

An Empirical Study of Orphan DNS Servers in the Internet

Andrew J. Kalafut^{*}
School of Computing
and Information Systems
Grand Valley State University
Allendale, MI, USA
kalafuta@gvsu.edu

Minaxi Gupta
School of Informatics
and Computing
Indiana University
Bloomington, IN, USA
minaxi@cs.indiana.edu

Christopher A. Cole
School of Informatics
and Computing
Indiana University
Bloomington, IN, USA
coleca@cs.indiana.edu

Lei Chen
School of Informatics
and Computing
Indiana University
Bloomington, IN, USA
lc6@cs.indiana.edu

Nathan E. Myers
School of Informatics
and Computing
Indiana University
Bloomington, IN, USA
natmyers@cs.indiana.edu

ABSTRACT

An *orphan DNS server* is a DNS server which has an address record in the DNS, even though the domain in which it resides has no DNS records itself and hence does not exist. For example, the DNS server `ns.foo.com` would be an orphan DNS server if it had an address record, but the domain `foo.com` did not exist. In this paper, we undertake the first systematic study of the prevalence of orphan DNS servers in the Internet. We also examine who is using them and what they are used for. We find that certain top-level domains (TLDs) account for a disproportionate number of orphans. We also find that some orphans are used for malicious activities and as placeholders for records from deleted domains, while others likely only exist due to simple configuration errors. Our study points to the need for better scrutiny of orphan DNS servers so they cannot be misused.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols—DNS; C.2.5 [Local and Wide-Area Networks]: Internet

General Terms

Measurement, Security

Keywords

DNS, Top Level Domains

^{*}This research was conducted while Andrew Kalafut was a Ph.D. student at Indiana University, Bloomington, IN, USA

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC'10, November 1–3, 2010, Melbourne, Australia.

Copyright 2010 ACM 978-1-4503-0057-5/10/11 ...\$10.00.

1. INTRODUCTION

The Domain Name System (DNS) [11, 12] is an important component of the Internet's infrastructure. Though it carries out a multitude of functions, its primary task is to translate host names to IP addresses. Internet users rely on the DNS every time they wish to address a host by its name.

The DNS is a hierarchical system, beginning at the root domain. Directly below the root, are the top-level domains (TLDs), such as `.com`. Below each TLD are domains falling within that TLD, such as `example.com`. The DNS servers for each TLD are responsible for storing records used to locate the DNS servers for each domain within that TLD. Clients need these records in order to find the addresses of hosts within those domains. For example, the DNS server for `.com` will store records for the name server¹ for `example.com`. Commonly, these records are of two types, those that contain the host name of the domain's name server (associating `example.com` with `ns1.example.com`), and those that contain its IP address (associating `ns1.example.com` with an address). This paper focuses on a special case of IP addresses records known to the TLD DNS servers: those where an IP address record exists for a DNS server whose parent domain has ceased to exist. An example of such a case is if there were an IP address record for `ns1.example.com` at the `.com` TLD when `example.com` itself was non-existent. Servers such as `ns1.example.com` have come to be referred to as *orphan name servers* in the operational community.

Orphan DNS servers have been found in use by malicious web sites [17]. This paper is motivated by curiosity about the existence of this phenomenon. Accordingly, we focus on studying the prevalence of orphan name servers in the Internet and to understand what, if anything, they are used for. We are also interested in understanding the impact of removing orphan servers. Specifically, if they are only used by malicious domains, it is actually beneficial to remove them. However, if they are used by legitimate domains as well, users of these name servers would see their DNS service degraded or interrupted if such servers were removed. To the

¹The terms *DNS server* and *name server* are used interchangeably.

best of our knowledge, ours is the first work to investigate this phenomenon systematically. The questions we ask are:

- How prevalent are orphan DNS servers?
- How long do orphan name servers live?
- Are orphan name servers being used for purposes other than as DNS servers?
- Are orphan name servers commonly used for malicious purposes?

We examine an average of 106 million domains (60% of domains in the Internet at the time of our measurements) on a daily basis for a period of 31 days to investigate aspects of orphan DNS servers. Searching for orphans in these domains, we find that 1.7% of DNS servers we examine are orphans. Almost 1/3rd of these are being used as DNS servers by domains we can identify. Many others are used as web and email servers. We confirm that each of the 1% of domains using orphans as their DNS servers are web sites listed in phishing and malware blacklists, and a significant portion of IP addresses associated with orphans, 4.4%, are known to send spam. Benign uses of orphans include being used as placeholders for records from deleted domains. Due to a mix of uses, we conclude that orphans could not simply be removed but need to be monitored closely.

