Department of Informatics
Technical University of Munich

Technical University of Munich

Department of Informatics

Bachelor's Thesis in Informatics

Analyzing the Effect of Domain Parking on DNS Based Research

Steffen Deusch

Technical University of Munich

Department of Informatics

Bachelor's Thesis in Informatics

Analyzing the Effect of Domain Parking on DNS Based Research

Analyse der Auswirkungen von Domain Parking auf DNS-basierte Forschung

Author:	Steffen Deusch
Supervisor:	Prof. Dr.-Ing. Georg Carle
Advisor:	Johannes Zirngibl
	Patrick Sattler
	Juliane Aulbach
Date:	February 15, 2022

I confirm that this Bachelor's Thesis is my own work and I have documented all sources and material used.

Garching, February 15, 2022

Location, Date

Signature

Abstract

Ever since the commercialization of the domain name ecosystem, domains are a valuable resource as only one individual or company can own a single domain. This makes it possible to buy domains in large quantities and later sell them for a higher price, reaching up to tens of millions of dollars in sales. Instead of leaving these domains unused, domain parking services offer a way to monetize them, primarily by showing advertisements or redirecting users to an advertiser's landing page.

By analyzing the domain parking ecosystem and collecting DNS configurations for 82 different parking services, we are able to identify the presence of domain parking in a dataset of more than 130 B DNS records and show that these services hold a substantial amount of domains with the possibility of impacting research based on domain names massively. Out of 267 M unique resolving domains in our dataset, nearly 62 M can be classified as parked. We examine the development of domain parking over the year 2021, seeing a cumulative amount of more than 106 M parked domains. Not only are parked domains found in large quantities in the new gTLDs, but even more so in the important TLDs *.com*, *.net* and *.org*, as well as in top lists that are frequently used in scientific research.

Due to their content, showing only advertisements, sales pages or generic placeholders, parked domains differ greatly from other web content and should therefore be considered as a possible source of bias. We discuss the implications of parked domains on DNS based research and show possible areas where further research is needed.

Zusammenfassung

Seit der Kommerzialisierung des Domainökosystems sind Domains eine wertvolle Ressource, da nur eine Person oder Unternehmen eine einzelne Domain besitzen kann. Dadurch ist es möglich, Domains in hoher Stückzahl zu kaufen und diese später, für einen höheren Preis, weiterzuverkaufen. Manche Domainverkäufe erreichten bereits zweistellige Millionenbeträge. Anstatt diese Domains ungenutzt zu lassen, bieten Domain Parking Dienste die Möglichkeit, diese zu monetarisieren, primär durch das Schalten von Werbeanzeigen oder der Weiterleitung auf die Seite des Werbetreibenden.

Durch Analyse des Domain Parking Ökosystems und Sammeln von DNS Konfigurationen für 82 verschiedene Parking Dienste, sind wir in der Lage, die Präsenz von Domain Parking in einem Datensatz von mehr als 130 Milliarden DNS Einträgen zu identifizieren und zu zeigen, dass diese einen beträchtlichen Anteil an Domains besitzen, mit der Möglichkeit, auf Domainnamen basierende Forschung massiv zu beeinflussen. Von 267 Millionen eindeutigen, auflösenden Domains in unserem Datensatz klassifizieren wir 62 Millionen als geparkt. Zudem untersuchen wir die Entwicklung von Domain Parking über das gesamte Jahr 2021 mit insgesamt 106 Millionen gesehenen geparkten Domains. Geparkte Domains sind nicht nur in den neuen gTLDs in großer Zahl zu finden, sondern vor allem auch in den wichtigen TLDs *.com*, *.net* und *.org*, sowie in Toplisten, die häufig als Basis für wissenschaftliche Forschung verwendet werden.

Durch ihren Inhalt, welcher lediglich aus Werbeanzeigen, Verkaufsseiten oder Platzhaltern besteht, unterscheiden sich diese geparkten Domains stark von anderen Webinhalten und sollten daher als mögliche Quelle für Verzerrungen in Betracht gezogen werden. Wir diskutieren die Auswirkungen von geparkten Domains auf DNS basierte Forschung und zeigen auf, in welchen Bereichen weitere Forschung nötig ist.

Contents

List of Figures

List of Tables

Chapter 1

Introduction & Motivation

The Domain Name System (DNS) is one foundation of the Internet, responsible for resolving domain names to other useful information, like IP addresses. Usually, these domains are registered on a "first come, first served" basis and can be registered and used by everyone, from businesses to private home pages. However, they are sometimes acquired without the intent to actually use them. There are "domainers", a term used for people that own a huge portfolio of domains, speculating to sell these domains later for a higher price. Instead of leaving these domains completely unused, *Domain Parking* is a strategy to use these domains, most often for showing advertisements targeted at the domain name, thus earning passive revenue for the domain owners. As an example, the domain parking and marketplace provider *Sedo* claims to sell 19 M domains on their about us page [1].

For reasons of profit, these domain parking services tend to host a large amount of domains on the same infrastructure, having the same properties. While this is technically similar to large shared hosting providers and content delivery networks, parked domains only contain advertisements, sales pages or generic placeholders with questionable value for Internet users. Depending on their share of managed domains, these services might impact Internet-wide service studies based on domains massively.

1.1 Goals

Existing work [2] already showed that domain parking occurs in DNS scans. This thesis expands this analysis with a focus on three main aspects.

1. Identifying domain parking services, their mode of operation, similarities and differences among each other as well as the company structures behind them, resulting in a comprehensive overview of the most important domain parking services and the domain monetization ecosystem in general.
2. Finding a way to reliably detect these services in large-scale DNS scans and building an analysis pipeline that is extensible for future work.
3. Analyzing parked domains and their impact on DNS based research, e.g., Transport Layer Security (TLS) scans based on domain lists.

1.2 Outline

The thesis starts with necessary background knowledge in Chapter 2 and presents important related work in Chapter 3. We introduce and explain our methodology in Chapter 4, focusing on the identification of domain parking services. In Chapter 5, we evaluate our results concerning domain parking in the chair's dataset and continue in Chapter 6 to discuss effects of domain parking on further research. We conclude the thesis in Chapter 7.

Chapter 2

Background

Domain Parking relies on the DNS for its operation, therefore we present a brief overview of the DNS and introduce domain parking, its terminology and monetization schemes as well as the existing DNS scan procedure that this thesis relies on for acquiring input data.

2.1 DNS

The DNS is a hierarchical naming system used for identifying services on the Internet. It can be described as a key-value storage for resolving *domain names*, labels separated by dots, into values like Internet Protocol (IP) addresses. The DNS hierarchy is based on a tree structure of *zones*, beginning at the *root zone* (`.`), served by 13 root servers, delegating authority over top-level domains (TLDs) like `com.` to other *name servers.* These *authoritative* name servers for `com.` then delegate authority over second-level domains like `google.com.` to the respective authoritative name servers, which are responsible for storing the *resource records* for their zone. For this thesis there are four important record types:

NS (name server) records that are used to delegate authority over a zone to other name servers.

A (address) records that specify the IPv4 addresses for a specific domain.

AAAA records that specify the IPv6 addresses for a specific domain.

CNAME (canonical name) records that specify an alias domain name. A resolver will restart the query using the CNAME value.

Querying the DNS is done by using *resolvers*, either using an iterative, or a recursive lookup. An iterative lookup starts at the root zone and repeatedly asks the authoritative name server for each zone, until the answer is found. A recursive query asks a recursive resolver to do the lookup, which can cache records based on their time to live (TTL) value.

2.2 Domain Parking

Domain names (also shortly referred to as *domains*) are usually assigned on a "first come, first served" basis, where the domain registrant owns the domain until they stop paying a registration fee. This practice, combined with the fact that each fully qualified domain name (FQDN) is unique, can make domain names a highly valuable and contested commodity, with names like `carinsurance.com` changing owners for multiple tens of millions of dollars [3]. Therefore, domain name monetization can be a profitable business, with domain investors (also called "domainers") buying and selling domains in large quantities. Making profit with domains can be achieved in multiple ways.

One option is to sell a domain to a new owner for a higher price. To advertise the possibility to buy the domain, special sales pages can be used, often with the option to contact the domain owner or directly buy it through a domain marketplace. Another option to monetize otherwise unused domains is using advertisements when a user visits the domain in a web browser.

The process of showing pages with no real, valuable content on these domains is referred to as *domain parking*, as the domains are usually "parked" at a service provider, further referred to as *parking service*, that takes care of it until the owner decides to sell it or put it to better use. These services provide the DNS and HTTP infrastructure to serve web requests directed at the parked domains.

As long as enough users click on the advertisements to generate more revenue than the registration fee for the domain, the domain owner makes profit, while the parking provider takes a commissioning fee. One usually distinguishes between two types of advertisement based monetization options.

Pay-Per-Click (PPC) is a monetization scheme where revenue is generated once a user actively clicks on an advertisement banner or keyword and is then redirected to the advertiser's landing page.

Pay-Per-Redirect (PPR) is a monetization scheme where a user is immediately redirected to the advertiser's landing page without any manual action.

Parked domains usually employ a PPC based monetization scheme and sometimes fall back to PPR monetization when the service's advertisement partner does not want to receive specific traffic. Most parking services that use advertisement based monetization also offer a way to inform visitors about the possibility to buy the parked domain using varieties of sales banners, while services specialized on domain sales typically do not offer advertisements. This is the case because large scale advertising usually requires upstream advertisement providers like *Google*, *Yahoo* or *Bing*.

Domain owners that want to park their domains at a parking service have multiple ways of doing so. Most frequently, DNS records are used for directing their domains to the service's infrastructure. Sometimes HTTP redirects or HTML features such as *iframes* are employed. Using the DNS, the options are:

1. Delegating authority of the domain to the parking service's name servers using `NS` records, thus giving the parking service complete control over the domain.
2. Setting the parking service's IP addresses using `A` / `AAAA` records.
3. Pointing the domain to the parking service using a `CNAME` record.

Concerning the third option, it is technically invalid to use a `CNAME` record at the root of a zone. According to RFC1034 [4] a `CNAME` record cannot coexist with other data at a specific node, which would be violated by the `SOA` record that is mandatory for the root of a zone. Still, we see this in practice, not only when looking at parked domains. Some DNS servers provide a mechanism referred to as *CNAME flattening*, *alias records* or *ANAME records*, that resolves the alias before returning an answer, which is not visible to a client.

In the past, parked domains often attracted traffic by users directly typing the domain name into their browsers. With the popularity of search engines nowadays, this so-called "type-in traffic" [5] dwindled in significance. Google itself announced a classifier for parked domains in 2011, as "[p]arked domains are placeholder sites with little unique content for our users and are often filled only with ads" and it "prefer[s] not to show them" [6]. Parked domains that were previously registered can benefit from backlinks on other popular sites that still point to the domain. Another, highly controversial method of attracting users is by utilizing typing errors and registering so-called "typo domains", e.g., `youtubee.com` as a typo of `youtube.com`. This is also referred to as "typosquatting". A similar method that generally refers to registering domain names in bad faith, e.g., by consciously infringing trademarks, is called "cybersquatting" and also often comes up in the realm of parked domains [7].

2.3 DNS Scans

This thesis heavily relies on existing, regular large-scale DNS scans executed at the chair's *Global INternet Observatory* project. Therefore, we briefly present the used input data, result files and scan procedure itself.

2.3.1 Used Input Lists

The current DNS scan setup at the chair uses multiple sources for obtaining domains to be resolved, further called *input lists*. These lists can be categorized into the categories *zone files*, *top lists* and *other*. A full overview of the used input lists is given in Table 2.1.

Table 2.1: Used input lists. Lists that are updated regularly are marked with ↻, lists that are static are marked with ☐.

	Name	**Size (domains)**[1]
	Zone files	
↻	.com, .net, .org	187.0 M
↻	other gTLDs	35.4 M
↻	.ch	2.6 M
↻	.se	1.4 M
↻	.nu	247.8 k
	Top Lists	
↻	Majestic Million	1.0 M
↻	Cisco Umbrella	1.0 M
↻	Alexa Top 1 Million	567.5 k
↻	Alexa Country	2.4 k
	Other	
☐	Certificate Transparency Log	158.2 M
☐	ccTLD[2]	98.1 M
↻	Blocklists	4.3 M
☐	Chrome UX	3.3 M
↻	Chromium Preload	153.0 k

[1] The size of the lists was calculated by averaging the values from 2021-11-01 to 2021-11-30.
[2] This is only a domain list, **not** a zone file.

A zone file offers a complete view of a single DNS zone, e.g., all records at the root of a single TLD. Zone files for most generic top-level domains (gTLDs) can be acquired from the Centralized Zone Data Service (CZDS) [8], a service offered by the Internet Corporation for Assigned Names and Numbers (ICANN) where researchers and other parties can get access to zone data.

These zone files include the most popular gTLDs `.com`, `.net` and `.org` as well as most other new gTLDs. They are essential for obtaining a good overview of second-level domains (e.g. `google.com`) that are actively in use and therefore have an entry in the DNS, but because of the hierarchical structure of the DNS, they do not include subdomains like `www.google.com`.

Top lists like the *Alexa Top 1 Million* [9] provide a ranking of popular domains and are frequently used as a basis of scientific research for acquiring domains visited by real users. The details how these top lists obtain their ranking is not important for this thesis, but it is different for each list [10].

To get a more complete picture of domains in use, other lists are also included in the DNS scans. These include a static list of country code top-level domains (ccTLDs) that are not available from the CZDS, as well as a combined list of domains obtained from *block lists*.

Another way of acquiring domains in use is looking at *Certificate Transparency* logs. These logs contain information about issued certificates from all major *certificate authorities*, therefore also including domains from the *subject alternative name* field. Complementary to domains obtained from zone files, domains in these *other lists* often also include subdomains, like `www.google.com`.

2.3.2 Existing Scan Procedure

Starting with the *input lists*, the obtained domains are then resolved using *MassDNS*[1] and a local *Unbound*[2] DNS resolver. We call the results of these DNS scans *resource record* files in the following sections and chapters.

The resource record files obtained from these scans have a format which is similar to the output of the *dig* DNS lookup utility and the format described in `RFC1034` [4]:

```
...
www.example.tld. IN 300 A 10.0.0.1
foo.example.tld. IN 3600 CNAME bar.example.tld.
bar.example.tld. IN 60 A 10.0.0.2
...
```

[1] https://github.com/blechschmidt/massdns

[2] https://www.nlnetlabs.nl/projects/unbound/about/

Every line in the output file is one DNS record, which consists of the following five fields. The first field is the *owner* (i.e., domain name) the record belongs to. Next is the *record class*, which is always the value `IN` for these scans. The next field is the TTL, which specifies how long a record may be cached by resolvers, followed by the *record type*, which specifies how to interpret the following *data* field.

For every *input list* and every *record type* there are multiple output files that are then further processed for other purposes. For this thesis, the relevant files are the unfiltered *resource records* for the `A`, `AAAA` and `NS` record types. While `NS` records are only resolved for some specific input lists, there is a combined *merged* scan for `NS` records that we use in this thesis.

As domains can also use `CNAME` records to specify a domain alias which can recursively use a `CNAME` record as well, there exists a *CNAME chasing* procedure, which starting at the resource record files processes these chains and generates a new output file with the effective domain to IP address mapping, which we will further call *ip-domain* files.

Chapter 3

Related Work

This chapter reviews important related work in the field of domain parking and the domain ecosystem.

3.1 Domain Parking

In 2014, Alrwais et al. [11] analyzed domain parking from the customer perspective and focused on possibly illicit activities. They registered accounts at several parking services and advertisement networks. Parking some of their own domains and then creating advertisement campaigns that targeted the traffic from their parked domains, they were able to analyze the complete monetization chain, uncovering potential cases of click fraud, traffic spam and traffic stealing.

To collect their data, the researchers obtained a list of parked domains from an external reverse NS dataset. Then they used a browser extension to automatically crawl the parked domains, gathering HTML content and HTTP traffic into a database with each visit corresponding to a unique *monetization chain*, the sequence of URLs recorded on each visit.

According to the authors, fraudulent activities are only prevalent on parked domains that use PPR monetization and of their 24 M visits on parked domains, only 1.2 M lead to redirections. They presume that this is related to the classification of domains into "primary" and "secondary" classes, depending on their acceptance by top-tier search networks like *Google*. When it comes to malware distribution, the authors assume that the parking services are probably not intentionally involved, but only because of a lack of verification of placed ads by the used advertisement networks. By using estimates, they conclude that these illicit activities could account for up to 40 % of some services' revenue. Further, they conclude that because of the complexity of the observed redirection chains, even when malicious activity is confirmed, it often cannot be clearly attributed to a responsible party.

While the researchers did not actively quantify parked domains themselves, they used numbers from a service called *DailyChanges* to estimate the total number of parked domains, acknowledging that this list is not comprehensive.

Vissers et al. [12] focused on the analysis of parked domains from the user perspective visiting a parked domain. Gathering a list of 15 parking services from sources like search engine results, forums and surveys, they then proceeded to collect the DNS configurations, including `NS`, `A` and `CNAME` records required for setting up a domain with the corresponding services. To gather these, they also registered accounts at multiple services. Matching these configurations against the DNS Census dataset [13] and manually resolving domains to verify that they are still parked, they were able to gather a list of 8 064 914 parked domains. They acknowledge that this list is outdated and incomplete and therefore conclude that there are *at least* eight million domains associated with domain parking.

