Missing the Wood for the Trees: Comments on the January 2009 ICANN Working Party Report on FastFlux

Richard Clayton, 15th February 2009

I am deeply unimpressed with the "Initial Report of the GNSO Fast Flux Hosting Working Group" (ICANN, 26 January 2009).

- I do not believe that it provides an adequate (or accurate) description of the problem that is currently being faced;
- The report fails to describe the various roles played by ICANN, the registries and the registrars, and this obscures the policy issues;
- The report does not consider the issues abstractly enough, but narrowly concentrates on some aspects of current criminal behaviour, which could change at a moment's notice;
- The report fails to give any hard data as to the size of the problem, nor how it is changing over time, preventing its actual importance being judged.

In short, the report fails to provide any basis for policy development and should be completely reworked before any other actions are considered.

Before setting out my detailed criticisms, I will explain my credentials for making them. I am presently a security researcher in the Computer Laboratory of the University of Cambridge. I developed one of the earliest integrated packages for dial-up access to the Internet, then spent the second half of the 1990s working at a very large ISP. In 2000 I took the opportunity to study for a PhD, and my thesis on "Anonymity and Traceability in Cyberspace" was completed in 2005. I have remained an academic, and have been studying phishing for over two years, and along with Dr Tyler Moore have published a number of peer-reviewed (and award-winning) papers that examine the lifetime of phishing websites and the various factors that affect this. Indeed, one of our earlier results is quoted on page 26 of the report.

What is the "fast-flux problem"?

The key problem with the report is that it does not really provide a suitably general definition of what fast-flux hosting amounts to – relying instead on characterising existing attacks and deferring to the honeynet.org description of one particular gang's methodology in mid 2007.

The specific distinguisher of a fast-flux attack is that the dynamic nature of the DNS is exploited so that if a website is to be suppressed then it is essential to prevent the hostname resolving, rather than attempting to stop the website being hosted.

The report attempts to characterise fast-flux attacks in many other ways, and gets itself into a considerable tangle by doing so. For example, the "fast" changing of the IP addresses is also performed by legitimate users to provide resilience, load balancing or to meet latency requirements. Yet, there's no need for the criminals to operate this way, if they did a "wide-flux" attack – spreading IP addresses in space rather than time (so that different people accessed different botnet machines) it would be still be necessary to stop the hostname resolving rather than attempt to clean-up all the compromised machines. I will not labour this point, since explaining to the criminals how they might improve their systems is undesirable.

Once it is realised that the actual issue is that of preventing DNS resolution, then most of the current report can be discarded. There are no technical ways to proceed which are effective and avoid collateral damage. Hence it is essential to suspend the domain names, and attention must be paid to the relationships between ICANN, the registries and the registrars.

At present, in most TLDs, decisions about whether to suspend domain names are made by the registrars. Issues only arise when the registrars are not making these decisions correctly, or in a timely manner. The explanations for inaction may come down to ignorance, laziness, incompetence or conspiracy; and the correct way forward in each case will depend on identifying what the issue is and fixing it.

The banking industry, through the Anti-Phishing Working Group (APWG), has already made considerable strides in dealing with ignorance. Laziness must be dealt with by Service Level Agreements (SLAs). It is unacceptable for registrars to fail to provide a 24 × 7 × 365 abuse handling team. Incompetence and conspiracy are a matter for the registries to handle, with ICANN having a role in promoting consistent standards and contractual arrangements.

The difficulty that needs to be addressed is to establish when it is appropriate to suspend a domain name. When a phishing website is hosted on `geocities.com` we want Yahoo! to take action against their customer – but when it is hosted on `id-07i.eu` (a new phishing domain on 14th Feb 2009) we want the domain suspended. Establishing guidelines and principles for how the two cases are to be told apart, and arranging compensation for any innocent domains caught in the cross-fire, would be a useful role for an ICANN report.

What of the technical suggestions in the report?

The report currently contains numerous red herrings.

For example, it discusses the age of the domains that are being used by the criminals, and the quality of their "whois". This is because at present the domains are purchased, limited contact details provided, and the new domain is almost immediately pressed into use. There is no inherent reason for this – the criminals could keep new domains on ice for a long period, and provide genuine contact details of an innocent third party. Hence setting up data sharing arrangement for whois would compromise privacy for the law-abiding and have no impact at all.

Similarly, the suggestion that a TXT record would record the number of name server changes for a domain shows a lack of appreciation of the devolved nature of the DNS. The criminals would just indirect via a server they controlled and/or arrange to fast-flux with names with one extra label (one extra dotted section). Most of the other suggestions made in #5.7 have similar flaws; and they all tackle the symptoms rather than the disease.

How important is fast-flux and is it changing?

The report quotes a 2007 paper written by Tyler Moore and myself where we found that fast-flux was extending phishing website lifetimes. Since then we have done considerable further work. In particular we have found that we receive information from so many sources, that we know about and measure phishing attacks that are not known about by the bank that is being targeted. Since the bank is unaware of the attack, there is no attempt to remove the website and it stays up for a long period.

