

From Deletion to Re-Registration in Zero Seconds: Domain Registrar Behaviour During the Drop

Tobias Lauinger
Northeastern University

Ahmet S. Buyukkayhan
Northeastern University

Abdelberi Chaabane
Nokia Bell Labs

William Robertson
Northeastern University

Engin Kirda
Northeastern University

ABSTRACT

When desirable Internet domain names expire, they are often re-registered in the very moment the old registration is deleted, in a highly competitive and resource-intensive practice called domain drop-catching. To date, there has been little insight into the daily time period when expired domain names are deleted, and the race to re-registration that takes place. In this paper, we show that .com domains are deleted in a predictable order, and propose a model to infer the earliest possible time a domain could have been re-registered. We leverage this model to characterise at a precision of seconds how fast certain types of domain names are re-registered. We show that 9.5 % of deleted domains are re-registered with a delay of zero seconds. Domains not taken immediately by the drop-catch services are often re-registered later, with different behaviours over the following seconds, minutes and hours. Since these behaviours imply different effort and price points, our methodology can be useful for future work to explain the uses of re-registered domains.

CCS CONCEPTS

• **General and reference** → **Measurement**; • **Networks** → *Naming and addressing*; *Public Internet*;

KEYWORDS

Domain Name System (DNS), domain name, expiration, pending delete, deletion time, drop-catch, re-registration delay, registrar

ACM Reference Format:

Tobias Lauinger, Ahmet S. Buyukkayhan, Abdelberi Chaabane, William Robertson, and Engin Kirda. 2018. From Deletion to Re-Registration in Zero Seconds: Domain Registrar Behaviour During the Drop. In *2018 Internet Measurement Conference (IMC '18), October 31–November 2, 2018, Boston, MA, USA*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3278532.3278560>

1 INTRODUCTION

Re-registration of expired Internet domain names can be quite competitive. Drop-catchers, who attempt to re-register a domain in the

same instant the expired registration is deleted, consume a considerable share of resources of the domain registration ecosystem. In prior work, we showed that three large drop-catch services control 75 % of all domain registrar accreditations and are responsible for at least 80 % of domain creation attempts, even though the vast majority of their attempts fail and no more than 9.5 % of new .com domains are created by drop-catch services [10].

Due to a lack of metadata with a better time resolution, we previously approximated drop-catch domains as all those re-registered on the deletion date [10]. However, as we illustrate in this work, the truly competitive re-registration period lasts only about one hour. Consequently, our prior work could not distinguish between true drop-catch domains, and others re-registered many hours later with less competition and normal resource consumption. For instance, there is evidence that some actors engage in “home-grown” drop-catching using desktop software and domain reseller APIs [5, 14] to avoid the fees of traditional drop-catch services, which can be two to ten times more expensive than a regular domain registration. This difference in price may have an influence on future uses of re-registered domains. For example, one may suspect that the pricier drop-catch domains are less likely to be used for malicious purposes than delayed re-registrations.

In this paper, we improve the precision of drop-catch research [10, 13, 17] to a time scale of seconds, instead of days. We show that .com domains become available in a predictable order, and develop a model to infer the earliest time a domain could have been re-registered. Using this model, we paint a more detailed picture of the competitive landscape. We calculate domain re-registration delays with respect to the earliest possible instant, and classify re-registration behaviour according to the order of magnitude of this delay. We show that there exist qualitative differences in the domains re-registered, and in the domain registrars active within each of these delay periods. For example, we find that drop-catch accounts for approximately 86.1 % of domains re-registered on the deletion day. This means that the approximation from prior work based on the date alone misclassifies 13.9 % of domains, which are in fact delayed re-registrations with different domain and registrar characteristics, and a different cost. Our methodology for distinguishing re-registration types can be useful for future studies that explore the registration intent and uses of domains.

Our work makes the following contributions:

- We provide the first detailed look at .com domain deletion and re-registration times at a precision of seconds, as opposed to the daily aggregates used in prior work.
- We infer and present the first model of the earliest possible re-registration time of .com domains on their deletion day.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '18, October 31–November 2, 2018, Boston, MA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5619-0/18/10...\$15.00

<https://doi.org/10.1145/3278532.3278560>

- We characterise re-registrations based on the delay from the earliest possible instant, and show that there are qualitative differences in domain attributes and registrar behaviour.

2 BACKGROUND & RELATED WORK

Domain names need to be assigned to a registrant or domain owner before they can be used on the Internet. In each zone (under a top-level domain), domain registrations are recorded by a central registry, such as Verisign for .com domains. Typically, customers cannot access the registry system directly. Accredited registrars or resellers manage the domain registration on behalf of the customer.

