

Quality of Name Resolution in the Domain Name System

Casey Deccio, *Member, IEEE*, Jeff Sedayao, Krishna Kant, *Senior Member, IEEE*,
and Prasant Mohapatra, *Fellow, IEEE*

Abstract—Name resolution using the Domain Name System (DNS) is integral to today’s Internet. The resolution of a domain name is often dependent on namespace outside the control of the domain’s owner. In this article we review the DNS protocol and several DNS server implementations. Based on our examination, we propose a formal model for analyzing the name dependencies inherent in DNS, and experimentally validate the model with actual domain names. Using our name dependency model we derive metrics to quantify the extent to which domain names affect other domain names. It is found that under certain conditions, more than half of the queries for a domain name are influenced by namespace not expressly configured by administrators. This result serves to quantify the degree of vulnerability of DNS due to dependencies that administrators are unaware of. When we apply metrics from our model to production DNS data, we show that the set of domains whose resolution affects a given domain name is much smaller than previously thought. However, behaviors such as using cached addresses for querying authoritative servers and chaining domain name aliases increase the number and diversity of influential domains, thereby making the DNS infrastructure more vulnerable.

I. INTRODUCTION

The Domain Name System (DNS) is one of the systems most critical to the Internet. Nearly all Internet applications rely on the DNS for proper function. Although Internet hosts communicate using the Internet Protocol (IP) and are identified by IP addresses, the DNS abstracts IP addressing from users and clients, so they can identify Internet hosts with domain names representative of human-friendly words, titles, and abbreviations. In addition to IP address lookup, the DNS is also necessary for email delivery, service discovery, and host identification.

Since DNS details are often left to the client resolver and abstracted at the application level, its integrity and security are critical. While temporary failures due to misconfiguration may cause inconvenience, targeted attack by malicious parties could be much less discernible, and the repercussions more severe. Tainted DNS responses may transparently redirect user applications to servers within adversarial control, where sensitive information can be stolen.

This research was supported in part by the National Science Foundation under the grant CNS-0716741.

C. Deccio is with Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000 (email: ctdecci@sandia.gov).

J. Sedayao and K. Kant are with Intel Corporation (email: jeff.sedayao@intel.com; krishna.kant@intel.com).

P. Mohapatra is with UC Davis (email: pmohapatra@ucdavis.edu).

While the concept of name resolution is relatively simple, the overall system is complex and its effects far-reaching. Name resolution for a domain is often dependent on servers well outside the control of the domain’s owner and managed by third parties. A network of inter-organizational relationships overlays the DNS infrastructure, and configurations that create a dependency on peer organizations are in turn affected by the security and accuracy of namespaces linked through this network. An understanding of a domain’s context in the entire system will allow architects and administrators to better design and configure their DNS services to maximize the reliability of DNS name resolution.

In this work we present a quantitative analysis of name dependencies in DNS. We review pertinent aspects of the DNS protocol and derive a model to probabilistically determine the namespace influencing resolution of a domain name. Based on our model, we define several metrics to assess the quality of name resolution for a domain name, based on the other names that affect its resolution. The behaviors of name server implementations, some of which are configurable, can affect these metrics. We analyze a large sample of recent DNS data in light of the presented model. The results show the impact of name dependencies in terms of the size of the namespace that influences a domain name, and in the percentage of queries that will reach namespace not explicitly configured by DNS administrators.

We list the following as primary contributions of this article:

- A detailed study of DNS protocol specification and name server implementation, with respect to DNS name dependencies
- A formal model for analysis of DNS name dependencies, based on specification and implementations
- Metrics for quantifying the influence domain names have on other domain names

Our analysis shows that 92% of influential namespace of domain names is explicitly configured by domain administrators. However, certain caching behaviors reduce that figure, and increase the probability that resolution of a domain name is influenced by names or organizations not explicitly configured by administrators. A diverse set of dependencies has the potential to affect the reliability of name resolution for a domain name. Based on our findings, we discuss best practices for design and administration of DNS services to contain influence of domain names.

We begin our analysis with a review of the DNS protocol in Section II. We describe previous work in this area in Section III. In Section IV we introduce the concept of DNS




Fig. 1. The zone hierarchy for the the zone data shown in Table I.

name dependencies, present a graph model for their analysis, and derive methodology for quantifying influence. We describe methodologies employed for data collection and an analysis of the observed quality of name resolution in Section V. We explore considerations that increase outside influence in Section VI, and we conclude our analysis in Section VII.

II. DNS BACKGROUND

The Domain Name System (DNS) [1]–[3] provides a name-based lookup service for the Internet. It was designed as a robust and scalable database distributed among servers spread across the globe. The DNS namespace is organized hierarchically. Each *domain name* is comprised of dot-separated, alpha-numeric *labels*, which describe the ancestry of the name itself, from left to right. For example, the ancestry of *www.soccer.com* is: *www.soccer.com*, *soccer.com*, *com*, and the *root* domain (represented by a single dot with no labels: “.”). All domain names are descendants of the root.