The found domains were then analyzed for *typosquatting abuse*, where a domain is registered with the intent to attract traffic from users mistyping a legitimate website's address. All parking services except one had at least a small amount of typosquatting domains associated with them and to verify, the researchers tried to park a typosquatting domain themselves with success at every parking service that accepted their account registration.

The researchers actively crawled the domains using an automated web browser and saved the resulting HTML content as well as the HTTP requests done when loading each page. Therefore, they were able to measure not only the parking services, but also the used advertisement networks, e.g., *Google Adsense*[1]. Some parking services also used PPR monetization where a user is directly redirected to the advertiser's landing page. Similar to the paper by Alrwais et al. [11], the researchers were able to show malicious redirects from parked domains to pages containing malware, scams or inappropriate sexually explicit content at most parking services that used PPR monetization.

Motivated by their results, the researchers built a domain parking classifier that automatically detects parked domains based on multiple generic features, including average

[1] `https://www.google.com/adsense/start/`

link lengths, link-to-text ratio and request characteristics. They propose that this classifier could be used as a browser extension to automatically warn users when visiting parked domains.

While both Alrwais et al. and Vissers et al. quantified parked domains, they both acknowledge that their numbers are not comprehensive. One focus of this thesis is to quantify parked domains in depth, trying to get a comprehensive picture of the extent parked domains have across multiple DNS zones and their development over time. We explicitly do not look into the parking monetization chain in depth, as this is already covered in great extent. In comparison to domain parking from the customer or user perspective analyzed by the shown related work, this thesis focuses on parked domains in general and also discusses their implications for further research, which, to the best of our knowledge, have not yet been addressed in related work.

3.2 DNS and Domain Registrars

Halvorson et al. analyzed the top level domains `.biz` [14], `.xxx` [15] as well as many new gTLDs [16] and always found a significant amount of domain parking. While the largest amount of `.xxx` domains in 2013 were registered, but non-resolving, over 50 % of registered resolving domains were classified as parked. According to the researchers, the `.biz` TLD had a similar amount of domain parking in 2011 as the `.com` TLD with about 23 % of resolving domains. When analyzing new gTLDs in 2015 using the CZDS they found that about 32 % of all domains in the zone files are parked. While this thesis also analyzes gTLDs, we do not focus on any specific TLD, but try to analyze parked domains in general, comparing our results to the existing work when applicable.

Concerning domain registrars, Lauinger et al. [17] analyzed practices of so-called *drop catching* where domains are immediately re-registered after their expiration. Out of a manual sample of 50 domains, 23 showed parking pages after their re-registration. They conclude that most of these *drop-catch domains* are used for domain parking and provide "limited value to the Internet community as a whole".

In 2022, Zembruzki et al. [18] analyzed the centralization of the hosting industry and found a single company, *GoDaddy*, that moved 27 M parked domains from their own servers to the Google Cloud, making Google the top autonomous system for web hosting domains. While they did not look at these parked domains in detail in their paper, they show the potential impact that large parking services can have on other industries. This is consistent with results we are able to obtain in our analysis.

Chapter 4

Methodology

This chapter presents our approach for analyzing the results of the DNS scans in order to classify domains as parked. We also introduce specific indicators that we use to find parking services in the utilized datasets.

4.1 Scan Processing

To enable easy querying of the data, we use *ClickHouse*[1], an open source database with a focus on analytical processing using columnar data storage. We first tried using a *PostgreSQL* database, but found that its performance was not sufficient for the queries we wanted to execute. A further advantage of using a columnar storage is a high compressibility of the columns, allowing us to store and query many days of DNS scan results at once.

As the *input list*, *resource record* and *ip-domain* files all use different formats and are stored using a multitude of different archive formats, we created a Python script for importing the different files into the database.

4.2 Classification of Domains

As described in Section 2.2, domain parking services can offer three ways of parking a domain using the DNS:

1. Using the parking services' name servers (→ `NS` records)

---

[1] `https://clickhouse.com`

2. Pointing the domain to the parking services' web servers using an IP address (→ `A` records)
3. Pointing the domain to the parking services' infrastructure using a domain alias (→ `CNAME` record)

We use the following methods to determine a match between a domain and a parking service:

Exact value match
: The record matches a specific value exactly, e.g., an `NS` record with the value `ns1.myservice.tld.`. This is useful for all three methods when looking for a specific known value.

Case insensitive match
: The same as an exact match, except that text casing is ignored, therefore a value of `Ns1.MyService.TLD` would also match when looking for `ns1.myservice.tld`.

IP network match
: The `A` or `AAAA` record matches an IP network in CIDR notation, e.g., assuming `192.168.1.0/24` is the IP network we try to match, all `A` records from `192.168.1.0` to `192.168.1.255` would match.

Regular expression match
: The record matches a specific regular expression, e.g., the regular expression `ns([0-9]{1,2})\.myservice\.com\.` would match values starting with `ns`, then having either one or two digits and ending in `.myservice.com.`.

With the right set of rules we can therefore classify each record as belonging to a parking service or not, although this method has its limitations, the most obvious one being that only services that we know of can be identified. More limitations will be discussed in Section 5.6.

4.3 Identification of Parking Services

Finding domain parking services can be achieved in multiple ways, the most obvious one is probably using search engines and browsing blog posts or forums. This way, we already found 14 services, including most of the large providers. Another way is to use indicators that are common for domain parking in the DNS *resource records* themselves, which we will focus on in this section.

We presume that domain parking companies are profit-oriented and try to host a large amount of domains on a comparably small set of infrastructure. Therefore, our first indicator are IP addresses that host a very large amount of domains. Finding these top IP addresses is straightforward using our database approach. We limit our search to the top 500 IPv4 addresses and the top 100 IPv6 addresses, as we need to manually verify the resulting IP addresses. For this verification, we use 25 domains that we sample randomly from all domains pointing to each IP address and visit them in a web browser. An excerpt of the top 50 IP addresses is shown in Table A.1. We used the same procedure to rank the top `CNAME` aliases and evaluate them for domain parking.

We can extend this assumption to parking services that operate their own name servers. For serving a parking page, these name servers need to resolve the parked domains to an IP address. While name servers of registrars provide ways for customers to configure their own DNS records, thus leading to an arbitrary set of IP addresses resolved by these name servers, we do not expect this behavior when looking at parking name servers. Assuming that a parking service only uses its name servers for resolving parked domains, we again anticipate that parking services try to host as many parked domains on their infrastructure as possible, with only a few IP addresses.

Therefore, the second indicator we use are name servers that are resolving domains to a small amount of IP addresses, while still resolving a large amount of domains overall. We use the *Public Suffix List* [19] to group domains by the shortest private suffix of their corresponding `NS` record. As an example, a domain with `NS` records `ns1.parkingcrew.net.` and `ns2.parkingcrew.net.` would be grouped under `parkingcrew.net`. We limit our search to name server domains that have at least 5000 domains associated with them. Then, we join the list of domains using these name servers with their `A` records to count the number of IPv4 addresses resolved by each name server domain.

Figure 4.1: An exemplary setup of parked domains using name servers of *ParkingCrew*, resolving to *ParkingCrew* IP addresses.

As we are interested in name servers resolving many domains to few IP addresses, we define the *domain-ip ratio* as the number of domains hosted by a name server, divided by the number of resolved IP addresses. We look at the candidates with the highest *domain-ip ratio*. Again, we need to manually verify the potential candidates. By taking a random sample of 25 domains connected to each name server, we manually verify that all the sampled domains exhibit parking behavior. An excerpt of the top 75 name server domains is shown in Table A.2.

As third indicator, we assume that legitimate parking providers are likely to use parking terminology in their DNS infrastructure. For example, *ParkingCrew* uses (among others) the name servers nsX.**parking**crew.net and Namecheap uses the alias **parking**page.namecheap.com. Name server records or DNS aliases containing words like *park*, *parking*, *buy* or *sell* therefore could indicate domain parking. Again, a manual verification of found name servers and aliases is necessary. A similar assumption was successfully used by Kührer et al. in 2014 [20] to create a list of parking name servers.

These methods are all complementary. When finding a conspicuous IP address of a parking provider, it is likely that this provider also operates its own name servers, which should then also stand out in the analysis of the name servers. In general, matches based on `NS` records are preferable since IP addresses are more volatile, especially when the services are using cloud providers, but not all services operate their own name servers and some do not provide a response to `NS` queries for domains delegated to them.

For each found DNS configuration linked to a parking service it is useful to check for corresponding records, e.g., after finding an IP address used for parking, finding the name servers of domains resolving to this specific IP address. This can provide new indicators for recognizing the service. The same applies the other way round, finding all the IP addresses that are resolved by a specific name server.

Chapter 5

Evaluation

In this chapter, we present our results focusing on the found parking services and their presence in the DNS scans.

5.1 Domain Parking Services

As briefly mentioned in Section 4.3, finding the most prominent domain parking services can be achieved by browsing the web for search terms like "best domain parking services" and visiting forums specialized on domain parking like namepros.com [21]. Using this approach, we already found the biggest players on the market.

GoDaddy is the largest registrar worldwide [22] and does not only offer domain registration, but also, among others, hosting services and a domain marketplace for selling domains. For domain parking, GoDaddy offers a free parking page [23] as well as a paid service they call *CashParking* [24]. In 2013, GoDaddy bought the domain marketplace *AfterNic*, the parking service *SmartName* and name generator *NameFind* [25]. In 2020, GoDaddy bought the *Uniregistry* registrar [26] which also offers domain parking.

Newfold Digital is a company group that focuses on buying companies in the hosting and registrar business. It was founded in 2021, merging *Web.com* and the *Endurance International Group* (EIG) [27]. For domain parking Newfold uses *Skenzo* which belongs to *Directi* and was seemingly bought by EIG in 2014 [28]. *Skenzo* parking pages are used in many of the registrar and hosting businesses belonging to Newfold, e.g., *Network Solutions* and *Bluehost*.

United Internet is a German company group that owns the domain monetization service *Sedo* which offers a domain marketplace as well as domain parking pages since 2002 [1].

CentralNic is a holding company that owns many companies in the domain registrar and monetization business. In 2018, CentralNic bought *KeyDrive* [29], the parent group of *Key-Systems* which runs the parking and traffic monetization service *traffic.club*. In 2019, CentralNic bought the Munich located *Team Internet AG* which is the company behind the parking service *ParkingCrew* [30].

One option to categorize the found services is to look at the customers and the services offered to them. Using this approach there are three main categories:

Parking Services that offer domain parking using advertisements or advertised redirects and pay out the profit to their customers. These include among others *Sedo*, *ParkingCrew* and GoDaddy's *CashParking*.

Marketplaces that primarily offer customers ways for buying and selling domains, showing generic or custom for-sale pages when visiting the domains for sale. These include among others *dan.com* and *AfterNic*.

Other services, for example web hosting companies that offer placeholder pages for not yet developed domains or domain registrars that show advertisements on expired domains. The main difference to the previous categories is that these services do not provide a way for customers to earn profit using these domains. Examples for these services include among others *Namecheap*, *Hostinger* and GoDaddy's free parking service (where only GoDaddy earns a revenue from the shown advertisements).

As our analysis focuses on the domains found in DNS scans, we propose another categorization that focuses only on the typical content we find on the parked domain:

Parking
Services that monetize domains using PPC advertisements or advertised redirects (PPR) monetization.

Marketplace
Services that do **not** show advertisements on their domains, but "for-sale pages" that indicate that the domain is for sale.

Placeholder
Services that show any kind of placeholder page with no real content.

Mixed

Services that cannot be clearly classified, as they offer multiple of the previously mentioned service categories on the same infrastructure.

In further analyses we will use this classification, as, e.g., a user that is visiting a domain cannot distinguish if it is parked by a customer of a domain parking service or an expired domain parked by a domain registrar. When talking about "parked domains", we not only refer to domains that show advertisements, but also to domains that are in the other categories.

5.2 Domain Parking in our Dataset

For our analyses, full resource record results are only available starting 2021-01-05, older results are only available in the *ip-domain* format mentioned in Section 2.3.2. We will therefore limit our analysis over time to the period from the beginning of January to the end of December 2021. Furthermore, the merged `NS` records are only available starting 2021-05-23, which we also have to consider when looking at results over time.

In the following, we will always use the shortest private suffix according to the Public Suffix List [19] when counting domains, i.e., the domains `www.example.co.uk.` and `mail.example.co.uk.` are only counted once as their public suffix is `co.uk.` and their shortest private suffix therefore being `example.co.uk`. We do this because domains are usually parked or not parked, without specific subdomains used for legitimate purposes. Furthermore, our input sources are different, while zone files do not contain any subdomains, the domains on certificate transparency logs are mostly subdomains. Excluding subdomains from the domain count leads to more comparable results.

The Tables 5.1, 5.2 and 5.3 give an overview of all the services classified as *parking*, *marketplace* and other (*placeholder* / *mixed*) respectively, as well as the number of domains identified at the specific service. In total, when looking at the average numbers across 15 scans from December 2021, there are 388 M domains in our input sources, 267 M are actually resolving to an `A`, `AAAA` or `CNAME` record and 62 M are classified as parked. When counting resolved domains, we exclude candidates that solely resolve to an `NS` record as these cannot be used for actual services, e.g., web servers. From now on, we will use this definition when counting resolved domains.

When looking at the *parking* category (Table 5.1), i.e., services that show advertisements when visiting their domains as described in Section 5.1, the largest service by far is *GoDaddy's* free parking service [23] with more than 30 M parked domains. These also include domains that are available in GoDaddy's domain auctioning system with

TABLE 5.1: The found parking services using primarily advertisements, along with their presence in our dataset. The NS, IP and CNAME columns show how we can detect these services. Tier 1 indicates if this parking service provides its own parking pages. Services that use the parking pages of another service or only redirect using PPR are not classified as tier 1. □ indicates that we found the service in our manual research, while ◊ means that we first found the service in our DNS data.

	Name	Tier 1	NS	IP	CNAME	Domains[1]
□	GoDaddy (Free Parking)	✓	×	✓	×	30.3 M
□	Skenzo	✓	✓	✓	✓	2.9 M
□	GoDaddy (CashParking)	✓	✓	✓	×	2.2 M
□	ParkingCrew	✓	✓	✓	✓	1.6 M
□	Bodis	✓	✓	✓	✓	1.1 M
◊	survey-smiles.com[2]	×	✓	×	×	968.4 k
□	above.com	×	✓	✓	×	851.5 k
◊	Namecheap	×	✓	×	✓	791.5 k
◊	Bluehost	×	✓	×	×	613.9 k
◊	Dynadot	×	×	✓	×	391.9 k
◊	TrafficMotor	×	✓	×	×	391.7 k
□	ParkLogic	×	✓	×	✓	352.4 k
◊	snparking.ru	×	✓	×	×	264.0 k
□	DomainSponsor	×	✓	✓	✓	190.3 k
◊	123 Reg	×	×	✓	×	138.0 k
□	Voodoo	✓	✓	✓	×	68.0 k
◊	Expiereddnsmanager[2]	×	✓	×	×	37.6 k
◊	traffic.club	×	✓	×	✓	26.7 k
◊	Domainpower	×	✓	×	×	22.3 k
◊	regtons	×	✓	✓	×	14.5 k
◊	Fabulous / Directnic	×	✓	×	×	8.7 k
◊	Internetvikings	×	✓	×	×	5.4 k
◊	Domaincntrol[2]	×	✓	×	×	5.0 k
□	The Parking Place	×	✓	×	×	4.4 k
◊	Tucows	×	✓	×	×	90.0 [4]
	Total					43.1 M[3]

[1] The number of domains was calculated by averaging the results across 15 scans from 2021-12-03 to 2021-12-31.
[2] We could not identify the company behind this service.
[3] Some domains are parked with multiple services, therefore this number can be lower than the sum of all services above as we count unique domains.
[4] *Tucows* seems to block or rate limit our scan starting 2021-10-11. Before, the amount of parked domains was around 200 k.

the only difference being one sentence that says that the domain "is available on GoDaddy Auctions". Next, we have *Skenzo* with nearly three million domains, followed by GoDaddy's paid *CashParking* service, *ParkingCrew* and *Bodis*, all with more than a million parked domains. These largest services have in common that they not only offer their own parking pages to end users, but also provide parking pages to other services, such as registrars or hosting companies. We call these services "tier 1". The only other tier 1 provider is *Voodoo* with 68 k domains. In total, we identified only six services (not counting GoDaddy twice) that actually implement their own parking pages and are therefore classified as tier 1. This also includes *Sedo*, which is categorized as mixed and therefore not included in Table 5.1. All the other services sit on top of these tier 1 providers or use PPR monetization.

While we list *survey-smiles.com* as a parking service, we were not able to identify the company behind this service. We identified 25 different name server domains used by this service (for example `nsX.rentondc.com`), that all have in common that the name servers themselves also host a web server that redirects to the domain `survey-smiles.com` when doing an HTTP request (see Listing 5.1).