Domain names are registered for a limited time and must be renewed regularly in order to remain active. When registrants do not renew their domain, the registration expires and after proceeding through a number of stages [11], the domain will eventually be deleted. The current state of a domain can be observed through the WHOIS protocol or its designated successor, the Registration Data Access Protocol (RDAP, RFC 7493 [4]). Available information includes the time the domain was registered, when the information was last updated, the identifier of the sponsoring registrar, as well as (often obfuscated) contact data of the registrant.

Once a domain has been deleted, it can be re-registered by any interested party on a first-come, first served basis. In the .com zone, between 66 k and 112 k domains are deleted each day, as shown in Figure 1. Some of these domains may be valuable, thus prospective registrants race to be the first to “catch” a “dropping” domain.

When exactly an expired domain is deleted depends on a number of factors. At a macro level, we showed in prior work [11] that a variable number of days elapse after a domain’s expiration date before the domain is deleted, which is due to differences in how registrars manage domains. At a micro level, we studied domain re-registrations on the same day the prior registration is deleted [10]. For .org and .biz domains, we found that the majority of same-day re-registrations occurred during a one to five-minute interval beginning at the same time each day. This interval is known as the *Drop*. Due to a lack of timestamps in WHOIS data at the time of that study, we were not able to conduct the same type of analysis for the far more popular (and competitive) .com domains.

Verisign does not publicly disclose details about how .com domains are made available for re-registration on the deletion day. From a number of anecdotal reports by participants of the ecosystem [6, 7, 9, 16], we infer the following likely characteristics: The Drop starts every day at 2 pm Eastern Time and lasts for around one hour, depending on how many domains are scheduled for deletion. Domain deletions are spread out in time over the duration of the Drop. They occur in a predictable order [1, 16], but to the best of our knowledge, how the order is created is not publicly known.

Drop-catch services allow customers to backorder domains that are about to be deleted, and compete with each other to re-register domains in the very instant they become available. To do so, drop-catch services send large quantities of speculative domain registration requests to the registry. In prior work, we found that registrars associated with drop-catch services have domain creation success ratios as low as 0.05 %, whereas regular registrars may exhibit success rates above 99 %. In order to be able to submit more requests,

some drop-catch services maintain large numbers of registrar accreditations. Three large drop-catch services together control 75 % of registrars, worth millions of dollars in accreditation fees [10]. A key element to allocating these resources efficiently is to precisely predict the time when a domain will become available. Therefore, it is likely that drop-catch services possess proprietary models of domain deletion times during the Drop.

Domain backorder fees from drop-catch services can be two to ten times more expensive than regular domain registrations. In order to avoid these fees, prospective registrants can attempt to re-register domains on their own. It may appear difficult to compete with the resources and expertise of drop-catch services, but only around 10 % of deleted domains are re-registered on the deletion day [10]. This leaves a large number of domains that can be re-registered at a lower cost. Software tools such as DropKing [5] cater to the niche market of “homegrown” drop-catching; they utilise reseller APIs of domain registrars for automation purposes.

In related work, domain name re-registrations have been discussed as potential [15, 18] or actual [8, 12] attacks. Salvador and Nogueira [17], and Miramirkhani et al. [13] studied how registrants select domains to re-register. Miramirkhani et al. found that shorter length, higher age, more traffic, and prior maliciousness all resulted in a higher re-registration probability. Furthermore, fewer than 11 % of re-registered domains hosted content, while the remainder was found to be re-registered for speculative or malicious purposes.

3 DATA COLLECTION

In this paper, we aim to characterise registrar behaviour during the Drop, and assess whether deleted domains are re-registered as early as possible. In contrast to prior work [10, 13], we measure at a time scale of seconds instead of days. The measurement is based on knowledge of the date when a domain is deleted, metadata about the prior registration, and the time when the domain is re-registered. Conceptually, our data collection methodology is similar to prior work [10], which we simplify by using more robust data sources.

Pending Delete Lists: In an effort to promote registrations, Verisign’s DomainScope [2] service publishes lists of domains that are scheduled to be deleted within the next five days. Since this service is offered by the .com and .net registry, we assume that the pending delete list is authoritative. We downloaded the list each day for a duration of eight weeks in the beginning of 2018; Figure 1 shows the date range and number of domains on the lists.

