




National Cyber
Security Centre
a part of GCHQ



Active Cyber Defence

The Fourth Year

Contents

Introduction	3
Takedown Service	6
Mail Check	25
Web Check.....	33
Protective DNS	36
Dangling DNS	43
Routing and Signalling	45
Host Based Capability.....	50
Vulnerability Disclosure	51
NCSC Observatory	53
Suspicious Email Reporting Service	54
Exercise in a Box	56
Logging Made Easy	59
Cyber Threat Intelligence Adaptor	60
MyNCSC.....	61
Conclusion	62

Introduction

"How can we use it to help..." was the theme of our Active Cyber Defence (ACD) efforts in 2020. As the pandemic took hold, we looked at ways in which we could use the ACD tools, services and projects to support people and organisations moving to a working from home model, and also to protect the health, retail, and other sectors during that critical period.

As criminals sought to exploit the situation by adapting their cyber attacks, we moved to frustrate their efforts. Throughout the year we adapted our offerings as circumstance changed and supported organisations critical to COVID-19 response, including the vaccine supply chain. The year ended with a very different challenge, as we took an important role in responding to the SolarWinds Orion compromise and its impact on the UK.

The aim of the Active Cyber Defence (ACD) programme is to "Protect the majority of people in the UK from the majority of the harm caused by the majority of the cyber attacks the majority of the time." We do this through a wide range of mechanisms, which at their core have the ability to provide protection at scale. In most cases that requires a lot of automation, and hence data is generated as part of the process. One of the goals of this paper is to provide transparency using that data, to demonstrate what works (and what doesn't). In many cases we've developed these approaches with the public sector in mind. We're actively looking at how we can broaden the reach of these efforts directly, but there will be limits. While we won't be able to share our implementations with everyone, we're keen that the lessons we've learned are broadly understood and applied.

The report is broken down by individual effort, but these aren't siloed efforts: each service and project influences, supports, and guides the others. This can be seen in data from the Vulnerability Reporting Service highlighting the scale of the 'Dangling DNS' problem, and the team addressing that problem using data from Protective DNS (PDNS) to improve their detection capability. Similar supportive flows exist across the portfolio, and are growing as the efforts mature.

We've also made this report less "numbers heavy", and instead have tried to focus on key stories and important trends that we have discovered in the course of our work. Despite the uniqueness of the events in 2020, we've included comparisons to 2019 as we have in previous years. Sometimes there were significant differences in scope in 2020, and where that is the case we've included an explanation and/or comparable data.

As is mentioned throughout, this is a team effort, including UK public sector, commercial and international partners without whom we wouldn't be able to implement these national scale cyber security defences.

We welcome feedback on this report, particularly ideas for improved approaches, data that would be useful to include future reports, and comparisons or pointers to similar efforts. Please contact us at ACDenquiries@ncsc.gov.uk or via our social media and other contact channels.

Broadening

Until 2020, the tools and services developed under the ACD programme predominantly focused on the public sector. The benefits they have provided to public sector organisations has been proven many times over, as described in previous ACD annual reports.

The ACD Broadening project was established in 2020, with the aim of expanding the impact of ACD beyond the public sector. The aspirations of the project are to:

- provide ACD services to a broader audience where possible
- inform and advise in sectors and areas that our services do not cover
- incentivise the cyber security market to provide quality, useful products
- encourage adoption of cyber security products - both ACD and other - for all organisations

The project started by assessing the challenges faced by organisations outside the public sector. This involved analysis of data collected from numerous sources and included surveys of private sector organisations, reporting on common incidents and threats, NCSC experience, and so on. This analysis will inform how ACD helps provide solutions to challenges faced by organisations outside the private sector, and is intended to be an iterative process.

The project has considered a variety of options to help address challenges organisations face, including:

- **Existing ACD tools** - expanding the eligibility of existing services to allow private sector organisations to subscribe to them.
- **New ACD tools** - developing new tools to meet the needs of private sector organisations.
- **Guidance** – writing guidance and blog posts based on NCSC expertise and lessons learned from roll out of ACD services within the public sector. This information will allow organisations to be more informed when acquiring and rolling out cyber security services.

Throughout the individual sections of this report, you will see reference to expansion of services into new sectors, because of ACD Broadening.

For much of 2020, the project has focused on drawing together data, understanding legal and policy challenges, and working out the best way to bring the benefits of ACD to organisations outside the public sector. This will provide the foundations for greater impact next year.

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

Note that we use the Government Digital Service (GDS) service manual guidelines for describing the phases of our projects and services; that is, the discovery, alpha, beta, or live phases. When a service goes live, this means it has passed all the relevant standard assessments and gates, not that it's only just been turned on. For more information, refer to the [Agile delivery section](#) of the service manual.

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 and telephony systems are based: the Border Gateway Protocol (BGP) and the Signalling System No. 7 (SS7). This includes setting up initiatives such as the SMS SenderID Protective Registry, which helps organisations protect their brand from abuse in SMS phishing campaigns.

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.

Cyber Threat Intelligence Adaptor

The Cyber Threat Intelligence (CTI) Adaptor is a software program, designed by the NCSC's Threat Detection and Response team, that enables authorised organisations to receive a high-quality, contextually rich, cyber threat intelligence feed from the NCSC.

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

About the service

2020 is the fourth 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.

Progress

2020 total takedowns

In total, 700,595 campaigns (1,448,214 URLs) were taken down in 2020: a massive fifteen-fold increase in campaign takedowns on the figure for 2019 (45,603 campaigns and 192,256 URLs). This increase was possible because we invested in a wider set of takedown measures during the year, allowing us to address different categories of campaigns.

These different categories have large numbers of URLs associated with them, such as the fake celebrity endorsement scams (286,213 campaigns, 731,080 URLs) and fake shops (139,522 campaigns, 222,353 URLs), and this is the principal reason for the overall increase in takedowns. That said, we still saw an increase in the number of campaigns (from 45,603 in 2019 to 91,629 in 2020), URLs, and IP addresses when we excluded these new takedowns. There was also a decrease in the percentage of attacks taken down within 24 hours, from 64.6% in 2019 to 55.5% in 2020 (excluding new attack types). The details are shown in Table 1.

Table 1. Comparative overall annual takedowns, 2019-2020

Measure	2019	2020 (excluding new attack types)	2020 (including new attack types)
Total number of takedowns (attack URLs)	192,256	311,363	1,448,214
Total number of IP addresses	21,111	40,789	301,938
Total number of attack groups (campaigns)	45,603	91,629	700,595
Median attack availability (hours)	11	17	44
URLs down within 24 hours (%)	64.6	55.5	33.4

In the Outcomes section, we will cover [UK government-themed attacks](#), [Takedowns in UK-delegated IP space](#) and a number of new takedown initiatives ([The COVID-19 response and new types of takedown](#) and [Active fraud defence \(from Nov. 2020\)](#))

Outcomes

UK government-themed takedowns

We have continued to provide brand protection for UK government departments and services this year. In 2020, we removed a total of 27,611 campaigns that used UK government branding in some way, not all of which were phishing sites. The details are shown in Table 2.

We noted that the number of UK government-themed phishing campaigns (11,286) and URLs (59,435) more than doubled compared with 2019's figures (4,471 campaigns and 25,741 URLs). The median availability also increased from 15 to 21 hours, although this trend was not mirrored in other attack categories, such as web shells where the median fell from 50 to 23 hours.

Table 2. UK government-themed attacks by type, 2020

Attack Type	Number of attacks (URLs)	Number of attack groups (campaigns)	Median Availability (hours)
Phishing URL	59,435	11,286	21
Phishing URL mail server	4,913	4,913	25
Malware attachment mail server	2,890	2,890	25
Advance fee fraud	2,310	2,310	11
Advance fee fraud mail server	463	463	25
Malware infrastructure URL	918	448	2
Malware distribution URL	388	341	1
Web shell	586	296	23
Instagram brand infringement	258	258	22
Phishkit archive	187	148	32
DKIM signed email domain	111	111	1,858
Malware command and control centre	91	67	31
Facebook brand infringement	62	62	34
Phishkit email	56	56	14
Fake mobile app	44	44	43
Twitter brand infringement	38	38	311
Brand infringement	35	32	245

UK government-themed phishing

In 2020, we took down 11,286 UK government phishing campaigns, a total of 59,435 URLs. These attacks were hosted all over the world and the median availability of these attacks was 21 hours, with 52% taken down within 24 hours of discovery.

Table 3. Comparative UK government-themed phishing attack availability, 2019-2020

Measure	2019	2020
Mean (hours)	207.3	177.3
Median (hours)	14.7	21.2
Skewness	7.8	6.3
25 th percentile	1.8	3.5
75 th percentile	62.2	111.6
Down in 4 hours	33.6%	26.6%
Down in 24 hours	60.0%	51.9%

As in previous years, although the majority of attacks are taken down within the first 24 hours (52%), there are a number of takedown requests that take a long time to be actioned, with 19% taking more than 200 hours to be removed.

Figure 1 shows that NameCheap became the most popular host of UK government-themed phishing during 2020. By December 2020 we found that it hosted in excess of 60% of phishing in this category.

Figure 1. Top 10 hosters of UK government-themed phishing campaigns, highlighting NameCheap and GoDaddy who saw greater volatility in their monthly totals, 2020





Table 4 compares the change in the top 10 hosters from 2019 to 2020.

Table 4. Comparison of top 10 hosters of UK government-themed phishing attacks, 2019-2020

2019			2020		
Hoster	Share (%)	Median availability (hours)	Hoster	Share (%)	Median availability (hours)
GoDaddy	15.7	29	NameCheap	28.8	47
Shinjiru Technology	5.8	40	GoDaddy	11.2	37
Amazon	4.7	1	OVH	4.8	6
Cloudflare	4.6	21	Amazon	4.7	4
Microsoft Corporation	4.6	1	Endurance International	3.9	23
Endurance International	4	9	Shinjiru Technology	3.6	14
OVH	3.5	14	Cloudflare	3.0	12
Velocity Servers Network Exchange	2.9	19	Alibaba Group	2.7	28
NameCheap	2.5	20	DigitalOcean	1.7	18
Alibaba Group	2.3	29	Hostkey	1.6	16

Looking specifically at the number of campaigns hosted by NameCheap against its monthly median attack availability, we see that by mid-year the median takedown times were consistently in excess of 60 hours. This undoubtedly made NameCheap an attractive proposition to host phishing and may explain the rise in monthly hosted campaigns that followed for UK government-themed phishing.

Figure 2. Rise in UK government-themed phishing campaigns hosted by NameCheap coincided with a period of median attack availability for NameCheap consistently in excess of 60 hours, 2020



We have continued to track which UK government departments are used as a phishing lure. As in previous years, we found that HMRC-themed attacks were the most popular, attracting over 4,000 campaigns in 2020.


Of particular note, we saw elevated levels of phishing that used NHS branding in 2020 (see [Case Study 1](#)):

Table 5. Top 10 most targeted brands of UK government-themed phishing attacks, 2020

Government brand	Number of attacks (URLs)	Number of attack groups (campaigns)	Median availability (hours)
HMRC	22,148	4,249	20
Generic gov.uk	16,945	3,322	23
TV Licensing	13,658	3,035	23
DVLA	4,157	882	16
Council Tax	864	275	26
Government Gateway	778	138	16
NHS	159	122	13
BBC	523	112	24
UK University	46	23	30
Student Loans Company	34	11	6
All UK government-themed phishing attacks (total)	59,435	11,286	21

We have continued to track the global phishing ranking of HMRC, TVL and DVLA throughout 2020. Both DVLA and TVL fell out of the top 100 in the latter parts of 2020, with HMRC's brand phished more consistently throughout the year, as shown in Figure 3.

Figure 3. Variation in global phishing target rank for HMRC, TVL and DVLA, 2017-2020



Case Study 1: Protecting the NHS throughout the pandemic

The NHS is one of the departments that the NCSC provided support to during the pandemic. By looking for attacks that use NHS branding we identify phishing forms which seek to harvest NHS credentials. Such stolen credentials could be used to compromise critical systems, which could have serious consequences at an incredibly important time for the NHS.

Figure 4. Example of a phishing website that cloned NHSmail to steal NHS credentials

Sign In

<http://www.oohlimited.one/nhs/index.php>



In 2020, there were 122 phishing campaigns which used NHS branding. These attacks had a median attack availability of 13 hours. This was an increase from 2019, when we noted 36 campaigns, with a median attack availability of 4 hours.

In December 2020, we saw the first campaign that used the COVID-19 NHS vaccine rollout as the lure, and it differed because it did not seek to harvest NHS credentials. Like many other phishing attacks, the purpose of this attack was to harvest victim personal information for use in fraud. However, these attacks also undermine public confidence in the NHS vaccine rollout. As we moved into 2021 we saw this trend increase with campaigns delivered by email and SMS as fraudsters attempted to capitalise on the vaccine rollout.

We also looked for fake or unofficial copies of the NHS Test and Trace app. We took down 43 instances of NHS apps which were hosted and available for download outside of the official Apple and Google app stores.

Case Study 2: TV Licensing

In last year's paper, we noted that TV Licensing (TVL) phishing increased on par with the levels of abuse that HMRC experiences. We noted these attacks reached a peak that corresponded with news of changes to TV Licensing entitlements for UK pensioners during July 2020, which could be because attackers were trying to exploit the situation. As shown in Figure 5, plotting the sum of TVL-themed phishing mail servers with TVL-themed phishing campaigns over 2020 shows a clear peak around July.

