

On the Effects of Registrar-level Intervention

He (Lonnie) Liu* Kirill Levchenko* Mark Felegyhazi^{o†} Christian Kreibich[†]
Gregor Maier[†] Geoffrey M. Voelker* Stefan Savage*

*Dept. of Computer Science and Engineering
UC San Diego

[†]International Computer Science Institute
UC Berkeley

^oLaboratory of Cryptography and System Security (CrySyS)
Budapest University of Technology and Economics

Abstract

Virtually all Internet scams make use of domain name resolution as a critical part of their execution (e.g., resolving a spam-advertised URL to its Web site). Consequently, defenders have initiated a range of efforts to intervene within the DNS ecosystem to block such activity (e.g., by blacklisting “known bad” domain names at the client). Recently, there has been a push for domain registrars to take a more active role in this conflict, and it is this class of intervention that is the focus of our work. In particular, this paper characterizes the impact of two recent efforts to impact scammers’ use of domain registration: CNNIC’s blanket policy changes for the .cn ccTLD made in late 2009 and the late 2010 agreement between eNom and LegitScript to reactively take down “rogue” Internet pharmacy domains. Using a combination of historic WHOIS data and co-temporal spam feeds, we measure the impact of these interventions on both the registration and use of spam-advertised domains. We use these examples to illustrate the key challenges in making registrar-level intervention an effective tool.

1 Introduction

Operational computer security is reactive, ever responding to new attacks with corresponding interventions. However, while there has been considerable effort in characterizing our own reaction times (e.g., how quickly do phishing sites appear on blacklists), less is understood about how our adversaries react to such interventions in kind. Our paper explores this second question in one particular intervention context: domain name registrars.

Domain names are central to the broad range of Internet scams that seek to attract user traffic to particular Web sites; for example, to advertise goods (e.g., counterfeit pharmacies), install malware (e.g., drive-by downloads), or defraud users of their credentials (e.g., phishing). While there are many different lures used to attract this

traffic — including email spam, search-engine optimization, social network abuse, typo-squatting and so on — virtually all rely on domain name resolution to direct recipients to the site being advertised. For example, a spam email containing the URL `http://toppillls.com` can only be monetized if the user both clicks on the link *and* the domain name can be correctly resolved to the site being advertised.

Given this critical role, domain names have become a key battlefield in the fight between scammers and defenders. As scammers advertise new domains, defenders in turn blacklist them (e.g., blocking inbound mail in mail servers or outbound requests in Web browsers) and pressure ISPs to disable the associated name servers. Unfortunately, blacklisting approaches are limited to protecting only their subscribers and name server takedowns can be technically bypassed (e.g., using double-flux [22] techniques or simply changing the domain’s NS records manually). Thus, many have argued that for domain name interventions to be effective, registrars themselves must take on more oversight responsibility. While there is some controversy over whether registrars should take on this role, it is clear that they are increasingly doing so in response to external complaints.

Given this state of affairs, our goal is to understand the impact that registrar-level interventions have had on scammers use of domain names, how and why scammers have adapted in response, and ultimately how to reason about the use of this approach as a general anti-abuse tool. In particular, we have empirically characterized how domains seen in email spam have been affected by two distinct large-scale interventions — one, a blanket shift in obligations for registering .cn top-level domains (TLD) and the other a more targeted commitment by the registrar eNom to take down counterfeit pharmacy domains. We show that both interventions have real effect, although the pattern and strength of the impact varies considerably. Moreover, we show that, in the current environment, scammers appear quite resilient to

these efforts; both changes lead to displacements in domain use (e.g., between TLDs or between registrars) but not to appreciable reductions in overall activity. Finally, we discuss the challenges that must be overcome to make registrar-level takedown a more effective tool.

The remainder of this paper is structured as follows. Section 2 provides an overview of domain-oriented interventions, how the registration ecosystem operates, and explains the particular interventions characterized in this paper. We describe our data sources and methodology in Section 3 followed by analyses of each intervention in Sections 4 and 5. We conclude with a discussion of the key challenges in using registrar-level intervention as an effective anti-abuse tool.

