Understanding Implications of DNS Zone Provisioning

Andrew J. Kalafut
akalafut@cs.indiana.edu

Craig A. Shue
cshue@cs.indiana.edu

Minaxi Gupta
minaxi@cs.indiana.edu

Computer Science Department
Indiana University
Bloomington, IN

ABSTRACT

DNS is a critical component of the Internet. This paper takes a comprehensive look at the provisioning of Internet domains and its impact on the availability of various services. To gather data, we sweep 60% of the Internet's domains for *zone transfers*. 6.6% of them allow us to transfer their complete information. We find that carelessness in handling DNS records can lead to reduced availability of name servers, email, and Web servers. It also undermines anti-spam efforts and the efforts to shut down phishing sites or to contain malware infections.

Categories and Subject Descriptors

C.2.2 [**Network protocols**]: Applications—*DNS*

General Terms

Measurement

Keywords

Domain Name System, DNS, Zone transfer

1. INTRODUCTION

The primary role of Domain Name System (DNS) is to map human-friendly domain names to IP addresses. Over time, this role has expanded to serve as an Internet-wide distributed database, providing support for applications ranging from simple mail delivery to advanced applications, such as spam filtering, voice over IP (VoIP), and other multimedia services. A typical unit of administration in DNS is a second-level domain name, such as `example.com`. A *zone* file corresponding to it stores information about the hosts, services, and sub-domains contained in that zone. While typical DNS queries inquire about a single host or service, some use-cases require complete information contained in a DNS zone. An instance of this occurs when DNS servers for a domain need to synchronize with each other in their view of the zone. The DNS provides a special query for that, called the *zone transfer* query. In this work, we leverage the zone transfer query to capture detailed information about DNS zones in the Internet. During a three month period, we swept 60% of the Internet for zone transfers. In order to increase our data beyond those zones allowing zone transfer, we *walked* the zones of the second-level domains known to deploy DNSSEC [2] (DNS Security Extensions). This is a slow process since it involves making a large number of queries, but its net effect is the same as a zone transfer. Using the two data sets, we examined the DNS zones in our two data sets. The key findings of our study are the following:

1. 6.6% of the second-level domain names allowed us to perform a zone transfer of their zones in spite of the well-known fact that the zone transfers are a security risk [7]. Surprisingly, this included a large percentage of DNSSEC-deploying zones, which may be expected to disallow them.

2. While 88-97% of the zones have at least the prescribed two DNS servers, 82% of them have all their DNS servers in the same AS, 61% have them all in the same IP prefix, and 91% have them all in the same second level domain in the first data set. This may diminish name server availability. The corresponding numbers were an order of magnitude better for the zones deploying DNSSEC, indicating that the zones deploying DNSSEC are more careful in ensuring the availability of their name servers.

3. 0.5-12% of the zones are likely using the same DNS server for internal and external clients, when it is recommended to have them separate for security reasons [7].

4. A small fraction of zones deploying anti-spam technology make mistakes in configuring the relevant records. Fortunately, the email programs at the recipients can be enhanced to account for these mistakes.

5. We find several errors in record contents of a small fraction of zones which may affect the availability of DNS servers, mail servers, and other hosts within a zone. Zones deploying DNSSEC have these errors in fewer numbers.

The rest of this paper is organized as follows. Section 2 provides background on DNS zones. Sections 3 and 4 detail

IMC'08, October 20–22, 2008, Vouliagmeni, Greece.
Copyright 2008 ACM 978-1-60558-334-1/08/10 ...$5.00.

the data collection methodology and present an overview of the data we collected. An analysis of DNS zones is presented in Section 5. Finally, the related work is outlined in Section 6 and Section 7 concludes the paper.

2. BACKGROUND

The behavior of the DNS is specified in a series of Internet Engineering Task Force (IETF) Request for Comments (RFC) documents, dating back to the 1980s. While there are many DNS-related RFCs,the key RFCs describing the basics are RFC 1034 and 1035 [11, 12].

The DNS is organized as a tree, with branches at each level separated by a ".". The entire DNS space is divided into various *zones*. Each zone consists of a connected portion of this tree under the same administrative control. A typical unit of administration in DNS is a second-level domain name, such as `example.com`. A *zone* file corresponding to this second-level domain name stores information about the hosts, services, and sub-domains contained in that zone.

