



National Cyber
Security Centre



Active Cyber Defence

The Sixth Year

Contents

Foreword	3
Takedown	4
Suspicious Email Reporting Service	11
Mail Check	13
Vulnerability Checking	17
Protective DNS	21
Exercise in a Box	27
Early Warning	30
MyNCSC	32
Routing and Signalling	34
Host Based Capability	37
Vulnerability Reporting and Disclosure	38
Logging Made Easy	40
Cyber Threat Intelligence Adaptor	41
Conclusion/forward look	42

Foreword

Now in its sixth year, this report illustrates how the NCSC's Active Cyber Defence (ACD) programme continues to make the UK measurably safer from cyber attacks.

Our rationale for producing the report has remained constant during this time; a commitment to transparency, and basing our interventions on unbiased data and evidence to better understand the reality of cyber attacks, as well as the efficacy of our products and services.

The specifics change over time, of course. Threat actors come and go, and the types of vulnerabilities being introduced and exploited continue to evolve. However, most of our ACD initiatives address enduring cyber security challenges: sharing knowledge of threats, closing down vulnerabilities, and responding to breaches. The need to tackle these challenges through automation will persist, because as things stand, that's the only realistic way of generating the scale and reach required.

For all these reasons, we see ACD as a core part of how the NCSC will improve the UK's cyber resilience over the coming years, as we continue to build services designed – as Dr Levy put it – “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.”

When ACD was launched in 2016, we developed services with the protection of government organisations specifically in mind. However, at the core of the UK's [National Cyber Strategy](#) is a 'whole of society' approach, which is why we've broadened the utility of ACD services to a wider range of users, from small business owners to the education and charity sectors. This conscious shift to designing and developing 'radically simple' digital services (with accessibility and ease of use as core design principles) will help provide the benefits of vulnerability checking to those individuals and organisations that do not have a dedicated security function.

We also want to make it simple for users to find, sign-up to and manage our services, whilst reducing duplication and providing a smoother, more integrated user experience. We built the [MyNCSC platform](#) to turn that vision into reality.

[Last year's ACD report](#) noted the challenges of developing new services, which included improvements in levels of defensive capability, the need to deliver a more dynamic commercial market, and the growing sophistication of commodity threats. This has meant embracing different ways of 'getting things done', whether that's building services ourselves, contracting with market-leading UK companies, or engaging with collaborative projects. Looking beyond ACD, we've also 'badged' the assured industry services to help users differentiate quality. We'll keep investing in proven delivery models, but stay attuned to new approaches as the consumption of IT services shifts (for example, through cloud provision).

As with previous reports, we have tried to focus on key findings and important trends. We highlight the successes but we're honest about the gaps in the evidence base that still make it hard to be definitive about impact. The underpinning message is that 'cyber security is a team sport', involving the public sector, commercial and international partners...

...which just leaves me to thank all of our partners who contribute to the success of ACD, without whom we would not be able to implement these UK-wide cyber security defences. It is great to see how far we have come over the last six years, delivering interventions – at scale – that help tackle high-volume commodity attacks that affect people's everyday lives.

As always, 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](#).

Jonathan Ellison
NCSC Director for National Resilience and Future Technology

Takedown

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

About the service

The Takedown service finds malicious sites and sends notifications to the host or owner to get them removed from the internet before significant harm can be done using them. The types of attacks we select are based on judgment of what causes the most harm to UK interests (and our progress against these attacks in 2022 is discussed below); all types of malicious activity hosted in the UK is also targeted. The NCSC manages the service centrally, so departments automatically benefit without having to sign up.

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; and
- an ‘attack group’ is how we refer to all the URLs that form part of a campaign.

Progress in 2022

The first 5 years of the Takedown service saw significant year-on-year growth in the total number of takedowns conducted. 2022 is the sixth year and for the first time we have seen a drop in the number of takedowns compared to the previous year.

Table 1: Total takedowns by campaign group and URLs

Year	Campaigns	URLs
2022	1,800,000	2,400,000
2021	2,700,000	3,100,000
2020	700,595	1,448,214

Most of the reduction in takedowns can be attributed to extortion mail servers (528,000) and cryptocurrency investment scams (459,000), whilst the frequency of other attack types has either grown or remained static.

These two attack types have some of the shortest uptimes on average, which could explain the reduction in prevalence as attackers concentrate on areas where their return on investment is greater. Mail servers and cryptocurrency investment scams have a median availability of 25.5 and 1 hour respectively, whereas the next top five attack types have a combined median of 56.29 hours.

Table 2: Total takedowns by attack campaign group

Attack Type	2021	2022
Extortion Mail Server	1,867,439	1,338,718
Cryptocurrency Investment Scam	610,621	151,343
Fake Shop	123,359	100,311
Phishing URL	54,671	56,632
Web Shell	26,326	30,312
Brute Force Attack	-	40,890
Advance Fee Fraud	21,168	3,116
Malware Infrastructure URL	5,270	18,337
Technical Support Scam	14,486	-

Attack Type	2021	2022
Web-Inject Malware URL	1,466	6,287
Advance Fee Fraud Mail Server	6,632	365
Facebook Brand Infringement	331	5,277
Malware Distribution URL	2,284	3,310
Phishing URL Mail Server	3,437	1,554
Malware Attachment Mail Server	2,580	2,294
Fake Pharmacy	884	3,367
Vulnerable Application	4,128	-
Malware C2 IP	1,902	645
Shopping Site Skimmer	962	1,540
Malware Command and Control Centre	719	745
Instagram Brand Infringement	728	71
Google Adwords	1	516
Phishing Dropsite	4	408
TikTok Brand Infringement	89	313
Clone Firm Email	352	-
Cryptocurrency Miner	138	153
Twitter Brand Infringement	206	62
Clone Firm URL	231	-
DKIM Signed Email Domain	149	44
Survey Scam	138	33
Phishkit Archive	118	29
Fraudulent Use of PayPal on Fake Shops	13	127
JavaScript Resource	29	100
Brand Infringement	49	37
Fake Mobile App	81	-
Skimmer Credential Dropsite	35	41
Advance Fee Fraud Phone Number	65	1
Other URL	65	-
Clone Firm Phone Number	45	-
Phishkit Email	22	11
WhatsApp Brand Infringement	5	28
Technical Support Scam Phone Number	32	-
Telegram Brand Infringement	9	21
Other Email	2	12
Malware URL Mail Server	-	14
Fake Bank URL	4	-
Malware Payment URL	-	3
Business Email Compromise	1	1
Other Phone Number	2	-
Fake Bond Comparison Site	1	-
LinkedIn Brand Infringement	-	1
Blocked Ownsite	-	1

Cryptocurrency investment scams

We started commencing takedowns against this attack type in 2020. Takedowns against this attack type peaked in January 2021, with a consistent downward trend into December 2022.

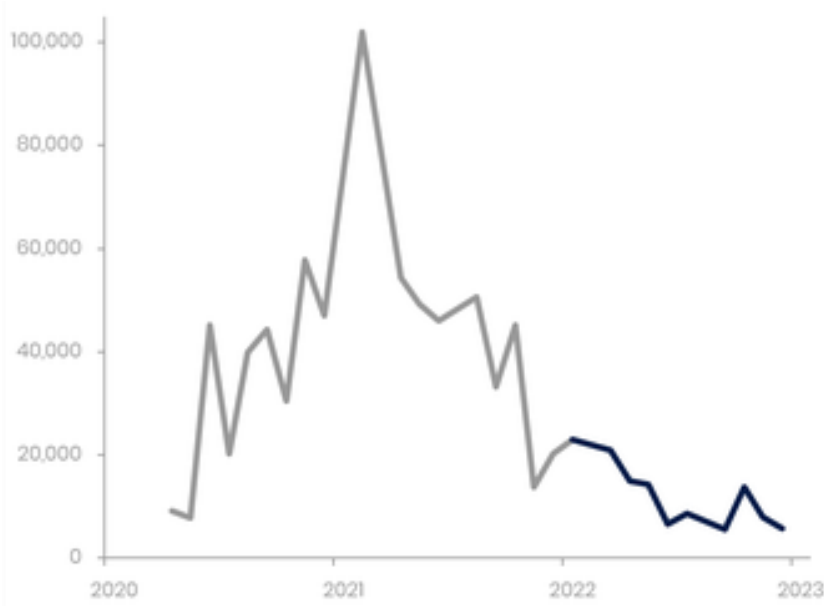


Figure 1: Number of takedowns against cryptocurrency investment scams

Despite the fall in takedowns, cryptocurrency investment scams continue to be a high-volume attack type. These attacks usually use celebrities or well-known brands to appear more legitimate.

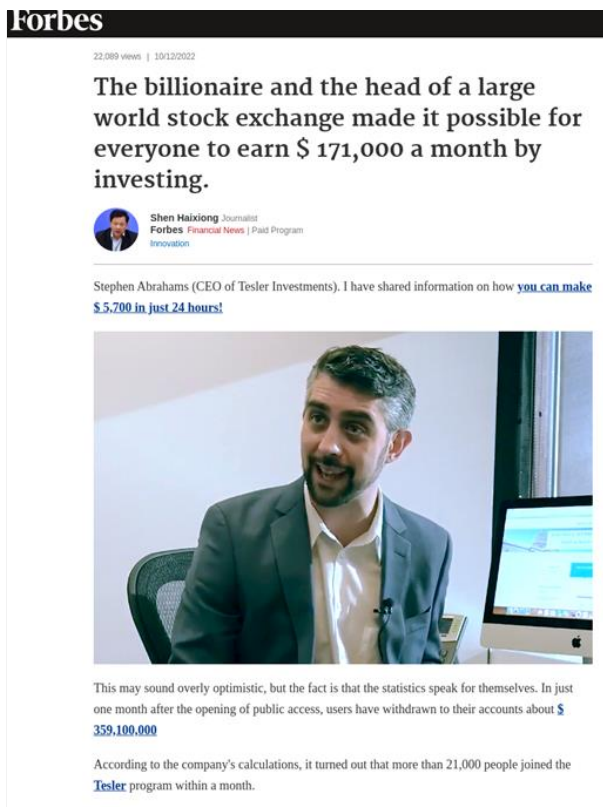


Figure 2: Example cryptocurrency scam featuring fake endorsement

Government-themed scams

In 2022, the number of takedowns against UK government-themed phishing attacks continued to reduce from its peak in early 2021. Figure 3 shows the top UK government brands used in phishing attacks, and the reduction we've seen over the last two years.