2 Background and Related Work

Domain names are central to virtually all Internet applications (Web, mail, instant messaging, etc.) as they provide both a human-readable name space as well as a distributed resolution service. There are two key processes that govern the use of domain names: resolution and registration. Name resolution is widely understood in the security community and involves iteratively querying name servers, starting with the root server for the top-level domain, until an associated network address can be found. However, domain registration — the process by which domain names are reserved and installed into these top-level servers — is less widely appreciated and so we review it briefly here.

2.1 Domain registration

Today there are roughly 20 generic top-level domains (such as .com or .org) and roughly 250 country-code top-level domains (such as .ru or .cn). Each of these TLDs is overseen by a singular registry (e.g., VeriSign for .com, CNNIC for .cn, etc.) that manages the right to install, modify and remove names in the associated registry database and TLD root name server zone file. These registries in turn may contract with one or more registrars to sponsor domains under the particular TLD extension. For example, cnn.com is sponsored by CSC Corporate Domains, who acts as a registrar for .com under contract to Verisign who runs the .com registry. Registrars are required to be accredited by ICANN (today there are roughly 1,000 entities accredited to act as registrars) and required to meet the distinct contractual obligations of the associated registries they sponsor for (the terms of which can vary considerably between registries). In turn, registrars may offer their services on a retail basis to the public or to resellers on a wholesale basis (each relationship including its own contractual requirements). The precise pricing terms and payment flow (i.e., to re-

sellors, registrars, registries and ICANN) varies between registries and registrars and is generally not transparent to the public.

Finally, the relationship between a given domain name and its registrar can be dynamic over time. The registrar used to sponsor a particular name is known as the “designated registrar” for that name and is the only registrar allowed to modify or delete the entry from the associated registry database (technically accomplished through the Registry Registrar Protocol [8, 9] or Extensible Provisioning Protocol [7, 6]). However, the owner of the domain may elect to transfer the name between registrars subject to both broad policies set by ICANN as well as operational policies set by the registrars themselves.

2.2 Domain-level intervention

Domains are widely used by scammers to convert user “clicks” (e.g., in a browser) into Web traffic directed to a particular site. This activity includes phishing scams, email spam, search engine spam, blog spam, twitter spam, kinds of click fraud and so on. These uses have in turn engendered a range of interventions, each seeking to undermine the use of such domains and thus prevent a given scam from succeeding.

By far the most common form of abuse intervention is blacklisting, a process in which each newly discovered offending item is “broadcast” to subscribers. Blacklisting schemes differ based on the nature of the items being listed (i.e., IP addresses, domain names or full URLs), how quickly they discover new items, how they are then used to mitigate exposure (e.g., to filter inbound email, outbound Web requests or search engine results), and how widely they are deployed (e.g., users of the Firefox browser vs. all email recipients at hotmail.com). IP-based blacklists, particularly those used for identifying spam senders, date back to 1997 and are in widespread use today. However, the size and sophistication of modern botnet operations has made it challenging to track the full set of abusive senders. Indeed, one recent study of such lists found false negative rates between 35–98% [21].

Thus, a newer approach focuses on *advertised* Web sites and blacklists individual domains or even full URLs. Well known domain-based blacklists include Spamhaus’ Domain Block List (DBL), the URIBL and the Google Safe Browsing lists (focused on phishing and malware sites). Similar studies of effectiveness have demonstrated high “catch rates” for such mechanisms (e.g., as high as 90% for phishing sites [14]). Kreibich et al. [11] showed that in 2008 domains used for advertising pharmaceuticals via the Storm botnet were registered in batches weeks before use, blacklisted within 18 minutes on average, and moved from .cn to .com and .eu.

On the other hand, Sheng et al. reported that it can also take hours for new sites to appear on blacklists [20]. To address this issue, several recent efforts have explored “predictive” blacklisting, a technique premised on the observation that particular features or combinations of features are highly correlated with the eventual outcome of appearing on a blacklist [5, 15, 19]. Related to this paper, these approaches commonly observe that the registrar and registry can be a particularly salient feature since, for example, spam-advertised domains tend to be disproportionately sponsored by a few registrars.

However, such approaches only protect the subset of potential victims who are blacklist subscribers, still allowing a scam site to harm others. An alternative intervention which can protect all potential victims is “take-down,” whereby pressure is exerted on service providers to shut down the offending Web servers, name servers, or domain name registrations. The best-known study of such approaches is due to Moore and Clayton [17, 18, 16] who have focused primarily on the effectiveness of take-down at addressing advertised phishing sites. Among their findings, they show that fast-flux techniques (where a scammer makes use of a range of Web servers, and possibly name servers as well, for a given domain name) are successful at significantly increasing takedown response time (suggesting that hosting providers may be more responsive than domain registrars).