The data within each zone is stored in the form of *resource records* which consists of four basic parts: a *name*, a *class*, a *type*, and *data*. All DNS records relating to the Internet are in `IN` class. 59 different types of records exist for storing various types of data. A zone is defined by two types of records. The first, `SOA` (Start of Authority), indicates the start of a DNS zone. Each zone should have a `SOA` record. The contents of the `SOA` record are the email of an administrator, the domain name of the primary name server, and various timers. The second, one or more `NS` (Name Server) records, also should exist in each zone. These records indicate the set of name servers for the zone and can also indicate the delegation of sub-zones.

3. DATA COLLECTION METHODOLOGY

We use two data sets in this paper. The first, `zone_transfer`, was obtained by attempting to transfer the zones listed under `.com` and `.net`. There were 65,101,733 second-level domains under `.com` and 9,224,482 under `.net`. Combined, these 74,326,215 domains represented about 58% of the 128 million zones registered at the time [19]. For each zone, we had the list of name servers. We looked up the IP addresses corresponding to each of these name servers in order to be able to contact them. We used our own custom software, written using the `Net::DNS` Perl library [9], to zone transfer each of these DNS zones in random order. This process took three months. We attempted a zone transfer from each name server for a zone until we either successfully transferred the zone, or the zone transfer failed for all its name servers. Additionally, if two zone transfers from the same IP address failed, or upon request from the DNS server's administrator, we discontinued making further attempts to transfer any zone from that IP address. Upon connection establishment failure, we retried once. In order to process the records in a timely manner, we used five machines, each with one hundred processes issuing zone transfer requests. We succeeded in transferring zones for 4,947,993 (6.6%), indicating that many DNS servers willingly distribute their information to outsiders. While our data set was confined to the `.com` and `.net` TLDs, it still contained geographically distributed sites.

Our second data set, `dnssec`, is from zones which may be thought of as more security focused, the ones that deploy DNSSEC [2]. DNSSEC adds security to the DNS. It is a set of extensions to the DNS which provide origin authentication and integrity to DNS data, and authenticated denial of existence. We obtained the `dnssec` data set through walking DNSSEC records. This process is slow but allows retrieval of all the records in a zone, much like a zone transfer does. This data set is limited by the low deployment of DNSSEC. To build this data set, we began with a list of 862 zones with DNSSEC in production usage from the SecSpider DNSSEC Monitoring Project [13]. We limited this to the second level zones within the `.com` and `.net` TLDs to allow a fair comparison with the zones we transferred data from in the same TLDs. This yielded a total of 124 zones. Surprisingly, we also found 161 zones deploying DNSSEC in our zone transfer data. Since 96 of the zones listed under SecSpider already existed in our zone transfer data, we only had to obtain data from the rest of the 28 zones that did not allowed us a zone transfer. (We excluded those 96 zones from the first data set.) To obtain data from the 28 new zones in the SecSpider data, we used the DNSSEC Walker tool [8]. This tool relies on the presence of `NSEC` (NextSECure) or `NXT` (NeXT) records which should be present in zones deploying DNSSEC. These records provide a way to discover all of the records from within a zone without using zone transfer. Of the 28 zones we attempted to walk, 4 were only partially walkable due to missing some `NSEC` or `NXT` records. The remaining 24 were completely walkable allowing us to get the same information as we would though zone transfer without actually using the zone transfer query. Our final `dnssec` data set consists of 189 total zones.

4. OVERVIEW OF COLLECTED DATA

We sanitized the data by removing repeated records, records with empty name field, records that exhibited failed attempts at commenting, and records that were not supposed to have been transferred (such as those belonging to a sub-zone). We now present a combined overview of the collected data.

Total `.com`/`.net` zones	74,326,215
Name servers by name	1,611,145
Name servers by IP	820,547
Zones successfully transferred	4,947,993
Record types defined	59
Record types seen in data	42
Valid record types seen in data	40
Record types seen in > 10 zones	31
DNSSEC zones walked	28

Table 1: Aggregate statistics for the combined data sets.

