National Cyber Security Centre

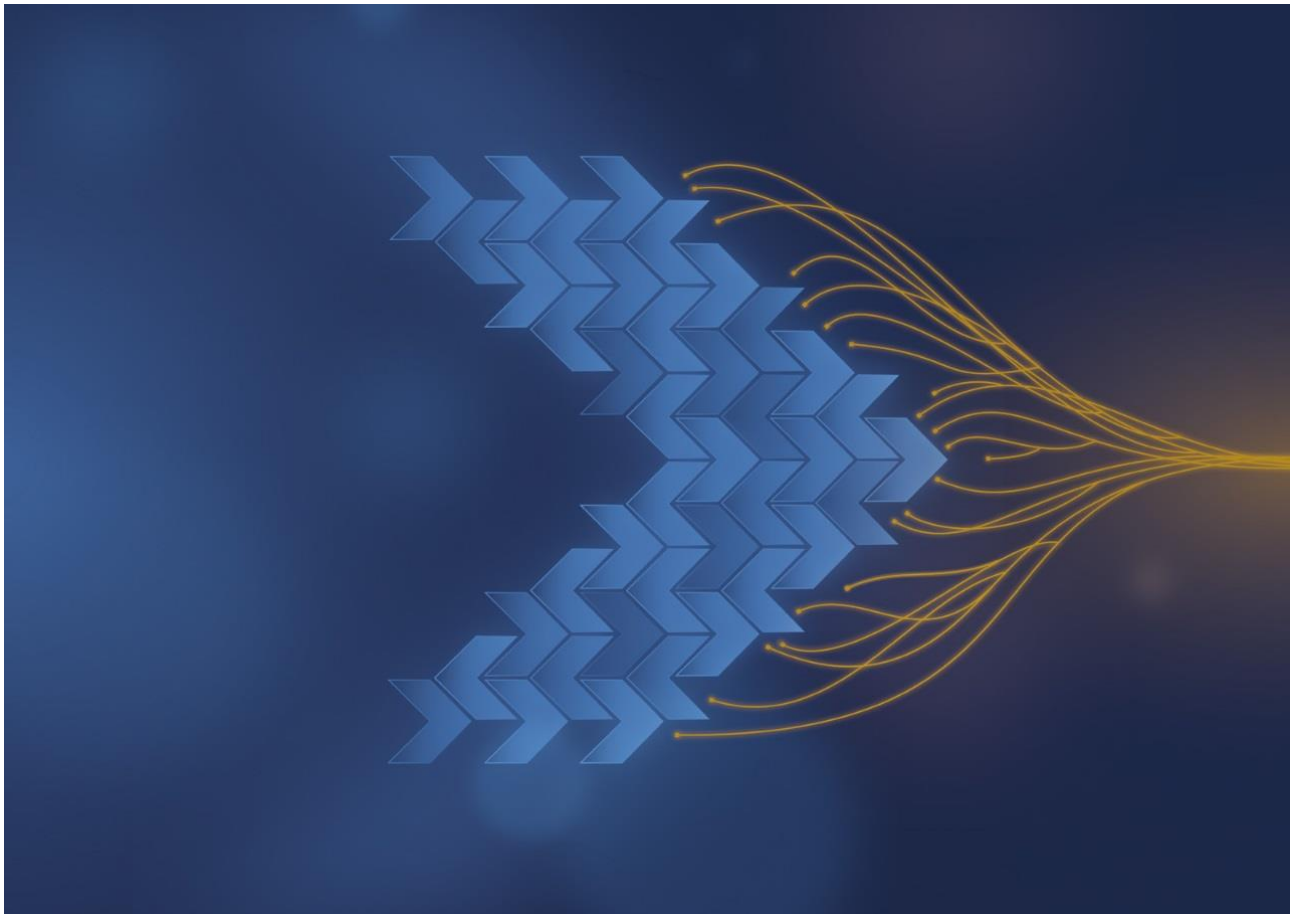

Active Cyber Defence

The Fifth Year

Contents

Introduction

The NCSC's Active Cyber Defence (ACD) efforts in 2020 were centred around helping with the COVID-19 response. In 2021, this driver has continued, as the UK has endeavoured to establish a 'new norm'. This report captures what ACD has been focused on during 2021, and the range of services that we have provided. It has been a year of consolidation and learning how to respond quickly and effectively to matters arising from the pandemic (whilst encouraging innovation and - where possible – finding new ways of doing business).

The aim of the 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, aiming to be resilient to commodity attacks, through the ability to provide protection at scale. In most cases that requires significant automation. We have an evidence-based approach and a considerable amount of data is generated as part of the process.

One of the goals of this paper is to use this data to provide transparency, and to demonstrate what works (and, indeed, what doesn't). In many cases we've developed these approaches with the public sector in mind, but in recent years we've looked at how we can broaden the reach of these efforts directly, and we will continue to do so, noting that there will inevitably be limits. Sharing our implementations with everyone will not be possible, but we're keen that the lessons we've learned are broadly understood and applied.

Several ACD services are already widely available, such as the Suspicious Email Reporting Service (SERS) which is open to the public to use. Others, such as Protective DNS, can have an impact beyond the public sector by sharing their threat intelligence (rather than having to broaden their direct service to further end users). We are exploring how other services can have a more direct impact on a much greater number of organisations.

The ACD Broadening project, established in 2020, aims to expand the impact of ACD beyond the public sector. The aspirations were:

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

Throughout this report, you will see reference to expansion of services into new sectors due to ACD Broadening. The project has supported further concerted work with the NCSC's resilience and customer teams to understand demand, and to pilot services (including with charities, schools, and town and parish councils). We commissioned guidance to support issues which public sector CNI organisations told us they needed (specifically '*What does good logging look like?'* and '*How do we secure internet-facing services?')*.

Working with our Mail Check team, we ran a secure email campaign focusing on how organisations can ensure their emails are protected in transit and reduce the risk of their email domains being spoofed. We ran an additional campaign with key CNI organisations to encourage them to onboard NCSC threat intelligence sharing mechanisms, such as the Malware Information Sharing Platform (MISP) and CiSP.

The report is broken down on a service-by-service basis, but the effort expended is not siloed as each service influences, supports and guides the others. Supportive flows exist across the ACD portfolio and are growing as the efforts mature. As with last year's report, we have tried to focus on key stories and important trends that we have discovered in the course of our work, with the underpinning message that 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 in future reports, and comparisons or pointers to similar efforts. Please contact us at ACDenquiries@ncsc.gov.uk or via our social media and normal contact channels.

Takedown

www.ncsc.gov.uk/information/takedown-service

About the service

The NCSC's Takedown Service reached a five-year milestone in 2021. The service, run by Netcraft, finds 'bad stuff' hosted on the internet and seeks to have it removed (the goal being to remove cyber security threats quickly to minimise the harm to members of the public - or organisations - who could fall prey to them).

As with previous papers 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 a group:

- an 'attack' is a single URL involved in a campaign
- an 'attack group' is how we refer to all the URLs that form part of a campaign

Progress in 2021

2021 total takedowns

In total, 2.7M campaigns (3.1M URLs) were taken down in 2021. This is a significant increase when compared with 2020's tally (700,595 campaigns and 1,448,214 URLs) and is principally due to the prolonged period we have been performing takedowns against extortion mail server and celebrity-endorsed investment scams throughout 2021. These attacks are widely distributed and generate a proportionally large number of takedown records.

Table 1 Total takedowns by attack campaign group, 2020 and 2021

Attack Type	2020	2021
Extortion Mail Server	179,008 (Nov-Dec)	1,867,435
Celeb Endorsed Investment Scams	290,345 (Apr-Dec)	607,723
Fake Shop	160,295 (Apr-Dec)	107,251
Phishing URL	33,964	54,382
Web Shell	5,323	26,060
Advance Fee Fraud	27,346	19,197
Technical Support Scam	1,450 (Nov-Dec)	14,448
Advance Fee Fraud Mail Server	2,686	6,635
Malware Infrastructure URL	4,755	4,668
Vulnerable Application	-	4,050
Phishing URL Mail Server	6,849	53,437
Malware Attachment Mail Server	7,839	2,580
Malware Distribution URL	5,198	2,188
Malware C2 IP	880	1,767
Web-Inject Malware URL	1,616	1,358
Fake Pharmacy	797	881
Shopping Site Skimmer	1,505	875
Instagram Brand Infringement	266	723
Malware Command and Control Centre	720	682

Attack Type	2020	2021
Facebook Brand Infringement	65	327
Clone Firm Email	161 (Nov-Dec)	319 (Jan-May)
Clone Firm URL	132	224
Twitter Brand Infringement	38	206
Cryptocurrency Miner	135	133
Survey or Affiliate Scam	-	133
Phishkit Archive	150	117
TikTok Brand Infringement	-	71
Fake Mobile App	67	67
Other URL	6	65
Advance Fee Fraud Phone Number	186 (Nov-Dec)	65 (Jan-June)
DKIM Signed Email Domain	314	64
Clone Firm Phone Number	124	45
Brand Infringement	33	44
Skimmer Credential Dropsite	66	33
Technical Support Scam Number	-	32
JavaScript Resource	74	24
Phishkit Email	84	18
Fraudulent Use of PayPal on Fake Shops	-	12
Telegram Brand Infringement	-	8
Fake Bank URL	4	4
Phishing Dropsite	4	4
WhatsApp Brand Infringement	3	2
Other Phone Number	1	2
Other Email	3	2
Business Email Compromise	-	1
Fake Bond Comparison Site	-	1
Google Adwords	-	0
Malware URL Mail Server	11	0
Defaced Website	4	0
LinkedIn Impersonation	3	0
Malware Payment URL	3	0
Code Repository Sensitive Data	2	0
Credential Drop URL	2	0
Domain	1	0
Targeted Attack	1	0

Outcomes

UK government-themed takedowns

In 2021, we removed a total of 24,612 campaigns that used UK government lures. This was a reduction compared to 2020's figure of 27,611. As with previous years, phishing was the most prevalent attack type in this category.

Table 2 Government-themed takedowns

Attack Type	Number of attacks (URLs)	Total Groups	Median availability (hours)
Phishing URL	49,690	11,001	14
Advance Fee Fraud	4,753	4,753	1
Phishing URL Mail Server	3,250	3,250	25
Advance Fee Fraud Mail Server	1,777	1,777	25
Malware Attachment Mail Server	1,531	1,531	25
Web Shell	1,984	608	35
Instagram Brand Infringement	553	553	19
Facebook Brand Infringement	306	306	28
Twitter Brand Infringement	206	206	1,548
Malware Distribution URL	484	184	2
Malware Infrastructure URL	531	181	1
Phishkit Archive	123	116	21
TikTok Brand Infringement	71	71	1,445
Fake Mobile App	67	67	32
Brand Infringement	40	39	27
DKIM Signed Email Domain	27	27	23
Malware Command and Control	99	26	4
Phishkit Email	17	17	18
Other URL	6	6	74
WhatsApp Brand Infringement	2	2	1,105
Business Email Compromise	1	1	284
Google Adwords	1	1	40
Other Phone Number	1	1	34

UK government-themed phishing

In 2021, we took down 11,001 phishing campaigns, a total of 49,690 URLs with a median availability of 14 hours. The median availability of these campaigns was 7 hours shorter than last year.

Table 3 Comparative UK government-themed phishing attack availability, 2020-2021

Measure	2020	2021
Mean (hours)	177.3	139.5
Median (hours)	21.2	14.3
Skewness	6.3	9.4
25th Percentile	3.5	2.0
75th Percentile	111.6	59.8
Down in 4 hours	26.6%	32.9%
Down in 24 hours	51.9%	57.7%

In January 2021, UK government-themed phishing hosting was dominated by Namecheap with nearly 70% of attacks hosted on their infrastructure. Slow response times to takedowns may have made Namecheap an attractive option for phishing actors. However, this trend was reversed during 2021 as illustrated in the data in the figure below.

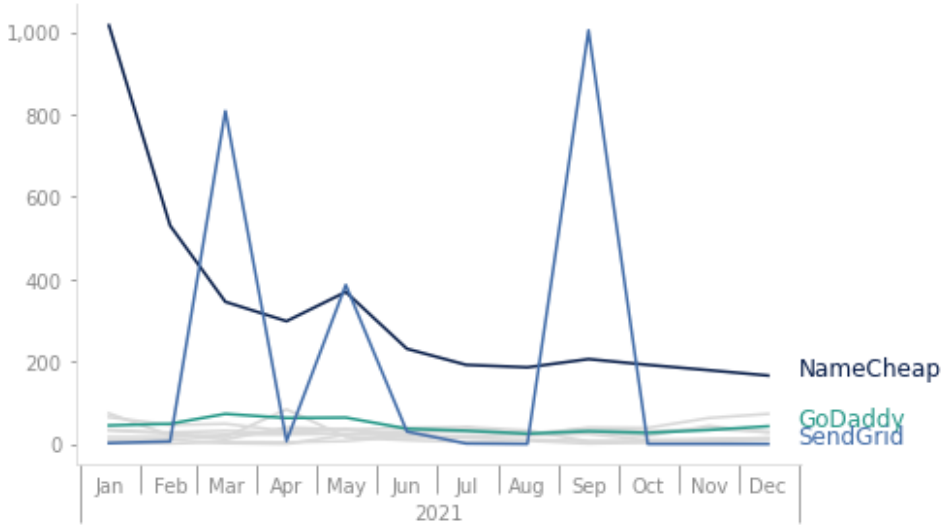

Figure 1 UK government-themed phishing campaign hosting 2021

In late January 2021, the NCSC engaged with Namecheap directly. We shared statistics relating to the number of phishing campaigns and how long it took for these attacks to be removed (median availability). Namecheap agreed to help and provided escalation routes for both COVID-19 and UK government-themed attacks. At that time, they were also changing their abuse processes and, by February 2021, we could see signs that Namecheap were managing phishing abuse more efficiently.

Namecheap's efforts to promptly remove the phishing content reported to them resulted in attack group medians dropping from 66 hours to less than 2 hours by April 2021. Noting the drop in attack availability, we can speculate that the value proposition for phishing actors using Namecheap was dramatically lowered. A more diverse set of UK phishing hosting began to evolve throughout the rest of the year as threat actors moved away from their platform. We hope that other brand owners are seeing similar improvements with Namecheap and their abuse management.

Figure 2 Takedown median (hours) for Namecheap hosted UK government-themed phishing in 2021

Comparing 2021's hosting with 2020, we see some new companies entering the top 10 hosters of UK government-themed phishing such as Sendgrid.com, Newfold Digital and MultiDC.net.

Table 4 Top 10 hosters of UK government-themed phishing

2020			**2021**		
Hoster	**Share (%)**	**Median availability (hours)**	**Hoster**	**Share (%)**	**Median availability (hours)**
Namecheap	28.8	47	Namecheap	30.2	9
GoDaddy	11.2	37	Sendgrid.com	20.4	15
OVH	4.8	6	Alibaba Group	4.2	21
Amazon	4.7	4	GoDaddy	3.8	24
Endurance	3.9	23	Amazon	3.1	13
Shinjiru	3.6	14	MultiDC.net	2.5	417
Cloudflare	3.0	12	Google	2.50	493
Alibaba Group	2.7	28	Digital Ocean	1.80	18
Digital Ocean	1.7	18	Endurance	1.8	49
Hostkey	1.6	16	Newfold Digital	1.8	7

The top 10 most targeted UK government brands reflects how phishing actors have continued to use lures relating to the COVID-19 pandemic.

Table 5 Top 10 UK government-phished brands

Government brand	**Number of attacks (URLs)**	**Number of attack groups (campaigns)**	**Median availability (hours)**
Generic 'gov.uk'	18,037	5,257	15
HMRC	12,516	2,592	11
NHS	5,513	1,405	8
TV Licensing	4,412	1,089	154
DVLA	6,418	1,013	17
Office for National Statistics (Census theme)	572	354	1

Government brand	Number of attacks (URLs)	Number of attack groups (campaigns)	Median availability (hours)
Government Gateway	1,334	267	21
BBC	103	51	20
Council Tax	94	39	3
Generic 'HMG'	29	27	30
All UK government-themed phishing attacks	**49,228**	**11,001**	**14.3**

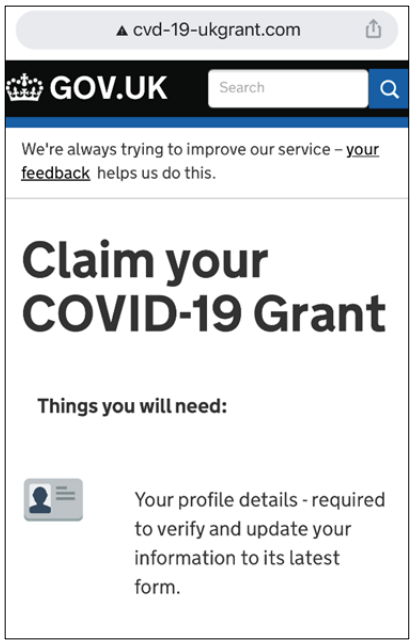

Figure 3 Generic 'Gov.UK' phishing with a COVID-19 grant lure

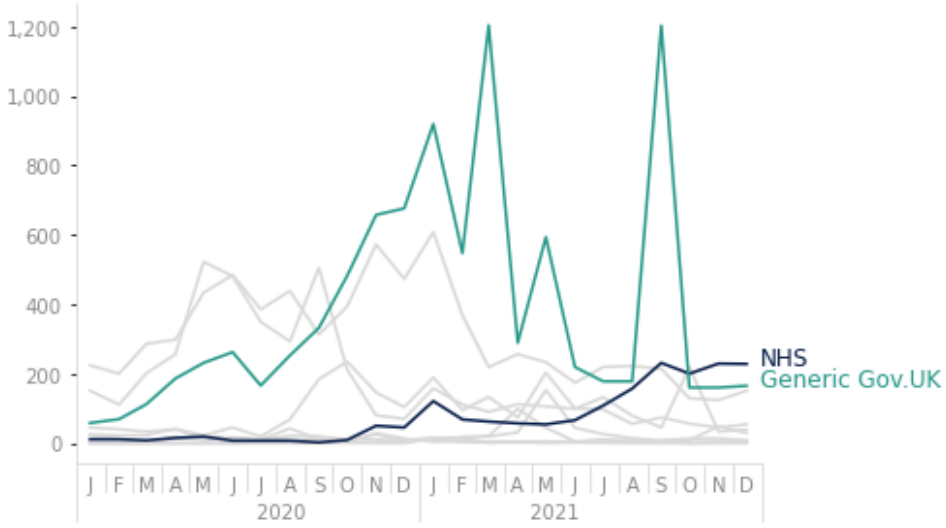

Figure 4 UK government phishing by brand in 2021

The first campaigns to use the NHS vaccine lure were noted in late December 2020 and were delivered in email and SMS campaigns with over 70 throughout January 2021. These attacks tailed off in numbers until summer when vaccine certification became a popular topic for lures.