Figure 3: Top UK government brands used in phishing attacks

Energy bill scams

Scammers continue to use topical events to make phishing attempts more believable, and to target vulnerable people. In September and October, we saw an influx of phishing attempts targeting the UK government's Energy Bills Support Scheme. These URLs typically included key words such as 'rebate', 'grant' and 'scheme' in an attempt to sound like a legitimate source.

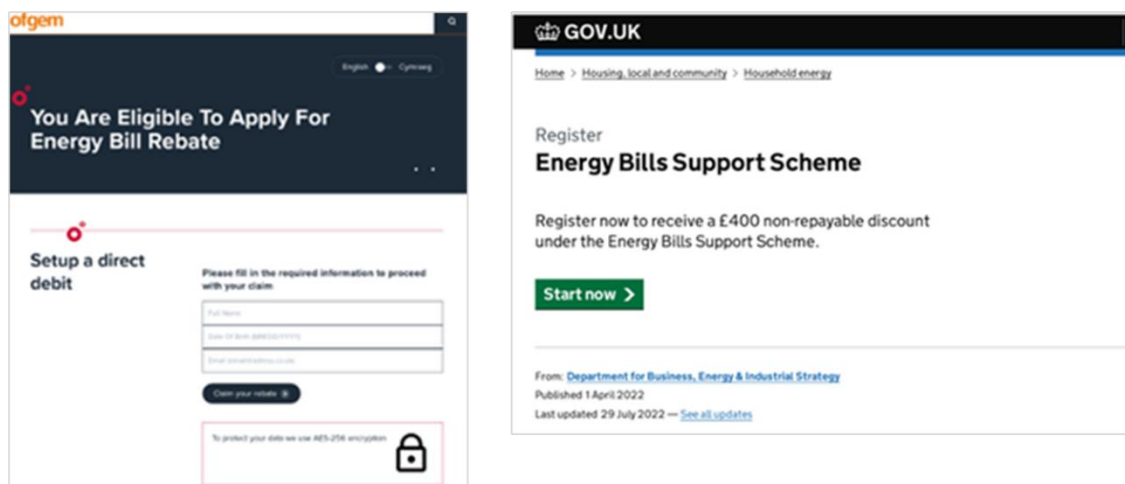


Figure 4: Phishing attempts targeting the UK government's Energy Bills Support Scheme

Web shells

Web shells are created by attackers using malicious scripts to install control panels on compromised servers. These servers can then be used as a launch pad for malicious activity such as hosting phishing sites and sending fraudulent emails.

The number of web shells we have discovered and taken action against has increased in 2022 by around 15%. The most prevalent hosting providers of web shells are listed below.

Table 3: Most prevalent hosting providers of web shells (2022)

Hoster	Takedown groups
Newfold Digital	4,666
Cloudflare	2,074
GoDaddy	1,787
NameCheap	1,266
OVH	1,242
Hostinger Group	1,135
Amazon	1,020
DigitalOcean	894
Other Hosting Providers	10,323
Totals	24,407

Brute force attacks

In August 2022, we started using honeypots to expose commonly attacked protocols to the internet in order to discover more targets for takedowns. SSH is the protocol which led to the most takedowns, with Exchange being targeted the least frequently.

Table 4: Most common protocols targeted in brute force attacks

Category	August 2022	September 2022	October 2022	November 2022	December 2022	Total
SSH Brute Force	6,477	5,110	6,169	8,643	5,831	3,2231
RDP Brute Force	946	947	1,171	1,311	1,459	5,869
WordPress Brute Force	669	500	619	540	414	2,742
Exchange Brute Force	6	10	9	4	22	51

Ukraine war cryptocurrency donation scam emails

Scammers use current events to make their attacks both more believable, and to elicit an emotional response so people are more likely to be tricked by a scam. In March, we saw attackers start to use the crisis in Ukraine to convince people to send them cryptocurrency donations. This remained a consistent type of attack throughout the rest of 2022.



Figure 5: Number of cryptocurrency scams relating to the war in Ukraine (2022)

These attacks are usually sent from compromised mail servers in large numbers. They often impersonate public figures and high-profile organisations to appear legitimate.

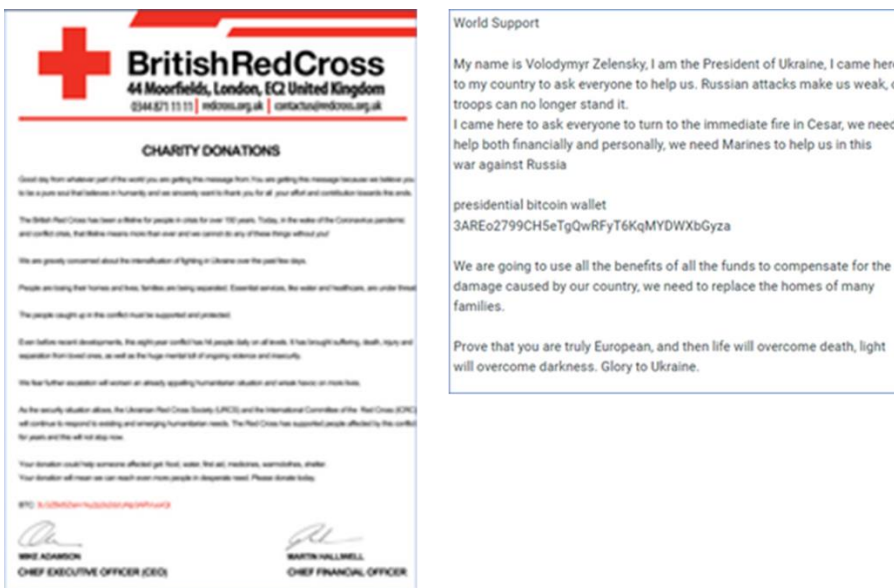


Figure 6: Example of scams exploiting the war in Ukraine

Outcomes

Takedowns in UK delegated IP Space

We continue to target any malicious activity hosted in the UK regardless of the brand targeted. Overall, the number of attacks we discovered that were hosted in the UK has decreased by over 25%. The median availability of the top three attack types has also decreased, reducing both the likelihood of people falling victims to these scams and the return on investment for the attacker.

Table 5: Takedowns by attack type in the UK delegated IP space

Attack Type	Number of attacks		Median availability (hours)	
	2021	2022	2021	2022
Phishing URL	113,457	77,471	10	7
Web Shell	12,969	9,020	47	31
Fake Shop	5,815	7,212	1,113	414
Web-Inject Malware URL	1,517	5,725	88	104
Fake Pharmacy	3,386	376	33	27
Shopping Site Skimmer	1,246	1,448	85	69
Malware Infrastructure URL	278	562	82	48
Cryptocurrency Miner	324	232	78	94
Malware Distribution URL	341	120	30	27
Phishing Dropsite	4	374	85	12
Malware C2 IP	47	15	223	200
Technical Support Scam	65	0	34	N/A
JavaScript Resource	17	41	527	2,017
Skimmer Credential Dropsite	30	15	79	545
Malware Command and Control Centre	21	14	72	21

Phishing continues to be the most prevalent attack type hosted in the UK, despite a 30% reduction in takedowns in 2022 compared with 2021. Since we started the Takedown service in 2016, we have been measuring the proportion of global phishing hosted in UK IP address space. This has significantly and consistently decreased over time, from a high of 5.3% in June 2016 to a low of 1.7% in December 2022. While we cannot directly attribute this trend to the action of the Takedown service, we have made the UK a less attractive jurisdiction for scammers to host malicious content.

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 and web sites to the NCSC. These reports are sent on to our takedown provider for analysis and, when links to malicious sites are found, we seek to remove those sites from the internet to prevent them doing further harm.

Progress in 2022

SERS reports were responsible for the removal of over 72,000 malicious URLs across 40,000 scam campaigns in 2022. Malicious URLs reported to SERS were removed from the internet, on average, within 6 hours.

It should be noted that most of the malicious content alerted via SERS has already been discovered by the Takedown service provider (Netcraft) using other means, which is why the number of URLs attributed to SERS is generally low. We are investigating why an increase in reports is not leading to an increase in takedowns, but it is likely we are seeing the same attacks more times as the number of reports grows.

In 2022, SERS received over 7.1 million reports from members of the public, an average of over 19,500 a day. This is an increase of over 33% on the number of reports received in 2021.

Considering the number of reports received on a monthly basis, there were approximately 30,000 in May 2021, which was the low point. Since then, there has been consistent growth culminating in a record of nearly 80,000 reports received in December 2022.



Figure 7: Reports submitted to SERS

This increase in submissions is partly explained by an increase in the public awareness of SERS. In 2021, we received reports from around 375,000 unique users, whereas in 2022 this increased to around 437,000.