Listing 5.1: The name server ns1.rentondc.com redirecting to survey-smiles.com using HTTP.

```
$ curl -v http://ns1.rentondc.com
*   Trying 172.98.192.34:80...
* Connected to ns1.rentondc.com (172.98.192.34) port 80 (#0)
> GET / HTTP/1.1
> Host: ns1.rentondc.com
> User-Agent: curl/7.77.0
> Accept: */*
>
* Mark bundle as not supporting multiuse
< HTTP/1.1 302 Found
< cache-control: max-age=0, private, must-revalidate
< connection: close
< content-length: 11
< date: Thu, 20 Jan 2022 17:03:54 GMT
< location: http://survey-smiles.com
< server: nginx
< set-cookie: sid=f2db814e-7a12-11ec-a49e-0589d34f483b; path=/; domain=.rentondc.
    com; expires=Tue, 07 Feb 2090 20:18:01 GMT; max-age=2147483647; HttpOnly
<
* Closing connection 0
Redirecting
```

We manually checked the top domain disputes listed on *dndisputes.com*[1] and in many cases name servers associated with survey-smiles.com were used. The domains also always use WHOIS masking services like *Fundacion Privacy Services LTD* or *Whois Privacy Corp*. In a specific case it is stated that "[i]t is therefore a matter of disappoint-

[1] https://www.dndisputes.com/case/respondent/

ment and concern that the Registar [sic] refused to disclose the underlying registrant details in this case, either in response to the Center's verification request or in response to the Panel's express request that it do so. The only explanation offered for that refusal is a claim that it 'cannot disclose the details of the registrant behind privacy', but no reason is offered as to why that this is the case. This statement also appears to be untrue or at best misleading." [31]. As this service also uses many hosting providers for its infrastructure, in contrast to other services usually using one hosting or cloud company only, we conclude that this service is involved in dubious activities.

Similarly, the companies behind the services named *Expiereddnsmanager* and *Domaincntrol* could not be identified, the homepage of the latter one simply reading "We are a domain parking company..."[1] and using a typo of GoDaddy's `nsXX.domaincontrol.com` as name servers (`nsX.domaincntrol.com`, note the missing `o`).

We have observed parking services sending traffic to other parking services, for example *ParkingCrew* and *Sedo* using *Skenzo*. This practice is probably used to monetize traffic that is not accepted by the used advertisement network. In contrast, *above.com* is a parking manager that does not provide its own parking pages, but parks its domains solely using other parking services, while trying to optimize the generated revenue. When visiting a domain parked through *above.com*, the service uses HTTP redirects to a subdomain with the form `wwXY`, e.g., `ww16.example.tld` where the subdomain is then parked at *ParkingCrew*, *Sedo* or *Bodis*, among others. Similar behavior can be observed at *Expiereddnsmanager*, *regtons*, *survey-smiles.com* and *Domaincntrol*. Another approach to park a domain with multiple parking services that we observed is to set `NS` records to multiple parking providers, thus using DNS to load-balance between them. We presume that domains showing such behavior are not systematic for a specific service, but rather domain owners that want to park their domains with multiple services at once.

Table 5.2 shows the found marketplaces. In our initial research, we only found *AfterNic*, but there we focused on advertised parking. The other services were discovered because they were either in the list of top IP addresses (see Table A.1) or the list of name servers with a high *domain-ip ratio* (see Table A.2). Many of these services can be found when specifically searching for domain marketplaces.

It is noteworthy that the largest service by the number of domains, *HugeDomains.com*, only sells domains that are owned by the service itself. There is no option for customers

[1] `https://domaincntrol.com`

Table 5.2: The found domain marketplaces along with their presence in our dataset. The NS, IP and CNAME columns show how we can detect these services. □ indicates that we found the service in our manual research, while ◊ means that we first found the service in our DNS data.

	Name	**NS**	**IP**	**CNAME**	**Domains**[1]
◊	HugeDomains.com	×	×	✓	4.5 M
◊	dan.com	✓	✓	×	2.0 M
□	AfterNic	✓	×	×	976.7 k
◊	domain.com	×	✓	×	769.1 k
◊	Aliyun	×	×	✓	253.5 k
◊	DomainMarket	✓	×	×	230.3 k
◊	Squadhelp	✓	×	×	211.5 k
◊	DomainProfi	✓	×	×	211.3 k
◊	epik	✓	×	×	204.3 k
◊	domainname.de	✓	×	×	198.2 k
◊	REG.RU	×	✓	×	184.0 k
◊	efty	✓	×	×	179.3 k
◊	Sonexo DNFS24	✓	×	×	140.2 k
◊	Domainparking.ru	✓	✓	×	135.8 k
◊	BrandBucket	✓	×	×	131.1 k
◊	domainmarkt.de	×	✓	×	87.2 k
◊	PerfectDomain	✓	✓	×	75.9 k
◊	22.cn	×	✓	×	58.4 k
◊	WangGuai.com	×	✓	×	40.8 k
◊	domainrecover.com[2]	✓	×	×	40.6 k
◊	LinkUWant	×	✓	×	37.4 k
◊	Alter	✓	✓	×	32.9 k
◊	DomainOrder	✓	×	×	30.5 k
◊	Domainist	✓	×	×	20.3 k
◊	DNSPod	×	×	✓	17.5 k
◊	Domain Brokers Sweden	✓	×	×	17.5 k
◊	TRUSTEDNAMES	✓	×	×	16.4 k
◊	nameprovider.net	✓	×	×	14.6 k
◊	domain.io	✓	×	×	12.1 k
◊	Flippa	×	×	✓	6.3 k
	Total				10.6 M[2]

[1] The number of domains was calculated by averaging the results across 15 scans from 2021-12-03 to 2021-12-31.
[2] We could not identify the company behind this service.

Table 5.3: Other services with a behavior similar to parking, along with their presence in our dataset. The NS, IP and CNAME columns show how we can detect these services. □ indicates that we found the service in our manual research, while ◊ means that we first found the service in our DNS data.

	Name	NS	IP	CNAME	Domains[1]
	Placeholder				
◊	Alibaba	×	✓	×	394.9 k
◊	123 Reg	×	✓	×	272.0 k
◊	Hostnet.nl	×	✓	×	245.2 k
◊	transip	×	✓	×	208.2 k
◊	Hostinger	×	✓	×	187.0 k
◊	Hostpoint	×	✓	×	174.4 k
◊	domainname.shop	×	✓	×	173.9 k
◊	fasthosts	×	✓	×	155.7 k
◊	Namebright	×	×	✓	145.7 k
◊	Namecheap	✓	×	×	122.6 k
◊	one.com	×	✓	×	122.1 k
◊	HostGator	×	✓	×	117.8 k
◊	Domainbox	✓	×	×	109.5 k
◊	west.cn	×	✓	×	99.0 k
◊	domaindiscount24	✓	×	×	90.6 k
◊	forpsi internet CZ	×	✓	×	78.1 k
◊	Enom	×	✓	×	77.6 k
◊	101domain	×	✓	×	60.3 k
◊	gabia	×	✓	×	54.7 k
◊	Markmonitor	×	✓	×	53.7 k
◊	wedos	×	✓	×	38.7 k
◊	aruba.it	×	✓	×	31.7 k
◊	Turkticaret.Net	×	✓	×	26.6 k
	Total				3.0 M
	Mixed				
□	Sedo[2]	✓	✓	✓	3.1 M
□	sav.com[3]	✓	×	✓	992.0 k
□	Uniegistry[2]	✓	✓	×	936.5 k
◊	dne.com[2]	✓	×	×	20.3 k
	Total				5.1 M

[1] The number of domains was calculated by averaging the results across 15 scans from 2021-12-03 to 2021-12-31.
[2] Primarily advertisements, but also for-sale pages.
[3] Primarily for-sale, but also placeholder pages.

to sell their own domains through *HugeDomains.com*. Among others, *domain.com*, *DomainMarket* and *DomainProfi* also only sell domains seemingly owned by themselves, while the biggest services that offer domain sales for customers are *dan.com* and *AfterNic*.

When looking at the services in the placeholder category in Table 5.3, there are mostly hosting companies serving default informational pages that the domain is registered, but there is no content (yet). We list *Namecheap* and *123 Reg* in the parking and placeholder category, as we can differentiate the domains based on distinct DNS configurations. As a concrete example, *Namecheap* uses the IP address `99.83.154.118` and alias `parkingpage.namecheap.com` for domains in the parking category, while the name servers `failed-whois-verification.namecheap.com` and `verify-contact-details.namecheap.com` are used for domains in the placeholder category, showing a page about the mandatory WHOIS contact verification.

The biggest service categorized as *mixed* is *Sedo*, a German domain parking service that also offers its own domain marketplace and domain brokerage service. We classify it as mixed, as the parking and sales pages are both hosted on the same IP address space, therefore we cannot distinguish them using DNS data alone. The same applies to *Uniegistry*. While *sav.com* primarily serves sales pages, it also runs placeholder pages on the same infrastructure.

5.2.1 Results by Input List

As explained briefly in Section 2.3.1, we use multiple different data sources to gather domains. Therefore, we are interested in the distribution of parked domains among these input lists.

Figure 5.1 shows the amount of resolved and parked domains on each list in relation to the total domains on the input list. We can see that the zone files for *.com, .net, .org* and the other gTLDs have a high ratio of parked domains. About 25 % of all domains on the *.com, .net, .org* list are parked (30 % of the resolving domains). Looking at the other gTLDs acquired from the CZDS, 17 % of all (26 % of the resolving) domains on the input lists are classified as parked. The low resolve ratio of the zone files, especially the new gTLDs probably comes from defensive registrations, as found by Halvorson et al. [16] in 2015. We manually checked a random sample of 1000 domains that were in the input list and missing in the resolved data and about 89 % of the missing domains failed resolving with a `SERVFAIL` error or returned no data when querying for `A` records. Therefore, we suppose that this, among possible rate-limits, timeouts or packet losses while scanning probably explain this low resolve ratio.

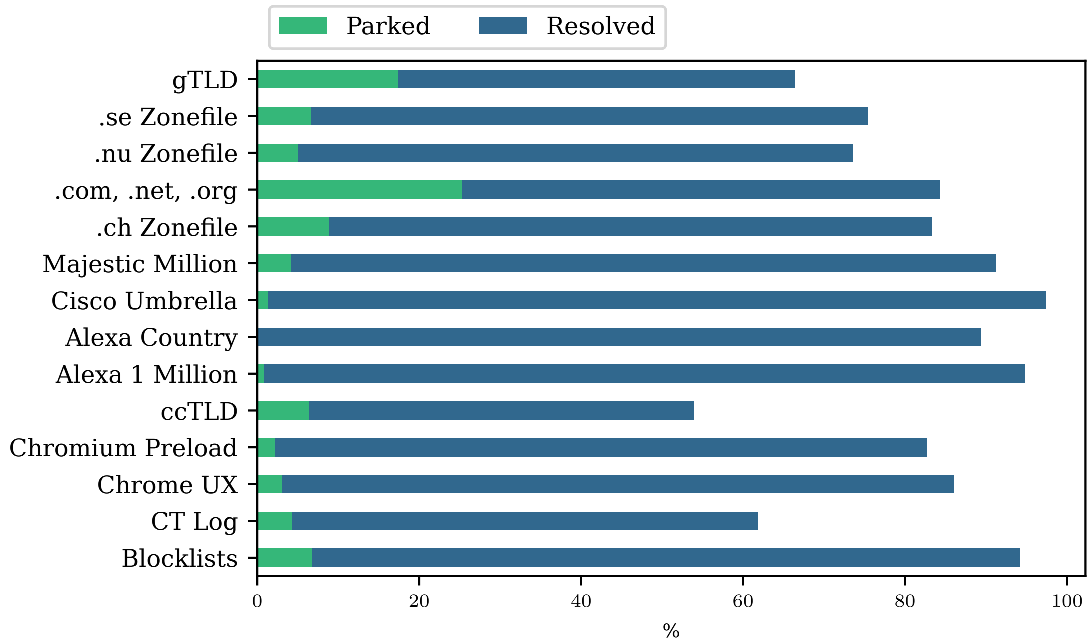

Figure 5.1: Domain Parking by Input List. Data averaged across 15 scans from 2021-12-03 to 2021-12-31.

The parking ratio in the top lists are low, which is what we expected, but with 4.5 % of the resolving domains on the *Majestic Million*, we still find a rather high amount of parked domains on this list, comparing to only 0.9 % on the *Alexa Top 1 Million* and 1.3 % on the *Cisco Umbrella* list. It is important to note that while the *Cisco Umbrella* list has 1 M entries, by our counting method it averages only 215 k unique private suffixes. For example, on 2022-01-28, there are 27 k domains ending in `.googlevideo.com`, 12 k `.fbcdn.net` and 11 k `.webex.com` subdomains on the list. Overall, 65 % of private suffixes are on the list more than once. This behavior is explained by the DNS based collection method used by the *Cisco Umbrella* list. Analyzing a sample of 1000 parked domains on the *Majestic Million* list in December 2021, we find that the majority (67 %) of parked domains are already parked since the beginning of our records in January 2021 and also found on the list continuously. In contrast, the same method applied to the *Alexa Top 1 Million* yields a very different result with 33 % of parked domains being seen only once on the list. This is consistent with the stability results of these top lists by Scheitle et al. [10]. Compared to the zone files, the top lists consistently have a high resolve ratio, around 90 %, which indicates that most of the listed domains are actively in use.

Table 5.4: Parked domains and their ranking on top lists. Data from 2022-01-28.

	# Domains			# Parked in Top		
List	**Total**	**UPS[1]**	**Parked**	**10k**	**100k**	**250k**
Alexa Top 1 Million	590 025	582 343	5333	1	378	1432
Majestic Million	1 000 000	997 740	40 359	15	891	6133
Cisco Umbrella	1 000 000	267 910	3735	1	147	439

[1] Unique Private Suffixes.

To see how popular parked domains are on top lists, we analyze the ranks of parked domains on these lists (Table 5.4). We use data from 2022-01-28, but verify that the amount of parked domains is in the usual range for the lists. On the *Alexa Top 1 Million*, out of a total of 590 025 entries, we find 5333 parked domains. Only one parked domain is in the top 10 k (`iyfbodn.com`, a domain used by *Skenzo*, rank 9898). The median rank of parked domains is 358 280 and the average 343 911.0. Looking at the *Majestic Million*, we find 40 359 parked domains (1 M total), with a median and average rank of 532 270 and 540 335.6 respectively. The *Cisco Umbrella* list shows the least amount of parking, with only 3735 (1 M total) parked domains and also the lowest median and average rank, 721 389 and 651 676.4 respectively. When counting unique private suffixes, the amount of parked domains is similar, although a little higher than on the *Alexa Top 1 Million*, around 1 %. As expected, while parked domains do occur in these lists, we can conclude that it does not play a major role in the upper ranks of top lists.

The static list of ccTLDs shows the lowest resolve ratio of all lists. This time we assume that this is not only because of the previously mentioned reasons of registered, non-resolving domains, but also because the list is from July 2020 and therefore probably contains many domains that are not registered anymore. Concerning parking, about 11.8 % of the resolving domains on the *ccTLD* list are classified as being parked, which is considerably lower than on the gTLD zones. We presume that ccTLDs are not as attractive for domain parking, especially when they are not marketed for specific purposes, e.g., `.co` being marketed as referring to the keyword *corporation* or *company*, while actually being the TLD for Colombia. Some ccTLDs have special contact verification requirements or are not available for everyone, e.g., `.eu` requiring registrants to be citizens of the European Union or members of the European Economic Area. Similar low values of parking can be seen on the ccTLD zone files of `.se`, `.nu` and `.ch`.

Domains on the certificate transparency logs have the second-lowest resolve ratio, but as this list is also not regularly updated, we presume that many domains are not in use anymore. The parking ratio is low at 4 % of all and 7 % of resolving domains.

This suggests that most domain parking services do not use TLS, which we will further analyze in Section 5.3.

When looking at parked domains on block lists, we find that about 7 % of domains on these block lists are parked. The high resolve ratio of 94 % indicates that these lists mostly contain active domains. Related work showed that malicious domains of traffic distribution systems often end up parked after being taken down [32]. In contrast, we do not see a noticeably high amount of parked domains on the used block lists. We acknowledge that this can greatly differ based on the lists used, and it could also mean that parked domains are removed from the block lists after being initially taken down. For this thesis, we decided not to look into the block lists in more detail.

5.2.2 eTLD Analysis

Related work already analyzed the presence of domain parking in the `biz.` [14] and `xxx.` [15] TLD as well as the set of new gTLDs [16] in 2015. Seven years later, we compare our results and provide an updated view on domain parking across all effective top-level domains (eTLDs) in our dataset according to the Public Suffix List [19]. While Halvorson et al. [16] used more than just DNS data to classify domains as parked, we focus only on domains that we can clearly classify solely based on their DNS records, therefore our results are lower bounds and not directly comparable.

In our dataset of 23.6 M resolving new gTLD names, 26 % are parked, as shown in Figure 5.1. Halvorson et al. found that in 2015, out of 3 M domains on the new public gTLDs zone files about 32 % were parked [16]. We presume that this difference is caused by the large growth of these gTLDs since 2015 as well as the differences in classification methods.

In Table 5.5 we see the top eTLDs sorted by their number of unique shortest private suffixes, i.e., `www.example.co.uk.` and `mail.example.co.uk.` would only be counted as one domain under the suffix `.co.uk`. The most important eTLD by far is `.com` with almost 139 M domains and a parking ratio of nearly 31 %. This is already the majority of the 62 M total parked domains. The ratio of parked domains varies strongly between `.cf`, the ccTLD of the Central African Republic with only 2.5 % of resolving domains classified as parked and 34 % in the `.info` zone. The high number of `.cf` domains is most likely explained by the free availability of `.cf` domains through the *Freenom* registry.

We use the Shannon entropy to get an impression of the distribution of services among each eTLD. This entropy is calculated with the following formula:

Table 5.5: The top eTLDs by their total number of unique resolved domains in their shortest private form according to the Public Suffix List. Note that we only count a domain as resolved, when it resolves to an `A`, `AAAA` or `CNAME` record. Data from 2021-12-31.