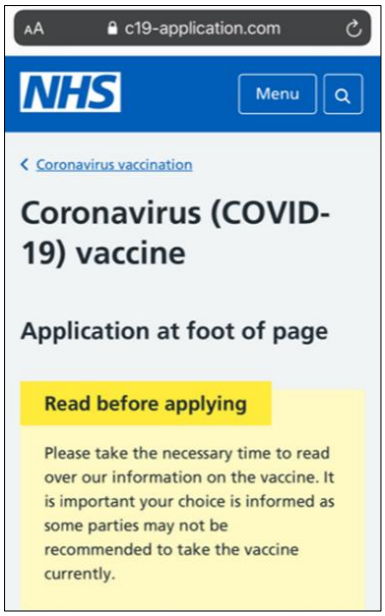

Figure 5 Fake NHS COVID-19 Vaccine Booking (January 2021)

Figure 6 Fake NHS Digital Vaccine Passport Phishing (June 2021)

These campaigns were designed to harvest personal and financial information from victims. The phishing sites falsely offered vaccine booking appointments in return for a small 'fee', but as is common with other phishing idioms, personal and financial information posted into these phishing forms was subsequently used by phishers to enable further fraud, often contacting victims directly purporting to be from UK banks.

By summer 2021, phishing actors modified the lure to offer fake vaccine passports which falsely claimed to support international travel requirements. Some even provided a QR code which looks authentic but when scanned simply redirected the victim to a free QR code generation site. In March 2021, the Office for National Statistics (ONS) ran the first digital census for England and Wales. Phishing actors exploited the period after this event with campaigns that falsely threatened fines for late census returns. Again, the goal was to harvest personal and financial information for use in further fraud.

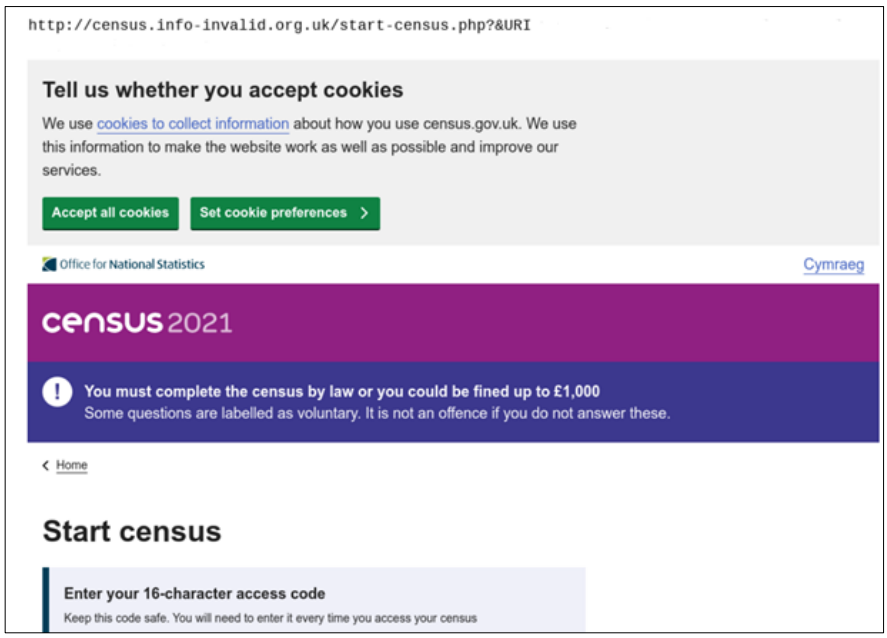

Figure 7 ONS-branded census phishing attack

85% of these campaigns were hosted on Namecheap's infrastructure and, due to a combination of timely reporting from both the UK public (via SERS) and ONS's own Security and Information Management team, these were reported into Takedown with Namecheap removing the content promptly. The median attack availability for the 354 campaigns was an impressive 1 hour and 5 mins throughout 2021.

UK government-themed advance fee fraud

The UK government continued to be a popular theme in many advance fee fraud scams in 2021, with the National Lottery the most targeted. We took down 4,753 attacks in 2021 with a median attack availability of 1 hour.

Table 6 Top 10 UK government brands found in advance fee fraud attacks

Government brand	Number of attacks (URLs)
National Lottery	511
Financial Conduct Authority	459
Bank of England	392
Ministry of Justice	370
British Broadcasting Corporation	123
Metropolitan Police	44
Department for Exiting the European Union (Brexit)	43
HM Treasury	37
National Crime Agency	25
Prudential Regulation Authority	18

Even the NCSC was not immune from being used as a lure this year, as the following advanced fee fraud example shows.

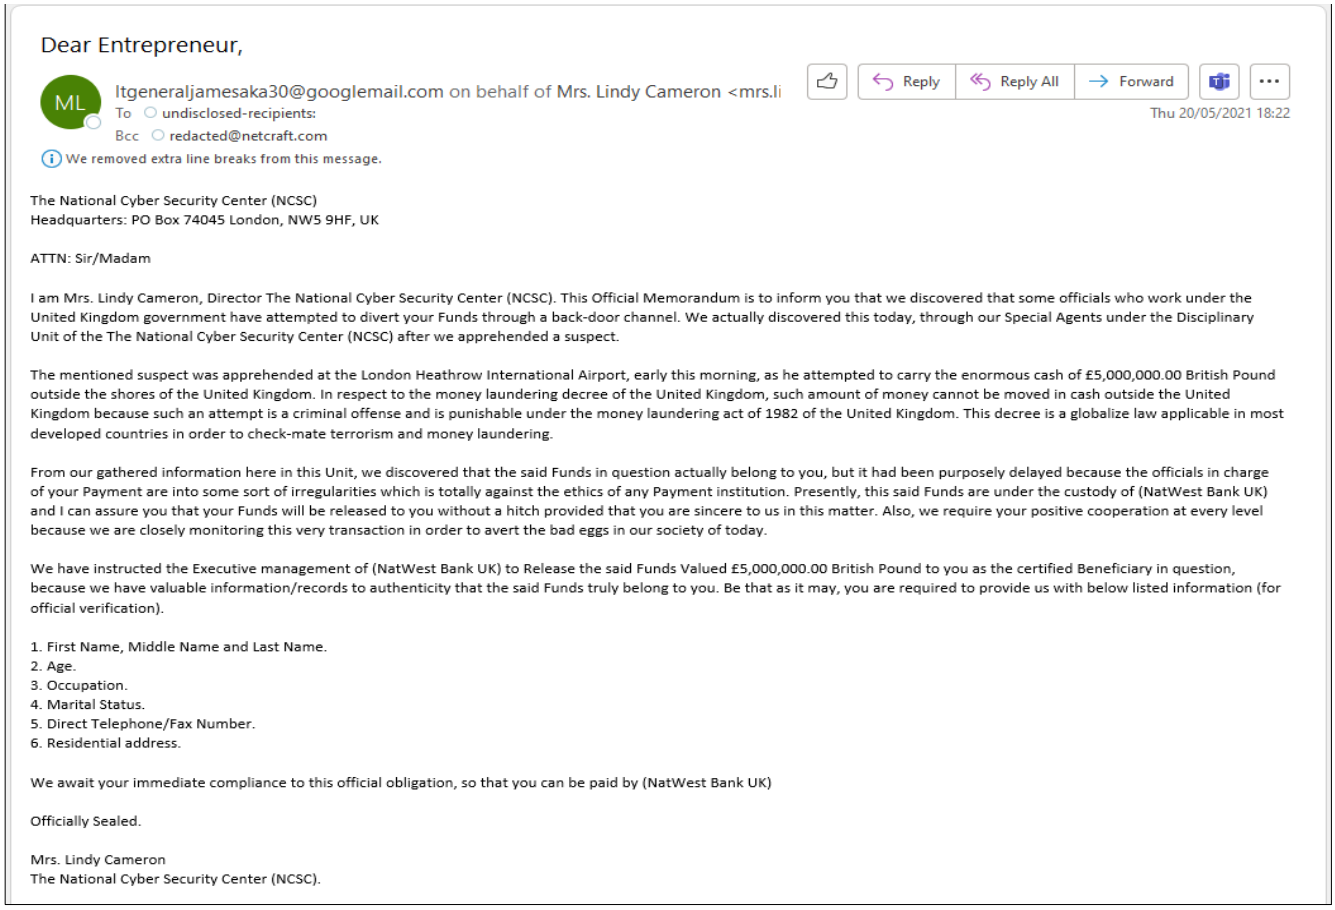

Dear Entrepreneur,

ML

ltgeneraljamesaka30@googlemail.com on behalf of Mrs. Lindy Cameron <mrs.li

To undisclosed-recipients:

Bcc redacted@netcraft.com

Reply | Reply All | Forward

Thu 20/05/2021 18:22

We removed extra line breaks from this message.

The National Cyber Security Center (NCSC)
Headquarters: PO Box 74045 London, NW5 9HF, UK

ATTN: Sir/Madam

I am Mrs. Lindy Cameron, Director The National Cyber Security Center (NCSC). This Official Memorandum is to inform you that we discovered that some officials who work under the United Kingdom government have attempted to divert your Funds through a back-door channel. We actually discovered this today, through our Special Agents under the Disciplinary Unit of the The National Cyber Security Center (NCSC) after we apprehended a suspect.

The mentioned suspect was apprehended at the London Heathrow International Airport, early this morning, as he attempted to carry the enormous cash of £5,000,000.00 British Pound outside the shores of the United Kingdom. In respect to the money laundering decree of the United Kingdom, such amount of money cannot be moved in cash outside the United Kingdom because such an attempt is a criminal offense and is punishable under the money laundering act of 1982 of the United Kingdom. This decree is a globalize law applicable in most developed countries in order to check-mate terrorism and money laundering.

From our gathered information here in this Unit, we discovered that the said Funds in question actually belong to you, but it had been purposely delayed because the officials in charge of your Payment are into some sort of irregularities which is totally against the ethics of any Payment institution. Presently, this said Funds are under the custody of (NatWest Bank UK) and I can assure you that your Funds will be released to you without a hitch provided that you are sincere to us in this matter. Also, we require your positive cooperation at every level because we are closely monitoring this very transaction in order to avert the bad eggs in our society of today.

We have instructed the Executive management of (NatWest Bank UK) to Release the said Funds Valued £5,000,000.00 British Pound to you as the certified Beneficiary in question, because we have valuable information/records to authenticity that the said Funds truly belong to you. Be that as it may, you are required to provide us with below listed information (for official verification).

1. First Name, Middle Name and Last Name.
2. Age.
3. Occupation.
4. Marital Status.
5. Direct Telephone/Fax Number.
6. Residential address.

We await your immediate compliance to this official obligation, so that you can be paid by (NatWest Bank UK)

Officially Sealed.

Mrs. Lindy Cameron
The National Cyber Security Center (NCSC).

Figure 8 NCSC-themed advance fee fraud

Takedowns in UK delegated IP space

In this section we will look at the types of attacks we see in UK-delegated IP space. As in previous years, we have continued to discover and remove attacks hosted in the UK, regardless of the brand targeted.

Table 7 Takedowns in UK IP space

Attack Type	Number of attacks (URLs)	Number of attack groups (campaigns)	Median availability (hours)
Phishing URL	113,457	19,939	10
Web Shell	12,969	4,649	47
Fake Shop	5,815	1,737	1,113
Fake Pharmacy	3,386	881	33
Web-Inject Malware URL	1,517	746	88
Shopping Site Skimmer	1,246	522	85
Malware Distribution URL	341	265	30
Malware Infrastructure URL	278	139	82
Cryptocurrency Miner	324	89	78
Malware C2 IP	47	53	223
Skimmer Credential Drop site	30	22	79
Malware Command and Control Centre	21	18	72
JavaScript Resource	17	16	527
Technical Support Scam	65	10	34
Phishing Drop site	4	4	85

Phishing was again the most prevalent attack type we found in UK IP space, and the service continued to monitor the UK share of global phishing throughout the year.

Figure 9 UK percentage share of Global Phishing 2016-2021

The UK share of global phishing varied around the ~2% mark throughout 2021.

Table 8 Comparative brand-agnostic UK-hosted phishing attack availability 2020-2021

Measure	2020	2021
Mean (hours)	122.7	116.5
Median (hours)	14.5	10.0
Skewness	8.4	8.4
25th Percentile	2.2	2.0
75th Percentile	57.5	49.2
Down in 4 hours	31.9%	34.9%
Down in 24 hours	58.3%	64.7%

Web inject malware in UK IP space

This year we took down 1,358 instances of web inject malware present on compromised UK hosted web sites. This is an increase of 471 on from last year. Median availability of these attacks also increased from 66 to 99 hours.

Shopping site credit card skimmers

We take down credit card skimming malware code in UK IP space, and also sites which offer financial transactions in UK sterling (which could be hosted anywhere).

- In 2021, we took down 875 instances of skimming code (575 less than 2020).
- Of this total, 522 were UK hosted with a median availability of 85 hours and a further 353 were hosted overseas where we noted drop in median availability to 95 hours (from 206 hours in 2020).
- As with previous years, Magento Cart was the most commonly noted platform in 2021 skimmer takedowns.

Cryptocurrency miners

We found 89 instances of non-consensual cryptocurrency mining code on UK hosted websites during 2021 with a median availability of 78 hours. This type of attack exploits visitors to the website to mine cryptocurrency without the visitor or site-owner's permission. Analysis of the embedded mining keys in the code showed that CoinImp and WebMinePool keys were the most prevalent in our takedowns in 2021.

Poisoned JavaScript resources

By modifying JavaScript code snippets which are used by multiple websites, attackers can 'poison' a JavaScript resource with malicious code which will subsequently infect the systems of any organisation that uses that code. We took down 24 instances of poisoned JavaScript resources, with a median attack availability of 192 hours during 2021. 17 were hosted in the UK, with the remainder hosted overseas.

Web shells and control panels

The service has made good progress on detecting and taking down malicious web shells and control panels in recent years. These panels are often seen active on compromised servers and enable the attacker to maintain access to a compromised machine and orchestrate malicious activity. By removing these shells and panels we hope to disrupt their use of compromised infrastructure by removing the common tools they use.

- In 2021 we took down 4,649 web shells (12,969 URLs) in UK IP space, with a median availability of 47 hours.
- We also took down a further 608 web shells (1,984 URLS) associated with UK government-themed phishing, with a median availability of 35 hours.
- We also took down an additional 20,617 web shells (76,665 URLs) that were associated with other phishing attacks hosted overseas, with a median availability of 25 hours.
- March 2021 was the busiest month with 7,709 web shell takedowns associated with UK hosted Microsoft Exchange Servers vulnerable to CVE-2021-26855.

Post exploit web shell infection of vulnerable infrastructure

In March 2021, Microsoft patched a vulnerability (CVE-2021-26855) in Microsoft Exchange Server (on-prem). The vulnerability enabled attackers to bypass authentication and gain unauthorised access to them. This vulnerability was exploited quickly by attackers; tens of thousands of servers were compromised globally prior to administrators applying patches to mitigate this vulnerability. As an attacker, finding a vulnerable Microsoft Exchange Server is simple, and deploying a web shell would be a direct way of maintaining access after a patch was applied.

Between March and the end of 2021, the Takedown service found 2,920 UK hosted exchange servers compromised with web shells, attempted to notify the owners (through emails to postmaster@ and admin@ accounts), and monitored for their removal.

Additionally, the NCSC was able to take the hostnames found in our takedown data and use them to seed further notifications for organisations who had subscribed to the NCSC Early Warning service.