We continue to have a number of 'super users' who are regularly sending us high volumes of reports. We received over 5,000 reports from each of 46 different contributors, with the top 10 being responsible for sending over 213,000 between them. Over 96% of reports are received via email. The remainder are from a combination of the NCSC website form and [O365 'report phishing' function](#).

Outcomes

Takedowns resulting from SERS reports show the extent to which attackers use legitimate brands to lure in victims. Table 6 shows the brands with the most takedown groups resulting from SERS referrals.

Table 6: Top brands used in takedowns from SERS referrals

Category	% of SERS takedowns (groups)
BT	10.1
Santander UK	7.0
Yahoo	5.2
PayPal	5.1
Webmail Users	3.9
Other Brands	40.9
N/A	27.7

Attackers continue to send mass email campaigns (purporting to be from reputable brands) to trick victims into clicking on links to malicious websites. These emails are often very believable and difficult to distinguish from the real thing, as the following example shows.



Figure 8: Example scam email sent as phishing campaign

SERS reports resulting in takedowns against UK government brands are much smaller in number. We believe this is because the NCSC **Takedown Service** is actively searching for scams involving these brands, and so most instances are found before any reports are received.

Table 7: Top UK government brands used in takedowns from SERS referrals

Category	% of SERS takedowns
National Health Service	1.9
TV Licensing	1.0
HM Revenue & Customs	0.9
Gov.uk	0.9
DVLA	0.7

Mail Check

www.ncsc.gov.uk/information/mailcheck

About the service

Mail Check is the NCSC's service 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.

Email Security Check

emailsecuritycheck.service.ncsc.gov.uk

Email Security Check (ESC) is the lightweight version of **Mail Check** which is publicly accessible to anyone. **ESC** provides a quick and simple way of checking your organisation's email security (by providing understanding in areas such as anti-spoofing and email encryption) and acts as a gateway to more advanced services.

Progress in 2022

Mail Check is a mature service with a large existing user base. However, 2022 saw a growth in both the number of users and organisations using the service. This was primarily driven by increased usage in the academia sector, with significant numbers of schools signing up in response to a marketing campaign run in conjunction with the Department for Education, and a pilot in the charity sector.

Table 8: Organisations using Mail Check, by sector

Sector	Organisations using Mail Check Dec 2021	Organisations using Mail Check Dec 2022	Change
Central government departments and arms-length bodies	123	131	+8
Local government	352	302	-50
Health	190	191	+1
Police and fire and rescue services	63	56	-7
Devolved administrations and their agencies	66	69	+3
Academia (universities, colleges and schools)	515	1,258	+743
Charities	188	410	+222
Other	33	35	+2
Total	1,530	2,452	+922

DMARC adoption

In order to make a domain as difficult to spoof as possible, the NCSC recommends a DMARC policy of either 'reject' or 'quarantine'. Table 9 shows the **Mail Check** user base adoption of this standard by sector.

In 2022, we reached the landmark of 100% of all central government organisations adopting a strict DMARC policy. We could not have achieved this without the very close working relationship that we have with the Government Security Centre for Cyber (Cyber GSeC) and the invaluable level of tactical support that they provide across government.

Table 9 highlights that a school subscribed to Mail Check is over four times more likely to have securely configured its DMARC policy (EDP = email data protection).

Table 9: Organisations adopting a strict DMARC policy, by sector

Sector	Subset of organisations tracked	31 Dec 2021 % orgs with EDP	31 Dec 2022 % orgs with EDP	Change (%)
Central government	44 (government departments + 10 Downing Street)	91	100	9
Central government	223 arm's length bodies	60	71	11
Local government	404 (UK principal councils)	81	87	6
Health	279 NHS Trusts and key central functions	35	42	7
Devolved administrations and their agencies	Includes local authorities, health services and emergency services in devolved administration regions	48	62	14
Police and fire services	51 police forces and 54 fire and rescue services	76	87	11
Universities	164 universities, university colleges and other degree-awarding bodies	23	29	6
Schools	Benchmark sample of 1,000 schools across England	6	9	3
Schools	Schools using the Mail Check service	n/a	41	n/a
Charities	Top 3,000 charities in UK by income	13	17	4
Charities	Charities using the Mail Check service	n/a	48	n/a

Integration with MyNCSC

During 2022, we completed our integration with the **MyNCSC** service, and the majority of **Mail Check** users have been migrated across. This migration brings the benefits of being on MyNCSC, and ensures the service is sustainable and supportable for years to come.

Email Security Check

As **Mail Check** is currently only available to certain eligible sectors (the UK public sector, academia and charities), we launched a lightweight version of the service that could be offered to all organisations, called **Email Security Check**. Since its launch at CYBERUK in April 2022, it has scanned over 54,000 email domains.

Outcomes

Spoofted emails

Mail Check's DMARC insights feature shows to what extent an organisation's email domains are being spoofed, and how many of these emails are being blocked. By looking at this data holistically, we can see which sectors are the most targeted by fraudsters and prioritise which organisations we should work with to implement strict DMARC controls.

Table 10 shows the top five sectors with the most spoofed emails detected in an arbitrarily-chosen period (a 30-day period taken as the report was being prepared). Central government departments are the most targeted domains, but the implementation of a strict DMARC policy has ensured that almost all of the 83 million spoofed emails are blocked before reaching their target.

Table 10: Top five sectors with the most spoofed emails detected

Sectors	Org count	Spoofed emails sent over the previous 30 days	Malicious emails blocked by DMARC	% malicious emails blocked
HMG: ministerial departments	15	83,373,427	83,020,928	99.6
Academia	762	4,112,819	2,039,385	50
HMG: local government	266	538,435	466,467	87
Health	147	428,047	105,300	25
HMG: agencies and public bodies	81	154,810	66,363	43

Large spoofing campaign

In 2022, a large public sector organisation approached the **Mail Check** team for help with implementing DMARC reporting in order to combat their email domain being spoofed.

This organisation's domain was regularly being spoofed, with up to 40,000 emails per day being fraudulently sent in their name to unsuspecting members of the public. By enforcing a strict DMARC policy we managed to stop this abuse, which can be seen by comparing the number of emails sent which are categorised as 'untrusted' and 'rejected' in Figure 9.




Figure 9: Sent emails categorised as 'untrusted' and 'rejected' in a large public sector organisation

MTA-STS uptake

In 2021, **Mail Check** implemented full support for the new internet standard MTA-STS (Mail Transfer Agent Strict Transport Security) and provided guidance for how to implement it.

Throughout 2022, we have continued to see an increase in the uptake of this standard, noting that organisations which have implemented it correctly will have greater protection against MITM attacks; the levels of increase in government are primarily due to the efforts of Cyber GSeC, through encouragement and hands on guidance and support (for example in how to implement on AWS/Microsoft/Google).

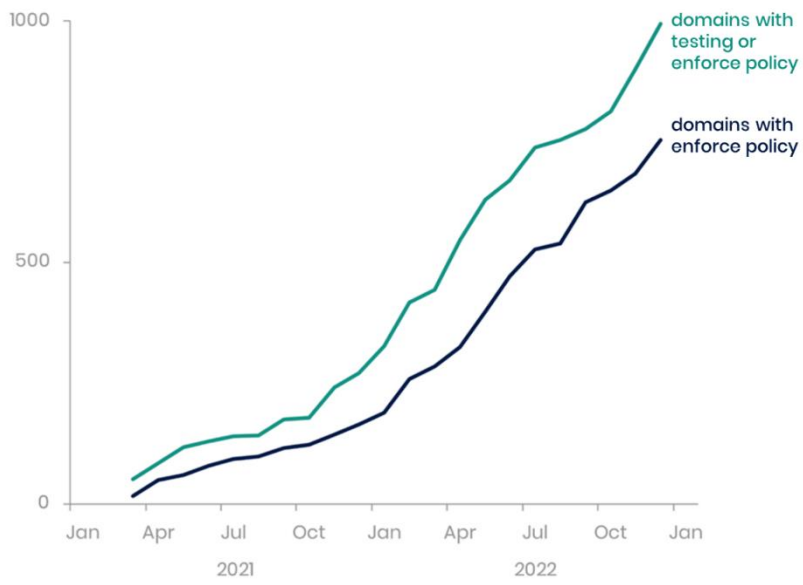


Figure 10: Domains implementing MTA-STS

Vulnerability Checking

About the service

The NCSC has been offering vulnerability checking for some time, but in 2022 we refreshed our offer to provide a new two-tier approach to:

- deliver improved findings to the existing Web Check user base; and
- encourage more organisations to try our new, simplified Check Your Cyber Security service.

Check Your Cyber Security

checkcybersecurity.service.ncsc.gov.uk

Check Your Cyber Security is a tool that is built to be radically simple to enable non-technical users find and fix some of their most important cyber security issues without requiring ongoing support from the NCSC. It is geared towards empowering non-technical users to fix their vulnerabilities, exploiting the NCSC's data and expertise at scale. It is simple, scalable, cost-effective and cyber security outcome-focused.

Web Check

ncsc.gov.uk/information/web-check

Web Check helps 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 MyNCSC interface, together with appropriate and clear mitigation advice.

Subdomain Takeover Alerting and Reporting Service (STAR)

STAR notifies users of subdomains that are potentially vulnerable to misuse due to a lack of maintenance in Domain Name System (DNS) records. When a DNS record points to a site or other resource that no longer exists (a 'dangling DNS'), there is a vulnerability as these resources can be hijacked (registered by another party) resulting in the dangling DNS pointing towards a site under the control of an attacker, which can make the web site seem trustworthy and hence exploitable as part of phishing attacks.

