

XXXtortion? Inferring Registration Intent in the .XXX TLD

Tristan Halvorson
trhalvorson@cs.ucsd.edu

Stefan Savage
savage@cs.ucsd.edu

Kirill Levchenko
klevchen@cs.ucsd.edu

Geoffrey M. Voelker
voelker@cs.ucsd.edu

Department of Computer Science and Engineering
University of California, San Diego

ABSTRACT

After a decade-long approval process, multiple rejections, and an independent review, ICANN approved the xxx TLD for inclusion in the Domain Name System, to begin general availability on December 6, 2011. Its sponsoring registry proposed it as an expansion of the name space, as well as a way to separate adult from child-appropriate content. Many independent groups, including trademark holders, political groups, and the adult entertainment industry itself, were concerned that it would primarily generate value through defensive and speculative registrations, without actually serving a real need. This paper measures the validity of these concerns using data gathered from ICANN, whois, and Web requests. We use this information to characterize each xxx domain and infer the registrant's most likely intent. We find that at most 3.8% of xxx domains host or redirect to potentially legitimate Web content, with the rest generally serving either defensive or speculative purposes. Indeed, registrants spent roughly \$13M up front to defend existing brands and trademarks within the xxx TLD, and an additional \$11M over the course of the first year. Additional evidence suggests that over 80% of annual domain registrations are for purely defensive purposes and do not even resolve.

Categories and Subject Descriptors

C.2.m [Computer Communication Networks]: Miscellaneous;
K.4.1 [Computers and Society]: Public Policy Issues—*regulation, use/abuse of power*

Keywords

DNS; TLDs; Internet Economics; Abuse; Registration Intent

1. INTRODUCTION

Domain names are a unique and exclusive resource. Once a name is registered it is no longer available for others. Thus, names that correspond to popular trademarks, brands, proper names, common words, acronyms and so on, can accrue tremendous commercial value for those who hold them. On the theory that expand-

ing the name space would therefore increase competition and create more naming opportunities, ICANN has created a number of new "top-level domains" (TLDs) in recent years (e.g., biz, aero, xxx, etc.) and is today considering a range of bids to dramatically increase the domain name space further. However, there has been considerable debate about the extent to which new TLDs will actually serve a real need, or whether they will instead exploit existing brand holders who will feel compelled to protect their marks on a defensive basis.

In this paper, we examine this question through the unique lens of one particular TLD: xxx. The xxx TLD was created to serve the needs of the adult entertainment industry and its approval by ICANN was a matter of great controversy (both from within the adult industry and without). Critics argued that the TLD did not serve any real need and that its creation would largely be monetized via speculators and defensive registrations. With the benefit of hindsight (xxx opened for general availability in December of 2011) we examine this question empirically. While it is impossible to measure the intent of a registrant directly, we argue that such intent can be inferred through how each of their domains are used (i.e., does it resolve?, serve content?, what kind of content?, etc.)

By combining public ICANN records, analysis of the xxx zone files, active whois requests and crawling the content of live sites, we have constructed what we believe is the most comprehensive picture to date of how xxx is used in practice. We find that almost half of all xxx domains do not appear in the zone file at all, and of the approximately 107k domains that do resolve, three-fourths are objectively defensive registrations. A significant fraction of the remainder point to parked sites, suggesting speculative use. Finally, our calculations suggest that the clear majority of registration revenue, both during the pre-registration period and for ongoing operation, is driven by defensive concerns rather than entrepreneurship. While the xxx TLD situation is clearly unique, we believe it is a telling data point as the Internet community considers a range of far broader expansions in the name space.

2. BACKGROUND

The Domain Name System (DNS) is widely used to provide memorable identifiers for Internet hosts and services. At a technical level, DNS simply translates human-readable strings to machine addresses. In the Internet's modern incarnation, the DNS is overseen by the Internet Corporation for Assigned Names and Numbers (ICANN) who in turn licenses individual *registries* to operate distinct name spaces. Thus, each registry maintains exclusive management rights over its names, and is responsible for licensing *registrars* (who sell registration services on a retail basis) and for ensuring authoritative name resolution for such names.

2.1 Domain Value

DNS technically treats all names equally, but in practice different names accrue significantly different value. Since names are human readable, by design, certain words more naturally resonate with particular commercial categories (e.g., `books.com`) and others alias existing trademarks (e.g., `mcdonalds.com`). This aspect, combined with the “first come, first served” model of domain registration, has effectively turned domains into property rights with the attendant problems thereof. Thus, some popular domains can re-sell for millions of dollars [21], and a whole cadre of “domainers” purchase choice domains speculatively with no intent to use them. Still others, in a practice called “typo-squatting”, register domains that are lexically similar to popular domains to poach their traffic from users who mistype [16]. In response to these behaviors, ICANN introduced the Uniform Domain-Name Dispute-Resolution Policy (UDRP) in 1999 to resolve domain disputes between trademark holders and domain name speculators through non-binding arbitration. However, the UDRP process can be both time consuming (six weeks or more) and costly (typically \$1,000 plus any attorney’s fees) and thus many brand holders register domains on a prophylactic basis simply to minimize risk.

2.2 Top-Level Domains

The DNS is administered hierarchically, split first into registry-managed Top-Level Domains (TLDs), which may then allow direct registration or broken up further in a TLD-specific manner. In its original incarnation, the DNS included seven distinct TLDs. Of these, four (`mil`, `edu`, `gov`, and `arpa`) require users to have specific registration purposes (e.g., the `edu` registry, EDUCAUSE, only registers names for educational institutions). The remaining three domains, `com`, `net`, and `org`, were also intended for different purposes, but today all allow general registrations, giving them the designation of *generic TLDs* (gTLDs). In principal, this provides three times as many naming opportunities, but it is not uncommon for brand holders to replicate their marks in each TLD (e.g., `microsoft.com`, `microsoft.org` and `microsoft.net` are all owned by Microsoft Corporation).

ICANN and its predecessors have since approved the addition of many other TLDs, including a full set of country code TLDs (ccTLDs), and alternative gTLDs like `biz`, `aero` and `info`. Each introduction of a new TLD includes a clear purpose. For instance, ccTLDs allow countries to operate their own local DNS trees, providing independent sovereign control on part of the name space, while `biz` was meant as an alternative to `com` for businesses with competing trademarks in different industries [3].

Not surprisingly, custom and convention have made some TLDs more popular than others. For example, `com` had 108 million domains registered as of October 2012, including 57% of the Alexa Top 100,000 Web sites. In contrast, `biz` had only 2.3 million domains at the same time, and covers only 0.3% of the Alexa ranks. Since different TLDs serve different purposes, the sheer number of registrations is not an ideal metric of success. For example, since `edu` only serves educational institutions it is artificially constrained in size. Yet it serves its function effectively and is widely supported by its stated community. Instead, another way to evaluate the “health” of a TLD is the extent to which its registrations host unique content (under the assumption that the names are therefore being used to provide direct access to this information). A TLD that instead has a high fraction of speculative or defensive registrations offers much less value to the Internet community. However, crisply determining intent can be difficult in practice.

The `xxx` TLD provides an excellent case study in the current benefits of a new TLD, as the rollout periods make determining

registration intent significantly easier than for other TLDs. In the remainder of this section we review the history of the `xxx` TLD, its complex multi-phase rollout process and the subsequent legal challenges and criticisms faced by its registry.