	# Domains			
eTLD	**Total**	**Parked**	**Parking %**	**Service Entropy**
com	138 724 243	42 687 909	30.77	3.25
de	12 221 972	1 488 215	12.18	2.94
net	10 825 963	3 051 885	28.19	2.75
org	9 071 889	2 642 694	29.13	2.42
nl	3 636 374	430 983	11.85	2.88
ru	3 412 825	492 150	14.42	1.80
cf	3 405 761	84 468	2.48	0.15
info	3 285 206	1 121 458	34.14	1.96
cn	2 691 823	167 372	6.22	2.90
xyz	2 656 560	572 643	21.56	3.79

$$H(X) = -\sum_{i=1}^{n} P(x_i) log_2(P(x_i))$$

In our case, n is the number of parked domains and $P(x_i)$ is the probability that a parked domain is parked at a specific service. As an example, if there are two services among ten parked domains, four parked at service A and six at service B, this would lead to an entropy of

$$\begin{aligned} H(X) &= -(P(x_1)log_2(P(x_1)) + P(X_2)log_2(P(X_2))) \\ &= -(0.4log_2(0.4) + 0.6log_2(0.6)) \\ &\approx 0.97 \end{aligned}$$

The maximum value the entropy can have is $log_2(n)$ when the possibilities are uniformly distributed. A value near zero means that the set is largely dominated by one value and generally, a value of n is equivalent to a uniformly distributed set of 2^n different values, indicating that the dataset is largely dominated by about 2^n values. As indicated by the low entropy of 0.15, almost all parked `.cf` domains are parked with one specific provider, *DomainSponsor*.

When looking at the top eTLDs by the ratio of parked to non-parked domains (Table 5.6), we almost exclusively see new gTLDs in the top 100 except for a few ccTLDs that are marketed for other purposes, like `.co` or `.cm`, the ccTLD of Cameroon on rank 6 being known as a typo of `.com`. As indicated by the low entropy of 0.23, a closer

Table 5.6: The top eTLDs by the ratio of parked to total domains in their shortest private form according to the Public Suffix List. Note that we only count a domain as resolved, when it resolves to an `A`, `AAAA` or `CNAME` record. Data from 2021-12-31.

	# Domains			
eTLD	**Total**	**Parked**	**Parking %**	**Service Entropy**
realty	13 682	12 546	91.70	0.23
xn–tckwe[1]	2189	1392	63.59	1.09
bond	6253	3868	61.86	1.67
app	594 971	343 064	57.66	0.98
firm.in	1250	719	57.52	1.46
cm	10 747	6080	56.57	2.05
miami	14 192	7910	55.74	1.70
hair	1841	1006	54.64	2.85
boston	3047	1650	54.15	1.86
makeup	1114	602	54.04	3.16
attorney	5052	2706	53.56	1.94

[1] This is `.com` written in Japanese.

look at the outlier on rank 1, `.realty`, shows that nearly all the parked domains are showing a sales page by *epik*, a domain registrar that, according to *nTLDStats* [33], has registered about 13 000 `.realty` domains at the end of January 2022, which means that almost all of these domains are parked. We therefore presume that this is either the registrar itself, or one single customer of *epik* that is responsible for most of the parked domains in the `realty.` zone. The `.app` TLD has the highest amount of parked domains across the top parking eTLDs with more than 57 % of its 595 k domains associated with parking. This is consistent with the high amount of parking already found on some new gTLDs by Halvorson et al. in 2015, although the `.app` TLD is only open for registration since 2018. In this ranking, the `.com` suffix with its 30.8 % parking ratio is far behind on the 202nd place.

In Table 5.7, we analyzed the number of eTLDs among parking services. We find that the big providers also have the highest amount of eTLDs, with *ParkingCrew* and *Sedo* having domains from nearly 900 different eTLDs on their platform and an entropy of around 3.5. The service *101domain* has a noticeably high entropy of 6.63, indicating that it hosts a lot of different domains, compared to most parking services where `.com` and only a few more TLDs are dominant. The entropy of 6.63 is equivalent to a uniform distribution of $2^{6.63} \approx 99$ eTLDs. On the other side we mainly find marketplaces that focus on selling specific domains, e.g., *domainmarkt.de* which only sells `.de` domains, therefore having an entropy of zero.

Table 5.7: The top ten and bottom ten services by the number of distinct eTLDs found, including the entropy of eTLDs. Data from 2021-12-31.

Service	# eTLDs	eTLD Entropy
ParkingCrew	899	3.384
Sedo	883	3.663
dan.com	802	3.333
Bodis	762	2.479
101domain	756	6.630
GoDaddy (Free Parking)	720	2.275
Skenzo	707	1.900
above.com	670	3.452
Uniegistry	579	1.862
GoDaddy (CashParking)	548	1.361
DomainOrder	8	1.725
DomainMarket	6	0.029
LinkUWant	6	0.009
nameprovider.net	6	1.777
HugeDomains.com	6	0.049
Domainist	4	0.005
WangGuai.com	4	0.156
BrandBucket	4	0.163
domainrecover.com	2	0.035
domainmarkt.de	1	0.000

5.2.3 Development over Time

In the following, we will analyze the development of domain parking in our dataset over the course of the year 2021. This dataset consists of 199 scans starting with the availability of detailed resource records on 2021-01-05 until 2021-12-31. We exclude the certificate transparency (CT) scan results, as our storage is limited and an analysis of the data from December 2021 showed that including the CT scans only increases the number of detected parked domains by about 365 k in this time range. The dataset used in this analysis consists of 132 B DNS records, 14 B classified as parked. Because of errors before or during the scan, some invalid or incomplete scan results at the end of June and October as well as the beginning of December are excluded.

Figure 5.2 shows the number of unique domains in the scan input, the resolved records as well as the number of domains classified as parked over all conducted scans in 2021. The drop of about 5 M detected parked domains in March 2021 comes primarily from one provider, namely *HugeDomains.com*, with more than 4 M domains that are almost all missing in our resolved dataset. We presume that this was either an error with

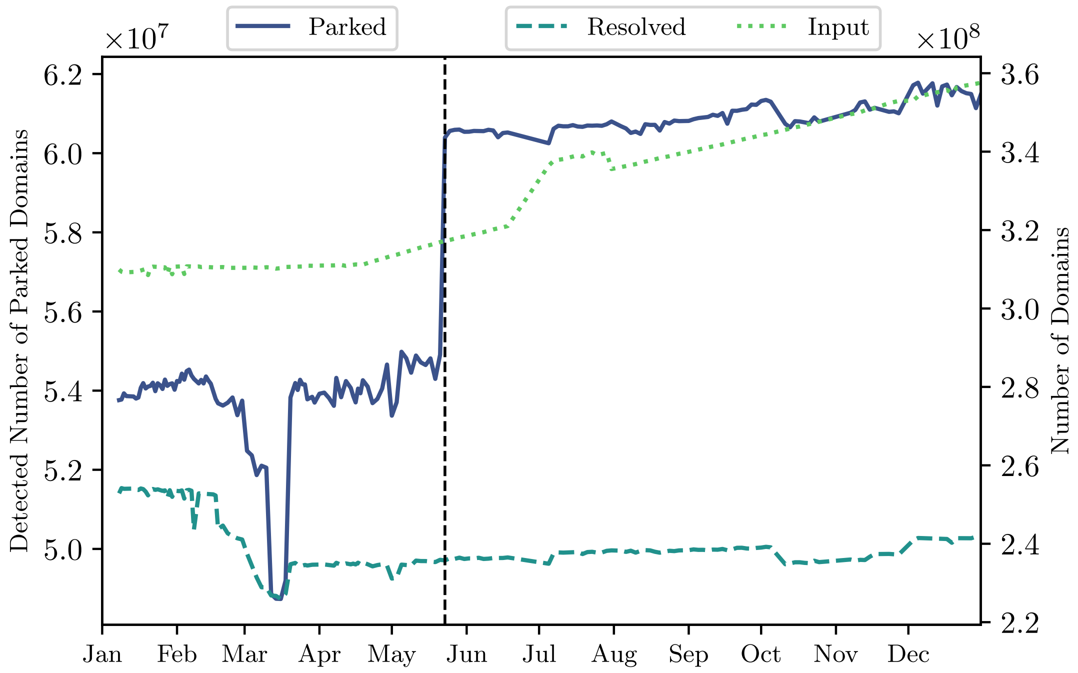

Figure 5.2: The number of domains (unique private suffixes) over the year 2021. The left axis shows the parked domains, the right axis the domains in the input set and the resolved domains. The vertical line shows the date when the first *merged* `NS` scan is available, 2021-05-23.

specific name servers or that the scanner was being blocked for this period of time. Similarly, we can see a drop in overall resolved domains, when comparing January to April 2021 and thereafter, that is not reflected in the input or the number of parked domains. This is because 15.6 M domains from the `.tk` zone are missing, probably because of rate-limits, expired domains, or a block of our scan. We cannot determine the exact reason, as there is no zone file of `.tk` available.

The jump in detected parked domains in May 2021 comes from the inclusion of the merged `NS` scan that is only available since 2021-05-23 and is needed for detecting parked domains based on their name servers on all input lists except the top lists that have a separate `NS` scan. This shows that the `NS` results are very important to reliably detect parked domains, as it is not feasible to always use IP addresses, especially when a parking service uses cloud providers where IP addresses and hostnames can frequently change. As a concrete example, we have seen the domain marketplace *HugeDomains.com* changing its `CNAME` aliases on the *AWS*[1] cloud multiple times in late 2021 before switching

[1] `https://aws.amazon.com`

to using aliases of the format `traf-XY.hugedomains.com`. Afterwards, the number of parked domains is relatively steady, hovering around 61 M after the addition of the `NS` scans.

To analyze the stability of parked domains, we examine the fluctuation of domains across our scans in Figure 5.3. There we can see the amount of parked domains detected at each scan as well as the cumulative number of unique parked domains seen in the top graph. The bottom graph shows for each scan the number of unique domains missing, i.e., not classified as parked anymore compared to the last scan, including the new-added and re-added domains. New-added domains were never classified as parked before, therefore the second scan only has new-added domains.

The already explained drop of 5 M domains on 2021-03-12 can clearly be seen, along with the re-addition of these domains on 2021-03-20. The large spike of added domains on 2021-05-23 through the addition of the *merged* `NS` scan is also clearly visible. Not only do we see 4.5 M new-added domains, but also 1.4 M domains that are re-added, which indicates that these re-added domains were possibly parked at a service that we only recognize using `NS` records during the time they were missing.

At the beginning of July, we see an increased number of new-added and missing domains that can be explained by the missing excluded scans at the end of June. Similar patterns are visible after missing scans in October, November and December. The median of added domains is 376 k, with the median of missing domains at 363 k. This means that there is a relatively stable set of domains that are consistently parked throughout our measurements. Still, the fluctuation of parked domains between each scan, usually every second day, is not to be neglected at more than 350 k domains, indicating that the domain parking ecosystem is very active.

These fluctuations are probably caused in large parts by newly expired domains being parked, which is consistent with related work by Lauinger et al. [17] that showed that special *drop-catch* services, specialized to quickly try to re-register newly expired domains, are responsible for about 80 % of domain creation attempts with lots of re-registered domains ending up being parked. On 2021-12-31, the cumulative amount of parked domains seen since 2021-01-05 is more than 106 M.

5.3 TLS Support Among Parking Services

We briefly analyzed the usage of TLS at the identified parking services, showing that while there are some providers that have a working TLS setup, including valid certificates

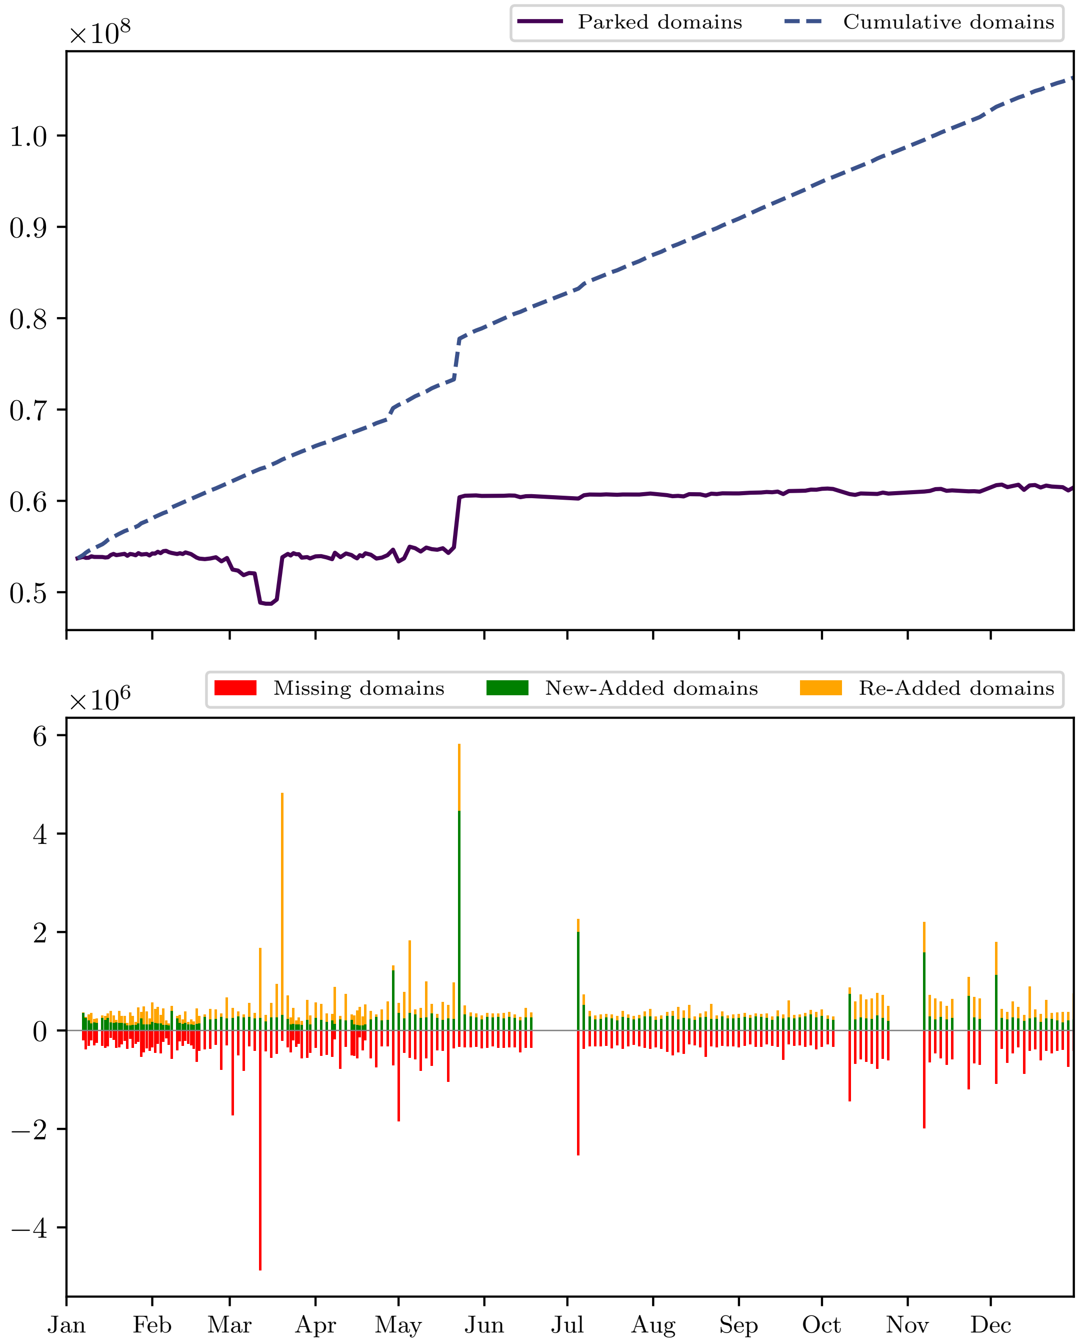

Figure 5.3: The number of parked domains as well as the cumulative number of all unique parked domains seen in 2021. Below, the number of added, removed and equal parked domains from one scan to the next. Note that the y-Axis does not start at zero.

Table 5.8: The top ten parking services by the ratio of total domains with a valid TLS certificate. **Total** shows the total number of parked domains detected at this service. **In Scan** is the number of domains that are in the TLS scan. **Success** shows the number of domains that the scan was able to successfully connect to. **Valid** is the number of domains that presented a certificate that is considered valid. Data from 2022-02-03.

	# Domains				%		
Service	**Total**	**In Scan**	**Success**	**Valid**	**Suc.[1]**	**Valid[2]**	**Total[3]**
epik	185 255	184 664	184 225	183 638	99.8	99.7	99.1
BrandBucket	134 792	134 729	134 291	132 366	99.7	98.6	98.2
TrafficMotor	396 231	394 448	394 360	385 291	100.0	97.7	97.2
Squadhelp	201 878	201 716	201 572	195 891	99.9	97.2	97.0
Alter	42 330	41 736	41 707	40 730	99.9	97.7	96.2
dan.com	2 161 796	2 150 940	2 146 854	2 069 087	99.8	96.4	95.7
domain.io	20 249	19 743	19 740	19 158	100.0	97.1	94.6
regtons	30 356	30 221	30 197	28 397	99.9	94.0	93.5
Sedo	3 166 135	3 138 553	2 876 576	2 871 788	91.7	99.8	90.7
above.com	923 337	622 329	578 640	573 889	93.0	99.2	62.2
Total	61 868 096	21 167 295	13 930 624	7 604 765	65.8	54.6	12.3

[1] Domains with successful connection / total domains in the TLS scan.
[2] Domains with valid certificate / domains with successful connections.
[3] Domains with valid certificate / total parked domains.

for parked domains, the majority of domains either do not support TLS at all (including notably the largest service, *GoDaddy*) or serve self-signed or expired certificates.

