey to them.

When discussing takedowns, we will talk about attacks and attack groups. The major distinction here is how we count associated URLs related to a single campaign into an attack group. An *attack* is a single URL involved in a campaign, while an *attack group* is how we refer to all the URLs that are used to launch that campaign.

2019 total takedowns

In total, 217,173 URLs were taken down in 2019. The breakdown is compared with 2018 in Table 1.

Table 1. Comparative overall annual attack takedowns, 2018-2019

Measure	2018	2019
Total number of takedowns (attack URLs)	192,256	217,173
Total number of IP addresses	24,320	21,111
Total number of attack groups (campaigns)	51,569	45,603
Median attack availability (hours)	9	11
Down within 24 hours	63.7%	64.6%

As shown in the table, the total takedowns in 2019 were distributed across 21,111 IP addresses. This is the third year in succession where the number of IP addresses hosting our takedowns has fallen, though not significantly compared to 2018. Again, the number of attack campaigns (groups) has also fallen. Last year we suggested that the reason for this might be that infrastructure used to conduct attacks is harder to acquire. However, similarly to last year, we do not have evidence that points to any systemic changes which might support this hypothesis.

Attack availability, that is, the amount of time a compromised domain can be used as part of an attack, has increased slightly, from a median of 9 hours in 2018 to 11 hours in 2019. The percentage of campaigns taken down within 24 hours has also increased slightly, from 63.7% in 2018 to 64.6% in 2019. Overall, response times are consistent in comparison.

UK government-themed takedowns

We have continued to provide brand protection for UK government departments and services this year. Again, we take a wide view to include brands that might not immediately identify as government, such as universities and TV Licensing. In 2019, we removed a total of 17,399 campaigns that used UK government branding in some way, not all of which were phishing sites. The details are shown in Table 2.

Table 2. UK government-themed attacks by type, 2019

Attack type	Number of attacks (URLs)	Number of attack groups (campaigns)	Median availability (hours)
Phishing URL	25,741	4,471	15
Phishing URL mail server	3,772	3,772	25
Malware attachment mail server	3,473	3,473	25
Advance fee fraud	2,954	2,954	2
Malware distribution URL	973	885	10
Malware infrastructure URL	977	530	10
Instagram brand infringement	348	348	10
Malware URL mail server	328	328	24
Malware command and control centre	363	209	40
Web shell	265	167	50
Phishkit archive	189	145	15

Attack type	Number of attacks (URLs)	Number of attack groups (campaigns)	Median availability (hours)
Phishkit email	97	97	20
Facebook brand infringement	40	40	23
Brand infringement	29	29	60
DKIM signed email domain	13	13	633
Twitter brand infringement	6	6	42
Malware payment URL	3	3	0.47
Phishing credential dropsite	2	2	0.32
Credential drop URL	1	1	77
Other URL	1	1	57

UK government-themed phishing

In 2019, we took down 4,471 UK government phishing campaigns, a total of 25,741 URLs. These attacks were hosted all over the world and the median availability of these attacks was 15 hours, with 60% down within 24 hours of discovery. Comparatively, this is a slight increase from the 2018 values (13 hours, 60% down within 24 hours) but response times are consistent in comparison, as shown in Table 3.

Table 3. Comparative UK government-themed phishing attack availability, 2018-2019

Measure	2018	2019
Mean (hours)	197.3	207.3
Median (hours)	12.7	14.7
Skewness	5.8	7.8
25th Percentile	1.4	1.8
75th Percentile	78.0	62.2
Down within 4 hours	37.0%	33.6%
Down within 24 hours	60.0%	60.0%

As in previous years, although the majority of attacks are taken down within the first 24 hours (60%), there are a number of takedown requests that take a long time to be actioned. Figure 1 shows the distribution of attack availability with the last 'bucket' containing anything that takes more than 200 hours to be removed.

Figure 1. Skewed distribution of UK government-themed phishing attack availability, 2019

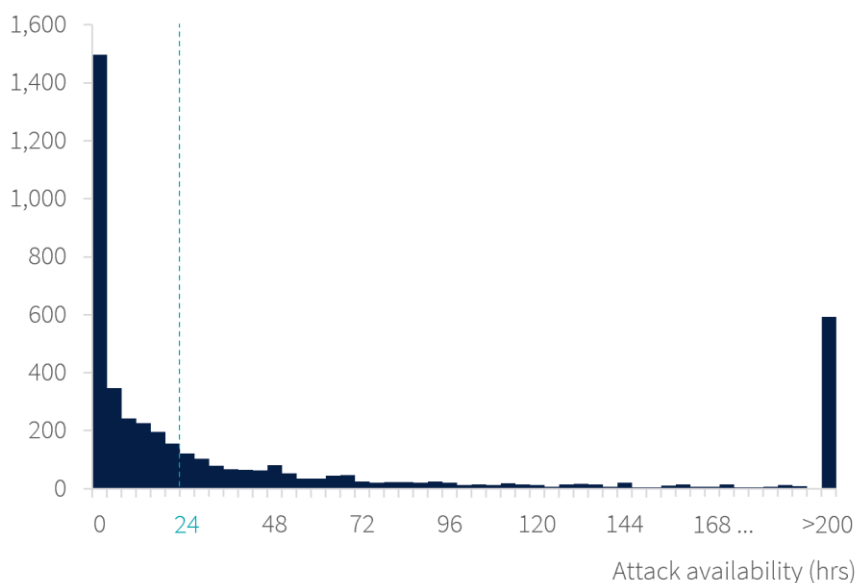



Figure 2 identifies the top 10 hosters of UK government phishing attacks. It is evident that UK government-themed phishing is relatively consistent across hosters, except for GoDaddy and Microsoft, who saw greater volatility in their monthly totals than others.

Figure 2. Top 10 hosters of UK government-themed phishing campaigns, highlighting Microsoft and GoDaddy who saw greater volatility in their monthly totals, 2019



We did see some interesting variation in how quickly hosting companies responded to abuse, with a range of monthly median attack availability for all hosters between ~6 and 21 hours (see Figure 3). Figure 3 shows the fastest responses were found to be during August and October, which coincided with peaks in Microsoft-hosted UK government phishing, where the median attack availability was just 1 hour (see Table 4).

Figure 3. Spikes in UK government-themed phishing campaigns hosted by Microsoft coincided with periods of lower attack availability across all hosters, 2019




Table 4 shows that GoDaddy remains the top hoster of UK government-themed phishing attacks, but Shinjiru Technology doubled their share in 2019 to rank second with 5.8%. Cloudflare also increased their share of these attacks, and median attack availability also increased from 5.5 to 21 hours, which could make their hosted domains a more attractive target for attackers. New names in the top 10 hosters of UK government-themed phishing include NameCheap and Alibaba Group.

Table 4. Comparison of top 10 hosts of UK government-themed phishing attacks, 2018-2019

2018			2019		
Hoster	Share (%)	Median availability (hours)	Hoster	Share (%)	Median availability (hours)
GoDaddy	17.3	53	GoDaddy	15.7	29
Endurance International	10.5	2	Shinjiru Technology	5.8	40
OVH	4.1	26	Amazon	4.7	1
Amazon	3.4	9	Cloudflare	4.6	21
Cloudflare	3.2	6	Microsoft Corporation	4.6	1
Shinjiru Technology	2.9	69	Endurance International	4	9
Webafrica	2.7	51	OVH	3.5	14
United Internet	1.5	18	Velocity Servers Network Exchange	2.9	19
Digital Ocean	1.4	49	NameCheap	2.5	20
Velocity Servers Network Exchange	1.3	15	Alibaba Group	2.3	29

As illustrated in Table 5, despite HMRC takedowns having the most attack URLs, it was in fact TV Licensing which we found had the most phishing campaigns (attack groups) during 2019.


Table 5. Top 10 UK government-themed phishing attacks by brand, 2019

Government brand	Number of attacks (URLs)	Number of attack groups (campaigns)	Median availability (hours)
TV Licensing	6,035	1,313	20
HMRC	9,990	1,277	15
Generic gov.uk	3,118	774	17
DVLA	3,872	618	18
BBC	1,169	248	1
Council Tax	542	176	13
Government Gateway	605	93	19
UK Police	39	39	3
NHS	127	36	4
UK University	198	90	4
All UK government-themed phishing attacks (total)	25,741	4,471	15

HMRC, DVLA and the rise of TV Licensing phishing

Throughout the period that the service has been operating, we noted that HMRC's ranking as a phishing target has dropped out of the top 20 most phished brands, and continues to fall largely due to the implementation of HMRC's anti-spoofing controls and the continual takedown of attack infrastructure, as shown in Figure 4.

It is natural enough to speculate that successful countermeasures against the abuse of the HMRC brand may lead criminals to start targeting other departments. It is possible that they might also target other organisations who have not implemented a DMARC/DKIM/SPF policy, or other anti-spoofing controls.

Figure 4. Variation in global phishing target rank for HMRC, TVL and DVLA, 2017-2019


We also saw a corresponding rise in DVLA related 'vehicle tax' and TV Licensing (TVL) campaigns. Notably, DVLA made several control changes relating to their email, moving from their GSi mail domains to the [dvla.gov.uk](https://www.dvla.gov.uk) domain in 2019. DVLA also implemented other controls to protect their SMS SenderIDs in April 2019. It is interesting to note that DVLA campaigns dropped significantly in mid-2019 coinciding with these controls.

UK government sending malware

Since 2017, we have been performing takedowns associated with campaigns that send emails with malicious attachments purporting to be from the UK government. Attackers do this to trick victims into installing malware. For example, the email could take the form of a legitimate-seeming communication from HMRC encouraging the recipient to open the attached tax refund form. When the attachment is opened, the malware infects their device.

Typically, malware calls out to various places on the internet to perform further malicious activity, such as linking to a Command and Control (C2) server that the attacker uses to control the malware. Installation of keyloggers, banking trojans, or other tools can then follow. By discovering the domains and URLs that the malware connects to and taking them down, we can mitigate the effects of such attacks.

During 2019 we noted consistently lower mail server takedowns for malware campaigns which used the UK government brand, as shown in Figure 5. This illustrates that we have not experienced attack campaigns using a high number of distributed mail servers to launch these attacks since September 2017.

Figure 5. UK government-themed mail server malware attacks consistently low since 2017 spike, Jun 2017 - Dec 2019

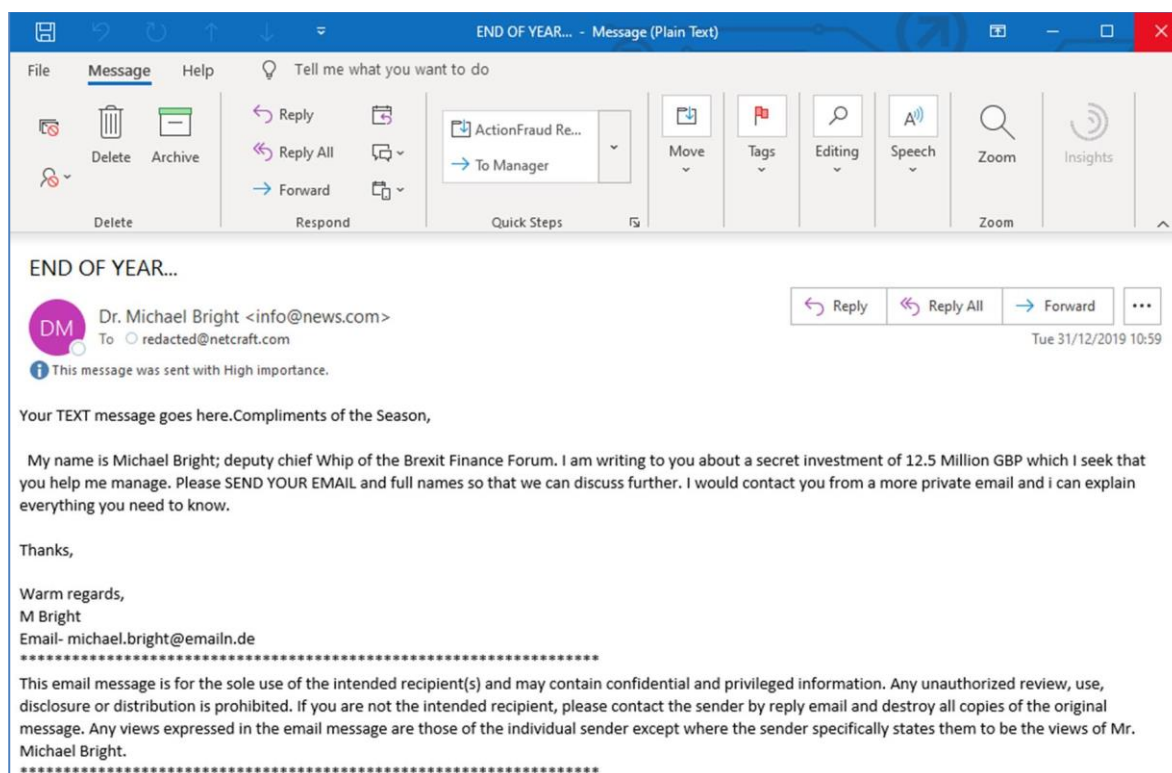
Nearly half of the campaigns we noted in 2019 that used UK government branding were associated with the Emotet family of malware. As a general trend, over 50% of the malware infrastructure taken down in this category was found to be associated with banking trojans.