2. BACKGROUND

Here, we discuss the DNS records important to understanding orphan name servers. We also present a scenario which may lead to the creation of an orphan name server. Finally, we discuss related work.

2.1 DNS Background

The DNS server at each TLD maintains a *zone file* containing two types of records relevant for studying orphan name servers. These are the NS (name server) and A (address) records for each domain falling within that TLD. The NS records map a domain to the names of its DNS servers and A records provide IP addresses of those DNS servers. Figure 1 shows a sample portion of a zone file for .com TLD. Here, two name servers each for domains `fake.com` and `example.com` are listed in the NS records. However, names do not suffice. Clients need the corresponding IP addresses to contact name servers. Since the authoritative source for this IP address is the server itself, A records corresponding to the NS records are also needed in the zone file. These records are referred to as *glue* records. Notice, however, that the A records for `ns1.fake.us` and `ns2.fake.us` are not listed in Figure 1 because they may be obtained from the name servers for the `fake.us` domain, or possibly the name servers for `.us`.

2.2 Orphan Creation

We describe how orphans are created through a scenario. A domain, `orphan.com`, is created, leading to the addition of an NS record for this domain to the zone file: `orphan.com NS ns1.orphan.com`. The host `ns1.orphan.com` is in `orphan.com`, so its address record, `ns1.orphan.com A 1.2.4.4`, is also added to the zone file. Other domains may also use this server, resulting in the addition of NS records for them as well. If `bad.com` uses this server, then `bad.com NS ns1.orphan.com` will be added.

<code>fake.com</code>	NS	<code>ns1.fake.us</code>
<code>fake.com</code>	NS	<code>ns2.fake.us</code>
<code>example.com</code>	NS	<code>ns1.example.com</code>
<code>example.com</code>	NS	<code>ns2.example.com</code>
<code>ns1.example.com</code>	A	1.2.3.4
<code>ns2.example.com</code>	A	1.2.3.5

Figure 1: Sample portion of a TLD zone file

<code>fake.com</code>	NS	<code>ns1.fake.us</code>
<code>fake.com</code>	NS	<code>ns2.fake.us</code>
<code>example.com</code>	NS	<code>ns1.example.com</code>
<code>example.com</code>	NS	<code>ns2.example.com</code>
<code>bad.com</code>	NS	<code>ns1.orphan.com</code>
<code>ns1.example.com</code>	A	1.2.3.4
<code>ns2.example.com</code>	A	1.2.3.5
<code>ns1.orphan.com</code>	A	1.2.4.4

Figure 2: Sample portion of a TLD zone file containing an orphan name server record

Now imagine that some time later `orphan.com` is removed. In this case, the above mentioned NS record for the domain is deleted from the zone. However, what happens of the A record? One option is to delete it. Another option is to let this record live on, to account for cases when other domains in `.com` or other TLDs depend on it, for example, the above mentioned `bad.com`. Deleting the `ns1.orphan.com` A record would harm `bad.com`, perhaps taking it completely offline if this is its only DNS server. Keeping this A record prevents the deletion of this single domain from affecting others which may rely on the same server. This is an example of a common situation that leads to the creation of orphan DNS servers. The sample zone file from Figure 1 after these additions and deletions is shown in Figure 2.

The above scenario is of course not the only way an orphan server may appear. Some may arise out of configuration errors, such a simple typo where the name associated with an A record is not what it should have been and so this incorrect name does not match any existing domain. Others may occur due to lax registrar policies. While glue records are the only reason an A record *should* exist in a TLD zone, mechanisms may not be in place to actually enforce this. Lack of enforcement of such policy would allow such records to intentionally be added directly to the TLD zones.