In the chair's TLS scan from 2022-02-03, we find 21.2 M parked domains. This does notably not include the 27 M domains that *GoDaddy* hosts on a single IP address, as this address is manually excluded from the scan for performance reasons. A manual check showed that *GoDaddy* does not support TLS at all on their parking infrastructure. While the well-known port 443 is open, the connection is immediately closed with a *FIN* packet after receiving the initial *SYN/ACK* packet.

Table 5.8 shows the top ten parking services by the ratio of valid TLS certificates served for their parked domains. This includes mainly services in the marketplace category, but also *Sedo* and *above.com* that both offer advertised parking. Overall, out of 62 M parked domains, only 7.6 M allow an encrypted connection with a valid certificate (12 %).

When looking at the issuing organizations for the used TLS certificates (Table 5.9), we find that the dominant organization is *Let's Encrypt*, followed by *DigiCert*. The issuer *testexample* is used by *Skenzo* for their self-signed certificates.

Table 5.9: The top ten certificate issuer organizations used by parked domains. Data from 2022-02-03.

Issuer Organization (**Common Name**)	**Domains**
Let's Encrypt (R3)	4 870 530
DigiCert Inc (Encryption Everywhere DV TLS CA - G1)	2 879 132
testexample (testexp)	2 293 409
DigiCert Inc (Thawte TLS RSA CA G1)	1 502 387
Sectigo Limited (Sectigo RSA Domain Validation Secure Server CA)	1 141 466
Let's Encrypt (E1)	511 412
GoDaddy.com, Inc. (Go Daddy Secure Certificate Authority - G2)	295 312
Sectigo Limited (Sectigo RSA Organization Validation Secure Server CA)	175 953
ZeroSSL (ZeroSSL ECC Domain Secure Site CA)	115 453
GlobalSign nv-sa (GlobalSign RSA OV SSL CA 2018)	58 041

5.4 Comparison to Rapid7 Dataset

Rapid7 is a US based cybersecurity company that offers Internet-wide scan data from their *Project Sonar*[1], including among others DNS, HTTP and TLS. We use the data provided in their forward DNS (FDNS) dataset [34] to get an impression of how comparable our results are.

The Rapid7 data is provided in *JSON* format which enables us to conveniently and quickly load the data into our database. While Rapid7 offers data for multiple DNS record types, we are only interested in the `NS`, `CNAME`, `A` and `AAAA` data, as this is relevant for our parked domain detection. For comparison, we use the Rapid7 data from 2022-01-28 and our scan results from the same day.

The Rapid7 dataset contains 3.4 B records, while our dataset consists of 2.2 B records. Counting unique domains in their shortest private suffix form, the Rapid7 dataset consists of 332 M resolved domains, compared to our 267 M resolved domains. Again, we exclude domains that solely resolve to `NS` records.

A brief analysis of these missing domains in our dataset shows that 35 M are hostnames from *AWS*, a major public cloud provider, nearly 11 M `.blogspot.com` and related domains as well as about 3 M domains connected to network attached storage devices from *Synology*, *Western Digital* and *AVM* routers. Notably missing from our dataset are 4 M `.co.uk` domains, 3.8 M `.com` and 1 M domains from the `.tk` and `.ga` zone each. Concerning `.com`, we presume that these domains are missing because of timeouts and

[1] `https://www.rapid7.com/research/project-sonar/`

other errors during our scan, as the used zone file is a source of truth and there are only 20 k domains that the Rapid7 dataset contains and that are missing from our input list. The `.uk` zone is excluded from our scans, and we presume that the small number of resolved domains comes from `CNAME` records that are implicitly resolved. The other reason for differences in the seen ccTLDs is probably that we do not have an up-to-date list of these domains. Interestingly, our dataset contains 2.6 M `.de` domains missing in the Rapid7 data. Table A.3 contains a generic comparison of the datasets.

Table 5.10: The top eTLDs by the number of parked domains missing in our dataset (2022-01-28).

	# Domains		# Missing Domains	
eTLD	**Our Data**	**Rapid7**	**Our Data**	**Rapid7**
co.uk	19	585 195	585 176	0
com	42 868 119	43 208 024	580 239	240 334
co	658 940	445 642	128 992	342 290
uk	1	77 330	77 329	0
de	1 465 099	1 172 283	63 241	356 057
net	3 052 578	3 094 595	57 090	15 073
in	325 771	207 696	56 297	174 372
us	471 324	476 981	53 458	47 801
org	2 658 971	2 682 318	38 226	14 879
xn–p1ai[1]	2775	37 401	34 982	356
Sum	61 977 012	60 995 056	2 221 981	3 203 937

[1] This is `.ru` in Cyrillic.

Despite these differences, the total number of parked domains that we find in the Rapid7 dataset is 60 995 056, which is only about one million less than the result obtained using our dataset for the same day, 61 977 012. Analyzing the differences, we find that there are about 3.2 M parked domains in our dataset that are missing from the Rapid7 dataset and 2.2 M domains in Rapid7's dataset missing from ours. Analyzing the domains missing in our dataset (Table 5.10) shows, that there are about 585 k `.co.uk` domains (out of the 4 M we have already identified as missing above) and 580 k `.com` domains (22 % of all missing `.com` domains). On the other side, we find that `.de` has the largest amount of missing domains in Rapid7's dataset compared to ours (356 k), followed by `.co` (342 k), `.ca` (304 k) and `.com` (240 k).

Breaking this down by the parking services (Table 5.11), we find that there are about 273 k domains missing from *Sedo*, followed by *dan.com* (264 k) and GoDaddy's free parking service (209 k). Notable here is the absence of 194 k domains from *Tucows* that

Table 5.11: The top services by the number of parked domains missing in our dataset (2022-01-28).

	# Domains		# Missing Domains	
eTLD	**Our Data**	**Rapid7**	**Our Data**	**Rapid7**
Sedo	3 148 337	3 097 746	273 954	257 906
dan.com	2 232 563	2 470 374	264 403	23 933
GoDaddy (Free Parking)	30 041 322	28 920 270	209 767	1 331 002
Tucows	77	199 820	194 309	0
Namecheap	794 768	920 516	158 106	31 860
Skenzo	2 816 850	2 771 508	150 321	203 322
ParkLogic	355 603	477 675	137 097	17 071
above.com	919 787	1 028 183	127 372	19 099
survey-smiles.com	1 063 798	1 178 811	119 759	7512
Bodis	1 330 161	1 133 548	98 119	87 581
Sum	61 977 012	60 995 056	2 221 981	3 203 937

are all in our input dataset, but missing in the resolved records. Looking at historical data shows that *Tucows* seems to block our scans starting 2021-10-11.

Taking the union of both datasets, we find a total number of 64.2 M parked domains, showing that our results are only lower bounds.

5.5 Parking Services Faking Authority

While experimenting with DNS queries, we noticed that domain parking services that operate their own name servers often claim to answer authoritatively to all queries and resolve all domains to their parking infrastructure.

For example, Listing 5.2 shows that the *ParkingCrew* name servers answer a query for `google.com` with one of their parking servers. Of 17 advertisement based parking services using their own name servers, only four services, *Domaincntrol*, *snparking.ru*, *survey-smiles.com* and *Tucows* do not exhibit this behavior. Some services seem to have a block list of specific domains, presumably brands. For example, *Sedo* only answers with `127.0.0.1` when queried for `google.com` and `facebook.com`, but answers with its parking IP addresses for `tum.de`. We observe a similar behavior at *above.com*, *DomainSponsor* and *Skenzo*, although they answer with a `REFUSED` error or return no data instead of answering with `127.0.0.1`.

To further determine the extent of this practice, we conducted a scan including all the name servers we already know of from the *merged-NS* scan. In this scan we queried these name servers for three domains, `google.com`, `facebook.com` and `measr.net`, the

Listing 5.2: ParkingCrew answering a DNS query for google.com

```
; <<>> DiG 9.10.6 <<>> google.com @ns1.parkingcrew.net
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 43463
;; flags: qr aa rd; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0
;; WARNING: recursion requested but not available

;; QUESTION SECTION:
;google.com.                    IN      A

;; ANSWER SECTION:
google.com.             600     IN      A       185.53.177.51

;; Query time: 30 msec
;; SERVER: 13.248.158.159#53(13.248.158.159)
;; WHEN: Thu Jan 20 11:25:42 CET 2022
;; MSG SIZE  rcvd: 54
```

latter one not resolving at the time of the scan. The scan consists of 1.3 M name server IP addresses (2.7 M domains). Only 53 k name servers (287 k domains) replied with an answer out of which 42 k (261 k domains) answered with the authority flag set. We find 795 IP addresses of name servers resolving the queries to an IP address already known for parking (62 k domains). Out of these 62 k name server domains, 25 k are subdomains of `orderbox-dns.com` that apparently only answer with parking IP addresses for domains they are not authoritative for and answer with benign IP addresses otherwise. Out of the 1.3 M name server IP addresses, only 9973 answer a query for `measr.net` with the authority flag set, 37 001 for `google.com` and 21 532 for `facebook.com`.

Focusing on the name servers that respond to IP addresses not known for parking (Table 5.12), we find lots of Chinese IP addresses that only respond to queries for `google.com` and `facebook.com`, but not `measr.net`. We presume that these may be related to DNS censorship by the Chinese Firewall, as presented in an anonymous conference paper in 2014 [35]. The top AS, 46 606 is used by *Bluehost* and *HostGator*, that seemingly decided to just resolve all domains to their hosting servers, which in turn are also used as name servers.

Another interesting behavior that we noticed is that the *ParkingCrew* name servers resolve subdomains of any parked domain starting with `ns`, followed by an arbitrary number to its name servers. Therefore, a name server domain that is parked at *ParkingCrew*, e.g., because it recently expired, ends up parking every domain that still references this name server domain. An example of this is shown in Figure 5.4.

While we saw similar behavior with a few other services, *ParkingCrew* is the only service that deliberately resolves such domains to other IP addresses, while the other services

TABLE 5.12: The top autonomous systems by the number of name servers answering with the authority flag set to queries to `google.com`, `facebook.com` or `measr.net`. Name servers answering with known parking IP addresses are excluded. Data from 2021-12-03.

			# IP addresses	
ASN	**Country**	**Organization**	**Name Servers**	**Targets**[1]
46 606	US	UNIFIEDLAYER-AS-1	6476	2432
4837	CN	CHINA169-BACKBONE	3397	1063
37 963	CN	ALIBABA-CN-NET	3084	1027
4134	CN	CHINANET-BACKBONE	2563	750
4538	CN	ERX-CERNET-BKB	1038	487
55 960	CN	BJ-GUANGHUAN-AP	866	180
45 090	CN	TENCENT-NET-AP	816	456
135 629	CN	WESTCLOUDDATA	790	254
23 724	CN	CHINANET-IDC-BJ-AP IDC	783	214
38 283	CN	CHINANET-SCIDC-AS-AP	737	213
4812	CN	CHINANET-SH-AP	711	854
9808	CN	CHINAMOBILE-CN	684	205
203 391	BG	CLOUDNSNET	622	1
397 239	US	ULTRADNS	587	2
50 599	PL	Data Space Sp	516	3
58 466	CN	CT-GUANGZHOU-IDC	506	117
6939	US	HURRICANE	474	12
4808	CN	CHINA169-BJ China Unicom	467	264
16 276	FR	OVH	463	118
16 509	US	AMAZON-02	415	159

[1] The IP addresses resolved by the name servers in this autonomous system.

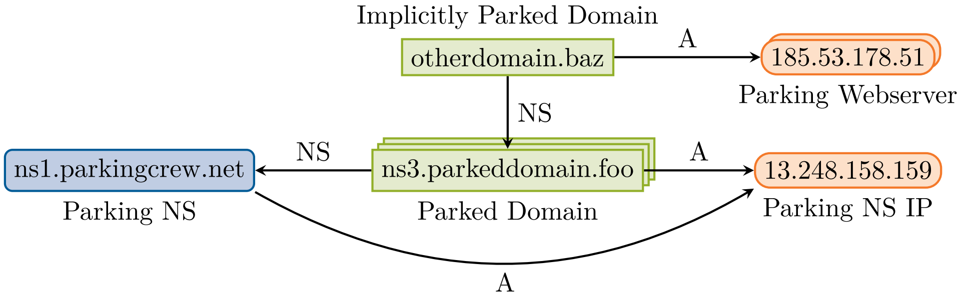

FIGURE 5.4: An exemplary setup of a domain that is parked because its name server is parked. The fictional domain `parkeddomain.foo` is parked and ends up implicitly parking `otherdomain.baz`, as it uses `ns3.parkeddomain.foo` as name server.

coincidentally have their web and DNS infrastructure on the same IP addresses and simply resolve all subdomains to the same addresses. Related work by Akiwate et al. investigated the renaming practice of expired domains [36] that are still referenced in the *Extensible Provisioning Protocol* and found that this can also lead to domain hijacking, stating that "parking sites dominat[e] the sample". In our scan, we find 308 domains parked at *ParkingCrew*, still being used as name server for up to 989 other domains. While we do not suggest that this is a case of domain hijacking (it could also be an intentional configuration), we note that this behavior could in fact lead to domains being hijacked.

5.6 LIMITATIONS

While our approach already shows that there is a substantial amount of unused, parked domains on the Internet, it does this purely by evaluating first-level DNS records directly tied to a specific domain. As the rules for matching records to parking services were manually assembled, changes in the used IP addresses or DNS names of name servers or aliases require manual adaptation. One way to improve this is to automatically link related records, for example IP addresses learned from known parked domains and using this transitive relationship to classify domains resolving to these new IP addresses as parked. The reason why we did not pursue this method further is that the used `NS` and `A` / `AAAA` scans do not run exactly at the same time and as the fluctuation of parked domains is quite high, we always have entries in our dataset where, e.g., the `NS` records indicate parking while the `A` record still points to a legitimate hosting provider. This can lead to false positives that we want to avoid, therefore we decided to not look into this further.

It is important to note that the used DNS scans are run only from one physical location, therefore we may not see all the IP addresses used by parking services, e.g., if they use the DNS to route traffic to different data centers based on geolocation.

Another aspect is that we found domains that use HTTP redirects from their apex to the `www.` subdomain and only this subdomain showed parking behavior. As our dataset heavily relies on zone files that do not include subdomains, these domains would not be classified as parked. An example where we found multiple thousands of domains with such a behavior is *Namecheap*, using the same infrastructure for their legitimate URL forwarding services, therefore making these domains indistinguishable from legitimate redirects to their customer's web pages with our DNS data alone. More sophisticated approaches using HTTP scans could detect these cases.

Chapter 6

Discussion: Impact on Further Research

Due to their content, showing only advertisements, sales pages or generic placeholders, parked domains differ greatly from other web content. Combining this with the massive amount of domain parking we were able to identify, especially the concentration of over 30 M domains at a single provider, *GoDaddy*, this can clearly affect DNS based research in multiple ways.

For example, Holz et al. analyzed the deployment of TLS 1.3 in 2019 [37] and used domain names from the `com.`, `net.` and `org.` zone files as well as the *Alexa Top 1 Million*. While the researchers briefly acknowledge that domains in the new gTLD zones are known for parking and defensive registrations to protect brands, they do not seem to consider the amount of parked domains in *com/net/org* and find that of 144 M resolving domains in the *com/net/org* zones only 5 % support TLS 1.3. Knowing that nearly 30 % of resolving domains in the *com/net/org* zones are parked domains, where the majority does not support TLS, changes the significance of this result. Similarly, in the unlikely event that all our found parking services suddenly adopt TLS 1.3, this amount could rise by up to 30 % without any real impact for most Internet users from a security perspective. It would be enough if GoDaddy decided to deploy TLS 1.3 to the 30 M domains using its free parking service to change the result of Holz et al. by up to 20 %. Similarly, research focusing on other web based protocols, such as *QUIC*, are affected as long as domain names are involved. In 2021, Zirngibl et al. analyzed the deployment of *QUIC* [38] using, among others, domains from zone files. Again, knowing that nearly 30 % of domains in the *com/net/org* zones are parked can help when interpreting results. With the ever rising importance of encrypted web traffic and browsers declaring plain HTTP connections as insecure [39], it is only a matter of time until the major parking services roll out TLS across their domain inventories.

We do not want to object any results, but explicitly only point out the importance of keeping parked domains in mind when performing DNS based research. The impact that a single parking service can have has been unintentionally shown by Zembruzki et al. in 2022 while researching the centralization of the hosting industry [18]. They found that *GoDaddy* migrated 27 M domains from its own autonomous system (AS) to the *Google Cloud*, making *Google* the largest AS for web hosting in 2021. Investigating these domains, they noticed that all domains were hosted on the same IP address, showing a parking page. This is consistent with the parked domains that we observe hosted with *GoDaddy*, also finding 27 M domains on a single IP address in *Google's* AS.

Chapter 7

Conclusion

By analyzing domain parking in existing DNS scans, we were able to show that there is a considerable amount of parked domain in all major DNS zones across a multitude of different domain parking services.