Domain Status: Three days before the scheduled deletion date of a domain, we requested the metadata of the expired registration. Since our last study, Verisign increased the precision of registration, update and expiration dates to timestamps. We collected them from Verisign’s RDAP test deployment [3], and fell back to WHOIS lookups in the rare case of errors. For example, domains sponsored by Papaki Ltd (registrar IANA ID 1727) resulted in HTTP 500 errors from the RDAP server, but had the expected data available using WHOIS. At least 8 weeks after the deletion date, we repeated the same lookups to collect metadata about any possible re-registration. Of a total of 4,599,802 .com domains from the pending delete lists, 512,802 (11.2 %) were re-registered on the deletion day and form the basis of our dataset. Because the volume of .net domains on the pending delete list was nearly an order of magnitude smaller,

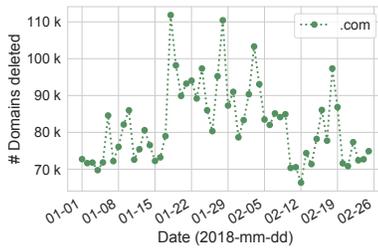


Figure 1: Expired .com domains deleted each day during our measurement period according to the pending delete list published by Verisign.

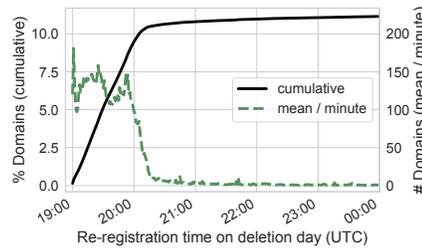


Figure 2: Re-registrations of .com domains on their deletion day. Most re-registrations happen during the Drop, from 19:00 until approx. 20:00 UTC.

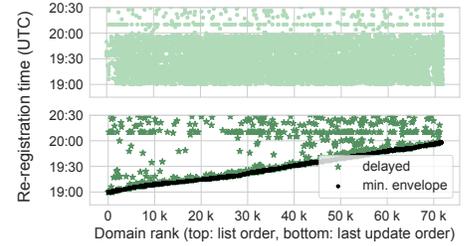


Figure 3: Scatterplots of re-registration time by domain rank in the pending delete list (above, incorrect order) and by last update time (below, likely deletion order).

we restricted our lookups to .com domains. In light of our findings in Section 4.1, this decision turned out to be unfortunate.

4 ANALYSIS

On the deletion day, most domain re-registrations occur during the time of the Drop. Figure 2 shows that no .com domain is re-registered before 7 pm UTC. One hour later, around 9.4% of all deleted domains have already been re-registered; this percentage increases to 11.2% by the end of the day. The one-hour period from 7 to 8 pm accounts for 84% of all same-day domain re-registrations.

During the Drop, same-day re-registration rates are high at over 100 domains per minute. Towards the end, the (aggregate) rate slowly decreases until it reaches just 3 re-registrations per minute at 9 pm on an average day. On any individual day, re-registration rates drop sharply immediately after the end of the Drop. The slower decrease in Figure 2 is due to aggregation of data over 56 days.

The duration of the Drop depends on the number of domains deleted. As shown in Figure 1, the length of Verisign’s pending delete list for .com ranges from 66 k to almost 112 k domains during our measurement period. The Drop may end before 8 pm on some days, and can last much longer on other days with more deletions. From the last observed drop-catch re-registration detected using the methodology from Section 4.2, we estimate that the longest Drop occurred on 18 January, the day with most domain deletions, and lasted until at least 20:49:48 UTC. Since the domains deleted last are not necessarily re-registered, the Drop is likely longer than our estimate. For example, 11 February had the shortest estimated duration (until 19:56:32 UTC), whereas the following day had fewer domain deletions, but an estimated duration until 19:58:29 UTC.

4.1 Domain Deletion Order

Online sources suggest that the order in which expired domains become available for re-registration during the Drop is predictable [1, 16]. However, we are not aware of any public source containing details about such a deletion order, presumably because this knowledge confers a competitive advantage to participants in the drop-catch race. If we plot domains in deletion order against their re-registration timestamp, we expect re-registrations to occur on or above a diagonal line, where the line corresponds to the earliest possible time, and the area above to delayed re-registrations. The area below the line should remain blank, as no domain is expected to be re-registered before the predicted earliest possible time. The upper part of Figure 3 shows such a plot using the order of

Verisign’s pending delete list. Since the re-registration points cover the entire area corresponding to the Drop period from 7 to 8 pm, it is evident that they are not in deletion order.

After similarly ruling out domain deletions ordered by domain ID, registrar ID, creation date, expiration date or alphabetical order, we find that deletions likely occur ordered by the domain registration’s “last updated” time. Since timestamps in our dataset are at a second precision and many registrars update large batches of domains at the same time, this ordering can be ambiguous. A secondary sorting criterion of creation timestamp or domain ID appears to work well; we opt for domain IDs because they induce a total order. The lower half of Figure 3 uses this order to visualise deletion day re-registrations on 2 January 2018. Around 80% of same-day re-registration points appear on the highlighted diagonal, and none below. We conclude that a deletion order exists, and that it is closely approximated by the update time and domain ID ordering.

