



Passive Detection of Misbehaving Name Servers

Leigh B. Metcalf, Jonathan M. Spring
CERT Network Situational Awareness Group
netsa-contact@cert.org
Publication NetSA-2012-01

January 2012

Executive Summary

In this paper we demonstrate that there are name servers that exhibit IP address flux, a behavior that falls outside the prescribed parameters. We demonstrate this flux in two types of data: passively collected DNS messages and the contents of several large, top-level domains' official zone files. The community of name server operators has previously indicated that there is no benign use case for such behavior and has attempted to quash it. The continued existence of such behavior is an indicator of malicious name server activity and the inadequacy of attempts to control it.



This work was created with the funding and support of the Department of Homeland Security under the Federal Government Contract Number FA8721-05-C-0003 between the U.S. Department of Defense and Carnegie Mellon University for the operation of the Software Engineering Institute, a Federally Funded Research and Development Center.

CERT® is a registered mark owned by Carnegie Mellon University.

NO WARRANTY THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Copyright 2012 Carnegie Mellon University.

Notice: Pursuant to Contract Number FA8721-05-C-0003 the Government is not authorized to publish, or allow others to publish data or software first published for sale by CARNEGIE MELLON UNIVERSITY but retains unlimited rights to use it for its own purposes.

Contents

1	Introduction	4
1.1	Related work	4
1.2	Motivation	5
1.3	Data Sources	5
2	Method	6
3	Results	7
4	Discussion	11
5	Future Work	12

List of Tables

1	<i>Unique active name servers per day and total, June 12 - July 9 2011 .</i>	8
2	<i>Number of name servers that change IP address and ASN a given number of times in 28 days in zone files (of 2,796,128 name servers)</i>	8
3	<i>Number of name servers that change IP address and ASN a given number of times in 28 days in SIE records (of 1,926,367 name servers)</i>	9
4	<i>Distribution of TTLs in NS record sets (760,796,769 total record sets)</i>	14

List of Figures

1	<i>The average and minimum number of times an IP address change of a name server also changed the ASN, normalized by the number of IP changes. Each point is binned by IP change frequency for each data source. The maximum for both sources is 1 for all bins.</i>	10
2	<i>Log-log plot of observed NS record TTL value frequency out of total NS records observed. There are 760,796,769 total records represented. The figure does not show an outlier: 0.035% of the records had a TTL value greater than 604,800 and less than the maximum value of 4,294,944,960.</i>	11

1 Introduction

Detecting malicious domains is becoming an important task in limiting the impact of ne'er-do-wells on the internet. It has become a race between defenders developing new detection methods and adversaries developing new evasion methods. There are only a few properties a domain is required to have, so malicious domains can leave few traces. But all domains must have a name server. Focusing on the characteristics of a domain that are necessary for its operation limits the ability of malicious domain controllers to avoid countermeasures implemented by security personnel. The domain name system (DNS) requires only two associations for a domain: its location and whom to ask about its location. Network administrators and security personnel have pursued restricting location, mainly via IP (internet protocol) addresses blocking, with some success. Blocking name servers, the entities that provide location information about domains, has been pursued with much less energy than IP blocking.

Distinguishing between suspicious and benign name servers is one step in the process of categorizing malicious domains. This step can allow the name servers themselves to be acted against, given sufficient evidence. Because name servers are one of the primary components of a well-functioning domain name system, they do not change IP addresses frequently. Domains that change IP addresses quickly or often are said to exhibit IP flux. IP flux makes services more resilient because it circumvents the defender's attempts to block the IP addresses. Web pages that deliver malicious content can use IP flux, for example. Name servers exist on domains, and those domains may also exhibit flux to evade attempts to block them.

1.1 Related work

Significant work has investigated how to find and block domains that demonstrate malicious activity. This work assumes that malicious domains will somehow behave differently than benign domains. Some of these efforts are operational, and some are still in the research phase. Most specialize in a particular brand of malicious activity. Some operational blocking lists conceal their selection criteria to prevent adversaries from exploiting the rules. Some notable operational lists include Spamhaus, PhishTank, and Google Safe Browsing. Spamhaus maintains a few operational lists, each targeting aspects of malicious email (Spamhaus, 2011). PhishTank publishes lists of phishing URLs that are consumed by popular browsers to block phishing pages from reaching users. PhishTank takes a community-based approach, providing a site where "anyone can submit, verify, track and share phishing data" (PhishTank, 2011). Google also derives lists of phishing and malicious sites while it crawls the web, and it makes these lists accessible to the public using an API (Google, 2011). While the effectiveness of blocking phishing seems to vary (Rasmussen and Aaron, 2010, 2011; Spring, 2010), efforts to take down or block phishing sites have been demonstrated to shorten their lifetime (Moore and Clayton, 2007).