DNS data is managed within the DNS *zone* to which it pertains. A zone includes all the subdomain namespace below its *zone cut* or *origin*, except that whose management has been explicitly *delegated* to child zones. Fig. 1 illustrates a zone hierarchy in which the *com* zone delegates management of the *soccer.com* namespace to the *soccer.com* zone. In this case, *www.soccer.com* is part of the *soccer.com* zone because there is no further delegation to a *www.soccer.com* zone. For each zone a set of name servers are configured and advertised as *authoritative* to provide answers for non-delegated namespace in that zone. Servers may be authoritative for multiple zones. Logistics for delegation are covered subsequently in this section.

A. Zone data

DNS zone data consists of *resource records* (RRs), each identified by an owner name (e.g., *www.soccer.com*), a class (e.g., IN)¹, and a type (e.g., A). Among the RR types defined, the following are pertinent to this article: A (IP address)², CNAME (canonical name), and NS (name server). Each RR contains record data specific to its RR type. For example, the record data for an A RR is an IP address (e.g., 192.0.2.1). The record data for RRs both of type CNAME and NS is

¹We abstract RR class from the remainder of our analysis, as our research focuses on specific RR types in the IN class.

²This work focuses on DNS name dependencies, which affect name resolution to both IPv4 and IPv6 addressing using RRs of types A and AAAA, respectively. For simplicity, however, our analysis references only RRs of type A.

\$ORIGIN soccer.com. (SoccerMania, Inc.)			
	Name	Type	Value
1	soccer.com.	NS	ball.soccer.com.
2	soccer.com.	NS	racket.tennis.com.
3	soccer.com.	NS	ns1.sports.net.
4	ball.soccer.com.	A	192.0.2.1
5	www.soccer.com.	CNAME	www.tennis.com.

\$ORIGIN tennis.com. (Tennis Pro, Inc.)			
	Name	Type	Value
1	tennis.com.	NS	ns1.tennis.com.
2	tennis.com.	NS	ball.soccer.com.
3	tennis.com.	NS	ns1.sports.net.
4	ns1.tennis.com.	A	192.0.2.2
5	www.tennis.com.	A	192.0.2.3
6	racket.tennis.com.	A	192.0.2.4

\$ORIGIN athletics.com. (Sports Central, Inc.)			
	Name	Type	Value
1	athletics.com.	NS	ns1.athletics.com.
2	ns1.athletics.com.	A	192.0.2.8

\$ORIGIN com. (VeriSign, Inc.)			
	Name	Type	Value
1	com.	NS	ns1.com.
2	ns1.com.	A	192.0.2.5
3	athletics.com.	NS	ns1.athletics.com.
4	soccer.com.	NS	ball.soccer.com.
5	soccer.com.	NS	racket.tennis.com.
6	soccer.com.	NS	ns1.sports.net.
7	tennis.com.	NS	ball.soccer.com.
8	tennis.com.	NS	ns1.tennis.com.
9	tennis.com.	NS	ns1.sports.net.
10	ball.soccer.com.	A	192.0.2.1
11	ns1.tennis.com.	A	192.0.2.2
12	ns1.athletics.com.	A	192.0.2.8

\$ORIGIN sports.net. (Sports Central, Inc.)			
	Name	Type	Value
1	sports.net.	NS	ns1.sports.net.
2	sports.net.	NS	ns1.athletics.com.
3	ns1.sports.net.	A	192.0.2.6

\$ORIGIN net. (VeriSign, Inc.)			
	Name	Type	Value
1	net.	NS	ns1.net.
2	ns1.net.	A	192.0.2.7
3	sports.net.	NS	ns1.sports.net.
4	sports.net.	NS	ns1.athletics.com.
5	ns1.sports.net.	A	192.0.2.6

TABLE I

THE ZONE DATA FROM SEVERAL FICTITIOUS ZONES, WHOSE HIERARCHY IS SHOWN IN FIG. 1. ANY COINCIDENCE WITH ACTUAL ZONES OF THE SAME NAME IS UNINTENTIONAL.

another domain name—a *target* for further lookup (e.g., *www.tennis.com*). The result of a DNS lookup for a domain name and type is the set of all RRs matching that criteria, known as an *RRset*.

Table I contains zone data for some fictitious zones.

One or more servers are designated as *authoritative* for a zone by including their domain names as the targets of RRs comprising the NS RRset corresponding to the zone origin. In order to properly delegate the namespace, this RRset must not only exist in the child zone, but also in its parent—as so-called *delegation records*. This RRset forms the link between the parent and child, and proper function requires careful

QUESTION			
www.soccer.com.		A	
ANSWER			
www.soccer.com.	3600	CNAME	www.tennis.com.
www.tennis.com.	3600	A	192.0.2.3
AUTHORITY			
soccer.com.	3600	NS	ball.soccer.com.
soccer.com.	3600	NS	racket.tennis.com.
soccer.com.	3600	NS	ns1.sports.net.
ADDITIONAL			
ball.soccer.com.	3600	A	192.0.2.1
racket.tennis.com.	3600	A	192.0.2.4
ns1.sports.net.	3600	A	192.0.2.6

Fig. 2. A response for `www.soccer.com` issued by `ns1.sports.net`, which is authoritative for `soccer.com`.

coordination between administrators of both zones. In Table I the delegation and authoritative NS RRsets for `soccer.com` are found on lines 4–6 of the `com` zone and lines 1–3 of the `soccer.com` zone, respectively.

