Active Cyber Defence - tackling cyber attacks on the UK

Created: 01 Nov 2016
Updated: 01 Nov 2016
Author: Ian Levy
Part of: Cyber strategy, Government strategy, The NCSC

National Cyber Security Centre
a part of GCHQ

On 1st November, the Chancellor announced the Government's new National Cyber Security Strategy and, as part of that, our Active Cyber Defence (ACD) programme. Bits of this have been alluded to before, but I thought it would be a good idea to put down a high level overview of what we're planning to do. We'll be publishing more detail as the programme progresses including - critically - data around what the different measures achieve. This (long!) blog isn't intended to be a full technical deep dive, but I expect we'll be publishing some of those soon.

There's a common complaint from industry to governments about cyber security. It's generally that governments tell them they're not doing enough and must do more, often without really understanding the real-world impacts or commercial implications of their demands. Well, our strategy is to use government as a guinea pig for all the measures we want to see done at national scale. We'll be eating our own dog food to prove the efficacy (or otherwise) of the measures we're asking for, and to prove they scale sensibly before asking anyone else to implement anything.

The ACD programme is intended to tackle, in a relatively automated way, a significant proportion of the cyber attacks that hit the UK. Automation means the measures scale much better. It's not a panacea but should help us mitigate the impact of a significant proportion of the attacks we see. It won't affect the really targeted attacks (at least initially) but we're hoping that we can reduce the noise enough to make the defenders' jobs easier when tackling those very targeted attacks.

The programme breaks down into a number of strands:

Fix the underlying infrastructure protocols

This is about changing the implementation of Border Gateway Protocol (BGP), the protocol used to sort out IP routing between carriers, and SS7, the international telecoms signalling protocol, so that we can stop trivial re-routing of UK traffic and make some more bold statements. If the BGP work succeeds, we should be able to say that hijacking a UK prefix by BGP is harder.

The other thing I'd like to be able to say is that UK machines will not be able to easily participate in a scaled DDoS attack. Once we have proved this works, we intend to work with the international ISP and IX community to have similar protections built in other major exchanges - in order to make DDoS and prefix hijacks globally much harder prospects. The SS7 hardening work should allow us to make traffic re-routing harder but also to make smishing (that's phishing over text message if you've not heard of it) harder in the UK for certain SMS TPOAs (Transmission Path Originating Address - think 'from address'). That'll all be through working with the relevant companies to get the implementation standards written and implemented.

Make email mean something again

Email is the main attack vector in the majority of cyber attacks – regardless of the actor involved – and usually relies on an abuse of the trust in the sender of the email, be that a well-known brand (HMRC, PayPal, etc) or a person known to the victim. The current advice given to people 'don't open an attachment or click a link in an email unless you trust it' is dumb. Most people can't reason through complex internet headers and there's really no other way today of determining trust, so we need to fix that and stop blaming the users.

There exists already a number of internet standards that can help tackle spoofing, including SPF, DKIM and DMARC. We've already published with GDSan email security standard that includes, among lots of other things, DMARC and that's going to become mandatory soon for government. We're already pushing hard to get all the domains in the gov.uk (and in due course other domains that public sector uses) namespace to have DMARC records, which will stop people spoofing gov.uk email addresses. We think 'doing DMARC for government' is a pretty good proof that anyone can implement (if they try), and we'll then be talking to the major industry sectors who have brands with high public trust and confidence to get them to do the same, at scale. That then leaves us with the rest of the deceptive domains and emails (things like ncsc.secure-government.honestly-we-wouldnt-lie.xyz).

The big mail platforms are pretty good at not delivering things that are obviously bad. So, what's in your inbox is generally either OK or possibly dodgy - the problem is most people can't tell the difference. We're talking to industry about a new standard that would present high quality risk information to the end user to help them make a judgement. For example, if you've got an email from a random looking domain that was registered yesterday and it's the first time you've ever seen it, you probably don't want to be opening the attachment that claims to be an invoice.

Basically, we're talking about a reputation system for email domains and addresses, run by the industry. There's a lot of work to do in this area. The hard bit of DMARC and other things like it is the processing of the failure reports and we're centralising that for public sector. The idea is this central processing function (that should only process bad messages that fail some validity check) will be able to pull out things like the sending mail server, any attachments or links in the message, which brand is being abused etc and we can automatically take action with that data.

Go looking for badness and take it down

We've already started some experiments in this area with pioneering UK SME Netcraft. They're off looking for phishing hosted in the UK, webinject malware hosted in the UK and phishing anywhere in the world that targets a UK government brand. When they find it, they ask the hosting provider to take down the offending site. It's surprisingly effective and again generates data we can use. We'll definitely do more in this space.

Filtering DNS to manage impact

Ciaran mentioned this in his Billington speech in September and it generated a lot of attention - including some claims that we'll be using this for things other than reducing the impact of cyber attacks against the UK.