42 different record types exist in our data sets. Table 1 presents aggregate statistics about the combined data sets. Some records such as `SOA` (Start of Authority), `NS` (Name Server), and `A` (Address) exist in nearly all zones. Interestingly, the `SOA` record, the only record type absolutely required for a zone to exist, is the only one we saw in every zone. Even the vital `NS` record was not present in a 0.2% of zones, even though it is required by the DNS specification, and despite the fact that we know every one of these

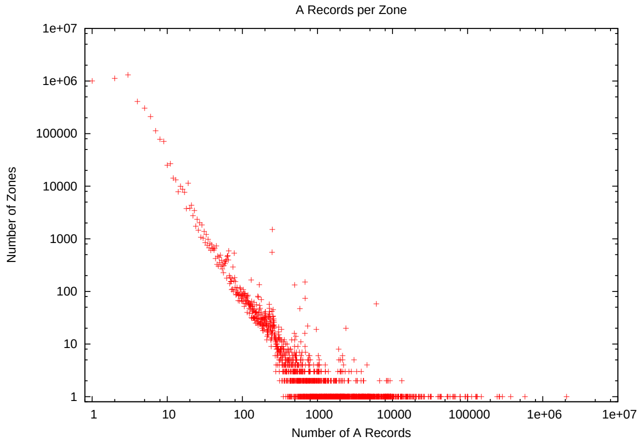

Figure 1: Number of A records per zone in the combined data set (log-log scale).

zones has at least one name server: the one we used to obtain the zone transfer. The next most popular record type is `MX` (Mail eXchange). Several record types showed up infrequently, with some appearing only in a single zone. In fact, three of the record types in our data are obsolete. We also found several email-related experimental resource records in our data. Finally, two of the record types we saw were not even allocated record types.

The zones in our data sets vary in sizes. To examine zone sizes, we looked at the `A` records contained in them. While not present in all zones, most zones have these records. Since all hosts must have an `A` or an `AAAA` (IPv6 Address) record, and IPv6 is not widely used, the number of `A` records in a zone should roughly correspond to the number of hosts in the zone intended to be accessible though DNS. Figure 1 shows the number of `A` records per zone. As seen in the figure, a majority of zones are small, containing only one `A` record. Some have more, but it was surprising how much more. The largest one has 2,073,715 `A` records. This is because it has an `A` record for each address in the 10.32.0.0-10.63.255.255 private IP address space in addition to enumerating every address in another public prefix. Further, there are 14 others with over 100,000 `A` records, although no others with over 1,000,000. These zones either have an `A` record for every address in a prefix or they have a large number of domain names all pointing to the same IP address. Some of these appear to be hosting providers. Most record types that were used by more than just a few zones followed the same trend as Figure 1 but with smaller proportions.

5. ANALYSIS OF DNS ZONES

5.1 Impact on Name Servers

The `NS` records in a zone indicate the name servers for that zone and for its sub-zones. Problems in these records could slow down DNS queries to the zone or even make the sub-zones inaccessible. We find that many zones have `NS` records that point to host names which are not externally accessible. In our `zone_transfer` data set, 35,618 zones (0.72%) have `NS` records with host names consisting of a single label (a host name with no dots in the name). These cannot be a host within any domain because a valid host name must have at least two dots in it. Further, 3,437 zones (0.07%) have `NS` records indicating name servers with host names in the `.local` TLD, which is not a valid TLD. Neither of these errors occur in any `dnssec` zone.

We also see problems in the hosts pointed to by the `NS` records. In 24,457 zones (0.5%) in the `zone_transfer` data and one zone in the `dnssec` data, there are `NS` records pointing to hosts for which no `A` records exist. In a further 3,337 zones (0.07%) in the `zone_transfer` data, there are `NS` records pointing to `CNAME` (Canonical Name) records, which will cause extra queries to be issued on every DNS query for the `NS` records [5], slowing down the resolution for these records.