B. Name resolution

DNS servers have two primary roles: *resolver* and *authoritative server*. The resolver queries authoritative servers to find the answers to domain name lookups. The answers, comprised of RRsets, may be stored in the resolver’s cache for later retrieval. Each RRset includes time-to-live value (TTL), which designates the lifetime for which it may persist in a resolver’s cache. Authoritative servers respond with answers to names corresponding to zones for which they are authoritative (e.g., `ball.soccer.com` for the `soccer.com` zone). If an authoritative server receives a query for a domain name in namespace which has been delegated, then it responds with a referral to the servers authoritative for the child zone.

DNS messages are comprised of several sections, each built of RRsets. The *question* section contains a single RR describing the name and RR type being queried but does not include the corresponding record data. Answers to the question are returned in the *answer* section. Information about the servers authoritative for the zone of the name being queried are contained in the *authority* section. The *additional* section contains supplemental information that may be helpful or necessary for the resolver. The anatomy of a response that might be issued by a server authoritative for `soccer.com` in response to a query for `www.soccer.com` is shown in Fig. 2.

A resolver begins the name resolution process by issuing a query to one of the servers authoritative for the root zone. This root server responds by populating the authority section of the reply with the NS RRset for the delegated zone that is in the ancestry of the name in question. The resolver re-issues the query with the same question, this time directing it to one of the servers corresponding to the NS RRset provided in the authority section of the response. The resolver iteratively continues this process until it receives a response from an authoritative source containing either an answer or an indication that the requested RRset does not exist.

The NS RRset in a referral specifies the authoritative servers for a delegated zone by name, but a resolver must have IP addresses to actually query the servers. Any addresses already

known by the resolver initially populate the authoritative server list for the zone. The resolver initiates requests in parallel to acquire the addresses for any others. For any NS targets that are subdomains of the delegated zone the referring server must provide the addresses in the additional section of its response to avoid a chicken-and-egg problem with resolving the server names. These addresses are included as *glue records* in the referring parent zone, so name resolution can be properly “bootstrapped” [1]. The A glue record for `ns1.tennis.com` is on line 11 of the `com` zone in Table I.

If a server issuing a referral to a delegated zone has pertinent non-glue A RRsets available locally, it may include them in the additional section of its response as a resource for the resolver. Such RRsets might be available if the referring server is also authoritative for the zones to which the names belong or if it has in its cache the A RRset for one of the names in question [1]. However, owner names of any such RRsets that are not subdomains of the referring zone are considered *out-of-bailiwick* (i.e., outside its jurisdiction). Resolver implementations should independently obtain an authoritative answer for the out-of-bailiwick target names before querying such servers.

We reference several name server implementations throughout the remainder of this work. The Berkeley Internet Name Daemon (BIND, version 9.6) [4] implements both a resolver and an authoritative server. `unbound` (version 1.0.2) [5] and `dnscache` (version 1.05) [6] implement resolvers, and the Name Server Daemon (NSD) [7] and `tinydns` [6] implement authoritative name servers.

1) *Trust ranking*: An understanding of how resolvers and authoritative servers rank the trustworthiness of DNS data is relevant to our analysis of DNS name dependencies. RFC 2181 [8] outlines a relative ranking of trustworthiness of data for name servers (resolver and authoritative server) to consider as part of operation. Among the total ranking are the following (in decreasing order of trustworthiness):

- Data from a zone for which the server is authoritative, other than glue data
- The authoritative data included in the answer section of an authoritative reply
- The data in the authority section of an authoritative reply
- Glue from a zone for which the server is authoritative
- Data from additional section of a response

This ranking affects behavior of both resolvers and authoritative servers. A resolver must initially use the delegation records provided in the referral from a server authoritative for the parent. However, when it receives an answer from a server authoritative for the zone itself, the authoritative server often includes the authoritative NS RRset for the zone in its authority section. Because these two sets may differ, the resolver will use the NS RRset returned in the authoritative response in preference to the delegation records received in the referral.

If a resolver receiving a referral has the A RRset for an NS target in its cache, previously received in the answer section of an authoritative response, then the resolver deems the cached data more trustworthy than any data received in the additional section of the referral. Thus, it will use the locally cached data in preference to the additional data, which includes glue records.

Both the BIND and unbound resolvers respect the notion of trustworthiness described in RFC 2181, using the NS RRset from an authoritative response in preference to that received in the authority or additional section of a referral. However, `dnscache` updates its cache with the most recently received RRset, regardless of the section in which it was received.

The combined service of cached and authoritative data by a single name server is allowed by specification, although best practices suggest that these functionalities not co-exist [9]. BIND’s implementation allows an authoritative server to additionally maintain a cache, unlike NSD and `tinydns`, which only serve locally available zone data. With such a configuration, if an authoritative server providing a referral has in its own cache the A RRset for an NS target, which it received in the answer section of an authoritative response from another server, it will trust this data more than a corresponding glue record from its local configuration. Thus the authoritative server will use the data from cache to populate the additional section of a referral to a resolver, rather than the glue record, which might be different.

Section VI describes practical examples of these caching behaviors in more detail.