UK government advance fee fraud

Advance Fee Fraud (AFF), aka the 419 scam, is not only the oldest but also one of the most well-known attacks. Despite its simplicity, it still offers a valuable return on investment with respect to time and effort a fraudster might put into such campaigns. The Takedown Service finds these attacks and seeks to have the associated mail account immediately suspended.

The numbers of AFF takedowns between 2018 and 2019 are broadly similar; however, we did see scams that used a 'Brexit investment' opportunity as a lure, seeking to capitalise on the financial uncertainty felt by many at that time, as shown in Figure 6. This is a good example of how fraudsters adapt their campaigns according to the current issues of the day.

Figure 6. Example of a Brexit-themed advance fee fraud email, Dec 2019



The National Lottery remains the most targeted brand in this category (as shown in Table 6), and it is no surprise that criminals continue to use money as a principal theme in scams of this type.

Table 6. Comparison of most popular UK government-themed advance fee fraud attacks by brand, 2018-2019

Government brand	2018	2019
	Number of attacks (URLs)	
National Lottery	1,198	1,039
Financial Conduct Authority	813	604
Bank of England	731	520
Ministry of Justice	23	392
BBC	183	98
Metropolitan Police	195	72
HMRC	79	60
Dept for Exiting the European Union	0	45

Government brand	2018	2019
	Number of attacks (URLs)	
National Crime Agency	17	49
Prudential Regulation Authority	40	37
HM Treasury	54	29
All UK government-themed AFF attacks (total)	3,246	2,954

UK government deceptive domains

Much of the work done by the Takedown Service can be described as reactive in nature; we find evidence of a bad thing, we notify the hoster of it, and then the bad thing is removed. However, there is a subset of attacks that use a deceptively named domain to appear authentic (for example <http://hmrc-tax-claim.co.uk>). By proactively monitoring for deceptive domain registrations, we can identify when their content changes from the domain parking page to a phishing form, and then act to have it taken down.

In 2019 we discovered approximately 8,700 such deceptive domains. Of that total, 99 subsequently hosted phishing content related to UK government departments or services, up from 21 the previous year, and we noted a further 500 instances that infringed on UK government branding. We will continue to monitor domains in this way as it is a useful early intervention measure.


Takedowns in UK-delegated IP space

The NCSC Takedown Service cleans up phishing and/or malware found to be hosted on IP addresses delegated to the UK. Cyber security is a team sport and introducing any control in isolation is of course limited. As always, we hope that openly sharing strategies and techniques that we find to be effective will encourage other nations to introduce similar measures, so that the benefits can be global.

Brand agnostic UK-hosted phishing

Figure 7 shows that in 2016, the UK hosted 5.3% of all phishing attacks. By 2018, this share had fallen to 2% despite a rise in phishing attacks globally during the same time frame. In 2019 we took down a total of 18,202 phishing campaigns hosted on UK IP addresses. These are brand-agnostic takedowns; if it is hosted here in the UK, we will take it down. At the end of the year, we found that the UK share had risen slightly to the same level it was at the end of 2018, despite dropping below 2% briefly during the year. The number of phishing campaigns hosted in the UK have largely stayed relatively static since 2016, but our overall share has continued to drop because of the growth of phishing outside of the UK.

Figure 7. UK share of global phishing, Jun 2016 - Dec 2019



The campaigns used a total of 155,319 URLs with a median attack availability of 12 hours. A total of 62.9% were reported down within 24 hours, as shown in Table 7.

Table 7. Comparative UK-hosted phishing attack availability, 2018-2019


Measure	2018	2019
Mean (hours)	125.6	121.3
Median (hours)	9.1	12.0
Skewness	9.3	8.5
25 th percentile (hours)	1.3	1.7
75 th percentile (hours)	41.0	50.1
Down in 4 hours	39.0%	35.2%
Down in 24 hours	69.1%	62.9%

Over 61% of these attacks were hosted on compromised sites with phishing content uploaded by the attacker. 11% were under full domain ownership/control of an attacker; that is, purposefully registered for use in malicious activity. We were less confident about approximately 28% of other domains found in these attacks, but they were most probably owned by the attacker concerned.

Web-Inject malware in UK IP space

Web-inject attacks seek to inject code which runs in a victims browser as they visit a website. This code could do a variety of bad things such as steal credentials, insert bogus payment forms and modify legitimate content with something malicious. In 2019 the service took down 1,823 generic web-inject attack groups (3,175 URLs), compared to 571 (1,287 URLs) in 2018. We see that the number of these attacks over time are consistently low, with one notable spike in late 2019, as shown in Figure 8. However, we have not been able to attribute this spike to a campaign, single host, registrar, or other event. November is a busy month for e-commerce, so other factors may be involved.

Figure 8. Spike in UK-hosted web-inject malware attacks in late 2019 but otherwise remained consistently low, 2018-2019



As shown in Table 8, comparisons with 2018 data remain favourable, median attack availability has dropped despite the number of attacks increasing from 561 in 2018 to 1,823 in 2019.

Table 8. Comparative UK-hosted web-inject malware attack availability, 2018-2019

Measure	2018	2019
Mean (hours)	885.3	288.6
Median (hours)	54.7	36.0
Skewness	2.4	4.9
25 th percentile (hours)	4.0	4.5
75 th percentile (hours)	715.2	159.2
Down in 4 hours	25%	22%
Down in 24 hours	40%	45%

Web shells

One of the principal ways an attacker can maintain a foothold on a compromised server to support a variety of malicious activity is by installing a web shell. Web shells can enable lateral movement on a network, password cracking, send phishing emails, or can simply be used to control or administer compromised infrastructure. As with other malware takedowns, we found web shells on UK-delegated IP addresses, and they are sometimes active alongside some phishing attacks. In 2019 we took down 1,613 attack groups with a median availability of 17 hours. For web shells associated with UK government-themed attacks hosted overseas we took down 167 attack groups with a higher median availability of 32 hours.

When JavaScript turns bad

We have already talked about generic web inject malware that causes harm to visitors to compromised websites. There are of course other bad things that can happen when you browse the web.

Attackers can inject code that collects credit card details or deploys cryptocurrency mining code via your browser, which will steal your compute cycles and perhaps noticeably slow down your computer. In such circumstances, the attacker can sit back and reap the benefits whilst most site owners and victims are unaware of what is happening.

There is also some downstream infrastructure which can support these attacks. For example, an attacker can use a separate domain or email/IP address to extract stolen (skimmed) credentials. Also, an attack might rely on a JavaScript library that has been tampered with. By injecting code into a trusted JavaScript dependency in common use, an attacker can compromise multiple downstream websites. We are looking for and taking down these sites too.

Credit card skimmers

We took down 1,393 credit card skimmers during 2019; 861 were hosted here in the UK, and 532 were hosted overseas. The overseas sites offered transactions in UK sterling, and therefore met our criteria for takedown. These attacks continue to be a problem and we find many site owners have issues cleaning them up, based on how long these attacks can last. The median availability of the UK hosted skimmers was 109 hours (compared to 115 hours in 2018) and 318 hours for overseas hosted skimmers (compared to 192 hours in 2018).

Taking down credit card skimmer sites was a new type of takedown introduced last year. The service initially addressed skimming code that exploited unpatched versions of the Magento e-commerce platform. Since then, we have taken down skimmers found in other e-commerce platforms; however, Magento remains the most common. We also perform takedowns against the domains where the skimmed credentials are sent (skimmer credential dropsites).

From monitoring these attacks, we see some interesting patterns relating to the time it takes for site owners to remove them. The point in time when the skimming code is first removed (initial clean up) typically falls within the first few hours of an attack lifecycle. We noted however that there was a much longer duration between initial cleanup, and the malicious code being completely removed (final outage).

With site compromises of this type we typically see the attack stopping and restarting several times before the attacker loses access to the site permanently. This would perhaps indicate a more systemic problem with the security of these sites, which most likely points to poor security practices such as weak or default passwords used in server content management software.

Attacks on sports organisations

In July and August 2019, our Takedown Service detected two separate attacks against an e-commerce business that provided personalised merchandise for several sports organisations, including English and Scottish Premiership football teams. Many of these teams run their own e-commerce stores but link to a controlled subdomain to cater for personalised items, such as replica shirts.

The first attack used a familiar method and was cleaned up within 7 days by the site administrator. The timing of the attack was perhaps designed to coincide with the start of the 2019/20 season, as new replica shirts would be very popular at this time.

In mid-August, a second, customised attack was detected. This attack inserted a separate payment form into the website prior to checkout. This form sent victim credentials to a domain controlled by the attacker. The injected form and skimming code were removed within 3 to 4 days.

This illustrates how attackers will target services that serve multiple businesses to increase their return on investment. Unchecked, these attacks would likely run for many months and become a lucrative source for an attacker to sell stolen credit card details.


Non-consensual cryptocurrency miners

There has been a steady decline in the takedowns of non-consensual cryptocurrency mining since we began our work to take them down in March 2018.

Notably, in March 2019 we noted the demise of Coinhive, which ran a service to monetise websites by serving *Monero* cryptocurrency mining code to visitors. For legitimate sites who use Coinhive, it is an alternative to serving ads, with visitors clicking to accept the mining code as they browse. Monero coins would then be mined by visitors with the proceeds going to a legitimate site owner via their unique site key.

However, in the case of a compromised (non-participating) site, an attacker will embed their own site key in the injected code, so the cryptographic spoils of such endeavours go to the attacker and not the site owner. Following Coinhive's demise we saw the code of other web mining providers being used in a similar way. Overall numbers have continued to fall throughout the year, as shown in Figure 9. This may reflect the fall in cryptocurrency value and exchange rates in general.

Figure 9. Continued decrease in UK-hosted non-consensual cryptocurrency mining attacks, Mar 2018 - Dec 2019



Again, looking at attack availability versus the initial clean-up phase, we noted some long periods between the initial mitigation steps and the final outage, mirroring what we see with credit card skimmers.

We also attempted takedown against the site keys being employed. Takedown notifications were sent to Coinhive directly as the keys were clearly being used in a way that broke their terms of service. None of these came to fruition, with Coinhive Support helpfully suggesting that "...deleting the Coinhive snippet from the source code" was adequate. Sadly, there was no evidence that any action was taken by Coinhive against the account holder within our records.

Poisoned libraries

In 2019, the service also began takedowns against malicious third-party JavaScript libraries that were found to contain skimming code. Similarly, there have been plenty of other examples where fraudsters have sprayed skimming code into unsecured cloud instances normally used for website content delivery. The code is subsequently injected into other websites and, in a similar way to the attack on sports organisations, an attacker can scale their ill-gotten gains quickly.

In 2019, we took down over 70 poisoned libraries that were hosted in the UK, and a further 41 overseas.

Reinfection rates

We have also looked at the reinfection rates to ascertain whether the same JavaScript is being deployed repeatedly. We found that across shopping site skimmers, cryptocurrency miners, and poisoned JavaScript libraries:

- 61% of affected sites did not get reinfected with the same JavaScript resource
- 19.8% experienced one reinfection
- 7.4% had 2 reinfections
- 3.8% had 3 reinfections
- 2.3% had 4 reinfections
- 5.7% had 5 or more reinfections

In short, the data supports some systemic issues with how some small business/enterprises manage their web presence and e-commerce systems, as a significant number of them experience multiple reinfections. Patching to the latest versions of the e-commerce product would likely have the greatest impact on the numbers of successful attacks.

SSL certificates

We checked our HMG phishing takedowns to determine which of them used SSL certificates. As in 2018, we noted that a high proportion of the attacks leveraged Domain Validated (DV) cPanel and Let's Encrypt certificates. The former is likely a by-product of many compromised sites using poorly administered hosting, while the latter lends itself to ubiquitous, free SSL services. Similarly, there are high proportions of cPanel and Let's Encrypt certificates used in UK-hosted phishing attacks, as shown in Table 9 and Table 10.

Table 9. Top 10 certificates used in UK government-themed phishing attacks, 2019

Certificate issuer	Number of attacks (URLs)
No Certificate	9,714
Let's Encrypt	8,602
cPanel Inc.	3,880
Cloudflare	1,178
GoDaddy Inc.	1,086
Sectigo Ltd.	460
Comodo CA Ltd.	324
DigiCert Inc.	314
GlobalSign nv-sa	64
SSL.com	43

Table 10. Top 10 certificates used in brand-agnostic UK-hosted phishing attacks, 2019


Certificate issuer	Number of attacks (URLs)
cPanel Inc.	55,583
Let's Encrypt	47,328
No Certificate	43,546
Sectigo Ltd.	3,190
Comodo CA Ltd.	1,856
DigiCert Inc.	1,725
Starfield Technologies Inc.	1,029
GlobalSign nv-sa	519
Microsoft Corporation	281
GoDaddy Inc.	114

Blocking

The other aspect of the Takedown Service is that bad URLs end up in *block lists*, which browser vendors and security appliances use to prevent visits to malicious sites. These lists are updated every few minutes, so they offer some instant protection to users once an attack is discovered. Google Safe Browsing (GSB) and Microsoft's Smart Screen are good examples of how users get some measure of protection from browser-based blocks.