Figure 5. Mid-year peak in phishing campaigns targeting TVL, 2020

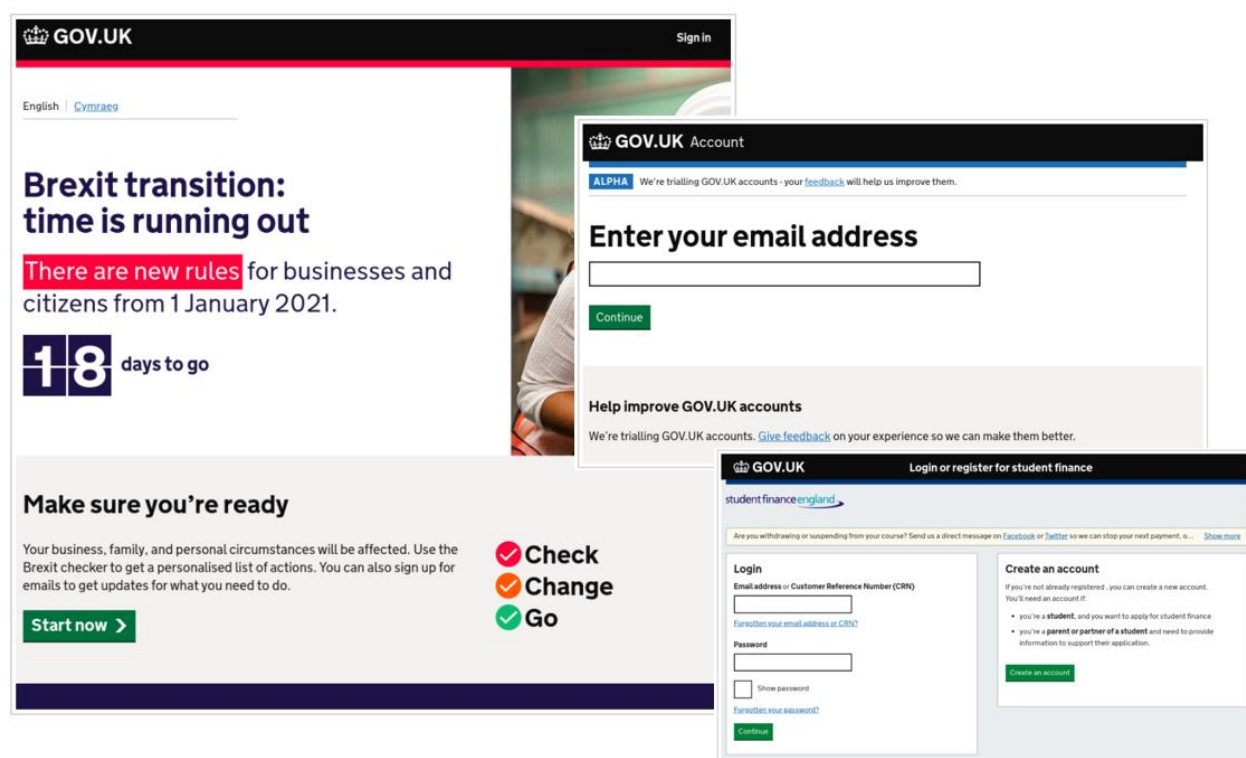


In total, we took down 3,604 TVL-themed phishing mail servers and 3,035 TVL-themed phishing campaigns in 2020.

Case Study 3: Brexit-themed UK government phishing

The level of Brexit-themed UK government phishing was surprisingly low during 2020, perhaps eclipsed by the coronavirus pandemic. However, on the 14th of December the Takedown service discovered such an attack that cloned large parts of the official gov.uk website. The attack was hosted on grant-gov.com and victims were directed to fill out various forms to sign up to various services; all pretty standard phishing idioms. However, on this occasion the site acted as a reverse-proxy, whereby it set up real accounts for victims on gov.uk for a number of services, which included Brexit transition, Student Finance England, and Universal Credit. It was a clever way of creating and harvesting many different types of credentials for UK government services.

Figure 6. Example of a phishing website that cloned gov.uk content to create and harvest credentials for various UK government services



Thankfully this attack was taken down promptly and additional steps were taken to warn relevant departments that accounts created via the grant-gov.com IP address should be flagged as compromised.

Use of SSL certificates in UK government-themed phishing

We continued to monitor HMG phishing takedowns to determine which of them used SSL certificates. As in previous years, we noted that a high proportion of the attacks leveraged Domain Validated (DV) cPanel (8,154) and Let's Encrypt (15,149) certificates, but the largest rise in 2020 was found to be phishing sites using Sectigo (15,620) certificates. This increase was unusual, and we were not able to determine the reason behind it.

UK government sending malware

The service took down a total of 2,890 campaigns during 2020, a fall from 3,473 in 2019. Typically, these campaigns send malicious attachments or embed clickable URLs that will attempt to distribute malware. We have not seen significant numbers in this takedown category since 2017 (see the [ACD - One Year On](#) paper for more information), which suggests that the application of DMARC, DKIM and SPF can consistently reduce domain spoofing.

UK government advance fee fraud

Advance fee fraud attacks using UK government lures continued during 2020, with a total of 2,310 campaigns taken down. Similar to previous years, the National Lottery was the most popular UK government brand targeted by fraudsters.

Table 6. Top 10 most targeted brands of UK government-themed advance fee fraud attacks, 2020

Government Brand	Number of attacks (URLs)
National Lottery	659
Bank of England	484
Financial Conduct Authority	451
Ministry of Justice	418
British Broadcasting Corporation	114
Metropolitan Police	61
HM Revenue & Customs	30

Government Brand	Number of attacks (URLs)
Department for Exiting the European Union	25
National Crime Agency	24
HM Treasury	24
All UK government themed AFF attacks (total)	2,310

UK government deceptive domains

As in previous years, we continued to monitor suspicious or deceptive domain registrations that could be used to impersonate UK government departments and services. Examples of this sort of domain include www.govuk-councilrefund.com and ukcontactnumberservices.co.uk, both of which were configured with subdomains that were used in phishing attacks during 2020.

We discovered 1,016 UK government-themed attacks on domains that were purposely registered by and under the control of the attacker. In these scenarios we only contact the hoster or domain registrar, as the webmaster or site owner is likely to be the attacker themselves.

A further 73 attacks were found on domains where we could not be sure that the attacker was in control, and so we contacted the site owner.

Takedowns in UK delegated IP space

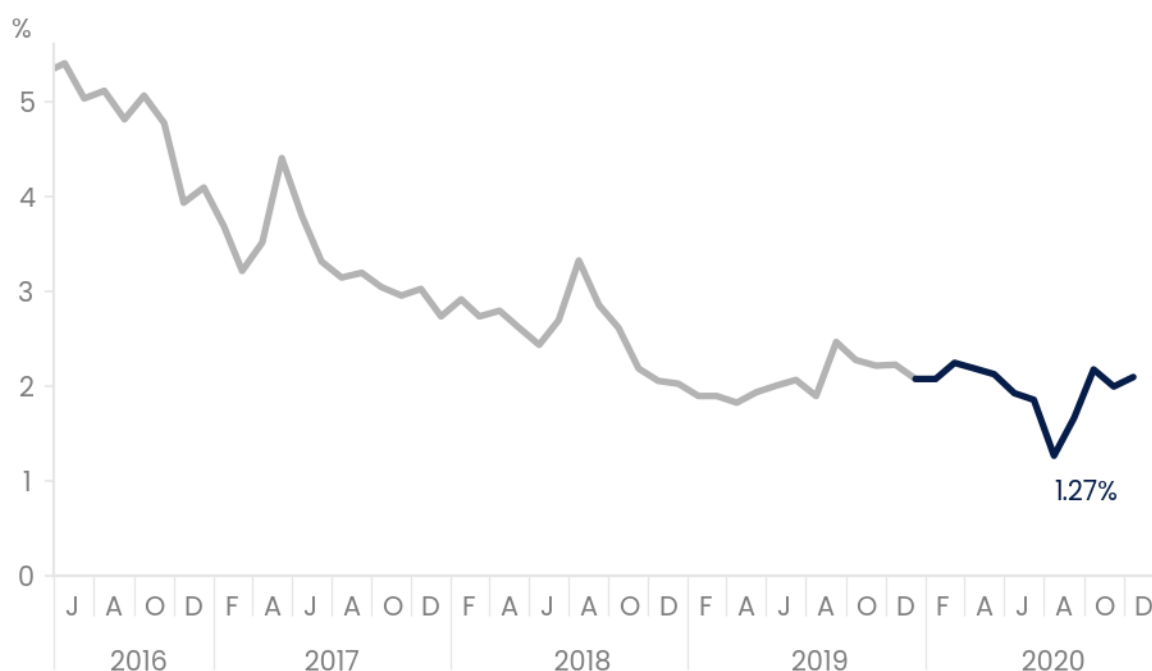
In this section we will look at the types of attacks we see in UK-delegated IP space. There are some new categories of attack here that are covered in more detail later, such as fake shops. As in previous years, we have continued to take down attacks hosted in the UK regardless of their brand. Unsurprisingly, phishing remains the most common attack type hosted in the UK.

Table 7. Brand agnostic UK-hosted attacks by type, 2020

Attack Type	Number of attacks (URLs)	Number of attack groups (campaigns)	Median Availability (hours)
Phishing URL	122,109	17,947	15
Fake shop	6,263	3,542	255
Web shell	5,963	1,532	21
Web-inject malware	1,490	1,082	66
Shopping site skimmer	2,634	973	94
Fake pharmacy	3,317	784	11
Malware distribution URL	1,250	735	12
Fake celebrity endorsement scams	1,394	412	15
Malware infrastructure URL	1,451	306	12
Support scam	934	287	28
Cryptocurrency miner	267	96	79
Malware command and control centre	74	57	19
JavaScript resource	50	46	350

Brand agnostic UK-hosted phishing

In August 2020, the UK share of global phishing hit 1.27%, which is the lowest we have seen since the Takedown service began. The number of phishing campaigns hosted in the UK have largely stayed relatively static since 2016 as shown in Figure 7. UK share of global phishing, Jun 2016 - Dec 2020, 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 2020

In 2020, we took down 17,947 phishing campaigns hosted on UK IP addresses, a total of 122,109 URLs. The median availability of these attacks was 14.5 hours, with 58% down within 24 hours of discovery. Comparatively, in 2019 we took down 18,202 phishing campaigns hosted on UK IP addresses, a total of 155,319 URLs, with a median availability of 12.0 hours, and 63% were down within 24 hours.

Table 8. Comparative brand-agnostic UK-hosted phishing attack availability, 2019-2020

Measure	2019	2020
Mean (hours)	121.3	122.7
Median (hours)	12.0	14.5
Skewness	8.5	8.4
25 th percentile	1.7	2.2
75 th percentile	50.1	57.5
Down in 4 hours	35.2%	31.9%
Down in 24 hours	62.9%	58.3%

Web injects in UK IP space

We took down 1,082 instances of web inject code in 2020, comprising 1,490 URLs. Median is up to 66 hours, from 36 hours in 2019, but the number of sites compromised in this way is down from the 2019 figure of 1,823 sites.

Shopping site credit card skimmers

We continued the work begun in 2019 to take down credit card skimming malware, and we noted some improvements regarding attack median availability. In 2020 we took down 973 skimmers that were hosted in the UK, and a further 477 skimmers hosted overseas on ecommerce sites that offered transactions in UK sterling.

For the UK-hosted skimmers, the attack availability median fell to 94 hours (from a 2019 value of 109 hours) and the overseas hosted median fell to 206 hours (compared to 318 hours in 2019). This shows that these attacks are being remediated much faster both in the UK and on sites hosted overseas. As with last year, the vast majority of skimming malware exploited the Magento eCommerce platform.

Cryptocurrency miners in 2020

In 2020, we took down 96 campaigns of non-consensual cryptocurrency miners running on UK-hosted websites with a median attack availability of 79 hours. 67 of the campaigns mined cryptocurrency using CoinImp embedded code and keys.

Poisoned JavaScript libraries

Sometimes, attackers will "poison" a software library with malware, so that the malicious code will infect the systems of any organisation that uses the library. We took down 65 instances of poisoned JavaScript libraries, with a median attack availability of 197 hours.

Web shells and control panels

Between January and October 2020, the service took down 1,870 web shells (6,745 URLs). Some shells were found alongside UK government-themed phishing (7%), and others in isolation on compromised sites (14%), but the vast majority were found in UK-delegated IP addresses. However, having attracted extra funding from NCA, we doubled down on shells and panels that were associated with phishing and took down as many as we could identify, regardless of the brand of the attack they were associated with or hosting location. See the [active fraud defence](#) section for further details.

The COVID-19 response and new types of takedown

In late February 2020, the NCSC was evaluating a number of new takedown initiatives for the new financial year. The final decision on which initiatives would be implemented was made in March, just as the UK was entering the first lockdown.

The NCSC noted the link between commodity cyber crime and subsequent fraud that may follow a breach of credentials or personal information through phishing or similar attacks. In an attempt to do more to prevent this, we decided to see whether the Takedown service could lower the value proposition for other types of high-volume internet-enabled fraud.

So, we began takedowns against the following types of scams:

- **COVID-19 themed cyber crime:**
 - phishing
 - malware
 - fake shops
 - advance fee fraud
 - vaccine fraud
- **Fake online shops** - websites offering heavily discounted goods, which either:
 - were hosted in the UK, or
 - offered transactions in UK sterling (hosted anywhere)
- **Fake celebrity endorsement scams** - scams using bogus content and fake endorsements from well-known figures, hosted on sites that claim to be UK newspapers or similar publications.
- **Remote access trojans (RATs)** - malware infrastructure takedowns.
- **Banking trojans** - malware infrastructure takedowns.