2) *Name server selection*: RFC 1035 [2] dictates that a resolver associate historical statistics, such as response time and success rate, with each server address. A resolver will avoid using an authoritative server twice until all other authoritative servers for the zone have been tried at least once. After that, it prefers the server with the best performance record, thus fine-tuning the performance for lookups for a zone [2].

Server selection behavior varies by resolver implementation. The BIND and unbound resolvers both follow the performance-based selection guideline, giving preference to authoritative servers with historically better responses. The `dnscache` resolver, however, selects an authoritative server uniformly at random for each query.

Although individual resolvers following specification may gravitate towards a particular authoritative server, we assume that requests for a zone arrive from resolvers in diverse network and geographic locations with varying proximity to authoritative servers. Based on this assumption we use a uniform distribution to describe queries issued to authoritative servers for a zone, such that any server has an equal chance of being queried for names within the zone.

C. NS target distribution

In a simplistic zone configuration, the names of authoritative servers are in the zone itself. In this case, a resolver determining addresses corresponding to the authoritative servers has no need to query outside the namespace of the zone to determine their IP addresses. Also, if server names are within the zone, we may infer that the authoritative servers are under administrative control of the same organization that governs the zone. For example, `ball.soccer.com` is in the `soccer.com` zone. If it alone were authoritative for `soccer.com`, then the administration of `soccer.com` would be self-contained, and the dependency on external entities is minimized.

However, such a self-contained setup may not attain the network and geographic diversity to achieve the DNS robustness

desired by an organization. Servers administered by a single organization may result in a single point of failure for the zone. Topological diversity may be achieved by outsourcing DNS administration to an organization specializing in network diversity. A less expensive solution to gain diversity is partnership with other organizations with similar interests to host each other’s zone data on their servers. Either configuration is typically achieved by including the names of authoritative servers from the other organizations as targets in the zone’s NS RRset. This is exemplified in Table I, in which the `racket.tennis.com` server is authoritative for the `soccer.com` zone. This sharing provides both load balancing and resilience in the case of network failure or other disaster, but it creates a dependency on an external entity.

D. CNAME aliasing

The CNAME RR type is defined to provide aliasing of one domain name to another. For example, `www.soccer.com` is an alias for `www.tennis.com` in Table I. If a resolver receives a response indicating that the name in question is an alias to another name, it must subsequently resolve the target of the alias, and so on until an address is returned. This functionality is mostly provided for ease of maintenance or for seamless transition of namespace.

CNAME RRs are prohibited from co-existing with other RRs of the same name [8], [10], with the exception of the RRSIG and NSEC RRs, introduced as security extensions [11]. It is also recommended that CNAME RRs not be used as targets of NS RRs, MX (mail exchange) RRs, or other CNAME RRs [10]. However, these practices are not completely disallowed by existing implementations. The misuse of aliases results in added complexity, overhead, and potential vulnerability to name resolution.

III. PREVIOUS WORK

In this section we review related research. First we summarize previous studies analyzing DNS dependencies, misconfiguration, and availability. We then describe some of the solutions which have been proposed to protect name resolution from unwanted influence.

A. Analysis

Ramasubramanian, et al. [12] demonstrate the far-reaching effects of DNS dependencies by surveying the DNS namespace and tracing the dependencies of a large number of domain names. They identify the set of all authoritative name servers potentially involved in resolution of each name, which comprise the potential attack target for that name. Their results show that a domain name relies on 44 authoritative servers on average and 6% of names depend on more than 200 servers. We perform further examination of name resolution behaviors to create a formal model of name dependencies in DNS and quantify the significance of such dependencies.

A survey of common DNS misconfigurations is presented by Pappas, et al. [13]. The authors identify lame delegation, diminished server redundancy, and cyclic dependencies as

problem areas, and an empirical study is performed to show their pervasiveness in DNS. The formal model proposed in this article presented in this article creates a foundation to systematically identify and quantify the impact of such misconfigurations on security and availability.

B. Solutions

The DNS Security Extensions (DNSSEC) [14] are the leading standard for cryptographically validating DNS queries. DNSSEC extends the DNS protocol [1], [2] to allow resolvers to authenticate DNS responses using cryptographic public keys and signatures embedded in new RR types [11], [14].

The IKS Jena (Information, Communication, and Systems) [15] and SecSpider [16], [17] projects have monitored the deployment of DNSSEC, maintaining an online listing and status of signed zones. Signed zones are discovered through Web crawling techniques, by names submitted by users via the project Web sites, and by traversing NSEC RRs in known signed zones. Zone data collected from the SecSpider project was used to perform an assessment of availability, verifiability, and validity of DNSSEC deployments [18]. The SecSpider and IKS project sites report the number of production signed zones at roughly 24,000 and 44,000, respectively, at the time of this writing [15], [17]. Only 0.02% of the zones we used in our analysis (see Section V) were signed with DNSSEC.

Proper application of DNSSEC foils attempts at cache poisoning and other attacks on DNS integrity. However, just like unsigned names, DNSSEC-signed names also rely on the availability of their dependencies. In fact, the added maintenance and protocol complexity of DNSSEC actually increases the reliance on existing dependencies, such as the “chain of trust” following a zone hierarchy. The model presented in this article is thus a useful analysis for both signed and unsigned domain names.