Between June and December 2019, the service monitored whether the attacks we discover and report were finding their way into the GSB lists deployed to user devices, and whether reported URLs get blocked on an array of test devices covering desktop and mobile. Typically, our service would check a live phishing attack every hour to see if a block ensues. Figure 10 shows the data we collected.

Figure 10. A higher proportion of phishing URLs are blocked on desktop PCs compared to mobile devices, Jun - Dec 2019



What is most apparent is the large difference between blocks on desktop and mobile devices. This illustrates the trade-off between security and performance on mobile devices. Desktop computers are generally more capable devices, while running large blocklists on a less powerful mobile device will affect performance and the user experience. Therefore the lower block percentage on both Android and iOS devices may be reflective of this.

SMS-related attacks (aka smishing)

During 2019, the NCSC looked at the kind of attacks that the UK public received via SMS, also known as 'smishing'. There is some good industry practice in place to tackle such attacks in the UK, including centralised reporting, work to block attacks at aggregator level, and initiatives to act on the attacks that get through to users.

For acting on the attacks that get through, we can scan SMS content, automate analysis of the URLs, identify attacks and inject them into the takedown workflow. Blocking at aggregator level is covered in [Routing and Signalling](#). Between December 2018 and May 2019, the NCSC looked at malicious SMS reporting kindly produced by BT's Everything Everywhere (EE) Team, which is derived from public reports to short code 7726.

These reports summarise malicious SMS campaigns that have been noted. A total of 294 reports were processed. 264 of the reports contained malicious URLs (phish and/or malware). These reports contained 54 URLs which were unknown to the service and subsequently credited as new takedowns following inspection.

Additionally, the NCSC purchased a threat intelligence feed from an industry partner to enhance our knowledge of smishing attacks. Between September 2019 and the end of the calendar year, this feed contributed to 166 new takedowns related to UK government or UK-hosted phishing attacks. There is more information on how we are working to counter the threats regarding SMS later in the section [The SMS SenderID Protective Registry](#).

Conclusions

Throughout the year, the service has continued to demonstrate harm prevention via takedown and blocking at scale. It has been gratifying to note the interest in this service from overseas. Other CERTs and government cyber security organisations have taken interest in our approach to takedown. We hope that soon similar services will be adopted outside the UK at scale.

Also, it is worth acknowledging the fantastic contributions that various organisations make to ensure the internet is a safer place. Hosters who efficiently deal with their respective abuse queues make a huge difference, as do the domain registrars who suspend accounts quickly. Furthermore, the efforts by the major webmail providers to prevent much of the bad stuff from reaching inboxes is also significant. All these countermeasures prevent harm.

Timely reporting of phishing and malware campaigns has many benefits and it has been gratifying to see many UK government departments, their IT teams, and other employees and individuals forwarding malicious mail to the service for analysis and possible takedown.

Mail Check


Mail Check is the NCSC's platform for assessing email security compliance. It helps domain owners identify, understand, and prevent abuse of their email domains. It helps organisations assess their email security and implement secure email standards.

Progress in 2019

In 2019, Mail Check was developed into a mature alpha product ready for progress to the beta stage. We initially focused on assisting approximately 40 mainly central government departments. During the course of the year, we expanded our efforts to include more than 400 local government organisations. Not only did this help these organisations improve email security, it also gave us a rich source of data that we could use to develop and improve the functionality and user interface of the tool.

By the end of 2019, 7,059 mostly public sector domains were being monitored in Mail Check, an increase of 380 over the year. The number of monitored domains with a DMARC enforcing policy of 'quarantine' or 'reject' to prevent illegitimate emails from being delivered increased from 700 to 1,246, or from 10% to 18% of all monitored domains. Figure 11 shows the monthly progress through 2019.

Figure 11. Cumulative increase in Mail Check monitored domains with a DMARC enforcing policy, 2019



Functional improvements

Development efforts focused on two key areas of Mail Check functionality: TLS Reporting and the new DMARC Report Presentation.

TLS reporting (TLS-RPT)

We added TLS-RPT configuration checking to Mail Check. If TLS-RPT is misconfigured (or not configured at all), users now receive advice on how to configure it properly. This standard allows all other domains sending email to your domain to report whether the email was sent using TLS encryption in transit. If it was not, the reason is reported.

New DMARC report presentation (aka 'drilldown')

User feedback indicated that we needed to improve the way in which Mail Check displays DMARC report information. Email traffic is now broken down by originating provider; we identify the provider by referring to a number of public data sources (such as DNS records) but only when the reverse has a matching forward address to prevent spoofing. If that fails, we refer to the internet routing table to find the origin Autonomous System, and then look up the name in a Regional Internet Registry.

Mail Check also curates a list of the IP addresses used in non-internet accessible public sector networks, such as PSN and N3. Often there is a large amount of 'noise' in the DMARC reports that can mask problems that need attention, so we took the decision to categorise emails from reported spammer IP addresses into a separate pseudo provider, allowing the more interesting data to emerge from the noise. We use public databases - known as DNS Blocklists (DNSBL) - to make this decision, as these are used by spam filters throughout the world.

Other developments

Aside from functionality improvements, we also completed a move to a pure microservice event-driven architecture, where each service can use the most appropriate technology for its use case. For example, we use PostgreSQL for storing information about IP addresses, as it supports subnetting in queries, and DynamoDB for storing a read model for the status of each domain.

As we approached the end of 2019 our technical focus had shifted to the following:

Improvements to the user experience

For most new users, it can be a time consuming and complex experience to progress their email security from a policy of 'none' through to 'quarantine', and eventually to 'reject'. We conducted user research into how this could be improved, and will be revisiting the resulting design improvements.

Self-service functionality

As well as working to simplify the Mail Check user experience, we had to develop self-service functionality to support the new influx of users in 2019. Enabling users to 'self-serve' for most administrative tasks became a focus for design efforts in the last quarter of the year.

Integration with MyNCSC

Efforts to standardise and optimise the user experience across the suite of NCSC services gathered pace in the last quarter of 2019, with the result that most Mail Check development changes were coordinated with the wider technical delivery within the [MyNCSC](#) platform.

Understanding implementation challenges

MTA-STS and TLS-RPT implementation trends and issues

Implementing Mail Transfer Agent Strict Transport Security (MTA-STS) and TLS-RPT for a domain can be tricky, so we started gathering reports from Mail Check users about how they were getting on with it. As a quick reminder:

- MTA-STS is a relatively new standard that can be used to advertise that your mail server supports TLS (Transport Layer Security)
- TLS-RPT allows other mail servers to send reports about whether or not your mail server is responding properly to TLS

In other words, the primary reason for implementing MTA-STS for your domain is to ensure that confidential email that is sent to you is transmitted securely over TLS.

Mail Check started receiving reports on 14th May 2019. By the end of the year the number of reporting domains stood at 71, which had reported on 6,562 TLS sessions. Of these, 6,144 sessions were successful and 418 had failed because they could not be validated. We also found that:

- only Google and Comcast were sending TLS reports, showing that vendor support has some way to go
- 5 of the 71 domains sending TLS reports to Mail Check by the end of 2019 had configured MTA-STS; of these, four were in testing mode and one was in enforced mode

In summary, these are emerging standards that we expect - and will encourage - to be more widely adopted over the next couple of years.

Findings from analysis of user support requests

In September we took steps to better understand our customers' support needs. We had always offered a support function, but by tracking the type and number of enquiries being received into our external facing mailbox (our 'helpdesk') we gained valuable insight into the ongoing health of the tool, the sorts of recurring issues, and a better understanding of the benefits of introducing self-service features. A breakdown of the different types of user support requests received from September to December 2019 is shown in Figure 12, and a more detailed description of the types is shown in Table 11.

Figure 12. Mail Check user support requests by type, Sep-Dec 2019

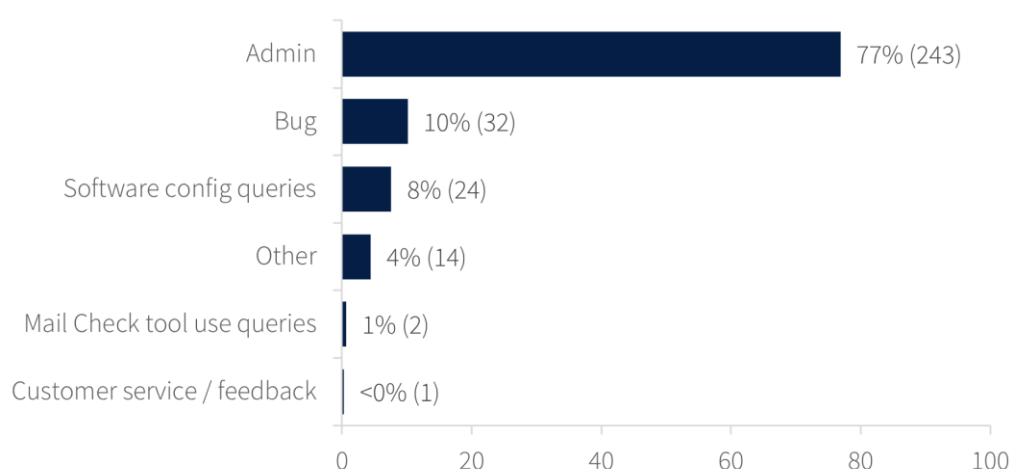


Table 11. Description of Mail Check user support request types, Sep-Dec 2019

Type	Description
Admin	Adding/deleting domains, users, workspaces.
Bug	An error/ flaw with Mail Check (e.g. spelling mistake or a user's records not showing in their user interface).
Software config queries	Technical queries about DMARC, SPF, DKIM and end-to-end encryption.
Other	All other queries.
Mail Check tool use queries	Queries about how to use Mail Check ('why is Mail Check saying x or y?', 'how do I sign up?').
Customer service / feedback	Positive/negative comments about Mail Check, suggested improvements to the tool.

Our main conclusions were that:

1. Enabling users to add and manage users and domains without the need for administrator intervention would make the account setup process quicker and reduce the workload on Mail Check administrators. These tasks accounted for 77% of user support issues. Self-service is a key focus and will enable considerable broadening of the user base.
2. Users generally understand how to use the tool, as only 1% of queries were from people who had usability issues. This indicates that Mail Check has good usability and that the guidance provided is effective. We will monitor this closely as we add new functionality and adjust the interface.
3. The 24 queries about software configuration suggests that the whole area of email security remains complex for some of those who want to improve it. Just as interesting would be to know how many are too daunted to start the process.

Campaigning for the adoption of secure email policies

We invested much more effort into promoting the adoption of secure email policies generally, and Mail Check in particular across the UK public sector. In 2018 we started supporting the Prime Minister's Office and 45 ministerial and non-ministerial departments. We continued this in 2019, recently working alongside the Cyber Team from the Transforming Government Security Programme. We also threw our net much wider to support the 418 UK councils (later 408 as a result of local government reorganisation in April 2019).

Measuring progress

Good progress was achieved by both central and local government sectors, though the question of how to demonstrate that progress proved interesting. We could measure progress by the percentage of domains with a DMARC policy of 'reject', but what if an organisation has four domains with this policy, and their SPF records are misconfigured? Or what if they changed the policy to 'reject' too quickly, resulting in many of their legitimate emails not being delivered?


There are many quantitative checks we could use to measure progress: the number of domains, sign-ups to Mail Check (or a similar tool), DMARC records, SPF records, DKIM records, and others. We settled for 'best DMARC policy per organisation'. This doesn't give the full picture, but it offers two advantages:

- it enables a simple way to show progress across multiple organisations
- it provides a target for organisations that are new to adopting secure email policies

If an organisation could begin the journey to a policy of 'reject' with, for example, their main email sending domain, it would demonstrate that they understood what was required for all their other domains, including the parked ones. Also, if we saw at the outset that an organisation had at least one domain with a policy of 'reject', it would indicate that they already knew about DMARC and how to implement it. Hence we could focus more effort on other organisations that had yet to move from a policy of 'none', or no policy at all.

By the end of the year, all 45 central government ministerial and non-ministerial departments and the Prime Minister's Office had a DMARC policy, and 57% had a DMARC enforcing policy of 'quarantine' or 'reject'. Figure 13 shows the progress during the first 12 months of campaigning within local government.

Figure 13. Increase in councils using a DMARC enforcing policy, Feb 2019 - Jan 2020



The chart shows a steady month-on-month improvement. Even the local government reorganisations in April did not seem to hinder the progress; there was an increase in councils implementing a DMARC enforcing policy of 'quarantine' or 'reject' from 201 in April to 235 in May.

We believe a contributing factor to this was our short note to chief executives to raise awareness of the risks and what they could do about them, and linking this to an entry in the Society of Local Authority Chief Executives (SOLACE) newsletter. We worked closely with the councils' IT teams via our SOCITM, LGA and WARP networks to create and distribute the note.