For each deletion day, we re-order Verisign’s pending delete list and assign each domain a rank based on the inferred deletion order. The heatmap in Figure 4a aggregates the result over all 56 days. Diagonal lines of different lengths correspond to days with different quantities of domains being deleted. The diagonal lines do not perfectly overlap, appearing to have different slopes, and they are not entirely straight as one may expect. We hypothesise that Verisign may have a single domain deletion process for both .com and .net domains combined, as suggested by DnMeter [1]. We did not collect data for .net domains, and our computed domain ranks do not account for .net domains that may be interleaved with the .com domains. Even though their volume is comparatively low, they may appear in batches in the deletion order and make the curve deviate from an ideal straight line.

Most of the domains appear to be re-registered as early as possible, as the highest density is on the diagonal lines. Delayed domain re-registrations (above the diagonal) occur at a rate at least one order of magnitude lower than during the deletion instant.

To observe the behaviour of services across all of the registrars they control, we reuse the methodology from prior work [10] and cluster registrars according to their contact details. For re-registrations on the deletion day, DropCatch (not shown) and SnapNames (Figure 4b) are by far the most popular registrar clusters. Both are well-known drop-catch services, and appear to focus all their efforts during the Drop on early re-registrations, as indicated by the dark diagonal lines and the empty area above. The horizontal lines around 8.30 pm and later indicate domains re-registered in

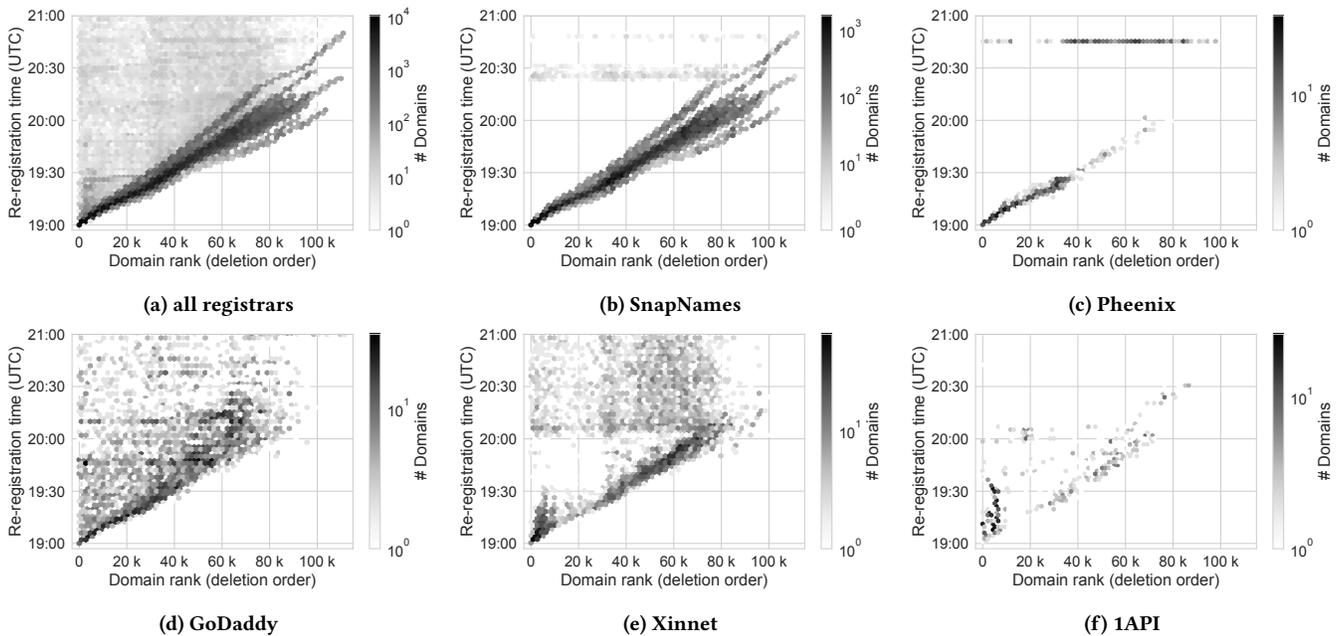


Figure 4: Heatmaps of re-registration times by the domain’s deletion order rank, aggregated over 56 deletion days; intensity indicates number of re-registrations per bin (log scale, different for each subplot). Domains re-registered as early as possible appear on diagonal lines, delayed re-registrations above the lines, and re-registration batches independent of original deletion times as horizontal lines. Most re-registrations are as early as possible. Drop-catch services such as SnapNames (b) or Pheenix (c) dominate re-registrations and exhibit behaviour distinct from regular registrars (e.g., d) or API providers (f).

batches independent of the original deletion order; the timing suggests that such delayed re-registrations are held back until the end of the Drop. Figure 4c shows similar trends for Pheenix, a drop-catch service that appears to be less active.