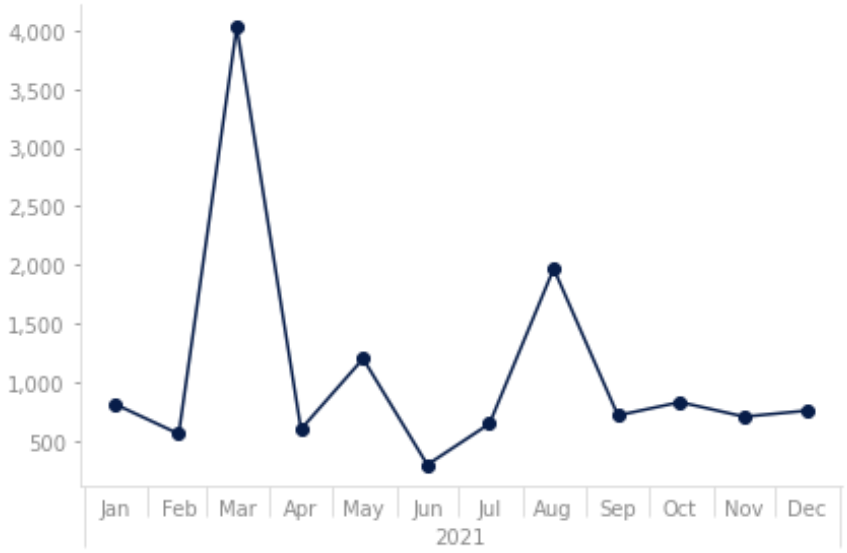

Figure 10 Web shell takedowns per month in 2021 (peak in March due to Microsoft Exchange vulnerability)

COVID-themed takedowns

We continued takedowns against generic COVID-themed attacks until the end of September 2021. Similar to 2020, the most prevalent type of attack which used this type of lure was advance fee fraud (aka the 419 scam). A high percentage of these attacks used Google Mail accounts which were promptly removed upon notification with a median of just 2 hours throughout 2021.

Table 9 COVID-themed takedowns in 2021 by attack type

Attack Type	Number of attack groups (campaigns)	Median availability (hours)
Advance Fee Fraud	14,453	2
Advance Fee Fraud Mail Server	4,855	25
Malware Attachment Mail Server	1,049	26
Phishing URL Mail Server	187	25
Instagram Brand Infringement	170	144
Survey Scam	133	198
Malware Infrastructure URL	72	49
Malware Distribution URL	50	19
Phishing URL	45	24
DKIM Signed Email Domain	37	22
Facebook Brand Infringement	21	502
Fake Shop	15	615
Cryptocurrency Investment Scam	13	1
Malware Command and Control Centre	13	20
Telegram Brand Infringement	8	730
Brand Infringement	4	131
Other URL	2	80
Other Phone Number	1	74
Phishkit Email	1	8
Other Email	1	1

Fake online shops

We've been discovering and taking down fake shops since April 2020. These sites don't map to real businesses and offer large discounts on popular goods. Victims ordering goods via these sites will likely be defrauded having posted payment details and personal information.

In 2021, the number of fakes shop campaigns has fallen to 107,251 (from 139,522 in 2020). However, we note long median attack availability for these sites. In 2021 the median attack availability for fake shops was 481 hours (compared with 341 hours in 2020).

Hosting of these sites has been consistent since we began takedowns with Fibergrid Group having the largest share.

Table 10 Top 10 hosters of fake shops in 2021

Hoster	Number of attack groups (campaigns)	Share (%)	Median availability (hours)
Fibergrid Group	34,165	31.9	723
Cloudflare	32,984	30.8	121
Internet Bilgisayar	14,021	13.1	694
Google	4,281	4.00	166
Scaleway	3,563	3.32	868
plexus-network.com	2,868	2.67	667
CloudRadium	2,342	2.18	700
Rebel Hosting	2,108	1.97	204
Alibaba Group	1,698	1.58	812
ProHoster	986	0.92	2,180

In terms of site registration, Alibaba Group was the registrar of choice with 30% of domain registrations in this attack category.

Fake celebrity endorsement scams

We started these takedowns in 2020 and they have continued to be a widespread attack promoted in spam, SMS and via online adverts. As before, none of the celebrities featured in these scams are aware or involved in any way. During 2021 we took down 319,365 campaign groups (607,723 URLs). These attacks had a median attack availability of just 1 hour.

Figure 11 Example of fake celebrity endorsement scams

Throughout 2021, the monthly takedown totals for this type of scam have shown a downward trend.

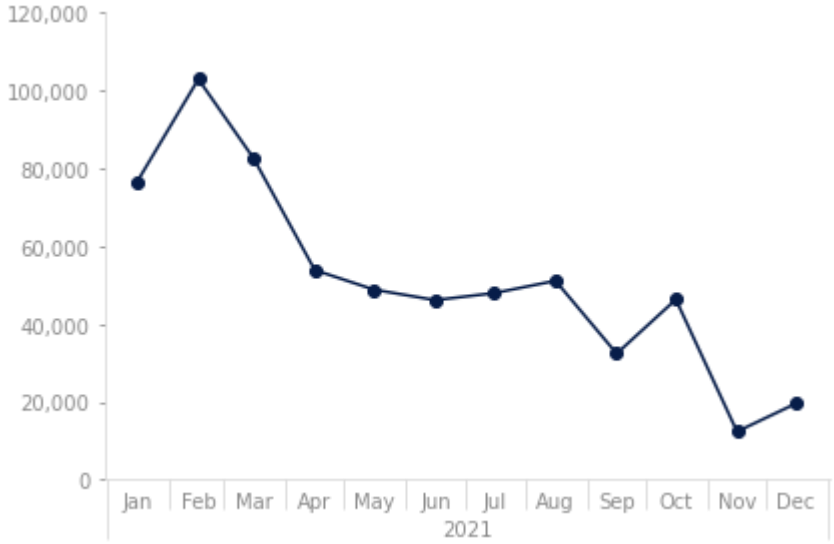

Figure 12 Celebrity endorsed scams in 2021 by campaign groups

The median attack availability has fallen dramatically with many hosters removing these pages quickly and at scale.

Table 11 Top 10 hosters of celebrity-endorsed scams in 2020 and 2021

2020			**2021**		
Hoster	**Share (%)**	**Median Availability (hours)**	**Hoster**	**Share (%)**	**Median Availability (hours)**
Amazon	46.4	19	Amazon	76.6	0.5
OVH	5.8	292	Google	4.9	68
Digital Ocean	5.8	11	Cloudflare	2.5	56
Cloudflare	5.2	123	Namecheap	1.6	72
Google	4.3	340	OVH	1.0	0.5
GoDaddy	2	347	Bitly.com	0.8	1
Endurance International	1.1	36	uCoz Web Services	0.7	1
Alibaba Group	0.7	156	GoDaddy	0.6	60
Namecheap	0.7	245	Endurance International	0.6	51
Velocity Servers Network	0.1	125	Digital Ocean	0.6	124

Extortion emails

Like the celebrity endorsed scams, the extortion mail server takedowns are another example of a very prevalent attack which many readers may recognise. These attacks falsely claim that recipients have been *hacked* and some will attempt to add authenticity by quoting a password which a recipient may recognise. These passwords are often derived from publicly pasted breaches which the recipient may not realise they were part of. The mail goes on to urge recipients to purchase and send cryptocurrency to a crypto wallet or risk the release of 'compromising' material.

In 2021, we took down 1.87M servers sending these attacks with a median attack availability of 26 hours. Looking at the hosting of these servers we can see that they are widely distributed across the internet. The top 10 hosters of extortion mail servers in 2021 are below.

Table 12 Top 10 hosters of extortion mail servers in 2021

Hoster	Number of attack groups (campaigns)	Share (%)
Bharti Airtel	15,6516	8.4
PTCL	10,1788	5.5
Vodafone Group	57,436	3.1
Telefonica	52,092	2.8
iam.ma	31,736	1.7
ideacellular.com	28,084	1.5
America Movil	27,082	1.5
Antel	21,487	1.2
Hathway Cable & Datacom	21,316	1.1
Deutsche Telekom	21,193	1.1

Remote access trojans

Remote access trojans (RATs) are used by attackers to conduct more detailed reconnaissance of a victim device or network. These provide a backdoor to enable other capabilities, which an attacker could subsequently deploy.

In 2021, we were able to take down 5,944 instances of infrastructure used for RAT distribution and command and control (9,069 URLs). This is a large increase on last year's total of 1,733 (2,954 URLs). The main reason for this increase is due to improvements in detection and monitoring of these attacks. The median availability increased this year from 39 hours to 106 hours. Cobalt Strike C2 was the most prevalent RAT in our takedowns this year with a 60% share.

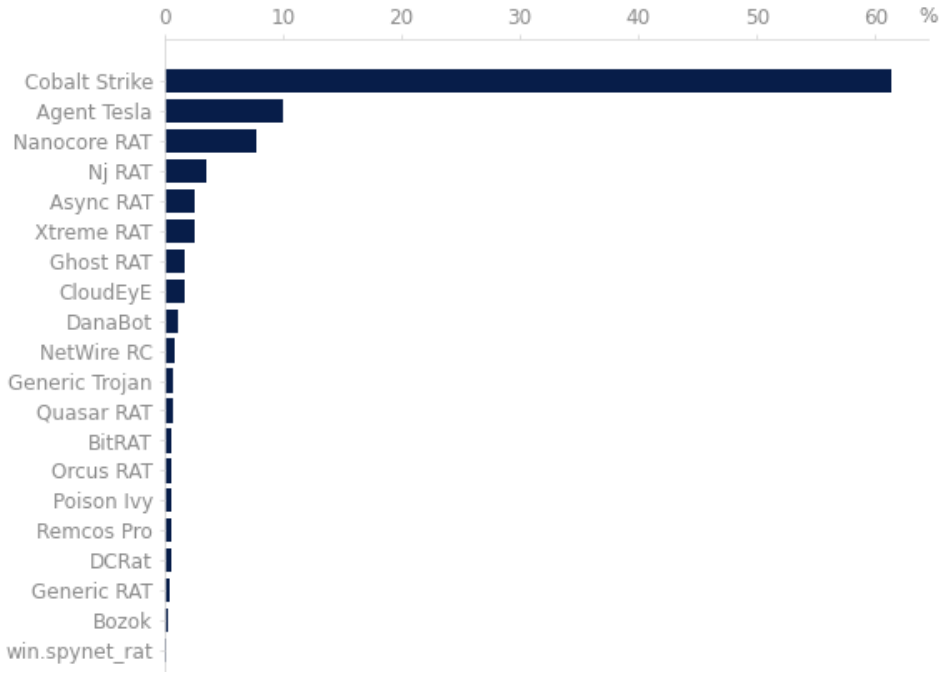

Figure 13 Remote access trojan takedowns by prevalence in 2021

Banking trojans

Upon installation on a victim's device or network, banking trojans (as the name suggests) have a more customised payload and will harvest banking credentials and exploit financial transactions. The service took down 1,956 banking trojan instances (4,874 URLs).

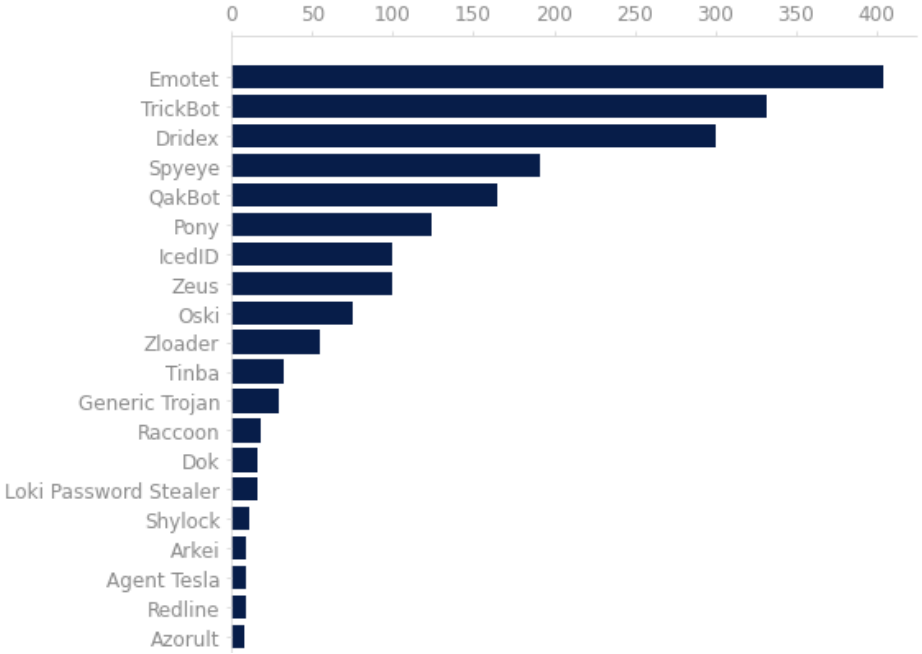

Figure 14 Banking trojan takedowns in 2021 by family

Emotet was one of the most prevalent banking trojans in 2020 and it still featured heavily in 2021 despite a coordinated international effort to disrupt it during January. The effect of this coordinated effort can be seen in our takedown data across 2021. However, by November we began to see an increase in Emotet deployments once more.

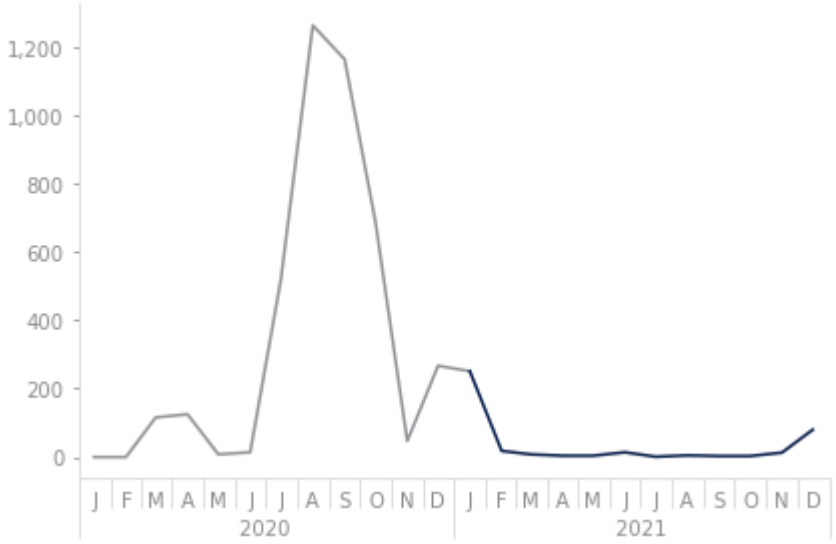

Figure 15 Emotet infrastructure takedowns in 2021

Active fraud defence

Following renewed funding from National Economic Crime Centre, we continued takedowns aimed at disrupting cyber-enabled economic crime and fraud.

Computer software service fraud (aka tech support scams)

Technical support scams will often appear when browsing the web in the form of an unwanted pop-up window. It will falsely claim there is an issue such as a virus affecting a victim's computer. Victims who engage with the pop-up will then be invariably contacted via telephony (often VoIP), email or online chat with scammers who will purport to be representatives from various reputable companies such as Microsoft or Apple to attempt fraud.

We took down 14,448 technical support scams in 2021 between January and end of October with a median availability of 15 hours. Microsoft were the most commonly spoofed in this category (78%). Having established that Microsoft were also acting upon these attacks, we decided to focus on other categories.

Table 13 Tech support scams by brand in 2021

Brand	Number of attack groups (campaigns)	Median Availability (hours)
Microsoft	11,279	18
Apple	3,661	8
McAfee	626	33
Norton	76	13
Google	20	11
Support Scams (HMG)	5	116

FCA investment warnings

In November 2020 we started takedowns against fake or clone investment companies that had been flagged by the Financial Conduct Authority (FCA) as fraudulent. We continued with these takedowns until end of May 2021. The FCA decided at this point to procure their own takedown countermeasures and formally took this activity on. We consider this to be an excellent outcome.

In 2021, we performed 775 takedowns in this category.

Table 14 FCA investment warning takedowns in 2021

Attack Type	Number of attack groups (campaigns)	Median Availability (hours)
Clone firm Email	319	14
Clone firm URL	224	121
Other URL	57	123
Clone firm phone number	45	26
Fake bank URL	4	2

UK telephony found in advance fee fraud

Between January and end of June 2021, we continued to extract UK mobile and landline numbers from advance fee fraud emails in order to perform takedowns (that is, get the UK numbers taken out of service). During this period, we took down 65 numbers, including one which called the Takedown service directly (purporting to be from HMRC demanding a tax return). That particular number belonged to Virgin Media and was taken out of service just 23 minutes later, though the median for takedowns of this type was generally much higher at 40 hours.

Additional takedowns from Suspicious Email Reporting

Having enabled public reporting to the Suspicious Email Reporting service (SERS) (report@phishing.gov.uk) in 2020, we acted against as many malicious referrals to SERS as we could. In 2021, we took down an additional 24,972 phishing campaigns across several different industry sectors. The median availability of these campaigns was 7 hours 30 minutes.

Takedowns derived from 7726 via participating UK mobile networks

Since 2020, we have been receiving URLs derived from public/consumer reporting to short code 7726. These reports are from UK mobile consumers on participating UK networks. These UK reports are aggregated by cyber security vendor Proofpoint. Our Takedown service receives URLs from these reports and generates new takedowns when malicious content is found. In 2021, we took down 4,350 campaigns from this source.

Conclusions after 5 years

The NCSC's Takedown service began in June 2016 and continues to evolve against new types of attacks. The integration of public reporting via SERS (and 7726) with Action Fraud means we are now able to better drive our takedown activities to match the commodity cyber crime categories associated with reported financial losses.