2.3 Registrars and takedown

Because of their central role, registrars have come under increasing pressure to aid in takedowns. However, such requests can present a quagmire for registrars. While they have a vested interest in protecting their brand from negative publicity, at the same time they are motivated to maximize registration revenue. Moreover, some argue that what is considered “abuse” can vary from one locale to another, sometimes placing registrars in the uncomfortable position of making judgements that may be at odds with their own economic incentives. Consequently, policies (both explicit and implicit) between registrars can vary considerably and thus scammers have tended to migrate to those TLDs and those registrars that are both cheapest and viewed as most “friendly” to their activities [2, 10]. For example, in the widely read “Rogues and Registrars” report from LegitScript and KnujOn [12] the authors document that far more domains hosting counterfeit pharmacies were registered through eNom than any other registrar and that eNom was one of the key registrars that were unresponsive to complaints about such domains.

In the remainder of this paper we will examine the concrete effects seen when registration policies change and how scammers react in turn.

3 Data Sources

In this paper we make use of two kinds of empirical data involving scam domains: (1) their appearance in URLs in spam and on blacklists, and (2) the management activity of such domains as recorded in WHOIS databases and DNS zone files. We describe here how we constructed such datasets to cover two time periods surrounding key policy actions by CNNIC (the registry for .cn) and eNom (a major wholesale and retail registrar).

3.1 Spam-advertised domain data

We focused on spam-advertised domains to represent instances of scammer domains impacted by registrar actions. Between mid-November 2009 to mid-March 2010, we combined data from three spam feeds, two of which use MX honeypots (capturing all SMTP traffic to otherwise unused domains with an active MX record) while the other is driven by individual email account honeypots distributed over live domains. We also collected contemporaneous data from the URIBL,¹ a popular domain blacklist used primarily to filter email spam. We reconstructed the URIBL “black” list (“domain names belonging to and used by spammers”) as well as the “gold” list (“proactive black listings”) from daily snapshots. In these feeds our main concern is the TLD of domains used in URLs in spam messages, and it is this feature that we track across time.

For the measurement period between August and October of 2010 inclusive, we also added four additional URL feeds from third-party partners (including multiple commercial anti-spam providers) as well as URLs harvested from our own managed botfarm environment. Additionally, for this data we were also able to “classify” the URLs by visiting a subset of 6 million (biased to select as many different domains as possible), downloading their Web content, and then using manually-generated regular expressions to identify the type of storefront (i.e., pharmacy, replica, software) as well as the *particular* affiliate program.² Focusing on .com domains in this measurement period, we identified 14,286 domains associated with on-line pharmacies.

3.2 Registration data

We also collected information about domain registrations over the same periods. For the 2009 measurement period, we randomly sampled 1,000 of the spam-advertised domains and obtained their WHOIS [3] records, extracting the domain registration date. We relied on the Domain-

¹<http://www.uribl.com>

²This dataset and analysis are part of a much larger effort focused on analyzing how spam is monetized [13].

Tools commercial service,³ which enabled us to obtain historic WHOIS data from the period in question but effectively limited the number of domains we could query.

For the 2010 measurement period, we focused specifically on .com domains, which are very popular among spammers. We gathered daily snapshots of the DNS zone files for .com. In addition, we collected WHOIS records for all .com domains up to September 20th, 2010 and WHOIS records for all .com domains after September 20th, 2010 whose start date had changed in the zone files. Finally, for each of the 54,142 domains classified in our spam analysis that used eNom as a registrar, we again used the DomainTools service to gather historic WHOIS information. Together, these data sources allow us to identify when a domain is first registered, when it is inserted and removed from the zone file (used to control accessibility), when its name server is changed, when it is transferred between registrars, and when it is put in a “client hold” state (used to prevent transfers).

4 CNNIC Policy Change

We first examine the impact of a late 2009 registry-level change, one that impacted all registrars for the .cn TLD and hence all potential .cn registrants as well.