2.3 The `xxx` TLD

ICM Registry created the first proposal for a `xxx` TLD in 2000, which ICANN rejected later that year. In their report, they cited the complex political climate and the lack of need for such a TLD, especially in their initial rollout [18]:

ICM Registry’s application for an `xxx` TLD does not appear to meet unmet needs. Adult content is readily available on the Internet. To the extent that some believe that an `xxx` TLD would segregate adult content, no mechanism (technical or non-technical) exists to require adult content to migrate from existing TLDs to an `xxx` TLD.

Despite these points, they also indicated their willingness to reconsider the proposal in the future.

After ICANN rejected their third and newest proposal in 2007, ICM Registry appealed for an independent review of their application. The review found in favor of ICM Registry but left the ultimate decision up to ICANN, urging it to “evaluate the continued uncertainty and risk associated with its decision, including risks to ICANN resulting from potential legal actions” [9]. In 2010, against the recommendation of its Governmental Advisory Committee (GAC), ICANN accepted ICM Registry’s newest sponsored proposal, to be rolled out in 2011 for the express use of the “adult entertainment industry”.

One of the largest concerns around an adult-oriented TLD was the trademark resolution process, an issue ICM Registry acknowledged even in their earliest proposals. In particular, due to the social stigma associated with the adult entertainment industry, brand and trademark holders wanted assurance that their marks would not be unfairly associated with the `xxx` TLD. For this reason, among others, ICM registry introduced a multi-phase rollout for `xxx` registrations, described in Section 2.5.

2.4 Criticisms and Challenges

Criticisms of the `xxx` TLD have come from both the adult entertainment community and general brand holders.

The former class has been very suspicious that the `xxx` would both incur additional costs to their businesses and that it might be a harbinger of future regulations forcing them to register exclusively under `xxx` where they could be easily filtered.

In November 2011, shortly before the opening of General Availability, Manwin Licensing International, the owner of many popular adult entertainment sites in other TLDs, filed a complaint against ICM Registry [13, 14]. In their complaint, they claim ICM Registry uses “monopolistic conduct, price gouging, and anti-competitive and unfair practices.” Manwin also describes the necessity of defensively registering one’s names, and further claimed issue with the entirety of the defensive registration process, described in the next section. This includes ICM Registry’s naming restrictions, which require registered trademark status on *exact strings* to participate in the Sunrise periods, and its “monopolistic pricing.”

In a settlement announced in May 2013, the companies settled under terms that temporarily reduce the cost of `xxx` registrations (\$8.99, comparable to `.com`, for up to 10 years) for the remainder of the month, with similar periodic discounts in the future.¹

¹While not indicative of all registrations, the `xxx` zone shows a

<i>Registration Period</i>	<i>Domains</i>	<i>Fee per Domain (USD)</i>	<i>Total Revenue (USD)</i>
Registry Reserved Names	15,000	0	0
Approved Performer Program	3500	0	0
Founder's Program	1524	Variable	4.2 M
Sunrise A	12,496	210	2.6 M
Sunrise B	66,442	200	13.3 M
Landrush	5022	200	1 M
General Availability	55,367	100	5.5 M
Totals	159,351		26.6 M

Table 1: Domain registrations on or before December 6, 2012 (the opening of general availability) and associated first-year fees.

2.5 xxx Deployment Phases

Partly due to the controversy surrounding xxx and partly to reduce the number of speculative registrations, the xxx TLD initially had many options for users to choose from when registering or defending their names. This section describes the rollout period in detail, while Table 1 summarizes the registration options. We use GoDaddy as our source for all pricing information [5] because it is the largest registrar for xxx domains with a quarter of all registrations [15], as well as the one most prominently advertised by ICM Registry [4].²

ICM Registry made the very first xxx domains available through its *Founder's Program*. Users could register these names by bidding on them, beginning in late July 2011. ICM Registry's goal in auctioning off the initial batch of xxx domain names was to ensure that the most highly desirable adult-relevant domains would be purchased by someone who wanted to grow a business and actually use the name, rather than for speculation purposes. These proposals required an accompanying description of the intended monetization model. ICM Registry auctioned 1,524 domain names in this way, including names like `casting.xxx`, for a total of roughly \$4.2 million USD [10].

In addition, ICM Registry reserved some xxx domains from registration. These domains fell into two distinct categories. The first category, domains in the *Approved Performer Program (APP)*, contains the names of specific celebrities, as well as actors and actresses in the adult entertainment industry [24, 28]. The registry held on to these domains until February 2012, a few months after the beginning of general availability. These then became available only to their respective persons, with the first year of registration included. ICM Registry claims this program included 3,500 domains at its inception [28].

The other category, *Registry Reserved Names (RRNs)*, includes domains made unavailable by ICM Registry on behalf of others, including "words of cultural and/or religious significance" [29] and words requested by the government [22]. These domains will never

14x increase in the number of new domains per day during this discounted registration period. Within a few weeks of its end, the churn rate returned to its pre-discount levels, giving some indication of price sensitivity on the registration process.

²We use retail prices in our calculations since this reflects the gross revenue paid into the xxx TLD. However, it is important to remember that this money is shared between the registrar, the registry, the registry's sponsor and ICANN. Our understanding is that the *wholesale* price for xxx is effectively \$62, with \$10 per domain going to the International Foundation For Online Responsibility (IFFOR), who is the sponsor for xxx, and \$2 going to ICANN.

be registrable and they resolve to a Web page with the text "This domain has been reserved from registration." While some of these names can be guessed (e.g., the first and last name combination of every U.S. Senator or Representative as of December 2011), ICM Registry does not supply an exhaustive list. Consequently, we are unable to empirically measure the correct number of these (see Sunrise B below), but ICM Registry's chief executive officer is reported to have claimed 15,000 domains were so reserved [17].

After this auction, ICM Registry opened the *Sunrise A* registration period on August 28, 2011. During this phase, users could get first pick of domains relating to their existing trademarks, or domains they already owned under a different TLD. Simultaneously, ICM Registry also opened the *Sunrise B* phases, which provided the first (and only) chance for brand holders to *permanently* reserve their documented trademarks by paying a one-time fee. No potential registrant can later register a Sunrise B domain, not even the original applicant.

Domains reserved from registration through Sunrise B follow the same process as Registry Reserved Names; both resolve to the same Web page and have the same whois record, using a specific contact email at ICM Registry. We are able to empirically measure the aggregate of these two categories, 81,442 domains, but cannot distinguish between them. Thus, we derive the number of Sunrise B registrations, 66,442, by subtracting the claimed 15,000 RRNs from this total.

At the close of both Sunrise periods, ICM Registry announced the total number of sunrise domains as 78,938 [11]. The 66,442 Sunrise B registrations thus implies 12,496 Sunrise A registrations.

A one-week *landrush* phase followed the sunrise period, where companies with competing applications could bid for a name. Finally, on December 6, 2011, xxx opened up to *general availability*. Like other TLDs, names are now available on a first-come, first-served basis. ICM announced the number of first-day general availability registrations as 55,367 and a total TLD size of 159,351 domains [23]. Based on that total and the sizes of the categories, we calculate the size of the only remaining category, *landrush*, as 5,022 domains.