GoDaddy, depicted in Figure 4d, does not have such a strong focus on the diagonal; instead, re-registrations appear all over the area above, suggesting that domains are re-registered seconds, minutes and hours after the deletion time. This may be explained by GoDaddy’s role as the globally most popular domain registrar. Xinnet (Figure 4e) combines features of drop-catch and regular registrars, with a distinguishable, though more spread out diagonal line, suggesting re-registrations close to the earliest possible instant, but not as timely as those made by SnapNames, for instance. Re-registrations delayed by 30 minutes or more occur even during the Drop, although not as frequent as for GoDaddy; more delayed re-registrations start shortly after 8 pm and last well beyond 9 pm. Some registrars such as 1API offer APIs that customers could use for their own drop-catch scripts. Figure 4f confirms that customers indeed use the API for same-day re-registrations, but (in this case) not with the same scale or timeliness as drop-catch services. Precise prediction of when domains will be deleted appears to be a well guarded “secret” of the drop-catch trade.

4.2 Earliest Re-Registration Times

To analyse the timeliness of domain re-registrations, we need to infer the first moment when a domain can be re-registered. We assume that competition causes drop-catch services to attempt re-registrations immediately upon deletion, which means that our dataset contains many observations of earliest possible re-registrations.

Furthermore, we assume that domains are deleted in a predictable order, which implies that the subset of domains re-registered as early as possible must have monotonically increasing timestamps, as visible in the dark diagonal in the lower half of Figure 3. To predict the earliest possible re-registration time of a domain given its rank in the deletion order, we need to model this diagonal.

The diagonal suggests a linear relationship between the deletion rank and the earliest re-registration time, which could be modelled using linear regression. Yet, the deviations from a straight line seen in Figure 4a can cause errors in the order of minutes. Furthermore, a priori prediction is not necessary for our purposes; it is sufficient if we can use the observed data a posteriori. Instead of attempting to fit a straight line, we identify the domains that have been re-registered as early as possible by computing the “minimum envelope” curve of the scatterplot. Specifically, we look for a sequence of domain re-registrations in deletion order such that their re-registration timestamps are monotonically increasing, and minimal. Iterating over ranks from right to left, we retain any re-registration if its delay from 7 pm is no larger than the value previously added to the curve. This approach could lead to outliers at the right end of the curve. For example, if the domain with the numerically highest rank is re-registered with a large delay, it should not become part of the curve because it does not indicate an as-early-as-possible re-registration instant. To address this issue, we additionally truncate the right end of the curve wherever the time delay between two consecutive points on the curve is larger than one minute, which has proven to be a good indicator for the end of the Drop.

We calculate this minimum envelope curve separately for each deletion day to obtain better accuracy, given our observation of

different Drop durations, different slopes and imperfect linearity in Figure 4a. As an illustration, the lower half of Figure 3 highlights the minimum envelope for 2 January 2018. Each day has a median of 7.6 k points on the envelope curve; the delay between two consecutive points is 3 s or less for 99 % of points on the curves, with a maximum of 38 s. Figure 7 shows that nearly all re-registrations on the curve come from drop-catch services, which makes us confident that our strategy closely models the actual domain deletion curve.

To infer the earliest possible re-registration time for any domain given its rank in the deletion order, we return the actually observed re-registration time if the rank corresponds to a point on the deletion day’s minimum envelope curve, which occurs for 52 % of the re-registered domains in our dataset. For 48 % of domains, we apply linear interpolation between the two neighbouring points with the closest smaller and higher ranks on the curve. Typically, these points are no more than 3 s apart, and we round the interpolated time to the closest second in order to remain consistent with the precision of the original timestamps in the RDAP data. Only 0.02 % of re-registered domains in our data have ranks outside the range of their minimum envelope curve, which causes us to use the first or last re-registration time from the curve. The *re-registration delay* is the time difference between the inferred earliest possible re-registration, and the actually observed re-registration timestamp.

4.3 Re-Registration Delays

About 9.5 % of deleted .com domains are re-registered with a delay of 0 s, that is, in the instant we predict as their deletion time. This percentage grows to 13 % for domains re-registered with a 24 h delay (600 k domains in our dataset). As the detail in Figure 5 shows, growth is fast for the first 30 s, but then flattens out at a microscopic scale. Between 3 h and 8 h after deletion, there is a second relatively fast period, adding re-registrations for another 1 % of deleted domains. These effects are the aggregate of registrar and customer behaviour, and correspond to different re-registration strategies.