2.3 Related Work

To our knowledge, there has been only a single other work that has looked into the problem of orphan name servers. This resulted in a recent presentation at an ICANN meeting [17], and was based on data collected by Karmasphere [8] and Internet Identity [7]. Their work begins with lists of malicious domains and checks if any name servers used by them are orphans. They find 3.4% of phishing domains and 59% of fast flux [10] domains use orphans. However, because they start with malicious domains, they are not able to identify orphans used for benign purposes. We believe that in deciding how best to solve the problem of orphan DNS servers, it is necessary to determine if they have legitimate uses. Therefore, we take a different approach, starting

with TLD zone files instead of just with malicious domains, in order to find the general prevalence of orphans.

3. DATA AND METHODOLOGY

In this Section, we describe our methodology for locating orphan name servers and for finding who is using them.

3.1 Finding Orphan DNS Servers

We use three different techniques for discovering orphan name servers in order to be exhaustive where we have the necessary data, and still find some orphans where we have less data available. These techniques are direct analysis of zone files, DNS queries based on zone file information, and DNS queries based on malicious feeds.

3.1.1 Analysis of Zone Files

Analysis of the zone file for a top level domain (TLD) is the ideal way to find all orphans in that TLD. We have the TLD zones for 6 TLDs: `.asia`, `.com`, `.info`, `.mobi`, `.net`, and `.org` [1, 3, 14, 18, 25]. These TLDs accounted for 60% of the 177 million domains on the Internet at the time of our experiments [24]. We obtain these zones through file transfers from the organizations that administer them, so they should be up to date.

To find an orphan in a zone file, we use a simple algorithm. We first examine the entire list of `A` records in the zone file. For each host name listed in an `A` record, we extract the domain name. Since all of the TLDs considered here delegate zones to customers at the second level, directly under the TLD, the domain name in this context is always the TLD plus a single label. We then check if an `NS` record exists for the domain in order to verify if the domain itself exists. If such a record does not exist, then we classify the name from the `A` record as an orphan name server. Returning to Figure 2, when we come across the `A` record `ns1.orphan.com` A 1.2.4.4, we then search the `NS` records in the zone file for any `NS` records for the domain `orphan.com`. Since no such records exist, we classify `ns1.orphan.com` as an orphan.

To locate which domains are using the orphan name server, we search the zone files for `NS` records that point to orphans. In the case of Figure 2, we would find `bad.com`.

3.1.2 DNS Queries from Zone Files

While our TLD data contains the largest TLDs, a notable deficiency is that it does not contain any country code ccTLDs because operators of these TLDs are not willing to share them even for research purposes. Finding orphan name servers exhaustively for these TLDs is not possible. However, we can still find potential orphans through the zone files available to us.

Instead of looking at the `A` records, as we did earlier, we now look at the name servers pointed to by the `NS` records. If the name pointed to is in one of the TLDs we have a zone file for, we ignore it because if it were an orphan, it would have been identified through the previous method. If it is not, as is `ns1.fake.us` in Figure 2, we perform two DNS queries. We first extract the domain name from the name server’s host name (in this case `fake.us`). For this purpose, we use the Public Suffix List [13], taking any suffix on the list, plus one additional label, as the domain name. We then and perform an `NS` query on this domain name. This query checks if the domain name that the server’s host name is contained in exists. Orphans can only be present if the domain name

does not exist. If this query fails, we perform an `A` query on the name server’s host name (`ns1.fake.us`). If this query succeeds, the host name is an orphan. Otherwise, it is just a misconfiguration of a `NS` record pointing to a server that does not exist.

This method directly finds users of the orphans as well since the domain on the `NS` record used to find the orphan is a user of that orphan.

3.1.3 DNS Queries from Malicious Feeds

To get a glimpse of the use of orphans in hosting malicious web sites, including domains not seen by the above two methods, we bring in one more data source. We examine five different live feeds of phishing and malware-hosting sites [2, 4, 9, 15, 16]. We perform `NS` queries to identify their name servers, and then we perform DNS queries to check if these servers are orphans exactly the way we did in Section 3.1.2.

3.1.4 Data Collection Limitations

The above mechanisms allow us to find all orphan name servers contained in or used by domains in the DNS zones we can access. It also allows us to find ones used by domains in our malicious feeds. However, it is still not a complete picture of all orphans. Some orphans likely exist in zones we do not have access to. Some more may be used by domains not present in our feeds or zone files. The effect of this limitation is that we underestimate the prevalence of orphan name servers. Similarly, while all of these mechanisms automatically locate users of orphans, they will not identify any users not contained in the zone files or malicious feeds we have access to. This leads to an underestimation of the users of orphans as well.