Overall, there was a steady improvement by the local government community during the first 12 months of the campaign, with 90 councils having progressed from a policy of 'none' and 'no policy' to a DMARC enforcing policy of 'quarantine' or 'reject'.

Campaign approach

Compliance with the UK government's minimum cyber security standards is offered as the basis for promoting secure email policies for central government departments. However, there remained a question over whether these standards were directly applicable to the autonomous local government community.

We understand that organisations need to set their own priorities, and merely stating that they can adopt these or other standards (such as ISO27001 or our own Cyber Essentials standards) does not necessarily equate to action, let alone DMARC adoption. Therefore we decided to take a change management approach to supporting the councils.

Our dedicated Secure Email Campaigns team, together with the NCSC Engagement Team, tapped into the invaluable expertise of several existing networks, including:

- Ministry of Housing, Communities and Local Government (MHCLG)
- Devolved administrations
- NI Cyber Security Centre
- Society of IT Managers (SOCITM)
- Local Government Association (LGA)
- Regional Warning, Advice and Reporting Points (WARPs)
- Society of Local Authority Chief Executives (SOLACE)

Using the Prosci® change management method, and the ADKAR® tool in particular, we took time to raise awareness of the need to change, and to build desire for it. Thereafter, we focused our efforts on sharing technical knowledge, working in close collaboration with the external change network. Collectively, we helped build up the internal ability in councils to make the necessary secure email policy changes. At the same time, we invested considerable time and effort into refreshing our online guidance and reinforcing the progress made with ongoing one-to-one support where necessary.

We refrained from setting one-size-fits-all timelines of (for example) 'x' days to reject, on the understanding that the time taken for an organisation to get to a position where they could set the DMARC policy to 'reject' would depend on a great many factors (such as resources, priorities, and the complexity of their email infrastructure). Instead, we favoured a high-level functional roadmap, without set timings and deadlines.

What we learned

As well as improving adoption rates for central government and local government organisations, we also learned a great deal about how to improve on our work.

Alignment is key

If we had to choose one word that cropped up persistently in 2019, it was 'alignment'. As well as being the technical concept in SPF and DKIM that continues to cause the most confusion, it is also critical for successful email adoption in organisations. Success is not achieved by signing up to a tool or implementing a DMARC policy of 'reject'; nor is it about correctly syntactically configuring SPF, DKIM, and DMARC records. It's about being aligned, at every level.

An organisation's leadership team need to be aligned, and the organisation's public DNS record for DMARC, DKIM and SPF also need to be aligned. Moreover, they need to be aligned for every single source of legitimate email. There are now hundreds of third-party apps for bona fide organisational processes that can, and do, legitimately send email on behalf of an organisation. It is critical that all legitimate sources of email are identified, and therefore the leadership team need to be aware and motivated to liaise with the IT team, and other teams that run email systems.

Treat each organisation as a new case

We recognise that each organisation has its own set of challenges, resources, infrastructures, and priorities. It follows that one organisation's timeline or implementation approach will not necessarily work for another, or indeed any other.

Achieving p= 'quarantine'

Getting to, and staying for a while at, a policy of 'quarantine' helps to ensure that spoofed email is correctly identified and handled, and also helps to ensure that legitimate emails that have not been set up correctly are also delivered, albeit to the junk folder.

Involving the campaigns team

The campaigns team has a tried and tested approach to help organisations with secure email adoption. It generates motivation, and feedback gathered from workshops, meetings and one-to-one support activities indicates that organisations are as keen to work with us as we are with them.

Not just an IT problem

The IT part of all secure email adoption (including, for example, setting up the public DNS records for DMARC/SPF/DKIM) is complex and time consuming. However, equal if not greater effort must be made to raise awareness within an organisation and to gain support for the time-consuming process of gathering email source information. This is not just an IT team responsibility, and if not done properly, it will lead to legitimate emails not being delivered and nefarious activity not being stopped.

Web Check

Web Check helps organisations identify and fix common security issues in their websites. It works as follows:

1. A user signs up to use the service on behalf of their organisation.
2. The user adds URLs to their watchlist.
3. URLs on the user's watchlist are checked regularly by the service.
4. The user can view the most recent results, with appropriate and clear mitigation advice.
5. Changes to results are listed on their Notifications dashboard and, if the user so chooses, sent by email.

Web Check went live in September 2019. This landmark followed a few years of agile development through the alpha and beta phases, with a strong focus on regularly consulting users. The final changes were less visible as we addressed some key items of technical debt 'under the bonnet', with a view to ensuring ongoing robustness. However, this is not the end of the journey, and our focus has moved to continual service improvement.

By the end of 2019, the service had over 2,500 users, representing over 1,200 customer organisations. For the most part, these organisations were in the central government, local government, emergency services, and health sectors. Together they were using the service to regularly scan over 30,000 URLs per day.

The user survey perspective

We surveyed our users to gather feedback about the service, and received a large number of responses.

63% of respondents said that they had been using the service for over a year. That feels like a healthy percentage in two ways; it shows that user retention is high and so the service is perceived as worthwhile, and also that the service continues to attract new users.

The influx of new users was from turnover of staff within existing customer organisations, and also new organisations signing up to Web Check. Many of the new users arrived in a steady trickle, having heard about the service from ourselves or other contacts. One area from which we received a burst of interest was Northern Ireland, on the back of our friends in the NI Cyber Security Centre promoting Web Check and other ACD services.

Furthermore, 88% of respondents asserted that the service continues to be useful and relevant to them. The survey included some open-ended questions to allow respondents to give more details about their experiences. Here we were pleased to see evidence of a number of points of our design intent being realised. Some of the key themes were as follows:

- The service was simple to use.
- Small teams who do **not** consider themselves cyber security experts appreciated having access to a basic set of checks from an external source of expertise.
- The regularity of the URL scans was reassuring, even where the websites had not been recently changed. This was because scans could pick up existing issues that had been missed (such as an approaching expiry date of an X.509 certificate) or the need to update software to the current version.
- The mitigation advice that accompanied the identification of issues reduced the time spent to research solutions.
- Some users who also used commercial web scanning tools appreciated the fact that Web Check performed daily checks, a frequency that they could not replicate with their commercial tools. Others from organisations with limited budgets appreciated Web Check being provided to them at no cost.

Web Check itself can only do so much, and the organisations that use it need to address the identified issues. So it was also pleasing to see that 87% of respondents felt that Web Check played a role in driving improvements to security within their organisation. Supporting comments here noted the value of having objective evidence to help improve the security culture within a user's organisation.

Responses to the survey also included some constructive suggestions for making improvements:


- The service currently offers only an individual user basis of operation. However, it is natural for users within an organisation to want to share watchlists and findings with their colleagues. This has also been the service improvement most frequently requested in other user feedback.
- A frequent theme was that of integration with other ACD services. Our users would, for example, value a single dashboard, drawing their attention to the most significant findings across the various ACD services they use. Again, we are pleased to report that we are on the case, as outlined in the [MyNCSC](#) section of this review.
- The user survey listed a few common service improvement suggestions, and the one that won most votes was improved filtering for watchlists and findings. The work the NCSC is undertaking to integrate the ACD services should offer a good environment for progressing such improvements.
- Surprisingly, 90% of respondents expressed interest in more technically detailed findings. While this is at odds with other feedback and with our service vision of providing entry-level web security checks, we must take account of this strong response. We have options to consider, such as making changes to the checks performed by Web Check, or even how looking at how other ACD services can help.

The internal metrics perspective

We review internally-generated data to help us understand the health of and benefits delivered by the service.

Drilling down into the overall number of findings to see which type of concerns are raised most frequently is illuminating. Figure 14 does so for urgent and advisory findings flagged to users, and gives a comparison of results between 2018 and 2019.

Figure 14. Comparison of urgent and advisory Web Check findings by type, 2018 -2019



The types of checks were broadly unchanged from 2018, with the numbers indicating that the associated concerns remain relevant, and the general increase largely reflecting the greater numbers of URLs checked. For an explanation of what these checks involve, please see the [previous year's report](#).

These numbers show an ongoing need for our customer organisations to remain vigilant to ensure their website security. Most significantly, the largest numbers display the need to:

- Properly implement HTTPS and TLS protocols, supported by use of X.509 public key certificates, to give those visiting the websites confidence both in the authenticity of the websites and in the confidentiality and integrity of the data flows involved in communicating with them.
- Maintain the currency of the chosen web server software and, where used, Content Management System (CMS) applications, thereby reducing the risk of known vulnerabilities in those products being exploited by attackers.


- Consider whether there is a genuine requirement for use of XML-RPC. This introduces security weaknesses so should either be disabled or, if necessary, access to it should be limited.

Another metric we consider is the number of the most urgent findings resolved each month. These are issues that Web Check detected as no longer existing, having previously flagged them up to users. We cannot be certain that it was Web Check that prompted our customer organisations to address these issues, but it seems fair to assume that this was so for many of the cases, and that is corroborated by responses to the user survey.

Figure 15 shows the month-by-month results. The differences appear to be natural monthly variations, which one would expect from a statistical viewpoint. The key message is that the service consistently delivered benefit, with an average of almost 550 urgent issues resolved per month. Examples of findings categorised urgent include:

- imminent (that is, in 7 days or less) expiry of an X.509 certificate (which would then block operation of HTTPS and undermine confidence in the customer website)
- running a version of a CMS that it is no longer supported (and therefore carries significant risk of vulnerabilities being present)
- the WordPress CMS being configured with directory listing enabled (thereby carrying the risk of an attacker obtaining detailed data about the site)

Figure 15. Urgent Web Check findings resolved per month remained fairly consistent, 2019



Conclusion

Web Check is a relatively mature service, offering usable features that deliver ongoing benefit to our customers. Our user research has revealed that our user base has an appetite for further improvements to the service, and we are also keen to develop it further to ensure its ongoing health.

Protective DNS

The Domain Name System (DNS) is the address book of the internet. Your computer relies on DNS to find out exactly where 'example.com' (a domain) is located (its IP address) so it can connect to it. Anyone can register a domain so that everyone else can find the IP address associated with it, to enable them to connect to it.

Unfortunately, 'anyone' includes those who wish to cause harm. Attackers often use seemingly legitimate domains as part of malware and phishing attacks.

PDNS exists to combat that malicious activity for public sector users. PDNS prevents the successful resolution of domains associated with malicious activity, while enabling the rest of the internet to remain accessible.

DNS, numbers and scale

As always, quantifying *good* DNS protection is challenged by the nature of DNS and varied behaviour of devices on a network. False positives skew numbers and the changes we made to our threat feed providers also impact the statistics. Notably in 2019 we saw large numbers of false positives generated by security appliances pre-emptively resolving their own 'bad' lists. This isn't necessarily bad practice, but it does present a challenge to DNS security products and puts additional burdens on customers to get the configuration right to ensure their security tools, security information and event management (SIEM) products and their PDNS dashboard report back good data.

On a much more positive note, we have been quite successful in using the data we collect across the whole service to model average or typical behaviour, and use it to allow customers to see how they rank against others. This enables us to plot percentiles, provide red/amber/green scoring, and feed back to customers at times when they may want to look more closely at their logs.

PDNS is maturing, and as our active users grow, our visibility across the public sector is allowing us to make observations, provide more meaningful metrics and feedback, and identify the areas most needing attention.

Quantifying the protection we provide does still rely on old fashioned 'total blocks' and 'unique blocks', customers protected, and so on. However, we can now increasingly show examples of how we leverage the scale of PDNS, in which we make observations, inform policy, and take mitigative action across our entire customer base (now representing 2.2M users). This allows us to truly deliver the original impact-at-scale vision and see the unique value of a centrally run service for the public sector.

Protective DNS in 2019

In 2019, PDNS increased its estimated number of protected UK public sector employees by 57%, from 1.4 million to 2.2 million. In total, the service handled 142 billion queries over the 12 month period, more than double the 68.7 billion queries made in 2018, with a peak query rate of 43,726 queries per second at one point in October. The highest peak query rate seen in 2018 was 27,109 queries per second in November.

Early in the year we made some changes to our threat feed sources with the aim of improving our intelligence data. This was successful and we are pleased with the outcome, but the change in source data makes direct comparisons with data from previous years difficult. However, it is still useful to look at the breakdown of threat types for this year.

Of the 142 billion queries handled in 2019, we blocked 80 million queries to 175,000 unique domains. When we look closer at these numbers, we find that 25 million blocks were related to algorithmically generated domains (AGDs); 16 million blocks were related to botnet C2, 14,000 for indicators related to exploit kits, and 3,200 for ransomware. PDNS continues to build a rogues' gallery of malware. Common culprits in 2019 were:

- Emotet
- Necurs
- Kraken
- Sphinx
- Neutrino
- Cerber
- CryptoLocker


- GandCrab
- Wannacry
- NotPetya
- BadRabbit
- Ramnit
- Tiny Banker
- Conficker

This year also saw the launch of our new web portal, which provides alerts and information about DNS events, as well as access to logs of blocked DNS requests, and integration with SIEM solutions.