Progress in 2022

Check Your Cyber Security

Web Check is one of our services which focused on the public sector and, although we have expanded where it has been practicable, currently it is only available to certain sectors. After conducting extensive user research, we introduced a product that can be used by all UK organisations called **Check Your Cyber Security**. This is a digital service that is built to be radically simple to enable non-technical users to find and fix some of their most important cyber security issues, without requiring ongoing support from the government. It is highly performant, cost effective and built using serverless tooling and modern architecture.

Our user research (we had ~1,500 pilot users and conducted user research with over 570 participants) showed that small organisations appreciate the need for cyber security but had low confidence in the adequacy of commercial offerings given the lack of an authoritative voice to advise them on what was 'good' and 'not good'. The NCSC does not believe it is feasible to assure such a plethora of offerings, but respondents were very enthusiastic about the UK government, as a trusted authority, stepping-in to provide a tool that helps provide relevant cyber security advice and checks.

Provision of cyber security findings to non-technical users is a relatively unexplored market for data vendors. Most of those are geared towards selling their products to groups that have higher technical maturity to understand more raw findings data. Our premise is that we can harness the collective expertise of the NCSC and deploy that at a large scale with easy-to-use guidance and language for non-technical users, supported by rigorous user testing and behavioural science experiments. We can leverage existing commercial data and harness the collective expertise.

What vulnerabilities does it look for?

Check Your Cyber Security currently looks for seven of the most common vulnerabilities in the UK that we could identify in existing commercial scan data (as opposed to needing to scan for these vulnerabilities ourselves).

In addition to these seven vulnerabilities, there is a separate check of the user's internet browser version to ensure they are using the most up-to-date version available. This check has been an aspiration of the NCSC, with code having been previously developed but not deployed until now.

Use of data sets

From commercial reviews completed, the bulk data set offers in the market are designed for enterprise-level technical users rather than small organisations. The cost (c£100k) coupled with the user limitations (for example cost barriers, lack of time, need for confidence) means that few organisations in the target set would use a commercial offering. Some bulk scanning providers offer free or freemium solutions. However, these generally act as a product trial, and our user research suggests organisations would not use these offers without technical staff to explain the scan results.

Check Your Cyber Security queries data sets from two commercial data providers (Censys and Driftnet) to provide the vulnerability information that is presented back to users.

Automation of processes

We have built automated load testing, cost projections, and disaster recovery exercises into the continuous deployment pipeline for this service. As these are now part of the regularly running automated processes, we have a higher degree of assurance of the resilience of the service. In particular, the automated weekly disaster recovery tests that build up environments and deploy from code have proven helpful in ensuring we are always able to restore to a recent state, catching configuration drift early as opposed to annually, when these recovery tests are often done when manually performed. We believe this to be the first public-facing project within the NCSC to do this and hope to establish this as a new norm for development.

Future aspirations

We plan to issue communications about the product more widely. We will be adding additional checks over the coming months to cover more technologies that should not be connected to the public internet.

Web Check

Scanning capabilities

Initially **Web Check** only scanned for the most common vulnerabilities. Since then we've increased its scanning capabilities to look for over 60 common vulnerabilities and exposures (CVEs), as detailed on the [CISA](#) (Cyber Security and Infrastructure Security Agency) catalogue, using open source tools.

We have also decommissioned scans that were no longer providing value to our users. These include a scan that checked for a vulnerability in the Citrix Application Delivery Controller and Citrix Gateway that, if exploited, could allow an unauthenticated attacker to perform arbitrary code execution, as it hadn't returned any findings in a very long time. We also removed a scan that checked if the website in question was using a reserved government second-level domain such as [gov.uk](#), [nhs.net](#), [ac.uk](#) etc. This change reflects the broadening user base of **Web Check** and the assets they are scanning.

We have also been testing **STAR** capability within **Web Check** which resulted in over 80 dangling DNS records, which we reported to users via manual routes. In 2023, we plan to begin looking at how we can introduce these as findings within **Web Check**, presenting them via the **MyNCSC** service. We have also completed a comparison exercise of the findings produced via **STAR** versus those that would be achieved by using other dangling DNS tools to solidify the decision that **STAR** is the right product to use.

Findings

Table 11: 'Urgent' Web Check findings in 2022, with resolutions

Month	Total Unique URLs (cumulative figure)	Urgent Findings Detected	Urgent Distinct Findings Detected	Urgent Findings Resolved	Urgent Distinct Findings Resolved
January	48,734	1,179	982	1,220	1,012
February	50,434	960	837	1,003	862
March	52,453	871	726	981	795
April	54,559	747	633	696	577
May	57,026	1,286	1,048	1,183	945
June	58,206	748	627	739	641
July	59,034	1,122	979	1,037	929
August	61,382	967	776	997	821
September	61,758	1,310	1,071	1,253	1,081
October	62,031	1,237	1,016	977	871
November	64,072	1,242	982	1,096	952
December	65,253	922	747	824	723
Total	N/A	12,591	10,424	12,006	10,209

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 organisations to respond to the findings presented to them, thereby improving the security of their websites. Throughout 2022, it presented over 12,000 urgent findings to users; 95% of these were actioned.

Web Check has also had an increase in the number of unique URLs that it is scanning, from ~46,500 in December 2021 to ~65,200 in December 2022.

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. For example:

- in May, we broadened out to the schools sector, which resulted in many new assets being subscribed to Web Check which had their websites' certificates expiring; and
- in September, there was a major Content Management System version update which led to a spike in findings due to users having not updated to the latest version.

MyNCSC migration

Web Check, alongside **Mail Check**, is migrating all of its users and assets onto the **MyNCSC** platform. As of December 2022, over 65% of users had migrated onto MyNCSC with the majority of users having received their initial invitation email to migrate. The **Web Check** team have spent time supporting users by migrating assets on their behalf to increase the speed of uptake and improve the experience of onboarding new users onto the platform. The migration completed early in 2023.

Increasing Web Check's user base

Web Check has seen an increase in the number of users, from ~3,700 by the end of 2021 to ~4,900 in 2022. This has been achieved through further take up in sectors already served by **Web Check** and by broadening to additional sectors such as schools and multi-academy trusts (MATs), social housing, parliamentary parties and overseas territories.

Wider impact & future aspirations

The **Vulnerability Checking Services** are working alongside another part of the NCSC's work to build a data-driven view of 'the vulnerability of the UK'. This work, which includes active scanning of internet-accessible systems hosted within the UK, will inform the future direction of our ACD Vulnerability Checking Service, ensuring we prioritise the most important and impactful features to protect users.

In alignment with this work, the **Vulnerability Checking Service** will look to continue to expand its vulnerability detection during 2023, firstly by surfacing dangling DNS records to users via **MyNCSC** using the **STAR** service. We also plan to expand into infrastructure checks, helping to identify vulnerabilities within users' public-facing infrastructure.

Outcomes

By Web Check raising 12,000 urgent findings to users, and 95% of those being resolved, the chances of these vulnerabilities being exploited has been reduced and we are actively contributing to the resilience of eligible organisations.

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. 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 UK public sector users. It prevents the successful resolution of domains associated with malicious activity, while enabling the rest of the internet to remain accessible. We encourage organisations who are not eligible for **PDNS** to take advantage of similar services available in the market.

Progress in 2022

PDNS has continued to grow in terms of number of UK organisations it protects, the number of queries and blocks it performs, and the protection it offers to users from an ever-changing threat landscape.

PDNS helps protect over 1,200 UK organisations, with 228 new accounts added in 2022, including a pilot for UK-registered social housing providers and management organisations. Whilst the number of organisations using PDNS steadily increases month on month, there was a significant jump in February 2022 when we brought onboard a group of town and parish councils, which accounted for over 60% of the new users in 2022. By sector, we saw the most growth in government agencies and public bodies with the support of colleagues at the Cyber GSeC.




Figure 11: Organisations using PDNS (2022)

As **PDNS** is a mature service, 2022 has been about improving the ‘user experience’; focusing on the improvements that mean the most to our users, and ensuring that we maintain a high quality of service.

User feedback and participation in shaping **PDNS** are critical to the service. Our delivery partner, Nominet, has continued to grow the [user community](#) to make it easier for our users' voices to be heard. Users can create and vote on ideas that would improve **PDNS** for them, see the development roadmap, and communicate with the **PDNS** product team directly. Feedback from the user community has led to several improvements on the [PDNS portal](#) to make it easier to navigate through the areas that are most used.

The two areas that we have had most feedback on are:

- protecting roaming devices that aren't always on an enterprise network; and
- identifying devices making blocked DNS queries.

We are currently developing integration tools that will make it easier to ingest the PDNS data into the most popular Security Information and Event Management (SIEM) tools.

Roaming devices

In September 2020, we launched PDNS Digital Roaming, an app for Windows 10 that directs DNS to **PDNS** when the device is not connected to its enterprise network. By the end of December 2022, it was deployed to 55,000 devices, up from 23,000 in December 2021. Since it launched, demand for macOS and iOS has been high, and we heard loud and clear that users want deployment and management mechanisms to be simpler. In response to this feedback, a trial is underway with pilot organisations and a new version of PDNS Roaming is due to be launched in 2023.

PDNS Roaming will bring a simpler deployment mechanism for Windows devices and add iOS and MacOS support via configuration profiles. It will allow users to tag a device, or a group of devices, which is passed through into the block data if a device makes any blocked DNS requests. In trials, this has been shown to help to identify the exact device that has made a blocked DNS request.