5.1.1 *Diminished Name Server Redundancy*

The `NS` records also shed light on name server redundancy. Every zone is required to have at least two name servers [11] and recommended to have at least three [6]. This ensures availability of records when attacks or outages occur. 1,665 zones (0.03%) in the `zone_transfer` data list no name servers at all even though they are required to. Note, however, that this does not make them inaccessible. Clearly, they are accessible since we transferred their zone. Instead, it implies that their `NS` server records existed in their parent zone, but not in the zone itself, as well as they are required to. This problem does not occur in the `dnssec` data. Further, we find that 11.9% of zones have less than the required two name servers, and 22.1% with three or more in the `zone_transfer` data. In the `dnssec` data, we find 3% of zones with less than the required two name servers, and 66% with three or more, showing that the `dnssec` zones provide much better redundancy.

By separating name servers, both physically and in the network topology, zones can ensure that redundancy provides greater resiliency [6]. We examined name server redundancy at several granularities: according to the BGP prefix advertisements, by autonomous system (AS) they belong to, and across second-level domain names (the final two components of a domain name). In Table 2, we show the extent of redundancy at each granularity. *We note that 82% of the name servers in the `zone_transfer` data set are within the same AS, 61% within the same BGP prefix, and 91% within the same second-level domain. These results indicate that while name servers may be redundant, they are not physically or topologically distributed for many zones. This could potentially make them susceptible to single points of failure. Correspondingly, 7% of `dnssec` zones are in the same AS, 5% in the same prefix, and 12% in the same second-level domain.* Clearly, the `dnssec` zones pay more attention to the quality of redundancy in their name servers.

5.1.2 *Co-located Internal and External Name Servers*

In 0.5% of `zone_transfer` zones and 22 (11.6%) `dnssec` zones, we find `A` records pointing to private IP addresses. Since these records cannot be used by hosts external to the domain, their presence in a zone may be an indication that the zone is running the same DNS server for internal and external clients, and not separating them as is recommended. This has the unfortunate consequence of exposing the internal DNS server to attacks when separating the two would

#	Percent of Zones					
	zone_transfer			dnssec		
	AS	Prefix	Domain	AS	Prefix	Domain
1	82.3%	61.0%	90.7%	6.9%	4.8%	12.2%
2	15.6%	22.3%	8.4%	87.3%	33.3%	84.1%
3	1.9%	3.0%	0.5%	3.2%	58.2%	2.7%
4	0.2%	13.6%	0.2%	2.1%	3.2%	0.5%
5	0.04%	0.06%	0.00%	0.00%	0.00%	0.00%

Table 2: Number of ASes, BGP prefixes, and second-level domains name servers of the zones contained in the two data sets belong to.

normally make it hard for an adversary to even know the whereabouts of the internal DNS server. (Notice that the NS records are for external DNS servers only.)

5.2 Impact on Email and Anti-spam

MX records are used to indicate the email server for a domain. Problems with these records could lead to mail for the domain becoming undeliverable. 24% of zones in our zone_transfer data, and 9% in out dnssec data do not contain an MX record with a name matching their domain name. Thus, no email addresses could exist at the domain name of those zones. This in turn means that either the zone is not in use or its email services are provided by some other zone.

The MX records contain two data fields: a priority, and the host name of an email server. We found no discernible errors in the priority field. The types of errors we saw in the host name field were similar to the errors we saw in the data portion of NS records. In 4,452 zones (0.09%) in the zone_transfer data, there were MX records pointing to a host name which consists of a single label, and in 17 zones in this data, MX records point to mail servers in the .local TLD. The net result of these errors is the unavailability of mail for the domain name of the record if these are the only MX records for a domain, or delays in mail delivery if there are others. Neither of these errors occur in the dnssec data.

Many zones with valid MX record types have issues with the hosts those records pointed to. In the zone_transfer data, 18,376 zones (0.37%) had MX records pointing to nonexistent hosts, in that no A records to map those hosts to IP addresses exist. Further, 47,186 (0.9%) zones point to CNAME records, causing extra DNS look-ups to be necessary. These problems occur in 2 and 1 zones in the dnssec data respectively.