3. DATA COLLECTION

We use data gathered from the xxx zone file and ICANN reports, as well as actively gathered Web, whois, and name server responses. Additionally, we use Alexa and SIE passive DNS to examine the xxx TLD visit patterns of Internet users. This section describes our data collection methodology, as well as the data collection methodology of each external service we rely on.

3.1 Zone File

As per the ICANN .XXX Registry Agreement [25, 30], ICM Registry allows users to download complete copies of the xxx zone file, updated once every 12 hours, except for certain purposes (such as bulk advertising). The xxx zone file we download contains only NS and A records for second-level domains. We largely use the zone file as a list of resolving xxx domains for any particular day. The zone file contained 106,789 unique domains on April 12, 2013, the date of our latest Web crawl.

3.2 ICANN Monthly Reports

An addendum to the ICANN .XXX Registry Agreement also specifies that the registry must compile domain registration reports, once per month [26]. ICANN publishes these reports for all gTLDs on their Web site [15]. These reports provide unique domain registrations, adds, transfers, and deletes per registrar.

These reports only include *registered* domains, i.e. those sponsored by a registrar. For each registered domain, some registrant must pay an annual fee or the domain becomes available again. Registered domains do not necessarily resolve. Additionally, ICM Registry includes some domains in the xxx TLD themselves; since no registrar sponsors these domains, ICM Registry does not include them in the ICANN reports.

ICANN publishes these reports once per month with a three-month delay. Partly based on our measurements (Section 4.1.2), we believe the information in the reports reflects data as of the last day of the month. While the size of these reports varies, the April 2013 report contains summary information for 88 unique registrars and 108,337 domains.

Together, the zone file and ICANN reports provide at least summary information about every registered or resolving domain in the xxx TLD. We expect most non-registered and non-resolving second-level domain strings to be those that are still available for registration, although other categories do exist. The only other domains we believe to fall in this set are those names reserved from registration by ICANN in all new gTLDs (Appendix 6 of the Registry Agreement [27]), such as all names two letters or fewer and ICANN- or IANA-related names.

3.3 Web Crawl

For each registered domain appearing in the xxx zone file, we collect the Web content served with the top-level URL (e.g., `http://icm.xxx`). We save the DOM, a screenshot, the HTTP status code, and any associated headers per domain.

Our Web crawler is based on a custom extension for Mozilla Firefox [12] with supporting code written in Python. By using a browser with significant market share, a real JavaScript engine, and widespread plugin support, we are able to follow even convoluted redirects and parse a wide variety of content. We can also conveniently record iframe contents, which are popular on both parked and redirected Web pages.

For this study we use two crawls of the xxx zone file, one on January 10, 2013 and the other on April 12, 2013. While most of our analysis focuses on the most recent, having two Web crawls allow us to investigate category shifts over time.

During the January 2013 Web crawl, the connection between the network file server and two of our crawler instances failed. For 7.8% of our Web visits with HTTP 200 (OK) status codes, we recorded headers and associated data, but no DOMs or screenshots. We crawled the zone in order with each crawler instance processing 100 domains at a time, causing these failures to cluster into bursts alphabetically. This outage had no effect on our April data set.

Table 2 gives a breakdown of the HTTP and DNS resolutions for both web crawls. This table represents every domain in the TLD; data for domains that do not appear in the zone are calculated with the ICANN reports using the methodology described in Section 4.1.5. The zone contains significantly fewer domains in April than it did in January, which reflects domains that were not renewed at the end of the year. The registration period corresponds to a little over a year and a month due to ICM Registry’s grace period for renewing registrations, a minimum of 35 days [8].

3.4 Active Whois

We actively crawled whois records to identify Registered Non-Resolving (RNX) domains (Section 4.1.5). RNX domains return valid whois records, while non-registered domains return the string “NOT FOUND”, thereby providing an easy method to differentiate between them. We also used whois to identify reserved domains (Section 4.1.1), which all return the same contact email address.

Category	2013-01-10	2013-04-12
In Zone	116,861	106,790
DNS Errors	3,961	3,176
Other Errors	730	741
HTTP 200	112,170	102,873
DOM Stored	103,436	102,873
No DOM (bug)	8,734	0
No Zone	90,156	86,710

Table 2: The Web and DNS resolution results of xxx domains.

We used a modified version of `pywhois` for whois crawling and parsing. We crawled the xxx versions of `edu` domains and domains in the Alexa top 750,000 at a rate of 10,000 per day between January 24, 2013 and April 11, 2013. In addition, we crawled all domains in the xxx zone file on April 17, 2013, as well as each xxx domain in our SIE data set (Section 3.6) that was never in a xxx zone file.

3.5 Alexa

Alexa [1] makes domain name rankings and traffic information publicly available. Alexa bases their rankings on data collected by the Alexa toolbar, an opt-in browser add-on. Besides performing Alexa’s data collection, the toolbar provides the user with traffic rankings and reviews of the sites they visit, along with an analysis of how traffic is driven to the site. This toolbar is available for Chrome, Firefox, and Internet Explorer. We use the Alexa rankings as one gauge of domain popularity (Section 6) and when searching for Registered Non-Resolving domains (Section 4.1.5).

3.6 SIE Passive DNS

We use passive DNS data collected by SIE [20] to identify those domains that users resolve in practice. This data contains DNS queries performed by recursive DNS resolvers. Organizations provide data to SIE on an opt-in basis, and it includes data from a variety of different providers, including multiple residential ISPs and universities. SIE groups queries with responses whenever possible.

Due to caching effects, it is difficult to use passive DNS data to accurately estimate the visit rate of any particular domain. Every registered domain in the xxx zone has a TTL of one day. Barring early cache evictions, a resolver will not look up a registered domain more than once per day even if many users configured with that DNS resolver visit the site during the day. As a result, we limit our use of SIE data to very coarse-grained measurements. To alleviate concerns over two resolvers in our data set that account for a significant portion of all xxx queries, and to remove highly infrequent lookups, we include only those domains queried by at least three resolvers. Apart from this requirement, we use SIE as a binary data source; either the domain is included or it is not.

For each A or AAAA query of a xxx domain, we keep the timestamp at date granularity and the registered domain. We distinguish between *live* and *dead* domains. Live domains have a valid response, no error code, and a non-empty answer section. These domains appear in the zone file, and the existence of such a record suggests that some user visited the domain in search of content. Any domain for which SIE contains a query but the domain is not live is a dead domain. These domains may be registered and not resolving, or they may not be registered at all. Dead domain lookups suggests that some user would have visited the corresponding domain, if it were registered and populated with some sort of content. We use dead domains when searching for domains that are Registered Non-Resolving (Section 4.1.5).

4. CONTENT CATEGORIES

As a first step towards categorizing the registration intent of each domain, we categorize their contents. In this section we describe our domain classification methodology, including the indicators we use to identify each type, which data sources we draw from, and the number of domains in each category. We also examine how domain classifications have changed over time.

4.1 Classification

After crawling all domains, we cluster them based on resolution status and Web content using text shingles [2] to find the most prevalent page types. For each large cluster, we then create regular expressions to match and classify the content. For example, all Web pages corresponding to a particular domain parking service may use a single CSS file served off a different domain. Any pages with this particular CSS file likely belong to the same parking service, regardless of whether or not they are all in one cluster. We apply a *content tag* to any Web page matching a number of regular expressions for the tag, with the threshold varying by tag.