Several papers have described blocking techniques using passive DNS; some

techniques also utilize zone files, the official lists of domain-to-name-server and name-server-to-IP-address mappings maintained by the registry. Antonakakis, Perdisci, Dagon, Lee, and Feamster developed a reputation-based classification system, called Notos, that utilizes passive DNS (pDNS) monitoring data (Antonakakis et al., 2010). Its classification scheme divides its features into the broad categories of network-based, zone-based, and evidence-based. Bilge, Kirda, Kruegel, and Balduzzi designed the EXPOSURE system, which also uses pDNS as the data input. EXPOSURE introduces features based on time series and time to live (TTL) (Bilge et al., 2011). Preceding these efforts, Felegyhazi, Kreibich, and Paxson used zone files to predict which domains would be used maliciously. This prediction is based on the previous evidence of malicious activity by other domains using the same name server (Felegyhazi et al., 2010). Stoner finds malicious activity, specifically malicious fast flux and domain flux, using simpler methods. That research uses only two features of a domain: the IP addresses it maps to and the associated autonomous system numbers (ASNs) in which the IP addresses reside (Stoner, 2010).

1.2 Motivation

The ultimate goal of the current work is to hinder criminals' free use of domain names as an accessory to their crimes. Other than Felegyhazi et al. (2010), current techniques can only reactively hinder criminals. All the operational lists currently in use (see 1.1) are reactive, so they can at best take away names only after some damage has been done. This is important in limiting damage, but it is not ideal. Detection via name server behavior can improve the current state-of-the-art deterrence because it preempts at least some domains and is comprehensive in that all domains must be served by a name server. The particular aspect of name server behavior we measure is name server IP flux.

1.3 Data Sources

The general categories of data sources used in this work are TLD zone files and pDNS traffic. The contents of zone files are generally reported to the registry by the registrars. New generic TLD (gTLD) operators are required to make this file available under certain conditions (ICANN, 2011). Because gTLD files are more readily available than other TLDs, we demonstrate our analysis using the *com*, *org*, *net*, *biz*, *info*, and *mobi* zone files. Zone files have a required form, so this analysis is applicable to any other TLD zone file (Mockapetris, 1987). Not having visibility into any country code TLDs (ccTLD) does bias the data set. However, *com* is by far the largest TLD and represents a large percentage of the total domains for the internet (Stoner, 2010). By Verisign's terms of use, one zone file download is permitted every 24 hours, which limits the granularity of the detectable changes in this analysis.

Passive DNS collection was first described in Weimer (2005). We use a large pDNS source, the Security Information Exchange (SIE), for pDNS data. While the coverage of the SIE sensor array is incomplete and biased, there is evidence

that it is wide, and it processes many tens of millions of distinct messages per day (Spring et al., 2011). The SIE is the best generally available source of pDNS data. The SIE data is delivered in resource records sets (RR sets). One record set is all the resource records in a single message that share the record name, class, type, and TTL, with the record data sorted by the standard DNS ordering and stored as the final fields of the record set.

Routers use Border Gateway Protocol (BGP) information to find paths to IP addresses throughout the internet by way of associated ASNs. A digest of this data can associate observed IP addresses with ASNs. This is valuable because ASNs represent blocks of effective control and help distinguish organizational boundaries. The current analysis uses the University of Oregon Route Views Project to associate IP space with ASNs (Route-Views, 2012).

2 Method

Various characteristics of name server behavior were collected. All calculations are on data collected during the four weeks from June 12, 2011 through July 9, 2011.

Active Name Servers. The number of unique active name servers per day from both data sources is calculated. This is a straightforward counting operation on each day of the data, as well as a count of unique entries for the total duration of the observation period.

Zone File Name Server IP Changes. The zone files (the zones analyzed are named in section 1.3) contain the name server's IP address. Each day's file is processed to make a list of unique name server-IP pairs. The name servers present on the first day of the observation period are checked for movement throughout the observation period. The unique addresses for these name servers are collected from the data. The number of changes, both per server and by all servers, is then analyzed.