Several non-cryptographic approaches have been proposed for defending against response spoofing attacks. Dagon, et al. [19] present a technique in which a resolver mixes the case of the name being queried and requires the authoritative server to preserve the case of the name in the question section of the response. This increases the probability space that attackers have to work with. Another technique for increasing entropy is WSEC DNS [20], which introduces wildcard DNS records into a zone, which aliases the actual DNS records. Clients generate a random string which is prepended to the wildcard name, and attacks require a larger effort to spoof responses for the randomized queries.

Additional entropy makes guesswork by a third-party more difficult. However, if an authoritative server on which a domain name is dependent for its resolution is run by a nefarious organization or otherwise compromised by offenders, then no guesswork is needed. An understanding of DNS dependencies will aid administrators in employing practices that will help them maintain control of the name resolution for their domains.

Other systems for next-generation name resolution have been proposed for enhancing performance and security. Yuan, et al. [21] propose DoX, a peer-to-peer solution in which a

single DNS lookup is performed by multiple collaborating peers in parallel, and the answers are compared by the requester for accuracy. Cache poisoning is detected by an inconsistent set of responses. CoDoNS [22] is another peer-to-peer approach in which the namespace is distributed uniformly across the peer-to-peer network, rather than hierarchically. The hash of each domain name maps to a “home node” in the network, which is the authority for that name. Name lookups are cryptographically verified using certificates issued on a per-name basis.

IV. DNS DEPENDENCY MODEL

In this section we introduce the notion of domain name dependencies and establish a formal model for identifying and quantifying domain name influence.

A. Name dependencies

Various DNS dependencies stem from the DNS protocol specification. Three specific components introducing dependencies are the following:

- *Parent zones*: A resolver must learn the authoritative servers for a zone from referrals from the zone’s hierarchical parent.
- *NS targets*: The NS RR type uses names, rather than addresses, to specify authoritative servers, so a resolver must resolve the names before it can query the authoritative servers.
- *Aliases*: If a name is an alias, then to obtain an address, a resolver must subsequently resolve the alias target.

We say that name u *depends on* domain name v if resolution of v may *influence* resolution of u . Influence is categorized into two classes: *active* and *passive*. If domain name u is actively influenced by domain name v , then with some non-zero probability resolution of v will be *required for* resolution of name u . If domain name u is passively influenced by domain name v , then although v may not be required for resolution of u , resolution of v may *affect* resolution of u with some probability.

Parent zones and aliases are both examples of active influence because both are required for resolution of the dependent name. However, NS target dependencies may either be the cause of passive or active influence, depending on the state of server cache and server implementation behavior.

We establish some notation for our discussion of NS target dependencies. To distinguish NS RRsets from different sources, we use NS_z and NS'_z to respectively denote the sets of NS target names for zone z , as configured in zones z and $Parent(z)$, respectively. Let NSA_z and NSA'_z denote the sets of IP addresses corresponding to the server names in NS_z and NS'_z , respectively.

Suppose $v \in NS_z$ is a subdomain of $Parent(z)$, $Parent(v) \neq z$, and $Parent(z)$ is properly configured with a glue record for v . If an authoritative answer for v has previously been resolved and cached by either authoritative server $s \in NSA_{Parent(z)}$ or resolver c , then z is affected by v and its name dependencies, based on the trust ranking discussion from Section II. This behavior describes passive

Term	Definition
r	The root name “.”
$I_u(v)$	The measure of name v 's influence on name u
$I_u(D)$	The aggregate influence of names in set D on name u
$Parent(d)$	The nearest ancestor zone of name d
$Cname(d)$	The alias target of name d
NS_z, NS'_z	The set of NS target names authoritative for zone z , as configured in z and $Parent(z)$, respectively
NSA_z, NSA'_z	The set of addresses corresponding to the names in NS_z and NS'_z , respectively
$P_{NS}(z)$	The probability that the resolver has the NS RRset for z cached from an authoritative source
$P_{\{s,c\}}(v)$	The probability that either s or c has in cache and uses NS target name v from an authoritative source
$G_d = (V_d, A_d)$	Name dependency graph for name d
$G'_d = (V'_d, A'_d)$	Active influence dependency graph for name d
$P_q(z, v)$	The probability that NS target v is used to resolve z
$w(u, v)$	The weight of edge (u, v) in A_d
S_u	The set of addresses corresponding to name u
$U'_d \subseteq U_d \subseteq Z_d$	The sets of first-order, non-trivial, and all zones in V_d , respectively

TABLE II
NOTATION USED IN THIS RESEARCH.

influence of v on z . The probability of passive influence, $P_{\{s,c\}}(v)$, is the combined probability of $P_s(v)$ and $P_c(v)$, the likelihood that either s or c has and uses a cached authoritative answer for v . Since the probabilities are independent of one another, $P_{\{s,c\}}(v)$ is calculated:

$$P_{\{s,c\}}(v) = P_s(v) \vee P_c(v) = 1 - (1 - P_s(v))(1 - P_c(v))$$

A resolver is responsible for resolving any names from NS'_z which are out-of-bailiwick or not included in the additional section of a response from an authoritative server. Such induced queries indicate active influence of the resolved names on z , since it is directly dependent on their resolution.

