Poster: On the Resilience of DNS Infrastructure

Haya Shulman*
Fachbereich Informatik
Technische Universität Darmstadt

Shiran Ezra†
Computer Science Department
Bar-Ilan University

ABSTRACT

We study the operational characteristics of the DNS infrastructure: *transitive-trust*, *coresidence* and *servers placement*. We discuss how these factors impact resilience, stability and security of the DNS services. As our study indicates, common configuration choices, that domain operators make, result in a fragile DNS infrastructure, susceptible to malicious attacks and benign failures. We provide recommendations for improving robustness of DNS.

Categories and Subject Descriptors

C.2 [**Computer-Communication Networks**]: General, Network Architecture and Design, Distributed Systems

Keywords

DNS privacy; DNS resilience; coresidence; zone file security;

1. INTRODUCTION

The resilience and availability of the Domain Name System (DNS), [RFC1034, RFC1035], are critical to the stability and functionality of the Internet. Initially designed to translate domain names to IP addresses, the DNS infrastructure has evolved into a complex ecosystem and is increasingly utilised to facilitate a wide range of applications and constitutes an important building block in the design of scalable network infrastructures.

Best practices for ensuring availability and security of the DNS infrastructure recommend (1) defining a number of name servers for each domain, (2) configuring these name servers under at least two different parent domains and (3) placing the physical name servers, hosting the zone files for the domain, in separate networks. The redundancy provides for stability of the domain and prevents single point of failure. In particular, if one of the parent domains is not accessible, the domain will remain functional via the other parent domain; in case one of the networks, hosting the name servers, is under attack, the other name server, located in available networks, can be reached.

On the flip side, while ensuring availability, this redundancy introduces new dependencies which can be utilised to attack the domain. Specifically, if vulnerability exists in a network or a name server hosting the domain, it can be exploited to attack the domain, e.g., inject spoofed DNS record for domain hijacking.

The security and availability of domains in DNS also depends on other domains via a transitive trust. Transitive trust dependencies were introduced in [4], which observed that resolving a single domain name often requires traversing multiple other domains. Failure to resolve the domains in the transitive trust, e.g., due to misconfiguration or attack, or resolution to an incorrect address, may impact all the dependant domains.

In this work we study the operational characteristics of the DNS infrastructure and their impact on resilience, availability and stability of the DNS services. Specifically, we measure *transitive trust dependencies*, *coresidence* and *name servers placement*. Our investigation comprises top 50,000 Alexa domains[1] (`www.alexa.com/topsites`) and 568 TLDs. We show that: (1) resolution of names in many domains are susceptible to high latency and even failures due to multiple transitive trust dependencies; (2) high coresidence rate can disrupt services to multiple domains during benign failures or attacks on a single name server; (3) high concentration of name servers in certain geographical locations can facilitate censorship of (and attacks on) a high volume of DNS requests.

Domain Name System. DNS is a client-server protocol. The client side of the DNS infrastructure is composed of resolvers, which lookup records in zones by sending DNS requests to corresponding name servers. The resolvers communicate to the name servers using a simple request-response protocol (typically over UDP); for instance, (abstracting out subtleties) to translate `www.foo.bar` resolvers locate the name server `ns.foo.bar`, authoritative for `foo.bar`, and obtain the IP address of the machine hosting the web server of the website `www.foo.bar`; see sample resolution process in Figure 1. Resolvers store DNS records, returned in responses, in their caches for the duration indicated in the Time To Live (TTL) field of each record set. The resource records in DNS correspond to the different services run by

*`haya.shulman@gmail.com`
†`shiran.ezra@gmail.com`

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).
CCS'14, November 3–7, 2014, Scottsdale, Arizona, USA.
ACM 978-1-4503-2957-6/14/11.
http://dx.doi.org/10.1145/2660267.2662376.

[1]Alexa website provides a list of a million top ranked domains on the web.

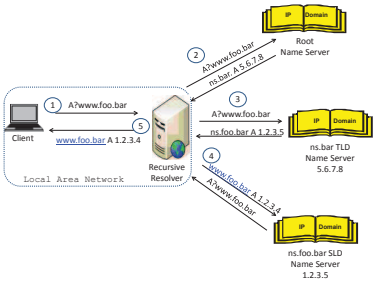

Figure 1: A sample DNS resolution process for `www.foo.bar`, initiated in step (1) when a client sends a request to the recursive resolver. In steps (2-4) the resolver queries the name servers, and in step (5) returns the response to the client.

the organisations and networks, e.g., hosts, servers, network blocks. The zones are structured hierarchically, with the root zone at the first level, Top Level Domains (TLDs) at the second level, and millions of Second Level Domains (SLDs) at the third level.

2. MEASURING DOMAIN NAME SYSTEM

We study resilience of the domains to failures and attacks by measuring: (1) inter-domain dependencies via transitive trust, (2) zones coresidence due to name servers sharing and (3) diversity of name servers' placement.