3.2 Data Overview

We receive new zone files for each zone on a daily basis, and new malicious feeds every hour. We repeat the process described above each day on the new data. For this analysis we use 31 days of zone files and 14 days of the malicious feeds. Table 1 provides an overview of the data used to locate orphans.

Start date of zone files	2009-04-01
Days of zone files used	31 Days
Domains represented in zone files	106 million
A records in zone files	2.2 million
NS records in zone files	249 million
Out of zone name servers looked up	241,179
Start date of malicious feeds	2009-04-16
Days of malicious feeds used	14 Days
Hosts in malicious feeds	242,752
Domains in malicious feeds	9,554
Name servers for domains in malicious feeds	48,042

Table 1: Overview of data used for orphan detection. Numbers presented are daily averages

We see that as expected, there are over double the number of `NS` records than there are domains, since most domains use multiple name servers. Also as expected, there are far fewer `A` records than domains, since many domains share name servers. A majority of the name server IP addresses were contained in the zone files themselves. Only about

10% of them were out of the TLD zones and required a DNS look-up. The malicious feeds provided fewer servers to look up than the zone files, but are useful for providing more diversity in the potential orphan users.

4. CHARACTERISTICS OF ORPHANS

We now examine the prevalence, distribution, and lifetimes of the orphans name servers we found.

4.1 Prevalence of Orphan DNS Servers

We begin by examining the prevalence of the orphans. Overall, we find a total of 46,369 orphans. Each day, we record an average of 15,962 orphans. Many are present in our data for several days. 39,443 (85%) of orphans are in the TLDs we have zone files for, representing 1.7% of the A records in our TLDs, a small but significant percentage. For these, we know that they are all the orphans in these TLDs.

4.2 Distribution of Orphan DNS Servers

Now, to determine if certain entities on the Internet are responsible for disproportionate numbers of orphan name servers, we examine the distribution of orphans in terms of the domains, TLDs and autonomous systems (ASes) they belong to, as well as their IP addresses.

4.2.1 Distribution in Domains and TLDs

Although we saw 46,369 orphan name servers, they are likely not independent of each other. We can quantify relationship among them by looking at the number of zones that contained orphans. We find orphan servers in 23,153 domains, an average of 2.0 orphan servers per domain that contains them.

Since TLD policies may affect the presence of orphan servers, we examine the number of orphans contained in each TLD for which we have zone files. These results are seen in the top part of Table 2. Somewhat surprisingly, we see that *.info* has a much higher percentage of its A records as orphans than any other TLD. 26,111 (18.8%) of its 139,126 A records were orphans. In fact, *.info* contains four times as many orphans as the much larger *.com* which has an order of magnitude more A records. *.org*, also much smaller than *.com*, has more orphans than it as well. Notably, the two zones run by Verisign, *.com* and *.net*, have by far the least percentage of their A records as orphans as compared to the other TLDs. This indicates that perhaps Verisign makes policy decisions that discourage orphans.

We also checked the TLD distribution of domains containing orphans, to determine whether these high numbers were caused in part by many orphans in a few domains. The TLD distribution of these domains was similar to the distribution of the orphans themselves, including the same ordering of TLDs accounting for the most domains containing orphan name servers.

Owing to lack of zone files, we can not determine what percent of A records in other TLDs are orphans. We can also not determine the true number of orphans in other TLDs. However, even with the information we have, a few other TLDs stand out. Notably, 11 TLDs each have over 1,000 orphans, shown in the bottom part of Table 2, the most being 3,931 in *.us*. Since our methodology under-counts the number of orphans in these TLDs, the actual numbers are likely to be higher. In total, we find 123 TLDs containing orphans, although many of these contain only a few.

TLD	# orphans	A records	%
.info	26,111	139,126	18.8%
.mobi	433	4,062	10.7%
.asia	99	1,313	7.5%
.org	7,715	206,513	3.7%
.com	6,428	1,566,392	0.4%
.net	530	3,31896	0.2%
.us	3,931	n/a	n/a
.cm	2,771	n/a	n/a
.cn	2,715	n/a	n/a
.br	2,495	n/a	n/a
.de	2,413	n/a	n/a
.ws	2,344	n/a	n/a
.kr	2,118	n/a	n/a
.ru	1,801	n/a	n/a
.ca	1,368	n/a	n/a
.in	1,340	n/a	n/a
.jp	1,182	n/a	n/a

Table 2: Orphan name servers by TLD, as compared to the total A records in each TLD zone

4.2.2 IP Address Distribution

As with orphans in the same domain, orphans pointing to the same IP address are also related since these are using the same physical machine. Collectively, we find 14,411 IP addresses in use by orphans, an average of 3.2 orphan names pointing to the same physical machine.