Table 9. Newly targeted attacks by type, Mar-Dec 2020

Attack Type	Number of attacks (URLs)	Number of attack groups (campaigns)	Median Availability (hours)
COVID-19 themed cyber crime	33,313	29,959	25
Fake shops	213,147	134,755	354
Fake celebrity endorsement scams	729,686	285,801	32
Remote access trojans (RATs)	2,954	1,733	39
Banking trojans	39,255	6,303	31

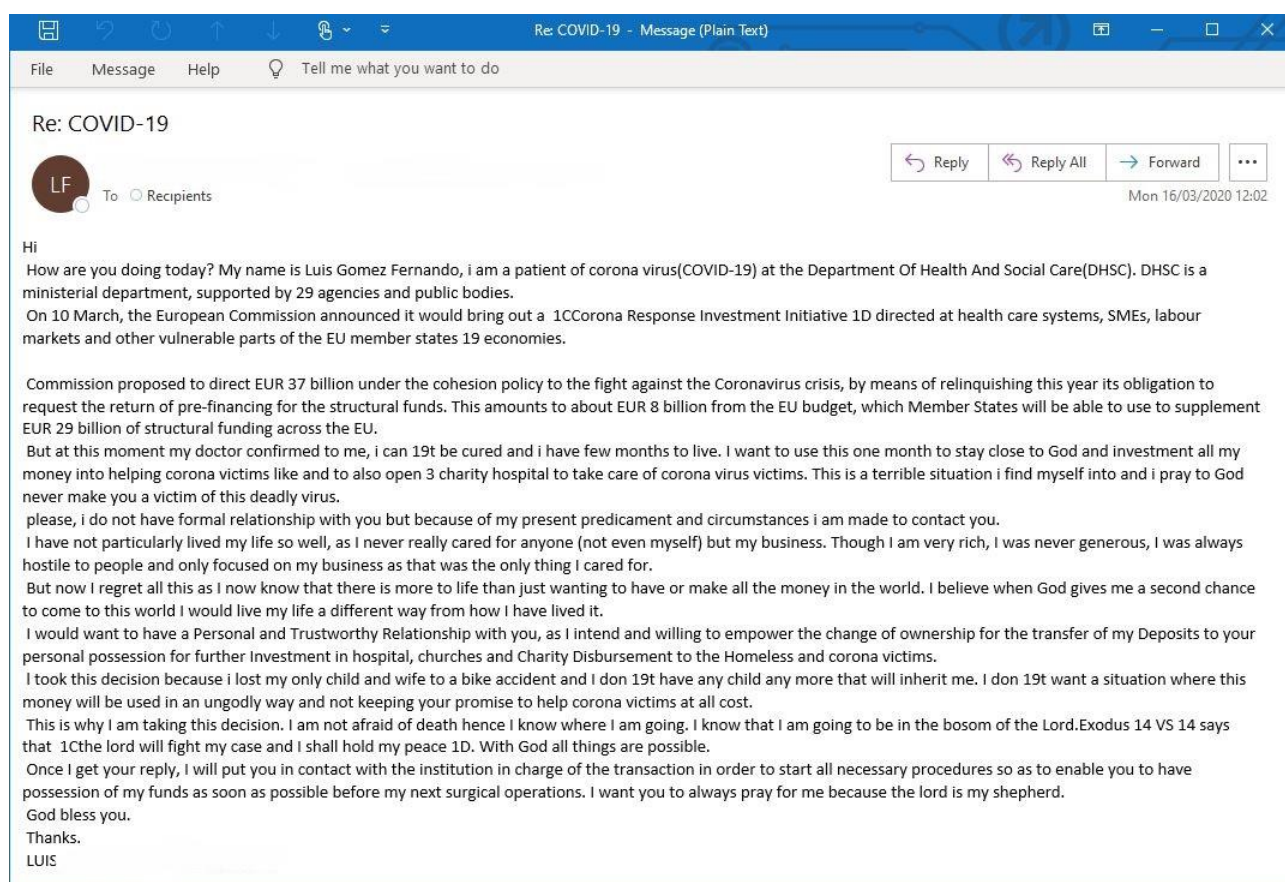
COVID-19 themed cyber crime

One of the methods of discovering commodity cyber attacks, such as phishing, is to deploy regular expressions (regex) to analyse spam mails that match known patterns and expressions. This is one method we used to discover COVID-19 themed cyber crime, and unsurprisingly there was a lot of data to inspect. Terms such as 'coronavirus', 'pandemic', 'covid' and so on were commonly found in spam email content.

Over the year the types of attacks that used COVID-19 themes varied; however, it was particularly attractive to 419 scammers, and so we found advance fee fraud to be the most popular attack in this category.

Our COVID-19 response began with this, our first takedown of an advance fee fraud attack on 16 March:

Figure 8. The first COVID-19 takedown, an advance fee fraud attack



Despite the challenges of removing false positives and authenticating these attacks, we took down 29,959 COVID-19 themed attack groups (comprising 33,313 URLs) between March 2020 and the end of the calendar year, as shown in Table 10.

Table 10. COVID-19 themed attacks by type, Mar-Dec 2020

Attack Type	Number of attacks (URLs)	Number of attack groups (campaigns)	Median Availability (hours)
Advance fee fraud	18,547	18,547	9
Malware attachment mail server	4,930	4,930	25
Advance fee fraud mail server	2,220	2,220	25
Phishing URL mail server	1,742	1,742	25
Fake shop	2,943	1,225	79
Malware infrastructure URL	1,844	411	52
Phishing URL	339	269	11
Malware distribution URL	350	242	70
DKIM signed email domain	133	133	68
Malware command and control centre	127	102	25

Fake online shops

These sites typically offer outlandish discounts on popular items to attract victims. They do not map to a real business and if a victim were to try to purchase an item, they would probably be charged for counterfeit goods or receive nothing at all.

Between April 2020 and the end of the calendar year, the service identified and took down 139,522 fake shops (222,353 URLs). Hosters were slow to remove these attacks and we noted a high median attack availability of 341 hours. The top 10 hosters of fake shops are shown in Table 11.

Table 11. Top 10 hosters of fake shops, Mar-Dec 2020

Hoster	Share (%)	Median availability (hours)
Fibergrid Group	29.4	463
Cloudflare	15.8	122
S.C. Parfumuri-Femei.com	9.9	1571
Inter Net Bilgisayar	8.6	869
Iliad	7.9	40
Unknown	3.8	426
Alibaba Group	2.8	783
Gigahost	2.7	42
Psychz Networks	2.4	2
Scaleway	1.4	747

Notably, we see a very different set of hosting companies in this category when compared to phishing or malware hosting.

Figure 9. Example of a fake shop offering discounted clothing

The North Face Jackets Online | Save Up To 50%
<http://www.safeplace.org.uk/>

Wish List (0) Checkout £ GBP Login Register

safeplace.org.uk Search 0 item(s) - £0

HOME DELIVERY INFORMATION PRIVACY POLICY TERMS & CONDITIONS

Categories

- > The North Face 100% Authentic
- > The North Face Big Discount
- > The North Face Cheapest
- > The North Face Classic Style
- > The North Face New Arrival
- > The North Face Recommend

Latest

- Zoomie Jacket in Baby Blue. £39 ~~£106~~
- Zipline Rain Jacket (Toddler) £37 ~~£130~~
- Zaneck Insulated Parka in Black £38 ~~£120~~

New product

- Vestes & Vestes sans manches | Mountain Jacket 1985 Seasonal £37 ~~£110~~
- THE NORTH FACE Women's Arctic Parka Down Coat Black Green £37 ~~£116~~
- Himalayan Hoodie Medium Grey Heather £38 ~~£118~~

The fake shop takedowns began shortly after the UK's first COVID-19 lockdown and we noted over 1,200 fake shops selling fake PPE and secret "cures".

Figure 10. Example of a fake shop selling N95 Masks

N95 Coronavirus Mask For Sale, Medical Face Mask
<https://www.n95coronavirusmaskforsale.com/>



Figure 11. Example of a fake shop selling a 'coronavirus cure'

<https://coronavirus-cure.store/>





Figure 12. The rise and fall of COVID-19 themed fake shops at the start of the COVID-19 pandemic, 2020

Fake celebrity endorsement scams

In April, we began performing takedowns against a relatively new type of scam, which takes the form of web pages or online articles that purport to be from well-known publications and feature celebrities or other figures. The articles themselves feature fake endorsements for cryptocurrency investment schemes, which are linked within the article. Though the theme of these attacks is cryptocurrency, the template could be used to promote any type of fake investment opportunity and we expect this theme to change over time.

We found that links to these fake articles were heavily promoted in mass mail campaigns and also via SMS and online adverts on many websites. Much of the article content is broadly similar from URL to URL, and is distributed widely across the internet. We also took down large numbers of redirect URLs which sit between the links found in spam and the end content (article).

By distributing the final landing page content widely and by obfuscating the path to them via the redirects, the actor hopes to disrupt takedowns and keep the attack campaigns active longer.

As a visitor to these online articles, the content served will vary depending on your IP address and location. For example, UK visitors will see curated content with recognisable figures such as Sir Richard Branson (as part of a fake Daily Mirror article), but if you were to view the same article from France, the content served would likely appear as an article in a fake French newspaper featuring a French celebrity. The brand (or celebrity) abused is dynamic in nature, and customised for the site visitor. This technique could make pursuing brand infringing content challenging. By serving customisable content based on location, the scam is very much a global issue and not exclusive to the UK.

Figure 13. Examples of fake celebrity endorsement scams



Note: to be absolutely clear, the people whose images are used in these scams are not aware of their use nor involved in the scam.


Between April 2020 and the end of the calendar year, we took down 286,322 campaigns (731,080 URLs) of this type, with a median attack availability of 32 hours.

Table 12. Top 10 hosters of fake celebrity endorsement scams, Mar-Dec 2020

Hoster	Share (%)	Median availability (hours)
Amazon	46.4	19
OVH	5.8	292
DigitalOcean	5.8	11
Cloudflare	5.2	123
Google	4.3	340
GoDaddy	2.0	347
Endurance International Group	1.1	36
Alibaba Group	0.7	156
NameCheap	0.7	245
Velocity Servers Network Exchange	0.1	125

The median attack availability for these investment scams has steadily improved throughout 2020, as shown in Figure 14.

Figure 14. Steady improvement, following an initial peak, in median availability of fake celebrity endorsement scams, 2020




Remote access trojans

Remote access trojans (RATs) are used by attackers to conduct more detailed reconnaissance of a victim device. They support several covert capabilities that a victim will not be aware of, such as keyloggers, screengrabs of documents, and password theft. These provide a backdoor to enable other capabilities, which an attacker could subsequently deploy.

In April we began takedowns against RATs, starting with UK-hosted infrastructure, expanding soon after to those on overseas-hosted infrastructure. By performing extended malware analysis of RATs and scanning IPV4 IP space for RAT Command & Control (C2), we were able to take down 1,733 RATs (2,954 URLs), with a median attack availability of 39 hours.

The Agent Tesla RAT was the most commonly found trojan, with 36% of RAT campaigns associated with it, as shown in Figure 15.

Figure 15. Malware families associated with RATs, Mar-Dec 2020




Banking trojans

Upon installation on a victim's device, banking trojans (as the name suggests) harvest banking credentials and exploit financial transactions. Ultimately, a victim could have their bank accounts cleared out by an attacker after opening a malicious file. The Takedown service disrupts these attacks by taking down the infrastructure that the hosts the malware.

Between April and end of December 2020, we took down 6,303 banking trojan campaigns (39,255 URLs), with a median attack availability of 31 hours.

We found that 75% of the URLs were identified as associated with the Emotet family of malware. Subsequently, Emotet was subject to coordinated international takedown efforts during January 2021, spanning government, law enforcement and industry (www.bbc.co.uk/news/technology-55826258) and 2021's data will no doubt reflect this.

Figure 16. Malware families associated with banking trojans, Mar-Dec 2020



Active fraud defence (from Nov 2020)

Following the relative success of the other counter-fraud initiatives, we worked with the National Economic Crime Centre, National Fraud Intelligence Bureau, the National Crime Agency, and the Financial Conduct Agency on a number of new takedown initiatives to further combat fraud. The hope was to tackle high volume fraud and disrupt the types of crime reflected in public reporting to Action Fraud and other parts of UK law enforcement with a number of trial takedown initiatives.

Table 13. Active fraud defence attacks by type, Nov-Dec 2020

Attack Type	Number of attacks (URLs)	Number of attack groups (campaigns)	Median Availability (hours)
Computer software support fraud	3,575	1,120	15
Clone firms subject to FCA investment warnings	400	367	74
UK telephony numbers found in advance fee fraud	181	181	15
Extortion mail server takedowns	179,008	179,008	26
Additional web shell countermeasures	11,820	2,990	37
Additional takedowns via SERS	11,302	4,452	7

Computer software support fraud

This type of scam typically involves a technical support theme, where victims are tricked into thinking they are interacting with official technical support staff of well-known brands and corporations. The attacker will offer to fix or provide software to remediate a problem that they claim to have found on a victim's computer. The scams use a combination of hosted content, remote access tools, and telephony to trick victims, impersonate tech support, and gain access to victims' computers.

Between November and the end of December 2020, the service took down 1,120 support scams (3,575 URLs) with a median attack availability of 15 hours. This was in addition to a further 287 campaigns we took down that were hosted in UK-delegated IP address space earlier in 2020. In terms of the impersonated brands, we found Microsoft, Apple, McAfee, and Google to be the most prevalent.

Clone firms subject to FCA investment warnings

In another initiative, the service looked at fake investment firms that had been subject to investment warnings by the UK Financial Conduct Authority. The FCA regularly publishes such warnings on their website (<https://www.fca.org.uk/scamsmart/warning-list>). This initiative sends notifications in a bid to take down a clone firm's fake site, its mail domains and servers, and any telephony connected to them.

Between November and end of December 2020, we took down (by group):

- 130 clone firm email addresses
- 118 clone firm phone numbers
- 117 clone firm URLs
- 2 fake investment banking sites

UK telephony numbers found in advance fee fraud

Analysis of advance fee fraud mails over time has shown a prevalence of UK landline and mobile numbers in use by attackers. This initiative sought to identify campaigns that used +44 numbers and the service provider concerned. We would then contact the service provider to notify them that the number is being used for fraud. Between November and end of December 2020 the service took down 181 phone numbers involved in advance fee fraud, the median availability of which was 15 hours.

Extortion mail server takedowns

Extortion mails seek to trick potential victims into believing their device (a phone, tablet or PC) has been compromised by a sophisticated hacker, who then seeks to extort victims by threatening to expose their online habits and personal information to all of the victim's contacts.

The extortion mail seeks further authentication by customising the email with "evidence", such as a password or phrase that the victim may recognise as their own. Such information is often derived from credential breaches and is therefore recognisable by the victim. The emotional impact of this is high, and for some will tip the balance toward paying the ransom, usually via anonymous payment to a cryptocurrency wallet referenced in the mail.

Between November and the end of December 2020, the service identified and took down 179,008 servers sending extortion scams. The NCSC published guidance on [how to protect yourself from such scams](#) in 2018.