We created tags to match the largest 40 clusters, including every cluster with 50 or more domains. We also tagged some smaller clusters, typically because they were popular earlier in our study.

One consequence of our content classification methodology is that we are unlikely to find primary registrations with the clustered Web content. Domains will only cluster together if they have the same or similar Web content, but few content creators will likely need many domains for the same content. Instead, this technique is much more useful for identifying defensive, reserved, or parked domain registrations.

We classify each domain into one of eight categories:

- Reserved** domains are permanently unregistrable.
- Approved Performer Program** domains correspond to celebrity names and have a separate registration process.
- Premium Domain Names** are potentially desirable names with larger up-front registration costs.
- Parked** domains are advertised as for sale, often with additional automatically-generated advertisements as well.
- Non-Resolving** domains have been registered but do not resolve, and are usually not even in the xxx zone.
- Content** domains give meaningful Web responses.
- Unused Web** domains host little or no Web content.
- Unknown** domains could not be automatically classified.

The rest of this section discusses each classification in detail.

4.1.1 Reserved Domains

As described in Section 2.5, there are a few categories of reserved domains. This section deals with two of those, Registry Reserved Names and Sunrise B reservations. We group these together because we are unable to distinguish between them.

Reserved domains are straightforward to identify. Each returns an identical ICM-generated Web page over HTTP stating “This domain has been reserved from registration”. These pages all have the NS records `NSB{1,2}.ICMREGISTRY.NET`, which host all of ICM Registry’s domains. The whois information for these domains also specifies ICM Registry as the contact point for the domain, consistent with their intent as stated in [29].

We match whois instead of Web content to classify a domain as reserved. The whois record for each reserved domain contain the contact email address `registryescrow@icmregistry.com`, a different address than for other kinds of ICM Registry domains. Additionally, we found 172 instances of registrants copying ICM

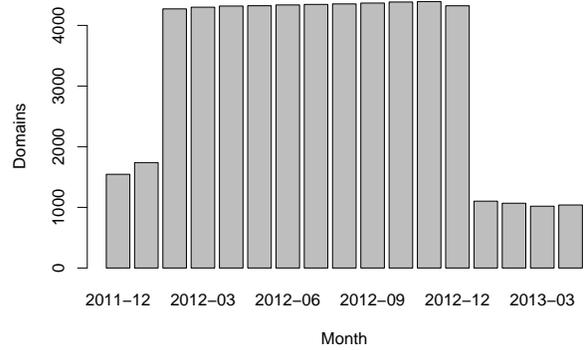


Figure 1: xxx registrations with name.com by month.

Registry’s domain reservation Web page and hosting it on their own Web server with no changes, making the web content ambiguous.

Using this methodology, we find 81,442 reserved domains. Even after checking for distinguishing characteristics in their whois records, we are unable to empirically separate Sunrise B and RRN reservations. ICM Registry reported 15,000 RRNs [17], which we use as an estimate, and classify the 66,442 remaining domains as Sunrise B reservations. These domains resolve but are not registered, so they appear in the zone file but not the ICANN reports.

4.1.2 Approved Performer Program

Domains in the Approved Performer Program, described in Section 2.5, resolve to a Web page indicating their participation in the program and specifying a process by which the person in question may take ownership of their domain name. Since they resolve, they also appear in the zone file. We identify them using our tagging methodology described in Section 4.1.

ICM Registry registered these domains through name.com to ease transferring ownership, and also to provide the rest of the first year free to the registrant. According to whois, this registration was valid from February 3, 2012 to January 24, 2013, at which time the sponsoring registrar switched back to ICM Registry. Thus they appear in the ICANN monthly reports during this time period, but not otherwise. We use this set of domains to discover whether the ICANN reports provide the unique number of domains per registrar across the entire month, or from a single snapshot.

Figure 1 shows the total domains registered through name.com each month, as reported to ICANN. As expected, we see an uptick of name.com registrations between January and February 2012, and a drop-off between December 2012 and January 2013. Since these domains switched registrars on January 24, we conclude that the ICANN reports must only use data from the last seven days of the month. We make the simplest assumption, and assume ICM Registry reports totals for the last day of the month.

We identify 2,415 Approved Performer Program pages as of April 12, 2013. The increase of 2,534 name.com registrations between January and February 2012, which include both normal domain registrations and those via the Approved Performer Program, suggests that ICM Registry only sold a very small number of these domains in the intervening 15 months. When ICM Registry started the program in October 2011 [28] they claimed withholding 3,500 of these names [24], making the first few months of the program significantly more successful.

4.1.3 Premium Domain Names

ICM Registry also withholds potentially desirable domain names from registration as part of their Premium Domain Names (PDN) program. These domains require potential registrants to purchase them for larger sums prior to registering the name. Premium Do-

main Names successfully resolve, returning ICM Registry’s search portal for xxx domains when visited via HTTP. They appear in the zone file, but do not appear in the ICANN monthly reports since they are reserved directly by ICM Registry. We find 991 of these domain names.

4.1.4 Parked Domains

Many domains are parked with domain parking services, often but not always established through a registrar. Domain registrants speculatively monetize parked domains through ad revenue and domain resale. While there are many templates for parked pages, most contain these features:

- ❖ Text indicating the domain’s availability (“This domain for sale”).
- ❖ Advertisements relating to words in the domain name, often laid out to look like search engine results.
- ❖ The name and/or logo for the domain parking service.

Our results show domains parked through many different parking services. While each parking service can set up their infrastructure in any way they choose, we expect all parked domains to have Web content. This expectation clearly holds for parked domains monetized through ad revenue, as these registrants can only make money if they actually serve advertisements. Speculatively parked domains do not necessarily require Web content, but likely get more offers if they advertise the domain as being for sale. For these reasons, we identify parked domains through their Web content, not through their whois records or hosting infrastructure.

Though we note that parked domains can be motivated by either domain resale or ad revenue, we consider domain resale to be the most likely motivator behind parked domains. Remember that xxx domains are significantly more expensive than domains in other TLDs, costing roughly \$100 per year. Ad revenue is unlikely to scale with domain registration costs, which intuitively suggests the other strong motivator, domain resale, as the most likely reason for these domains’ registrations.

These domains are both registered and resolving, so they appear in both the zone file and the ICANN monthly reports. We find 8,262 domains parked through 21 different programs using our content tags, as described in Section 4.1, including 3,561 domains parked through GoDaddy and 1,923 through Sedo.

4.1.5 Non-Resolving Domains

Some xxx domains are configured with name servers but do not successfully resolve when queried through the DNS. Upon further investigation, we discovered that many domains in this category are hosted on otherwise working name servers, but return REFUSED or SERVFAIL responses when queried with the domain in question. For example, Google has registered `picasa.xxx` with the name server `ns1.google.com` (and sequentially numbered variants), but returns REFUSED for lookups of this domain (which recursive servers often report as SERVFAIL). Other name servers return no error, but no answer, authoritative, or additional sections either. We see some less common failure modes, such as temporary failures.

Further, some xxx domains are *Registered Non-Resolving (RNX)*. These domains do not appear in the zone file because they have no associated DNS records, but some registrant is still paying the normal domain registration fee. The consequence of this kind of registration is that no other party will be able to successfully register that domain name, making this a clear case of a defensive domain registration. ICM Registry and some registrars cite this method as the correct way to defend your trademark after the beginning of general availability [7, 19].