4.1 Description

The .cn ccTLD is overseen by the China Internet Network Information Center (CNNIC). As the registry of record, CNNIC in turn accredits registrars who are authorized to sell .cn domains. For much of 2007 and 2008, .cn domains were made available for a single Yuan (\approx US\$ 0.10–0.15) as part of CNNIC’s “Experience .cn Domain Name for One Yuan” Campaign.⁴ Due to their low cost, broad availability and minimal documentation requirements, .cn domains became popular among spammers, particularly for quasi-random “throw away” domains used to foil blacklists. In response, CNNIC issued significant new regulations — taking effect on December 14th, 2009 — requiring formal paper documentation and validation, limitations on customers of non-Chinese registrars and limitations on individual registrations.⁵ In addition to the challenges placed on scammers by these overheads, the policy changes also increased the price offered with the lowest reported minimum prices at the time rising to 69 Yuan (\approx US\$ 10).⁶

³<http://domaintools.com>

⁴<http://www.cnnic.net.cn/html/Dir/2007/12/27/4953.htm>

⁵<http://www.cnnic.net.cn/html/Dir/2009/12/12/5750.htm>

⁶<http://www.techsecuritychina.com/2010/01/07/9093-price-increase-for-cn-chinese-website-domain-names/>

4.2 Analysis

Given this change in effective per-domain registration cost, we next consider how this new burden affected spammer behavior.

We start by examining new spam domain registrations over time, illustrated in Figure 1(a). Using the WHOIS registration dates of the sampled domain set (Section 3.2) we split the sampled domains into three groups: .cn domains, .ru (Russia) domains (for reasons that will become clear), and all .com, .net, and .org domains combined. Each curve shows the 7-day average of the number of domains first appearing on that day. Each domain is counted once — on the day it first appears — and we use the average to smooth the very bursty daily values. The two dashed vertical lines denote the CNNIC policy change announcement (December 11, 2009) and enforcement (December 14, 2009). The x -axis spans four months from November 1st 2009 to March 15th 2010, covering the periods leading up to and following the date of the CNNIC policy change. Finally, the y -axis shows the number of distinct sampled domains in a TLD group scaled by the total number of domains relative to the size of our sample (for visual consistency with the remaining graphs, which do not show sampled data).

From this graph we see that the CNNIC policy change had a swift and dramatic effect on spammer domain registration. Spammers stopped registering the now more expensive domains in .cn *immediately*, eventually moving their business wholesale to .ru where spam domain registrations rose to the volumes previously observed for .cn.⁷ Against a steady backdrop of new spam domains registered in .com/net/org, a rise in .cn domain registrations in late November 2009 quickly fell around the CNNIC policy change announcement, nearly disappearing days later. In contrast, spammers in our feeds rarely registered .ru domains before the CNNIC change, registered a burst of .ru domains immediately after abandoning .cn in mid-December 2009, and then steadily reached their original registration volume — now in .ru domains — starting in mid-January 2010.

The flight of domain registrations from .cn to .ru is also clearly apparent in the use of such domains in spam messages. Figure 1(b) shows the number of distinct domains that appear in URLs over time in our spam feeds, while Figure 1(c) provides the same information from the perspective of spam domains appearing in the URIBL blacklist. In both cases, we show curves for the same groups of TLDs over the same time period as in Figure 1(a), and again use a 7-day moving average to smooth daily bursts. The large volumes of .com/net/org do-

⁷We note that a large portion of the rapid reduction in new .cn domain registrations actually *precedes* CNNIC’s policy announcement, suggesting that some may have had inside knowledge of the imminent change.