To determine if certain network operators are responsible for disproportionate numbers of orphans, we translate these IP addresses into the corresponding AS numbers using the service offered by Team Cymru [23]. We find that 14,291 of these IP addresses belong to 1,960 ASes, an average of 7.3 IP addresses of orphans per AS that has them. This implies that about 6% of the approximately 30K ASes that originate routes on the Internet contain orphans. The remaining 120 are IP addresses with no corresponding AS, such as private IP addresses. These are in a sense doubly misconfigured, since they are orphans and point to addresses which are not publicly accessible.

Most ASes with orphans only have a few, with 89% having 10 or less, and 43% having just one. In these cases, it is likely that the operators of the AS themselves are not responsible for the presence of these orphans, at least not intentionally. Instead, these are either the result of mistakes, or set up intentionally by some other party. However, we also see four ASes each containing over 100 IP addresses of orphans. In these few cases with many orphans pointing to the same AS, the operators may be responsible for the behavior, but there may be another explanation. These four top ASes belong to commercial web-hosting companies; it is possible that this behavior is instead due to their customers creating such records pointing to the hosted address. However, if this were the case we would expect to see the same behavior in many more ASes.

4.3 Lifetimes of Orphan DNS Servers

Next, we examine the lifetimes of the orphan name servers in our data. Each day we check if orphans name servers recorded on the previous day are still present to examine how long each orphan persists. The results are shown in Figure 3.

We expected orphans to live for long periods of time, ow-




Figure 3: Orphan Lifetimes (CDF)

ing to the negligence in noticing them. In reality, many are shorter-lived. Specifically, 12% of them only live for a single day. These orphans may have been created unintentionally and then removed once somebody noticed the error. Less than 2% lasted for our entire 31-day data window. The median lifetime for an orphan is between 8 and 9 days, suggesting that many orphan name servers might exist for tasks that are completed in a short amount of time, such as hosting malicious web sites which are known to live only for a few days [10].

5. USES OF ORPHANS

In this section, we examine how orphan name servers are being used. We approach this question in a few ways. First, we examine their use in the most expected way, as a DNS server. Next, we examine the services running on each orphan. We also examine domains with a large number of orphans to learn how they are using orphans. Finally, we look specifically for malicious uses.

5.1 Use as DNS Servers

We now examine how many domains are actually using orphans as their DNS servers. In total, we see 212,850 domains using orphans as DNS servers, an average of 4.6 per orphan. Figure 4 shows that *almost 1/3rd of the orphans are actually in use as a name server by at least one domain*. We also find a few orphans being used by a large number of domains (not shown in the figure). Specifically, two orphans are each used by 30K domains and two by 15K. The rest of the orphans, 69%, are not being used as DNS servers by any users in our data, and thus appear to be present as a result of an oversight in keeping zone file contents up-to-date. We investigate if they are in use in some other manner in Section 5.2.

5.2 Running Services

Several host names of the orphan name servers begin with `www`. Because `www` is a common label for hosts that run an `http` server, we check if port 80 is open on each orphan. If so, the orphan may be used as a web server. We also check if port 25 is open to see if they may be running an `smtp` server. We find that 33% of the unique IP addresses of orphans have port 80 open, and thus appear to be running




Figure 4: Number of domains using each orphan name server (CDF)

an `http` server. Additionally, 22% have port 25 open and appear to be running an `smtp` server, implying that they can potentially be used for sending e-mail.

From orphans with an open port 80, we select 1,000 which we do not see used as a name server, in order to determine if those not used as name servers had some other purpose. From these orphans we request a web page. We find that 42.9% of orphans return nothing despite having an open port. While we can not determine what content they are serving, we can not say they are serving nothing, since they may be using cloaking or ignoring our request since it does not have a standard user agent. 27.9% of the orphans we checked displayed the Apache welcome page. Although they may be serving content on a URL somewhere other than the root of their domain, but there is none at the root of the domain. Although some web sites may be set up in this manner, we have no good way of guessing which path may be used in order to check it.