Our study encompassed top 25K Alexa domains and 568 TLDs. We also measured all the domains depending on these Alexa domains and TLDs via a transitive trust, which resulted in a total of 150K domains. These domains are served by 48K name servers; these 48K name servers have 65K different IP addresses, since sometimes a single name server is assigned a number of IP addresses.

2.1 Dependencies via Transitive Trust

A transitive trust dependency can be twofold: (1) a name server can appear in a number of transitive trust chains (i.e., how many domains can be impacted by a failure of a specific name server), and (2) a domain can depend on multiple domains for its resolution (i.e., a failure of a single server can impact the latency or availability of a domain). The former impacts the resilience of the DNS infrastructure and the later the resilience of a specific domain. Ideally, both should be low.

Our study shows that on average a domain in 25K-top Alexa depends on 43.5 other domains via transitive trust chains, and on average a domain in TLD depends on 43.7 domains via transitive trust chains. The maximal number of transitive trust dependencies in Alexa domains is 220 and in TLDs is 183. For instance, domain `sigcomm.com`, ranked 373097 on Alexa, is hosted at `dnsmadeeasy.com`, coresiding with 400 other domains.

Figure 2 plots the cumulative distribution function (CDF) $F(x) = \Pr[X \leq x]$ of the *number of name servers appearing in transitive trust chains* of 25K-top Alexa domains and TLDs; this encompasses the first dependency and increases the traffic volume to the name servers and resolution latency for the clients. The CDF curves for the Alexa domains and TLDs represent the number of transitive trust chains that a name server in Alexa domains (resp. TLDs) appears in. Approximately, 50% of name servers appear in two or more

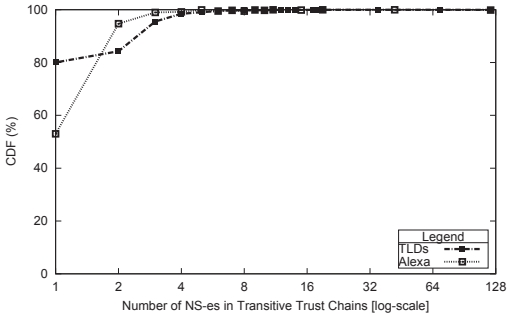

Figure 2: Name servers appearing in transitive trust chains of 25K-top Alexa domains and TLDs.

transitive trust chains. More than 90% of the name servers appear in 8 and less chains. Some name servers appear in more than 128 transitive trust chains. Figure 3 plots the CDF of the transitive trust dependencies of 25K-top Alexa domains and TLDs; this expresses the second dependency (2) - and increases resolution time for the clients. Approximately 50% depend on 20 or more other domains for their resolution, and more than 90% depend on more than 128 domains.

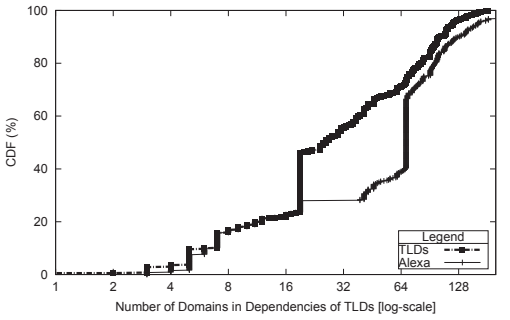

Figure 3: Transitive trust dependencies in 25K-top Alexa and TLDs.

Name servers with high dependencies via transitive trust, i.e., whereby multiple domains depend on them for their resolution, have two side effects: (1) they can become a lucrative target for attacks. For instance, recently, [2], showed how to launch a DNS cache poisoning attack using fragmented DNS responses, to replace the authentic IP address of a victim name server with a spoofed one, and ran this against `sns-pb.isc.org` name server, which appears in 69 transitive trust chains of other domains.

(2) Another notable side effect is that large transitive trust chains introduce more latency to resolution of records within domains depending on many other domains, and increase the queries rate to name servers appearing in multiple transitive trust chains. Our study measured an increase of 50ms for every transitive trust chain of 3 links, when measured with a cold (empty) cache. Resolutions of larger chains, e.g., 200, can often result in timeouts and unnecessary retransmissions, overloading the network and the name server, and increasing the latency for clients' queries.

Transitive-trust dependencies also nullify effectiveness of DNSSEC, [RFC4033-RFC4035], and impede its adoption [1, 3]. In particular, if name servers or other resources of a signed zone are placed under unsigned domains, the DNS re-

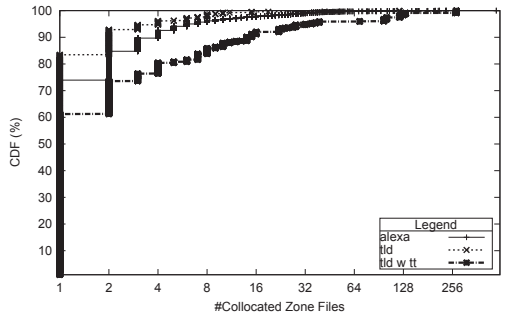

Figure 4: Zones coresidence on name servers among 25K-top Alexa domains and TLD, and servers in transitive trust dependencies of Alexa domains and TLDs.

solver will not be able to establish the security of the signed records, and the security will depend on the the security of the weakest link in a transitive trust chain.