Over the last 5 years, the service has taken down 3,729,404 campaign groups (5.8M URLs covering 2,019,550 IP addresses). We have halved the UK share of global phishing whilst significantly reducing the lifecycle of commodity cyber attacks.

From the outset, the NCSC has been transparent regarding the effect that takedown countermeasures have on attack lifecycles. Our continued hope is that other nations, National CERTs, and other organisations employ similar services to amplify the effect of this work.

Suspicious Email Reporting Service

www.ncsc.gov.uk/collection/phishing-scams

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 forwards these potentially suspicious emails onto our takedown provider who analyses the emails, and when links to malicious sites are found, seeks to remove those sites from the internet to prevent them doing further harm.

Progress in 2021

SERS transitioned into a Live phase in March 2021. SERS progressed from Public Alpha in April 2020, through Beta, to Live in less than a year, marking significant progress in development. As a result, SERS is now in a sustain stage with regular reviews for service improvements. This year, we have made multiple enhancements to the service.

Firstly, we have amended the number of times that a submitter receives the auto response email. 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. The submitter was previously receiving this email after every submission. We have now reduced this to one email being sent within a 24-hour period, irrespective of how many reports were sent.

Secondly, we released guidance for Office 365's 'Report Phishing' add-in for Outlook. This allows users, at the click of a button, to send potential scam emails directly to SERS as well as their own organisation's IT team. In addition, we also offer the functionality of users being able to report a website directly to us through the NCSC website.

Thirdly, SERS now has the capability to share the URLs collected from the public reports with other parties for their own independent analysis and to contribute towards wider protection of their own customers. These third parties are carefully vetted takedown providers that receive the data stream directly from the NCSC following signature of relevant data sharing agreements.

The service has received substantial public attention with multiple mentions within mainstream media which has contributed to SERS receiving over 10 million reports since its inception in April 2020. SERS has also completed a user survey, which was open for a week and was only promoted via our auto response email. We had over 4,000 responders to the survey, of which, 78% said they would recommend the service.

SERS has continued to maintain its strong relationships with external partners and law enforcement to provide vitally important data that helps them to understand the strategic threat to the UK, alert the public to phishing trends, and identify investigative leads

Outcomes

Between January and December 2021, SERS received just over 5.4 million reports, an average of 14,800 a day. These reports identified just over 2M malicious URLs, many of which were also identified by the Takedown Service. SERS was credited in instigating the removal of more than 44,000 scams, involving 82,000 malicious URLs, not previously identified by the Takedown Service. SERS was used by 375,000 unique submitters with 62.8% of submissions being spread across BT Internet, Hotmail, Gmail, and Yahoo domain users.

We have had 6 'super users' who have each submitted over 10,000 reports to SERS, and a further 23 users who have each submitted over 5,000 reports. You know who you are; thank you.

Mail Check

www.ncsc.gov.uk/information/mailcheck

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 and MTA-STS) - keeping messages encrypted and private as they are sent over the internet

Progress in 2021

Scaling up whilst reducing costs

During 2021, we grew our user base by 27%; growing the number of organisations we support from 1,209 to 1,530. This growth is primarily from the schools and charity sectors.

We also focussed on reducing costs as Mail Check continues to progress as a mature live service, making changes to our platform to reduce infrastructure costs, in addition to reducing our team size by 30%.

Table 15 Types of organisation using Mail Check in 2020 and 2021

Sector	Organisations using Mail Check end Dec 2020	Organisations using Mail Check end Dec 2021	Change[1]
Central government departments and arms length bodies	138	123	-15
Local government	409	352	-57
Health	184	190	+6
Police and fire and rescue services	82	63	-19
Devolved administrations and their agencies	59	66	+7
Academia (universities, colleges, schools)	256	515	+259
Charities	34	188	+154
Other	47	33	-6
Total	**1,209**	**1,530**	**+321**

1 Reductions are a result of disabling dormant accounts (ie where users have not logged into Mail Check for 12 months).

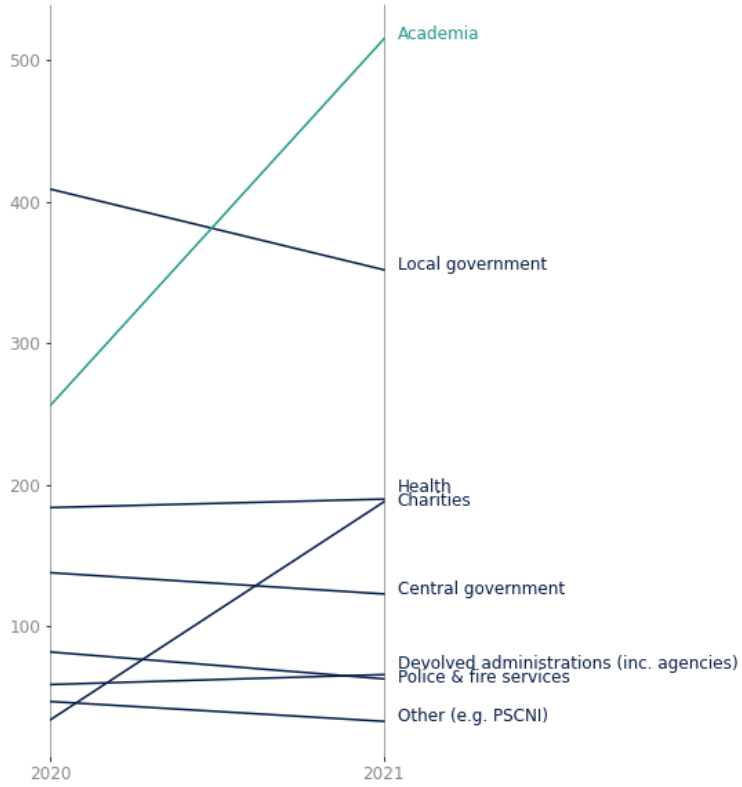

Figure 16 Organisations using Mail Check

Protecting email privacy from downgrade attacks: MTA-STS

In 2021, Mail Check introduced full support for the internet standard MTA-STS (Mail Transfer Agent Strict Transport Security) and the supporting standard TLS-RPT. We also published updated guidance to include these protocols.

Well-configured mail servers will send email across the internet using secure connections encrypted using Transport Layer Security (TLS). However, there remain vulnerabilities, whereby a malfeasant-in-the-middle (MITM) can trick incoming connections to send to another server and/or send information in the clear. MTA-STS is a standard designed to address these vulnerabilities.

It is early days still for a standard like MTA-STS, with not all email providers supporting the standard yet, but working with colleagues in the Government Security Centre for Cyber (Cyber GSeC), we have identified and progressed early adopters to utilise the standard, so that we can learn and prepare for a broader push in 2022. Through a campaign delivered by the Cyber GSeC there has already been an uptake of 14 ministerial / non-ministerial departments and an additional 30 arm's length bodies / public bodies either with, or testing, the standard.

The graph below illustrates the growth in the number of domains protected by MTA-STS throughout 2021.

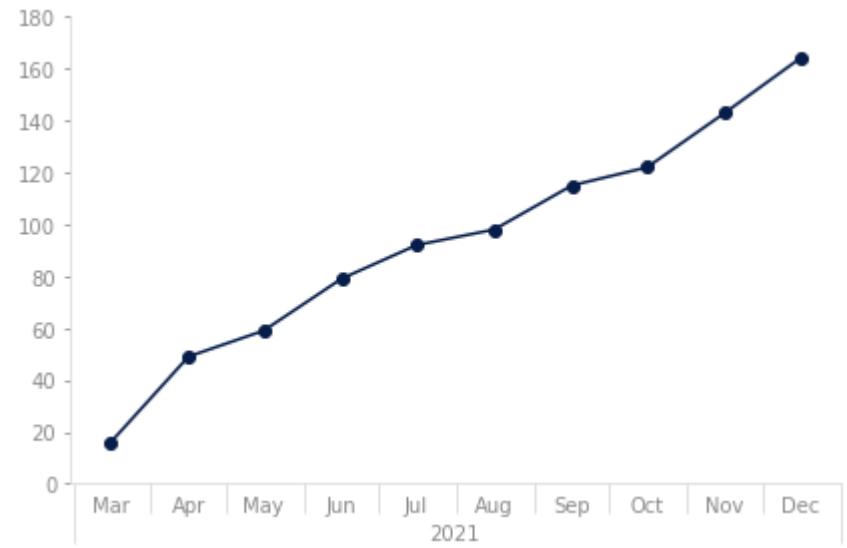

Figure 17 Domains protected by MTA-STS in 2021

Taking DMARC analysis to the next level: DMARC Insights

In 2021, we set out to improve the value and adoption of the most complex area of Mail Check: the analysis of DMARC reporting. All DMARC analysis tools like Mail Check provide an enriched, interactive graphical interface to help users identify and fix issues across the email sending systems. The challenge with this approach is that it takes a degree of skill and knowledge to really use the tool effectively, knowing how to filter out noise in the data, and to explore the graphs and data provided.

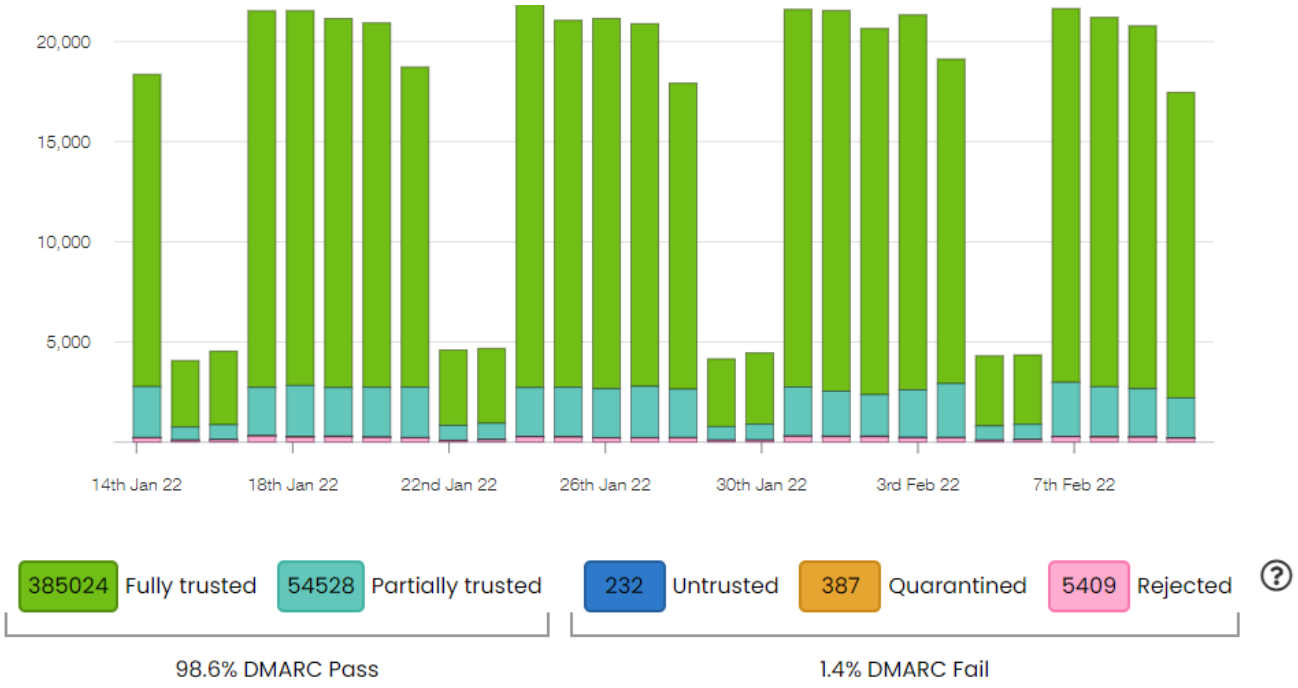

Figure 18 DMARC reporting graphical interface in Mail Check

That is where Mail Check DMARC Insights feature comes in. We developed an engine that could automatically pull-out key insights from the data. Types of insights extracted are:

- the volume of malicious email attacks that are using their domain name
- sub-domains we have discovered in their data that they need to configure
- the email sending systems we have discovered being used for their organisation, and whether these have been adequately secured

This really helps focus our users on the specific cyber security challenges, rather than spending time getting to grips with the data analysis. It also helps them sustain their cyber security position; in time we will switch on email notifications based on the insights we have derived from their data.

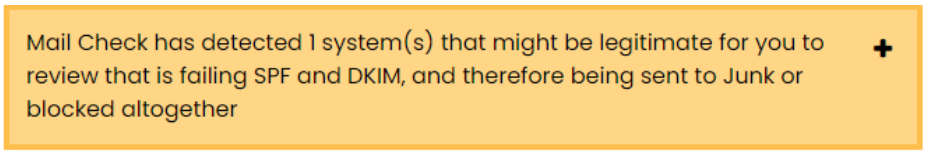

Figure 19 Example alert displayed in Mail Check

Expanding our services to charities and schools

In 2021, we expanded our user base to include the charity and school sectors. When looking at anti-spoofing controls like DMARC, we found that these sectors were at very low levels of adoption. Only 13% of the top 3,000 charities had effective DMARC policies in place, and only 6% of schools.

During 2021, we have onboarded approximately 200 organisations from each sector, with a significant number of both groups able to progress quickly to implement anti-spoofing controls. During 2022 we hope to significantly extend our reach and impact across these sectors, the challenge being scale with over 150,000 charities and over 30,000 schools.

Table 16 Organisations onboarded to Mail Check in 2021, based on the adoption of strong DMARC policies

	Number of organisations added in 2021	% of Mail Check organisations in this sector protected	% of sector protected
Charities	192	38%	13% (based on top 3,000 charities)
Schools	222	49%	6% (based on a sample of 1,000 schools)

Schools and charities both responded positively to Mail Check, with both sectors requiring different needs from the service.

Integrating the Mail Check service with MyNCSC

In 2021, we took the first steps to integrate Mail Check with the new platform MyNCSC. This involved:

- sending cyber security 'Findings' into MyNCSC so that users get an aggregated view with other cyber security services
- integrating with a central asset portfolio in MyNCSC, so that users can maintain a list of assets (domains in the case of Mail Check) that they want to check
- developing a model for migrating organisations, users and their list of domains to the shared platform
- supporting the migration of the first 250 organisations to MyNCSC

New 'Email security check' service – available to all

The Mail Check service is only available to certain eligible sectors: the UK public sector, academia and charities.

During 2021, we worked on developing the concept of a lightweight version of Mail Check that could be offered to all organisations. Something that provides a quick and simple way of understanding certain areas of email security (like anti-spoofing and email encryption), and acts as a gateway to NCSC services and tools (such as Mail Check and commercially provided alternatives).

The service was designed and built in 2021 and is due for release early 2022.

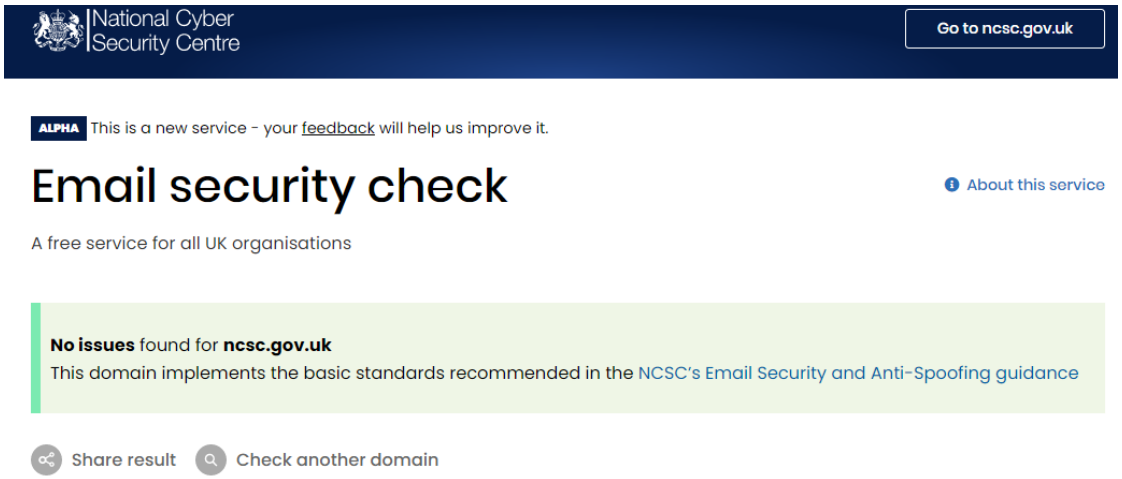

Figure 20 Screenshot showing the 'Email Security Check' service in development

Outcomes

The following table illustrates sector-based shifts from the end of 2020 to the end of 2021 for implementation of DMARC anti-spoofing controls.