5.2.1 Anti-Spam Technologies

Spam is undoubtedly one of the biggest security issues facing the Internet today. To avoid accepting spam, technologies that verify sender identity before accepting email have been proposed. Prominent examples of email verification systems are DomainKeys [4, 1], SenderID [10], and Sender Policy Framework (SPF) [20]. Each of these technologies use the DNS as a database. The DNS has a specially-formatted TXT (Text) record for each. Additionally, SPF has a special record type defined for itself which was introduced later. SPF is the most popular of the technologies but most zones have a TXT record for it instead of the SPF record. Specifically, 409,214 zones (8.3%) in zone_transfer data set contain TXT records related to SPF while only 50 zones have the SPF record. The corresponding numbers for dnssec data set are 31 (16%) and zero. Far fewer zones deployed DomainKeys (0.7%) and SenderID (0.02%) in the zone_transfer data set. 10 zones (5%) in the dnssec data set used DomainKeys and none used SenderID. Overall, these numbers indicate that DNS-based anti-spam technologies are being deployed by a significant fraction of zones in the Internet, with those deploying DNSSEC doing so even more.

Unfortunately, not everybody is configuring the TXT records for the anti-spam technologies properly which may render the entire effort useless for those zones. Specifically, 371 zones (0.0075%) in the zone_transfer data set mistakenly replace the "1" in the version of SPF with an "l". Similarly, 384 zones (0.008%) in this data set follow the SPF syntax for SenderID by mistake. While it may be non-trivial to get the zone administrators to fix these errors, the recipient's mailer program can account for them easily to make use of the intended records. None of the zones in the dnssec data set make these mistakes.

5.3 Impact on Host Availability

We now look at problems which effect all hosts in a zone irrespective of their functionality. There are two types of records related to hosts: 1) the A records, which map host names to IP addresses and 2) the CNAME records, which provide canonical names for host names. We see very few errors in the A records. Specifically, only one zone in the zone_transfer data has an A record with no IP address. The story is not the same for CNAME records, where we see more problems. We begin by examining the *data* portion of CNAME records in the zone_transfer data. A few of the errors we see here are similar to the ones we found in MX and NS records. In particular, we see CNAME records where the host name pointed to is a single label in 1,543 (0.03%) zones. 27 zones have CNAME records pointing to a host name in the .local TLD. Several other types of errors exist as well. First, we see 93 zones with CNAME records pointing nowhere, and 165 where it points to a URL instead of a host name. In 1,288 zones (0.03%), there are CNAMEs pointing to an IP address. It is unclear what these are meant to do different from the functionality of an A record. Upon examining the dnssec data set, we find that none of these problems exist in those zones.

The *name* portion of the CNAME records also have issues. In the zone_transfer data, we see 28,082 (0.57%) zones where chains of CNAME records exist, with one CNAME pointing to another. We see this in 5 zones in the dnssec data. This problem slows down all DNS resolutions involving these CNAMEs. *Loops of CNAMEs are seen in 9970 (0.2%) zones in the zone_transfer data and 1 zone in the dnssec data. These will cause the CNAME to be unresolvable, leading to unavailability.*

5.4 Impact on Reverse DNS Availability

Reverse DNS is used to map IP addresses to host names, the inverse of normal DNS operation. This function is used for spam filtering and in system diagnosis tools such as traceroute.

Reverse DNS is accomplished though the used of PTR records. However, these records are normally not contained in the same zone as other records for the organization. They are instead contained in separate zones under the in-addr.arpa domain. Since we did not attempt a zone transfer of any of these zones, our data should not contain reverse DNS records. The only exception to this is where a CNAME record is used to map from a domain under in-addr.arpa to a regular domain. However, this is easy to distinguish by looking at the name on the record.

In less than 0.01% of the zones in the zone_transfer data, and a single zone in the dnssec data, we see records with domain names that belong in the in-addr.arpa tree, sometimes with other errors as well. In these cases, the reverse DNS look-up will be unable to use these PTR records, so the reverse look-up will be unsuccessful unless the correct record is also present in the correct zone under in-addr.arpa.

5.5 Timers and their Implications on Zone Expiration

An SOA record indicates the start of a DNS zone. Each zone is required to have a SOA record. Among other things, the SOA records contain the values of four timers which are important in DNS zone operations. These are the *refresh*, *retry*, and *expire* intervals, and the *minimum TTL*. The refresh, retry, and expire intervals all control the behavior of secondary DNS servers with regards to updates. The refresh interval indicates the amount of time (in seconds) a secondary DNS server should wait before checking to see if its copy of the DNS zone is current. The retry interval indicates how quickly it should retry this operation if it is unsuccessful at the end of the refresh interval. The expire interval indicates the amount of time that can elapse without successfully refreshing the zone before a secondary name server can no longer give authoritative answers to DNS queries for the zone. The minimum TTL is the default duration for which records from this zone can be cached by resolvers.