Additional web shell countermeasures

The additional web shell and panel countermeasures implemented in November produced a dramatic uplift in the number of takedowns in this category. Between November and end of December 2020, the service took down 2,990 web shell campaigns (comprising 11,820 URLs).

Additional takedowns via the Suspicious Email Reporting Service

With NCA funding, we also performed additional takedowns against attacks found in public reporting to the Suspicious Email Reporting Service (SERS). We found that some attacks which targeted certain brands were not being dealt with promptly, which indicated that there was either no brand protection or the brand protection service responsible was not aware of the particular type of attack. Between November and the end of December 2020, we took down an additional 4,452 attack groups (11,302 URLs), of which 96% were phishing campaigns. The top three targeted brands during this time were PayPal, DPD Express Parcel Delivery, and Track-Trace (a parcel tracking site).

Reporting from the UK public to 7726

Last year we performed a limited trial whereby we injected URLs from public reports to the 7726 short code number into the Takedown service. We also purchased a feed of malicious SMS URLs from an industry partner to enhance our discovery of SMS campaigns.

In April 2020 we accelerated the rate at which we processed the URLs in 7726 SMS phishing reports. Between April and end of December 2020, these referrals were credited as the first reporter of over 22,000 URLs in the Takedown system. It has been a useful source of reporting throughout the year.

Reporting from the UK public to the Suspicious Email Reporting service

In late April 2020 the NCSC launched SERS, and referrals from the UK public created thousands of NCSC sponsored takedowns. SERS was credited as the first reporter for over 62,000 URLs in the Takedown system. In fact, approximately 32% of email submissions to SERS were classified as malicious by the Takedown service. See the [SERS](#) section for more details.

Conclusion

The service has delivered more takedowns in 2020 than all the previous years combined, which we hope has reduced the potential harms that malware, phishing and other scams could inflict on UK citizens. We also hope that 2020's new initiatives have genuinely lowered the value proposition for internet-enabled fraud in the UK (or that targets UK citizens). We're continuing to engage directly with hosting companies and other responsible organisations who can assist in taking the malicious sites down quickly and efficiently.

Mail Check

About the service

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.

In particular, Mail Check supports organisations in implementing the following controls:


- **Email anti-spoofing controls (SPF, DKIM and DMARC):** These standards help prevent various attacks (for example, phishing) that use an organisation's email domain to trick email recipients.
- **Email confidentiality (TLS):** Keeping messages encrypted and private as they are sent over the internet.

Progress

Maturing Mail Check to support over 1,000 organisations

Throughout 2020, we grew the Mail Check services from coverage of 800 to over 1,200 organisations. Much of the growth came from universities, colleges and schools signing up, while there was moderate growth in the more established customer sectors, as shown in Figure 17.

Figure 17. Increase in organisations signed up to Mail Check, highlighting academia, which saw the greatest growth, 2020



To achieve this growth, a significant effort has gone into maturing the Mail Check capability. The service went from an alpha capability, to beta, and finally to a live service throughout the year. The work to achieve this included:

- reducing the burden on first-line support by introducing self-service features for users.
- focused work on accessibility to work towards an AA rating.
- improving reliability.
- maturing our monitoring and support procedures.

Campaigning in 2020

Building on campaigns already run in both central and local government, we focused on making an impact in the following sectors:

- Academia; notably universities, further education colleges, and research institutes.
- Territorial police forces, and fire and rescue services.


Academia

We worked closely with the NCSC Economy and Society Engagement Team and JISC to get the message out and get organisations onto the platform. This involved communicating through a range of JISC channels (mailing lists, forums, blogs, twitter), as well as NCSC channels (Twitter, CiSP). The campaign ran alongside very significant transformation in these organisations as a result of COVID-19, in particular the move to online learning. Several institutions in the sector experienced cyber attacks in 2020, many of which were reported in the media.

From April to Dec 2020, we completed the first phase of the campaign, having onboarded 223 organisations: 116 universities, 72 further education colleges, and 35 other academic institutions, such as research institutes. We found that universities have a particular challenge when implementing DMARC anti-spoofing controls: their large sprawling IT estates mean that there is often no central view of which email sending systems need to be configured, and most have between 10-15 systems to identify and configure before progressing to an enforcing DMARC policy of quarantine or reject. We put additional effort into setting up coaching sessions with university teams to coach them through the technical and non-technical challenges of implementation in their organisation.

Size matters. Whilst we have fewer further education colleges on the platform, we have seen significantly swifter progress in implementation of anti-spoofing controls for them, as their IT estates are less complex. We are seeing early signs of a shift in the sector since we began the first phase in April 2020. At that time, 10% (17 out of 164) of universities had an enforcing DMARC policy of quarantine or reject on their main email domain. By December this had risen to 15% (25 of 164), as shown in Figure 18.

Figure 18. Increase in universities using a DMARC enforcing policy, Apr-Dec 2020



In July 2020 we confirmed that 13% of further education colleges had an enforcing DMARC policy of quarantine or reject on their main email domain. Tests run early in 2021 indicated that this number had risen to 18%.

Police forces, and fire and rescue services

In April 2020, we recognised a need to campaign for implementation by police forces, and fire and rescue services. Many of the organisations involved had signed up to Mail Check in previous years, but the sectors were in need of re-invigorating and supporting in the challenge to implement anti-spoofing controls.

The campaign was initiated with support from the NCSC engagement team and key influential sector bodies, such as the National Police Chiefs' Council (NPCC) and the National Fire Chiefs Council's (NFCC) ICT Management Forum. Thereafter, we focused our efforts on one-to-one engagement; sharing knowledge and coaching organisations in virtual sessions on both the technical and non-technical challenges in delivering strong email security. This included bringing together police forces and their outsourced IT providers to tackle the problem collaboratively.

The uptake by the sector was very rapid; those organisations that hadn't signed up did so promptly, and once on board initiated the activities required to deliver anti-spoofing quickly.

We saw a notable shift in DMARC adoption for emergency services, from 49% coverage in May 2020 to 67% by the end of the year.

Comments from our campaign sectors

"I feel it is a very positive move that the government have chosen to offer security-related products such as this to the public sector. Without these kinds of initiatives, it is unlikely we would've been able to afford to implement the service. Overall, this can only be a good thing to improve trust in the public bodies in the UK."

"Thank you for the information and all the work that's gone into assessing this with us, we found this informative and enormously helpful. We are a relatively small infrastructure and security team, and so the assistance and guidance you have provided here is very much appreciated. Mail Check is a very useful tool, helping to simplify and expedite the investigative process. We have found the service easy to use and well laid out."

"Many thanks for that meeting just now. Learnt more in 40 minutes about DMARC and email security in general than I have in a couple of years. Really appreciate your time and effort. Thank you."

Email spoofing still very much a threat


Synthetic DMARC still effective

In the [ACD - The Second Year](#) paper, we introduced 'Synthetic DMARC': wildcard SPF and DMARC records applied to *.gov.uk. This protects against spoofing of non-registered domains; for example, tax.gov.uk. The statistics we gathered demonstrated that this protection is both effective and that attackers would take advantage of such domains if they weren't protected.

Figure 19 shows that there is, in general, a low level of email hitting these DMARC checks; an average of less than 100 a day for most months of 2020. Some of this appears to be misconfigured automated email, but not all of it. So while there are some attempts at low-volume, likely targeted spoofing (which failed due to Synthetic DMARC), there's little evidence of large-scale campaigns.

There is a spike in September 2020, which indicates what appears to be a single campaign on a single day of roughly 30,000 messages, each using a unique non-existent subdomain of gov.uk. None of those messages were delivered to the intended inboxes, and the attackers did not try again the following day. This suggests that the attackers understood that the gov.uk domain is protected against this kind of abuse.

Figure 19. Spike in spoofing of non-registered '*.gov.uk' domains in September but otherwise remained consistently low, 2020



Local authority spoofing

Following adoption of Mail Check and successful campaigning by the NCSC, 76% of local authorities (principal councils) are protecting their domains with an enforcing DMARC policy of quarantine or reject. Figure 20 illustrates mass spoofing against a local authority domain, with spoofed emails in this case failing DMARC checks when received and being put into junk folders. These emails originated from a large number of countries; 22 countries listed in Figure 21 represented 80% of the spoofed traffic.

Figure 20. Spoofed emails failed DMARC checks during mass spoofing against a local authority domain while legitimate emails pass, Jan-Jun 2020





Figure 21. Twenty-two countries represented 80% of the spoofed traffic in mass spoofing against a local authority domain, Feb-Mar 2020



Spoofing of university domains

In 2020, we extended Mail Check to UK universities and other academic institutions. In October we conducted an analysis of spoofed emails for the first 40 universities to sign up, and found the following:

- An average of 20,000 emails per domain were sent from identified spam sources, or had failed other anti-spam tests, such as forward-confirmed reverse DNS (FCrDNS).
- 10% of the universities sampled were targeted by attackers sending emails from fake subdomains. These attacks sent emails from anywhere between 200 and 1,000 subdomains.

In the worst case, two universities in November 2020 saw simultaneous spikes in spoofed traffic, involving 1.69 million emails over a period of 11 days from 17,381 IP addresses, almost all of which were based in Russia. Typically, attackers limited the use of each IP address to a few days; 23% were active for just 1 day, whilst 67% were active for under 4 days.

In a separate case, when Mail Check monitoring was implemented by another university it was found that a significant proportion of emails sent from their domain were malicious. Furthermore, a significant proportion of these malicious emails were passing a traditional Sender Policy Framework (SPF) authentication check and had not been identified on a spam blacklist. After engagement with the NCSC, a DMARC policy of quarantine was put in place in the following weeks to better protect the university's domain.

In the two weeks before this change, 36,642 malicious emails were identified, which represented 30% of all email from their domain. Implementing an enforcing DMARC policy of quarantine appears to have had an immediate deterrence effect; in the two weeks after the policy was implemented only 5 malicious emails were identified and quarantined.

Introducing MTA-STS and TLS-RPT

We built in support for two new email security standards in Mail Check during 2020: SMTP Mail Transfer Agent Strict Transport Security (MTA-STS) and SMTP Transport Layer Security (TLS-RPT).

MTA-STS is a standard designed to address vulnerability in email TLS security, whereby a person-in-the-middle can trick incoming connections to send to their server and/or send information in the clear.

MTA-STS is designed to defeat these attack vectors by:


- advertising the MX server for a given domain on a separate secure web page, which means an attacker cannot just subvert the DNS entry to perform such an attack.
- enforcing TLS encrypted inbound communication; that is, if an attacker tries to get information sent in clear text the connection is refused.

TLS-RPT is a standard that gives feedback on whether TLS connections to your domain have been successful, and if not then why. This will give domain owners daily indications of issues, such as if there's a problem with TLS certificates. TLS-RPT is important because it provides the necessary feedback to domain owners before they tighten up policies, such as MTA-STS.

Supplier support for TLS-RPT


The market was still maturing throughout 2020 in support of MTA-STS and TLS-RPT. Figure 22 illustrates TLS reporting received in Mail Check throughout the year, almost all of which originated from Google. Microsoft announced support for the standard in October 2020.

Figure 22. Cumulative TLS reporting received in Mail Check, highlighting Google as the largest originator, 2020



Early growth in Mail Check domains adopting MTA-STS and TLS-RPT

We began offering TLS-RPT to Mail Check users, and by the end of 2020 it was applied to nearly 400 domains, as illustrated by the following chart. We plan to implement support for MTA-STS in 2021; currently, only 20 domains in Mail Check have adopted this standard.

Figure 23. Cumulative increase in Mail Check domains adopting TLS-RPT, 2020

Outcomes

Throughout the year we have seen a doubling in the number of domains with an enforcing DMARC policy (of quarantine or reject), as illustrated by the following chart:

Figure 24. Cumulative increase in Mail Check monitored domains with a DMARC enforcing policy, 2020

The following table illustrates sector-based shifts from the end of 2019 to the end of 2020. You can see the following trends:

- High levels of DMARC adoption in central government, local government, health and devolved administrations, and emergency services.
- Early growth in some sectors, such as universities.

Note that charities and critical national infrastructure (CNI) organisations were not eligible for Mail Check in 2020. We have included numbers for these sectors in the table to serve as a baseline should we decide to expand into them during 2021.

Table 14. Sector-based shifts in Mail Check monitored organisations with a DMARC enforcing policy, Dec 2019 - Dec 2020

Sector	Sub-set of organisations tracked	31 Dec 2019 % orgs with EDP*	31 Dec 2020 % orgs with EDP	Change	Comment
Central government	44 (government departments + 10 Downing Street)	66	86	+20	DExEU was dissolved in January 2020 while FCO and DfID merged to form FCDO in September 2020 - these figures treat these changes as if they had already happened by 31 December 2019
Central government	223 arm's length bodies	32	46	+14	Includes executive agencies, non-departmental public bodies, and similar organisations set out in Public Bodies 2019
Local government	404 (UK principal councils)	67	75	+8	
Health	279 NHS Trusts and key central functions	13	21	+8	Note that these numbers represent protection on nhs.uk domains. Many of these are no longer in use for email, as 80% of NHS trusts are now using NHS Mail (with the domain nhs.net) that is protected with an enforcing DMARC policy of 'reject'.
Devolved administrations and their agencies	Includes local authorities, health services and emergency services in Devolved Administration regions	29	49	+20	Averaged across Scotland, Wales and Northern Ireland
Police and fire services	51 police forces and 54 fire and rescue services	45	59	+14	Campaign focus for 2020 Includes organisations in Jersey, Guernsey and the Isle of Man
Universities	164 universities, university colleges and other degree-awarding bodies	10	16	+6	We began tracking DMARC uptake by universities in April 2020
Charities	Top 3000 charities	n/k	10	N/A	Not eligible for Mail Check in 2020
Private sector CNI organisations	Approx. 1000 organisations	n/k	25	N/A	Not eligible for Mail Check in 2020

*EDP = Enforcing DMARC Policy of quarantine or reject

Web Check

About the service

Web Check is a live ACD service that helps over 1,000 customer organisations identify and fix common security issues in their websites. Users can sign up on behalf of their organisation and specify URLs to be checked regularly for issues. The results of the scans are shared in the Web Check interface, together with appropriate and clear mitigation advice.