Block data

Block data can be very valuable to an organisation's security teams, and we recognised that we needed to improve access to this data. Whilst it is simple to access the block data from **PDNS**, it is not always as easy to ingest it into SIEM tools due to the various formats used and differences in those tools. We have prioritised the most popular SIEM tools used by **PDNS** protected organisations (Microsoft Sentinel, Splunk and Elastic) and are currently developing new integration tools that will make it easier to ingest the **PDNS** data.

Data analysis to optimise the service

Nominet research into newly registered domains and connections between domains led to the development of new algorithms for malicious domain detection. Previously unknown malicious domains using shared infrastructure were identified by clustering domains with shared characteristics.

Phishing is often the precursor to higher severity incidents, as stolen login credentials can give an attacker an initial foothold in a victim's network. Nominet analysts have developed several techniques which in combination identify potentially fraudulent or phishing domains for blocking by **PDNS**. Techniques including word fuzzy matching, tokenisation, and typo-squatting searches are used on newly observed domain (NOD) names to identify brand names, keywords, and other suspicious characteristics. The results are automatically combined with additional data, including DNS records, WHOIS data, and web content data to present analysts with domains that match phishing website patterns.

This research has been developed into a new threat feed for **PDNS**, using the new threat-detection algorithms in addition to a range of open source intelligence. In November and December, this feed blocked over 20,000 unique domains, not seen by any other feed provider used for blocking in **PDNS**, which were queried over 35,000 times.

As a side benefit, the dataset of confirmed malicious websites is used as training data for machine learning models that are used to automatically detect phishing websites not yet found within the dataset.

Finally, Nominet presented a series of webinars to inform and support **PDNS** users, as well as hosting several highly valuable face-to-face working groups to gather feedback on how we can improve the service. The engagement is reflected in the 'user satisfaction' and 'net promoter' (NPS) scores, which have been consistently high in the past year, as shown in Figure 12.

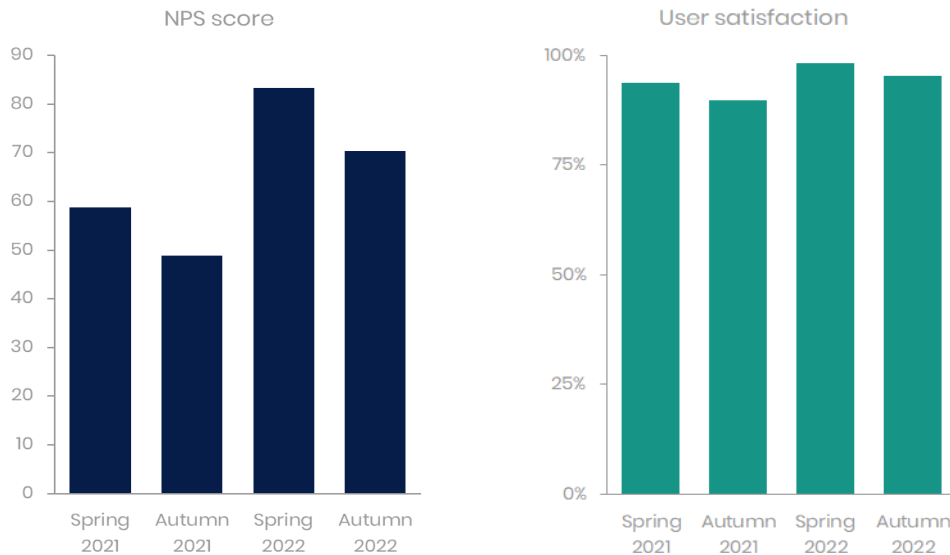


Figure 12: 'Engagement' measures for PDNS (2021, 2022)

Outcomes

In 2022, **PDNS** handled 0.81 trillion DNS queries, and blocked 11 billion DNS queries for 420,000 domains, corresponding to 2% of all queries (so, either **PDNS** is protecting users from an increased number of attacks, or it is catching more, or both). Many of these queries were from our largest organisations, which creates outliers in reporting, so some data has been excluded from this report.

On 25 February 2022, the day after Russia invaded Ukraine, we blocked a spike in queries to domains linked to the advanced persistent threat (APT) group known as Gamaredon (or Primitive Bear), which is known to carry out cyber attacks against Ukraine. These domains were registered under .ru top level domains (TLDs) and have between six and nine characters in the second level domain, which are characteristics that can be used to help identify these domains. Around the same time, we blocked an increase in domains from domain generation algorithms (DGAs) registered under .ru, .cn and .su TLDs.

Threat attribution is dependent on our sources, and there can be inconsistencies, however some of the most directly attributable threats are shown in below, along with 2021 for comparison. We have seen increases in all categories, but it is worth noting, as previously mentioned, that 228 new organisations were added during this time.

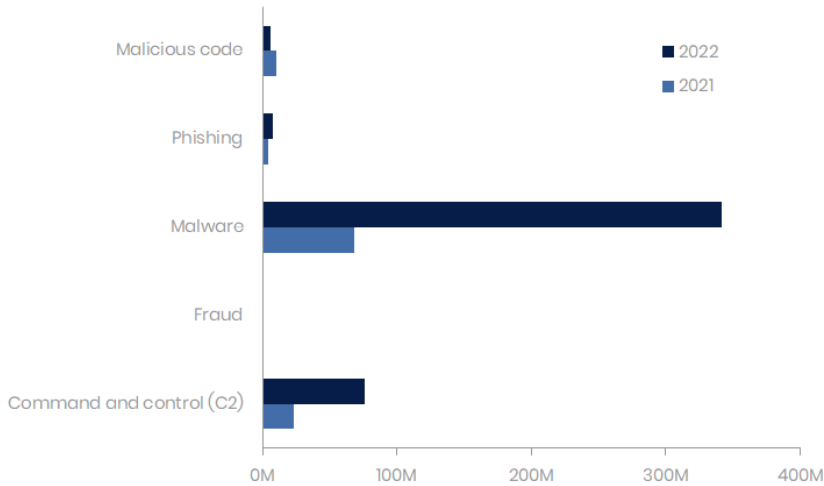


Figure 13 Top attributed threats by type

In 2022, the most blocked (attributed) threat was Cobalt Strike. Cobalt Strike is a penetration testing product that is frequently misused by malicious actors for command execution, lateral movement in a network, and dropping malware, among other malicious acts.

SUNBURST (associated to the SolarWinds compromise in December 2020) was the second most blocked attributed threat seen on **PDNS**. Due to its age, we suspect that these blocked DNS requests often come from organisations' security devices. The top attributed threats blocked in 2022 can be seen in Table 12.

Table 12: Most blocked attributed threats (2022)

Threat Name	Unique Domains	Total Blocks
Cobalt Strike	295	15,414,417
SUNBURST	29	7,109,940
Flubot	109,415	5,880,600
CryptoStealer	42	3,109,973

Flubot is a notable outlier and for more information on why that is, please refer to [ACD: The Fifth Year](#).

Table 13 shows the most commonly seen attributed threats blocked across PDNS organisations in 2022. The most widespread threat was SocGhosh, a delivery framework that enables drive-by-download watering hole attacks, which can lead to the delivery of ransomware and remote access trojans (RATs) to a victim's device.

Also notable is Emotet, which despite a widely reported takedown in January 2021, was still consistently amongst our most blocked threats. Emotet was originally developed as a banking trojan, but until January 2021 was commonly used as a dropper for other forms of malware such as ransomware.

Table 13: Most seen attributed threats blocked across PDNS organisations (2022)

Threat Name	Organisations (%)	Unique Domains	Total Blocks
SocGhosh	25	63	13,788
JSRedir-OE	24	313	42,027
Emotet	23	612	1,653,700

PDNS blocked over 5 million requests for domains associated with ransomware. Table 14 shows the top five of these threats by the total number of requests blocked. Though some blocked DNS requests were undoubtedly linked to malicious actors, we suspect that they often came from organisations' security devices. Conti Ransomware was the most blocked ransomware by this measure, which is unsurprising as the group behind this malware is believed to be one of the largest ransomware actors.

Table 14: Top five threats by the total number of requests blocked (2022)

Threat Name	Total Blocks
Conti Ransomware	1,350,986
Petya	641,135
NotPetya	529,442
FiveHands	522,434
Knot Ransomware	239,864

Table 15 shows that the most widespread blocked ransomware threat was Phoenix CryptoLocker. Phoenix CryptoLocker has been associated with the APT group known as EvilCorp (or Indrik Spider) and interestingly, all blocked requests were for a single domain.

Table 15: Top 5 most blocked ransomware threats (2022)

Threat Name	Organisations (%)	Total Blocks	Unique domains
Phoenix CryptoLocker	9.7	159,238	1
Wcry	9.5	150,114	15
HydraCrypt	8.6	99,629	2
Locky	3.5	81,300	39
Sodinokibi	2.8	116	22

Microsoft 365 continued to be a popular target for credential theft attacks on **PDNS** users. We observed multiple malicious spam emails containing links to domains containing fake Microsoft 365 sign-in pages. Another notable credential theft campaign observed in 2022 targeted the SAP Concur expenses platform. Over 30 previously unknown malicious hostnames associated with this campaign were detected by **PDNS** analysts and blocked.

We observed an increasing number of Android device infections, including Sharkbot, Octo, Gigabud, and Joker, although Flubot has decreased. As we looked at in [ACD: The Fifth Year](#), we observed a sharp rise in blocked DNS queries to domains associated with Flubot in 2021. It uses a DGA to search through potential command and control (C2) domains until an active server is found, which results in a *lot* of blocked DNS queries, with a single infected device often generating in excess of 20,000 queries per day.