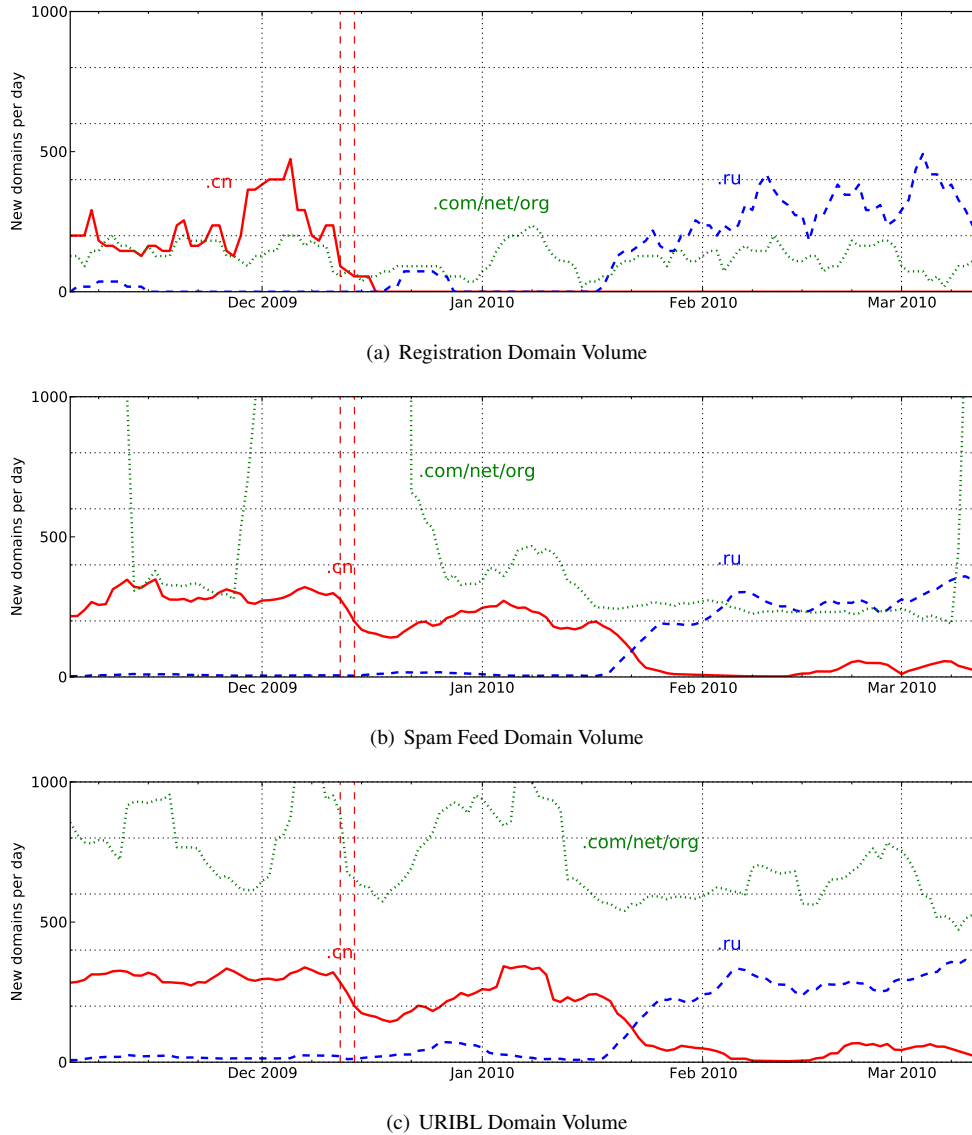


Figure 1: TLDs of spam domains over time: (a) when the domains were registered (estimated from a sample of 1,000), (b) when they appear in honeypot spam feeds, and (c) the distribution of TLDs appearing in the URIBL blacklist. The two dashed vertical lines denote the CNNIC policy change announcement (December 11, 2009) and enforcement (December 14, 2009).

mains appearing in spam in early November 2009, much of December 2009, and the middle of May 2010 correspond to surges in spam from the Rustock botnet⁸ specifically crafted to overwhelm and poison domain blacklists. Rustock flooded the Internet with spam using random, typically unregistered .com domains (note the lack of corresponding surges in registrations in Figure 1(a)).

The data shows that the CNNIC policy change had a dramatic effect on the *use* of .cn domains in spam as well as new registrations. In the span of just one week late in January 2010, spammers switched from using

their .cn domains in spam URLs and replaced them with .ru domains instead. However, the date of transition between registration and use are quite different; .cn domains continued to be used steadily in spam until six weeks after the CNNIC change, even though .cn registrations ceased immediately afterwards. This behavior suggests that these spammers had inexpensive .cn domains stockpiled, and it took six weeks for them to exhaust such .cn domains registered before the CNNIC change.