The remaining 29.2% of orphans in our sample had content at the root of their domain. We classify the type of content on these orphans by performing keyword searches on the returned pages and manual inspection. We find that they primarily served advertisements or contained blogs or scam sites. The latter included the well known “Canadian Pharmacy” scam [19] shown in figure 5. The Canadian Pharmacy scam is one instance of a popular type of scam which claims to sell pharmaceuticals online but may be unlicensed, may not ever deliver the ordered products, or may fill orders incorrectly [5]. Such scams were on 1.4% of the orphans we checked.

5.3 Domains with the Most Orphans

Some domains contain a high number of orphans. The ten highest all contain more than 50 and the largest has 1,198. Although we can not determine for sure what all of them are doing, we do see one surprising activity.

A distinct behavior is seen among seven of the top ten domains containing the most orphans. The domain names of these tend to contain substrings like “delete” or “old”. Such domains appear to contain name server records copied from domains that have or will soon expire. For example, we may see an orphan named `ns1.example.com.deleted.info`.



Figure 5: Screenshot of Canadian Pharmacy scam web site

Usually in this case, in addition to `deleted.info` not existing, the other domain implied in the host name, `example.com`, does not exist either. However, if we send this server DNS queries related to `example.com`, it responds as if the domain existed, and as if it was a name server for the domain. Checking IP addresses, if we ask this server for the IP address of `ns1.example.com` we get an address matching that of `ns1.example.com.deleted.info`.

Based on this behavior, we conjecture that registrars move name servers from expired domains to another domain as orphans to minimize disruption to other domains they administer. Other domains that use a server from the expired domain for name resolution can then have their NS record changed to point to the new orphan, perhaps automatically by the registrar, minimizing disruption to the service of these domains. This practice may help keep customers happy who own multiple domains using the same infrastructure, ensuring that the rest of their domains will still work after they cancel one. Since these are orphan name servers, they have the same resolution behavior as before the original domain expired. If they were not orphans, then an additional query would be needed to look up the authoritative name server for the domain to which the name server was moved.

5.4 Malicious Uses of Orphans

Given that orphans are being used to host scam sites, as we saw in Section 5.2, and the use of orphans by malicious sites found in [17], the connection between orphans and malicious activities deserves investigation. Orphan servers may be attractive to miscreants for a few reasons. For example, if they are using them for purposes other than DNS servers, having records for them on the TLD DNS server means the miscreants may not even have to run their own DNS server. Orphans related to malicious activity may also naturally arise from enforcement efforts by ISPs removing NS records of domains found to be malicious. The DNS server may already have been in use by other domains belonging to the same miscreant, and may continue to be used as such if the A record is left.

We look for malicious users of our orphan name servers in multiple sources. We use the Google Safe Browsing phishing and malware blacklists [6], three feeds of malware hosting sites [4, 9, 15], two of phishing sites [2, 16], and a feed of

scam sites [22]. Results in this section are based on a longer data window than our general results, 40 days instead of 31.

Table 3 shows how many domains using orphans as their DNS server are involved in malicious activities. There are 212,850 total domains using orphans in our data. Of these users, 2,250 (1.1%), using a total of 86 orphan name servers, are in our phishing feeds. Similarly, 2,838 (1.3%), using a total of 53 orphan name servers are in our malware feeds. Other sources of malicious data also have orphans associated with them, as seen in the table. A small fraction of domains containing orphans appear in blacklists as well, the most significant being 76 appearing in the Google malware blacklist. We conclude that although miscreants are using orphans, orphans themselves are not often blacklisted.

Data Source	% of Orphan Users	# Domains
Google malware	0.4%	897
Google phishing	0.0%	4
Phishing Feeds	1.1%	2,250
Malware Feeds	1.3%	2,838
Scam Feed	0.0%	22

Table 3: Domains using orphans appearing in sources of malicious data

We additionally look for the IP addresses of orphans in the Spamhaus SBL blacklist [20] which tracks IP addresses associated with sending spam, and the Spamhaus XBL blacklist [21] which tracks IP addresses with types of exploits commonly used by spammers. 4% of IP addresses of orphans appear in the SBL and 0.4% in the XBL. When we look only at the 500 IP addresses hosting the most orphans, the SBL results are significantly different, 13.4% are blacklisted. We conclude that a significant fraction of orphans send spam.