Although the main Flubot infrastructure was taken down in May 2022, we continued to block high volumes of DNS queries linked to Flubot. However, the number of organisations generating these queries dropped from 25 in 2021 to 7 in 2022, indicating that although infections may still be active on some Android devices, the number of infected devices connecting to **PDNS** is dropping.

Another threat analysed by the Nominet team was Gootloader, which uses compromised vulnerable domains split into two roles: C2 and phishing. The latter are used to host malware download links hosted on specific search engine optimisation (SEO) termed forums, which is a technique known as 'SEO poisoning'. Several techniques to obfuscate these activities are used, such as the page only appearing when clicking through from a search engine. When the obfuscation is successful, the visit will instead result in a fake blog page relating to the search term used.

The forums always have the same layout, however the wording and language changes. The 'thread' subject and download links are based on search engine query key phrases, for example "standardized United States sales tax 2020". Once a link is clicked, a file containing the name of the search phrase is downloaded from another compromised domain used as a C2 server. During monitoring of this attack chain, our analysts noticed that the initial search sites are available for a longer period than the download C2 compromised domains, which are changed frequently, often daily. All Gootloader domains we know of are blocked by **PDNS**.

PDNS analysts also observed queries to domains associated with 'dropper' malware, commonly used for enabling ransomware or RATs. Due to their severity, our work focused on improving protection against droppers.

IcedID is a banking trojan which was involved in several high-profile ransomware attacks in 2022 as a dropper for other threats. IcedID typically communicates with its C2 server using DNS, giving **PDNS** the ability to disrupt it before data exfiltration or encryption have taken place. During 2022, analysis identified over 200 previously unknown IcedID payload distribution and C2 domains that are now blocked by **PDNS**.

Bumblebee is a malware loader discovered in mid-2022, attributed to the Conti APT group and designed to replace the BazarLoader backdoor used to deliver various ransomware payloads. Our analysis discovered multiple, previously unknown, Bumblebee domains that are now blocked by **PDNS**.

We have continued to successfully use DNS-based blocking to disrupt a variety of high-risk attack chains, including those that later in the chain do not rely on DNS as a primary means of communication (for example, Cobalt Strike).

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 common cyber security incidents in a 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.

Initially aimed at non-technical audiences within both the public sector and SMEs, **EiaB** has also seen strong take-up amongst large organisations and cyber security professionals.

Progress in 2022

Early in 2022, we successfully took **EiaB** from 'Public Beta' to 'Live' status. We then prepared to rebuild the application with a different framework which also allowed us to simultaneously start work on a redesign and refresh of the look and feel of the content. The application was subsequently re-launched in November.

We also created two new exercises:

- Supply-Chain Ransomware Attack (table-top exercise)
- Securing Video Conferencing Services (micro exercise)

Our colleagues from the Australian Cyber Security Centre (ACSC) got in touch and were looking to host their own instance of **EiaB** on their own infrastructure. ACSC have now launched their own version to serve Australian organisations. We have a memorandum of understanding which allows for a 'federated' approach to content creation and are very much looking forward to working with them to co-create new exercises.

In 2022, just over 18,500 users worldwide signed up, which represents an increase of around 40% over the previous year. We continue to hold large-scale events to increase take up of the service, the largest one reaching the health supply sector with over 880 participants. In terms of the individual group sign-ups:

- public sector up 37%
- SMEs up 36%
- large businesses up 61%
- cyber security professionals up 50%

We also worked with the Scottish Business Resilience Centre (now the Cyber and Fraud Centre Scotland) to promote **EiaB** to Scottish businesses. They have done an impressive job by holding 46 events covering nearly 140 organisations.

We have continued capturing feedback from users concerning:

- the usefulness of the **EiaB** exercises
- whether they plan to make changes as a result of running the **EiaB** exercises

The results are summarised in Figures 14 – 16, which show that ‘Supply Chain Software’ and ‘Threatened Leak’ are the most useful of exercises, whilst ‘Supply Chain Ransomware’ and ‘Insider Threat’ are the most likely to illicit change to an organisation’s processes and procedures. However, it should be noted that the bias on all of the exercises is very much to the right in each of these Figures:




Figure 14: Overall ‘Usefulness’ of EiaB exercises (2022)




Figure 15: Usefulness of specific EiaB exercises (2022)




Figure 16: Intent to make changes after exercises (2022)

Early Warning

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

About the service

Any UK organisation with a static IP address or domain name can sign up to use Early Warning, which is a free NCSC service designed to automatically 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, as well as 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.

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.

Organisations using the Early Warning service can receive the 3 types of high-level 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.
- Potentially Malicious Activity: 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 your network, or potentially undesired applications are exposed to the internet (such as an exposed Elasticsearch service).

Progress in 2022

In 2022, 2,939 new user organisations signed up to the service, a 38% increase on the previous year, with a total of 7,819 organisations at the end of 2022.

- 570 organisations were warned about active malware on their networks.
- 2,270 were warned about vulnerabilities on their networks.
- 1,193 were warned about a host on their network scanning the internet (which might be - for example - an indicator of a possible compromise).

Early Warning ingested a total of 1.49 billion events from our data suppliers.

- We sent out 32.8 million events, approximately 2.2% of the total data we received; those events went out to 5,910 user organisations.
- We sent 41,000 daily email notifications of possible malicious activity such as malware infections or activity suggesting that a user device had been compromised; those emails related to 1.14 million IP addresses in total.
- We rolled out a new capability for the NCSC to send targeted 'compromised credentials' alerts to users. This is used when our partners obtain fresh information about user credentials obtained by criminals through phishing sites or other methods.

Outcomes

Every alert that Early Warning sends out is valuable and should be investigated (since it is highly likely to indicate some kind of incident has occurred). However, not all reports are as urgent as others. The most time-sensitive varieties of notification are those that relate to activity commonly seen prior to ransomware being deployed on a victim's systems. In 2022, the automated service sent notifications to 56 organisations to warn them about pre-ransomware malware infections.

Another common route of attack for ransomware actors is via the Windows Remote Desktop Protocol (RDP) service. Organisations frequently (and accidentally) leave this service exposed to the internet without multi-factor authentication in place, allowing criminals to brute-force passwords to gain access to victims' systems. In 2022, Early Warning sent alerts for 67,000 IP addresses which had the RDP service exposed to the internet. On average, Early Warning users receiving these alerts left the RDP service exposed for 19.7 days, whereas IP addresses that did not belong to our users left this service available for 49.3 days.

Overall in 2022, Early Warning notified users of malware infection on 823,000 IP addresses (out of 23.8 million that were reported to us). The average lifetime of malware on a user IP was approximately 70% as long as it was on non-user IPs.

Although not every alert we notify is urgent, there are some families of malware we're surprised we still see. [Conficker](#) was a very well-known and widespread self-propagating piece of malware back in 2008. In 2022, 14 years later, we still sent out notifications for Conficker being found on 2,869 IP addresses.

In 2022, Minerpanel, Avalanche and Cobalt Strike were the infections reported on the largest number of Early Warning user organisations. Ramnit, Citeary and Sality were found on the most IP addresses in total (whether those IPs belonged to an Early Warning user or not).

MyNCSC

www.ncsc.gov.uk/information/myncsc

About the platform

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 replace the ACD Hub as the single point of access to ACD services.

Progress in 2022

During 2022, the focus was on migration of **Mail Check** and **Web Check** users to the **MyNCSC** hosted versions of these services.

Our migration approach evolved during the year. We started with a self-service approach, making tools available for users to migrate their digital assets (previously registered under **Mail Check** and **Web Check**) to **MyNCSC**. But the user response was variable, often as users had limited time. So, we moved to migrating the digital assets for them.

Towards the end of 2022, we needed to chase up those organisations that had not responded to our migration invitation by joining their organisation on MyNCSC. Giving an end date encouraged many to take up the offer to migrate assets before **Mail Check** and **Web Check** services became unavailable. Those who have not taken up the offer can set themselves up afresh on MyNCSC if and when they wish to do so.

With the balance of **Mail Check** and **Web Check** usage moving towards MyNCSC, we also onboarded new users directly onto the MyNCSC versions. Of particular note were the several hundred schools, which took up the offer to use these services in response to a marketing campaign run in conjunction with the Department for Education.

Improving usability

A good design is required to deliver solutions which work well at scale to give our users a good experience and limit the number of queries coming through to our support team.

During 2022, we put a lot of effort into the 'Join organisation' user journey, one of the most complex in **MyNCSC**. Most users will only need to go through it once, but it is a key one to get right. It is one of the first journeys encountered on **MyNCSC** and the one which gets a user into the right organisation to work alongside colleagues when using the ACD services hosted on **MyNCSC**. The complexity is due to a number of factors:

1. There is no guarantee that two users will enter their shared organisation identically and fuzzy matching logic is tricky. Hence, **MyNCSC** invites users to search for their organisation in our database, which contains thousands of organisations, drawn from publicly available sources.
2. Within the database, there are organisations with the same name, so further differentiation is required, and hence we have provided context-specific advice to help the user make the right choice.
3. Finally, in a few cases, an organisation simply does not exist in our database (our sources might not have captured their organisation or perhaps the best fit is a working partnership, comprising multiple organisations, which isn't formally recognised). So the user has the option to request creation of their organisation where they have a genuine case for us to do so.