2.2 Dependencies via Coresidence

Hosting a number of zones files on the same name server enables DNS name server operators to optimise profit and reduce operational costs and management overhead. We measure and quantify the dependencies between zones, namely *the fraction of zone files residing on the same physical server*. We measure the coresidence among TLDs and 25K-top Alexa domains, including the coresidence between name servers appearing in their transitive trust dependencies. As our results, plotted in see Figure 4, indicate, the coresidence rate among the name servers are extremely high. We found that coresidence of multiple zones on the same name server is a common practice among Alexa domains and TLDs. In particular, more than 70% of the name servers of Alexa domains and more than 80% of the name servers of TLDs, host multiple zones. Some name servers host more than 500 zone files, such as the name server `pdns.ultradns.net`.

The implications of high coresidence rate is that a failure or a DoS attack against the availability of the name server or the network hosting it, impacts the availability of *all* of the zones hosted on it. An attack against a name server, e.g., exploiting a vulnerability in a DNS software or in the operating system, can enable the attacker to take control over the name server and inject records into the zone files hosted on it. We also find that high coresidence increases loss and failures. In particular, even a moderate queries rate can result in packets' loss. Of course, high coresidence rates, large requests rate or responses sizes further exacerbate the problem.

2.3 Name Servers Redundancy and Placement

We find that a large number of name servers share a geolocation. Name servers hosting most zone files are located in the U.S.; The side effects of this phenomenon are that: (1) this may make it easier for the U.S. to take down domains hosted on name servers within its boarders; (2) this may facilitate redirection of clients to incorrect hosts by injection of spoofed records into the zone files, e.g., for censorship; (3) the zones and domain operators have to be compliant with the policies and regulations of the U.S. government; (4) the latency increases with the fraction of coresident zone files, in particular, we observed that requests to name servers with more than ten zone files incur on average at least 10 milliseconds additional delay than requests to name servers hosting 2 or less zone files.

We studied location of name servers appearing in transitive trust chains of 25K-top Alexa domains and of TLDs, Figure 5. The question that we seek to answer is twofold: (1) which country has most name servers in transitive trust chains, and (2) which country hosts name servers appearing in most transitive trust chains.

The answer to the first question is that more than 30% of the name servers are located in the U.S. Then Canada and U.K. (with Canada leading among Alexa domains, and U.K. among TLDs). 'Others' stands for different countries hosting less than 1% of the name servers.

With respect to the second question, we found that the name servers appearing in most transitive trust dependencies of other domains reside in the U.K., e.g., a single name server appears in more than 121 different transitive trust chains (and this holds for a number of name servers located in the U.K.). The name servers in the U.K. constitute 16% of all the name servers appearing in transitive trust chains, and the name servers in the US constitute 26%.

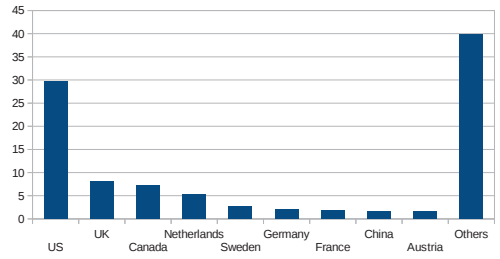

Figure 5: Location of servers in transitive trust chains.

Conclusions and Recommendations

Our study evaluated factors impacting availability, resilience and stability of DNS. Based on our findings we make the following recommendations for improving the resilience: (1) reduce coresidence rates by limiting the number of zone files that a name server can host, (2) reduce appearences of name servers in multiple transitive trust chains, by registering name servers under your own domain, (3) place name servers in diverse geo-locations.

Acknowledgements

This research was supported by projects EC SPRIDE (BMBF) and CASED (LOEWE).

3. REFERENCES

[1] A. Herzberg and H. Shulman. DNSSEC: Security and Availability Challenges. In *Communications and Network Security (CNS), 2013 IEEE Conference on*, pages 365–366. IEEE, 2013.

[2] A. Herzberg and H. Shulman. Fragmentation Considered Poisonous: or one-domain-to-rule-them-all.org. In *CNS 2013. The Conference on Communications and Network Security. IEEE*. IEEE, 2013.

[3] A. Herzberg and H. Shulman. Retrofitting Security into Network Protocols: The Case of DNSSEC. *Internet Computing, IEEE*, 18(1):66–71, 2014.

[4] V. Ramasubramanian and E. Sirer. Perils of transitive trust in the domain name system. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, pages 35–35. USENIX Association, 2005.