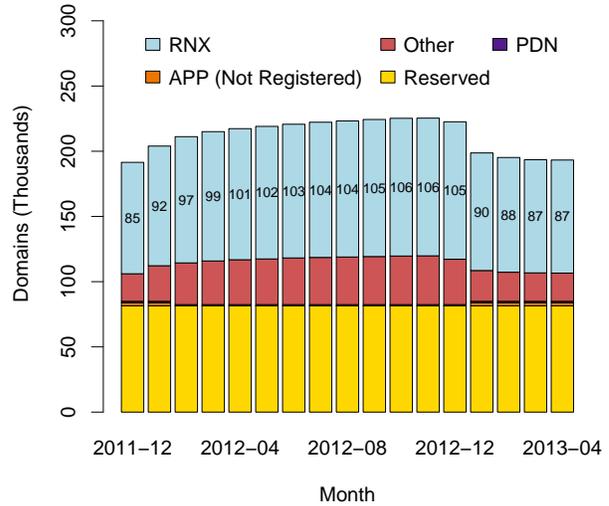


Figure 2: Breakdown of all domains in the zone and ICANN reports, which we use to estimate the number of RNX domains. The top two bars represent domains in the ICANN reports, while the bottom four bars represent domains in the zone. The red bar in the middle overlaps the categories.

Given that this category is the logical equivalent of domains in the ICANN reports but not the zone file, we can calculate how many of these exist at any given time by using the categorizations from the zone file. Some domains from the ICANN reports also exist in the zone file; let *Overlap* represent the total number of domains in this set. Then we can find the number of Registered Non-Resolving domains, *RNX*, using:

$$RNX = Report - Overlap$$

While this method will compute the number of RNX domains, the size of the overlap set changes over time and is difficult to measure directly. We can find more accurate historical values for *RNX* by rewriting this formula in terms of the size of the zone file, *Zone*, and the size of the more stable and measurable categories. This includes Sunrise B applicants, Registry Reserved Names, Premium Domain Names (PDNs), and sometimes the Approved Performer Program (Section 4.1.2).

Using this set of categories, we can calculate *RNX* as follows:

$$\begin{aligned}
 Report &= Overlap + RNX \\
 Zone &= Overlap + Reserved + APP + PDN \\
 Report - Zone &= (Overlap + RNX) - \\
 &\quad (Overlap + Reserved + APP + PDN) \\
 Report - Zone &= RNX - Reserved - APP - PDN \\
 RNX &= Report - Zone + \\
 &\quad Reserved + APP + PDN
 \end{aligned}$$

Figure 2 shows the number of RNX domains for each month from general availability of the xxx TLD through April 2013. Our methodology shows the existence of 86,710 RNX domains in April 2013. Added to the 3,176 non-resolving domains in the zone file, we find a total of 89,886 non-resolving xxx domains.

We were able to identify some of these RNX domains by actively crawling their whois records. Domains that have not been registered or reserved return “NOT FOUND”, while those that have been registered return valid whois information, regardless of whether or not they have any associated NS entries.

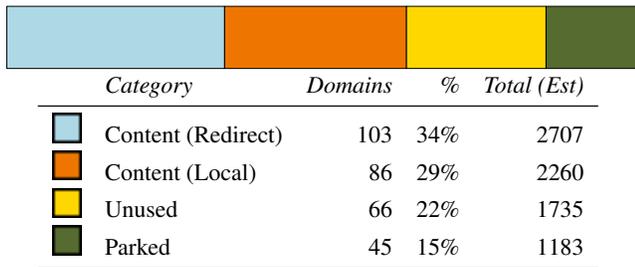


Figure 3: Categories of 300 random Unknown domains.

To preliminarily identify registered non-resolving domains, we performed whois queries for the `xxx` version of each domain in the `edu` zone file (such as `ucsd.xxx`), as well as for each domain in the Alexa top million Web sites. We also crawled whois information for each SIE lookup attempt for a dead domain. These sources of likely defenders identified 19,873 (23%) of the `RNX` domain names, but many other defenders remain unidentified.

4.1.6 Unused Web

We categorize domains as “Unused Web” if the domains resolve but do not return any useful Web content. Some `xxx` domains successfully resolve through the DNS but do not respond to HTTP requests on port 80. The remainder run a public Web server on the default port and return either a blank Web page, the default Apache Web page, a PHP error page, or a similar content-free response.

There are several possible explanations for these responses. First, this state could be transient; that is, other kinds of domains could fall into this state until their operators notice and rectify the problem. Second, these domains could signal defensive registrations; perhaps some brand holders are not familiar with the suggested defensive methods and misconfigure their servers or NS entries. Alternatively, these domains could be primary registrations, but intended to support services other than the Web.

Regardless of intent, we identify 570 domains as returning HTTP errors and 609 as returning successful but content-free responses, for a total of 1,179 unused Web domains.

4.1.7 Content

A handful of our large Web content clusters contain legitimate content that end users may find interesting, and thus we categorize them as “content” domains. Additionally, most adult pages have an 18 USC 2257 notice, which states their compliance with record-keeping requirements for each adult performer on the site. We classify 231 `xxx` domains as content due to our content tags on large clusters and 1,072 due to their 18 USC 2257 warning, for a total of 1,303 content domains.

4.1.8 Unknown

Finally, our automatic classification techniques do not correctly classify some `xxx` domains. Our classification criteria err on the strict side to avoid false positives, and often rely on manually-generated Web content tags as described in Section 4.1. The 7,885 domains we are unable to classify with the methodology above fall into this category.

We suspect much of this category contains legitimate content. Our content classification depends heavily on manually-generated tags, which we only create for our largest clusters. Many types of legitimate content will not cluster with any other `xxx` domains, while our other categories will tend to contain pages that cluster together better.

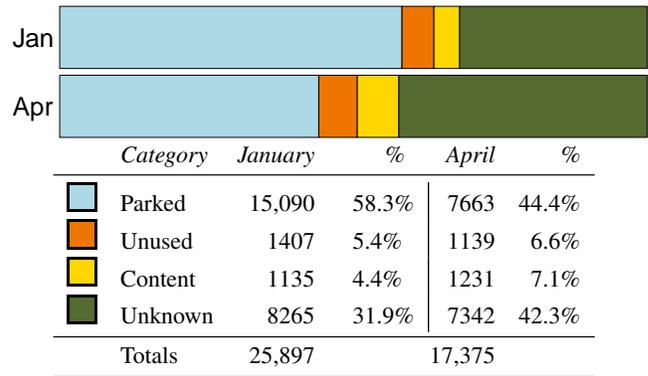


Figure 4: Content classifications for all registered and resolving `xxx` domains during both crawls.

To verify our suspicions about these unclassified domains, we performed a manual sample of the Web content for these pages. We examined the Web content for 300 domains in this category, chosen uniformly at random. We use the Web contents and screenshots from our April 12 Web crawl while classifying these domains, so the content matches exactly with the content we were unable to automatically classify. Figure 3 summarizes our results. In each case we were able to classify the domain into one of our existing categories. As expected, most domains (63%) that we could not automatically classify fell into the content category, with the rest dividing into unused and parked.