Table 17 Sector shifts in the implementation of DMARC anti-spoofing controls

Sector	Subset of organisations tracked	31 Dec 2020 % orgs with EDP[2]	31 Dec 2021 % orgs with EDP	Change	Comment
Central government	44 (government departments + 10 Downing Street)	86%	91%	+5%	
Central government	223 arm's-length bodies	46%	60%	+14%	Includes executive agencies, non-departmental public bodies, and similar organisations set out in the Public Bodies 2019 report.
Local government	404 (UK principal councils)	75%	81%	+6%	
Health	279 NHS Trusts and key central functions	21%	35%	+14%	These numbers seem low but note that most NHS organisations now use the NHS Mail service with a DMARC policy of reject. The numbers reflect legacy NHS domains.
Devolved administrations and their agencies	Includes local authorities, health services and emergency services in devolved administration regions	45%	48%	+3%	Averaged across Scotland, Wales and Northern Ireland.
Police and fire services	51 police forces and 54 fire and rescue services	59%	76%	+17%	
Universities	164 universities, university colleges and other degree-awarding bodies	16%	23%	+7%	

[2] EDP = Enforcing DMARC Policy of quarantine or reject

Sector	Subset of organisations tracked	31 Dec 2020 % orgs with EDP[2]	31 Dec 2021 % orgs with EDP	Change	Comment
FE colleges	375 further education colleges	17%	23%	+6%	
Charities	Top 3,000 charities	10%	13%	+3%	Still early days with only 200 charities on the Mail Check platform.
Schools	Sample of 1,000 schools (from 32,000)		6%	N/A	Still early days with only 200 schools on the Mail Check platform.
For comparison - these sectors below are not eligible to use Mail Check					
Private sector CNI organisations	235 organisations	30%	42%	+12%	Not eligible for Mail Check.

Web Check

www.ncsc.gov.uk/information/web-check

About the service

Web Check is an ACD service that helps over 1,500 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 in 2021

Product updates and improvements

The types of security issues seen in websites tend 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. Maintenance activities included reworking of scanning infrastructure for easier maintenance, detection of PHPinfo and reporting to the user, as well as updating the security.txt detections.

It is generally impractical for Web Check to scan for specific vulnerabilities; the focus being 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.

One such instance of this occurred in December when the Apache Log4j vulnerability (CVE-2021-44228) was widely publicised. The vulnerability allows unauthenticated remote code execution, which can be triggered when a specially crafted string, provided by the attacker through a variety of different input vectors, is parsed and processed by the Log4j 2 vulnerable component.

The Web Check team is exploring whether there are any further improvements that we can use via more enhanced scans. We are working to align Web Check with the most common vulnerabilities detected each year, and, where coverage provided, whether there was an opportunity to run scans that are more in depth to look for a wider range of vulnerabilities.

The Web Check team is also conducting a review of commercial market research to investigate enhancements to be able to react more quickly to new vulnerabilities and aid our ability to implement new scans which will contribute to the future vision for Web Check.

Increasing Web Checks' user base

Web Check has seen an increase in the number of users, from ~3,200 in 2020 to ~3,700 in 2021. This has been achieved through further take up in sectors already served by Web Check and by broadening to additional sectors.

The Government Security Centre for Cyber (Cyber GSeC) has been engaging with ministerial and non-ministerial departments, and has achieved 100% coverage, with all departments that host external websites being directly protected by Web Check. Cyber GSeC engagement has increased the number of government arm's length bodies protected through Web Check, from 151 to 173, and has established that a further 36 organisations do not require it because they have no website.

Throughout the year, Web Check has participated in a number of initiatives to increase the uptake of Web Check. These initiatives include a pilot to charities, an extension of the invitation to the academia sector, and a pilot to town and parish councils. Since these initiatives began, around 250 charities, 288 academic organisations, and 23 town and parish councils have signed up to use Web Check.

Migration onto MyNCSC

Web Check, alongside Mail Check, is migrating all of its users and assets onto the MyNCSC platform. So far, specific sectors of users have been invited to migrate, with the majority of users planned to be invited throughout 2022. The Web Check team spent time developing migration aids to provide clear guidance to users on how to migrate themselves and their assets. We are also using feedback from migrated users to help shape the roadmap of Web Check features and functionality not yet present within MyNCSC that will require integration.

Outcomes

The types of findings generated by Web Check are categorised by their severity, with 'Urgent' being the most significant level. The primary goal of Web Check is for customer organisations to act 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. We have assumed that Web Check was instrumental in prompting the customer organisations to resolve them.

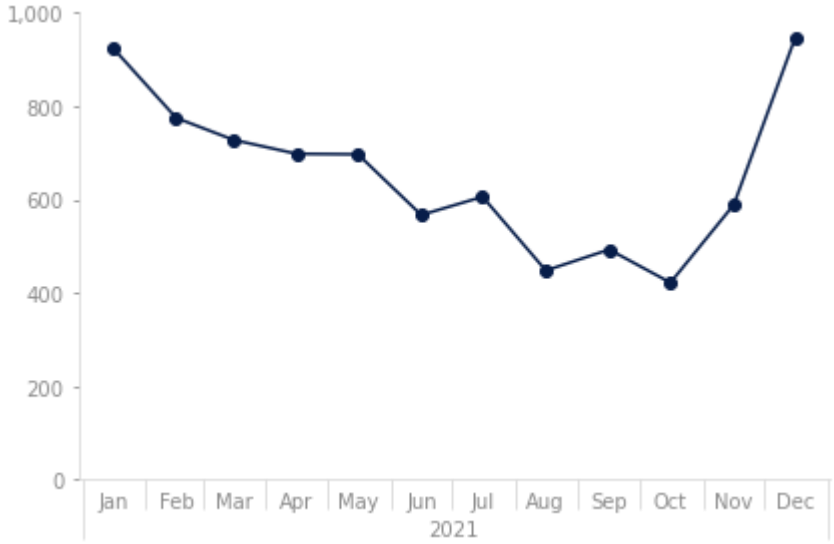

Figure 21 Urgent Web Check Issues Resolved

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. We saw an increase towards the end of the year of Urgent findings, two findings in particular showing a large increase. Firstly, the 'CMS out of support' finding. Throughout Oct, Nov and Dec, major version releases were made to products such as WordPress and Drupal, and support to previous versions were dropped.

Secondly, the 'x509 certificates expire soon' finding. After some analysis into why we had a spike of this finding, we concluded that many users renew certificates for multiple domains at the similar times, rather than spread throughout the year. Also, as SSL certificates are generated for domains (not for URLs), users add multiple URLs to Web Check with the same domain, of which all the URLs added will generate the 'expiring soon' finding, and its resolution at the same time.

Deduping the above finding resolutions by distinct domain names sees the number of finding resolutions roughly halving. This means that around 50% of these findings are for duplicate domains, this is an increase from 33% when the data was previously analysed in May 2021. In alignment with the increased number of Web Check users, we have seen a correlating number of unique URLs being scanned.

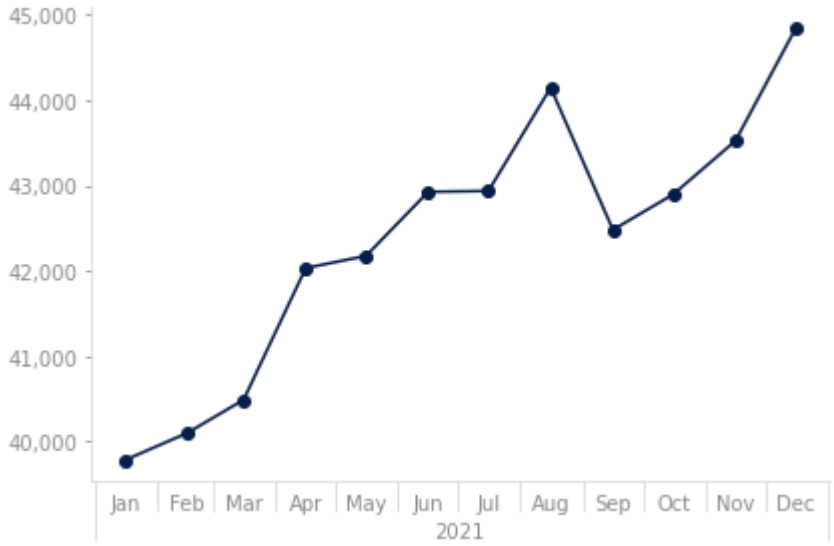

Figure 22 Number of unique URLs scanned

Protective DNS

www.ncsc.gov.uk/information/pdns

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 'example.com' (a domain) is located (its IP address) so it can connect to it. Anyone can register a domain so that everyone else can find the IP address associated with it, to enable them to connect to it.

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

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 in 2021

Five years ago, we began a project called Public Sector DNS with the aim of providing a protective DNS to public sector bodies. 'Protective' DNS was an undefined concept at the time, but as our public sector DNS evolved into our modern PDNS, we defined a world-leading capability and added a new term to the cyber security lexicon.

At its core, PDNS continues to fulfil its original brief by blocking access to known bad domains using a block list derived from a combination of commercial, open source and NCSC threat feeds, and notifying system owners so they can perform remediation. For the tens of thousands of hours since PDNS launched, it has achieved greater than 99.999% availability and has handled over 1 trillion DNS requests. From the beginning we recognised the value in analysing PDNS data to find security issues, and in late 2020 we had bittersweet vindication through our key contribution to responding to the SolarWinds incident (as described in last year's report).

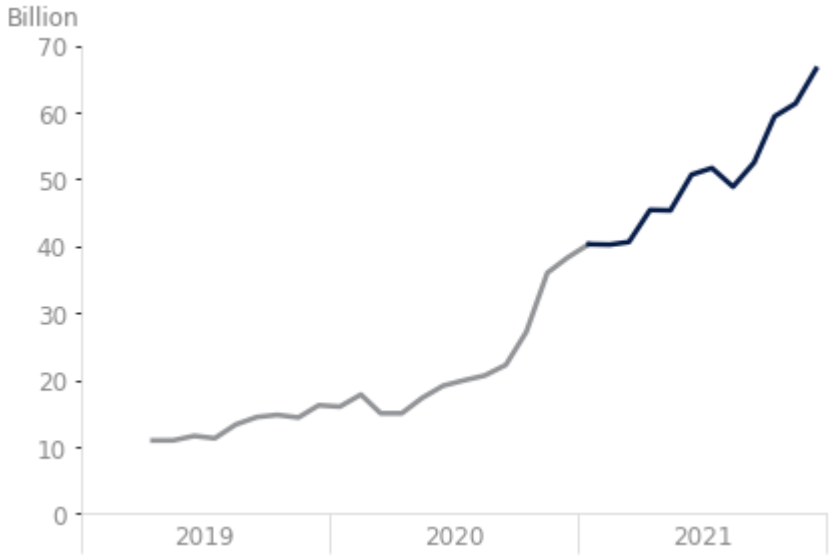

Figure 23 DNS Queries Resolved

The service has grown a lot over the last five years and the threats we face have changed. The user base now covers significant portions of central government, local government, law enforcement and the health sector. Back in 2017 we were often blocking DNS requests to domains associated with Wannacry, BadRabbit, Ramnit and Conficker, whereas today we are more likely to block domains associated with Ryuk, Conti, Flubot, Monerominer and still Conficker.

We owe much of the success of PDNS to our users, who have provided a steady stream of suggestions and feature requests through workshops, webinars and surveys.

In 2021 we were nominated for two national awards. We were honoured to be finalists for:

- Best Public Sector Project at the National Technology Awards
- Security, Defence or Law Enforcement IT Project of the Year at the UK IT Industry Awards

As seen in Figure 24, this year we have continued to welcome new organisations, with 146 joining the service.

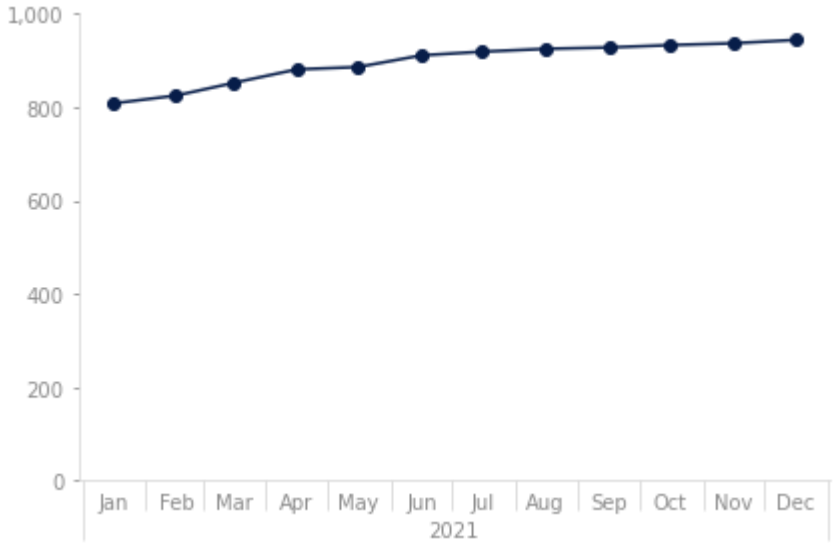

Figure 24 Number of organisations protected by PDNS

This year we continued running virtual events, with 10 webinars and 2 virtual working groups throughout the year, and they have proved more popular than our in-person events. In total, we engaged with 461 people from 259 organisations, or to put it another way, one in four PDNS organisations. From September, we included the ability to watch our webinars on demand after the event, which has proved popular, and we're now exploring the possibility of hosting hybrid events in 2022. Our events may be our most visible and interactive means of engagement, but we also use surveys, social media, and in 2021 launched both the PDNS Newsletter and the PDNS User Community.

This year we launched the PDNS User Community, an online space where our users can make suggestions for improvements to the service and share ideas. Importantly, these suggestions are viewable not just by us, but by all other PDNS users and anyone can 'upvote' an idea if they like it. We've had some great ideas, both big and small, from our users already and we read them all, not just the 'hot' ideas.

Since PDNS was launched in 2017, we have incorporated many suggestions from users into the service, such as:

- a customer portal and the dashboards it contains
- direct access to data on blocked DNS queries in STIX format
- useful customer tools (such as Block Check and a test page)
- and a solution for roaming devices (PDNS Digital Roaming)

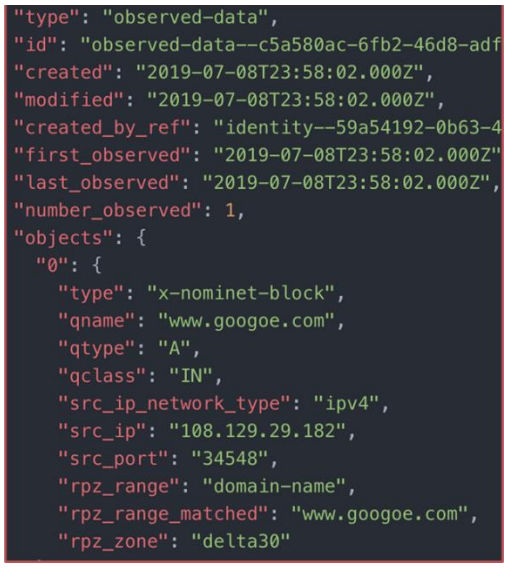

```
"type": "observed-data",
"id": "observed-data--c5a580ac-6fb2-46d8-adf
"created": "2019-07-08T23:58:02.000Z",
"modified": "2019-07-08T23:58:02.000Z",
"created_by_ref": "identity--59a54192-0b63-4
"first_observed": "2019-07-08T23:58:02.000Z"
"last_observed": "2019-07-08T23:58:02.000Z",
"number_observed": 1,
"objects": {
  "0": {
    "type": "x-nominet-block",
    "qname": "www.googoe.com",
    "qtype": "A",
    "qclass": "IN",
    "src_ip_network_type": "ipv4",
    "src_ip": "108.129.29.182",
    "src_port": "34548",
    "rpz_range": "domain-name",
    "rpz_range_matched": "www.googoe.com",
    "rpz_zone": "delta30"
```


Figure 25 User community: Block data

We launched PDNS Digital Roaming in September 2020 as an app for Windows 10 that routes a client's DNS traffic to PDNS when they are not connected through a traditional enterprise network. It is secure, encrypted through DNS-over-HTTPS (DoH), and deployable at massive scale through common mobile device management solutions. It has the added benefit of recording blocks with machine-level resolution, to aid incident handling when searching for the source of any malware found.

Thanks to the 72 organisations who started using PDNS Digital Roaming in 2021, it now protects around 23,000 devices. It has enabled us to work with cloud security vendors to extend PDNS protection further than ever before.

Customer feedback on PDNS Digital Roaming has been very constructive and, in some respects, we've been ahead of industry standards, especially in authenticated DoH. However, this year saw announcements from Microsoft, Apple and Google on native DoH support in their operating systems, which we're actively investigating to improve our product.

Outcomes

In 2021, PDNS handled more than 602 billion DNS requests. Of these, over 160 million requests were blocked. Our sources tell us that the most common reason for blocking a request was because the domain was associated with Flubot, Conficker, Monerominer, Ryuk, Doubleback, CobaltStrike Beacon, Conti, or another malware's command and control (C2).

Flubot

Flubot is malware that infects Android phones and devices. It is installed when a victim receives a text message asking them to install a tracking app due to a 'missed package delivery'. As the NCSC guidance points out, the tracking app is in fact spyware that steals passwords and other sensitive data. It will access contact details and send out additional text messages – further spreading the problem.

Flubot uses a domain generation algorithm (DGA) for C2 infrastructure, leading to infected devices making large number of DNS queries to candidate C2 domains, the majority of which are unregistered and don't actually exist. When it runs, the Flubot DGA creates 25,000 candidate C2 domains, based on the current year and month (older versions of Flubot used smaller candidate C2 lists of either 2,000 or 5,000) using the top-level domains (TLD) .com, .ru and .cn. This means that Flubot can be associated with DNS queries to a fixed set of candidate C2 domains, which remain static within each calendar month.

The research team at Nominet analysed queries to known Flubot domains from 1st May to 31st December. All queries were blocked. The diagram below shows the number of unique domains blocked per week across the reporting period.