However, in spite of these dramatic effects, the overall impact is one of displacement — driving spam domain registrations out of .cn and into .ru — rather than

⁸<http://garwarner.blogspot.com/2010/09/random-pseudo-urls-try-to-confuse-anti.html>

reduction; the total volume of spam after this action is unchanged. Thus, while this action undoubtedly imparts *some* additional costs to the spammer (e.g., slightly increased per-domain registration cost or operational overhead), in general the combination of low switching cost and readily available TLD alternatives undermine the global benefit of such changes.

5 LegitScript–eNom Agreement

In this section we attempt to quantify the effects of the agreement (and subsequent action) between LegitScript, an “Internet pharmacy verification service,” and eNom, a major registrar, to identify and terminate domains being used by “rogue” on-line pharmacies.

5.1 Description

eNom, a division of Demand Media, is one of the largest commercial registrars on the Internet, sponsoring domains for a wide variety of TLDs and commanding over 8% of the global domain registration market share (according to data compiled by WebHosting Info⁹). However, a number of studies demonstrated that scammers also made heavy use of eNom’s services and the registrar received significant criticism for being unresponsive to complaints about this subset of their sponsored domains [1, 10]. One of eNom’s most visible critics was LegitScript, an organization focused on addressing the problem of “rogue” Internet pharmacies. On September 21st, 2010 eNom entered into an agreement with LegitScript “by which LegitScript will assist eNom in identifying customers who are violating eNom’s terms of service by operating online pharmacies in violation of U.S. state or federal law” [4].

First anecdotal reports of eNom domains being seized as a result of this agreement appeared on September 23rd on Web forums covering pharmacy affiliate programs. A second, larger wave several days later on September 27th sent many affiliates into a panic. Reports of domain seizures continued throughout October, with registrants reporting that their domains were placed into “*clientHold*” state¹⁰ with the notification that “the domain will remain on hold and may not be transferred.”¹¹

5.2 Analysis

To better understand the effect of this action on the spam-advertised pharmaceutical business, we first set out to es-

⁹http://www.webhosting.info/registrars/reports/total_domains/enom.com?ob=gs&oo=asc

¹⁰A *clientHold* status value means that the domain is not published in the appropriate zone file. [6]

¹¹<http://www.gofuckbiz.com/showpost.php?p=382053>

timate the scope of the action. Our first challenge was to determine what happened. Taking the anecdotal forum reports as our starting point, we turned to our .com WHOIS data to determine if any eNom domains were placed into *clientHold* state on September 27th. As described in Section 3.2, from September 20th we queried the WHOIS records of every domain which had a name server change in the zone file the previous day. Thus, if an active domain (with at least one name server) had been placed into *clientHold* state, the change would be reflected in the .com zone file, allowing us to identify and obtain the WHOIS record for all such domains.

5.2.1 Estimating Scope

Figure 2 shows the number of eNom domains being placed into *clientHold* state between September 20th, 2010 and January 28th, 2011. There are notable peaks, including a peak on September 27, 2010 when 332 eNom .com domains were put into *clientHold* state. On manual inspection, we found most of these domains to contain pharmacy-related keywords (e.g., “pills,” “drug,” “rx”) or the name of a pharmaceutical (e.g., “viagra,” “cialis,” “sildenafil”). In fact, of the eight peaks exceeding 200 seized domains per day, for all two of the peaks 60% or more of the domain names contained such pharmacy-related keywords. The domains seized on October 4th and January 28, however, contain virtually *no* such keywords in their names, suggesting that the domain seizures on these two days are qualitatively different.

In support of the hypothesis that most these peaks corresponded to an action based on the LegitScript agreement, we were able to find posts on affiliate programs reporting eNom domain seizures on October 13th and December 22nd.

While we do not know the true number of domains reported by LegitScript to eNom nor all of the dates eNom took action, based on the observation that the offending domains were placed into *clientHold* state, we can upper-bound the number of .com domains affected. That is to say, we know that 7,110 eNom .com domains were placed into *clientHold* state between September 20, 2010 and January 28, 2011, of which a subset are domains reported to eNom by LegitScript.

5.2.2 Comparative Statistics

As described in Section 3.1, we also have at our disposal a set of 14,286 domains that have been advertised in spam between September 20th and October 31st, and which we have positively confirmed to lead to on-line pharmacy storefronts. Of these domains, 1,223 were registered through eNom. What effect did the LegitScript–




Figure 2: Number of eNom domains being placed into *clientHold* state daily from September 20th, 2010 to January 28th, 2011.

eNom action have on these domains?