For any domain that showed primary content, we also judged whether or not it was adult-oriented. We used a very liberal policy for assigning content to this category. For instance, `pureheaven.xxx` had broken image and product links for adult-oriented items during our crawl, which we considered adult. We found 76% of originally unknown content or redirect pages to be adult in nature.

4.2 Overall Classifications

Figure 4 shows the content class breakdown for each domain in the `xxx` zone file that is both registered and resolving (except Approved Performer Program domains in January). Both classifications ignore all `xxx` domains which had HTTP 200 response codes and we were unable to store the DOM in January (Section 3.3). While we believe these domains to be distributed across the content categories like the other `xxx` domains, removing them from both sets will make category size comparisons between the data sets more straightforward.

Even among only domains that both resolve and are registered, parked domains are the most prevalent. After breaking up the unknown category using the data from Figure 3, parked domains still make up the largest category of both classifications. Note, though, that the sum of all of these categories is still much smaller than either of the two largest categories of `xxx` domains in general, `RNX` and reserved domains.

When comparing the classifications across time, two features stand out. First, there are many fewer registered resolving domains in the April data set than in January. Second, there are far fewer parked domains in April than there are in January, even after accounting for the general dropoff. In the following sections, we investigate both of these features.

4.3 Elapsed Domain Names

Figure 4 shows that the `xxx` zone file contained fewer registered resolving domains in April 2013 than it had three months before, similar to the reduction in zone file size in Section 3.3. On January

Category	Domains		
	(Jan)	Elapsed	%
Parked	15090	7979	52.9%
Unknown	8265	1406	17.0%
Non-Resolving	3866	916	23.7%
Unused	1407	326	23.2%
Content	1135	85	7.5%
XXX Support	3302	56	1.7%

Figure 5: Elapsed domain names across categories. The last column shows the probability that a registrant chose not to renew a random domain from January.

Category	Total Domains	Total Changed	% Changed
Parked	15090	573	3.8%
Non-Resolving	3866	547	14.1%
Unused Web	1407	490	34.8%
Content	1135	99	8.7%
Unknown	8265	738	8.9%
Support	3302	18	0.5%

Table 3: Domains that changed content classifications between January and April.

10, 2013 the `xxx` zone file contained a total of 116,833 entries, only 3,048 fewer than its peak of 119,881 a month earlier; this drops off to 111,554 only ten days later, with 5,279 fewer domains. We see 10,768 registered and resolving domains on January 10 that do not appear in the April 12 zone. Since this time period falls roughly one year after the beginning of general availability, we can use this gap to approximate those domains the registrant chose not to renew after one year.

Figure 5 shows the content classification for each of the 10,768 `xxx` domains in the January 10, 2013 zone, but not the April 12 one. We actively crawled whois for these domains on May 6 and found that 9,584 of them were no longer registered, leaving 1,184 as registered but not resolving.

From the breakdown, most elapsed domains were parked, especially when compared to registered and resolving content in general. Using our content tags, we classified 15,090 domains as parked at the 21 roughly largest parking programs during our January classifications; over half (7,979, or 53%) of them disappeared over the next 92 days. This data suggests that domain parking in the `xxx` TLD has not been particularly profitable. Regardless of whether these domains were registered for ad revenue or resale potential, it appears as if their registrants no longer believed these domains were likely to pay off.

4.4 Classification Changes

When classifying domains with data collected over a small time interval, our classifications gain the appearance of being very stable. An alternate hypothesis is that some classifications, such as unused or parked domains, are transient domain states; that is, domains are only unused or parked for short periods of time while their owners set up real content.

Intent	Domains	%
Primary	6,270	5.9%
Speculative	9,445	8.9%
Defensive	89,886	85.1%
Total	105,601	

Table 4: Registration motivations of all *registered* `xxx` domains (55% of all domains) for which we could infer intent. The remaining domains are reserved domains, which are defensive by definition.

We use our January 10 Web crawl and a methodology similar to Section 4.3 to look for transient domain states. We classify domains from the January 10 set using the same methodology and content tags as we use for the April 12 data set. For each potential categorization, Table 3 shows the number of domains originally in that category, the subset of those that switched categories in the next 92 days, and the fraction of domains that switched categorizations.

For most categories, we see very little change in classification. Less than 4% of parked domains switched classifications. This result and that of Section 4.3 strongly suggest that parked domains in `xxx` tend to be parked for long periods of time.

The largest outlier is the category of domains with either no Web server, or that serve only empty Web content. Nearly 35% of domains in this category switch to a different one within 92 days. When we introduced these domains in Section 4.1.6, we suggested three different potential purposes for them: primary, defensive, or transient. While the motivations behind the other 65% are still unclear, this evidence supports transience as the explanation for a significant portion of these domains.

5. REGISTRATION INTENT

Upon categorizing each `xxx` domain, we would like to determine the registration intent, whenever possible. Similar to our previous work on the `biz` TLD [6], in this section we group `xxx` domains into one of three functional groupings: primary, defensive, or speculative. Table 4 summarizes our results. In particular, we find that an overwhelming number of registered domains (85%) are defensive in nature, with only 5.9% of domains registered for the purpose of serving content.

Before categorizing domain registrations by motivation, we must identify and separate all varieties of reserved domains. The primary purpose of further categorizing domains by motivation is to compare both the number and the amount spent on domains registered for different purposes. Because of this goal, we first need to identify and separate those domains with no recurring registration costs. This includes the reserved and unregistrable domains, Premium Domain Names, and those in the Adult Performer Program.³ We find a total of 84,848 domains with no recurring registration costs using these criteria, all of which are defensive in nature.

In Section 4.1.6, we proposed three different motivations behind `xxx` domains with unused Web content. While the classification change results in Section 4.4 were not conclusive in this regard, it did suggest the transient domain hypothesis is the most common. As a result, we do not include these domains in the rest of this

³While domains in the Adult Performer Program were registered through `name.com` for a year, we are unsure if ICM Registry obtained these through the normal domain registration process. Regardless, the majority of the domain registration cost would have returned to ICM Registry, so they should still not be included in this total.

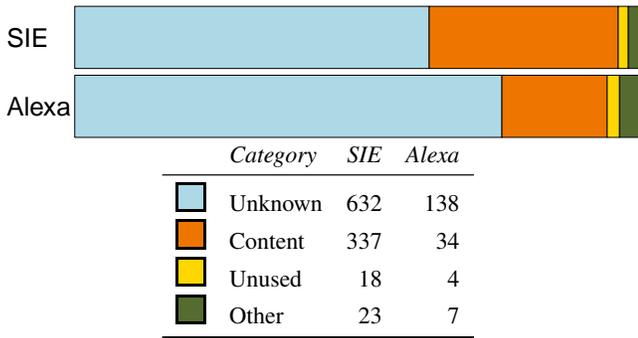


Figure 6: Classifications for domain names appearing in passive DNS (top) and Alexa (bottom) within three days of our Web crawl. The SIE data only includes domains resolved by at least three resolvers.

section. While it is possible that some of these domains should have definitive classifications, the size of the entire category is relatively small and so the misclassification of a subset of these domains will have little impact on the results.

The rest of this section discusses the registrant motivations assigned to the other 105,601 domains in the remaining content categories: “Parked”, “Non-Resolving”, and “Content”. We also describe any known issues with our decisions.