“Teams” functionality

The profile of ACD usage can vary from one organisation to another. For the majority, use of our services is small scale, with relatively few digital assets to register and the involvement of only a few users. However, for larger, more complex organisations, operating at the organisation level on **MyNCSC** would not be practicable. Teams can be set up within a given **MyNCSC** organisation to cater for different areas of the organisation’s IT estate. Users and digital assets are assigned to specific Teams, with findings for those Teams enabling individual users to focus on that part of the IT estate for which they have responsibility.

We remain aware of the contrasting needs of other organisations. Teams is less well suited to those cases where some users focus on email, whilst their colleagues are interested in web security. Filtering assets and resultant findings by subscribed ACD service offers a better approach in this context, although we need to do more to enable filtering of notifications.

For organisations with larger volumes of digital assets, we were also mindful of the need for users to manage those assets efficiently. Functionality released during 2022 included improvements for bulk upload. The volume of resultant findings requires similar consideration. Improvement work planned for 2023 includes a list-based presentation, which a number of users have expressed the desire for via our user research, and an improved search capability.

System admin functionality

As you would expect, **MyNCSC** requires system admin functions. A number were developed in 2022, which have enabled us to improve our support efficiency to focus on developing further functionality. One significant example is the ability to add a user to an organisation, which is achieved through the appointment of at least one administrator (Org Admin) for each organisation registered with **MyNCSC**. We overcame the issue of a lone Org Admin ceasing to act on behalf of an organisation, which had inhibited business (as there was no capability to approve further requests to join it).

Asset reverification

With a large number of user organisations moving to the platform, we watched for any systemic concerns around usage. One concern involved the need for asset reverification. **Mail Check** will only perform its DMARC checks where users have verified ownership of their digital assets and set up DMARC reporting. On migration to **MyNCSC**, digital assets previously verified via **Mail Check** were given a temporary verification status, which, after a period, required re-verification via **MyNCSC**. We noticed a number of organisations were not responding to notifications advising of upcoming verification expiry, so increased the comms on this with a fuller information page on **MyNCSC** and targeted emails.

Outcomes

At the end of 2022, 2,800 user organisations were using **MyNCSC**, thereby benefitting from a unified user interface to access Mail Check and Web Check, with the ability to perform some configuration functions just once at the platform level. We look forward to Early Warning joining **MyNCSC**, with much preparatory work for this undertaken during 2022.

It has not been entirely straightforward getting to this point. We have learned some lessons along the way about the challenges involved in porting mature services with different designs onto a common platform and may not repeat this for all other mature ACD services. With cognisance of **MyNCSC** integration patterns however, new ACD services can be developed to operate on it, thereby reducing their ‘time to market’.

Routing and Signalling

About the service

Fixing the underlying infrastructure protocols on which the internet is based has been a key strand of the NCSC's ACD work since inception. Traditionally, we have focused on two specific protocols: the Border Gateway Protocol and the Signalling System No.7. The latter was deprioritised at the end of 2021, so this section of the report focuses on the former.

We have also established the SMS SenderID Protective Registry, to help organisations protect their brand from use in SMS phishing attacks. Progress made in this area is also discussed.

BGP

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.

Progress in 2022

As a result of a Facebook incident ([Understanding how Facebook disappeared from the Internet \(cloudflare.com\)](#)), we have realised that it is critical that we process and understand "withdrawal" messages as well as the messages that we currently process. Several days were spent trying to work out how the Facebook address space had been hijacked, before it became apparent that Facebook had accidentally withdrawn the access to the address space.

We have noted previously that, in collaboration with BT, we have developed a proof-of-concept BGP Monitoring Platform, known as BGP Spotlight. In addition to developing an acceptable use policy for BGP Spotlight, improving the user experience and fixing a number of minor bugs, we made three significant improvements to the Raw Data Download summary, the Trace Routes Manual Trigger and to the Trace Routes ASN and IP detail.

- Raw Data Download is able to retrieve filtered or unfiltered results from across collectors so that users can view what the original BGP messages look like. While we keep unique values in the database (this only goes back 60 days), this feature allows users to go much further back and to get the unaltered versions of the data if they want to dig further into it.
- TraceRoutes Manual Trigger allows users to trigger traceroutes on an existing notification (to give an idea of what a traceroute looks like now, as opposed to when it went through the 1st and 2nd tests), whilst allowing users to run a traceroute on an event that wasn't monitored in the first place.
- TraceRoutes ASN and IP detail – allows users to see the IP addresses in a trace; the ASN version matches the IP addresses to ASN ownership, and hence determines what the BGP path looks like from a set of IP hops. This helps to highlight where BGP boundaries are in the path and potentially match where BGP anomalies might match up from the alert.

We have also made a number of improvements to the way the database holds information, improved the reliability of the data download and implemented scaling on the container-based processing to account for busy and quiet periods. The reference data displayed now includes PeeringDB data for ASN ownership and the addition of API lookups to fill gaps in registry data.

We have a multi-year strategy of continuous improvement of the product as well as taking feedback and suggestions for additional functionality from our users. Future plans include:

Research

We are considering ASN behaviour to identify “good” path changes from “bad” ones. A good change may be a legitimate change that happens frequently, or routinely, or maybe it affects an ASN that has frequent changes, whereas the opposite may hold true for a bad change. Initially this work was carried out by a summer student at BT but has now been taken on by the BT Data Science Hub.

Information that we display on ownership and location of ASNs and Prefixes is currently collected from sources such as Ripe, and is subject to the GIGO law. In the past, we have seen ASNs that we know belong to China Telecom named as “Mickey’s test network”. We are planning to research more reliable sources, or ways to verify the data to remove these anomalies.

Our final piece of planned research is around how a path behaves over its lifetime. By building up a knowledge base of how paths behave normally we hope to be able to more reliably highlight unexpected and therefore anomalous behaviour in paths across the internet.

Episode / Event re-work

Currently, we track episodes in 5-minute windows. This has the potential to result in a small amount of lost data. In the future we will be tracking all episodes from their start to finish, ensuring we have a complete picture of the episode. We also intend to allow episodes to be prefix based, ASN based or path based, giving the user a wider choice, and potentially a clearer picture of what they see displayed to them.

Website

As functionality has grown over the last few years, the web pages have become very cluttered, with information either off to the side, or not being visible until the screen has been scrolled. The menus have become very multifunctional and not at all intuitive. We will be changing the website to use a number of tabs or frames, each dedicated to a specific purpose, each self contained, and highly intuitive. This will significantly enhance the user experience, making the site easier to navigate, and therefore the tool easier to use.

Ingest and message flow

We will be using a new method for data ingestion and categorisation. Categorisation will happen later in the data processing, so that the actual ingest of the data does not have as much “work” to do, and will be able to ingest data in a more streamlined way. This will both speed up the process, as well as reducing the processing requirements and so potentially the cost.

We will be using data from ASN peers to de-duplicate input data. At the moment, if there are 3 peers, they will all have the same data about the ASN that they peer with, resulting in us ingesting 3 lots of data, which are all duplicates. In future we want to only ingest one lot of data, further speeding the ingest and reducing the processing/cost.

When we have finalised our research into what a withdrawal looks like we be creating withdrawal event types to display to users, providing a fuller picture of BGP activity than we are currently able to show.

Reference data

This is the information that we use to identify ownership of ASN, IP and Prefix data that we display. At the moment, we can only allow users to enter single, multiple or a conjoined range of IP addresses or ASNs. We are developing the system so that it will be possible to enter “France” as an example, and all ASNs or IPs known to originate there will be included. This will in some instances significantly simplify the input of alerts and will allow for rules such as “Tell me when UK traffic destined for US goes via Russia” for example.

With the de-duplication mentioned above we will also be improving the statistics that are displayed around the origin location of an IP/ASN. This will improve the accuracy of the data making it more reliable for users.

Outcomes

We now have 48 organisations signed up to BGP Spotlight, with 213 users between them. They are a mix of UK and international users, both telco and non telco.

On a daily basis, we typically ingest 800 million messages (but have seen that peak to 1.7 billion in one 24-hour period). These 800 million messages are processed down to 5 million events that we are interested in.

We know that our users are finding and addressing hijacks, but they are reticent to provide details due to confidentiality.

SMS SenderID Protective Registry

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

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)
- 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 the registry identifies valid sources for specific SenderIDs, 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 from unauthorised/invalid sources following a different path or bogus derivatives (such as DV1A) should not get delivered to users.

Progress in 2022

At the beginning of 2022, we launched our Business Communications Guidance, which is focussed on how organisations can help in the collective fight against fraud. This has been instrumental in facilitating conversations with a number of sectors such as finance, delivery companies, retail and other government departments who regularly feature as 'top-smished' brands according to available data.

Throughout 2022, we have been able to help protect a number of high profile campaigns through use of the SenderID registry and our relationships with the networks, often at short notice.

The capacity of the SenderID Registry has been increased to support more merchants, with 16 new merchants onboarded in 2022, including two new government departments.

We continue to work closely with MEF, UK Finance, the messaging providers, operators and merchants to lay the foundations for improvements to the service to help in the continual evolving nature of fraud against the citizen.

Outcomes

The SenderID registry supports 38 merchants and 30 aggregators in the UK and has also expanded to three territories (a term used by the MEF concerning other countries that logically fit together).

Host Based Capability

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

About the service

Host Based Capability (HBC) is a software agent 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 of the highest threat level, helping departments with their security via three service tenants:

- 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

In 2023, HBC will adopt a threat hunting posture within the operational work of the NCSC.

Progress in 2022

In 2022, HBC focused on sustainment of service provided to multiple departments within its existing capacity, with continued development of the software agent and sharing of threat surface information with our partners. Through this coverage HBC has continued protecting departments.