Throughout the year we have continued to increase the number of UK public sector organisations we protect. We are protecting over 200 more organisations and blocking 41,000 fewer queries per organisation than we were during 2018, which reflects an increase in network security amongst our customers and better control of false positives from the changes made to our feed providers.


Figure 16 shows the number of protected organisations on the last day of each month in 2019; we ended December six short of 500 at 494. That's still an impressive number considering that we have many grouped accounts, managed service providers, and public service networks on that list, so the true number of organisations protected is likely in the multiple thousands.

Figure 16. Cumulative increase in organisations and estimated customers protected by PDNS, 2019



A key measure of our effectiveness is what we call 'unique blocks', which are the unique domain names that are blocked, each one counted only once per period. Figure 17 shows the unique blocks per month per organisation in 2018 and 2019.

Figure 17. Unique blocks by PDNS for the average organisation per month began to level off, 2018–2019



As explained in [last year's ACD report](#), the 'per organisation' normalisation has limitations (as there is a wide variety in the size of organisation we protect - from a small local authority to a large central government department), but it is this bias that we can try to account for by looking at unique blocks. Also, we are ignoring the fact that organisations usually join part of the way through a month, but can justify that omission as the effect is likely to be small. In 2018 we saw a wide variation in unique blocks each month, but as the service has matured the number of unique blocks seems to have settled at about 40 for the average organisation.

Key uptake milestones

It's hard to know where to set expectations in terms of achievable uptake, especially in the hugely diverse sectors of local and central government. However, looking back on 2019 it is obvious we crossed a significant milestone in terms of uptake across our key government sectors. Most central government departments and the majority of local authorities are actively using the service, specifically, 21 out of 25 ministerial departments, 14 out of 20 non-ministerial departments, and 265 out of 408 local authorities are using PDNS.

Central government

Between January and December 2019, an additional 11 (24%) of the 45 central government departments joined PDNS, taking the total uptake from 53% (24 of 45 departments) to 78% (35 of 45 departments). Departments are at the heart of government and in one year PDNS uptake has increased from around half to more than three quarters. This represents a substantial improvement in the cyber security of central government.

Local authorities

Between January and December 2019, 102 local government organisations from across the UK joined PDNS. This means an extra 1 in 4 local authorities (as well as some organisations that provide shared services or carry out similar functions) are now protected by PDNS. This growth took PDNS coverage of local government from 40% (163 of 408 organisations) to 65% (265 of 408 organisations).

Also of note were particularly strong engagements from devolved governments in Scotland, Wales and Northern Ireland.

Infrastructure

Capacity increase

With the fantastic engagement and utilisation of PDNS from organisations across the public sector, we took the decision in 2019 to prepare for the future and doubled the capacity of the PDNS.

Infrastructure resilience

The infrastructure on PDNS remains a tempting target for malicious actors and unsurprisingly there were many DDoS attempts against it in 2019. Thanks to our multi-site, anycast configuration and automated provider level DDoS mitigation, none of these led to any degradation of service for any customer.

As in 2018, we saw both short, sharp attacks and more sustained attacks, but all were handled by the resilience of our infrastructure and our upstream providers. The peak load we saw was 16.9 Gbps in August, which although a new high, was comparable in magnitude to previous attacks. To date we have maintained 100% service availability, but again this shows the importance of building resilience into infrastructure and, for this service, the need to have secondary resolvers on standby.

Threat feeds

To enhance our blocking capability, we changed some of our threat feed providers early this year. This involved removing some feeds and adding others that we assessed to provide us with higher quality threat intelligence and, crucially, fewer false positives. For example, we now have much better data on command and control (C2) threats, which is reflected in a higher number of blocked domains identified as being C2-related, compared to 2018. Our threat feeds come from public and private sources and Table 12 shows the number of unique blocks we made in a single month of 2019, based on the information we get from each of our key providers.

Table 12. Threat feed performance in a month of 2019

Feed provider	Number of unique blocks	Feed used (%)
Feed 1	8,350	99
Feed 2	2,672	78
Feed 3	110	99
Feed 4	2,247	87
Feed 5	605	100

One of feeds in Table 12 is the NCSC's unique high quality, low volume threat feed, often containing the most serious threats. In December alone 5,622 known-malicious domains were passed from the NCSC to PDNS to be blocked. Overall, the year saw impressive leaps in technical capability, allowing higher volumes of declassification of sensitive threat intelligence from the NCSC's Threat Operations investigative teams.

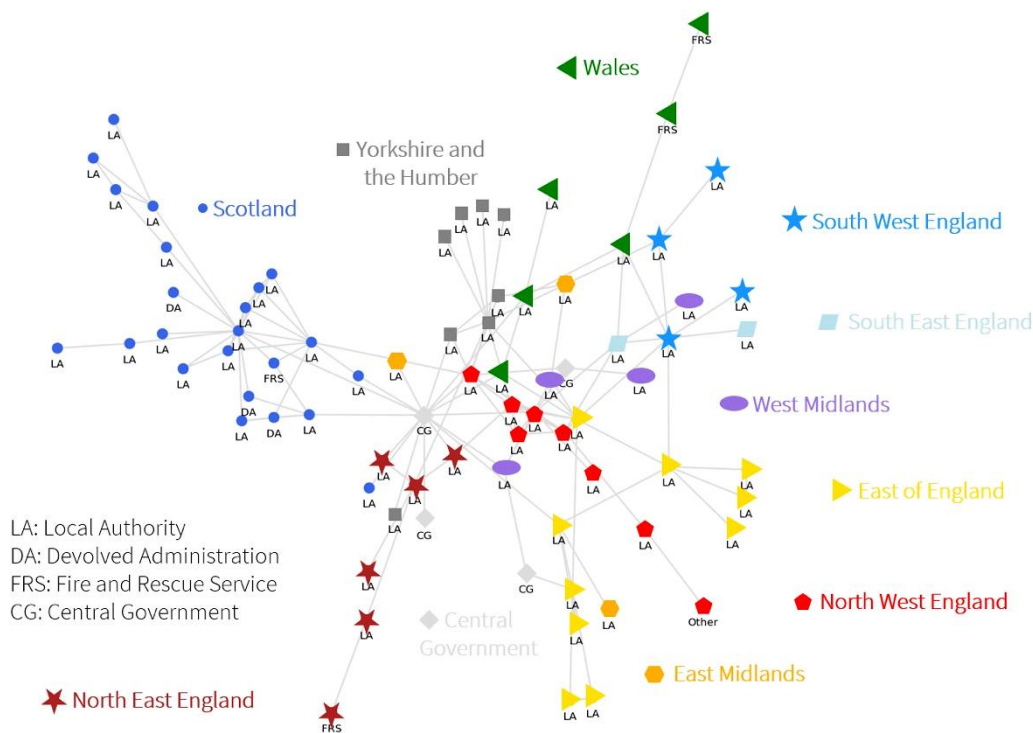
The changes to our threat feed providers early in 2019 led to a measurable improvement in our number of false positives, or domains that we are told are malicious, but in fact are not. We saw only 29 false positives (down from 62 in 2018) - a good sign of how we are improving in this area.

Case studies

2019 saw significant progress behind the scenes in how we share and exploit the PDNS data internally within the NCSC. This means that we're able to exploit this data in new ways to make observations at scale. Our analysts and data scientists are increasingly able to use PDNS to develop fascinating new ways of providing enhanced security to our customers, and across the public sector.

Improving our understanding of the government digital estate

PDNS data is a rich source of information for the NCSC's cyber security research projects. In 2019, the NCSC's data scientists developed a network model of government organisations using this data. The work involved identifying pairs of public bodies that request each other's domain names frequently. This simple metric is a powerful way to build up a picture of the working relationships between public bodies. We have not included organisations' names in the Figure 18, but it does indicate organisation types and geographic regions. You can see that (unsurprisingly) organisations that are near to each other geographically tend to interact with each other online. Some less obvious relationships are also revealed.

Figure 18. Network model of public bodies that query each other frequently via PDNS, Dec 2019

We were also able to identify new public sector domain names using a combination of PDNS data and machine learning techniques.

This research helps the NCSC spot new government digital estate as soon as it starts to be used, and gives us a better understanding of how public bodies interact in cyberspace. This helps us provide cyber security support to the organisations and infrastructure that will have the greatest effect on the overall cyber resilience of government. We are now starting to model and predict the propagation of incidents and vulnerabilities through government using the models developed in this project.

Tracking the extent of an incident across government

PDNS is a versatile tool that can provide early warning and monitoring of cyber security incidents affecting public bodies.

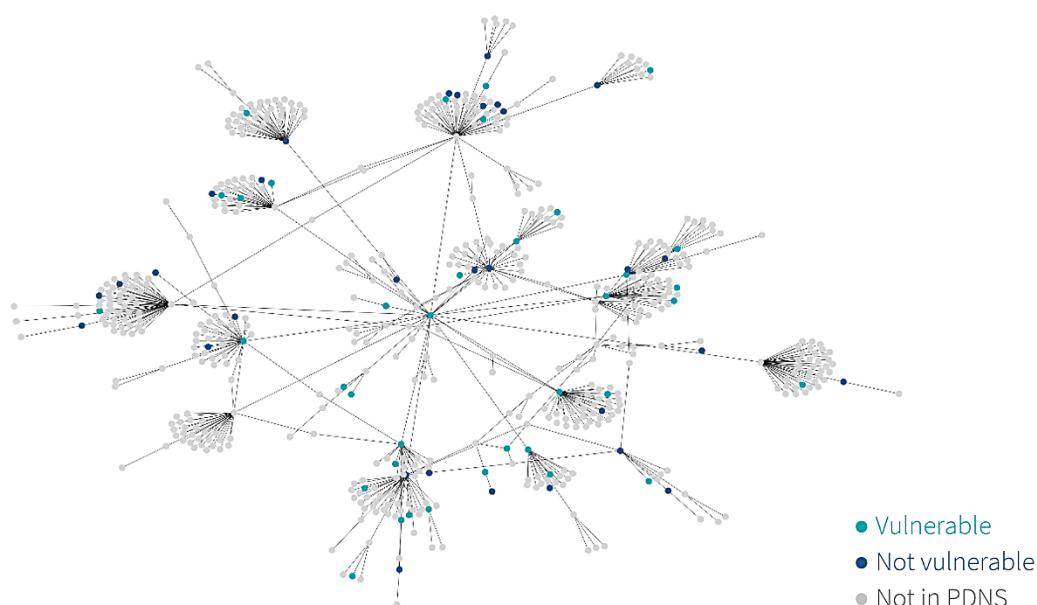
PDNS data was instrumental in quickly estimating how much of government was affected by several recent incidents, a model of which can be shown in Figure 19. As soon as the NCSC becomes aware of incidents affecting a particular type of infrastructure or service, such as Citrix or Cisco, we can look in the PDNS data for signs that participating organisations are affected.

The workflow includes the following steps:

- researching technical indicators, or fingerprints, of vulnerabilities and incidents
- looking for those fingerprints in PDNS and finding out which organisations show them
- during a confirmed incident, monitoring fingerprints continuously to track progress in remediation
- evaluating the effectiveness of different mitigations

Information on the estimated extent of government affected by an incident is passed from ACD data analysis teams to Incident Management, Cyber Operations and Engagement teams at the NCSC, who coordinate the response and support victims.

Figure 19. Network model of central government organisations with PDNS-based estimate of exposure to a vulnerability, Dec 2019



We used this capability to identify PDNS-utilising organisations vulnerable to the Citrix and Cisco vulnerabilities disclosed in 2019. We gave them the information required to fix or mitigate these vulnerabilities.

Machine learning AGD detection

There has been much research into detecting algorithmically generated domains (AGDs). These domains are used by malicious actors as rendezvous points with their C2 servers. Not all these domains will connect to the servers, but the large amount of them presents a challenge in terms of their identification and removal – more so when the domains are word-based. In this task, we sought to use machine learning and natural language processing techniques to try to identify three different malware types by their word-based AGDs.

We can extract certain features from a domain, such as:

- lexical features
- word similarity
- domain length
- number of digits
- top-level domain (TLD)

Note that we initially considered extracting every TLD, but we found that selecting the top 20 TLDs as independent features and the rest as a separate combined feature worked equally well.

As a control set, we used the Alexa top million dataset truncated to the first 10,000 results, as it was highly unlikely that malicious domains would appear to be popular.

There were two approaches to the modelling:

- combine all malware types and have one classification model
- consider each malware type separately and have three different classification models

Combination of malware types

When we grouped the three malware families, we found that we could not reliably identify AGDs. Table 13 shows the recall (our ability to find all positive examples) was 0.867, and the accuracy (the relevance of the returned examples) was 0.870. Both are better when closer to 1.

Table 13. Recall and accuracy figures for three malware families calculated as a group

Malware	Recall	Precision
Malware family 1	0.876	0.870
Malware family 2		
Malware family 3		

When we ran the analysis on each malware family separately, the classification of AGDs was nearly perfect, as shown in Table 14.

Table 14. Recall and accuracy figures for three malware families calculated individually

Malware	Recall	Precision
Malware family 1	1.000	0.998
Malware family 2	0.996	0.969
Malware family 3	1.000	0.993

These results show that it certainly is possible to accurately identify some types of algorithmically generated domains, and that having a rich dataset for both malicious domain samples (as well as an allow list) greatly aided in the analysis performed.