Figure 6 shows when a selection of registrar clusters re-register their domains during the first 24 h after deletion. Drop-catch services tend to re-register the majority of their domains with very short delays. DropCatch, the most active re-registration cluster, re-registers 99.3 % of its domains with a delay of 0 s (relative to all its domains re-registered within 24 h of deletion). XZ and Pheenix also re-register a majority of their domains at 0 s, but they keep adding more re-registrations during the following seconds. For example, XZ goes from 74.8 % at 0 s to 89.4 % at 3 s. This illustrates that drop-catch services are not equally timely in their re-registrations.

Our delay metric allows for a more precise detection of drop-catch re-registrations than the approximation used in prior work. If we consider as drop-catch re-registrations only delays of 3 s or less, about 86.1 % of deletion day re-registrations fall into this category. Consequently, labelling all domains re-registered on the deletion day as drop-catch, as done in prior work [10], results in “false positive” misclassification of 13.9 % of domains since they are not re-registered during the most competitive part of the drop-catch race. Another conceivable heuristic would be to label deletion day re-registrations as drop-catch only when they are re-registered during a typical Drop period (e.g., 19:00:00–19:59:59 UTC). However, this approximation is also problematic. Almost 9.5 % of deletion

day re-registrations are drop-catch re-registrations that happen after 8 pm due to the variable duration of the Drop, and would not be detected by this heuristic (false negatives). Another 7.4 % of deletion day re-registrations happen between 7 and 8 pm, but are not drop-catch re-registrations because their re-registration delay is larger than 3 s (false positives). Even though these domains are re-registered while the Drop takes place, they are not re-registered as early as possible, but at a time when the competitive re-registration race has already moved on to domains further down in the deletion order. This illustrates the utility of our re-registration delay metric.

Drop-catch services do not necessarily re-register all of their domains with short delays. Pheenix has another steep increase in re-registrations 30–90 min after domain deletion, likely corresponding to the horizontal line visible in Figure 4c. We do not know whether customers backordered these domains before the Drop and Pheenix postponed their registration; they could also correspond to regular domain orders placed during or after the Drop. Dynadot does exhibit some drop-catch activity, but the bulk of re-registrations appears at longer time scales, thus likely initiated directly by customers.

While GoDaddy re-registers some domains in the seconds after deletion, there does not seem to be significant drop-catch activity, and the vast majority of domains are re-registered hours later. Xinnet follows a similar pattern at a scale of hours; however, very few re-registrations occur until 10 s after deletion. Figure 4e suggests that Xinnet may hold back re-registrations with longer delays until the end of the Drop, which resembles characteristics of drop-catch services, albeit with a more modest outcome. Similarly, most re-registrations using 1API happen relatively early, with the median at 26 min, but they do not start until 30 s after deletion.

The examples above illustrate different re-registration strategies for deleted domains. Drop-catch services are the first to re-register domains, and there are visible differences in timing and quantity among them. To some extent, these differences are also reflected in pricing; backorders at Dynadot, for instance, are priced lower than at DropCatch. “Home-grown” drop-catching through APIs follows seconds to minutes later, yielding domains not taken by the drop-catch services at the cost of regular domain registrations. Hours later, remaining domains may be re-registered in batches.

4.4 Delay Interval Analysis

Domains re-registered through different means (and at different cost) suggest that domains may have different qualitative attributes depending on their delay. To analyse how those domains differ from each other, we group re-registration delays into intervals. We choose the duration of these intervals such that each one contains at least 8 k domains. Some intervals contain many more domains; we cannot subdivide them further due to the precision of the domain creation timestamps released by the registry. For each interval, we then compute the market share of different types of domains.

Figure 7 shows the market share of a selection of registrar clusters in each interval. DropCatch and SnapNames dominate the 0 s interval, but re-register a much smaller share of domains afterwards, until 8–10 min after deletion, when DropCatch reaches another momentarily high market share. Xinnet reaches a high market share of 50 % 1–9 h after deletion. No single registrar consistently dominates re-registrations across different delays.

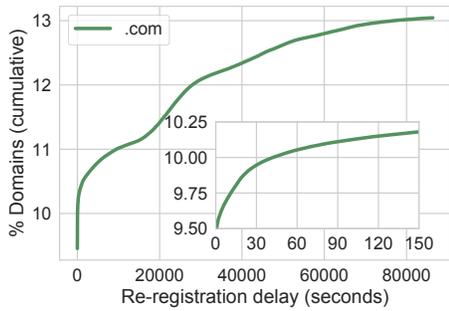


Figure 5: CDF of re-registration delays during 24 h after the deletion time (with 2.5 min detail). Around 9.5% of all deleted .com domains are re-registered instantly.