Progress

Evolution of the checks operated

At a higher level of abstraction, the types of security issues seen in websites have a tendency to evolve relatively slowly. In line with this, the majority of checks performed by Web Check were operated throughout the year, with regular maintenance where necessary.

By contrast, specific instances of these types of security issues occur more frequently as bugs appear in versions of software products run on web servers. It is generally impractical for Web Check to scan for specific vulnerabilities; the focus is rather on an entry level set of checks for general security concerns, with checks for product versions and patch levels used to encourage good security behaviours that reduce the risk of specific vulnerabilities remaining unaddressed.

However, where a vulnerability is considered particularly critical we include a specific check for it in Web Check. For example, in 2020 we included a check for the Tomcat AJP port vulnerability ([CVE-2020-1938](#)), which carried potential for remote code execution by an attacker.

A general improvement made to the service involved clearer presentation of recurring findings. It's important that Web Check's results are accurate and easy to understand. It is for this reason, for example, that one of our principles is to only include a check if we are confident that we can minimise the risk of false positive and false negative results from it.

We were therefore concerned at the presence of a small proportion of cases where findings were seen to repeatedly flip-flop between 'found' and 'not found' states from one day to the next. We felt this would cause many users to question the reliability of the service. In reality, there can be a number of explanations for why websites may exhibit this behaviour when scanned. For example, a site hosted on two or more load balancing servers with differing configurations may return different scans results over time, or caching effects may be responsible for this behaviour. This effect in itself can be a potential security concern.

Therefore, to give users the clearest possible presentation of results Web Check now labels flip-flopping findings as 'recurring', and gives a high level indication of what this might mean.

At times, an opportunity presents itself for one cyber defence initiative to enhance another. For example, adding a security.txt file to the root directory of a website containing details of how to contact the site owner in the event of a vulnerability is considered good security practice. This is recommended by the NCSC's [Vulnerability Disclosure](#) project, and therefore we have configured Web Check to look for this file and confirm whether it contains the required information.

Growing the user base

Throughout the year, Web Check's primary user base was to be found in public sector organisations. However, we did participate in two initiatives coordinated by the [ACD Broadening](#):


- A comms campaign was run to offer Web Check to universities and colleges of further education. Take up was good, and by the end of the year 91 such organisations had signed up.
- A pilot to investigate the value Web Check would deliver to charities was underway at the end of the year.

Support given by the NCSC to organisations playing key roles in the nation's COVID-19 response tended to be more bespoke than the Web Check offer of an entry level set of web security checks. However, in some cases it was considered appropriate to offer use of the service, with the organisations involved often being in customer sectors that would not normally be eligible to sign up.

Outcomes

The types of findings generated by Web Check are categorised by their severity, with Urgent and Advisory being the two most significant levels. The following chart shows the numbers of findings raised at these two levels.

Figure 25. Comparison of urgent and advisory Web Check findings by type, 2019-2020



The increase in numbers generally reflects a commensurate increase in the number of URLs being checked. As in previous years, issues are most frequently encountered in the following areas:


- Effective implementation of HTTPS and TLS protocols, supported by use of X.509 public key certificates, to give those visiting a website confidence both in the authenticity of the site and in the confidentiality and integrity of the data flows involved in communicating with them.
- Maintaining the version and patch level 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.

The biggest rise relative to 2019 (1,002 to 2,392) was in the area of RDP issues. These findings were reported to users when port 3389 was found to be exposed, leaving the website open to potential exploitation via a remote desktop connection.

The primary goal of Web Check is for customer organisations to take action in response to the findings presented to them, thereby improving the security of their websites. The following chart shows the number of Urgent findings resolved each month; these are issues detected as no longer existing, having previously been flagged up to our users. We have assumed that Web Check was instrumental in prompting the customer organisations to resolve them.

There is a natural degree of variation from month to month, but the broad picture is one of a significant number of issues being addressed, with this number rising in response to the rising number of findings being raised.

Figure 26. Urgent Web Check findings resolved per month remained fairly consistent, 2020



Conclusion

Web Check is an established ACD service, but work has continued to ensure ongoing relevance of the checks performed and the clarity of the resultant findings presented to users. Opportunities have also been taken to extend its reach into new customer sectors. Metrics demonstrate that it continues to deliver benefit in terms of security issues being addressed by the website owners.

Protective DNS

About the service

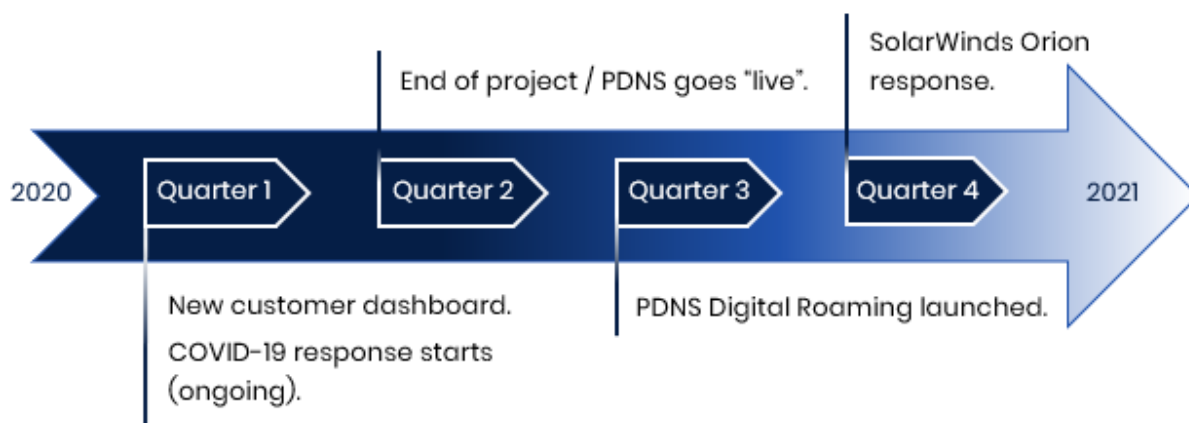
The Domain Name System (DNS) is the address book of the internet. Your computer relies on DNS to find out exactly where a domain (ncsc.gov.uk) is located (its IP address, for example 13.224.228.11) so it can connect to it. Anyone can register a domain so that everyone else can find the IP address associated 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.

The NCSC's Protective DNS (PDNS) service 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.

Progress

We were pleased that following the competitive tender in 2019, and against the backdrop of the first lockdown, we formally began the new contract with Nominet and moved PDNS to a fully live status. Improvements and updates will not stop or slow; this just marks a milestone where PDNS becomes a sustained service within the NCSC.

Figure 27. PDNS highlights of 2020



Here are some of the key things we achieved in the face of the challenges of 2020:

Onboarding healthcare organisations

Early in the year, our plan was to bring on board the Health & Social Care Network (HSCN), which comprises more than 1,000 organisations across over 12,000 sites in the UK. A plan that would roughly double the size of the PDNS user base. After that we would look at the wider health and social care sector, including NHS Trusts.

When lockdown began in March, we revised our plan. We began working with NHS Digital to onboard as many NHS Trusts as possible first, starting with those running Nightingale hospitals. This reprioritisation risked exceeding our capacity, meaning we would no longer be able to take on the HSCN. To mitigate this risk, we sought and secured funding for an additional capacity uplift (in less than 24 hours!), allowing us to pursue the onboarding of the NHS Trusts and critical healthcare sector organisations with immediate effect.

In October we returned our attention to onboarding HSCN organisations. In preparation for this we had been running capacity tests, "What if?" scenarios, and Service Desk rehearsals. After several false starts and a [Cybersecurity & Infrastructure Agency \(CISA\) alert](#) which warned that malicious actors were targeting US healthcare, we accelerated the process and implemented PDNS for these organisations within 24 hours.


At the same time, the maturity of PDNS allowed us to offer the service to the vaccine supply chain and Lighthouse Labs as part of the ACD Broadening project. This extended the protection of PDNS outside of our traditional base and saw us offering our service to private sector organisations for the first time.

As a result of our efforts and the increased capacity available to us, throughout the year we saw much stronger growth in the health sector than others, and now the majority of NHS organisations are actively using PDNS.

Onboarding other organisations

Besides HSCN, 302 organisations started using PDNS in 2020. By the end of the year, 799 organisations were using the service, 60% more than at the start of the year, as shown in Figure 28.

Figure 28. Cumulative increase in organisations protected by PDNS (excluding HSCN), 2020




Uptake of PDNS by government departments increased from 75% (33 of 44 organisations) to 84% (37 of 44) in 2020. Note that the number of government departments has changed since 2019 due to the closure of Department for Exiting the European Union (DExEU) and the merger of Foreign & Commonwealth Office (FCO) and Department for International Development (DfID) to form FCDO. We can now count 23 ministerial departments and 20 non-ministerial departments as onboarded and live!

Of the 404 local authorities in the UK, the number using PDNS increased from 265 at the start of 2020 to 304 at the end - corresponding to growth in coverage from 66% to 75%.

Usually, as the number of organisations using PDNS increases, so does the number of DNS requests we handle. However, as you can see in Figure 29 there was a noticeable decrease in the number of DNS requests we handled during the first COVID-19 lockdown. As organisations around the UK rallied we saw a quick recovery, and the upward trend resumed.

Figure 29. Dip in PDNS requests in March 2020 (first COVID-19 lockdown) but otherwise increasing consistently month on month, 2019-2020



It was here that we realised that this dip may have been due to the dramatic shift towards home working, and that customer network architectures may not have automatically supported PDNS protection of remote devices. It was essential we offer customers a way to extend PDNS to home working or roaming scenarios.

PDNS Digital Roaming

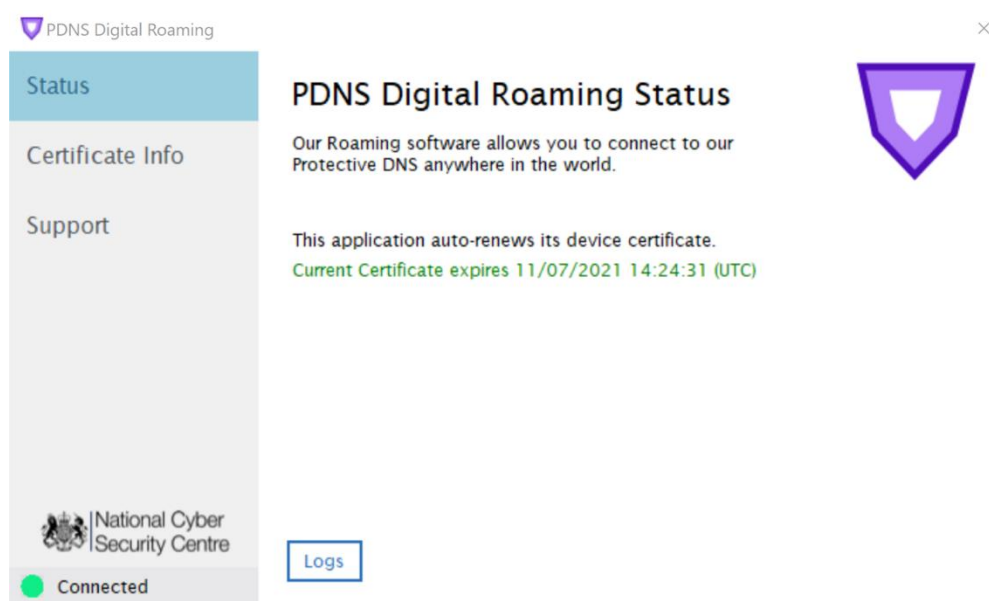
In 2020 we launched PDNS Digital Roaming.

Digging deeper into use cases, we found not all of our customers have always-on VPNs for off-site working, which would route DNS back into the enterprise network. In fact, as organisations move to cloud-managed enterprise networks PDNS can become harder to implement through traditional means (there is no enterprise recursive resolver to redirect DNS to our service!).

The PDNS Digital Roaming app is designed to route a client's DNS traffic to PDNS when they're not connected through a traditional enterprise network. It's secure, encrypted, and deployable at massive scale through common mobile device management (MDM) solutions. It also has the added benefit of recording blocks with machine-level resolution, to aid incident handling when searching for the source of any malware found.

Failure modes are permissive by default, which means end users will not be disrupted logging into captive portals (that is, the coffee shop login portal scenario), and all users will see is the PDNS 'shield' tray-icon indicating that PDNS protection is active. Clicking it will show users a little detail on the status of their protection, but no user input is required to enable protection.

Figure 30. PDNS Digital Roaming Client



Our thanks to the 12 customers who answered our call to be part of early development and testing. We will continue to work on our Windows app, adding new features and security, but will also be offering support for other common enterprise operating systems in 2021.

Take-up has been excellent so far. NHS Trusts and organisations across central and local government are already using it to extend the benefits of PDNS to areas of their IT estates that could not previously be reached, with more organisations planning implementation in 2021.

We also see customers plan their future, cloud-first enterprise networks, pleased that PDNS will continue to protect them without the need for a traditional recursive DNS resolver, and in a way that works alongside common cloud-based security and monitoring tools.

Protecting against malware

In 2020, PDNS handled more than 237 billion DNS requests. Of these, nearly 105 million requests were blocked, corresponding to 0.04% of all requests. These 105 million blocked requests were for nearly 160,000 distinct domains attributed to cyber crime Organised Crime Groups (OCGs) with ransomware-related malware featuring prominently, as shown in Table 15.

Table 15. Malware types associated with PDNS blocked domains/requests, 2020

Type of domain	Unique domains blocked	Total requests blocked
Domain generation algorithms (DGA)	46,468	18,052,375
Botnet command-and-control (C2)	31,921	18,921,845
Ransomware	3,152	10,472,963
Exploit kits	197	681,096
Total	158,664	104,517,035

As usual we should be careful how we interpret category or 'type of domain' labels, as naming and categorisation is often inconsistent across source data and isn't always reliable.

We use the count of 'Unique Source IP Addresses' as a proxy for commonality, as it shows the potential exposure to a malware threat across our customer base.