Changes in Name Server A Records. The IP address of a name server is reported through the DNS just like other IP addresses. A name server may be the answer to a name server query, but it is also the subject of A records. Using the SIE pDNS data, changes in these A records can be detected. First, a list of name servers for a day is calculated from the payload of all the unique name server RR sets. The A record sets for these names are then extracted from the data. This process yields a list of name-server-to-IP mappings for each day similar to that derived from the zone file as described above. The list of names for the first day in the observation period is then compared to the list for each other day to determine what name servers have changed IP addresses. Because this analysis consumes live data, a particular name server might not be requested every day, unlike the zone data, which is static data available on request. However, changes will by definition not be cached in the DNS

and so we should not miss any relevant data points, which for this analysis is only a change in the value of the name server's location.

Distribution of TTLs in NS records. The distribution is calculated using the SIE pDNS data. All unique RR sets for each day in the measurement period are stored. See 1.3 for a description of the SIE RR set format. The calculation operates on the unique name-server-type (NS) RR sets observed per day. The TTL value is extracted from each unique NS record set, and the instances of each value are counted. We provide some statistical evaluation of this data.

Name Server IP Flux. There are several previously documented methods of detecting IP flux in passive DNS traffic using A records (Passerini et al., 2008; Stoner, 2010). It does not seem that these methods have been applied to zone file behavior. Given a name server N , the following values are calculated to determine if its IP addresses exhibit flux:

- The number of unique IP address values N uses during the observation period.
- The number of ASNs to which those IP addresses belong. This is derived from the BGP mapping described in 1.3.

The IP addresses of name servers should have a different variability than those of domain names, so the threshold parameters for name server flux are different than those for traditional IP flux. The value of these parameters is derived from the contextual data and an understanding of the common practices of name server administration. Because the server needs to be well known for reliable zone operation, it should not change IP addresses frequently. We are limited to a resolution of one measurement per day due to the frequency we are permitted to download our zone file data.

3 Results

Table 1 on the following page displays statistics for the scale of the study. The number of name servers on the first day is relevant because those name servers were the ones checked for movement throughout the observation period. The ranges exclude outliers, whose data was not properly collected due to technical errors. The dates are June 13 in the SIE data and June 16, 20, and July 4 in the zone files. The total number of distinct names is the count of unique names observed over the 28-day observation period.

Name servers were observed to change their IP address. In the zone files, 61,801 name servers changed IP address at least once. Of these, 41,796 changes included at least one change in the ASN in which the IP address was located. Table 2 on the next page breaks down these results into counts of the observed occurrences of each number of changes. The maximum number of changes observed, in both IP and ASN, is 24 times in 24 measurements. This is a lower

	SIE	Zone files
First day (June 12)	1,926,367	2,796,128
Average # of NS	1,968,271	2,786,279
Range of # of NS	$1.8 * 10^6$ to $2.1 * 10^6$	$2.4 * 10^6$ to $2.8 * 10^6$
Total distinct observed	4,021,151	3,260,648

Table 1: *Unique active name servers per day and total, June 12 - July 9 2011*

ZONE FILE DATA				
IP Changes	NS changes IP	% of total	NS changes ASN	% of total
0	2734327	97.8%	2754332	98.5%
1	52741	1.9%	36645	1.3%
2	4855	0.2%	1846	0.1%
3	551	0.0197%	635	0.0227%
4	198	0.0071%	838	0.0300%
5	233	0.0083%	531	0.0190%
6	482	0.0172%	500	0.0179%
7	660	0.0236%	401	0.0143%
8	706	0.0252%	224	0.0080%
9	607	0.0217%	30	0.0011%
10	478	0.0171%	19	0.0007%
11	138	0.0049%	9	0.0003%
12	35	0.0013%	14	0.0005%
13	16	0.0006%	20	0.0007%
14	11	0.0004%	8	0.0003%
15	9	0.0003%	16	0.0006%
16	6	0.0002%	3	0.0001%
17	4	0.0001%	5	0.0002%
18	5	0.0002%	4	0.0001%
19	5	0.0002%	5	0.0002%
20+	61	0.0022%	43	0.0015%

Table 2: *Number of name servers that change IP address and ASN a given number of times in 28 days in zone files (of 2,796,128 name servers)*

bound; because observation frequency is limited to once every 24 hours for zone files, some of these name servers may have changed more frequently.

Table 3 on the following page displays data equivalent to that in Table 2, except it is from a different data source. However, these data are not simply comparable. The SIE data contains a different sampling because it can include all top-level domains and name servers for domains that are not just second-level domains. The SIE data also has a finer temporal resolution, so its domains may exhibit more changes per day. The most variable name server changed its IP address 82 times and its ASN 71 times during the 28-day measurement period. Otherwise, the values are calculated with the same methodology.