6. CONCLUSION

Our study identifies three reasons for the existence of orphan name servers. These are typographical errors and misconfiguration in TLD zone files, malicious activity, and as placeholders for deleted domains. We find that a significant number of the orphans are being used to send spam, and a smaller number of orphans are used for other malicious activity including hosting scam sites and serving as the DNS servers for phishing and malware sites. Although these malicious uses are not widespread enough that use of orphans can be taken as an indicator of maliciousness on its own, they can certainly contribute to the determination when combined with other factors.

Many orphan name servers are used in benign roles. Because of this, disallowing them from being created, and removing all currently in existence, may do more harm than good. However, since a large percentage of orphan name servers were not serving any domains from the largest TLDs, we hypothesize that many orphans could be removed without any negative effects.

Finally, the disproportionate prevalence of orphan name servers in some TLDs and ASes indicates that perhaps the benign roles they serve are not necessary, or can be accomplished otherwise. Owing to this, we recommend that orphan DNS servers be better scrutinized to prevent misuse and to ensure that they really are providing a benefit.

7. REFERENCES

- [1] Aflias Limited. How can I get access to Aflias' TLD zone file for .INFO domains? <http://www.info.info/faq/how-can-i-get-access-aflias-tld-zone-file-info-domains>.
- [2] APWG. Anti-phishing working group. <http://www.antiphishing.org/>.
- [3] DotAsia Organization Limited. .ASIA Zone File Access Agreement. <http://www.dotasia.org/info/DAO.ZONE-2007-10-24.pdf>.
- [4] eSoft Inc. <http://www.esoft.com/>.
- [5] Federal Trade Commission. Your health online. <http://ftc.gov/bcp/edu/pubs/consumer/health/hea12.shtm>.
- [6] Google Code Labs. Google safe browsing API. <http://code.google.com/apis/safebrowsing/>.
- [7] Internet Identity. <http://www.internetidentity.com>.
- [8] Karmasphere, Inc. <http://karmasphere.com/>.
- [9] Malware Patrol. Malware patrol - malware block list. <http://www.malwarepatrol.net/lists.shtml>.
- [10] D. K. McGrath, A. Kalafut, and M. Gupta. Phishing infrastructure fluxes all the way. *IEEE Security and Privacy Magazine Special Issue on Securing DNS*, September/October 2009.
- [11] P. Mockapetris. Domain names - concepts and facilities. IETF RFC 1034, Nov. 1987.
- [12] P. Mockapetris. Domain names - implementation and specification. IETF RFC 1035, Nov. 1987.
- [13] Mozilla Foundation. Public suffix list. <http://publicsuffix.org>.
- [14] mTLD, Ltd. dotMobi Zone File Access Agreement. <http://mtld.mobi/domain/zonefile>.
- [15] NETpilot GmbH. Viruswatch mailing list. <http://lists.clean-mx.com/cgi-bin/mailman/listinfo/viruswatch>.
- [16] OpenDNS. PhishTank. <http://www.phishtank.com/>.
- [17] D. Piscitello. Orphaned name servers. ICANN 35 Presentation, June 2009.
- [18] Public Interest Registry. .ORG Registry - ZOne File Access. <http://pir.org/index.php?db=content/Website&tbl=Registrars&id=7>.
- [19] Spamhaus Project. Register or known spam operations - canadian pharmacy. http://www.spamhaus.org/rokso/evidence.lasso?rokso_id=R0K8138.
- [20] Spamhaus Project. SBL. <http://www.spamhaus.org/sbl/index.lasso>.
- [21] Spamhaus Project. XBL. <http://www.spamhaus.org/xbl/index.lasso>.
- [22] Support Intelligence, LLC. <http://www.support-intelligence.com/>.
- [23] Team Cymru, Inc. IP to ASN mapping. <http://www.team-cymru.org/Services/ip-to-asn.html>.
- [24] VeriSign. Domain name industry brief, Feb. 2009. <http://www.verisign.com/static/044518.pdf>.
- [25] VeriSign, Inc. TLD Zone Access Program. http://www.versign.com/information-services/naming-services/page_001052.html.