Let's be *really* clear. No-one – not even me – is daft enough to suggest that GCHQ (through the National Cyber Security Centre) should be running the UK's DNS for everyone. Forget the technical stuff for a minute. The real question isn't about what technical stuff happens or who runs what. It's a much more basic question; is it OK for the infrastructure in the UK to allow users to unknowingly access sites that are known to do them harm? I think the answer to that should be 'no'. Let's be clear – this isn't about the nanny state or censorship. A DNS filtering service with an easy opt out for users is a pretty useless censorship tool to be honest – the people you're trying to censor would just opt out and be able to access whatever they want. The way we generate our list of threats will only be concerned with whether the site hosts malware, infrastructure, phishing or other cyber security threats. It won't care about content as seen by users.

Anyway, with GDS, we've partnered with Nominet to build a big anycast recursive DNS service for public sector. That's going to have a response policy zone (RPZ) on it that stops users of the service accessing things we know to be harmful. It'll also generate a load of telemetry data to help us understand what the state of public sector IT is. The RPZ will be created from industry feeds, analytics run on the recursive resolution data, our DMARC and deceptive domain processing and data from NCSC's sovereign capabilities.

Once we've proven the benefit, we'll be talking to ISPs about doing something similar for their residential customers by default. If they want our RPZ feed they can have it. If they want to use other data, that's fine too. And yes, we've thought about malware authors using their own DNS server as a response. Our intent is that, by default, the UK public is protected from things that would do them harm without their knowledge with an easy opt out if individuals want to. That should have a big impact on the scale and effectiveness of a lot of the attacks we see against the UK.

Drive the UK software ecosystem to be better

When you visit a website, your computer tells the site what software it's running (through the HTTP User-Agent string, for example) so the site can tailor content for your specific device. GDS already use that on gov.uk to warn people if their software is out of date and we want to work with the top websites accessed from the UK to do the same. We believe that widespread adoption will help nudge users to upgrade their software, when coupled with other messaging campaigns.

However, there are certain services and groups of users who are so high risk that we think that service differentiation based on software age is appropriate. We haven't got to exactly what this means yet, but as a hypothetical example tax accountants may not be able to submit new returns on their customers' behalf if they consistently use out of date software. Yes, we've thought about the attacker just using new software - there's some simple stuff that can be done to make this a sensible response.

We're really conscious of digital inclusion and will be working with the relevant experts in government and across the key industry sectors to work out how you do something similar for citizens that helps manage the impact of attacks, but doesn't disadvantage those most digitally vulnerable.

Help government get better

We've spent a long time berating both central and local government for not doing enough to secure their stuff. Maybe it's time we more actively helped them and did that in a way that lets them get on with delivering the great public services they provide rather than worrying about deeply technical minutiae. There's a lot to do in this area and we have to be careful that people don't think we're absolving them of their risk management responsibilities, but we have a set of things to try out.

The first one is our 'WebCheck' service. This is a relatively simple web vulnerability scanning service that we'll provide for free to all public sector organisations. It'll give the owners of public-facing sites and services (and only those owners!) a friendly report about any vulnerabilities or misconfigurations in their service and what to do about it. That's in alpha at the moment with 25 local authorities and is being well-received, even at this early stage. We want to build reputation services to help digital service owners make transaction risk decisions. Initially this service will give reputation information for IP addresses connecting to the service and credentials that are used, but we're looking to extend that over time.

We're also looking to experiment on government with novel cyber security techniques and capabilities. One example is a software agent that runs at low privilege on a government workstation and sends metadata back to a central processing facility for analysis. The question is can you detect unknown attacks and exploits using this sort of technique? We don't know yet, but there's some experiments happening to find out.

Encourage innovative alternatives for identity & authentication

Passwords are sub-optimal as an authentication mechanism, but there's not much incentive for industry to take the commercial risk in trying out new stuff. So we hope to stimulate research and development - and eventually a market - in novel ID&A techniques. We'll use government services to trial some new ID&A techniques out, once we've done the work to ensure the security. In the gov.uk Verify platform, we've got a great place where we can try these things out with very little impact on the actual services as each experiment becomes an ephemeral IDP for Verify. That's quite cool. Imagine being able to authenticate to HMG digital services using your face and Windows Hello or your Apple Watch or whatever else gets proposed. The idea here is to promote innovation and adoption of these technologies by de-risking the commercial piece and doing some security design and assurance work up front.

This also is closely aligned to our Secure By Default Partnership, which helps departments trial adoption of new technologies that they otherwise wouldn't see the benefit of.

We're going to provide more help for owners and operators of critical national infrastructure

It's often said that much of the industrial control system technology that's in use isn't designed with security in mind. Well, is that true and what does it mean?

We're thinking about how we provide good engineering and security risk information to CNI operators to help them make better risk management decisions. Probably more importantly, we'll end up with evidence of systemic and specific issues which we can then go tackle with the product vendors. This is probably a piece of work that will go well beyond this spending period and we're not 100% sure what it looks like yet, so there'll be more as we better define the response here.