SIE PASSIVE DNS DATA

Changes	NS changes IP	Ratio	NS changes ASN	Ratio
0	1846152	95.8%	1877654	97.5%
1	68401	2.4%	40422	1.4%
2	5134	0.2%	3276	0.1%
3	1420	0.0508%	1232	0.0441%
4	1177	0.0421%	966	0.0345%
5	1123	0.0402%	684	0.0245%
6	566	0.0202%	450	0.0161%
7	535	0.0191%	388	0.0139%
8	439	0.0157%	279	0.0100%
9	322	0.0115%	220	0.0079%
10	248	0.0089%	152	0.0054%
11	140	0.0050%	76	0.0027%
12	75	0.0027%	46	0.0016%
13	47	0.0017%	35	0.0013%
14	20	0.0007%	37	0.0013%
15	23	0.0008%	30	0.0011%
16	33	0.0012%	37	0.0013%
17	31	0.0011%	31	0.0011%
18	34	0.0012%	30	0.0011%
19	23	0.0008%	19	0.0007%
20+	424	0.0152%	303	0.0108%

Table 3: *Number of name servers that change IP address and ASN a given number of times in 28 days in SIE records (of 1,926,367 name servers)*

Tables 2 and 3 represent two related but different measurements. Each IP address change may or may not change the ASN the name server is in. ASNs are areas of logical control over internet routing, so a change in the ASN generally means the resource is changing areas of control. Tables 2 and 3 do not specify a relationship between IP address changes and ASN changes. They simply report the number of changes of each frequency. Figure 1 on page 10 indicates how frequently an ASN change was caused by an IP address change, with a maximum of each IP address causing one change in ASN. The figure normalizes the average and minimum number of ASN changes as a percentage of the maximum possible. This relationship explains why, in some cases in Table 2, there are more name servers that change ASN a given number of times than change IP that many times. Some subset of the servers that changed IP five times, for example, would have changed ASN three or four times.




Figure 1: The average and minimum number of times an IP address change of a name server also changed the ASN, normalized by the number of IP changes. Each point is binned by IP change frequency for each data source. The maximum for both sources is 1 for all bins.

Table 4 on page 14 presents the results for the distribution of TTLs in NS record sets over the four weeks of observation. They demonstrate a human predilection for round numbers: one hour, one day, and two days are inordinately common. This can make the results difficult to display. Table 4 displays the number of record sets less than or equal to the key value but greater than the previous key value. A key is chosen when the sum of records it will represent surpasses 0.1 percent of the total. If a single TTL key exceeds that value, the key before it is forced to be a key and the popular TTL stands by itself. Figure 2 summarizes these values and displays the broader trends. The log-log plot makes several groupings obvious: the very popular, the 1-5 percent, the 1-0.1 percent, and the rest.




Figure 2: *Log-log plot of observed NS record TTL value frequency out of total NS records observed. There are 760,796,769 total records represented. The figure does not show an outlier: 0.035% of the records had a TTL value greater than 604,800 and less than the maximum value of 4,294,944,960.*

4 Discussion

The first notable facet of the collected data is that name server flux is certainly alive and well. More than 41,000 name servers changed IP address and the ASN in which the name server is hosted. Because name server flux does not have any benign use cases, this is a problem.

The name server record TTL distribution is not surprising. People will naturally gravitate toward common values: two days, one day, one hour. These three values account for over 82 percent of the observed values. Variation from this pattern is not evidence enough to imply anything in particular, but coupling variation with some of the other measurements may allow inference of stronger implications.

The distribution of TTL values in name server records is also significantly different from that in records generally. The most common TTL in all records is much lower (Huth and Spring, 2012). Our research confirms that unsurprising notion. Because TTL affects caching, and caching affects the number of messages that are sent and detected, this difference matters to analysis techniques.

At some point, malicious name servers need to be removed and domains that use them should be blocked. The community must determine the threshold of the acceptable amount of damage before action is taken. The current rate of fluxing name servers is 41,000 out of 2,790,000, or 1.5 percent, in the zone files and 2.5 percent in SIE resolutions. The zone files' results indicate that a noticeable percentage of name servers are suspicious. This certainly seems to be a large enough value to warrant serious intervention.

5 Future Work

Evaluation of the precise extent to which known-malicious and benign domains use name servers that exhibit flux would be valuable information in determining the precise value and cost of blocking fluxing name servers. A complete evaluation of such behavior is not possible. However, estimations using existing lists of malicious domains should provide the necessary information.