Detect

HBC worked on four incidents in 2022, providing information that helped the departments targeted to understand the remedial action they needed to take. The HBC team also identified and notified departments of 56 suspicious activity observations (SAOs). These 'irregular' detections by HBC informed departments of suspected but unconfirmed malicious activity, for the respective department to conduct further investigation as needed.

The significance of these SAOs was demonstrated during a threat hunt, when the team detected activity that looked like reconnaissance against a domain controller at one department; specifically, they had seen a JavaScript file running processes to retrieve information regarding domain controllers. As this could have been legitimate admin activity, an SAO was issued. The department confirmed it was not admin activity, in response to which the device concerned was quarantined and the HBC team escalated the event to an incident.

Threat surface

HBC generated 364 threat surface reports in 2022 (281 were generated in 2021). 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

There were no new, major vulnerabilities in 2022 that met the threshold to instigate Forewarn checks and notifications. However, as part of the Threat Surface tenant, information has been provided to users of the number of products they have run that were vulnerable.

Vulnerability Reporting and Disclosure

ncsc.gov.uk/information/vulnerability-reporting

About the service

The NCSC Vulnerability Management Team works to mature the UK’s approach to vulnerability management, disclosure and remediation. We have three public projects:

1. 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.
2. Vulnerability Disclosure for Government Scheme: helps improve the UK government’s ability to adopt best practice disclosure processes by creating a Vulnerability Disclosure Programme that includes triaging the vulnerabilities, for any department that signs up.
3. 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 2022

Vulnerability Reporting Service

The NCSC runs the **Vulnerability Reporting Service (VRS)** in conjunction with [HackerOne](#) who provide the reporting platform and [NCC Group](#) who provide triage of all reported vulnerabilities. The **VRS** has had another outstanding year with nearly four times the number of reports than the first year we launched. We are proud to support security researchers who have taken the time to report vulnerabilities through our front door. By working closely with the system owners from across government, 74% of reported vulnerabilities are resolved within 30 days of being notified.

Analysing the reported vulnerabilities, we found that nearly 10% of all 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 17: Reports submitted to HackerOne (2022)




Figure 18: Resolution timeframes from NCC (from triage to closure)

Vulnerability Disclosure for Government Scheme

The scheme 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 2022, the scheme helped eight UK government departments launch their own Vulnerability Disclosure Programme (VDP). This brings the total number of VDPs to 30, enabling these departments to directly receive vulnerability reports so they can fix the issues before they cause harm.

Vulnerability Disclosure Toolkit

The NCSC's [Vulnerability Disclosure Toolkit](#) contains the essential components you need to set up your own vulnerability disclosure process. We updated the toolkit to include additional information on implementing a disclosure process, including validation and triage.

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. From 31st March 2023, the NCSC ceased its support of **LME**. The [US Cybersecurity and Infrastructure Security Agency \(CISA\)](#) have now taken on LME and relevant comms will be issued as their project progresses.

Progress in 2022

Following the release of Version 0.4 in 2021, further version updates were released with an update to Elastic (7.17.1) in March 2022 including:

- updated mapping files to the latest ECS version
- update of the relevant Winlogbeat install instructions to point to 7.17.1
- update of the Docker stack versions to 7.17.1
- update of the instructions for backing up **LME** logs to a separate drive to complement the latest version of Elastic and Docker being used

Outcomes

We have seen a steady uptake in **LME**, which was cloned up to 1,210 times in 2022, an average increase of 100 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 had been providing them with alerts about cyber threats.

Since its launch in 2019, we have seen 3,635 unique clones of **LME**, averaging 85 unique clones per month.

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 CTI Adaptor integrates with a variety of SIEMs, using user 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 user's logs.

The CTI Adaptor has been retired with effect from 31 January 2023. On considering the project outcomes against ongoing development of commercial products, we decided to cease this work to focus resources on other NCSC-specific capabilities.

Progress in 2022

During 2022, version 0.5 of the CTI Adaptor continued to be developed which included the development of the intelligent search feature (which enabled the CTI Adaptor to support larger threat intelligence feeds and prioritise searches based on severity and/or time age) and the addition of proxy support as a new feature.

We also updated support for all SIEM provider schemas and worked on the development of signature-based search queries using Sigma.

Updates in June 2022, as version 0.5.1, included Elastic updates and compatibility with Splunk on Premise SIEM. Further updates in August 2022, as version 0.5.2, included compatibility with Sentinel and Splunk Cloud, compatibility with Splunk Cloud SIEM and ECS/non-ECS configuration.

Outcomes

During the CTI Adaptor Pilot, we engaged with 30 organisations, mainly local authorities and government departments, with the support of the Department for Levelling Up, Housing and Communities, councils and housing to reach local authorities.

During 2022, 2.6 million sightings were detected, which were mainly signatures, with 20 IOCs. The majority of sightings were 'silent' or non-alerting. Silent advanced searches were seen by the NCSC and designed for testing new IOCs and signatures, and to further understand the threat landscape in the UK. Silent searches (non-alerting sightings) were recorded in an organisation's audit logs.

All alerts (non-silent) sightings were visible to the pilot organisations via their dashboards and mitigation actions accessible via the enriched sighting information. We averaged approximately 60 sightings per day to users.

CTI Adaptor was compatible with the following SIEM technology:

- **Logging Made Easy** (LME)
- Elastic – through native Elasticsearch and the NCSC **LME** project
- Splunk – on premise and on cloud, Azure Sentinel and LogPoint [version 4 only]

Sigma Rule detection was a large part of the signature-based sightings data. 233 Sigma rule files, written in YAML, have been created by the Threat Detection & Response (TDR) team. Each file contains one or more detections for a technique. As each file contains one or more detections, not a strict 1-to-1 mapping of each rule to detection opportunities, 297 different detection opportunities were created.

Conclusion/forward look

As the preceding pages demonstrate, an evidence-based approach remains central to everything we do, drawing on the considerable amount of data generated. The principal aim of this report, as with its predecessors, is to use this data to provide transparency, demonstrate what we have learned, and invite feedback and challenge from the cyber security community.

In this final section we also want to include some thinking about where we want to go next with ACD.

Let's start with the things we don't think will change...

These six years of reports tell us that combining digital tools, sensors, services, data and platforms has improved the UK's cyber resilience at a reach and scale that couldn't have been achieved by other means.

Most of our ACD initiatives address enduring cyber security challenges: sharing knowledge of threats, closing down vulnerabilities, responding to breaches. The specifics change over time, of course, but the overall need to tackle them through automation will persist, because as things stand that's the only realistic way of generating the scale and reach required.

One of the great things about digital services is the data they generate, which helps understand the impact we're having, the ability to make tweaks and measure the difference they made. It also means we can be transparent about this with the public and we strive to do that each year through the report we publish. This has not only helped us to understand what works, but inspired many governments around the world to undertake similar initiatives. We still want to understand the value our services provide in an empirical and data-driven way, and our commitment to transparency and openness to challenge is unchanged.

So we think the founding principles of ACD remain sound and are likely to remain so for the foreseeable future. But what does that look like in practice?

Over the next couple of years, we want to double down on the digital services where evidence, feedback from users and our own experience give us confidence that we're getting a good return on investment from a cyber security perspective. That includes working with industry partners on proven services that provide protection at scale (such as Takedown, and Protective DNS), including where they are fed by citizen reporting (through the Suspicious Email Reporting Service) and providing early warning of malicious activity based on the unique vantage point we have as part of GCHQ (Early Warning Service, Host Based Capability).

It also includes making it easier for all sorts of organisations to find and fix basic vulnerabilities, in the public sector (Web Check, Mail Check) and providing more services that are as simple as possible to use and universally available to access (like Check Your Cyber Security and Email Security Check).

Finally, all of our services exchange data, both internally and externally in a point-to-point fashion, from a policy and contractual point of view. We want to keep working on how to be more mature about data architectures and data engineering to be able to be more flexible and agile in our use of data.

Where are we making (or thinking about) changes?

In 5 years' time, ACD will resemble what we have today but with greater reach into organisations – either directly from us or perhaps via many other channels – to make the portfolio even more impactful. But we've also learned a lot about things we want to build on:

1. We need to reinvest in the earlier stages of the lifecycle, making the most of the exposure we get to the cyber security problems our users are facing now and will face in the future, and the technical brilliance of the innovators we have in our organisation and the partners we work with. We want to use that combination to generate a raft of ideas to feed the next generation of features and services.

2. We have a clearer 'checklist' now of the attributes that make a good cyber security service. For example:
 - scale (benefit is conferred to a wider group of users and/or beneficiaries)
 - sensitivity (ability to detect/block more things)
 - accuracy (user more likely to take notice)
 - usability (easier to action)
 - efficiency (in terms of the resources consumed)

We also want to use these, underpinned by data, to make sure new interventions will enhance the overall value of what we do.

3. We need to incorporate digital services much more closely with other things the NCSC does to help improve cyber resilience, such as improving reach and scale through industry (notably our assured industry services) and also the guidance we generate and other resources we produce.
4. We also see a lot of untapped potential in the way we're using the data generated by those tools and services, whether that's about deriving unique insight into threats and vulnerabilities or maximising the protective benefit we get from linking different classes of service together and sharing with industry providers.

Looking a bit further ahead we also see a lot of opportunity (and almost certainly some new challenges) in the shifts we already see in the way technology is developed and used, adoption of cloud services being an obvious example. But we also see potential in the way cyber security capacity and capability is maturing in parts of the public and private sectors, in the UK and overseas. That means more partners to work with, on reach, scale and data. If we get that right, it will mean significant step forward for our cyber resilience ambitions.