Among 267 M resolving domains in our input data sources, nearly 62 M can be classified as parked domains, showing ads, sales pages or merely placeholders with only questionable value for Internet users visiting them. In the popular zones `com.`, `net.` and `org.`, together more than 30 % of all domains resolving to an `A`, `AAAA` or `CNAME` record are parked. We found 25 services using advertisements on their parked domains, 30 services that show pages indicating that the domain is for sale and 23 services showing placeholder pages. Four services are hosting multiple categories of domains on the same infrastructure, therefore being categorized as mixed. The vast majority of parked domains are using one provider, *GoDaddy*, hosting nearly 27 M domains on a single IP address and over 30 M domains overall. We looked at the structure of parked domains and found that in the advertising category, there are only six services that provide their own parking pages, while all the others either use parked pages provided by these six services, or use a monetization method called PPR that directly redirects users from parked domains to an advertiser's landing page.

With a dataset of more than 132 B DNS records across the year 2021, we were able to show that there is a stable set of around 60 M parked domains and a high five-figure number of domains that are being newly parked or removed from parking from each scan to the next. By using open data available from *Rapid7*, we were able to analyze parked domains in a different dataset, confirming our results.

We discussed the possible impacts of these parked domains on further research and conclude that research based on domain names must consider these parked domains when drawing conclusions, in order to prevent bias that might occur purely based on the massive amount of parked domains found or simply to give a more nuanced conclusion about the significance of the presented results. For example, a sudden adoption of TLS by these parking services could drastically change results of TLS adoption studies.

While our work was purely based on DNS records, future work could extend this classification by also looking at HTTP scan results to detect parked domains using HTTP redirects. We also only looked briefly into the presence of domain parking in existing TLS scans. A more in-depth evaluation could for example focus on finding new parked domains based on the used certificates.

A major improvement could be made in the classification procedure itself, that currently relies on manually collected DNS configurations consisting of name servers, IP addresses and domain aliases that have been collected from manual research, as well as indicators found in the DNS scans themselves. While these indicators can provide a starting point for finding parking services, they always require verification, as similar patterns can be seen with big hosting providers or content delivery networks. This could be complemented by an analysis using automated web browsers and automatically evaluating the page's structure using a machine learning model.

Chapter A

Appendix

This chapter contains useful results that complement the thesis and topics that are not considered in-depth enough for the main content.

A.1 ClickHouse Experience Report

In this section, we provide a brief overview of our experiences using the *ClickHouse* database for analyzing DNS records and related data.

As briefly mentioned in Section 4.1 we use a *ClickHouse* database to process the data resulting from the DNS scans. Compared to databases like *PostgreSQL*, that are focused on transactional workloads and store the data for each row side by side, the *ClickHouse* database uses a columnar data storage where each column is saved (and also compressed) on its own. The approach followed by *PostgreSQL* and similar databases is also referred to as *OLTP* (Online Transaction Processing) and is very useful for applications where many transactions access few rows of data at once. The columnar approach is suitable for *OLAP* (Online Analytical Processing) workloads that mainly consist of aggregating data across large datasets. The *ClickHouse* does not support transactions at all, which is not a drawback for our use case. With the *ClickHouse* database, we are able to process more than 200 B of DNS records and domain lists and analyze them based on standard *SQL* queries. Compared to solutions like *Apache Spark*, the *ClickHouse* database is relatively easy to set up for users that already used a database like *PostgreSQL* or *MySQL*.

At the beginning of the thesis, we tried to implement our analysis based on a *PostgreSQL* database. Importing a full scan consisting of about 85 GB of uncompressed resource record files takes about 80 minutes and results in a *PostgreSQL* data directory size of roughly 100 GB. As the input files are line based, we use *GNU Parallel* [40] to ingest them into the database in parallel. The raw DNS records are filtered using a *PostgreSQL* trigger and records matching a known parking service are written into a separate table. The resulting table for a single imported day contains about 100 M rows and, without any indexes, it takes the *PostgreSQL* database six and a half minutes to count the number of distinct parked domains grouped by each service.

In comparison, we can import the same data into the *ClickHouse* database in only 17 minutes and this is not a fair comparison, as, due to disk space constraints, the *PostgreSQL* import skipped the *Certificate Transparency* scan results, which are the second-largest after the *merged-NS* scan. We presume that this improvement is largely due to the fact that the *ClickHouse* database uses an engine called *MergeTree*[1] that writes new data as fast as it can into temporary *parts*, that are then later sorted by a sorting key and merged in the background into bigger chunks. Comparing the query execution time of the grouped count, the *ClickHouse* database finishes in less than seven seconds, reporting itself:

Listing A.1: The *ClickHouse* database executing a grouped count for all parked domains on 2022-02-09.

```
:) select service, count(distinct domain) as num from records_parked where date =
    '2022-02-09' group by service order by num desc

SELECT
    service,
    countDistinct(domain) AS num
FROM records_parked
WHERE date = '2022-02-09'
GROUP BY service
ORDER BY num DESC

...

82 rows in set. Elapsed: 6.701 sec. Processed 112.90 million rows, 3.37 GB (16.85
    million rows/s., 503.43 MB/s.)
```

After importing, the *ClickHouse* uses about 20 GB of storage, before merging all temporary parts. Afterwards, the size of the data directory is only 11 GB using *zstd* compression for the table columns. This makes it possible for us to store a whole year of DNS scans using about 1.2 TB of storage, but notably excluding the *Certificate Transparency* scans.

A *ClickHouse* feature that was only released in November 2021 that we heavily rely on is *executable user defined functions*, that for example allows us to implement a function

[1] https://clickhouse.com/docs/en/engines/table-engines/mergetree-family/mergetree/

for getting the shortest private suffix of a domain using the `psl`[1] command line tool and using this function directly in our queries.

One downside of the *ClickHouse* database is high memory usage. Having enough memory, the database is happy to use it and when there is not enough, queries can and will fail with a *memory limit exceeded* error. During our analyses for this thesis, we ran queries grouping multiple days of records that used more than 500 GB of memory on the *galvos* testbed node. Performing joins can be particularly demanding. We were able to run out of memory using more than 980 GB while performing a left join on the parked domains table and the TLS scans. In this case, this could be mitigated by using *ANY* joins that do not match all entries, but only return the first matching entry for each key, instead of the default *ALL* join that performs like most database joins, returning all matching rows.

Considering scaling, the database is able to use all the available 24 CPU cores (48 threads) on the *intelexp0* testbed and seems to be able to scale well with more cores and memory available on the *galvos* testbed (64 cores, 128 threads, 1 TB of memory). Because of the background merges, it is important to have fast storage. We tried running the database on spinning hard disks and ran into issues when importing many scans at once, because the database was unable to merge the parts fast enough. With the *NVMe* storage on the *intelexp0* testbed this was not a problem.

For the results spanning an entire year of data, e.g., Figure 5.3, we first created temporary tables with filtered data for single days and then looped over every imported day. Overall, the *ClickHouse* database helped the analyses performed in this thesis tremendously. We can warmly recommend using it for similar purposes.

A.2 Supplementals

This section contains tables that were too large for being included in the main text, but are still considered useful to be referenced in the thesis.

Table A.1 shows the results of the top IP indicator explained in Section 4.3. For each IP address, we include the autonomous system and the corresponding organization as well as a description that explains what the specific IP address is used for.

In Table A.2 we show the second indicator from Section 4.3, the top name server domains by the *domain-ip ratio*. One can see that the top of the list consists almost completely of

[1] `https://github.com/rockdaboot/libpsl`

parking name servers, with the highest *domain-ip ratio* being more than 110 k domains per IP address. One important remark here is that due to the temporal difference in the `A` / `AAAA` and `NS` scans, not all IP addresses found connected to a name server domains must actually be related to it. For example, this is the reason that we find domains that according to our scans have a *Sedo* IP address, but *ParkingCrew* name servers. This domain would increment the number of IP addresses connected to the *ParkingCrew* name server domain and therefore decrease the *domain-ip ratio*.

Table A.3 compares our dataset to the *Rapid7* dataset with a focus on eTLDs where *Rapid7* sees more domains than we do. For each eTLD we include the total number of domains in each dataset and then use the set difference (`NOT IN` in SQL) to count the domains missing in each dataset. For our own scans we differentiate between the input and the resolved domains, while the *Rapid7* dataset only contains resolved domains.

The Tables A.4, A.5 and A.6 show the used DNS configuration for each parking service. We write the `NS` and `CNAME` configurations as regular expression and IP addresses in CIDR notation.

Table A.1: The top 50 IP addresses as described in Section 4.3. Data from 2022-01-02.

IP	# Domains	Parking	ASN	Hosting Organization	Description
34.102.136.180	26 806 853	✓	15 169	GOOGLE, US	GoDaddy Free Parking
3.33.152.147	9 926 949	×	16 509	AMAZON-02, US	GoDaddy Hosting
15.197.142.173	9 923 751	×	16 509	AMAZON-02, US	GoDaddy Hosting
185.230.63.171	5 541 710	×	58 182	WIX_COM, IL	wix.com Website Builder
185.230.63.186	5 522 809	×	58 182	WIX_COM, IL	wix.com Website Builder
185.230.63.107	5 522 561	×	58 182	WIX_COM, IL	wix.com Website Builder
198.185.159.144	3 346 137	×	53 831	SQUARESPACE, US	Squarespace Website Builder
198.185.159.145	3 293 154	×	53 831	SQUARESPACE, US	Squarespace Website Builder
34.98.99.30	3 292 346	✓	15 169	GOOGLE, US	GoDaddy Free Parking
198.49.23.145	3 291 531	×	53 831	SQUARESPACE, US	Squarespace Website Builder
198.49.23.144	3 288 302	×	53 831	SQUARESPACE, US	Squarespace Website Builder
35.186.238.101	2 198 098	✓	15 169	GOOGLE, US	GoDaddy CashParking
216.239.32.21	1 980 123	×	15 169	GOOGLE, US	Google Blogger / Blogspot
216.239.34.21	1 956 360	×	15 169	GOOGLE, US	Google Blogger / Blogspot
216.239.36.21	1 954 877	×	15 169	GOOGLE, US	Google Blogger / Blogspot
216.239.38.21	1 953 136	×	15 169	GOOGLE, US	Google Blogger / Blogspot
3.64.163.50	1 807 428	✓	16 509	AMAZON-02, US	dan.com Domains for Sale
64.190.62.111	1 720 644	✓	47 846	SEDO-AS, DE	Sedo Domain Parking
160.153.136.3	1 550 985	×	21 501	GODADDY-AMS, DE	Host Europe Webhosting
213.186.33.5	1 429 974	×	16 276	OVH, FR	OVH Webhosting
23.227.38.32	1 327 746	×	13 335	CLOUDFLARENET, US	Shopify
192.0.78.24	1 287 459	×	2635	AUTOMATTIC, US	wordpress.com
192.0.78.25	1 286 335	×	2635	AUTOMATTIC, US	wordpress.com
208.91.197.27	944 583	✓	40 034	CONFLUENCE-NETWORK-INC, VG	Skenzo Parking (NetworkSolutions)
23.236.62.147	929 260	×	15 169	GOOGLE, US	wix.com Website Builder
199.59.243.200	895 540	✓	16 509	AMAZON-02, US	Bodis Domain Parking
52.128.23.153	868 210	✓	19 324	DOSARREST, US	Uniregistry Parking (GoDaddy)
89.31.143.1	744 266	×	15 598	IPX-AS15598, DE	UnitedDomains Hosting
207.148.248.143	725 263	✓	29 873	BIZLAND-SD, US	domain.com Parking
99.83.154.118	687 652	✓	16 509	AMAZON-02, US	Namecheap (ParkingCrew)
23.227.38.65	679 867	×	13 335	CLOUDFLARENET, US	Shopify
217.70.184.38	679 437	×	29 169	GANDI-AS, FR	gandi.net Hosting + Placeholder
165.160.15.20	620 627	×	19 574	CSC, US	CSC Domain Protection and Redirects
165.160.13.20	616 248	×	19 574	CSC, US	CSC Domain Protection and Redirects
13.248.216.40	544 304	✓	16 509	AMAZON-02, US	afternic.com for sale (GoDaddy)
76.223.65.111	544 302	✓	16 509	AMAZON-02, US	afternic.com for sale (GoDaddy)
205.178.189.131	542 158	×	19 871	NETWORK-SOLUTIONS-HOSTING, US	Network Solutions Hosting
74.220.199.6	499 504	×	46 606	UNIFIEDLAYER-AS-1, US	BlueHost ParkingPage
62.149.128.157	476 886	×	31 034	ARUBA-ASN, IT	aruba.it Hosting
62.149.128.154	476 619	×	31 034	ARUBA-ASN, IT	aruba.it Hosting
62.149.128.160	476 476	×	31 034	ARUBA-ASN, IT	aruba.it Hosting
62.149.128.151	475 145	×	31 034	ARUBA-ASN, IT	aruba.it Hosting
75.2.26.18	465 353	✓	16 509	AMAZON-02, US	afternic.com for sale (GoDaddy)
99.83.153.108	465 223	✓	16 509	AMAZON-02, US	afternic.com for sale (GoDaddy)
192.161.187.200	459 460	×	8100	ASN-QUADRANET-GLOBAL, US	NameSilo Parking + Redirect
209.141.38.71	459 215	×	53 667	PONYNET, US	NameSilo Parking + Redirect
107.161.23.204	458 929	×	3842	RAMNODE, US	NameSilo Parking + Redirect
62.149.128.163	426 626	×	31 034	ARUBA-ASN, IT	aruba.it Hosting
62.149.128.166	424 822	×	31 034	ARUBA-ASN, IT	aruba.it Hosting

Table A.2: The top 75 name server domains by ip-domain ratio as described in Section 4.3. Data from 2022-01-18.

Domain	Parking	Service	# Domains	# IPs	Domains / IPs
domainmarket.com	✓	DomainMarket	220 167	2	110 084
domain-is-4-sale-at-domainmarket.com	✓	DomainMarket	220 167	2	110 084
panamans.com	✓	survey-smiles.com	438 353	8	54 794
thednscloud.com	✓	survey-smiles.com	160 574	3	53 525
domainprofi.de	✓	DomainProfi	221 284	7	31 612
power-dns.com	×		80 711	3	26 904
domainist.com	✓	Domainist	20 296	1	20 296
sonexo.eu	✓	Sonexo DNFS24	139 350	7	19 907
sonexo.com	✓	Sonexo DNFS24	139 350	7	19 907
cashparking.com	✓	GoDaddy (CashParking)	805 572	48	16 783
trustednam.es	✓	TRUSTEDNAMES	15 941	1	15 941
snparking.ru	✓	snparking.ru	110 284	7	15 755
com—type.in	✓	dan.com	27 866	2	13 933
nameprovider.net	✓	nameprovider.net	13 209	1	13 209
mytrafficmanagement.com	✓	TrafficMotor	369 762	30	12 325
123hjemmeside.dk	×		44 522	4	11 131
domain-for-sale.at	✓	dan.com	27 867	3	9289
this-domain.eu	✓	dan.com	27 867	3	9289
afternic.com	✓	AfterNic	973 755	105	9274
rookdns.com	✓	DomainSponsor	92 323	11	8393
dne.com	✓	dne.com	23 435	3	7812
gpk.eu	×		22 483	3	7494
expiereddnsmanager.com	✓	Expiereddnsmanager	37 295	5	7459
eftydns.com	✓	efty	182 705	25	7308
uniregistrymarket.link	✓	Uniegistry	856 540	119	7198
namefind.com	✓	GoDaddy (CashParking)	1 270 899	182	6983
domainhasexpired.com	✓	Domainbox	104 049	15	6937
domainrecover.com	✓	domainrecover.com	40 582	6	6764
emlakofisim.com	×		31 537	5	6307
brandbucket.com	✓	BrandBucket	132 478	22	6022
hostresolver.com	✓	survey-smiles.com	46 058	8	5757
undeveloped.com	✓	dan.com	370 258	65	5696
squadhelp.com	✓	Squadhelp	194 707	37	5262
dan.com	✓	dan.com	1 475 961	310	4761
parastorage.com	×		122 810	27	4549
domainsdirect.net	×		13 504	3	4501
parkingcrew.net	✓	ParkingCrew	1 349 318	307	4395
smartname.com	✓	GoDaddy (CashParking)	73 652	17	4332
dan.hosting	✓	dan.com	755 748	177	4270
dnsnuts.com	✓	survey-smiles.com	34 397	9	3822
rentondc.com	✓	survey-smiles.com	22 512	6	3752
dnslink.com	✓	Domainpower	22 136	6	3689
ack.de	×		10 441	3	3480
redirectdom.com	✓	domainname.de	194 117	58	3347
domain.io	✓	domain.io	23 250	7	3321
tacomadc.com	✓	survey-smiles.com	25 181	8	3148
marketo.co.uk	×		20 537	7	2934
takeaway.com	×		82 230	29	2836
jimdo.com	×		431 876	157	2751
domaintraffic.ch	×		14 641	6	2440
hastydns.com	✓	survey-smiles.com	43 957	19	2314
torresdns.com	✓	survey-smiles.com	45 666	20	2283
domainmx.com	✓	survey-smiles.com	47 689	21	2271
dsredirection.com	✓	DomainSponsor	33 320	17	1960
peoplebrowsr.com	×		32 805	17	1930
sedoparking.com	✓	Sedo	1 739 978	914	1904
shop-pro.jp	×		13 120	7	1874
domainorderdns.nl	✓	DomainOrder	29 896	16	1869
gridhost.com	×		10 963	6	1827
icasei.com.br	×		13 454	8	1682
nic.tel	×		26 352	16	1647
alphadnszone.com	×		21 019	13	1617
bodis.com	✓	Bodis	932 106	578	1613
nethit.de	×		15 921	10	1592
dns.ws	×		17 947	12	1496
day.biz	×		12 576	9	1397
smtmdns.com	✓	survey-smiles.com	13 734	10	1373
above.com	✓	above.com	613 379	451	1360
securetrafficrouting.com	✓	TrafficMotor	13 730	12	1144
tpiol.com	×		15 802	14	1129
milesmx.com	✓	survey-smiles.com	17 044	16	1065
promdns.net	×		29 380	29	1013
youcan.shop	×		30 415	32	950
emailverification.info	✓	domaindiscount24	82 592	91	908
f20.com	×		10 964	13	843

Table A.3: The top 50 public suffixes sorted by the number of domains missing in our dataset, as well as the total number of domains in our input and resolved dataset, compared to the Rapid7 data \ is the set difference. Data from 2022-01-28.