It was notable, then, that Emotet domains were seen across more customer IP ranges than all of the uncategorised domains we have received that are labelled as generic. Emotet was originally a banking trojan spread by email, but has evolved over time and in 2020 it was commonly used as a 'dropper' to deliver ransomware, such as Ryuk and Conti. Its popularity with malicious actors speaks to the insidious nature of cyber crime. Also notable is a resurgence in Newpoptab: adware that causes adverts in browser pop-ups, but is also a key delivery vector for various malicious applications. You will also see the usual suspects of malware threats old and new in Table 16.

Table 16. Exposure of malware threats across PDNS customer base, 2020

Threat name	Unique customer IP addresses	Sightings
Newpoptab	716	806,154
Emotet	641	884,480
Generic	612	5,883,328
Inor-AA	541	180,147
RefC-Gen	480	221,307
JSRedir-OE	456	58,802
Bongacams	455	102,374
Cryptor	419	421,727
Artemis	386	43,819
Crusewin	384	317,201
Zeus	355	448,717
RIG Exploit Kit	355	672,937
MalwareDownload_Emotet	351	140,546
Nemim	340	25,472
Malscript	336	25,429
Heodo	327	349,147
LokiBot (Android)	322	77,713
WP-VCD	320	23,017

Protecting against COVID-19 themed sites

In March, we responded to an increase in malicious domains related to COVID-19 by blocking these domains. Our process for triaging and distinguishing harmful domains from legitimate ones performed admirably, allowing us to protect PDNS users from these novel variants of traditional online threats.

For example, in March the COVID-19 related blocked domain that we blocked more than any other was a webpage that duplicated data on the spread of COVID-19 from a [GitHub repository at Johns Hopkins University](#). This legitimate data helped conceal the fact that the site hosted malware.

The second most blocked domain in March was very new (it appeared in mid-March) and was registered by a Chinese registrant. The site was in English and appeared to be selling health supplies, gloves, masks and hand sanitiser, however it was in reality a phishing site with a fake web shop.

We also saw examples of domains registered in March which, to begin with, did nothing more than redirect to genuine COVID-19 sites. In many cases there wasn't obvious malicious intent, however, in some cases the IP addresses hosting the domains were known to be malicious and the sites were blocked.

As these examples show, criminals are always willing to take advantage of a crisis to steal, defraud or extort money and PDNS has played an important part in countering that threat.

Building capability to support incident management

Throughout 2020, the NCSC developed its capability to use PDNS data more effectively when responding to incidents. NCSC data analysts and engineers created tools to fingerprint vulnerable technologies and attacker behaviour from PDNS queries, and to map these findings to potentially affected organisations. They also established closer working relationships with incident and vulnerability management teams.

PDNS findings are enriched with data from other ACD services - such as [Host Based Capability](#), [Web Check](#) and [Mail Check](#) - as well as open source data. The result is a rapid, flexible, and reliable data analysis pipeline, which the NCSC relied on when responding to incidents throughout the year.

Research and analysis with PDNS data

PDNS is a quick and convenient way to identify public bodies using vulnerable technologies or affected by malware. PDNS provides a record of historic activity for incident response and research. It allows the NCSC to monitor the remediation of incidents and vulnerabilities over time. It also provides a general picture of UK government cyber maturity by revealing the technologies in use in the public sector.

The PDNS dataset is a rich source of public sector domain names and IP addresses. NCSC data scientists combine this information with data gathered from other ACD services, and also commercial and open-source data. The data shines a light on the extent of UK public sector IP space and trends, such as the use of Domain generation algorithms (DGA) for malicious purposes. The NCSC uses this information to protect the UK and counter malicious actors more effectively.

In March and April, NCSC data scientists used the dataset to map NHS domains and IP addresses, identifying almost 1,000 IP ranges that were of interest. The NCSC used these ranges to identify 1.1 million NHS IP addresses that required protection.

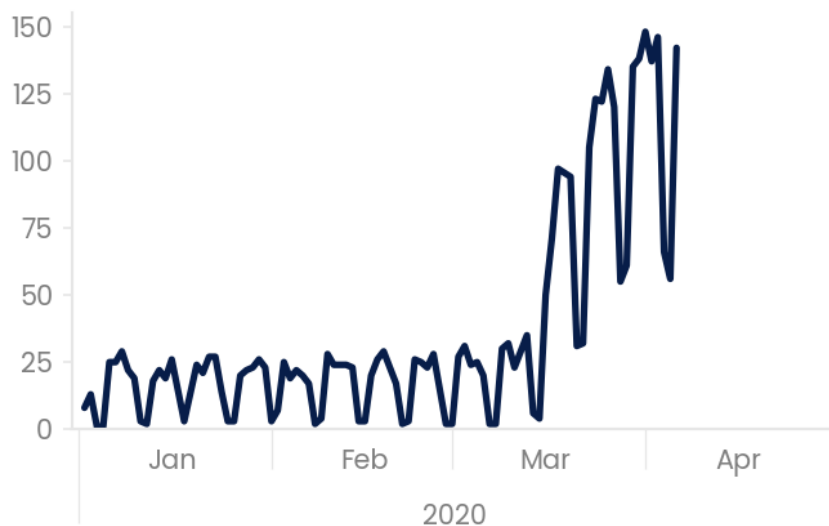
Figure 31. Anonymised example of network analysis techniques used to identify unknown NHS domains and IP addresses from known ones



At the same time, there was anecdotal evidence that Zoom had rapidly grown in usage in the public sector, and PDNS gave us the data needed to test this empirically.

Figure 32 shows the number of PDNS organisations using Zoom during the first three months of 2020, which increases from between 20 and 30 organisations in January to over 120 by the end of March. This data backed up the anecdotal evidence.

Figure 32. Spike in April of PDNS organisations using Zoom, Jan-Apr 2020



We used this analysis, and similar analyses of other platforms, to create advice on using the most popular platforms as securely as possible, and to provide our customers with advice tailored to their preferred platform.

SolarWinds

2020 drew attention to the value that could be found in PDNS data. By the end of the year our data had already helped with vulnerability detection, threat hunting, incident support, and with our COVID-19 response. If a more compelling argument were needed, we were unfortunately soon to get one, as throughout the year malicious actors made use of DNS for thoroughly underhanded reasons.

As mentioned previously, by the end of 2020 the number of organisations protected by PDNS had never been higher, and we covered the majority of key parts of the UK public sector. This coverage gave us a unique and fairly complete dataset to investigate should the need arise; and in December, it did.

In that month, FireEye disclosed that they had detected a sophisticated campaign which used a compromise of the software supply chain of the SolarWinds Orion product to distribute malware and gain access to victim networks. A key technology used in this activity was DNS, so the PDNS dataset was a primary data source for analysis, and critical to informing our situational awareness and our response.

PDNS's broad view of DNS activity across the UK public sector enabled NCSC analysts to rapidly measure how many public bodies were affected. As details of the incident became clearer, historic PDNS data revealed the extent of compromise in each affected organisation. This information helped the NCSC prioritise its support to organisations with the more concerning indicators of malicious activity and, just as importantly, to provide assurance to many core parts of government that they were not affected.

Customer events

We have always found that working groups and other similar events are extremely useful for learning from existing and prospective customers. As the service became more mature we planned to continue with them in 2020, and started strong with an in-person event during CyberScotland Week in February 2020.

For obvious reasons that was also our last in-person PDNS event of 2020, and in March we moved our programme of events online. Run by our delivery partners at Nominet, we hosted a series of webinars aimed at sharing knowledge with our users and gathering feedback to inform our continuous improvement programme.

Events, whether in person or virtual, are an important part of engaging with our users, and we know from feedback that customers use what they've learned to improve their cyber security. We hope to resume in-person events in 2021, but as there are benefits to virtual events we plan to use both formats in future. We would like to thank everyone who attended our events in 2020, and to everyone in the wider community for being active and engaged with PDNS.

Dangling DNS

About the service

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. As an attacker, having the ability to publish your own web site content using a legitimate domain name can make it easy to mislead people into believing the web site is trustworthy. As such, these vulnerabilities are often exploited as part of phishing attacks to make them seem more believable.

To put the scale of this problem into perspective, over 25% of all the reports received through the Vulnerability Reporting Service in 2020 were due to subdomains with dangling DNS records, making it the second most frequently reported type of vulnerability.

Progress

Renewing the approach

In 2020, the NCSC's Security Engineering team completely re-engineered the previous year's service solution to focus on a small set of services used in the majority of dangling DNS instances we had identified. These spanned three of the large cloud service providers (AWS, Azure, and Google Cloud), but also included more specialised services, such as GitHub Pages, WordPress, and Bitbucket.

It is important to understand that this is not a reflection on the security posture of these service providers, but merely shows that it is common for an organisation to configure its own custom subdomain to point to a resource offered by them. For example, the owner of my-registered-domain.co.uk may accidentally configure files.my-registered-domain.co.uk to point to a resource within AWS that had not yet been (or was no longer) registered. Whilst it's true that the service providers play a role in dangling DNS vulnerabilities by allowing anyone (including malicious actors) to register resources with arbitrary names, responsibility lies ultimately with the owner of the record.

The re-engineering project allowed us to draw on the successes of our previous solution while avoiding repeating the same mistakes. We wanted to maintain the ability to identify these vulnerabilities on a national or even global scale, whilst maintaining accuracy and cost-effectiveness. This meant turning to the Cloud once more, where resources could be instantly dialled up or down as needed to cope with the number of scans we might have to perform in any given moment. This time, however, we also wanted to ensure that the vulnerabilities we discovered could be mitigated automatically if required, and that this functionality would be incorporated into the solution early on.

Integration with external DNS sources

Finding dangling DNS records relies on having a plentiful source of records to check. DNS is designed to be distributed in a hierarchical structure of different "authoritative zones" and "delegations" with no single complete and authoritative source of records, so obtaining a set of data such as 'all the subdomains of any domain ending in .uk' is not at all trivial. We searched for resources where the hard work of discovery had already been completed for us, and in doing so settled on two: publicly curated and published DNS record sets, and our own protective DNS service, [PDNS](#).

Our first focus was on the publicly available sources, often in the form of lists of domains and subdomains that had been discovered and then made available - usually freely - to search or download. For example, for the domain ncsc.gov.uk this would include subdomains such as beta.ncsc.gov.uk and webcheck.service.ncsc.gov.uk. Although these lists are by no means exhaustive, by automatically downloading and scanning them on a weekly basis, we were able to maintain a regular global view of vulnerable instances of the services we were interested in.

The second focus was on leveraging PDNS. Scanning for dangling DNS vulnerabilities within the set of domains and subdomains queried by systems using this service gave us a unique 'inside looking out' perspective. This led to some important discoveries, including several dangling DNS records belonging to a large technology company that they resolved after we alerted them to the problem.

Outcomes

Whilst the automatic defensive registration of all vulnerable Dangling DNS resources still remains a distant possibility, the circumstances that unfolded in 2020 allowed us to 'test drive' the capability as part of the NCSC's COVID-19 response efforts in April and May. During this period, we were able to automatically identify and defensively register 95 subdomains belonging to organisations within the UK healthcare sector that could have otherwise been taken over and used maliciously.

Throughout the year, the Dangling DNS solution scanned over 10 million instances of services globally for vulnerabilities, and found over 90,000 vulnerable records pointing at resources amongst several popular service providers, as shown in Table 17.

Table 17. Results of Dangling DNS scanning of instances of popular services, 2020

Service	No. scanned	No. vulnerable	% vulnerable
Azure Cloudapp	141,153	53,306	37.8
AWS S3	355,166	25,448	7.2
GitHub Pages	561,175	4,789	0.9
Azure Traffic Manager	351,721	3,963	1.1
Surge	6,250	2,224	35.6
WordPress	9,434,833	888	0.009
Bitbucket	818	182	22.2

Conclusion

The complexities of administering DNS continue to be responsible for a large number of vulnerable dangling DNS records that can often be 'hijacked' and used for malicious purposes. Throughout 2020, we focused on our ability to automatically detect and mitigate the risks of such records pointing to several of the most commonly used online services. We also re-engineered our solution to support scanning for these at scale.

We were able to use this to help defend the UK public health sector during the first few months of the COVID-19 global pandemic, and throughout the year identified a total of over 90,000 vulnerable instances pointing to these common service providers.

Our vision for the future of the Dangling DNS solution remains ambitious yet clear:

- **Increase awareness** of dangling DNS vulnerabilities by extending and exposing the capabilities of a large-scale detection and alerting service.
- **Reduce harm** by minimising the window of opportunity for exploitation with defensive mitigation wherever possible and appropriate.
- **Drive change and ultimately eliminate this class of vulnerability** through developing and supporting the adoption of innovative and permanent solutions by service providers.

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 telecoms-related Signalling System No.7 (SS7). We have also set up the SMS SenderID Protective Registry, to help organisations protect their brand from use in SMS phishing attacks.

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 reroute large swathes of internet traffic accidentally or maliciously.

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. This only looks at how the internet moves the data packets around, not the data itself.

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.

In 2020, the NCSC worked closely with UK CSPs to produce and publish a [best practice guide](#) for BGP. The guide contains 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, which became known as BGP Spotlight. In 2020 this project progressed from the beta stage to a live environment, with a separate dedicated development environment.

During development, 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. In the live version, updates are collected via BGP peering; this is where two routers create a BGP connection in order to exchange information. BGP Spotlight has peerings with 4 other operators in addition to BT. These supplement the original BT and publicly available RouteViews and RIS-Live feeds.