5.1 Defensive Registrations

RNX domains are the most obvious category of defensive domain registrations. ICM Registry and GoDaddy, the largest registrar in xxx with a quarter of all domains [15], both specify RNX as the intended mechanism for defensive registrations after the conclusion of the Sunrise B phase [7, 19]. RNX domains also do not resolve, making it difficult for a devious user to use them even for obscure purposes.⁴

5.2 Primary Registrations

We assume all domains that point to content are primary registrations. We do not differentiate between direct and redirected content. We also ignore whether or not the content is adult in nature and therefore even belongs in the xxx TLD. At first this might seem strange, as an off-domain redirect might be a good indicator of a defensive registration.

To understand our policy, consider the definition of a primary registration. When one domain redirects to another, the primary domain is the one that the content owner actively advertises. If we could categorically say that xxx domains are undesirable to all registrants, redirects would strongly indicate defensive registrations. While valid for trademark defenders, it is not valid for content producers, who may want to use the connotations of the xxx TLD for marketing purposes. The registrant may serve content via a domain in a different TLD for historical reasons, but prefer and advertise the xxx version, thereby making it the primary registration.

We noticed a fair number of redirects in our manual sample (see Section 4.1.8); by extrapolating our manual sample results, we estimate 2,500 xxx domains whose Web contents redirect off-domain. This quantity makes up a large portion of domains with content but is dwarfed by the number of reserved (81,442) and RNX (86,710) domains. Given that there is no clear and simple policy for redi-

⁴These registrations could also be speculative, but we assume speculative registrants want to broadly advertise their domains as being for sale, at a minimum by resolving to a server hosting recognizable Web content.

Ranking	Domain
6,452	perfectgirls.xxx
12,450	hornytube.xxx
14,220	hdsextube.xxx
15,178	rule34.xxx
24,660	entire.xxx

Table 5: The five most popular xxx domains in Alexa.

rects, we assume they are all primary registrations and treat this classification as an overestimate.

5.3 Speculative Registrations

We consider parked domains to be speculative in nature. Parked domains incentivize the registrant to buy likely desirable domains without adding any useful content for their visitors, regardless of whether their monetization model favors ad revenue or resale.

6. VISIT PATTERNS

So far, we have focused mostly on the costs of the xxx TLD by showing that there are many defensive and speculative registrations. In this section we examine the benefits. We would like to examine the TLD from the perspective of its early rejections by ICANN and determine the degree to which the TLD serves any unmet needs. One way to examine this question is to look at how many users visit xxx domains. We answer this question using our two sources of visit information, SIE and Alexa.

Figure 6 shows the breakdown of xxx domains in our SIE passive DNS data set within three days of April 12, 2013, the date of our latest Web crawl. This date range is contemporaneous with our Web crawl, so we expect our classifications to remain accurate.

We classified nearly every xxx domain (96%) appearing in the passive SIE data as either “unknown” or “content”. This result matches our intuition, as we expect users to be looking up domains with real content on them. A small fraction of these domains fall into the “unused” category. These domains may have been registered for purposes other than Web content, such as IRC or email. Alternatively, these domains may offer Web content on a different port or content path than the default, and therefore not be represented in our data set. While xxx domains are significantly more expensive than domains in other TLDs, our manual sample showed us evidence of domains meant only for domain registrants and their close friends, such as blogs, a personal resume, and a handful of vanity domains. Some domains in our SIE data set may reflect this use case.

We performed a similar analysis on the top one million Alexa-ranked domains on April 12, 2013, of which only 184 xxx domains appeared. Figure 6 shows our categorizations for 183 of these domains; the last dropped out of the zone file on April 11. Table 5 shows the top-ranked xxx domains and their rankings. Only one domain appears in the top 10,000; for comparison, the name TLD also only has one domain in the top 10,000, but biz has 17 and info has 50.

The top Alexa domains predominantly fall into the unknown and content categories and closely resemble the SIE data, as we would expect. Together, unknown and content account for over 94% of xxx domains in the Alexa top million. We classified two Alexa domains as non-resolving; one of these appears to no longer be resolvable, while the other was likely a transient DNS failure during the time of our Web crawl. ICM Registry runs 5 domains in Alexa (classified above as “other”), such as search.xxx.

<i>Registration Intent</i>	<i>More Information</i>	<i>Peak Registrations</i>	<i>Initial Costs (USD)</i>	<i>Number of Domains</i>	<i>Yearly Fee (USD)</i>	<i>Recurring Cost (USD)</i>
Reserved	Section 5	84,848	13.3 Million	84,848	0	0
Defensive	Section 5.1	109,559	11.0 Million	89,886	100	8.99 Million
Primary	Section 5.2	8,666	5.7 Million	7,433	100	0.74 Million
Speculative	Section 5.3	22,313	3.3 Million	11,196	100	1.12 Million
Totals		225,386	33.3 Million	193,363		10.9 Million

Table 6: Upfront and yearly renewal costs for xxx domains by registrant motivation.

7. REGISTRATION COSTS

The domain classifications form a basis for measuring the impact of the xxx TLD in terms of its costs and benefits. The benefits are clear: primary domain registrants have a new domain name, likely one that is easier for users to remember or stumble upon, which might make their business more appealing. The cost lies in the money spent by domain registrants for primary, defensive, and speculative purposes.

Domain registration costs vary by registrar, but performing cost estimates requires us to choose a price. ICM Registry’s domain registration page features GoDaddy most prominently [4], so as in Section 2.5 we base our price estimates using GoDaddy’s registration costs of \$100 USD [5]. We use the registration cost instead of the \$62 wholesale price to capture money spent instead of any one particular party’s revenue.

Table 6 shows the estimated registration costs for each domain in xxx. Recurring cost estimates use the registration intent numbers from Section 5. The only exception is transient domain names; these names are still registered (i.e., someone is paying the registration cost for them), but we have not previously categorized them. For purposes of our cost estimates, we divide these between primary and speculative registrations, based on the ratios those domains otherwise exhibit. Since all measurements are from the period shortly after one year of general availability, the recurring costs section of this table reflects domains for which users chose to renew initial registrations after one year.

We estimate that in terms of initial costs, only 17% were for primary domains and the rest were speculative or defensive in nature. In terms of recurring cost, the situation is even more stark: only 6.8% of recurring costs are for primary domains. These estimates underscore the conclusion that the vast majority of registration revenue, both during the pre-registration period and for ongoing operation, is driven by defensive concerns rather than entrepreneurship.

Initial cost estimates use data from a variety of sources. The peak number of domain reservations uses the categories described in Section 5. The only category of domain reservations for which registrants paid money is Sunrise B, and that cost comes directly from Table 1. We were able to calculate the peak number of RNX domains, our only category of defenders, in Section 4.1.5, as well as the peak number of domains in xxx. Primary and speculative registrations make up the remaining 30,979 domains.

We use the January content classifications to estimate the ratio of primary to speculative registrations. We use this data instead of the April data because it is temporally closer to the peak. Additionally, the exodus of parked registrations seen in Section 4.3 makes the April classifications less reliable for this purpose. We expect our estimate to be slightly biased towards content, since the zone size had already fallen by roughly 3,000 domains by January 10, and because more speculative registrants choose not to renew than primary registrants.