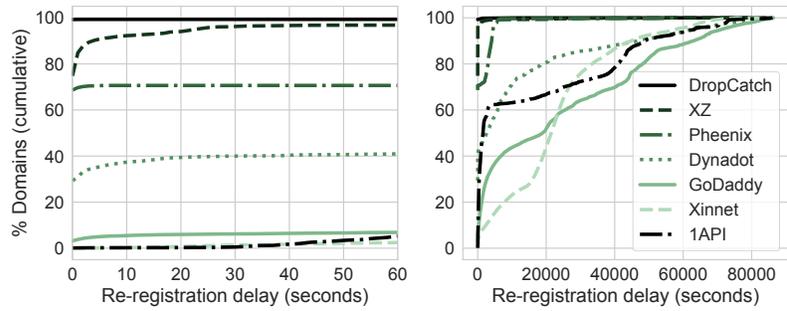


Figure 6: CDF of re-registration delays for registrar clusters, relative to 24 h after deletion (left: detail of the first 60 s). Drop-catch services re-register with short delays (DropCatch, XZ, Pheenix), other registrars peak hours later (GoDaddy, Xinnet), and some exhibit both behaviours (Dynadot, 1API).

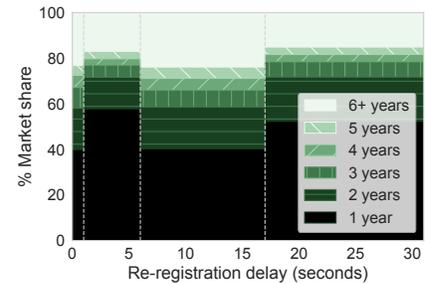
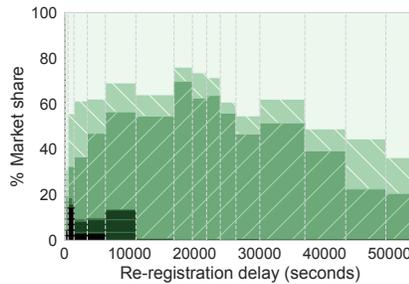
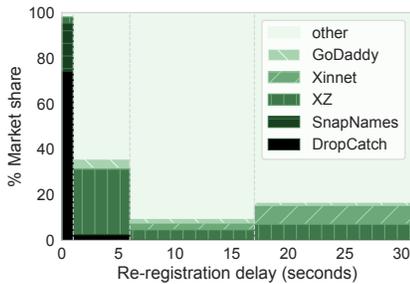


Figure 7: Market share of registrar clusters during variable-length intervals comprising at least 8 k total re-registrations (denoted by vertical lines). Two drop-catch services, DropCatch and SnapNames, dominate the first second (left detail), whereas Xinnet holds over 50% market share from 1–9 h after deletion.

Figure 8: Interval market share of the prior domain age of re-registrations. Older domains peak early, at delays of 0 s and 6–16 s, suggesting higher perceived value.

When we consider the market share in terms of the age of the domains being re-registered (Figure 8), we observe that re-registrations of older domains peak at delay 0 s. A higher prior registration age may be an indicator for a more valuable domain, as it may receive more traffic due to pre-existing inbound links [10, 13]. However, our data show that re-registrations of older domains peak not only at the beginning, but also at 6 s, 1 h, and between 13–14 h after deletion. We observed similar effects, albeit at different intervals, when comparing the number of keywords and English dictionary words contained in re-registered domain names. We hypothesise that these correspond to different actors registering lists of attractive domains that were not claimed by drop-catch services.

In the beginning of May, we looked up all re-registered domains using Google’s Safe Browsing API as an indicator for malicious behaviour. At that time, each domain had been re-registered for at least 9 weeks. We find that a majority of domains later labelled as malicious were registered by drop-catch services with a delay of 0 s. This may seem surprising given the higher cost, but Safe Browsing does not elucidate the origin of the maliciousness, nor can we infer the intent of the domain registration. It is a common occurrence that domains are parked after re-registration [13], and they may have inadvertently displayed malicious advertising campaigns without necessarily being re-registered for that purpose.

In terms of market share, however, the situation is different. Only 0.4% of domains re-registered with a delay of 0 s are labelled as malicious. From 30–60 s after deletion, the percentage of malicious domains registered during those intervals reaches 2%, but it corresponds to only around 250 domains. Overall, fewer than 0.5% of

domains re-registered within 24 h of deletion are labelled as malicious. Therefore, we caution that these results should be seen as preliminary and to be further investigated in future work.