Of these 1,223 pharmacy-labelled eNom .com domains, 825 (67%) were placed into *clientHold*, most of them during the action on December 22nd. If we optimistically credit LegitScript with all 7,110 eNom domain seizures between September 20, 2010 and January 28, 2011, then this amounts to a significant fraction of the labelled eNom .com domains associated with spam-advertised pharmacies.

What was the fate of the remaining 398 (33%) labelled eNom .com domains? They fell into three groups: 190 (16%) are still active, 57 (5%) are no longer listed in the .com zone file, and 151 (12%) transferred to another registrar. 100 of these 151 transfers all moved to another registrar, Realtime, on October 27–29. In fact, all of these transferring domains are associated with the RX-Promotion affiliate program. We hypothesize that the transfer was a pre-emptive move by RX-Promotion to avoid further eNom domain seizures.

It seems, then, that these domain seizures have had a noticeable effect. Unfortunately, only 1,223 (9%) of the 14,286 labelled pharmacy domains were registered with eNom. Of the 13,063 non-eNom domains, 6,952 (53%) are still active while 6,111 (47%) appear to have been removed from the .com zone file before their WHOIS expiration date.

6 Discussion

Domain names are central to attracting Web traffic and thus central to most Internet scams as well. Email spam, blog spam, search spam, phishing, drive-by downloads, click fraud and so on all generally require their victims to resolve domain names provided by scammers. It is this critical role that also makes domain-oriented defenses and interventions particularly attractive; if this domain resolution can be disrupted, then so is the underlying

scam. The key role that registrars and registries play in administering domains makes them an attractive “choke point” for shutting down such use and thus there has been increasing pressure for their participation in “domain takedown” operations.

In this paper, we have tried to characterize the effects of such interventions by examining how two particular policy actions have impacted the domains found in spam. In each case we see concrete effects, but also witness that the current ecosystem provides spammers with ample room to adapt.

The CNNIC example demonstrated that spammers are price-sensitive and that significant changes in overhead can be quite effective at evicting spammers from a TLD. However, because alternative TLDs (and registrars) are readily available and switching cost is low (virtually zero for new domains), spammers simply transitioned to the next lowest priced TLD (.ru). Moreover, because their existing .cn domains remained operational, spammers retained their existing domain investment and thus were able to buffer any financial impact.

By contrast, the eNom intervention was more selective and reactive. However, here too we witness spammer registrations simply being displaced from one registrar (eNom) to another (e.g., Realtime) — either through transfers or new registrations. However, unlike the situation with CNNIC, a large fraction of domains are killed and never transferred — representing real losses to spammers. We hypothesize that the subset of domains transferred represent those that for which spammers have invested sufficient value (e.g., they have been campaigned for search engine optimization) to warrant the delay, overhead and visibility of the transfer process, while the majority of domains have little value beyond their cost. Finally, roughly 15% of pharmaceutical domains persist over much of the measurement period without any action taken, highlighting another challenge in takedown

— identifying all the domains to be targeted.

Synthesizing these findings, we argue that while increasing the minimum *global* price for domains is likely a proactive drag on spammers (albeit imposing an additional cost on global domain users), simply changing local pricing is unlikely to have much operational impact. By contrast, reactive takedowns can be effective, but face multiple challenges. First among these, their reaction time must be shorter than the domain use lifetime (i.e., how long the domain is actually active in spam). For example, in one recent study of the Storm botnet, Kreibich et al. document that the average lifetime of spam advertised domains is under 6 days [11]. Thus, the existing lifetime is already “priced in” to the business model and any takedowns that take more time to effect may do little to undermine spammer economics. Second, the direct economic impact of takedown is limited by the availability of alternatives and quick provisioning (reducing the need to “warehouse” domains and risk capital) and thus takedowns are unlikely to reduce spam volumes so long as there are a range of non-participating registrars. Finally, takedowns require highly focused organizations that work diligently to discover new domains and also can establish credibility or effective pressure across registrars. Thus, while LegitScript provides this energy for domains advertising counterfeit pharmaceuticals, it has little impact on domains used to sell counterfeit software, replica watches, fake anti-virus and so on.