The refresh timer should be expected to have a value smaller than the expire timer. Otherwise, there will be a period where the DNS records cached at the secondary name server will be invalid before they are refreshed. During such a period, the availability of all the secondary servers will be reduced. We found that 14,003 (0.28%) of the zones in the zone_transfer data set have their expire timers set to values less than the refresh timers. Two of the dnssec zones have the same problem. In general, the zone administrators seem to be less careful about the expire timers. While the common values used for refresh and retry timers are mostly within the range of those recommended [3], the common values for the expire timer are 7 days and 41.6 hours. Both of these fall outside the recommended interval, which is 2-4 weeks.

5.6 Incomplete Contact Information

It is increasingly important that zone administrators be reachable. One example of such importance is phishing, where the process of shutting down phishing sites hosted at compromised servers belonging to reputable domains can benefit from being able to easily reach the domain administrators. Similarly, isolating members of bot armies or infected machines spreading malware can benefit significantly from the ability to contact their administrators. There are two places in the DNS records where such information is available. The first is the SOA record. Among other things, it contains the email address for the contact person for the zone. We found that all the zones in both of the data sets had an SOA record as required. While each of the SOA records we fetched in both data sets contained the email address, some in the zone_transfer data set contained this information in an incorrect format. The email address in SOA records is meant to be formatted with a "." replacing the "@" in the address. 29,946 (0.61%) of the zone_transfer zones fail to do so. Any automation to retrieve the email address in these records should thus account for the common mistake of failing to replace "@" with a ".".

The second place where the information about administrators can be present is the RP (Responsible Person) record. In fact, this information is meant to be more detailed, perhaps including phone numbers or full address of the contact person. Specifically, the RP record contains the email address of the zone administrator, and a pointer to a TXT record containing additional information. The email address in RP records should be formatted as in the SOA records. Unfortunately, a very small fraction of zones had this record: Only one dnssec zone and 6770 (0.14%) of the zone_transfer zones had it. Further, 2.6% of the RP records either contained no information or contained a single label that could not be an email address. Another 71.6%, including the one from the dnssec data set, just contained the email address and either pointed to an unusable TXT record or a non-existing one. This implies that 3/4th of the RP records at best contain as much contact information as the SOA record.

6. RELATED WORK

Wanrooij *et al.* [18], characterized DNS misconfigurations from a sample of the .NL cc-TLD. They did so by performing DNS ANY queries on 10,000 randomly zones mentioned in the .NL zone file. Their study had limited view of DNS provisioning because the ANY query, as they used, provides only a small subset of the records in a zone. Our analysis considers extensive information about orders of magnitude more domains.

Pappas *et al.* [15] examined the impact of three specific DNS configuration errors on the availability of name servers: lame delegation, diminished server redundancy, and cyclic dependency. We additionally examined the availability of other servers, including, mail server, Web server, etc. Our methodology for assessing diminished name server redundancy differed from theirs in several ways. First, we focused on security-conscious and less security-conscious domains in .com and .net TLDs while they sampled name servers of popular domains. Further, we examined redundancy in terms of ASes, BGP prefixes, and second-level domains when they focused on AS-level redundancy, geographic redundancy, and /24 prefixes. Due to these differences, we feel that the results are not directly comparable.

The Measurement Factory [17] has performed zone transfers on a small fraction of the .com and .net zones. They randomly sampled about 3.22% of .com and .net zones and attempted to transfer them. Though they had data similar to ours, they utilized it in ways that differ significantly from us. While we focus on information contained in zone records, they focused on the versions of DNS software in use (to infer possibility of cache poisoning), lame delegation, diminished server redundancy, and possibility of recursion (to infer potential misuse of such name servers by escaping detection). Surprisingly, they find that over 30% of the name servers allow a zone transfer. We find this percentage to be much lower – we were only able to transfer 6.6% of zones out of all the ones we attempted.