	# Domains						
Suffix	**Input**	**Resolved**	**Rapid7**	**Input \ Rapid7**	**Resolved \ Rapid7**	**Rapid7 \ Input**	**Rapid7 \ Resolved**
blogspot.com	76 468	76 552	9 407 131	41 147	41 121	9 371 809	9 371 699
us-west-2.compute.amazonaws.com	93	4535	8 531 894	0	0	8 531 801	8 527 359
eu-west-1.compute.amazonaws.com	27	4721	5 391 602	0	0	5 391 575	5 386 881
co.uk	2 865 098	3116	4 063 141	496 549	7	1 694 592	4 060 032
us-east-2.compute.amazonaws.com	42	2765	3 930 759	0	0	3 930 717	3 927 994
com	171 811 209	139 295 844	142 543 640	29 274 917	533 192	7348	3 780 988
ap-northeast-1.compute.amazonaws.com	45	1278	3 289 608	0	0	3 289 563	3 288 330
eu-central-1.compute.amazonaws.com	36	2620	3 208 973	1	1	3 208 938	3 206 354
herokuapp.com	927 653	934 536	3 491 313	86 991	87 319	2 650 650	2 644 095
ap-southeast-1.compute.amazonaws.com	16	1377	1 852 546	0	0	1 852 530	1 851 169
ap-southeast-2.compute.amazonaws.com	7	1283	1 670 968	0	0	1 670 961	1 669 685
eu-west-2.compute.amazonaws.com	10	612	1 509 359	0	0	1 509 349	1 508 747
ap-south-1.compute.amazonaws.com	11	798	1 444 061	0	0	1 444 050	1 443 263
us-west-1.compute.amazonaws.com	7	1092	1 376 522	0	0	1 376 515	1 375 430
myshopify.com	9571	686 117	2 036 544	2654	4617	2 029 626	1 355 043
ap-northeast-2.compute.amazonaws.com	18	330	1 249 762	0	0	1 249 744	1 249 432
ga	3 007 031	2 640 711	2 791 804	1 309 078	945 138	1 093 851	1 096 231
tk	29 810 435	755 383	1 407 224	28 758 056	409 216	354 845	1 061 057
remotewd.com	1 250 856	1 127 078	2 124 601	103 668	5815	977 413	1 003 338
direct.quickconnect.to	47 729	41 957	868 799	5781	23	826 851	826 865
myfritz.net	841 668	779 930	1 602 037	66 929	2124	827 297	824 230
ml	2 923 469	2 608 494	2 517 281	1 208 924	913 855	802 736	822 642
repl.co	221 573	223 013	1 039 321	217	221	817 964	816 528
ca-central-1.compute.amazonaws.com	1	370	786 944	0	0	786 943	786 574
sa-east-1.compute.amazonaws.com	6	581	756 630	0	0	756 624	756 049
gq	2 585 854	2 332 872	2 172 341	1 049 779	798 666	636 266	638 135
de	14 097 848	12 211 370	10 193 308	4 459 297	2 645 225	554 757	627 163
cn	8 003 254	2 822 902	1 868 691	6 691 500	1 540 225	556 937	586 014
com.br	3 311 079	2 001 521	2 311 764	1 496 352	247 459	497 037	557 702
ru	5 331 215	3 361 929	3 627 475	2 218 684	274 958	514 944	540 504
cf	4 668 312	3 321 444	2 509 488	2 609 653	1 326 546	450 829	514 590
kasserver.com	16	1390	433 907	3	4	433 893	432 520
net	14 064 735	10 815 342	11 189 086	2 878 129	32 852	2480	406 596
co.za	1 191 214	609 626	875 875	424 497	102 435	109 157	368 683
github.io	2046	120 379	468 158	554	1062	466 665	348 840
members.linode.com	12 064	8463	344 509	3897	1	336 341	336 046
nl	4 423 842	3 620 924	3 135 527	1 572 850	818 592	284 535	333 195
wixsite.com	1895	2071	333 606	583	589	332 294	332 124
fr	3 327 291	2 503 939	2 297 493	1 336 728	535 502	306 930	329 056
cn-north-1.compute.amazonaws.com.cn	76	159	327 619	0	0	327 543	327 460
co	2 330 786	1 503 520	1 167 991	1 466 960	655 810	304 165	320 281
blogspot.com.es	26 243	26 166	331 246	1125	1125	306 127	306 204
com.au	2 452 091	1 610 989	1 542 484	1 069 826	373 157	160 219	304 652
xn--p1ai	204 023	148 250	413 417	59 321	3989	268 715	269 156
uk	243 189	259	264 984	66 079	2	87 874	264 727
org	11 202 855	9 106 956	9 325 464	1 879 991	37 970	2600	256 478
blogspot.co.uk	42 983	1	252 430	19 653	0	229 099	252 428
cn-northwest-1.compute.amazonaws.com.cn	53	84	245 760	0	0	245 707	245 676
blogspot.com.br	51 970	51 854	271 946	22 548	22 509	242 523	242 600
cloudfront.net	3901	314 358	473 328	3050	81 983	472 477	240 953
Sum	380 287 669	267 735 212	332 957 412	125 873 007	22 401 375	79 092 113	87 623 575

Table A.4: The detailed DNS configuration for each service in the parking category.

<table>
<tr><th>Name</th><th>Name Servers</th><th>IP Addresses</th><th>Aliases</th></tr>
<tr><td>GoDaddy (Free Parking) [23]</td><td>×</td><td>34.102.136.180/32</td><td>×</td></tr>
<tr><td>Skenzo [41][42]</td><td>.*\.ztomy\.com
.*\.cnomy\.com
dns.*\.malkm\.com
dns.*\.parking-page\.net
.*\.searchreinvented\.com
ns.*\.nsresolution\.com</td><td>208.91.196.0/23
209.99.64.0/24
204.11.56.0/23
199.191.50.0/24</td><td>.*\.searchmagnified\.com</td></tr>
<tr><td>GoDaddy (CashParking) [43]</td><td>ns0(1|2)\.cashparking\.com
ns(1|2)\.smartname\.com
ns(1|2)\.namefind\.net</td><td>35.186.238.101/32</td><td>×</td></tr>
<tr><td>ParkingCrew [44][45][46]</td><td>ns(1|2)\.parkingcrew\.net
ns(1|2)\.parkingspa\.com
ns(1|2)\.fastpark\.net
ns(1|2)\.ibspark\.com</td><td>185.53.176.0/22</td><td>.*\.parkingcrew\.net
.*\.ndparking\.com
ndparking\.com</td></tr>
<tr><td>Bodis [47][48]</td><td>ns(1|2)\.bodis\.com</td><td>199.59.240.0/22</td><td>parking\.bodis\.com</td></tr>
<tr><td>survey-smiles.com</td><td>ns(.*)\.rentondc\.com
ns(.*)\.panamans\.com
ns(.*)\.torresdns\.com
ns(.*)\.thednscloud\.com
ns(.*)\.brainydns\.com
ns(.*)\.dnsnuts\.com
ns(.*)\.hostresolver\.com
ns(.*)\.kirklanddc\.com
ns(.*)\.tacomadc\.com
ns(.*)\.emu-dns\.com
ns(.*)\.milesmx\.com
ns(.*)\.smtmdns\.com
ns(.*)\.taipandns\.com
ns(.*)\.redmonddc\.com
ns(.*)\.domainmx\.com
ns(.*)\.commonmx\.com
ns(.*)\.wombatdns\.com
ns(.*)\.hastydns\.com
ns(.*)\.koaladns\.com
ns(.*)\.magpiedns\.com
ns(.*)\.dingodns\.com
ns(.*)\.namedynamics\.net
ns(.*)\.quokkadns\.com
ns(.*)\.weaponizedcow\.com
ns(.*)\.chookdns\.com</td><td>×</td><td>×</td></tr>
<tr><td>above.com [49]</td><td>.*ns.*\.above\.com</td><td>103.224.182.0/23
103.224.212.0/23</td><td>×</td></tr>
<tr><td>Namecheap</td><td>×</td><td>99.83.154.118/32</td><td>parkingpage\.namecheap\.com
parking\.namecheap\.com</td></tr>
<tr><td>Bluehost</td><td>×</td><td>74.220.199.6/32
74.220.199.8/32
74.220.199.9/32
74.220.199.14/32
74.220.199.15/32</td><td>×</td></tr>
<tr><td>TrafficMotor [50][51]</td><td>ns(1|2)\.mytrafficmanagement\.com
ns(1|2)\.securetrafficrouting\.com
ns(1|2)\.trafficcontrolrouter\.com
ns(1|2)\.searchfusion\.com</td><td>×</td><td>×</td></tr>
<tr><td>Dynadot</td><td>×</td><td>75.2.18.233/32
75.2.115.196/32</td><td>×</td></tr>
<tr><td>ParkLogic [52]</td><td>ns(1|2)\.parklogic\.com</td><td>×</td><td>pltraffic([0-9]+)\.com\.</td></tr>
<tr><td>snparking.ru [53]</td><td>ns([0-9]+)\.snparking\.ru
ns([0-9]+)\.salenames\.ru</td><td>×</td><td>×</td></tr>
<tr><td>DomainSponsor [54]</td><td>ns(1|2)\.dsredirection\.com
ns(1|2)\.dsredirects\.com
ns(.*)\.rookdns\.com
ns(.*)\.dnsspark\.com</td><td>66.81.199.0/24
141.8.224.195/32</td><td>.*\.fwdservice\.com
.*\.?dsredirects\.com
.*\.?dsredirection\.com</td></tr>
<tr><td>123 Reg</td><td>×</td><td>35.227.197.36/32</td><td>×</td></tr>
<tr><td>Voodoo</td><td>ns.*\.voodoo\.com</td><td>192.64.147.0/24</td><td>×</td></tr>
<tr><td>Expiereddnsmanager</td><td>ns.*\.expiereddnsmanager\.com</td><td>×</td><td>×</td></tr>
<tr><td>traffic.club [55]</td><td>.*\.sslparking\.com
ns.*\.ndsplitter\.com</td><td>×</td><td>cname\.sslparking\.com</td></tr>
<tr><td>Domainpower [56]</td><td>ns.*\.dnslink\.com</td><td>×</td><td>×</td></tr>
<tr><td>Fabulous / Directnic</td><td>expired-domain-ns.*\.fabulous\.com
expired-domain-ns.*\.directnic\.com</td><td>×</td><td>×</td></tr>
<tr><td>Internetvikings</td><td>parkdns(1|2)\.internetvikings\.com</td><td>×</td><td>×</td></tr>
<tr><td>Domaincntrol</td><td>.*\.domaincntrol\.com</td><td>×</td><td>×</td></tr>
<tr><td>The Parking Place [57]</td><td>ns(1|2)\.pql\.net
ns(3|4)\.tppns\.com</td><td>×</td><td>×</td></tr>
<tr><td>Tucows</td><td>ns.*\.renewyourname\.net</td><td>×</td><td>×</td></tr>
</table>

Table A.5: The detailed DNS configuration for each service in the marketplace category.

Name	Name Servers	IP Addresses	Aliases
HugeDomains.com	×	×	traff-([0-9]+)\.hugedomains\.com
dan.com [58]	ns(1|2)\.dan\.com ns(1|2)\.undeveloped\.com verification-.*\.ns.*\.dan\.hosting domain-for-sale\.at www\.undeveloped.com----type\.in your-browser\.this-domain\.eu ns(1|2)\.park\.do	3.64.163.50/32 52.58.78.16/32 2a05:d014:9da:8c10:306e:3e07:a16f:a552/128	×
AfterNic [59]	ns[0-9]+\.afternic\.com	×	×
domain.com	×	207.148.248.143/32 207.148.248.145/32	×
Aliyun	×	×	overdue\.aliyun\.com
DomainMarket	ns1\.domain-is-4-sale-at-domainmarket\.com ns2\.domainmarket\.com	×	×
Squadhelp [60]	ns(1|2)\.squadhelp\.com	×	×
DomainProfi	ns(1|2)\.domainprofi\.de	×	×
epik	×	45.88.202.115/32	×
domainname.de	ns(1|2)\.redirectdom\.com	×	×
REG.RU	×	194.58.112.165/32 194.58.112.174/32	×
efty	ns(1|2)\.eftydns\.com	×	×
Sonexo DNFS24	ns1\.sonexo\.eu ns2\.sonexo\.com	×	×
Domainparking.ru	ns(1|2)\.domainparking\.ru	31.31.205.163/32	×
BrandBucket [61]	ns(1|2)\.brandbucket\.com	×	×
domainmarkt.de	×	46.4.13.97/32	×
PerfectDomain [62]	ns(1|2)\.perfectdomain\.com	164.90.244.158/32 159.89.244.183/32	×
22.cn	×	119.28.128.52/32	×
WangGuai.com	×	43.128.56.249/32	×
domainrecover.com	ns(1|2)\.domainrecover\.com	×	×
LinkUWant	×	199.58.179.10/32	×
Alter	ns(1|2)\.alter\.com	3.234.55.179/32	×
DomainOrder	ns1?\.domainorderdns\.nl	×	×
Domainist	ns(1|2)\.domainist\.com	×	×
DNSPod	×	×	domainparking-dnspod\.cn
Domain Brokers Sweden	ns(1|2)\.domain-for-sale\.se	×	×
TRUSTEDNAMES	ns(1|2)\.trustednam\.es	×	×
nameprovider.net	ns(1|2)\.nameprovider\.net	×	×
domain.io [63]	ns(1|2)\.domain\.io	×	×
Flippa	×	×	webapp\.parkingpage\.production\.flippa.com

Table A.6: The detailed DNS configuration for each service in the placeholder or mixed category.

<table>
<tr><th>Name</th><th>Name Servers</th><th>IP Addresses</th><th>Aliases</th></tr>
<tr><td colspan="4">Placeholder</td></tr>
<tr><td>Alibaba</td><td>×</td><td>47.91.170.222/32</td><td>×</td></tr>
<tr><td>123 Reg</td><td>×</td><td>94.136.40.51/32</td><td>×</td></tr>
<tr><td>Hostnet.nl</td><td>×</td><td>91.184.0.100/32</td><td>×</td></tr>
<tr><td>transip</td><td>×</td><td>37.97.254.27/32</td><td>×</td></tr>
<tr><td>Hostinger</td><td>×</td><td>2.57.90.16/32</td><td>×</td></tr>
<tr><td>Hostpoint</td><td>×</td><td>217.26.48.101/32</td><td>×</td></tr>
<tr><td>domainname.shop</td><td>×</td><td>185.134.245.113/32</td><td>×</td></tr>
<tr><td>fasthosts</td><td>×</td><td>213.171.195.105/32</td><td>×</td></tr>
<tr><td>Namebright</td><td>×</td><td>×</td><td>comingsoon\.namebright\.com
expired\.namebright\.com</td></tr>
<tr><td>Namecheap</td><td>failed-whois-verification\.namecheap\.com
verify-contact-details\.namecheap\.com</td><td>×</td><td>×</td></tr>
<tr><td>one.com</td><td>×</td><td>46.30.211.38/32</td><td>×</td></tr>
<tr><td>HostGator</td><td>×</td><td>192.185.0.218/32</td><td>×</td></tr>
<tr><td>Domainbox</td><td>ns(1|2)\.domainhasexpired\.com</td><td>×</td><td>×</td></tr>
<tr><td>west.cn</td><td>×</td><td>103.139.0.32/32</td><td>×</td></tr>
<tr><td>domaindiscount24</td><td>ns(1|2)\.expirationwarning\.net
ns(1|2)\.emailverification\.info</td><td>×</td><td>×</td></tr>
<tr><td>forpsi internet CZ</td><td>×</td><td>81.2.194.128/32</td><td>×</td></tr>
<tr><td>Enom</td><td>×</td><td>98.124.204.16/32</td><td>×</td></tr>
<tr><td>101domain</td><td>×</td><td>52.60.87.163/32</td><td>×</td></tr>
<tr><td>gabia</td><td>×</td><td>121.254.178.252/32</td><td>×</td></tr>
<tr><td>Markmonitor</td><td>×</td><td>93.191.168.52/32</td><td>×</td></tr>
<tr><td>wedos</td><td>×</td><td>46.28.105.2/32</td><td>×</td></tr>
<tr><td>aruba.it</td><td>×</td><td>62.149.128.40/32</td><td>×</td></tr>
<tr><td>Turkticaret.Net</td><td>×</td><td>31.186.11.254/32</td><td>×</td></tr>
<tr><td colspan="4">Mixed</td></tr>
<tr><td>Sedo [64][65]</td><td>c?ns(1|2)\.sedoparking\.com</td><td>91.195.240.0/23
64.190.62.0/23</td><td>.*\.sedoparking\.com</td></tr>
<tr><td>sav.com [66]</td><td>ns(1|2)\.sav\.com</td><td>×</td><td>prod-sav-park-lb01-1919960993\.us-east-2\.elb\.amazonaws\.com
.*\.dsextra\.com</td></tr>
<tr><td>Uniegistry [67]</td><td>ns(1|2)\.uniregistrymarket\.link
.*\.expired\.uniregistry-dns\.com</td><td>52.128.23.153/32
34.102.221.37/32</td><td>×</td></tr>
<tr><td>dne.com</td><td>ns(1|2)\.dne\.com</td><td>×</td><td>×</td></tr>
</table>

Chapter B

List of acronyms

ccTLD	Country code top-level domain. A two-letter TLD intended for use by a country, e.g., *.de*.
CIDR	Classless Inter-Domain Routing.
CZDS	Centralized Zone Data Service.
DNS	Domain Name System.
eTLD	Effective top-level domain.
FQDN	Fully qualified domain name.
gTLD	Generic top-level domain. A domain without geographic or country desgination, intended for general use.
HTML	Hypertext Markup Language.
HTTP	Hypertext Transfer Protocol.
ICANN	Internet Corporation for Assigned Names and Numbers.
IP	Internet Protocol.
PPC	Pay-Per-Click. A monetization mechanism where the user has to actively click, e.g., on a keyword.
PPR	Pay-Per-Redirect. A monetization mechanism where the user is redirected directly to the advertiser's landing page.
TLD	Top-level domain. A domain at the highest level of the Domain Name System, e.g., *.com* or *.de*.
TLS	Transport Layer Security.
TTL	Time to live.
URL	Uniform Resource Locator.