B. Name dependency graph

To model the dependencies of domain name d we use a directed, connected graph, $G_d = (V_d, A_d)$. The graph G_d contains a single sink, r , which is the root zone. Each node in the graph $v \in V_d$ represents a domain name, and each edge, $(u, v) \in A_d$, signifies that u is directly dependent on v for proper resolution of itself and any descendant names. Each edge, $(u, v) \in A_d$, carries a weight, $w(u, v)$, indicative of the probability that it will be followed for resolving u . A name dependency graph for domain name *www.soccer.com* is shown in Fig. 3, built from the data in Table I.

Edges are placed on the graph from each domain name $u, u \neq r$ to its parent $Parent(u)$ with $w(u, Parent(u)) = 1$; a domain name is always dependent on its parent. If resolution of domain name u yields a CNAME RR, then an edge is placed between u and its target name, $Cname(u)$, with $w(u, Cname(u)) = 1$; the resolution of an alias is always dependent on the resolution of its target. Such edges in Fig. 3 are those between *www.soccer.com* and its parent, *soccer.com*, and between *www.soccer.com* and its canonical name, *www.tennis.com*.

Name	Type	Value	$P_q(z, v)$
foo.com.	NS	ns1.foo.com.	$\frac{2}{3} = 0.67$
foo.com.	NS	ns2.foo.com.	$\frac{1}{3} = 0.33$
ns1.foo.com.	A	192.0.2.9	
ns1.foo.com.	A	192.0.2.11	
ns2.foo.com.	A	192.0.2.10	
bar.com.	NS	ns1.bar.com.	$\frac{1+0.5}{2} = 0.75$
bar.com.	NS	ns2.bar.com.	$\frac{0.5}{2} = 0.25$
ns1.bar.com.	A	192.0.2.12	
ns1.bar.com.	A	192.0.2.13	
ns2.bar.com.	A	192.0.2.12	

TABLE IV
EXAMPLE ZONE DATA TO ILLUSTRATE QUERY DISTRIBUTION AMONG NS TARGET NAMES OF SERVERS AUTHORITATIVE FOR A ZONE.

Placement of edges and weights corresponding to NS target dependencies is somewhat involved and draws from the discussion in Section IV-A. The considerations are summarized in Table III.

We first identify the proportion of queries distributed among each of the NS target names in NS'_z , which we use as a base for calculating the weights of edges in A_d stemming from NS target dependencies. Resolvers prepare their lists of servers for selection from *addresses* rather than *names* of authoritative servers, as explained previously. Although a particular resolver instance will gravitate towards querying a single authoritative server for a zone, we reiterate our assumption that queries from arbitrary resolvers are distributed uniformly among the servers, due to network diversity. Thus, the probability, $P_q(z, v)$, of querying any NS target $v \in NS'_z$ for resolution of z will be some fraction of $|NSA'_z|$ that reflects the proportion of server addresses attributed to v . Let S_v represent the set of addresses to which $v \in NS'_z$ resolves. A naïve formula for determining query probability $P_q(z, v)$ is to simply calculate the fraction of total server addresses authoritative for z that correspond to v :

$$P_q(z, v) = \frac{|S_v|}{|NSA'_z|}$$

The zone data for *foo.com* in Table IV shows that an NS target name that resolves to multiple addresses, such as *ns1.foo.com*, has a higher probability of being queried for names in the zone than an NS target name that resolves to only a single address, such as *ns2.foo.com*.

It is possible that multiple NS target names in NS'_z resolve to the same address, so a single address in S_v may also be attributed to other names in NS'_z . A more complete approach to determining query probability therefore is to evenly divide the probabilistic weight attributed to a server address among all the names that resolve to that address:

$$P_q(z, v) = \frac{\sum_{s \in S_v} |\{u \in NS'_z | s \in S_u\}|^{-1}}{|NSA'_z|}$$

For example, in Table IV both *ns1.bar.com* and *ns2.bar.com* resolve to 192.0.2.12, so the weight of that server is split evenly among both names. The result is that *ns1.bar.com* is queried with 0.75 probability for *bar.com* because it also resolves to 192.0.2.13, and *ns2.bar.com* is queried with only 0.25 probability.