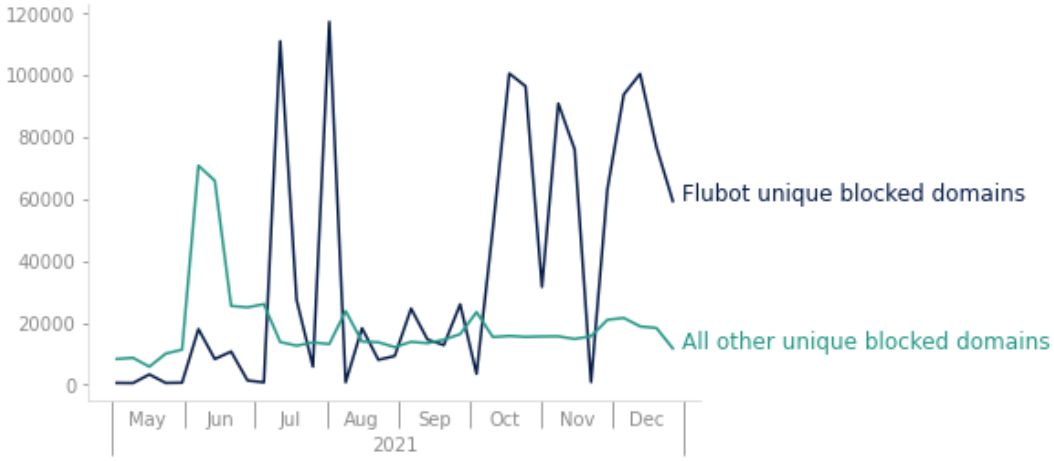

Figure 26 'Flubot unique' and 'all other unique' blocked domains for 2021

It is interesting to note that we do not see a constant high level of Flubot DNS queries, but rather intermittent periods of high volume. This intermittent pattern, together with the fact that this malware is known to target Android mobile devices, suggests that the spikes in Flubot DNS queries occur when Android devices connect to PDNS customer Wi-Fi networks. As the Flubot DGA cycles through a list of up to 25,000 candidate C2 domains monthly, it is possible that large numbers of the Flubot DNS queries we block may originate from a single, or small number of, Android devices connected to customer Wi-Fi networks for short periods of time.

Based on the trends seen in 2021, we expect to continue to block high levels of Flubot DNS queries over the short-to-medium term.

Malicious newly observed domain detection

We constantly observe DNS queries to domains that we have never seen before. Usually, this is because they have only just been registered, and while not all new domains are malicious, we know that most malicious domains are newly registered. These are known as Newly Observed Domains (NODs).

Earlier this year, our research team at Nominet analysed NODs seen in PDNS with the aim of defining procedures, methods and algorithms that could lead to automatic malicious NOD detection. Using a combination of PDNS data, security feed data, Whois data, zone file crawling, certificate transparency logs and reverse DNS lookup data, the team created several heuristic models and assessed their performance over a period of one month. All models were built with the human-in-the-loop approach.

One of our heuristic models looks for domains we call 'bad by association'. These are NODs that have near-identical characteristics to known malicious domains, and almost certainly should also be blocked. The results from our trials show that this method provides strong evidence for maliciousness in the NODs it identifies and is particularly good for highlighting domains that are part of ongoing campaigns.

When we focused on top level domains (TLDs), we developed a method of highlighting groups of NODs belonging to a single TLD for further investigation by analysts. In one trial, this method flagged three NODs with identical redirect behaviour, in this case a three-domain redirect chain. The domains:

- were all on the .info TLD
- were all 6-8 seemingly random characters long
- were all registered on the same day, with the same registrar and nameservers

When our analysts investigated, we found the IPs hosting these domains also hosted many other domains following the exact same pattern of redirects, with the exact same registration details. This resulted in a much larger investigation which found that each redirect chain ultimately resolved to a Chinese site hosting adult content and likely adware.

While investigating the maliciousness of the domains, we spotted that through the redirect chain there were many outbound connections from the browser to domains and IPs widely linked to Chinese adware, including some associated with the Chinese media company Sohu, which has been linked to spyware and adware in the past. This led us to categorise these NODs as a high risk and due to the extremely low likelihood of a legitimate need for PDNS users to connect to them a total of 307 domains were blocked.

Not all malicious uses of NODs rely on a single TLD, so we also analysed domain redirects between multiple TLDs. Using this method, we identified a cluster of NODs registered under a variety of TLDs including .club, .top, .com, .xyz and .info and flagged them for further analysis. Subsequent investigation by an analyst confirmed that the domains were malicious and they were blocked.

To investigate whether we could use heuristics to improve our protection against phishing we looked for a set of generic keywords in NODs. This analysis identified a number of domains almost certainly being used for phishing and flagged them to analysts as candidates to block. As an extension, we explored the detection of NODs imitating GOV.UK by using a specific set of keywords; however, improvements to this model are needed as most of the domains flagged for investigation by analysts were found to be legitimate.

The methods described above were complemented by analysis of web-crawl data, either HTML content or images gathered from websites. A keywords script was used to filter out noise, and then our analysts investigated the domains remaining each day.

At the end of the first month, 79.8% of domains analysed by a human after filtering algorithms were applied were found to be malicious (and subsequently blocked).

The above work confirmed the benefits of human-in-the-loop heuristic models for augmenting human analysis. For that reason, we are building a threat analyst tool (a reactive dashboard that serves user-friendly information about domains flagged by various algorithms for a quick assessment by threat analysts), and we continue to proactively analyse domains with the above methods.

Exercise in a Box

www.ncsc.gov.uk/information/exercise-in-a-box

About the service

Exercise in a Box (EiaB) 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, which includes links to relevant NCSC advice and guidance.

Primarily aimed at the non-technical audience within both the public sector and SMEs, 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 in 2021

We concentrated on new content this year with a number of exercises added to the existing suite comprising both Micro-Exercises, small bite-sized 15–20-minute learning experiences, as well as building upon the core Table Top Exercises (TTX). We added in the following exercises drawing on NCSC guidance:

- Passwords Micro-Exercise
- Securing Cloud Productivity Suites Micro-Exercise
- Supply Chain Attacks TTX (like last year's 'Home and Remote Working' exercise, this was designed specifically around 'real world events', in this case, the SolarWinds incident)

We also reacted to the recent Home Office focus on ransomware by reviewing and strengthening the content in our existing exercise (Ransomware delivered by a Phishing Email), as well as starting work on a new more technical Ransomware exercise which we expect to be delivered very early in 2022.

As well as new content, to support Singapore, who have signed an MoU with NCSC to run EiaB on their own infrastructure, we completed the packaging up of the code for Singapore and worked with them during some Technical Exchange meetings to help them become familiar with it. Having announced their partnership with us at their flagship Cyber Week event, we are eagerly awaiting their launch. We have also developed a 'stand alone' version which will run in an environment without an internet connection at the request of another partner who wished to evaluate the tool before committing to an MoU. And finally, we're working towards a GDS Live Service Review.

Outcomes

This year, we reached a milestone with EiaB with over 10,000 users worldwide and have now, at the end of December, reached over 13,000 users across the world, a circa 50% increase over the start of the year boosted by the social media campaigns we ran as we added exercises to the tool. We saw increases in our public sector audience by around 49%, SMEs by 45%, large businesses by around 84% and cyber security professionals by around 63% over the December 2020 numbers.

We have engaged with FCDO's capacity building programme, particularly in Africa. Other countries have also been in touch with a view to understanding how using the tool, or something similar, could be of benefit to their preparation for cyber incidents.

Within the tool itself, we have changed the way we ask for feedback. Below are two examples which show how organisations have responded. The first gauges their satisfaction, which is important because we need them to engage with the tool and value its outcomes, so they want to change behaviour. The second captures views on the value of individual exercises and whether they intend to make changes as a result of running the exercises. Note that some exercises were released part way through the lifetime of the tool.

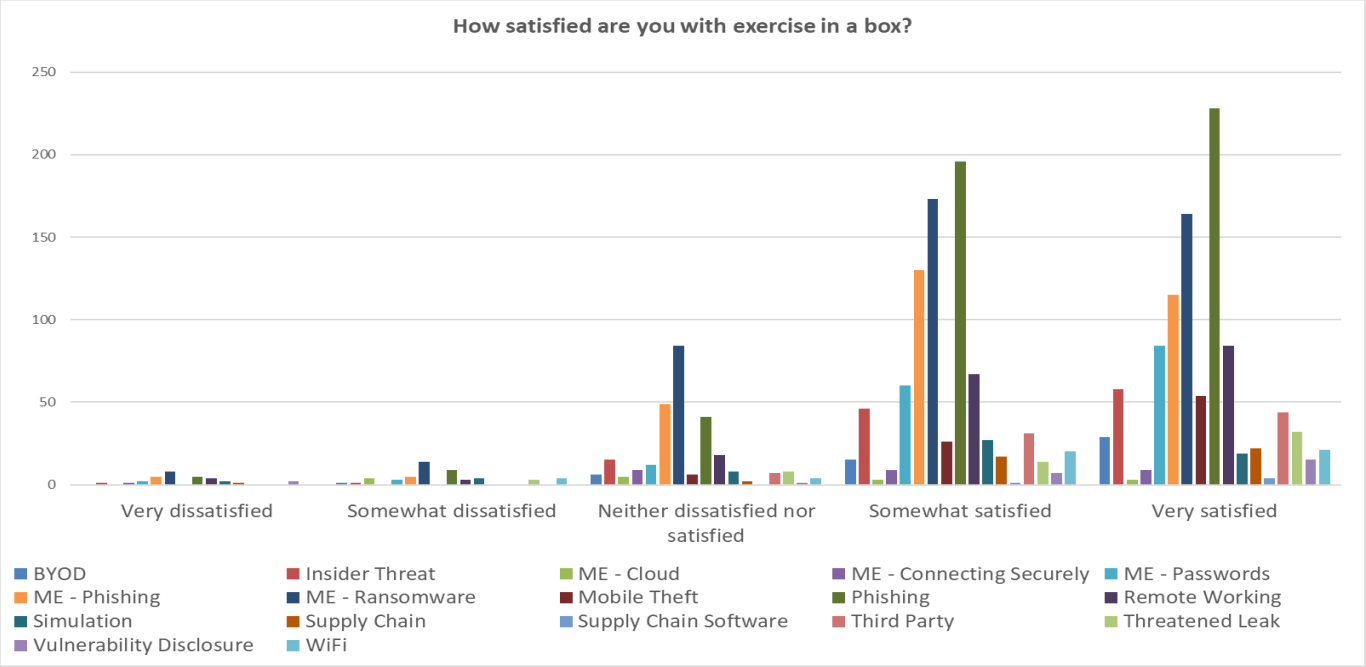

Figure 27 Customer satisfaction with EiaB by exercise.

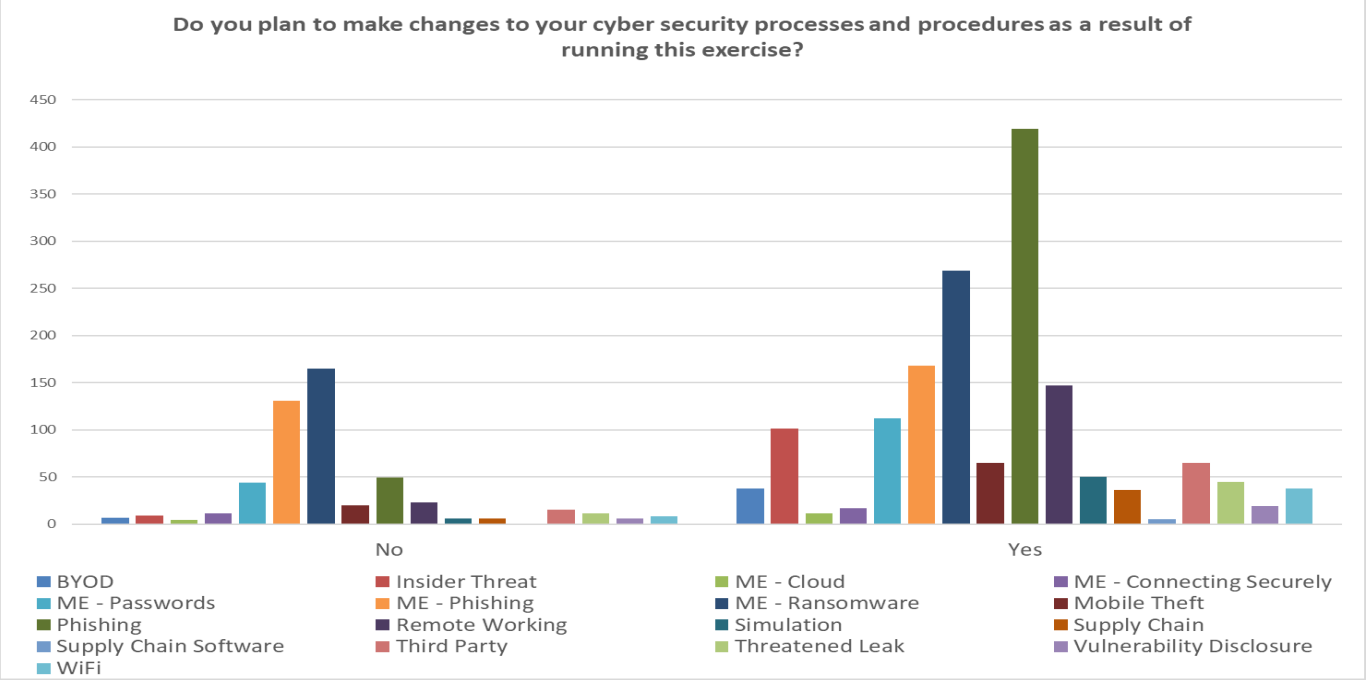

Figure 28 Intent to make changes as a result of running EIAB exercises

Early Warning

www.ncsc.gov.uk/information/early-warning-service

About the service

Early Warning is a free NCSC service designed to inform an organisation of potential cyber attacks on their network, as soon as possible. The service uses a variety of information feeds from the NCSC, trusted public, commercial and closed sources, which includes several privileged feeds which are not available elsewhere.

Early Warning filters millions of events that the NCSC receives every day and, using the IP and domain names provided by our users, correlates those which are relevant to their organisation into daily notifications for their nominated contacts.

Organisations will receive the following high-level types of alerts:

- Incident Notifications - Activity that suggests an active compromise of their system. For example, a host on their network has most likely been infected with a strain of malware.
- Network Abuse Events - May be indicators that your assets have been associated with malicious or undesirable activity. For example, a client on their network has been detected scanning the internet.
- Vulnerability and Open Port Alerts - Indications of vulnerable services running on their network, or potentially undesired applications are exposed to the internet. For example, they have a vulnerable application, or they have an exposed Elasticsearch service.

Early Warning does not conduct any active scanning of networks itself. However, some of the feeds may use scan derived data, for example from commercial feeds.

Outcomes

The Early Warning customer base grew from approximately 2,000 signed up organisations to 4,750 over the year, following the 2021 CyberUK launch.

- The service sent out 92,260 notifications to 4,610 organisations. This covered 36,027,426 events of which 7,496,610 events were about potentially malicious activity that the customer would have needed to fix (such as malware infections or other indications of a system having been hacked). This is around 20,000 potentially malicious events per day.
- 28,528,582 were about potential misconfigurations, including 3,126,220 notifications of Remote Desktop Protocol (RDP) open to the internet (which is a common vector for ransomware).

MyNCSC

www.ncsc.gov.uk/information/myncsc

About the service

The objective of the MyNCSC platform is to bring a number of the NCSC services together into a single, coherent experience tailored to each user and the organisation they are helping to defend. The intent is for MyNCSC to eventually replace the ACD Hub as the single point of access to ACD services.

Progress in 2021

2021 saw the MyNCSC focus shifting from continuation of initial user testing to migration of the first customer organisations onto the platform.

User testing during the early part of the year featured the first pilot release recognisable as a platform meeting the MyNCSC vision of:

- providing a unified user interface for accessing multiple ACD services, with the need to perform only once common activities needed to enable use of those services
- improved collaborative working for colleagues within the same organisation

Key functionality offered by this release included the abilities for a user to:

- join their organisation on MyNCSC (as either Org Admin or a member)
- manage the membership of their organisation (as an Org Admin) and view their organisation's membership (as any user)
- create and manage their organisation's shared asset portfolio (of domains and URLs)
- subscribe assets from the shared asset portfolio for checking by Web Check and/or Mail Check
- demonstrate the organisation's control over its assets through asset verification (as required for certain Mail Check functions)
- view the findings returned by Web Check and Mail Check
- receive notifications associated with findings and other system events

Positive results from the user testing contributed to MyNCSC passing a mid-year assessment gate to move from Alpha to Private Beta. With this success preparations were made to invite customer organisations to migrate their use of Web Check and Mail Check from the pre-MyNCSC to MyNCSC product versions. At the same time, development work remained ongoing. For example, whilst MyNCSC's collaborative working functionality at the organisation level is appropriate for customer organisations with fewer service users and fewer assets to check, we are mindful that other organisations will need functionality to assign users and assets to MyNCSC Teams within their organisation, to better reflect the responsibilities of different user teams for managing different areas of the organisation's IT estate (and thereby give those users an uncluttered view of the findings of interest to them).