Even though malicious actors must have a name server, they need not operate a second-level domain. Free DNS providers exist that give away domain names below domains they own, a service often called “dynamic DNS.” We cannot expect to obtain the zone files from these free providers. A future analysis should quantify to what extent domains using these services are disproportionately malicious. Our cursory investigations suggest that the extent is large. If so, automatic identification of these dynamic DNS services would be a useful area for future work. The list of known-malicious domains against which we test our methods should be extensive and should include multiple sources of known-malicious domains because no one list is comprehensive.

It would be beneficial to expand this study beyond the generic TLDs it investigated. If and how country-code TLDs would differ from gTLDs is not known, nor are the differences between gTLDs and dynamic DNS domains.

References

- Antonakakis, M., Perdisci, R., Dagon, D., Lee, W., and Feamster, N. (2010). Building a dynamic reputation system for DNS. In *19th Usenix Security Symposium*.
- Bilge, L., Kirda, E., Kruegel, C., and Balduzzi, M. (2011). EXPOSURE: Finding malicious domains using passive DNS analysis. *Proceedings of the Annual Network and Distributed System Security (NDSS)*.

- Felegyhazi, M., Kreibich, C., and Paxson, V. (2010). On the potential of proactive domain blacklisting. In *Proceedings of the 3rd USENIX conference on Large-scale exploits and emergent threats*. USENIX Association.
- Google (2011). Google safebrowsing FAQ, http://code.google.com/apis/safebrowsing/safebrowsing_faq.html.
- Huth, C. and Spring, J. (2012). The impact of passive DNS collection on end-user privacy. In *Securing and Trusting Internet Names 2012*. National Physical Laboratory, UK.
- ICANN (2011). *New gTLD Agreement Specifications*. <http://www.icann.org/en/topics/new-gtlds/agreement-specs-clean-19sep11-en.pdf>. Specification 4, section 2.
- Mockapetris, P. (1987). RFC 1035 - domain names - implementation and specification.
- Moore, T. and Clayton, R. (2007). Examining the impact of website take-down on phishing. In *Proceedings of the Anti-Phishing Working Group's 2nd Annual eCrime Researchers Summit*, pages 1–13. ACM.
- Passerini, E., Paleari, R., Martignoni, L., and Bruschi, D. (2008). Fluxor: detecting and monitoring fast-flux service networks. *Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 186–206.
- PhishTank (2011). FAQ, <http://www.phishtank.com/faq.php>.
- Rasmussen, R. and Aaron, G. (2010). Global phishing survey: trends and domain name use in 2H2009. Technical report, Anti-Phishing Working Group.
- Rasmussen, R. and Aaron, G. (2011). Global phishing survey: trends and domain name use in 2H2010. Technical report, Anti-Phishing Working Group.
- Route-Views (2012). University of oregon route views project, <http://www.routeviews.org>. <http://www.routeviews.org>.
- Spamhaus (2011). About, <http://www.spamhaus.org/organization/index.lasso>.
- Spring, J. (2010). Large scale DNS traffic analysis of malicious Internet activity with a focus on evaluating the response time of blocking phishing sites. Master's thesis, University of Pittsburgh.
- Spring, J., Metcalf, L., and Stoner, E. (2011). Correlating domain registrations and DNS first activity in general and for malware. In *Securing and Trusting Internet Names 2011*.
- Stoner, E. (2010). DNS footprint of malware. In *2010 OARC Workshop 2*, Denver, CO.
- Weimer, F. (2005). Passive DNS replication. In *17th Annual FIRST Conference on Computer Security Incident Handling*.

TTL value	% of NS RR sets \leq	TTL value	% of NS RR sets \leq
0	0%	28799	0.022%
30	0.110%	28800	0.413%
59	0.004%	36000	0.103%
60	0.156%	38399	0.020%
120	0.112%	38400	0.393%
299	0.042%	43199	0.046%
300	0.937%	43200	1.518%
599	0.096%	54973	0.100%
600	0.811%	66709	0.100%
899	0.033%	75901	0.100%
900	0.626%	84184	0.100%
1199	0.025%	86072	0.100%
1200	0.134%	86399	0.035%
1799	0.034%	86400	28.721%
1800	0.431%	89999	0.031%
3599	0.093%	90000	0.103%
3600	12.468%	116764	0.100%
7199	0.088%	146302	0.100%
7200	2.874%	171528	0.100%
10799	0.082%	172799	0.019%
10800	1.746%	172800	41.231%
14399	0.038%	259199	0.022%
14400	2.579%	259200	0.573%
17999	0.066%	345599	0.019%
18000	0.140%	345600	0.833%
21599	0.028%	604799	0.082%
21600	0.857%	604800	0.248%
25920	0.123%	4294944960	0.035%

Table 4: *Distribution of TTLs in NS record sets (760, 796, 769 total record sets)*