References

- [1] J. Armin. Demand Media—eNom: the World’s #1 Bad Host and Abusive Registrar. Technical report, HostExploit, Aug. 2010.
- [2] J. Armin. KnuiOn’s response to eNom statement. <http://hostexploit.com/blog/4-current-events/3514-knujons-response-to-enom-statement.html>, June 2010.
- [3] L. Daigle. WHOIS protocol specification. RFC 3912, The Internet Society, Sept. 2004.
- [4] Demand Media. eNom and LegitScript LLC announce agreement to identify customers operating illegal online pharmacies. <http://www.businesswire.com/news/home/20100921005657/en/eNom-LegitScript-LLC-Announce-Agreement-Identify-Customers>, Sept. 2010.
- [5] M. Felegyhazi, C. Kreibich, and V. Paxson. On the potential of proactive domain blacklisting. In *Proceedings of the Third USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET)*, San Jose, CA, USA, April 2010.
- [6] S. Hollenbeck. EPP domain name mapping. RFC 5731, IETF Trust, Aug. 2009.
- [7] S. Hollenbeck. Extensible provisioning protocol (EPP). RFC 5730, IETF Trust, Aug. 2009.
- [8] S. Hollenbeck and M. Srivastava. NSI registry registrar protocol. RFC 2832, The Internet Society, May 2000.
- [9] S. Hollenbeck, S. Veeramachaneni, and S. Yalamanchilli. VeriSign registry registrar protocol (RRP) version 2.0.0. RFC 3632, The Internet Society, Nov. 2003.
- [10] KnuiOn.com, LLC. Internet Security Report: Audit of the gTLD Internet Structure, Evaluation of Contractual Compliance, and Review of Illicit Activity by Registrar. http://www.knujon.com/knujon_audit0610.pdf, June 2010.
- [11] C. Kreibich, C. Kanich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage. Spamcraft: An inside look at spam campaign orchestration. In *Proceedings of the Second USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET)*, Boston, USA, April 2009.
- [12] LegitScript and KnuiOn. Rogues and Registrars—Are some Domain Name Registrars safe havens for Internet drug rings? <http://www.legitscript.com/download/Rogues-and-Registrars-Report.pdf>, April 2010.
- [13] K. Levchenko, N. Chachra, B. Enright, M. Felegyhazi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, H. Liu, D. McCoy, A. Pitsillidis, N. Weaver, V. Paxson, G. M. Voelker, and S. Savage. Click Trajectories: End-to-End Analysis of the Spam Value Chain. In submission.
- [14] C. Ludl, S. McAllister, E. Kirda, and C. Kruegel. On the effectiveness of techniques to detect phishing sites. *Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, pages 20–39, 2007.
- [15] J. Ma, L. Saul, S. Savage, and G. Voelker. Beyond blacklists: learning to detect malicious web sites from suspicious URLs. In *Proceedings of the 15th ACM SIGKDD Intl. Conference on Knowledge Discovery and Data Mining*, pages 1245–1254. ACM, 2009.
- [16] T. Moore and R. Clayton. Examining the impact of website take-down on phishing. In *eCrime Researchers Summit*, pages 1–13, 2007.
- [17] T. Moore and R. Clayton. The consequence of non-cooperation in the fight against phishing. In *eCrime Researchers Summit*, pages 1–14, 2008.
- [18] T. Moore and R. Clayton. The impact of incentives on notice and take-down. *Managing Information Risk and the Economics of Security*, pages 199–223, 2008.
- [19] P. Prakash, M. Kumar, R. Kompella, and M. Gupta. Phishnet: predictive blacklisting to detect phishing attacks. In *Proc. INFOCOM*, pages 1–5. IEEE, 2010.
- [20] S. Sheng, B. Wardman, G. Warner, L. Cranor, J. Hong, and C. Zhang. An empirical analysis of phishing blacklists. In *Sixth Conference on Email and Anti-Spam (CEAS)*, 2009.
- [21] S. Sinha, M. Bailey, and F. Jahanian. Shades of grey: On the effectiveness of reputation-based “blacklists”. In *3rd International Conference on Malicious and Unwanted Software (MALWARE)*, pages 57–64. IEEE, 2008.
- [22] The HoneyNet Project. Know Your Enemy: Fast-Flux Service Networks. <http://www.honeynet.org/papers/ff/fast-flux.pdf>, July 2007.