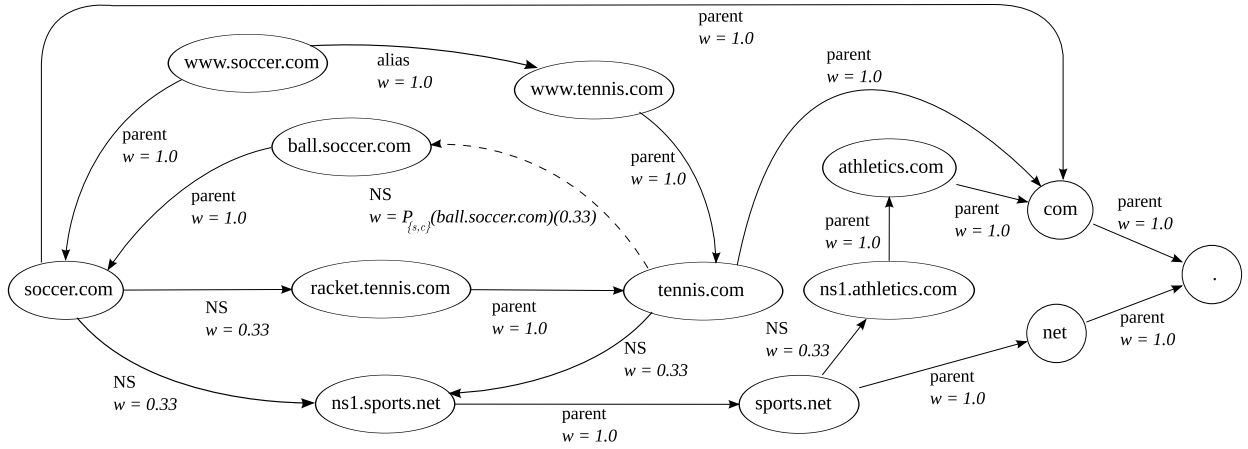


Fig. 3. The dependency graph for the domain name *www.soccer.com*, derived from the zone data in Table I. The solid lines represent active influence, and the dashed lines represent passive influence.

v subdomain of $Parent(z)$	Glue exists	$Parent(v) = z$	$w(z, v)$	Influence type	Example (Table I and Fig. 3)
no			$P_q(z, v)$	Active	<i>soccer.com</i> \rightarrow <i>ns1.sports.net</i>
yes	no		$P_q(z, v)$	Active	<i>soccer.com</i> \rightarrow <i>racket.tennis.com</i>
yes	yes	no	$P_{\{s,c\}}(v)P_q(z, v)$	Passive	<i>tennis.com</i> \rightarrow <i>ball.soccer.com</i>
yes	yes	yes	0	None	<i>soccer.com</i> \rightarrow <i>ball.soccer.com</i>

TABLE III

RULES FOR DETERMINING WHETHER OR NOT AND WITH WHAT WEIGHT $w(z, v)$ A EDGE IS PLACED BETWEEN A ZONE z AND AN NS TARGET $v \in NSA_z$.

When $NS_z \neq NS'_z$, the query probability of an edge to NS target v depends on v 's existence in NS_z and NS'_z and the likelihood that the resolver has cached the NS RRset for z from the answer or authority section of an authoritative response. We use $P_{NS}(z) \in [0, 1]$ to denote that likelihood:

$$P_q(z, v) = P_{NS}(z)P(v \in NS_z) \frac{\sum_{s \in S_v} |\{u \in NS_z | s \in S_u\}|^{-1}}{|NSA_z|} + (1 - P_{NS}(z))P(v \in NS'_z) \frac{\sum_{s \in S_v} |\{u \in NS'_z | s \in S_u\}|^{-1}}{|NSA'_z|}$$

For simplicity we assume that $NS_z = NS'_z$ unless specified otherwise.

If NS target $v \in NS_z$ is not a subdomain of $Parent(z)$, edge (z, v) is added to G_d with $w(z, v) = P_q(z, v)$. Resolution of v is required for (i.e., actively influences) resolution of z . An example is *soccer.com*'s dependency on *ns1.sports.net*.

If target name $v \in NS_z$ is a subdomain of z , the $Parent(z)$ zone should include a glue record for v . If no glue record exists for v in the $Parent(z)$ zone, then resolution of v is required for (i.e., actively influences) resolution of z , and an edge (z, v) is added to G_d with $w(z, v) = P_q(z, v)$. Such is the case with *soccer.com*'s dependency on *racket.tennis.com*.

If v is in-bailiwick, and a glue record for v exists, then resolution of v is not required for resolving z because the resolver will use the address provided from glue in the additional section of the referral from $Parent(z)$. When $Parent(v) = z$, there is no edge (z, v) in G_d ; all servers authoritative for z have the authoritative data for v , such as with *ball.soccer.com*'s relationship to *soccer.com*. However, when $Parent(v) \neq z$ an edge (z, v) is added with $w(z, v) = P_{\{s,c\}}(v)P_q(z, v)$; the name v passively influences z , dependent on the probability that either the resolver or the authoritative server has the ad-

dress for v cached from an authoritative source. An example is *tennis.com*'s dependency on *ball.soccer.com*. Passive influence is analyzed more closely in Section VI.

The active influence dependency graph, G'_d , of domain name d is the subgraph of G_d produced when $P_{\{s,c\}}(v) = 0, \forall v \in V_d$ and nodes with only zero-weight in-edges are removed from the graph. The active influence dependency graph for *www.soccer.com* would be created by eliminating the *ball.soccer.com* node in Fig. 3.

C. Level of influence

The *trusted computing base* (TCB) for domain name d is the set of all domain names in its dependency graph, G_d . An analysis of the raw size of domain name's TCB or even an enumeration of its constituents may be interesting, but insufficient for understanding the makeup of its influence. Using the well-formed dependency graph and its edges we can calculate *level of influence*, which is the probability that one domain name will be utilized for resolving another. Quantifying influence of name dependencies shows that some names may be much more influential than others. We let $I_u(v) \in [0, 1]$ denote v 's level of influence on u .