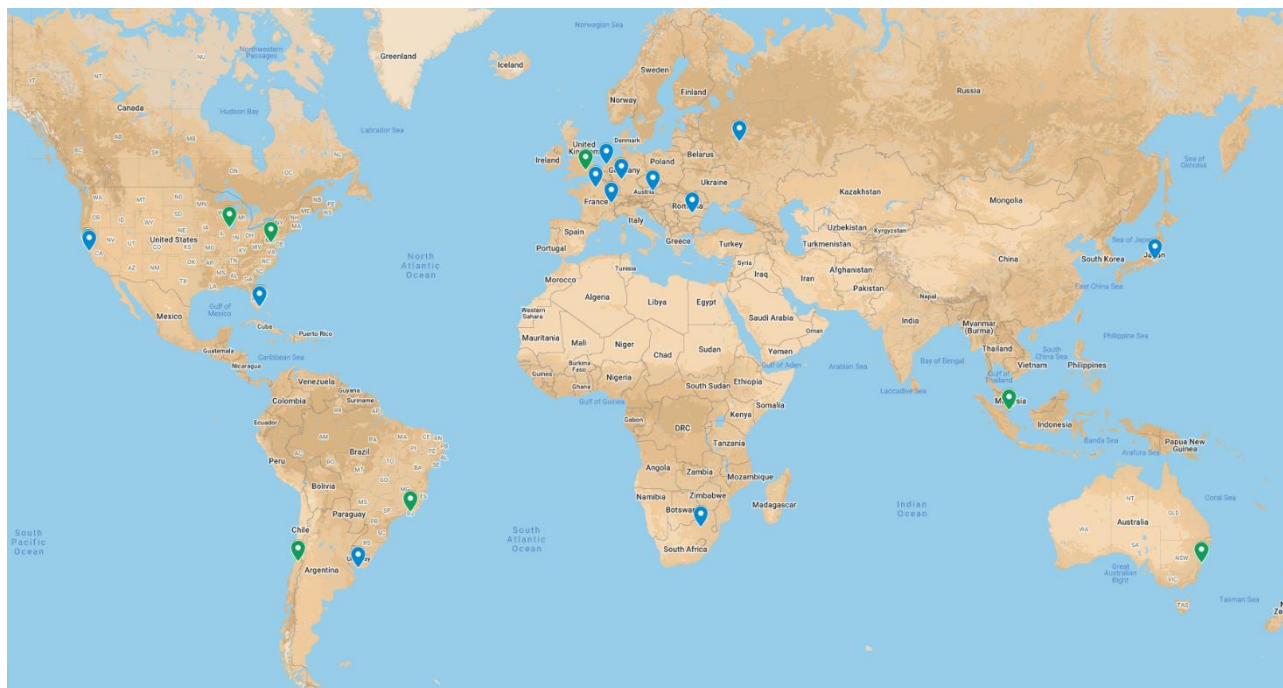
The collector locations used by the tool are shown in Figure 33. Diversity of data is key to getting multiple views of the traffic routing, and these additional peerings allow for a much wider view of BGP routing updates. As the data is ingested, analysis is performed to detect any unexpected or irregular 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, 11 operators are using the live system, as well as 4 universities, 2 technology companies, and the Canadian Centre for Cyber Security. There are 142 users in total, and in a typical 24-hour period BGP Spotlight ingests approximately 1 billion messages. These produce approximately 6m events of interest, of which approximately 4.5m are path changes, with most of the remainder being ownership changes. 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' events are the most significant, as they are most likely to lead to traffic outage or a hijack/rerouting scenario. 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 path changes could also be an indicator of malicious activity.

Figure 33. Locations of RouteView and RIS-Live collectors used by BGP Spotlight



Key: Green pins - RouteView collectors (9), Blue pins - RIS-Live collectors (12)

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 had similar vulnerabilities to SS7. So far, our tests have also revealed serious examples of these vulnerabilities in the UK's mobile networks.

In order to provide more flexible testing, we moved to a more automated approach. We've made a testing service available to the networks so they can run tests periodically, or when there are significant security improvements to their signalling 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.

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

In 2020 NCSC, together with DCMS and Ofcom have continued to develop the Telecom Security Requirements (TSRs) in discussion with UK telecoms operators and equipment vendors. These 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.

SMS SenderID Protective Registry

The NCSC, along with [UK Finance](#) and others, has part-funded an initiative to set up an SMS SenderID Protective Registry. SenderIDs are the text 'from addresses' you see on SMS messages. The Registry 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 SMS phishing campaigns.

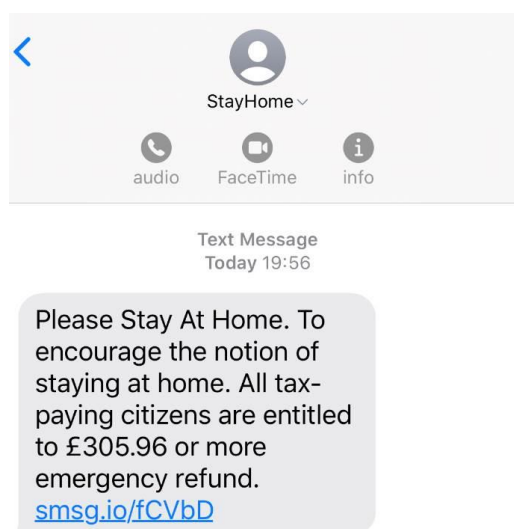
The Registry was created and is independently administered by the Mobile Ecosystem Forum (MEF). Participating SMS aggregators use the Registry to ascertain whether they should block or deliver SMS traffic that is routed via their networks. At a superficial level, you can think of the Registry as a codex, 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 if it follows the delivery path expected. Authorised SenderIDs following a different path or bogus derivatives (such as DV1A) should not get delivered to users.

Progress

Following on from the success in 2019 with DVLA and HMRC, TV Licensing became the next most abused government brand. TV Licensing joined the Registry in January and started populating it with data in March.

In the same month, the government embarked on a mass SMS campaign to the general public, asking people to stay at home. Working with MEF, we quickly added the SenderID 'UK_Gov' to the Registry along with derivatives that could be used to dupe citizens (known as typosquatting). At the same time, we added other reported SenderIDs associated with scams, such as 'COVID' and 'StayHome' to the block list, as shown in Figure 34.

Figure 34. Example of SMS phish using the 'StayHome' SenderID



Working with Government Digital Service (GDS), we set up a new account called 'Cabinet Office' and later another one called 'NHS' to enable us to manage all government initiatives as part of the COVID-19 response. This means that any government messaging implemented through GDS can now benefit from automatic SMS protection. With GDS we have been able to proactively protect SMS communications for the following additional services:

- Member Hub - voting service for MPs
- devolved administrations
- 119 and 111

HMRC


In April 2020, HMRC saw SMS referrals drop again (to 2,515), and volumes were back to the levels before they began working with the Registry. A review of the SMS phishes at this point revealed only a handful of scams, which appeared to be continuing to use the less credible SenderIDs. We also found that the majority of these scams had shifted to using MSISDN long numbers, also known as virtual mobile numbers (VMN).

In May, there was another spike of HMRC-themed SMS referrals (to 12,269). 5,023 of them contained the word 'rebate', and 5,082 contained the word 'refund'. All used a variety of different URLs, and many of them displayed the sender ID as a long number, which appeared to be unique.

HMRC worked with law enforcement colleagues to investigate further, and at the end of May a phishing scammer was arrested. HMRC's analysts identified 25 HMRC-branded templates used by this specific scammer, which may be why SMS referrals then dropped to 5,416 in the following month (June).

In August, referrals began to rise once more to a level of 11,192 in December, although no specific long number stood out as being used for a sustained attack. As always, any numbers identified were passed to the relevant provider for further investigation.

Figure 35. HMRC-themed SMS referrals, 2019-2020



Outcomes

We have seen a continued shift in SenderIDs to the use of long numbers, which hopefully should be easier for targets to identify as suspicious. Also, there has been global interest in the work we've been doing in this space, and we have supported MEF in discussions with governments and regulators wishing to understand more. We continue to work with MEF on discovering ways to improve the Registry, with some exciting opportunities being explored for next year.

Conclusion

One of the common complaints from merchants has been around the difficulty in identifying the legitimate routes for SMS. Many aggregators hide behind commercial sensitivity for their downstream routing and frequently change their providers to benefit from cheaper pricing.

We need to work with the messaging industry to improve transparency. Equally, merchants need to responsibly source their messaging providers and not rely on price alone. If a messaging provider can offer prices lower than an operator's cost price and is unwilling to share their downstream routing, this should warrant further investigation.

Host Based Capability

About the Service

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 departments understand their threat surface, and to forewarn those affected by new, major vulnerabilities:

- **Detect:** HBC has identified or assisted on a cumulative total of 18 incidents. By participating in HBC, departments can receive granular information at the device level. This can help improve incident response and remediation. HBC identified (and assisted) on a cumulative total of 23 Suspicious Activity Observations (SAOs) by the end of 2020. The HBC notifications identify when 'irregular' detections are made but are unconfirmed as malicious activity, and HBC supports departments where possible on further investigation.
- **Threat surface:** By the end of 2020, HBC had generated a cumulative total of over 250 threat surface reports (TSRs). These 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 has provided information to departments about their unique exposure to 6 new, major vulnerabilities since it launched, for example SIGRed and ZeroLogon. The information included comprehensive data to help departments react quickly to these emerging threats.

Progress

HBC is now towards the end of its third year of operation and has recently celebrated an exciting benchmark. At the end of 2020 it was installed on over 280,000 laptops, desktops, and servers across 24 UK government organisations; more than double the figure of around 130,000 at the end of 2019. This has been achieved in collaboration with the Canadian Centre for Cyber Security and other international partners.

HBC has been deployed to a number of organisations supporting the national response against COVID-19 in a number of supporting sectors, including public health agencies, national vaccine procurement efforts, and associated regulatory bodies. The service was used to collect technical metadata that could be useful for these organisations and move it offline for analysis. Note that user data was not collected or analysed.

HBC also played a strong role in the national effort to understand the exposure of the UK and its government to the extent of the SolarWinds attack, drawing on our Detect and Forewarn capabilities. We provided support by expanding server coverage of multiple departments, looking for potential compromises of UK government networks via the SolarWinds vulnerability.

Many more UK government departments have expressed an interest in HBC and are in the process of adopting the service. The HBC team held an inaugural All Partners Meeting in September 2020, to update our user community on the capability and to give them insights into the work we have done. This also gave them the opportunity to ask questions and to provide user feedback on the service and our roadmap of features.

Vulnerability Disclosure

About the service

The Vulnerability Disclosure project is focused on maturing the UK's approach to vulnerability disclosure and remediation. There are three 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 best practice disclosure processes by creating a Vulnerability Disclosure Programme (VDP) for any department that signs up.
- **Vulnerability Disclosure Toolkit:** a free online resource that organisations can download and use to implement the essential steps to establish a vulnerability disclosure process.

Progress

Vulnerability Reporting Service (VRS)

The VRS has grown rapidly during 2020 and has proven itself a valuable reporting service that continues to help affected organisations receive and remediate reported security vulnerabilities.

We received over 450 new reports during the year, and after triaging the reports we directly contacted over 150 separate organisations. Over 25% of the new reports were due to subdomains with dangling DNS records, making it the second most frequently reported vulnerability type. The most reported was reflected cross-site scripting (XSS), where user-provided data is used immediately by server-side scripts to parse and display a page of results without properly sanitising the content. In their simplest form XSS vulnerabilities can be used to display a nuisance pop-up message, but some attacks exploiting them can have much more serious consequences.

The most important aspect of the VRS is working with the system owners to help resolve these reported vulnerabilities. Once a report is triaged, we provide a summary to the system owners concerned. This includes a full description of the vulnerability, steps to reproduce, and recommendations on how to mitigate the issue.

To help highlight the importance of vulnerability disclosure we created a vulnerability disclosure scenario within the NCSC's [Exercise in a Box](#) toolkit. This scenario helps organisations to understand best practice when handling a vulnerability disclosure.

Vulnerability Disclosure Pilot

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.

During 2020, the pilot has helped launch 8 UK government VDPs, allowing departments to directly receive reports on, remediate, and fix issues before they caused harm, rather than relying on the VRS.

The developers of the NHS COVID-19 App were particularly keen to ensure security researchers could easily report any security vulnerabilities. Through our work with the Vulnerability Reporting Service and the Vulnerability Disclosure Pilot, we were able to quickly setup a disclosure process. This consisted of a [HackerOne vulnerability disclosure program](#) with a comprehensive disclosure policy, and a clear process for reporting vulnerabilities. Once reported, each vulnerability issue was triaged by [NCC Group](#) before being assigned to the development team.

Vulnerability Disclosure Toolkit (VDT)

The feedback we have received on VRS and the pilot has driven the publication of the [NCSC Vulnerability Disclosure Toolkit](#). This is aimed at any public or private sector organisations and provides them with three essential elements we want everyone to implement to make vulnerability handling easier, both for the organisation and the person trying to report a vulnerability.

The Toolkit provides:

- a simple way to communicate a vulnerability (such as a secure web form)
- advice on how to publish a vulnerability disclosure policy (we provide an example policy)
- the way to advertise the disclosure route (via publishing a security.txt file)

To coincide with the launch of the Toolkit, we have also added a scan for security.txt within [Web Check](#) to highlight the importance of a clear and secure vulnerability disclosure route.

NCSC Observatory

About the service

The NCSC Observatory focuses on generating data-driven insights to underpin the NCSC's research and strategy and allow an effective response to incidents. It analyses publicly accessible data, such as DNS records of UK domains, and uses this to track the uptake of DNS security protocols and the usage of different technologies and cloud providers. Additionally, it consumes and analyses data from [Protective DNS \(PDNS\)](#) to help track the usage of different technologies within the UK public sector. 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.

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.

Progress

Over the course of 2020 work has been done to re-write and re-design parts of the Observatory. This has allowed the system to scale in a reliable and safe manner. The Observatory can now process large quantities of data, drawing out insights and allowing further analysis.

Outcomes

The Observatory has continued to provide insights into the uptake of different security protocols in the UK, including the usage of the Domain-based Message Authentication, Reporting and Conformance (DMARC) protocol by domains registered into the UK. DMARC is a protocol used to prevent email spoofing, and the NCSC wishes to promote the uptake of this protocol to reduce the impact of phishing attacks. The Observatory allows us to measure our success towards this target.

We are also working on the following initiatives:

Measuring uptake of tools in the UK

The national lockdown in March 2020 resulted in increased use of a variety of remote working tools as more people worked from home, which brought with it an increased risk of cyber attacks. In response, the Observatory provided insights into the popularity of different tools and their uptake within the UK public sector. For example, there was a large uptake of the Zoom video conferencing service, as shown in Figure 32, while Google Meet and Webex had similar but smaller increases in usage.