For future research, we will explore DNS-specific features relating to domains, such as TTL, the query type (A, AAAA, TXT, etc.), frequency and other factors.

A limitation of this technique is that it relies on supervised learning; it requires knowledge of the malware type to be able to categorise it. While this requires us to pivot from known AGDs, we have still found this to be useful because threat feeds aren't able to enumerate (or don't share) all domains from an AGD. In practice this allows us to train on AGD domains reported, then identify many more which have not been reported.

Other future investigations will include clustering domains based on their characteristics, as described above, to identify the general notion of a word-based AGD rather than a specific malware type. The ability to distinguish between generally legitimate and malicious domains is the subject of future research and could, by extension, allow one to distinguish between malicious AGDs and legitimate ones (such as advertising and CDNs).

Phishing campaign

In May 2019 we identified queries for subdomains of a foreign internet service provider (ISP) originating from various PDNS customers. Queries from one of the affected organisations are shown in Figure 20.

Figure 20. Spikes in queries to a foreign ISP domain originating from a PDNS protected organisation, May-Jul 2019

The naming convention of the domains indicated that they belonged to the dynamic pool of IP addresses used for PPPoE (Point-to-Point Protocol over Ethernet) connectivity, which allows ISPs to divide the total bandwidth of their internet lines to give bandwidth to multiple customers. PPPoE is a popular protocol used on ISP's Digital Subscriber Lines (DSLs). Our sources indicated that these domain names were associated with spam campaigns, and that these queries were probably generated by a spam campaign that originated from IP addresses managed by the foreign ISP. We observed A/AAAA DNS record queries as well as pointer (PTR) queries associated with these domain names, which is expected behaviour from email servers that check the origin of incoming mail.

Ideally, mail servers should not trust emails coming from a server unless they can do a reverse DNS lookup. This means that email servers should do a forward lookup on the name that the sender's email server introduces itself as (such as mail[.]example[.]com) to find its IP address, and then do a PTR lookup on that IP address to make sure that it resolves to the same name.

In this case, the UK public sector organisations in question had been receiving a large increase in emails from a single foreign source, and we believe that they were victims of a targeted phishing campaign.

Infected virtual machine

In July, one of our customers saw 18,000 queries for domain names known to be associated with a certain malware family. Interestingly, the customer only had 21 unique blocks (meaning all of their blocked DNS queries were for only 21 known-malicious domain names), and the top 15 blocked domains were all related to this same malware family. We notified the customer and provided the logs of blocked queries to assist in an investigation.

The customer's analysis successfully identified the device that had been attempting to reach the malicious domains and found that it was a virtual machine (VM) being used by one of their developers; that is, an emulation of a computer system that runs in software on another computer. The malware seemed to be trying to use the infected VM to send spam.

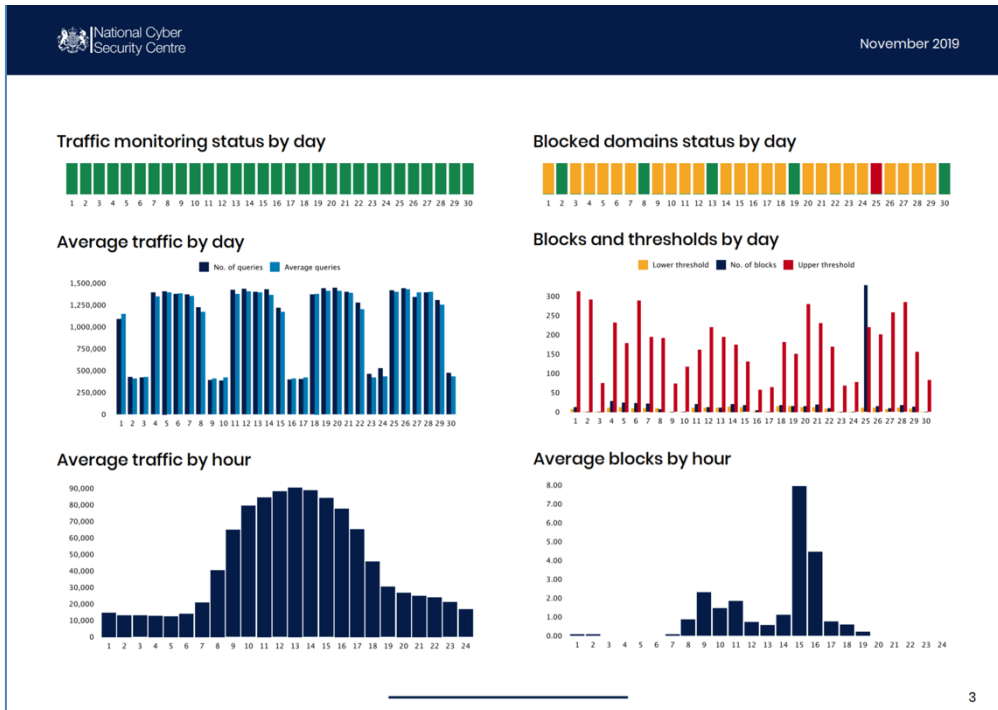
The VM was deleted, and the customer investigated all their VMs for evidence of malware. They also increased monitoring and the use of security products to limit the chances of similar events in the future.

New customer portal

In April we travelled to Glasgow for Cyber Scotland Week, incorporating our flagship event, [CyberUK](#). There, we launched three highly requested features; our new web portal with PDNS Customer Dashboard, access to logs of blocked DNS requests, and integration with security information and event management (SIEM) solutions.

The PDNS Customer Dashboard was designed to provide a high-level overview of an organisation's activity on the service and highlight areas for concern using traffic light indicators. Although high level, it contained useful information, such as the number of DNS queries made by the customer, and a comparison of block rate with the average organisation using the service. Organisations can now manage their account, query why a site is blocked (and request that it be unblocked), access the knowledge base, and view the service status through the portal.


That was just this first step, and the NCSC is committed to developing the PDNS Customer Portal to provide increasing amounts of actionable intelligence to illuminate the blocks and inform onward action to enable customers to identify and mitigate threats themselves. Later in the year we added Monthly Reports to the Customer Portal, based on input collected from our customers. It is now possible to download monthly statistics informed by the blocks made to see the positive impact of using PDNS, as shown in Figure 21.

Figure 21. Example of a Monthly Report from the PDNS Customer Portal

Block data and SIEM integration

With the launch of the portal at CyberUK, we included for the first time the ability for organisations to download the full logs of their blocked DNS queries, including all of the threat intelligence that informed the block. We ran a series of webinars to demonstrate this new feature and provide a 'how to' for its use. The aim was to make everything available, from using basic logging, to detailed data for analysis, all the way through to ingesting and analysing the data in a SIEM solution.

The first webinar was entitled 'How to Access your Block Data' and was held on 16th May. A recording can be accessed by registered users on the PDNS knowledge base. Immediately after it was held, a local authority submitted a request for their block data. Within an hour the AWS Bucket credentials had been supplied, and less than an hour later the council confirmed they had already found an indication of a ransomware infection. They successfully removed the ransomware infection and Figure 22 shows a clear drop in blocked queries for that organisation by the end of May.

Figure 22. Decrease in blocks per day for an unnamed local authority after May when they accessed their block data and successfully removed a ransomware infection, Apr-Aug 2019

Often requested alongside the ability to access logs, was support for SIEM solutions, the ability to pull logs in a standard format for ingest into SIEM tooling. This was a contributing factor to our choice of Structured Threat Information eXpression (STIX) as the markup language for logging, as it is a highly verbose, versatile and commonly supported format. Our new API provides access to a feed of blocked query events in rich STIX 2 format, including context and supporting metadata informing the threat. This enriches SIEM tooling, enabling prioritisation and investigation with the information needed to inform each step.

The choice of STIX was also made to ensure compatibility with our own [Logging.made.easy \(LME\)](#) service. If customers need a free to use, versatile SIEM tool, we have documentation on our knowledge base for registered users that describes how to pass PDNS Block Events into LME.

Knowledge base and training webinars

We invested in training materials for all the new features. The team at Nominet did a fantastic job putting together training documents and running regular webinars, walking customers through our new features and making those guides and videos available online in our knowledge base.

Like our other events throughout the year, we were amazed by the degree of interest shown, with some of our webinars attracting well over 100 attendees.

Customer engagement and feedback

Based on the successful engagement event we ran in London in 2018, we set our ambitions high for 2019 and planned workshops across the country to provide the opportunity to attend to our customers all over the UK.

Engaging with people is one way we encourage uptake to the service, but it is also incredibly useful to us for other reasons. All our new features are delivered to address customer needs, and without good input from our community there is no way of knowing what we need to deliver, and whether we're giving them something that's actually useful.

Nominet put the leg work in to deliver a number of engagements across a variety of formats, physical and digital.

PDNS workshops 2019

The PDNS Workshops were definite highlights of 2019, with over 300 attendees. This was in no small part due to the event management team at Nominet, and also the excellent engagement from our PDNS customers. Attendees joined us in Edinburgh, Manchester, Glasgow (CyberUK) and Bristol to learn more about the product, provide feedback, and to talk about new features that would be useful.


Webinar engagement

Although very useful, the workshops are a little cumbersome to use in an agile way. Enter the PDNS webinars. Throughout the year we used webinars to demo newly launched features, provide training, and to host informative sessions to explain the goings-on behind the scenes, and how to interpret data.

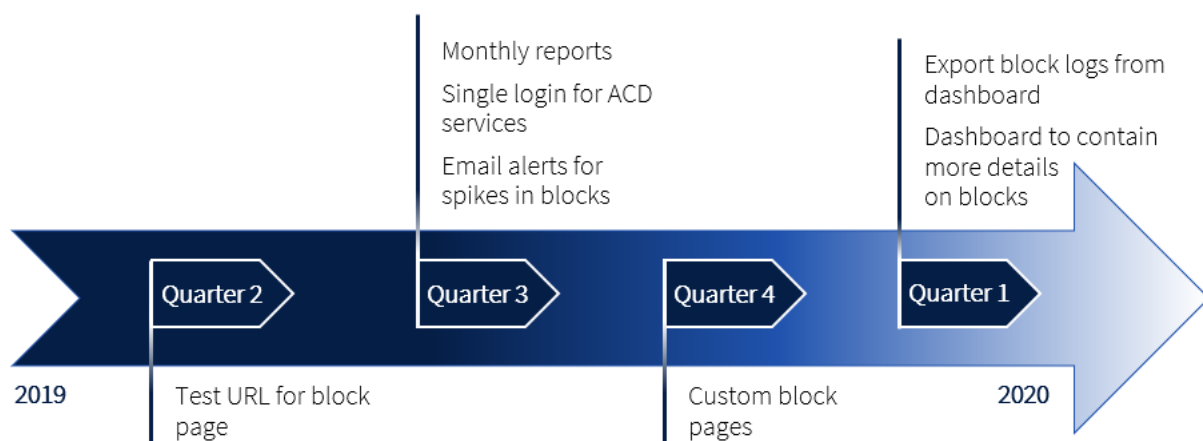
Again, Nominet did an excellent job running these webinars for hundreds of attendees, and based on the positive feedback from our customers and their invaluable use as training materials, we will continue running them in the future.

Customer feedback

The importance of feedback from our customers in shaping the service cannot be overstated, and in May 2019 we conducted a thorough cross-customer survey. Responses to key questions are shown in Figure 23.

Figure 23. Responses to key survey questions about recommending PDNS and ease of onboarding, May 2019

We are obviously very happy with the engagement numbers and the scoring against the key questions, but another key input we received was in feature requests. You can see when these were delivered in Figure 24.

Figure 24. Delivery times of features requested by customers in May survey, Q2 2019 - Q1 2020

We look forward to refining these features and sharing our development roadmap via the customer portal.

A new contract to supply a protective domain name service

A lot of hard work went on behind the scenes in 2019, and a huge effort went towards the landmark procurement that saw us approach the market for a bidder to take the service through the next 3-5 years.

We must express our gratitude to the 12 organisations that accepted the invitation to tender, and the final 4 that made the competition so close. It was encouraging to see the high standard of technical innovation and mature landscape of DNS security products, and the strength of quality options for DNS security available to the private sector. This all made for a very close competition. In the end, Nominet's experience meant they emerged as winners.

Signed in September 2019, the new contract is to commence in April 2020 and will introduce new features, increased technical research and innovation, whilst also providing the same highly available, geographically diverse and low latency service. Our high standards for the service will be maintained and we look forward to the development of new features, led by experts both here at the NCSC and Nominet, and most importantly by the experiences and feedback of our customers.

Nominet have been contracted to provide our customers with dedicated technical and customer support available 24 hours a day, 365 days a year by telephone, email and web portal, and will respond to all enquiries within one business hour. An important part of the customer support service is to 'unblock' a web address when PDNS does not resolve an address that an organisation needs to access. To support this Nominet will 'unblock' high priority domains within one hour inside business hours, and within four hours outside of business hours.