In a more recent paper looking at data from January 2008 we allowed for this effect and found that the mean removal time for known-about phishing websites on compromised machines was 3.5 hours, with a median of zero hours (viz: half were removed almost instantly). In contrast fast-flux sites had a mean lifetime of 96.1 hours and a median removal time of 25.5 hours. The very high mean reflects a very long tail of some domains being slow to remove. Table 1 summarises our results, the original can be found in [3].

	Period	Sites	Lifetime (hours)	
			mean	median
Child sexual abuse images	Jan-Dec 2007	2585	719	288
Phishing				
Free web-hosting (two brands)	Jan 2008	240	4.3	0
Compromised machines (two brands)	Jan 2008	105	3.5	0
Rock-phish domains (all brands)	Jan 2008	821	70.3	33
Fast-flux domains (all brands)	Jan 2008	314	96.1	25.5
Fraudulent websites				
Escrow agents	Oct-Dec 2007	696	222.2	24.5
Mule-recruitment websites	Mar 07-Feb 08	67	308.2	188
Fast-flux pharmacies	Oct-Dec 2007	82	1 370.7	1 404.5

Table 1: Website lifetimes by type of offending content.

In our most recent work, still under submission [6], we looked at data from the last week of September, finding that ordinary phishing websites now had a mean lifetime of 69 hours, median 26 hours (i.e. fewer are being instantly removed), whereas the situation with fast-flux sites was almost unchanged with a mean lifetime of 97 hours, median 21 hours. So by last September, the banks and their agents were removing the first half of the fast-flux sites more quickly than the first half of the ordinary sites; but the second half lasted longer on fast-flux domains.

We found that fast-flux is the most important mechanism for hosting phishing websites. During the week we studied, 120 domains were used for fast-flux, and 4250 websites were compromised and phishing sites installed. However 68% of the email spam sent was for the fast-flux domains, and the other 32% was for all the other websites put together.

The main lesson to be drawn from this is that fast-flux hosting is prolonging website lifetimes, but the situation is not getting worse, and there are signs of it getting a little better. Because the figures are dominated by the long-tails, improving the response of a handful of specific registrars will immediately result in a considerable improvement.

The above discussion related to phishing. We have also looked at other types of criminal content, some of which is also hosted on fast-flux domains. We do not have sufficient data to compare the take-down times for the different hosting mechanisms – however, we have compared the overall removal times (see the table above). All of these other types of content are being removed significantly more slowly than phishing websites. We conclude that the lack of incentives (and hence the lack of effective action) is far more important than whether or not fast-flux hosting is involved. That is to say, speeding up removal of fast-flux domain names will make very little difference for these other types of content.

Conclusions

The bottom line on fast-flux today is that it is almost entirely associated with a handful of particular botnets, and a small number of criminal gangs. Law enforcement action to tackle these would avoid a further need for ICANN consideration, and it would be perfectly rational to treat the whole topic as of minor importance compared with other threats to the Internet.

If ICANN are determined to deal with this issue, then they should leave the technical issues almost entirely alone – there is little evidence that the working group has the competence for considering these. Attention should be paid instead to the process issues involved, and the minimal standards of behaviour to be expected of registries, registrars, and those investigators who are seeking to have domain names suspended.

I strongly recommend adopting my overall approach of an abstract definition of the problem: The specific distinguisher of a fast-flux attack is that the dynamic nature of the DNS is exploited so that if a website is to be suppressed then it is essential to prevent the hostname resolving, rather than attempting to stop the website being hosted. The working group should consider the policy and practice issues that flow from considering how to prevent domain name resolution; rather than worrying about the detail of current attacks.

Dr Richard Clayton
Computer Laboratory
University of Cambridge
15th February 2009

References

[1] T. Moore and R. Clayton: Examining the impact of website take-down on phishing. In *APWG eCrime*, October 2007, pp. 1–13.

[2] T. Moore and R. Clayton: Evaluating the wisdom of crowds in assessing phishing websites. In *12th International Financial Cryptography and Data Security Conference (FC08)*, Springer Lecture Notes on Computer Science (LNCS), vol. 5143, February 2008, pp. 16–30.

[3] T. Moore and R. Clayton: The impact of incentives on notice and take-down. In *Seventh Annual Workshop on Economics and Information Security, WEIS08*, Dartmouth NH, USA, June 25–28 2008.

[4] T. Moore and R. Clayton: The consequence of non-cooperation in the fight against phishing. In *APWG eCrime*, October 2008, pp. 1–14.

[5] T. Moore and R. Clayton: Evil Searching: Compromise and recompromise of Internet hosts for phishing. In *13th International Financial Cryptography and Data Security Conference (FC09)*, February 2009.

[6] T. Moore and R. Clayton: Temporal correlations between spam and phishing websites. *In submission.*