The migration plan accordingly divided the Web Check and Mail Check customer organisations into tranches. All users from a given organisation were included in the same tranche, to enable them to experience the collaborative working opportunities from the outset. The incremental approach enabled us to better assure the viability of the platform through ongoing user research and performance monitoring. And we assigned to the earlier tranches those customer organisations well suited to the initial Private Beta functionality, leaving until later those expected to need additional functionality (such as MyNCSC Teams). When invited to migrate, users were given access to migration aids giving automated support for migrating the domains and URLs (checked by Web Check and Mail Check) to their organisation's MyNCSC asset portfolios.

Web Check and Mail Check are just the first two ACD services to be integrated with MyNCSC. Alongside the activity orientated around those two services, preparatory work was performed for the future integration of others. Some of this was specific to the individual ACD services, but opportunity was also taken to progress cross-cutting themes. For example, the current method of presenting findings will not scale well for use with a large number of services, so a model has been developed for presenting the findings in classes relating to different classes of ACD services.

Outcomes

In the middle of the year MyNCSC moved from Alpha to Private Beta.

By the end of the year approximately 340 of 1,500 customer organisations had been invited to migrate their use of Web Check and Mail Check from the pre-MyNCSC to MyNCSC product versions.

Further development is underway to move MyNCSC to Public Beta, suitable for offering to all Web Check and Mail Check customer organisations, and in preparation for integration of further ACD services.

Subdomain Takeover Alerts and Reporting (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 vulnerable DNS record pointing to a new resource. If an attacker targeting a domain discovers a vulnerable 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.

Progress in 2021

In 2021, as part of reinvigorating the existing solution, the product was renamed from Dangling DNS to STAR - Subdomain Takeover Alerts and Reporting. This marked a shift away from tackling the problem, to forming the solution.

We tested, improved and broadened the solution to allow us to draw on the successes of our previous solution whilst learning how to scale the capability to detect more vulnerabilities. We aimed to sustain the ability to identify these vulnerabilities on a national or even global scale, whilst maintaining reliability and cost-effectiveness.

We were able to expand our capability to detect vulnerable subdomains hosted by a larger set of cloud service providers. We looked at three of the largest providers - AWS, Azure, and Google Cloud - and a range of their services such as Azure Web Apps and Azure Traffic Manager, and also included more specialised providers, such as GitHub, Surge Cloud, UserVoice, WordPress, and Bitbucket. We determined the relative prevalence of subdomains vulnerable to takeover in each service and prioritised building our detection capability accordingly.

As we said last year, it is important to recognise that this is not a reflection on the security posture of these service providers, but simply shows that it is common for their customers to configure subdomains to point to resources they host. Whilst it is true that the service providers play a role in DNS vulnerabilities by allowing anyone (including malicious actors) to register resources with arbitrary names, responsibility lies ultimately with the owner of the record.

For example, the owner of dft.gov.uk may have accidentally configured charts.dft.gov.uk to point to a resource that has not yet been (or is no longer) registered. In November 2021, we saw the unfortunate consequences of this when that part of the Department for Transport website displayed explicit adult material instead of traffic statistics due to the subdomain being hijacked.

Putting the scale of this problem into perspective, over 25% of all reports to our Vulnerability Reporting Service in 2021 were subdomains with vulnerable DNS records, making it the second most frequently reported type of vulnerability.

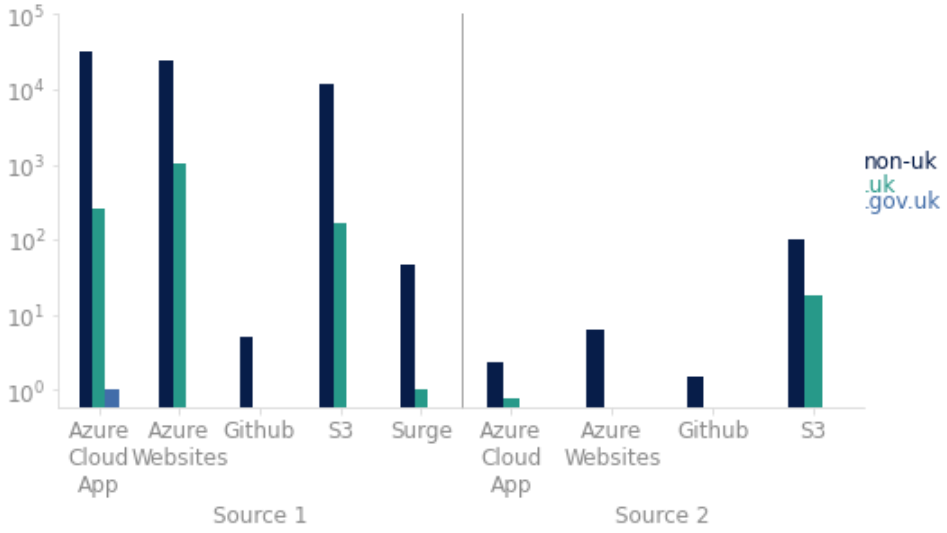

Figure 29 Vulnerabilities detected by STAR across December 2021

In December 2021, STAR detected 69,100 vulnerable subdomains within the datasets we routinely scan. To make best use of capabilities and intelligence already found within the NCSC, one of STAR's data sources (Source 2 in Figure 29) is the NCSC's Protective DNS (PDNS). The PDNS dataset contains a large number of domains and presented the opportunity for STAR to regularly scan the data for vulnerable domains. This has enabled us to find significantly more domains vulnerable to takeover, whilst providing us with a sizeable data source that can sufficiently test the scalability of the system, allowing it to be optimised accordingly.

Reporting the capability

Whilst STAR has matured to develop and test the capability to find vulnerable subdomains, it is important to be able to communicate them to the people responsible for their maintenance. As mentioned previously, the responsibility for the vulnerability primarily lies with the owner of the DNS record.

We have developed a reporting mechanism that can be shared with internal customers and aim to, in future, take domains from a customer and present to them a high-level summary of their domain vulnerability status.

In July 2021, we improved our ability to analyse ad hoc domain datasets by focusing on incorporating custom data sources, on demand. Specifically, this is benefiting NCSC teams, who can provide us with lists of domains on an ad hoc basis and be informed of any domains vulnerable to takeover.

Detection for external sources

The development of STAR initially focused on the NCSC's own domain estate, and then went on to provide the detection capability for other government departments and beyond. Therefore, as part of our continued effort to test and broaden the solution, we began to ingest data sources external to the NCSC to discover more vulnerabilities.

We used a source of UK domains to build a comprehensive dataset to scan. The benefit of this exercise was threefold: testing the solution; finding more vulnerabilities; and using our findings to improve the capability.

The figure below shows the number of UK subdomain takeover vulnerabilities we identified from a two-week period within this dataset.

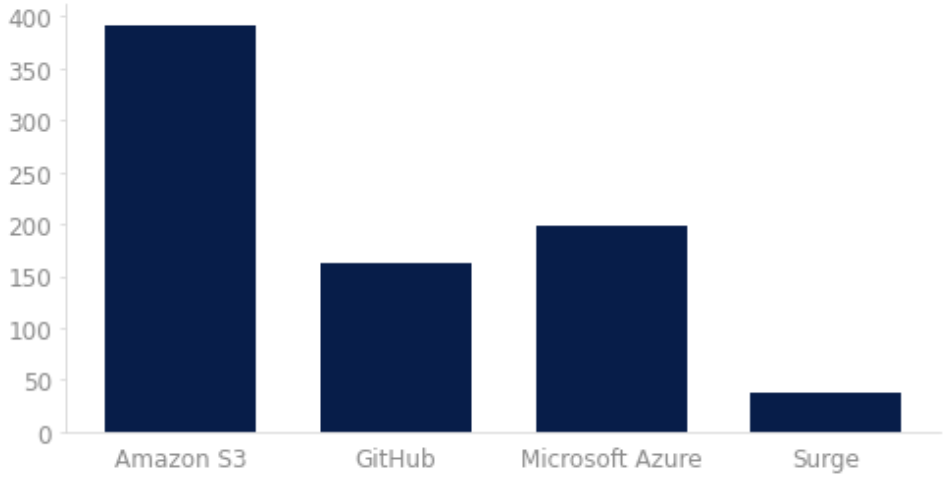

Figure 30 Two-week distribution of vulnerable domains across UK domain source

Outcomes

At this stage, STAR has been tested as a functional solution that is scalable and reliable. The next steps are for us to use the reporting mechanisms we have, and to develop new ones, to bring the vulnerabilities we find directly to the attention of the domain owners.

In late 2021, we began exploring additional methods to communicate vulnerable domains to users. The investigation began into scoping and exploring a platform that can readily present vulnerabilities to users.

Additionally, STAR continues to find new sources of domains, internal and external to the NCSC, to expand our detection capability and to ensure that our goal of reducing subdomain vulnerabilities is met.

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 and the Signalling System No.7. We have established the SMS SenderID Protective Registry, to help organisations protect their brand from use in SMS phishing attacks. In partnership with Mobile Network Operators, the NCSC have set out to create a National Telecoms Signal Monitoring Service to understand the threats to UK networks, inform defences and support incident investigation.

Border Gateway Protocol

About the service

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 fewer ASs, and has little authentication or integrity. Therefore, it is easy for any participant in the protocol to accidentally or maliciously reroute large swathes of internet traffic. There are cryptographic extensions to BGP that try to solve part of this problem. Unfortunately, the cost of implementation is high. 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. The NCSC has produced a best practice guide that covers responsible use of BGP for ISP interworking.

BGP monitoring

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

Progress in 2021

In the last twelve months, the BGP Spotlight platform has moved from Beta to a fully live environment, with a separate dedicated development environment. In its development stage, BGP Spotlight took in update feeds via an API with BT's UK, European, and rest-of-world networks, in order to collect different views of the internet routing table. Now, updates are collected via BGP peering; this is where two routers create a BGP connection in order to exchange information. BGP Spotlight has peerings with 4 other operators in addition to BT. These supplement the original BT and publicly available RouteViews and RIS-Live feeds. The RouteViews and RIS-Live collector locations used by the tool are shown in the figure below (we have not included peering details in this map). 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.

Over the last 12 months, there have been a number of changes to the Spotlight system, some to improve the presentation and detail of data provided to the users, and others to take advantage of newer technology to enhance the platforms performance.

The main database holding the BGP routing data has been migrated from MySQL to Aurora. This migration has improved performance and scalability, allowing us to capture and enhance additional data (as the database runs faster and is therefore able to process more data quicker). This enhanced data also makes it easier to identify activity on routes as well as ownership of ASNs and IP addresses. In turn this has given additional certainty to the identification of hijacked routes, and a reduction in false positives.

The system has 21 public data feeds, as well as three directly from UK telco BGP platforms. The system now automatically monitors for failed (dead) BGP routing feeds, and when one is identified, this is automatically restarted. This has improved the reliability and stability of the system.

The BGP Spotlight GUI has been significantly enhanced over the past year. We have upgraded the way that users specify Alert Rules so that it now allows users to create combined rules based on ASN, Prefix and Path. We have modified the Organisation and User Alerts pages to show the new 'Alert Rules'.

The Event, Episode and Notification searches have all been modified to cope with large volumes along with the introduction of a histogram to help users to navigate their results.

We significantly upgraded the Traceroute options and results so that the paths can now display either IP addresses or ASNs, there is information on the hops within the Traceroute, and we have added a Manual Traceroute option for users to trigger their own traceroute "now" for any of their alerts.

We have been working with Bristol University to apply Data Analytics and Machine Learning to the data that we have to try to identify more false positives, and predict routes that look likely to be hijacked.

Figure 31 Locations of RouteView (green) and RIS-Live (blue) Collectors used by BGP Spotlight

Outcomes

One recent example of a BGP Spotlight 'win' was when the tool allowed us to identify and resolve the hijack of traffic to the government PDNS service by a subsidiary of a Telco in India. We did this before Nominet, the provider spotted the hijack, and were able to advise them on exactly what had happened and when. Subsequent to the first hijack, Nominet spotted another one using BGP Spotlight in the same timeframe we did.

We have an additional 13 organisations now using BGP Spotlight, bringing the total of active organisations to 36, and increased the user base to 220 unique users. These are not just UK operators, but international operators and partners as well.

Signalling System No.7

About the service

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 2020, we moved to a more automated approach. We made a testing service available to the networks so they could run tests periodically, or when there are significant security improvements to their signalling networks, as well as facilitating the implementation of new equipment to mitigate some of these threats.

Progress in 2021

Operators have been slow to implement their mitigation strategies. As a result, the uptake of the tool has not been as successful as we had hoped. Benchmarking has shown little progress over the last year although some progress has been made since 2018.

Outcomes

It's clear that UK networks are poorly secured but there is little incentive or appetite for operators to improve. The Telecoms Security Act, however, will raise the bar and provide something to which we can hold them to account in the future. As a result, funding from this piece of work has been withdrawn.

SMS SenderID Protective Registry

https://mobileecosystemforum.com/sms-senderid-protection-registry/

About the service

The NCSC, along with UK Finance and others, has part-funded an initiative to set up an SMS SenderID Protective Registry. This allows brand owners to register authorised SenderIDs/alpha tags, define their SMS delivery chains (that is, the SMS aggregators they choose to deliver their traffic), and to provide a list of unauthorised SenderIDs that they have already seen abused in 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 simple 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 in 2021

Throughout the year we have continued to work with Government Digital Service (GDS) and the NHS as part of the ongoing COVID-19 response. At the beginning of the year, the focus was around vaccines, migrating later in the year to the COVID Pass as the vaccine programme progressed. We were also able to use the registry to protect any messaging relating to the Census.

Early in the year, we saw that shortcodes were being spoofed and used to add authenticity to an attack. Merchants taking part were advised to add their shortcodes to the SenderID registry.

Parcel delivery companies became a target throughout the year. Of particular note was the Royal Mail, which became one of the most abused brands with lots of negative publicity.

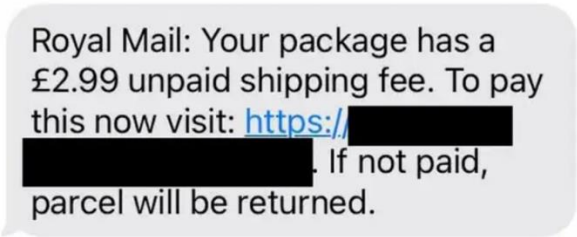

Figure 32 Royal Mail example scam messages

The Royal Mail were quickly onboarded to the Registry with the help of MEF. The NCSC were also able to help them with other protective measures to reduce their exposure.

The SenderID Registry, with its simplicity in design and speed of onboarding, has become extremely successful. Over the past year, the Registry has expanded into other territories, such as Ireland and Spain, with Singapore not far behind.

The time has now come to transition from a trial solution that now has a track record, to an industrialised solution. We continue to work with MEF on discovering ways to improve the registry, with some exciting opportunities being explored for next year.

7726

Customers using UK networks can text 7726 to report unwanted SMS messages or phone calls on a mobile. The number 7726 was chosen because it spells 'SPAM' on an alphanumeric phone keypad.

Since 2020, NCSC have been receiving reports of URLs (only) to feed into the Takedown Service.

MEF have been working with Proofpoint, who aggregate the UK public's reports, to trial the usefulness of the reported data in combatting fraud. Each merchant obtains redacted content related to their specific brand. NCSC receive this data for the government brands we are protecting in the Registry. Initial observations are:

- a lot of genuine messaging is reported as spam
- missing data reduces the usefulness (for example, telephone numbers may point to a threat actor, or the message being genuine)
- the data quality of the SenderID associated with the reported message can be of poor quality due to the manual process involved with supplying the information

However, it can be used to spot trends and areas to improve messaging. For example, we were able to identify that fraudsters were using the threat of fines to lure people into clicking on a link, which enabled us to improve the anti-fraud messaging. We were also able to identify specific government departments that had high levels of genuine messages reported as spam, and work with them to improve their communications.

Outcomes

A combination of the SenderID registry, the Suspicious Email Reporting Service (SERS) and the Takedown Service has seen a dramatic reduction in fraud aimed at government services. The DVLA and TV Licencing, for example, no longer appear in the 'most abused brands' statistics. The following table is broken down by attack type and it should be noted that nearly two thirds of those takedowns were phishing URLs and most of the remainder were cryptocurrency investment scams.

Table 18 NCSC sponsored takedowns that 7726 was first to report

Attack Type	Total Campaign Groups	Total Attack URLs	Total IP Addresses	Availability (hh:mm)	Percentage of Attacks
Phishing URL	2,888	8,875	1,175	13	66.50
Cryptocurrency Investment Scam	1,396	4,076	1,023	261	32.20
Fake Shop	45	45	29	375	1
Web Shell	8	16	8	10	0.20
Technical Support Scam	4	5	4	0:14	0.10

UK Smishing

The second half of 2020 saw 270% growth in smishing (SMS phishing) reports in the UK. Then in first half of 2021, there were over 536% reports of smishing in the UK compared to the second half of 2020. To provide context, during 2021, many months had in excess of one million combined smishing and FluBot reports.

National Telecoms Signal Monitoring Service

About the Service

Society is increasingly dependent on telecommunications. As this dependence grows, we need to improve the resilience of our networks. In partnership with Mobile Network Operators, the NCSC have set out to create a National Telecoms Signal Monitoring Service (NTSMS) to understand the threats to UK networks, inform defences and support incident investigation.