A few efforts have focused on developing tools for detecting misconfigurations present in DNS zone files. Pappas *et al.* [14] developed a tool to detect certain errors and inconsistencies by considering measurements from many vantage points. Many other tools to check for a variety of DNS problems are available online, for example at `dns.net` [16]. These tools analyze a single zone at a time and are not designed for the type of Internet-wide analysis we perform in this study. However, they are useful for administrators who wish to find and correct the errors in their own zones.

7. CONCLUSION

In this work, we investigated the intertwined relationships embedded in the various DNS records and the implications this relationship may have on the availability of services offered by the zone. The Internet-wide nature of our analysis allowed us to understand the common configuration mistakes that administrators make. Many of the problems we found occurred in a low number of zones, indicating that the DNS administrators are doing a fine job of maintaining DNS zones as a whole, with zones deploying DNSSEC doing so even better. The only notable exception to this conclusion was the lack of redundancy in the location of name servers, a problem which most of the zones we examined had, particularly among the ones that allowed us zone transfers. While an ideal approach will be to coax the zone administrators to correct the errors we found, we note that the consumers of this information can account for the errors themselves in a few cases, such as when the anti-spam records are misconfigured or when the email address of the zone administrators is not in the correct format.

Acknowledgments

Rob Henderson helped in running the zone transfers smoothly. Without his help, it would not have been possible to collect the data to conduct this research. Many thanks also to Mark Allman, who shepherded the paper and greatly helped its presentation.

8. REFERENCES

[1] E. Allman, J. Callas, M. Delany, M. Libbey, J. Fenton, and M. Thomas. DomainKeys identified mail (DKIM) signatures. IETF RFC 4871, May 2007.

[2] R. Ardnds, R. Austein, M. Larson, D. Massey, and S. Rose. Resource records for the DNS security extensions. IETF RFC 4034, Mar. 2005.

[3] D. Barr. Common DNS operational and configuration errors. IETF RFC 1912, Feb. 1996.

[4] M. Delany. Domain-based email authentication using public keys advertised in the DNS (DomainKeys). IETF RFC 4870, May 2007.

[5] R. Elz and R. Bush. Clarifications to the DNS specification. IETF RFC 2181, July 1997.

[6] R. Elz, R. Bush, S. Bradner, and M. Patton. Selection and operation of secondary DNS servers. IETF RFC 2182, July 1997.

[7] A. Householder, B. King, and K. Silva. Securing an internet name server. CERT Coordination Center Whitepaper, 2002.

[8] S. Josefsson. DNSSEC walker. `http://josefsson.org/walker/`.

[9] O. Kolkman, M. Fuhr, D. Franks, and C. Reinhardt. NET::DNS perl DNS reslover module. `http://www.net-dns.org`.

[10] J. Lyon and M. Wong. Sender id: Authenticating e-mail. IETF RFC 4406, April 2006.

[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] E. Osterweil, M. Ryan, and D. Massey. SecSpider. `http://secspider.cs.ucla.edu`.

[14] V. Pappas, P. Fältström, D. Massey, and L. Zhang. Distributed DNS troubleshooting. In *ACM SIGCOMM Workshop on Network Troubleshooting*, 2004.

[15] V. Pappas, Z. Xu, S. Lu, D. Massey, A. Terzis, and L. Zhang. Impact of configuration errors on DNS robustness. *ACM SIGCOMM Computer Communications Review (CCR)*, 34(4):319–330, 2004.

[16] A. Salamon. Tools to manage DNS. `http://www.dns.net/dnsrd/tools.html`.

[17] The Measurement Factory. DNS survey: October 2007. `http://dns.measurement-factory.com/surveys/200710.html`.

[18] W. van Wanrooij and A. Pras. DNS zones revisited. In *Open European Summer School and IFIP WG6.4/6.6/6.9 Workshop (EUNICE)*, 2005.

[19] VeriSign. Domain name industry brief, June 2007. `http://www.verisign.com/static/042161.pdf`.

[20] M. Wong and W. Schlitt. Sender policy framework (SPF) for authorizing use of domains in e-mail, version 1. IETF RFC 4408, Apr. 2006.