Bibliography

[1] *Sedo company details, the best place for domains is Sedo.com.* [Online]. Available: https://sedo.com/us/about-us/ (visited on 01/22/2022).

[2] H. Oßing, P. Sattler, and J. Zirngibl, „Analysis of Domain Parking“, Seminar Innovative Internet Technologies and Mobile Communications (IITM), Aug. 2021.

[3] J. Styler, *The top 25 most expensive domain names*, Jun. 2019. [Online]. Available: https://www.godaddy.com/garage/the-top-20-most-expensive-domain-names/ (visited on 01/22/2022).

[4] *Domain names - concepts and facilities*, RFC 1034, Nov. 1987. DOI: 10.17487/RFC1034. [Online]. Available: https://rfc-editor.org/rfc/rfc1034.txt.

[5] brooke, *Type in Traffic*, Nov. 2009. [Online]. Available: https://www.above.com/blog/type-in-traffic/ (visited on 02/13/2022).

[6] *Search quality highlights: new monthly series on algorithm changes.* [Online]. Available: https://search.googleblog.com/2011/12/search-quality-highlights-new-monthly.html (visited on 01/27/2022).

[7] T. Moore and B. Edelman, „Measuring the Perpetrators and Funders of Typosquatting“, in *Financial Cryptography and Data Security*, R. Sion, Ed., ser. Lecture Notes in Computer Science, Berlin, Heidelberg: Springer, 2010, pp. 175–191, ISBN: 978-3-642-14577-3. DOI: 10.1007/978-3-642-14577-3_15.

[8] *Sign In | Centralized Zone Data Service.* [Online]. Available: https://czds.icann.org/home (visited on 02/11/2022).

[9] Alexa. (2021). Top 1M sites, [Online]. Available: https://www.alexa.com/topsites.

[10] Q. Scheitle, O. Hohlfeld, J. Gamba, J. Jelten, T. Zimmermann, S. D. Strowes, and N. Vallina-Rodriguez, „A Long Way to the Top: Significance, Structure, and Stability of Internet Top Lists“, in *Proceedings of the Internet Measurement Conference 2018*, ser. IMC '18, New York, NY, USA: Association for Computing Machinery, Oct. 2018, pp. 478–493, ISBN: 978-1-4503-5619-0. DOI: 10.1145/3278532.3278574.

[11] S. Alrwais, K. Yuan, E. Alowaisheq, Z. Li, and X. Wang, „Understanding the Dark Side of Domain Parking“, in *Proceedings of the 23rd USENIX Conference on Security Symposium*, ser. SEC'14, San Diego, CA: USENIX Association, 2014, 207–222, ISBN: 9781931971157.

[12] T. Vissers, W. Joosen, and N. Nikiforakis, „Parking Sensors: Analyzing and Detecting Parked Domains“, in *Proceedings 2015 Network and Distributed System Security Symposium*, San Diego, CA: Internet Society, 2015, ISBN: 978-1-891562-38-9. DOI: 10.14722/ndss.2015.23053.

[13] *Home - DNS Census 2013*. [Online]. Available: https://dnscensus2013.neocities.org/ (visited on 02/11/2022).

[14] 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 *Passive and Active Measurement*, N. Taft and F. Ricciato, Eds., ser. Lecture Notes in Computer Science, Berlin, Heidelberg: Springer, 2012, pp. 221–230, ISBN: 978-3-642-28537-0. DOI: 10.1007/978-3-642-28537-0_22.

[15] T. Halvorson, K. Levchenko, S. Savage, and G. M. Voelker, „XXXtortion?: inferring registration intent in the .XXX TLD“, in *Proceedings of the 23rd international conference on World wide web - WWW '14*, Seoul, Korea: ACM Press, 2014, pp. 901–912, ISBN: 978-1-4503-2744-2. DOI: 10.1145/2566486.2567995.

[16] T. Halvorson, M. F. Der, I. Foster, S. Savage, L. K. Saul, and G. M. Voelker, „From .academy to .zone: An Analysis of the New TLD Land Rush“, in *Proceedings of the 2015 Internet Measurement Conference*, Tokyo Japan: ACM, Oct. 2015, pp. 381–394, ISBN: 978-1-4503-3848-6. DOI: 10.1145/2815675.2815696.

[17] T. Lauinger, A. Chaabane, A. S. Buyukkayhan, K. Onarlioglu, and W. Robertson, „Game of Registrars: An Empirical Analysis of Post-Expiration Domain Name Takeovers“, 2017, pp. 865–880, ISBN: 978-1-931971-40-9. [Online]. Available: https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/lauinger (visited on 02/04/2022).

[18] L. Zembruzki, R. Sommese, L. Z. Granville, A. S. Jacobs, M. Jonker, and G. C. M. Moura, „Hosting Industry Centralization and Consolidation“, *arXiv:2109.01187 [cs]*, Jan. 2022, arXiv: 2109.01187. [Online]. Available: http://arxiv.org/abs/2109.01187 (visited on 02/07/2022).

[19] *Public Suffix List*. [Online]. Available: https://publicsuffix.org/ (visited on 01/22/2022).

[20] M. Kührer, C. Rossow, and T. Holz, „Paint It Black: Evaluating the Effectiveness of Malware Blacklists“, in *Research in Attacks, Intrusions and Defenses*, A. Stavrou, H. Bos, and G. Portokalidis, Eds., ser. Lecture Notes in Computer

Science, Cham: Springer International Publishing, 2014, pp. 1–21, ISBN: 978-3-319-11379-1. DOI: 10.1007/978-3-319-11379-1_1.

[21] *advice - Which is the best domain parking service?* [Online]. Available: https://www.namepros.com/threads/which-is-the-best-domain-parking-service.989749/ (visited on 01/22/2022).

[22] *Registrar Stats: Top Domain Registrars - DomainState.com.* [Online]. Available: https://www.domainstate.com/top-registrars.html (visited on 02/03/2022).

[23] *Godaddy - Park a domain registered with GoDaddy.* [Online]. Available: https://www.godaddy.com/help/park-a-domain-registered-with-godaddy-23936 (visited on 01/22/2022).

[24] *CashParking | Turn Domain Parking into Cash - GoDaddy.* [Online]. Available: https://www.godaddy.com/domains/cashparking (visited on 01/22/2022).

[25] *GoDaddy buys Afternic and SmartName from NameMedia*, Sep. 2013. [Online]. Available: https://namemonetizer.com/2013/09/21/2013-9-21-godaddy-buys-afternic-and-smartname-from-namemedia-html/ (visited on 01/22/2022).

[26] *GoDaddy Acquires Uniregistry's Registrar, Marketplace & Portfolio.* [Online]. Available: https://aboutus.godaddy.net/newsroom/press-releases/press-release-details/2020/GoDaddy-Acquires-Uniregistrys-Registrar-Marketplace--Portfolio/default.aspx (visited on 01/22/2022).

[27] W. c. G. Inc, *Newfold Digital Launches as One of the World's Largest Web Solutions Providers*, May 2021. [Online]. Available: https://www.globenewswire.com/en/news-release/2021/05/19/2232456/9865/en/Newfold-Digital-Launches-as-One-of-the-World-s-Largest-Web-Solutions-Providers.html (visited on 01/22/2022).

[28] *Web Host Endurance International Bulks Up With Buys | Nasdaq.* [Online]. Available: https://www.nasdaq.com/articles/web-host-endurance-international-bulks-buys-2015-04-08 (visited on 01/22/2022).

[29] *CentralNic closes transformational acquisition of KeyDrive*, Jul. 2018. [Online]. Available: https://www.proactiveinvestors.co.uk/companies/news/200828/centralnic-closes-transformational-acquisition-of-keydrive-200828.html (visited on 01/22/2022).

[30] *Global domain registry provider CentralNic to acquire Team Internet for USD48 million | Private Equity Wire*, Section: Deals & Transactions, Nov. 2019. [Online]. Available: http://www.privateequitywire.co.uk/2019/11/19/280582/global-domain-registry-provider-centralnic-acquire-team-internet-usd48-million (visited on 01/22/2022).

[31] *WIPO Domain Name Decision: D2021-1514*. [Online]. Available: https://www.wipo.int/amc/en/domains/decisions/text/2021/d2021-1514.html (visited on 01/26/2022).

[32] Zhou Li, S. Alrwais, Yinglian Xie, Fang Yu, and XiaoFeng Wang, „Finding the Linchpins of the Dark Web: a Study on Topologically Dedicated Hosts on Malicious Web Infrastructures“, in *2013 IEEE Symposium on Security and Privacy*, Berkeley, CA: IEEE, May 2013, pp. 112–126, ISBN: 978-0-7695-4977-4 978-1-4673-6166-8. DOI: 10.1109/SP.2013.18.

[33] *.realty | Top-Level Domain Breakdown*. [Online]. Available: https://ntldstats.com/tld/realty (visited on 01/29/2022).

[34] *Forward DNS (FDNS) - Rapid7 Open Data*. [Online]. Available: https://opendata.rapid7.com/sonar.fdns_v2/ (visited on 01/23/2022).

[35] „Towards a Comprehensive Picture of the Great Firewall's DNS Censorship“, 2014. [Online]. Available: https://www.usenix.org/conference/foci14/workshop-program/presentation/anonymous (visited on 02/06/2022).

[36] G. Akiwate, S. Savage, G. M. Voelker, and K. C. Claffy, „Risky BIZness: risks derived from registrar name management“, in *Proceedings of the 21st ACM Internet Measurement Conference*, New York, NY, USA: Association for Computing Machinery, Nov. 2021, pp. 673–686, ISBN: 978-1-4503-9129-0.

[37] R. Holz, J. Amann, A. Razaghpanah, and N. Vallina-Rodriguez, „The Era of TLS 1.3: Measuring Deployment and Use with Active and Passive Methods“, *arXiv:1907.12762 [cs]*, Aug. 2019, arXiv: 1907.12762. [Online]. Available: http://arxiv.org/abs/1907.12762 (visited on 01/03/2022).

[38] J. Zirngibl, P. Buschmann, P. Sattler, B. Jaeger, J. Aulbach, and G. Carle, „It's over 9000: analyzing early QUIC deployments with the standardization on the horizon“, in *Proceedings of the 21st ACM Internet Measurement Conference*, Virtual Event: ACM, Nov. 2021, pp. 261–275, ISBN: 978-1-4503-9129-0. DOI: 10.1145/3487552.3487826.

[39] *A secure web is here to stay*. [Online]. Available: https://security.googleblog.com/2018/02/a-secure-web-is-here-to-stay.html (visited on 02/08/2022).

[40] O. Tange, *GNU Parallel 20220122 ('20 years')*, GNU Parallel is a general parallelizer to run multiple serial command line programs in parallel without changing them., Jan. 2021. DOI: 10.5281/zenodo.5893336.

[41] *Ztomy.com*. [Online]. Available: https://www.namepros.com/threads/ztomy-com.806046/ (visited on 02/13/2022).

[42] j, *WARNING! If You Own Domain Names, Avoid Web.com and Its Subsidiaries and Partners (Register.com, Network Solutions, and New Ventures Services – A.K.A. "New Vultures Services" – Among Others)*, Jul. 2014. [Online]. Available:

https://www.scamful.com/2014/07/warning-if-you-own-domain-names-avoid.html (visited on 11/27/2021).

[43] *Godaddy - Add domains to my CashParking account.* [Online]. Available: https://www.godaddy.com/help/add-domains-to-my-cashparking-account-2316 (visited on 01/22/2022).

[44] *ParkingCrew – Why is using ParkingCrew nameservers the best choice for parking?* [Online]. Available: https://www.parkingcrew.com/faq.php?mode=detail&detail=83 (visited on 01/23/2022).

[45] *ParkingCrew – Can I park my domains using a CNAME record?* [Online]. Available: https://www.parkingcrew.com/faq.php?mode=detail&detail=82 (visited on 01/23/2022).

[46] *ParkingCrew – Can I still park with Namedrive DNS?* [Online]. Available: https://www.parkingcrew.com/faq.php?mode=detail&detail=50 (visited on 02/13/2022).

[47] *BODIS. - How do I park my domain names with Bodis?* [Online]. Available: https://www.bodis.com/support#reamaze#0#/kb/parking-domains/how-do-i-park-my-domain-names-with-bodis (visited on 01/23/2022).

[48] *BODIS. - Does Bodis support CNAME forwarding?* [Online]. Available: https://www.bodis.com/support#reamaze#0#/kb/parking-domains/does-bodis-support-cname-forwarding (visited on 01/23/2022).

[49] *above.com - Add Domains - Portfolio Manager Manual.* [Online]. Available: https://www.above.com/manual/pm-add-domains.html (visited on 01/23/2022).

[50] DomainWatch, *WHOIS Information for mytrafficmanagement.com - DomainWatch.* [Online]. Available: https://domainwat.ch/whois/mytrafficmanagement.com (visited on 02/13/2022).

[51] *Trafficmotor.com.* [Online]. Available: https://trafficmotor.com/ (visited on 02/13/2022).

[52] *Registrar Auto Install DNS Settings - ParkLogic.* [Online]. Available: https://www.parklogic.com/images/content/Registrar_Technical_Implementation.pdf (visited on 02/13/2022).

[53] SNPARKING, *SNPARKING.RU.* [Online]. Available: https://www.snparking.ru/ (visited on 02/13/2022).

[54] *DomainSponsor.com | FAQ's*, Aug. 2021. [Online]. Available: https://web.archive.org/web/20210817082328/https://domainsponsor.com/FAQ/#toggle-id-8 (visited on 02/13/2022).

[55] *TrafficClub.com Domainparking.* [Online]. Available: https://www.sslparking.com/ (visited on 02/13/2022).

[56] *Frequently Asked Questions | DomainPower*. [Online]. Available: https://www.domainpower.com/home/faq (visited on 02/13/2022).

[57] *"Frequently Asked Questions | Monetizing Domains | The Parking Place"*. [Online]. Available: https://theparkingplace.com/faq (visited on 02/13/2022).

[58] Dan.com, *Your Own Domain For Sale Page*, Jul. 2021. [Online]. Available: https://blog.undeveloped.com/your-own-for-sale-page-adc1b66a3638 (visited on 02/13/2022).

[59] *Visit Afternic's FAQ Page for More Domain Help - Afternic*. [Online]. Available: https://www.afternic.com/domain-help#1 (visited on 02/13/2022).

[60] *Sellers: How to Update Your Name Servers*. [Online]. Available: http://helpdesk.squadhelp.com/en/articles/2770338-sellers-how-to-update-your-name-servers (visited on 02/13/2022).

[61] *BrandBucket: FAQ for domain sellers and brand inventors*. [Online]. Available: https://www.brandbucket.com/help/faq/sellers (visited on 02/13/2022).

[62] perfectdomain.com, *FAQs - PerfectDomain.com*. [Online]. Available: https://perfectdomain.com/faqs (visited on 02/13/2022).

[63] *Domain Portfolio Management - Domain.io*. [Online]. Available: https://domain.io/ (visited on 02/13/2022).

[64] *Domain parking | Earn money with parked domains | Sedo*. [Online]. Available: https://sedo.com/us/park-domains/ (visited on 02/13/2022).

[65] *Does Sedo offer an alternative method to name server parking?* [Online]. Available: https://faq-us.sedo.com/app/answers/detail/a_id/3527 (visited on 02/13/2022).

[66] *About Parked Domains*. [Online]. Available: https://help.sav.com/kb/en/article/about-parked-domains (visited on 02/13/2022).

[67] *Setting Uniregistry Market Nameservers*. [Online]. Available: https://www.uniregistry.help/hc/en-us/articles/115001726928-Setting-Uniregistry-Market-Nameservers (visited on 02/13/2022).