5 CONCLUSION

In this paper, we studied domain registrations during the Drop, when expired domains are deleted and can be re-registered for the first time. While we were not able to measure registration attempts made by drop-catch services, we characterised the final outcome of the race to re-registration. Leveraging our technique to infer the deletion time of domains during the Drop, we showed that most of the re-registrations happen at the earliest possible time. Instant re-registrations also tend to be the ones with the oldest domains and most keywords. We observed a range of re-registration behaviours, including the use of reseller APIs for “home-grown” drop-catching seconds to minutes after the Drop, and re-registration of deleted domains in large batches hours later. Based on the inferred deletion time, we proposed a precise metric to detect drop-catch re-registrations. Knowledge of when or how domains were registered may be useful in research into uses of domains.

ACKNOWLEDGEMENTS

The authors would like to thank Cameron Walters and Eric Case of Domainr for suggesting Verisign’s RDAP test deployment for data collection, as well as the anonymous reviewers and the shepherd Kensuke Fukuda for their valuable feedback. This work was supported by the National Science Foundation under grant CNS-1703454.

REFERENCES

- [1] [n. d.]. Domain Meter. <http://www.dnmeter.com/>.
- [2] [n. d.]. Verisign DomainScope. <https://www.domainscope.com/>.
- [3] [n. d.]. Verisign Registration Data Access Protocol (RDAP) Pilot. <https://rdap-pilot.verisignlabs.com/>.
- [4] 2015. JSON Responses for the Registration Data Access Protocol (RDAP). <https://tools.ietf.org/html/rfc7483>.
- [5] 2016. Drop King. <http://dropking.com/>.
- [6] Michael Cyger. 2016. A Drop Catching Programming Expert Discusses the Domain Name Expiration Process - With Chris Ambler. <http://www.domainsherpa.com/wp-content/pdf/Chris-Ambler-Expiration-on-DomainSherpa.pdf>.
- [7] DropCatch. [n. d.]. How it Works: Daily Drop Overview. <https://www.dropcatch.com/HowItWorks/Overview>.
- [8] Shuang Hao, Matthew Thomas, Vern Paxson, Nick Feamster, Christian Kreibich, Chris Grier, and Scott Hollenbeck. 2013. Understanding the Domain Registration Behavior of Spammers. In *Proc. of ACM Internet Measurement Conference*.
- [9] Ron Jackson. 2004. Inside a Drop Catcher's War Room: How Enom Arms Maker Chris Ambler Is Turning The Tide for Club Drop. <http://www.dnjournal.com/columns/cover080504.htm>.
- [10] Tobias Lauinger, Abdelberi Chaabane, Ahmet Buyukayhan, Kaan Onarlioglu, and William Robertson. 2017. Game of Registrars: An Empirical Analysis of Post-Expiration Domain Name Takeovers. In *Proc. of Usenix Security Symposium*.
- [11] Tobias Lauinger, Kaan Onarlioglu, Abdelberi Chaabane, William Robertson, and Engin Kirda. 2016. WHOIS Lost in Translation: (Mis)Understanding Domain Name Expiration and Re-Registration. In *Proc. of ACM Internet Measurement Conference*.
- [12] Chaz Lever, Robert J Walls, Yacin Nadji, David Dagon, Patrick McDaniel, and Manos Antonakakis. 2016. Domain-Z: 28 Registrations Later – Measuring the Exploitation of Residual Trust in Domains. In *Proc. of IEEE Symposium on Security and Privacy*.
- [13] Najmeh Miramirkhani, Timothy Barron, Michael Ferdman, and Nick Nikiforakis. 2018. Panning for gold.com: Understanding the Dynamics of Domain Dropcatching. In *Proc. of World Wide Web Conference*.
- [14] NamePros. 2015. DesktopCatcher software. <https://www.namepros.com/threads/desktopcatcher-software.873819/>.
- [15] Nick Nikiforakis, Luca Invernizzi, Alexandros Kapravelos, Steven Van Acker, Wouter Joosen, Christopher Kruegel, Frank Piessens, and Giovanni Vigna. 2012. You Are What You Include: Large-scale Evaluation of Remote JavaScript Inclusions. In *Proc. of ACM Conference on Computer and Communications Security*.
- [16] Dan Ramirez. 2013. A New Twist On Drop Catching Technology - Focus on the Drop Order. <http://www.domainafterlife.com/2013/08/a-new-twist-on-drop-catching-technology-focus-on-the-drop-order/>.
- [17] Paulo Salvador and António Nogueira. 2011. Analysis of the Internet Domain Names Re-registration Market. *Procedia Computer Science* 3 (2011), 325–335.
- [18] Johann Schlamp, Josef Gustafsson, Matthias Wählisch, Thomas C. Schmidt, and Georg Carle. 2015. The Abandoned Side of the Internet: Hijacking Internet Resources When Domain Names Expire. In *Proc. of International Workshop on Traffic Monitoring and Analysis*.