Conclusion

It has been a busy year for our PDNS team and for the supporting teams in the NCSC, not to mention our delivery partner Nominet. This summary has captured many (though not all) of the efforts made throughout 2019. It saw PDNS take the penultimate steps towards service maturity. It is now in a great place, and having secured critical mass across government departments and local authorities, we are actively delivering the impact, value and perspective to make the UK government and public sector as safe as we can.

Dangling DNS

When a Domain Name System (DNS) record points to a site or other resource that no longer exists, this is known as a 'dangling DNS' issue. This can happen for several reasons, the most common of which is an oversight or delay in registering resources at the start of a project, or forgetting to remove them in a timely fashion as part of the decommissioning process.

Sometimes these resources can be hijacked; that is, registered by another party, resulting in the dangling DNS record pointing to a new resource. If an attacker targeting a domain discovers a dangling DNS record, they could hijack it and cause it to point towards a site under their control.

To put the scale of this problem into perspective, 20% of all the reports received within the first year of our Vulnerability Disclosure Program were due to hijackable subdomains resulting from dangling DNS records.

Automated detection at scale

In 2019, we decided to investigate this problem. We had two goals:

1. To determine how easily such vulnerabilities could be detected continually on a large scale.
2. If feasible, to develop a working prototype system to prove it.

Our aim was to provide a critical insight into the services most commonly affected by these vulnerabilities, allowing us to identify and work together with key providers to prevent them from occurring in the future.

We also wanted to explore the feasibility of defensively registering the target resources automatically on behalf of domain owners to prevent them from being hijacked, with a secure and straightforward way of releasing them again afterwards. This would enable us to provide a unique service, offering UK-wide protection for businesses and organisations against some of these threats.

Outcomes

We identified the three types of DNS record most commonly found to be dangling, each of which could be exploited by attackers in different ways, but to achieve the same end result. These were:

- A: specifies an IP address for the domain
- CNAME: defines an alias for the domain
- NS: defines the authoritative name for the domain

These issues were by no means novel but had been popularised by Bug Bounty programs in recent years, and are often referred to as 'subdomain takeover' vulnerabilities. We experimented with applying detection methods for these record types using [Rapid7's extensive OpenData](#) sets as a source of both subdomains and DNS records.

By November 2019, we had developed a cloud-based prototype on AWS, and used it to scan over 44,000 subdomains of .gov.uk and .mod.uk for dangling A, CNAME, and NS records. The prototype could automatically schedule scans and scale to cope with large datasets, but it could not automatically register dangling records.

Nevertheless, we learnt a lot from this first prototype, including which were the most common service providers with dangling DNS records, and how to cope with numerous edge cases to avoid false positive results.

Routing and Signalling

Fixing the underlying infrastructure protocols on which the internet is based has been a key strand of the ACD's work since inception. We have focused on two specific protocols: the Border Gateway Protocol (BGP) and the Signalling System No.7 (SS7).

Border Gateway Protocol

The internet is comprised of nearly 90,000 networks, known as Autonomous Systems (ASs), and the Border Gateway Protocol (BGP) is used to determine how internet traffic is routed between them. BGP was developed when there were many fewer ASs, and has little authentication or integrity. Therefore, it is easy for any participant in the protocol to accidentally or maliciously reroute large swathes of internet traffic.

There are cryptographic extensions to BGP that try to solve part of this problem. Unfortunately, the cost of implementation is high. These problems are a good lesson about the importance of getting trust models right in distributed, global systems.

In an effort to improve security, the NCSC has been working on establishing best practices and developing a BGP monitoring platform.

BGP best practice

Whilst BGP is a ubiquitous protocol used extensively across the internet, there are many ways that it can be implemented. Although these implementations may all work from a routing point of view, they can also result in insecure deployments that can either put the Communication Service Provider (CSP) or its peers at risk.

Under the UK Network Security Information Exchange (NSIE), a working group has been established to look at BGP, and how it is used and deployed by the CSPs in the UK. This working group will produce a best practice guide, to be published on the NCSC website. The guide will give advice on how to configure BGP securely and efficiently, to protect both CSPs and their peers and customers. This is done by direct reference to recommendations, as well as providing links to other relevant BGP documentation and guidance already available.

BGP monitoring

Historically, there has been no foolproof way to detect BGP path update anomalies as part of normal operation of the protocol. In 2018, in collaboration with BT we began to develop a proof of concept BGP Monitoring Platform. This project is now known as BGP Spotlight, and the beta version - hosted by BT - is currently undergoing tests.

In its development stage, BGP Spotlight took in update feeds via an API with BT's UK, European, and rest-of-world networks, in order to collect different views of the internet routing table. Now, updates are collected via BGP peering; this is where two routers create a BGP connection in order to exchange information.

BGP Spotlight also has peerings with 3 other operators, and an ingest from a fourth is being developed. These supplement the original BT and publicly available RouteViews feeds. Diversity of data is key to getting multiple views of the traffic routing, and these additional peerings allow for a much wider view of BGP routing updates.

As the data is ingested, analysis is performed to detect any unexpected or irregular path updates or IP prefix advertisements. Updates are very context specific and are affected by real world events, such as networks going down or an accidental damage to a cable. The whole point of BGP and internet routing is to provide a self-repairing, resilient network, which can make it difficult to identify unexpected or irregular activity. The same update could be deemed normal in one context and anomalous in another.

Similarly, misconfigurations are commonplace, and can be hard to distinguish from malicious announcements. However, malicious and unintended routing anomalies can both be equally harmful to internet providers.

Currently four operators are using the beta system. Between them there are 120 users, and in a typical 24 hour period the platform ingests approximately 500m messages. These produce approximately 25m events of interest, of which approximately 24.8m are path changes, almost 200,000 are ownership changes, and a small number are 'first-seen' prefixes, which occur when a specific IP range has never been seen before.

The ownership changes and ‘first-seen’ are the most significant events, since they are most likely to lead to traffic outage. Path changes are typically just a result of dynamic routing, which is key to the resilience of the internet. However, they should not be discounted completely, as they can reflect a re-routing of traffic via a location that we would be concerned about.

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 messages, and allow users to roam between countries. SS7 was created in 1975 with no real security built in, and has changed very little since then.

Although it is impractical to change such a long-established standard, the NCSC believes it is possible to better protect users of UK networks from these sorts of attacks, while simultaneously ensuring that later generation telecoms signalling protocols (including DIAMETER) are better secured.

All UK networks using SS7 that we tested in 2018 contained vulnerabilities, some of them quite serious. In 2019 we extended the testing to cover DIAMETER (4G) and the GPRS Tunnelling Protocol (GTP), which have similar vulnerabilities to SS7. So far, our tests have also revealed serious examples of these vulnerabilities in the UK’s mobile networks.

The NCSC is continuing to work with the owners of the affected networks to resolve these issues. Most mobile operators are deploying new equipment to mitigate some of these threats, and in order to provide more flexible testing, we have now moved to a more automated approach. We’ve made a testing service available to the networks so they can test periodically or when there are significant security improvements to their signalling networks.

This also provides the NCSC with an ongoing and more up-to-date picture of the risk status of the UK’s signalling network. We are hoping to use this ongoing capability to monitor progress as the picture hopefully improves with time.

Telecom Security Requirements

During 2019 the NCSC began to develop the Telecom Security Requirements (TSRs). These establish a set of principles and requirements that, when followed, will establish a suitable baseline security framework for an operator’s telecoms network. The UK government is legislating, through the [Telecommunications \(Security\) Bill](#), to ensure adherence to the TSRs.

The TSRs are broken down into 13 sections, one of which specifically focuses on signalling security. As operators begin to adhere to the TSRs, this will help to mitigate the risk of basic signalling attacks and build effective monitoring systems to detect advanced attacks.

These operator monitoring systems will be able to either use, or build upon, the existing signalling security solutions developed as part of the NCSC’s ACD programme. Together, the Telecommunications (Security) Bill and the ACD programme provide the strategic means to significantly reduce signalling risks in UK networks.

The SMS SenderID Protective Registry

For the NCSC, tackling the abuse of SMS involves two approaches. In the [Takedown](#) section, we describe how we mitigate the URLs found in SMS phishing (‘smishing’) attacks using our Takedown Service. The second approach is about trying to prevent the messages being delivered in the first place.

The NCSC, along with [UK Finance](#) and others, has part-funded an initiative to set up an SMS SenderID Protective Registry. This allows brand owners to register authorised SenderIDs/alpha tags, define their SMS delivery chains (that is, the SMS aggregators they choose to deliver their traffic), and to provide a list of unauthorised SenderIDs that they have already seen abused in smishing campaigns.

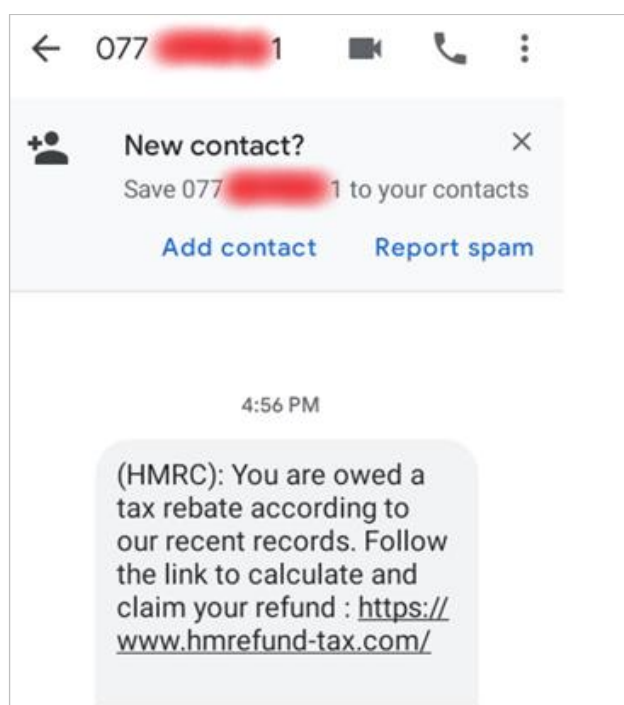
The registry was created and is independently administered by the Mobile Ecosystem Forum (MEF). Participating SMS Aggregators use the registry to ascertain whether they should block or deliver SMS traffic which is routed via their networks. At a superficial level, you can think of the registry as a codex, to illustrate whether an aggregator should block traffic or allow it to pass to the mobile network operators for onward delivery to their subscribers. In practice, an authorised SenderID (for example, DVLA) will be delivered, but a bogus derivative (such as DV1A) will not.

In 2019, both HMRC and DVLA began use of the registry. Though we cannot speak for other brand owners who use the registry, we can report some initial findings and observations from both DVLA and NCSC perspectives.

To their credit, HMRC enforced strict controls on their SMS usage in 2017, limiting how their brand and services interact with SMS. The numbers of malicious SMS reported by the public to HMRC dropped dramatically. That said, attacks did continue, albeit in much lower numbers. In particular, HMRC noted attacks using variations of unauthorised alpha tags such as 'Gov Tax', 'HMInfo' and 'HMRCUK', as shown in Figure 25.

Having logged these unauthorised senders in the MEF registry, they saw a 70% drop in public reports as the attacks were blocked by participating aggregators before they ever reached their intended targets. HMRC reported further drops in public reports throughout 2019, and notably many long number MSISDNs were identified as the originators of these attacks.

Figure 25. Example of a long number MSISDN smishing attack



The NCSC noted a dramatic fall in UK government smishing attacks using a SenderID in mid-2019, coinciding with the registry coming online. DVLA had a similar experience, noting long number MSISDN attacks becoming the norm from June 2019 onwards.

The last (malicious) use of SenderID 'DVLA' was reported to the NCSC in May 2019 and has not been seen since, which will have clearly lowered the authenticity of many DVLA-based smishing campaigns. Seeing the positive effect this had on DVLA we invited the TV Licensing agency to participate in the MEF registry in late 2019.

Host Based Capability

Host Based Capability (HBC) is a software agent that can be deployed on government OFFICIAL IT devices, such as laptops, desktops, and servers. It collects and analyses technical metadata to **detect** malicious activity, to help the department understand their **threat surface**, and to **forewarn** those affected by new, major vulnerabilities:

- **Detect:** HBC has identified or assisted on a cumulative total of 15 incidents. By participating in HBC, departments can receive granular information at the device level. This can help improve incident response and remediation.
- **Threat surface:** By the end of 2019 HBC generated a cumulative total of over 170 threat surface reports. These reports highlight cyber security strengths and weaknesses, helping departments make decisions about their security posture. Anecdotal evidence suggests that departments do implement changes to address the results of these reports.
- **Forewarn:** HBC provided information to departments about their unique vulnerability to [CVE-2019-0708 \(BlueKeep\)](#). This comprehensive, granular data can help departments react quickly to emerging threats.

Growth through 2019

HBC is now towards the end of its second year of operation and is in the beta phase of its service offering.

HBC celebrated an exciting benchmark: the number of government devices it is installed on has risen from 26,000 at the end of 2018 to over 130,000 by December 2019. Many more government departments have expressed an interest and are in the process of adopting the service.