Finally, for our initial cost estimates, we assume all of the most expensive domain names (those in Sunrise A and the Founder’s Program) all went to primary registrants. While some general availability registrations were certainly primary instead of speculative in nature, this assumption again lets us bias towards primary registrations. Additionally, while some domains in the Founder’s Program may serve parked content, the uneven and unknown cost of individual domain names in this category creates the potential for fine-grained estimates to be incorrect in unpredictable ways. Instead, we assume the most expensive registrations were for primary content, and treat the resulting number as a high estimate.

8. CONCLUSION

The introduction of new TLDs is typically meant to increase consumer choice with respect to second-level domain strings by opening up second-level domains that have already been claimed in other TLDs. As we have previously shown in a case study of the biz TLD [6], some people take advantage of such opportunities to register new content, while others feel compelled to defend their names, and still others seek to resell desirable domains without using them themselves. In general-purpose TLDs like biz, these registration types are all fairly common.

Aspects of the new xxx TLD amplify concerns about defensive and speculative registrations. Brand and trademark holders are particularly concerned about any potential association with the connotations of such a TLD, while previous members of the adult entertainment industry do not consider it to serve any real need. This paper empirically studies the validity of these concerns by measuring the costs and benefits of the xxx TLD.

By gathering public data from ICANN, zone file records, active whois and Web content, we are able to build a complete view of xxx domains. We show that concerns over defensive registrations are particularly valid: nearly 92% of all domains in the xxx TLD exist for solely defensive purposes, including 83% of all domains with recurring registration costs. Speculative registrations make up 60% of the remainder, and only 7,433 domains, or 3.8% of all xxx domains, serve real Web content. Additionally, we find that registrants spent \$24M in the first year of registration solely for defensive purposes.

9. ACKNOWLEDGMENTS

We would like to thank Neha Chachra and Andreas Pitsillidis for their tools that greatly ease Web content crawling and clustering, as well as Brian Kantor and Cindy Moore for managing our compute resources. We are also grateful to our anonymous reviewers for their feedback. This work was supported by National Science Foundation grants NSF-1237264, by the Office of Naval Research MURI grant N00014-09-1-1081, and by generous research, operational and/or in-kind support from Google, Microsoft, Yahoo, and the UCSD Center for Networked Systems (CNS).

10. REFERENCES

- [1] Alexa. <http://www.alexacom.com>.
- [2] A. Broder, S. Glassman, M. Manasse, and G. Zweig. Syntactic Clustering of the Web. *Computer Networks and ISDN Systems*, 29(8):1157–1166, 1997.
- [3] Domain Name Supporting Organization. Working Group C – creation of new gTLD. <http://www.dnso.org/dnso/notes/19990625.NCwgC.html>, June 1999.
- [4] Find and Buy Your .XXX Domain Name. <http://www.buy.xxx>.
- [5] GoDaddy Registration Prices on 2011-09-28. <http://web.archive.org/web/20110928063106/http://www.godaddy.com/tlds/xxx-domain.aspx?tld=xxx&ci=51450>.
- [6] T. Halvorson, J. Szurdi, G. Maier, M. Felegyhazi, C. Kreibich, N. Weaver, K. Levchenko, and V. Paxson. The BIZ Top-Level Domain: Ten Years Later. In *Proceedings of the Passive and Active Measurement Workshop*, Vienna, Austria, Mar. 2012.
- [7] How do future brands protect themselves? https://member-icmregistry.custhelp.com/app/answers/detail/a_id/26/.
- [8] How do I acquire NON-Renewed names? https://member-icmregistry.custhelp.com/app/answers/detail/a_id/85/.
- [9] ICANN Options Following the IRP Declaration on ICM’s .XXX Application. <http://www.icann.org/en/irp/icm-v-icann/draft-options-post-irp-declaration-26mar10-en.pdf>.
- [10] ICM Registry Announces the Successful Conclusion of the .XXX Founders Program. <http://www.icmregistry.com/press/icm-registry-announces-the-successful-conclusion-of-the-xxx-founders-program/>.
- [11] ICM Registry closes sunrise period with 80000 applications for XXX domain names. <http://www.icmregistry.com/press/icm-registry-closes-sunrise-period-with-80000-applications-for-xxx-domain-names/>.
- [12] K. Levchenko, A. Pitsillidis, N. Chachra, B. Enright, M. Félegyházi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, H. Liu, D. McCoy, N. Weaver, V. Paxson, G. M. Voelker, and S. Savage. Click Trajectories: End-to-End Analysis of the Spam Value Chain. In *Proceedings of the IEEE Symposium and Security and Privacy*, pages 431–446, Oakland, CA, May 2011.
- [13] Manwin Licensing International v. ICM Registry, et al. <http://www.icann.org/en/news/litigation/manwin-v-icm>.
- [14] Manwin Licensing International’s Complaint, Case CV 11-9514-PSG. <http://www.icann.org/en/news/litigation/manwin-v-icm/complaint-16nov11-en.pdf>.
- [15] Monthly Registry Reports. <http://www.icann.org/en/resources/registries/reports>.
- [16] T. Moore and B. Edelman. Measuring the Perpetrators and Funders of Typosquatting. In *Proceedings of the 14th International Conference on Financial Cryptography and Data Security*, 2010.
- [17] Porn domain .xxx blocks use of celebrity names. <http://www.guardian.co.uk/technology/2011/sep/07/porn-domain-xxx-celebrity-names>.
- [18] Reconsideration Request 00-15. <http://archive.icann.org/en/committees/reconsideration/rc00-15-1.htm>.
- [19] Registering .XXX Domain Names. <http://support.godaddy.com/help/article/5789/registering-xxx-domain-names>.
- [20] SIE Passive DNS. https://archive.farsightsecurity.com/SIE_Channel_202/.
- [21] T. Telegraph. Top 10 most expensive domain names. <http://www.telegraph.co.uk/technology/news/7412544/Top-10-most-expensive-domain-names.html>.
- [22] The ABC&Ds about .XXX. <http://www.icmregistry.com/about/policies/abcd/>.
- [23] Tweet by ICM Registry on the @dotXXX Twitter Account. <https://twitter.com/dotXXX/status/144495755946762240>.
- [24] .XXX Adult Performer Program. <http://performers.icmregistry.com/>.
- [25] .XXX Agreement Appendix 3: Zone File Access Agreement. <http://www.icann.org/en/about/agreements/registries/xxx/appendix-3-31mar11-en.htm>.
- [26] .XXX Agreement Appendix 4: Registry Operator’s Monthly Reports. <http://www.icann.org/en/about/agreements/registries/xxx/appendix-4-31mar11-en.htm>.
- [27] .XXX Agreement Appendix 6: Schedule of Reserved Names. <http://www.icann.org/en/about/agreements/registries/xxx/appendix-6-31mar11-en.htm>.
- [28] .XXX Announces Launch of Pre-Paid Adult Performer Program. <http://www.icmregistry.com/press/xxx-announces-launch-of-pre-paid-adult-performer-program/>.
- [29] .XXX Launch Plan and Related Policies. <http://www.icmregistry.com/about/policies/launch/>.
- [30] .XXX Registry Agreement. <http://www.icann.org/en/about/agreements/registries/xxx/xxx-agreement-31mar11-en.htm>.