We're still going to do things to demotivate our adversaries in ways that only GCHQ can do

That's euphemistic by design :-)

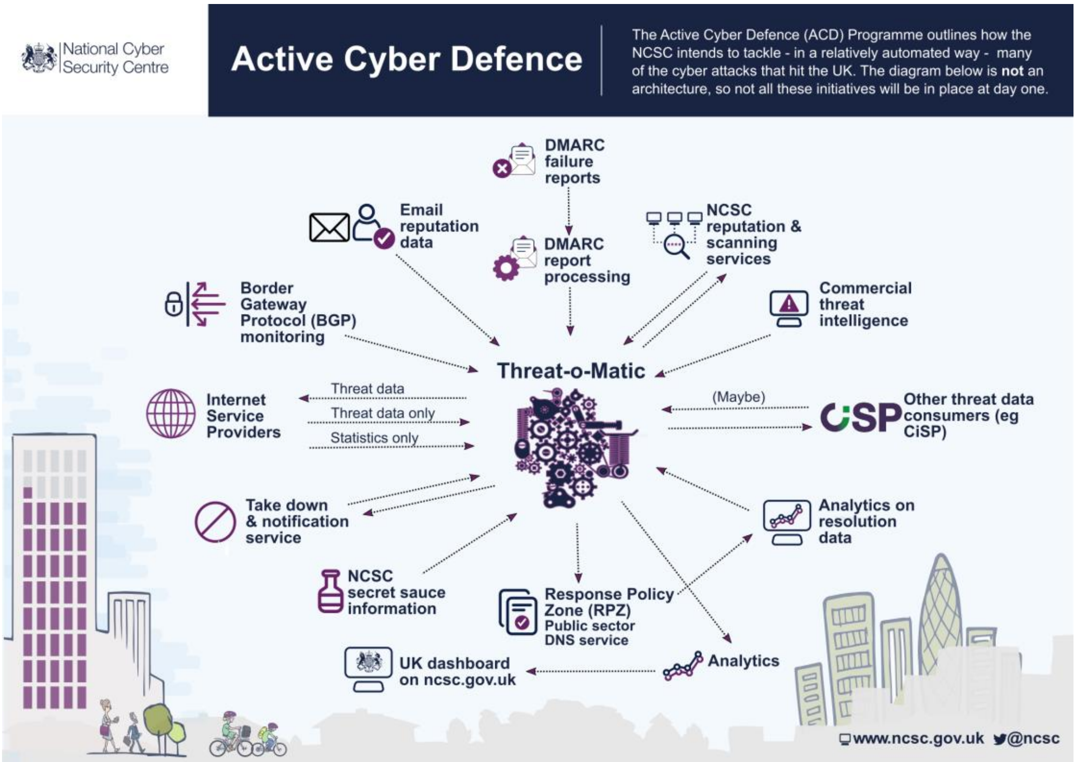

National Cyber Security Centre

Active Cyber Defence

The Active Cyber Defence (ACD) Programme outlines how the NCSC intends to tackle - in a relatively automated way - many of the cyber attacks that hit the UK. The diagram below is **not** an architecture, so not all these initiatives will be in place at day one.

DMARC failure reports
Email reputation data
DMARC report processing
NCSC reputation & scanning services
Border Gateway Protocol (BGP) monitoring
Commercial threat intelligence
Threat-o-Matic
Internet Service Providers
Threat data
Threat data only
Statistics only
(Maybe)
CiSP
Other threat data consumers (eg CiSP)
Take down & notification service
Analytics on resolution data
NCSC secret sauce information
Response Policy Zone (RPZ) Public sector DNS service
UK dashboard on ncsc.gov.uk
Analytics
www.ncsc.gov.uk
@ncsc

Infographic showing the interactions of the Active Cyber Defence programme

All of this will evolve. Some of it will work; some won't. We'll have to respond to adversaries as they respond to our defences. That's probably the new normal though. It's pretty obvious that sitting by and begging people, businesses and government to do better doesn't work, so this is at least trying to do something. It's not perfect - attacks will still get through and there will still be harm - just greatly reduced if we get this right. However, all the different facets of the strategy are designed to work together to be more effective than any one measure alone.

One thing that’s missing in cyber security is unbiased data and evidence. Many of the active defence measures are intended to generate useful data that will help us all understand much better the reality of cyber attacks and the efficacy of the various defences we’ll put in place over the coming years.

The NCSC intends to be a trustworthy and transparent organisation. We need to build that trust and so I intend us to publish as much as possible about what we’re doing and the results. I want to bring some science to cyber security. That needs data, evidence and most importantly peer review.

It’s time to stop talking about what the winged ninja cyber monkeys can do and start countering in an automated way the stuff we see at massive scale that causes real damage to citizens and businesses alike every day. That will include some things that some people class as APTs. However, the intent is to be in a place where the skilled network defender community are free to tackle the really nasty stuff. That’s what the UK’s active defence programme is about.