NTSMS aims to analyse signalling protocols used in mobile networks as all are susceptible to misuse and exploitation. Initially NTSMS is focusing on SS7, it being the most exploited of the mobile signalling types and the least likely to be fixed due to it being a legacy protocol developed to for 2G and 3G, however its continued widespread use leaves mobile networks and subscribers vulnerable.

Essentially the functionality in SS7 that allows subscribers to roam around the world whilst making calls, receiving and sending texts and accessing the internet also allows attackers to monitor conversations and track subscribers, often down to street level, in real time and without any indication on their device.

A common attack is to defeat second factor authentication (2FA) often used by banking and social media platforms. By exploiting the poor security in the SS7 protocol, it is possible to redirect SMS and voice calls which, instead of being delivered to the phone of the designated account holder, are diverted to a phone controlled by the attacker. The attackers then use the One Time Password (OTP) contained in the SMS or automated voice call to reset accounts and gain access.

It's not just subscribers that are vulnerable, mobile network operators are also under attack from SS7. A recent example being the Norwegian operator Telenor, where a handful of SS7 packets led to a network wide outage of around 3 hours. Operators are also susceptible to fraud with SMS and data being re-routed around charging systems and subscribers being signed up premium rate services without their knowledge which often leads to the operator waiving the customer's bill.

Progress in 2021

Working closely with a single network operator, we've jointly developed and deployed a range of analytics looking for specific Indicators of Compromise (IOC) which, when triggered, are used to evaluate and update firewall rules, preventing subsequent attacks. Cognisant that SS7 attacks continue to evolve, we've been working closely with the operator's Data Science hub to identify and characterise anomalous activity using data science techniques which will help identify new techniques trying to circumvent existing security measures.

This year we will be increasing protocol coverage to include either SIP (Voice and SMS in 4G and 5G) or Diameter (4G replacement for SS7) with the initial network operator and will be replicating the SS7 work to date with a second network operator.

Outcomes

Incidents are starting to be shared amongst the community via the Network Security Information Exchange (NSIE), Signalling subgroup and the GSMA Fraud and Security (FS) working groups, enabling operators to pre-empt attacks and pro-actively improve network security. A number of coordinated attacks across a number of UK mobile network operators have been identified and attempted geolocation tracking of UK citizens has been observed (although the attacks where not successful due to the operators' firewalls).

Host Based Capability

www.ncsc.gov.uk/information/host-based-capability

About the service

Host Based Capability (HBC) is a software agent that can be deployed on government OFFICIAL IT devices to enhance the security posture of our partners in government departments. It collects and analyses technical metadata to detect malicious activity, helping departments with their security via the three service tenets:

- detect: detecting malicious activity for departments to undertake remediation as required
- threat surface: providing security baseline reporting, informing departments of their cyber hygiene
- forewarn: notifying departments of detected exposure to the most serious of new vulnerabilities

Progress in 2021

In 2021 HBC increased its coverage by 90,000 endpoints across multiple departments, resulting in total coverage of an estimated 370,000 Central Government endpoints since being stood up. Through this coverage, HBC has continued protecting departments via:

- Detect: HBC worked on 18 incidents in 2021, a number of these having been detected by the HBC team. Working on the incidents HBC provided information that helped the departments targeted to understand the remedial action they needed to take. The HBC team also identified and notified departments of 31 Suspicious Activity Observations (SAOs); these 'irregular' detections by HBC informed departments of suspected but unconfirmed malicious activity, for them to conduct further investigation.
- Threat surface: HBC generated 281 threat surface reports (TSRs) in 2021, a cumulative total of 524 over 4 years. The reporting provides departments with information on their threat surface, as exposed by the devices running the HBC agent, contributing to monitoring and other information departments already collate to make decisions about their security posture.
- Forewarn: HBC issued an estimated 66 individual notifications on 2 major vulnerabilities identified in 2021, providing individual notifications to departments about their unique exposure to 8 new, major vulnerabilities since it launched.

While a significant focus in the past year has been development of the Service, HBC has continued to draw on its Detect and Forewarn capabilities to help protect UK government networks. Knowledge of MS Exchange and Log4shell have enabled notifications of these vulnerabilities, with HBC simultaneously building detection tradecraft into the capability to identify potential compromises utilising these vulnerabilities.

As we look to move from Beta to a Live service, ongoing work with select departments to explore improvements in information sharing as part of HBC coverage (among other areas) continues.

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 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 in 2021

Vulnerability Reporting Service (VRS)

The growth of the VRS has continued and the service received nearly four times the number of reports than the previous year. We have dedicated considerable time to improve our processes to ensure we can continue to provide the level of service both the finders and the affected system owners would expect.

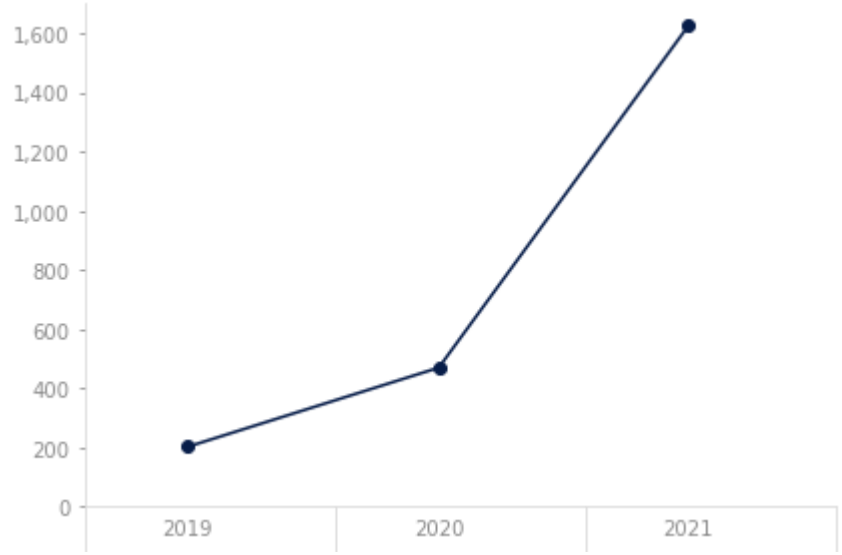

Figure 33 Annual number of vulnerability reports

We concentrated our improvements on ensuring we can contact affected system owners as quickly as possible. We have also ensured all the relevant information is included in the summary report so they can quickly remediate the vulnerability. By working closely with the system owners from across government, 79% of reported vulnerabilities are resolved within 30 days.

Through analysing the reported vulnerabilities, we found that 11% of the most commonly reported vulnerabilities were mitigated by updating to the latest version of the affected software. This highlights that keeping software up to date is a very important part of keeping systems secure.

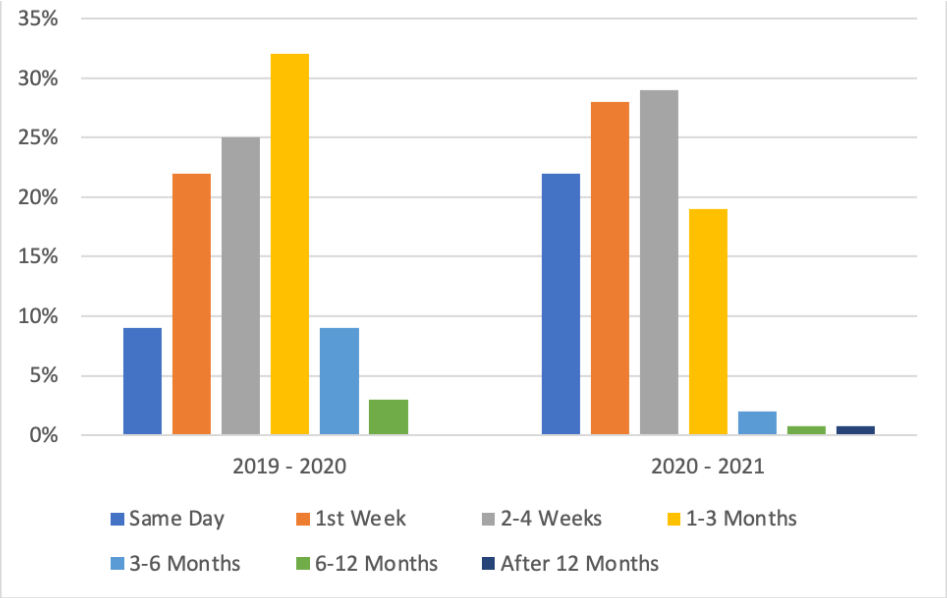

Figure 34 Resolution timeframes from system owners first contact

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. NCC Group triage all the reports and provide recommended mitigations to ensure that the vulnerabilities can be remediated as quickly as possible.

During 2021, the pilot helped 14 UK government departments launch their own Vulnerability Disclosure Programme (VDP). This brings the total number of VDPs to 22, enabling these departments to directly receive vulnerability reports so they can fix the issues before they cause harm. An additional 22 departments are following the onboarding process to join the pilot and embed the disclosure management process within their teams.

Vulnerability Disclosure Toolkit (VDT)

The NCSC VDT continues to be used by organisations to help them create a simple vulnerability disclosure process. We are always keen to hear feedback and this will go into a future update.

Outcomes

Building on the success of the NCSC VRS, the UK government will develop a coherent and joined up cross-government VRS. This will enable the mature handling of, and response to, vulnerabilities which have the potential to impact government. By providing this capability centrally, government will, for the first time, be able to holistically tackle cyber security vulnerabilities at scale and pace across the public sector.

Logging Made Easy

www.ncsc.gov.uk/information/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 in 2021

In May 2021, Version 0.4 was released onto GitHub. This included an overhaul of a number of areas of the LME installation script to make the product compatible with the Elastic Common Schema (ECS), and to fix integration with the Elastic SIEM. It also included additional functionality to upgrade existing users to the new schema and documentation updates explaining the steps needed.

Subsequent to the 0.4 release, additional version updates were also released:

- September 2021: update to the latest version of Elastic (7.13.4) including updating mapping files to the latest ECS version.
- December 2021: updated the relevant Winlogbeat install instructions and to point to Elastic (7.16.1), updated the Docker stack versions to 7.16.1, added additional instructions to the upgrade steps (if the latest mapping file is not supported by the version of Elasticsearch in use), and updated the required Elasticsearch and Kibana environment properties and added handling for the PublicBaseURL value.
- In December 2021, version updates also resolved the Log4j vulnerability, tracked as CVE-2021-44228, by updating to the latest supported version to Elastic 7.16.2.
- January 2022 a minor version bump to Elastic version 7.16.3 to resolve issues with vulnerability scanners flagging the service as exposed to the log4j vulnerability.

Outcomes

2021 has seen a steady uptake in LME which we can see has been cloned up to 1,048 times in the last 12 months an average increasing of 87 per month. This has provided organisations previously without a SIEM to have a basic logging capability; some of these organisations have consequently been able to participate in the CTI Adaptor pilot which has provided them with alerts about cyber threats. The log4j vulnerability patch also ensured that our customers have been updated to latest version requirements.

Cyber Threat Intelligence Adaptor

About the service

The Cyber Threat Intelligence Adaptor (CTI Adaptor) is a software program 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, sharing the information with both the system owner and the NCSC when an IOC is present in a customer's logs.

Progress in 2021

In the past year, API plug-ins have been developed for Logpoint, Sentinel and Splunk to add to the already compatible LME and Elastic SIEMs.

During 2021, the CTI Adaptor has been continually developed to provide enhanced features and functionality:

- CTI Adaptor version 0.3 was released in March 2021. This update included a plug-in for LogPoint and provided general bug fixes and improvements.
- Version 0.4 was released in June 2021. Improvements included CTI Adaptor plug-in for Azure Sentinel, development to enable schema-based searches (previously only free text searches were supported) for Splunk, Logpoint, LME, Elastic & Sentinel and a Canary was added for testing.
- Version 0.4.1 was launched in October 2021, which included some improvements and fixes for Logpoint and Sentinel users.

CTI Adaptor is planning to move the onboarding process onto MyNCSC. Discussions began in 2021; considerable progress has been made in creating the front-end and work with supporting services. Work is still ongoing but we should launch early in the next financial year, which will coincide with the release of CTI Adaptor version 0.5.

The development of version 0.5 continues into 2022. On release, this version will include new features such as the development of an intelligent search feature which enables the CTI Adaptor to support significantly larger threat intelligence feeds and prioritise searches based on severity and/or time (age). We have also updated support for all SIEM provider schemas and developed signature-based search queries using Sigma. The next step for release of this latest version is to finalise testing and risk acceptance.

In order to support HMG and demonstrate benefit to the government, the CTI Adaptor had a change in eligibility criteria in the autumn. The new version (0.5) will be targeted to all government organisations only, this includes central government and local authorities. We will however continue to support pilot organisations that have already engaged with our system and include public sector CNI organisations on a case-by-case basis only.

In 2021, 16 pilot organisations were onboarded to CTI Adaptor, with interest from many more. There were approximately 3,000 user sightings[3] in 2021, all of which were investigated and on occasion organisations were asked for further information to verify or provide more details. No sightings resulted in the need for further action. Further work to enhance detection rules and reduce false positives is underway to be released in the next version, early testing has shown that these changes have produced enhanced sighting detection.

[3] An IOC contained in the Cyber Threat Intelligence feed being present in the collected data of an organisation subscribing to CTI Adaptor

Outcomes

CTI Adaptor will enhance the NCSC's visibility of the threat landscape across the eligible sectors within the UK by returning sightings data to analysts. Current pilot work has seen us develop and enhance the CTI Adaptor breadth making it compatible with more SIEMs. Work to enhance the threat feed and sightings detection in the next iterations will evolve the capability to reduce the number of false positives. Additionally, integration with MyNCSC will give an easy onboarding platform to the Adaptor making it attractive to wider audiences.

The NCSC Observatory

About the service

The NCSC Observatory Service focuses on generating data-driven insights to underpin the NCSC's research and strategy and allow an effective response to incidents. This is done through its two flagship products: PULSE and DNS Insights. They provide analysis of publicly accessible data, such as DNS records of UK domains, and use this to track the uptake of DNS security protocols and the usage of different technologies and cloud providers. Additionally, the products consume and analyse 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. Our products have a united focus to analyse, track, and transform DNS data in a usable, valuable format to reduce cyber threats online and on technology devices. Our products find valuable use among analysts, researchers, policy-makers and more who want to make data-driven informed decisions fast and using a reliable, credible source.

Through its products, NCSC Observatory Service strives to ensure that DNS data it collects is displayed in an accessible, intuitive to understand, and compatible format, to incorporate with existing projects and other NCSC products.

Progress in 2021

During 2021, PULSE and DNS Insights became reliable contributors to internal research projects. Work was undertaken to refactor the systems to significantly reduce the operating costs of each and to implement an API to allow for the data sets to be queried programmatically by other systems.

We also ran a small trial to experiment with offering both products to external users to explore the potential benefits to them, this trial included Cabinet Office and HMRC. With DNS Insights, we were able to inform these departments about technology usage within their department based on interpreting their use of the PDNS service. This experiment helped us understand the potential benefits of incorporating this sort of functionality into our digital services and exposing it via MyNCSC.

Outcomes

Observatory has helped us understand the value of some of the data sets we produce as part of the ACD portfolio, and how we can use those data sets to make better informed decisions. Due to resource constraints and some of the other priorities in the portfolio, we are switching PULSE and DNS Insights into maintenance mode. Whilst the data sets curated by DNS Insights and PULSE will remain available to support internal decision making, we will be taking a break from development of new features, though we may re-start development in due course.

Conclusion

2021 was the 5th year of the NCSC's ACD 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, ranging from direct interaction with members of the public (SERS), to scanning and notification to system owners (Web Check, Mail Check), and from the detection and reporting of attacks to infrastructure providers (Takedown), to 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 terms of delivered outcomes, this five-year period has seen substantive change, from a series of largely technical experimental endeavours to a range of service development and delivery that covers the full lifecycle from concept to live service. The ACD portfolio has increased significantly due the range of target services, markets and delivering organisations, both internal and external to the NCSC, and we have placed an emphasis on proactively managing that portfolio to ensure optimal impact. It is also worth noting that the cyber security world has changed massively in that timescale:

- the sophistication and scale of the threat has increased, with yesterday's organised criminal threat being today's commodity attack
- the level of defensive capability, or the recognition of the need for it, has improved
- the capability available in the market has evolved, and the speed with which it is made available has increased

All of these developments make identifying the most appropriate capability or service to bring forward in a timely fashion a significant challenge, but we believe we have responded effectively to this, through collaboration with our partners, noting that, as with cyber security as a whole, ACD is a 'team sport'.

Since the very beginning of ACD, we have adopted a strategy of focusing on 'government and the public sector first', piloting our capabilities with such users, before considering whether to broaden to other sectors. Through the maturing of services and the approach we have adopted, we have also built in the ability to pivot effort and successfully focus ACD adoption in particular sectors as dictated by world events. In 2020, this was successfully tested in our response to the pandemic, and we have subsequently shown that our ability to be able to protect vital sectors at vital times, apply our technology and techniques in novel ways as required, is enduring. The first five years has also highlighted the benefit that ACD brings in contributing to the situational awareness picture through appropriate data collection and exploitation.

This 5th annual report documents many of the valuable contributions the ACD programme made to the cyber security of the UK in 2021. By providing data and case studies on our efforts, we hope that we can continue to demystify the challenges of large-scale cyber attacks, 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. The report also provides the opportunity to pause to reflect on the progress made in the first five years of providing ACD, and to consider the shape of the response in the coming years.