Vulnerability Disclosure

The Vulnerability Disclosure project is focused on maturing the UK's approach to vulnerability disclosure and remediation. There are two main strands of work:

- **Vulnerability Reporting Service** If someone finds a vulnerability in a UK government online service and is unable to report it directly to the system owner, they can report it to the NCSC.
- **Vulnerability Disclosure Pilot** Helps improve the UK government's ability to adopt vulnerability disclosure practices by creating a Vulnerability Disclosure Program for any department who signs up.

Outcomes

Vulnerability Reporting Service (VRS)

The VRS is continuing to mature and grow, and in 2019 we received over 200 reports from security researchers all over the world. These security issues could have caused harm to the UK government. As we continue to work with the security researcher community and UK government system owners, we are further developing how we manage and help remediate the reported issues.

One way we have done this is to work with other ACD projects. For example, we have collaborated with the [Dangling.DNS](#) project to automate the process of finding and registering dangling domains.

Vulnerability Disclosure Pilot (VDP)

The pilot provides government departments with a ready-made disclosure management process, and secure reporting and workflow management of received reports via the HackerOne platform. The triage service for reports is provided by NCC Group, who validate the initial report and severity rating. This ensures quicker adoption and implementation of industry standard approaches.

Following the success with early adoption, the pilot is focusing on rolling out to central government departments. We commissioned an independent baseline report of government departments and found that less than 1% of departments have an existing vulnerability disclosure management process. The pilot is continually learning and adapting, following feedback from both the participating departments and the security researcher community. In 2019 we launched five vulnerability disclosure programs and a further eight were going through the onboarding process

The pilot has helped government departments receive, remediate and fix issues before they cause harm.

NCSC Observatory

In [last year's report](#), we shared some early-stage research linked to an 'Internet Weather Centre' and described plans to develop a prototype. The work on this prototype has continued throughout 2019, albeit under the new name 'NCSC Observatory'. Whilst we share brief details of our work here to provide some context, the Observatory's real value is realised by quietly supporting the NCSC's other functions and services.

We remain focused on generating data-driven insights to underpin the NCSC's research and strategy. Modern networks, including the internet, are a web of complex interactions between systems, where issues with one system can affect a myriad of others. By identifying the deployment and use of technologies, how they are connected, and their use of particular security controls, we aim to illuminate systemic risks and vulnerabilities in the UK's digital economy.

Outcomes

Measuring adoption of security controls in the UK

We have completed a number of wide-scale studies into the adoption of particular security controls across the UK. For example, we investigated the adoption of the DMARC protocol by domains registered in several UK zones (such as co.uk). We will report on trends in next year's reports, once we've collected the data for a sufficient period.

We also hoped to compare the overall UK adoption rate against the UK public sector, where we are working to improve adoption of email security controls through our [Mail Check](#) service. However, whilst adoption is significantly higher within the public sector, we do not include the statistics here as we believe the results are heavily skewed by parked domains. These account for a large number of the domains counted, but typically do not implement these controls.

In addition to these studies, we've also worked with the NCSC's Platform Security team to conduct more targeted investigations. In one such study we analysed mail exchanger (MX) records, which specify the mail server to be used for incoming mail for a given domain. By looking at the [organisational domains](#) of these mail servers, we measured the relative market share of email providers used by UK domains, shown in Figure 26 and Figure 27.

Figure 26. Most popular MX records for [co.uk](#) domains, Nov 2019


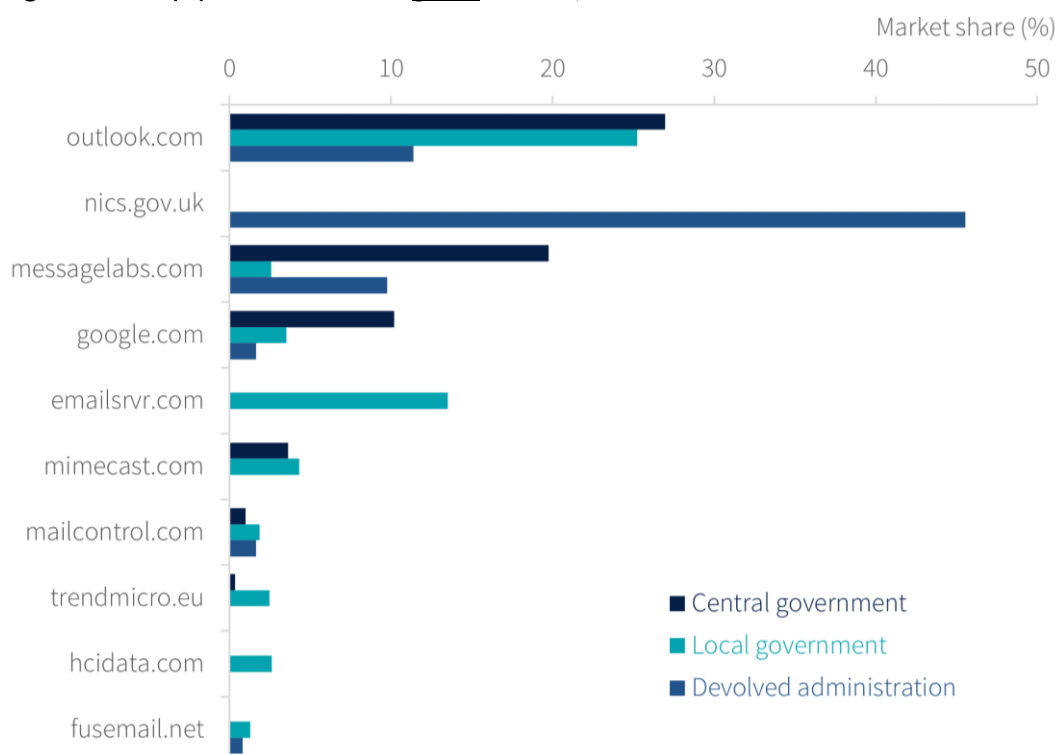


Figure 27. Most popular MX records for gov.uk domains, Nov 2019

Determining which providers have, or are gaining, significant market share helps our researchers prioritise their work. From this, we can see that whilst Microsoft (outlook.com) is the most commonly used email provider for both central and local government, 123 Reg is the most popular email provider for co.uk domains.

As a small extension to this study, we also measured adoption of the DKIM protocol by domains hosted on Office 365 and Google G-Suite. We wanted to understand whether one provider had a larger uptake, and if so, whether this might indicate that one provider's configuration options were more intuitive.

For public sector domains, we found that DKIM adoption was very similar at 33% for Office 365 and 35% for G-Suite. However, for .co.uk domains, uptake was observed to be 4.0% and 7.8% respectively; domains hosted on Google's service were almost twice as likely to have enabled DKIM than those hosted on Microsoft's service.

National-scale vulnerability management

Alongside these studies, we also used the Observatory to support the NCSC's vulnerability and incident management functions for the first time. In light of [CVE-2019-19781](#), a security issue affecting Citrix Application Delivery Controller (ADC) appliances, we identified the UK-linked hosts exposing an affected device on the internet. The NCSC's Incident Management function used this data to prioritise and direct their engagement with at-risk organisations.

Suspicious Email Reporting Service

In 2018, ACD created a proof of concept for a service that the general public could use to report suspicious emails: the Suspicious Email Reporting Service (SERS).

Development of this service alongside our law enforcement partners continued throughout 2019, in preparation for the planned launch in early 2020. When launched, members of the public will be able to report emails by forwarding them to report@phishing.gov.uk. These reports will be automatically analysed and, if malicious content is found, a takedown notice will be issued to the hosting provider requesting it removes the content.

Not only will SERS help to identify and take down malicious sites, it will also provide law enforcement with vitally important data that will help them to understand strategic threats to the UK and to alert the public to phishing trends.

Exercise in a Box

Prominent ransomware attacks within the UK, have brought cyber security into sharp public focus. Despite the varying skill levels of attackers, we have found that the most serious and impactful attacks only require a low-level of skill to execute. In some cases, the effects of such attacks could have been mitigated by taking steps to improve an organisation's defences, putting good cyber hygiene principles into practice, or by planning and implementing an incident response capability.

Running cyber exercises can help an organisation understand its preparedness for and resilience to various types of attack. Therefore, as noted in [last year's report](#), the NCSC's Exercise in a Box tool was formally launched at Cyber UK in April 2019. This is a free, publicly available, end-to-end tool that helps organisations to:

- establish the effectiveness of their current defence and response mechanisms
- test and check their existing policies and procedures
- improve their internal relationships and skills through the mechanism of cyber exercising
- identify areas for further improvement
- gain access to relevant advice and guidance from the NCSC


The Exercise in a Box tool is intended for use by the non-cyber expert within local government, or a small to medium-sized enterprise. As such, the exercises offered have been designed to include everything the facilitator needs to set up, plan, and deliver the exercise. Completion of the exercise generates an end report and links to relevant NCSC advice and guidance.

Based on current and emerging threats, topics are split into tabletop and simulation exercises. The currently available topics are:

- mobile phone theft and response
- threatened leak of sensitive data
- phishing attack leading to ransomware infection
- third party software compromise
- insider threat resulting in a data breach
- BYOD policy implications
- attack from an unknown Wi-Fi network
- cyber threat simulation exercise - 'responding to a mock threat'

Outcomes

From launch in April 2019 until the end of December 2019, the Exercise in a Box tool attracted over 4,500 registered users. These users were split across the public sector, small to medium-sized enterprises, large businesses, and cyber security professionals. Figure 28 shows the uptake from April 2019 to January 2020.

Figure 28. Cumulative increase in Exercise in a Box users across all organisation types, Apr 2019 - Jan 2020

Exercise in a Box has attracted publicity and praise not only from its target users, but also from the charitable sector, governing bodies, and numerous cyber security professionals. It has been endorsed by publications such as Computer Business Review and [Forbes.com](https://www.forbes.com):

‘Exercise in a Box’ addresses the knowledge gap by making security a business issue rather than levelling it at IT teams. ASOS CISO George Mudie said, “At ASOS we decided to incorporate the ‘Exercise in a Box’ content into our data security incident rehearsals. We found that the structure of the desktop exercises and simulation really helped to bring the rehearsals to life as well as encourage discussion and feedback”.

In addition, some organisations have begun to adopt use of the tool as a matter of best practice. For example, the NHS’s [Digital Data Security Continuity Planning Standard](#) documentation states, “It is highly recommended that you utilise National Cyber Security Centre (NCSC) ‘Exercise in a Box’ for desktop testing”.

Logging Made Easy

Logging is the foundation on which security monitoring and situational awareness are built. It is essential to be able to refer to logs in the event of a cyber security incident, in order to determine what has happened and to make the necessary changes to prevent it from happening again.

Logging Made Easy (LME) is an open source project that provides a practical way to set up basic end-to-end Windows monitoring of your IT estate.

It is not a professional setup, but it will give you a basic logging capability. This is infinitely better than no capability at all.

Alpha release

The alpha release of LME was launched at Cyber UK in [April 2019](#), and was greeted positively by both cyber professionals and potential users of the platform (that is, organisations without logging capability who are interested in getting one started).

LME combines multiple free and open source software tools, pre-made configuration files and scripts, and a tutorial to help people install it themselves. It is available from the [NCSC's GitHub page](#). Version 0.2 was released in October 2019, and included an Elastic, Logstash, Kibana (ELK) stack to enable customers to make use of the data through the analysis and visualisation capabilities enabled by ELK.

From launch until December 2019, over 400 unique clones of LME have been made from GitHub. Table 15 shows the unique clones and views per month, and calculates a conversion percentage based on these figures.

Table 15. Usage of LME from GitHub, Jul - Dec 2019

Period	Number of unique clones	Number of unique views	Conversion (%)
July 2019	60	1717	3.5
August 2019	53	835	6.3
September 2019	89	974	9.1
October 2019	69	845	8.2
November 2019	74	1179	6.3
December 2019	75	873	8.6

You can follow further developments of the project on the [LME home page](#).

MyNCSC

Since the inception of the NCSC, we've invested in building products and services that bring real value to our users. Along the way, we've seen repeating patterns emerge within those capabilities. These are centred around themes such as authentication and authorisation, asset management and data exploitation. Through extensive research we've also learned a lot about how our users would like to work smarter. We know that collaboration is important in helping users reduce duplication. We know also that centralised asset management can help users save time when working across multiple services.

Many ACD users make use of more than one service today, with over 40% of [Mail Check](#) users also registered for [Web Check](#). As we continue to increase the variety of services on offer, we expect that many more of our users will be multi-service.

In response to the things we've learned and the things users have been asking for, we are now consolidating our cyber security services within a single, highly-accessible platform delivered by the MyNCSC project, which was set up in late 2019. This new platform will be the single entry point to the public services already available today, and it will be where we make new services available in the future.

Behind the scenes, we're implementing some new architecture and enabling services. These will allow the technical integration of the cyber security services and support new, improved ways of working. We'll enable a coherent, consistent experience for users accessing NCSC services, helping them reduce duplication, save time and understand their security position across a range of services. Users will only need to sign in once to retrieve all their information and will be presented with service data, incident information and guidance to help them be proactive.