This information helped the NCSC prioritise research into the security of these tools, and provided an insight into the level of exposure to the UK should a vulnerability be exposed in one of these tools. More broadly the Observatory has worked to expand the number of products we can track usage of within the UK public sector.

National-scale vulnerability management

Throughout the course of 2020, the Observatory has supported the NCSC's vulnerability and incident management functions, assisting with the investigation of several high-profile incidents. For example, the Observatory helped identify users of SolarWinds products containing the [CVE-2020-14005](#) and [CVE-2020-13169](#) vulnerabilities, in the public and other sectors. This information helped the NCSC and those responsible for the affected systems act to resolve the issues.

Suspicious Email Reporting Service

About the service

The Suspicious Email Reporting Service (SERS) enables the public to report suspicious emails by forwarding them to report@phishing.gov.uk. The service analyses the emails, and when links to malicious sites are found we seek to remove those sites from the internet to prevent them doing further harm.

Progress

Following on from the original proof of concept in 2018 and further development in 2019, we partnered with law enforcement agencies, such as the City of London Police, to develop SERS into an automated system for handling reports of suspicious emails. A public alpha version of SERS was launched on 21st April 2020. It attracted national media coverage and was widely shared on social media. As a result, the service instantly proved very popular with the public, receiving over 500,000 reports in the first 30 days.

SERS automatically triages reports and forwards them to the [Takedown.Service](#) operated by Netcraft on behalf of the NCSC. Netcraft determines whether a report contains an email that is suspicious or benign, and reports back to SERS. When a person submits a report, they are sent an email that thanks them for their report and includes current statistics of reports and actions taken.

Not only does SERS successfully take down malicious sites, it also provides partners and law enforcement with vitally important data that helps them to:

- understand the strategic threat to the UK
- alert the public to phishing trends
- identify investigative leads


The successful uptake in reporting by the public provided useful information on the system and its further development needs. As a result, SERS successfully obtained public beta status in October 2020, and we have a clear roadmap to officially take it to live status in 2021.

Outcomes

From when the service launched on 21st April to 31st December 2020, we received just under 4 million reports from the public, an average of 15,800 reports a day. These reports identified just under 1.5 million malicious URLs, some of which were also identified by the [Takedown.Service](#). SERS was credited in instigating the removal of more than 26,000 scams, involving 48,000 malicious URLs, not previously identified by the [Takedown.Service](#).

Figure 36. Reports to SERS remained fairly consistent (~500,000 per month), 21 Apr - 31 Dec 2020



Figure 37. Resultant takedowns of SERS reported scams steadily increased per month, 21 Apr - 31 Dec 2020

Some recent user data analysis has shown that we have around 25,000 submitters a week. While most are from private email addresses, a proportion of reports come from company email addresses, which demonstrates that the suspicious emails can still get past company spam filters.

One surprising statistic we have noted is that 14% of submissions are from legacy commercial email services that no longer offer new accounts. We will be investigating this with providers to determine whether these users could be more at risk of receiving phishing mails, and to look for further ways to either protect the users or encourage migration to active domains.

Conclusion

Despite only being in service for a little over 8 months, SERS has already made a significant impact in improving the ability of the public to report suspicious emails, enabling the NCSC and our partners to take action against them, thus reducing the harm they cause. Ongoing support by the public in submitting all suspicious emails will continue to protect society as a whole, as more and more phishing sites are taken down before fraud can take place.

Exercise in a Box

About the service

Exercise in a Box is a publicly available tool that allows organisations to practise and refine their response to the most common and pressing cyber security incidents in safe and private environment.

Facilitators are given the tools they need to lead relevant staff within their organisation through a scenario that unfolds through a series of prompts. This is designed to stimulate discussion about an organisation's policies, processes and procedures, with attendees self-assessing their organisation's maturity and readiness against a sliding scale. At the end of the exercise, a downloadable 'End Report' is created, and relevant NCSC advice and guidance is linked.

Primarily aimed at the non-technical audience within both the public sector and small to medium enterprises, the service has also seen strong take-up amongst large organisations and cyber security professionals. The service has been designed to have a low barrier to entry, to be easy to navigate, and to be consistent with the NCSC look and feel.

Progress

During 2020, the Exercise in a Box team released additional table-top exercises in response to the growing number of people working from home: the 'Home and Remote Working' exercise, and separately a 'Vulnerability Disclosure' exercise. In addition, we developed and released a series of lighter-touch 'micro exercises': bite-sized modules designed to get people into exercising, or to form part of a regular cyber security meeting. Each new piece of content was developed in response to emerging threats and/or user demand, and the tool will continue to evolve in response to targeted user research.



As well as building on current content, we moved the service from public alpha to public beta status. This is not simply a name change; it was an opportunity to assess important aspects of the tool's evolution such as:

- **accessibility** - a detailed accessibility audit was conducted, and full remediation completed, to ensure that the tool can be used by as many people as possible.
- **support** - a support plan that considers the sustainability of the product in the medium to longer term, including integration with other relevant ACD products and services.

Over the course of the year we also worked with our supplier on a planned upgrade to the service's user interface. Grounded in user testing and with a design element that is formally aligned to other NCSC services - including [MyNCSC](#) - this work will manifest itself in an upgraded landing page with new features and notifications.

Outcomes

Take up in 2020 was good, as can be seen in the following charts. While the greatest concentration of users is in the UK, we were surprised by the international interest and take up of the tool. We have worked with many nations that have showed an interest in Exercise in a Box, some of which intend to create a similar tool. In particular, we signed a Memorandum of Understanding with Singapore, and have worked closely with them to provide Exercise in a Box in a form that will run on Singaporean infrastructure, which they expect to have running early in 2021. We have also created an offline model for interested nations wishing to evaluate the tool.

Figure 38. Cumulative increase in Exercise in a Box users across all organisation types, Jan-Dec 2020**Figure 39. Increase in completed exercises across all types, highlighting the ransomware exercise, remote working exercise, ransom attack micro exercise and phishing micro exercise, which saw the greatest growth, 2020**

Note: Exercises not labelled (in grey) are in order: Third Party, Threatened Leak, Simulation, BYOD, WiFi, Supply Chain, Vulnerability Disclosure. Some of these were not introduced until later in the year.

Conclusion

In the first full reporting year of Exercise in a Box, we have demonstrated the need for an easy-to-use exercising tool linked to world class NCSC guidance. We have reacted to world events by creating exercises to allow for remote and home working, and released micro exercises to allow for more 'citizen' focused engagement. We will continue to develop the tool in line with the feedback we receive from users and as further guidance topics are developed by NCSC.

Logging Made Easy

About the service

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.

Progress

Taking user feedback into account, we have continued to develop LME by improving the user experience, giving greater insights into data with metrics that support real world investigations.

The most recent release of LME included both Health Check and IT Security dashboards, which help customers understand important issues such as patch levels across their IT estate, or commands that were run with elevated privileges (by whom, and on which machines). These invaluable insights are key to improving management of your IT estate, responding to events, and informing your investigations.

Outcomes

From its initial release in Spring 2019 up to the end of 2020, 1,319 daily unique clones of the LME GitHub repository have been made, 899 of which were in 2020 alone. We deliberately do not capture analytics on the identities of the users, just on the numbers of clones, as is standard for GitHub. However, we do know that it has been shared and referenced internationally, and that in the UK it has been particularly embraced by local government.

As well as enjoying the advantages of insight into logging activity, the adoption of LME by users has enabled some UK organisations to benefit from the NCSC's Cyber Threat Intelligence (CTI) Adaptor. This provides a threat feed of advisories and alerts from the NCSC to a participating organisation, and highlights the sighting of a perceived threat in the logs. This functionality for LME users has been tested on a pilot basis since December 2020, and will be expanded to cover other Security Information and Event Management (SIEM) systems in 2021.

Conclusion

The continued adoption of Logging Made Easy has provided valuable logging capability for hundreds of users in a format that is free of charge and relatively simple to install and operate. In particular, it has provided an opportunity for some UK organisations to benefit from the protection of the [CTI Adaptor](#) pilot, which should be progressing to a live service in 2021.

Cyber Threat Intelligence Adaptor

About the service

The Cyber Threat Intelligence (CTI) Adaptor is a software program, designed by the NCSC's Threat Detection and Response team, that enables authorised organisations to receive a high-quality, contextually rich, cyber threat intelligence feed from the NCSC. The Adaptor integrates with a variety of SIEMs, using customer log data to detect known Indicators of Compromise (IOCs) contained within the feed.

Progress

Development of the CTI Adaptor progressed during the year. Initially, the focus was to produce a version of the Adaptor that was compatible with the NCSC's own Logging Made Easy (LME) service. A threat feed comprising IP addresses and suspicious domains was constructed, based on alerts and advisories published by the NCSC in recent months.

The IOCs contained within the feed are searched for in a customer's log data. A match between the customer's log data and an IOC, called a Sighting, is flagged to the customer for investigation, along with the Observable (that is, the actual log line that contained the IOC). The Sighting is also sent back to the NCSC.

Outcomes

The pilot version of the CTI Adaptor was installed at a small number of UK organisations in December 2020. The Adaptor quickly returned Sightings of potential threats, to both the organisations and the NCSC. The functionality of the Adaptor was demonstrated, as customers were able to get information immediately on the nature of the suspected threats, and the NCSC gained an increasingly broader view of the threat landscape in the UK.

Conclusion

The CTI Adaptor helps better protect the UK by pushing out the NCSC's threat knowledge to a wider audience. In return, the NCSC gains an improved awareness of threats arising within the UK. We have continued to roll out the pilot of the Adaptor to a wider user base, and the value of its ability to highlight sightings of potential IOCs has been demonstrated.

Further refinements have broadened the Adaptor's functionality to cover users with other widely used SIEMs. The success of the pilot so far has been promising, and a full evaluation will be undertaken before the service goes fully live and is offered to a wider user community.

MyNCSC

About the service

The NCSC's digital services, including but not limited to those delivered by [ACD](#), are designed to improve an organisation's approach to cyber security. Although each service can be used in isolation, the best outcomes are achieved when they are applied as part of a holistic approach.

The MyNCSC platform brings the NCSC services together into a single, coherent experience tailored to show the content, vulnerabilities, services and alerts most pertinent to each user and the organisation they're defending.

MyNCSC helps users reduce duplication, save time, and understand their security posture across a range of services. Users will only need to sign in once to retrieve all their service data, incident information, and the guidance required to help them be more proactive in improving the security of their organisation. MyNCSC will eventually replace the [ACD Hub](#) as the single point of access to ACD services.

Progress

2020 was an exciting year for the MyNCSC project. Early discovery work from 2019 soon evolved into low-fidelity wireframe designs of the platform, suitable for testing with future users. Concurrently, development of the enabling services got underway: those essential components that serve as the foundation of the MyNCSC vision, such as asset and organisation management, and reporting.

In the first half of the year we continued to build the platform in response to user feedback and requirements derived from a wealth of research from across the ACD portfolio. The MyNCSC platform itself remained largely under wraps, although some of the enabling services were made available to users through other channels.

By the summer of 2020 we had brought together enough of the component parts to start sharing the platform experience with some pilot users. These users helped us to test early versions of MyNCSC; challenging or validating assumptions we've made along the way and ultimately shaping future releases.

We added capability through a series of small releases; first user onboarding and personalisation, then asset management features, a Web Check integration, organisational collaboration, and much more.

We have conducted user research on the first three pilot iterations of MyNCSC. Since September, hundreds of ACD service users across all parts of the public sector have taken part in a variety of user research sessions.

User feedback was overwhelmingly positive, with all surveyed users commending the platform's ease of use. Both usability and accessibility are really important to us; the MyNCSC project has to overcome the technical challenges of integrating services and supporting user-preferred ways of working, while also providing a platform that users want to engage with.

We would like to thank all our ACD users for taking part. If you would like to join our ACD User Research panel, please email acduserresearch@digital.ncsc.gov.uk.

Outcomes

The overarching goal of MyNCSC is to improve cyber security outcomes for users and their organisations. As part of achieving that outcome, users will see the following changes when migrating from the ACD Hub:

- better user experience, with easier navigation of the NCSC services.
- enhanced ability to find and use the relevant tools and services available.
- consolidated automatic notification of the most important issues across all services.
- tools and services promoted to users based on individual and organisational needs.
- data reused across the NCSC services, so that users need only enter information once.
- users will be able to create and manage teams to assist data sharing, exploitation, and workflow.
- users will be able to choose the tools and services that best suit their organisation.

Conclusion

2020 was the fourth year of the NCSC's Active Cyber Defence programme, the aim of which is to make the UK objectively and measurably safer from cyber attack. Our efforts are focused on commodity attacks that affect the majority of the people in the UK which can be prevented at scale. We have many approaches, which include but are not limited to:

- direct interaction with members of the public ([SERS](#))
- scanning and notification to system owners ([Web Check](#), [Mail Check](#))
- detection and reporting of attacks to infrastructure providers ([Takedown](#))
- supporting network providers in their own attack detection and response processes ([BGP Spotlight](#))

As we focus on scale and commodity attacks, we do not expect our efforts to prevent every attack; rather, we seek to make life harder for attackers, and to raise their costs to a level that is difficult to sustain. Additionally, the data we generate (and the experience the teams gain through running these services) gives government a better understanding of the cyber threats currently facing the UK, including the best approaches to combat them.

In the years leading up to 2020, our tools, services, and teams had matured to the extent that we were able to rapidly respond when the pandemic took hold, and to apply our technology and techniques in novel ways. For example, as cyber criminals attempted fraud through [fake online stores purporting to sell PPE](#), the Takedown service was extended to address this challenge. As organisations moved their meetings online, [analysis from the Observatory](#) fed NCSC research and guidance on remote collaboration tools. And as cyber criminals and state actors targeted the health and vaccine sectors, the [PDNS service was extended](#) to protect them.

This fourth annual report documents many of the valuable contributions the ACD programme made to the cyber security of the UK in 2020. By providing data and case studies on our efforts, we are demystifying the challenges of large-scale cyber attacks, and shining a light on the ACD services and other solutions that work. As we continue our efforts to broaden adoption of the approaches and services we've developed, we hope this report provides evidence and inspiration for others to adopt, adapt, and copy across industry and foreign governments.