An analysis of the *dependency paths* in G_d is necessary to determine the level of influence of the domain names $v \in V_d$ on d . The dependency paths in G_d are modeled by performing a depth-first traversal of G_d , beginning with d . This depth-first traversal produces the exhaustive set of acyclic intermediate paths of name dependencies for resolving d . The level of influence is calculated by determining the probability that paths leading from d will reach v during resolution:

$$I_d(v) = P(d, \dots, v)$$

To calculate $P(d, \dots, v)$, the probabilities of encountering v in the dependency paths beginning with each of u 's direct dependencies must first be recursively calculated and aggregated. The probability of encountering v in a path beginning with edge $(u, j) \in A_d$ is calculated by multiplying the probability, $w(u, j)$, of following edge (u, j) by the probability of encountering v in a path beginning with j :

$$P(u, j, \dots, v) = \begin{cases} w(u, j) & \text{if } j = v \text{ (direct dep)} \\ 0 & \text{if } j = r \text{ (root)} \\ w(u, j)P(j, \dots, v) & \text{otherwise} \end{cases}$$

A particular domain name $u \in V_d$ may be directly dependent on multiple names, resulting in multiple out-edges from u . Some of these dependencies are mutually exclusive, and others are independent of one another. At most one NS target dependency of u is necessary for resolution (assuming the server queried is responsive), but alias and parent dependencies exist independently of the NS target dependencies. For example, when resolving names in *tennis.com* using the zone data from Table I, either *ns1.tennis.com*, *ball.soccer.com*, or *ns1.sports.net* will be selected for query, each with equal probability, and only the latter two result in a name dependency. However, resolution of u remains entirely dependent on its parent, *com*, regardless of which server is selected for query.

Aggregating the probability of encountering v in paths beginning with each of u 's direct dependencies is as follows. First the probability of encountering v through any NS-type dependencies is determined by calculating the sum of encountering it in each of the NS-type dependency edges because the probabilities are dependent on one another:

$$P(u, [NS \text{ dep}], \dots, v) = \sum_{j \in NS_u} P(u, j, \dots, v)$$

This probability is then combined independently with the probability of encountering v in paths beginning with any alias- or parent-type dependencies:

$$P(u, \dots, v) = 1 - \left(1 - P(u, Parent(u), \dots, v) \right) \left(1 - P(u, Cname(u), \dots, v) \right) \left(1 - P(u, [NS \text{ dep}], \dots, v) \right)$$

Using these expressions, we calculate the level to which *sports.net* influences *soccer.com*:

$$\begin{aligned} I_{www.soccer.com}(sports.net) &= \\ & 1 - \left(1 - P(www.soccer.com, soccer.com, \dots, sports.net) \right) \\ & \left(1 - P(www.soccer.com, www.tennis.com, \dots, sports.net) \right) \\ & \dots \\ & = 0.62 + 0.06P_{\{s,c\}}(ball.soccer.com) \end{aligned}$$

D. Dependency metrics

Calculating the individual influence of all dependencies of a domain name is computationally complex and ultimately not useful for a comparable analysis of two domain names. A more useful metric would aggregate influence of a domain name in a representative way such that can be meaningfully compared with that of other names.

```

1: procedure NONTRIVIALZONES( $d$ )
2:    $D \leftarrow \{Parent(d)\}$ 
3:   for all  $(u, v) \in A_d$  do
4:     if  $(u, v)$  is an NS target or alias dependency then
5:        $D \leftarrow D \cup \{Parent(v)\}$ 
6:   return  $D$ 

```

Fig. 4. The NONTRIVIALZONES algorithm identifies all non-trivial zones in the dependency graph for a domain name, d .

In this section we present several qualitative graph properties describing the makeup of the dependency graph for a domain name. We then use these properties to create normalized metrics for a quantitative analysis of domain names.

Influential zones: The set of zones influencing domain name d is a subset of all the influential domain names in d 's dependency graph, $Z_d \subseteq V_d$. It represents the diversity of the namespace influencing resolution of d . In Fig. 3 $Z_{www.soccer.com} = \{soccer.com, tennis.com, sports.net, athletics.com, com, net, \dots\}$.

Non-trivial zones: Non-trivial zones are the result of explicitly configured inter-zone dependencies. Included in this set are the parent zones of any NS or alias targets in A_d : $U_d \subseteq Z_d$. A non-trivial zone *foo.bar.com* that influences d may contribute up to four zones to Z_d . However, if no in-edges resulting from alias- or NS-type dependencies exist for any of its ancestor zones (*bar.com*, *com*, and “.”), then they exist in Z_d only because *foo.bar.com* is explicitly configured as a dependent zone and are thus trivial. The algorithm in Fig. 4 identifies non-trivial zones by iterating the set of edges A_d and adding the parent zones of NS and alias targets. In Fig. 3 $U_{www.soccer.com} = \{soccer.com, tennis.com, sports.net, athletics.com\}$.

First-order zones: A subset of non-trivial zones $U'_d \subseteq U_d$ are explicitly configured by d (or $Parent(d)$, if d is not a zone) and comprise *first-order zones*. This subset also includes the non-trivial zones in the ancestry of each explicitly configured zone. Fig. 5 finds all the alias (lines 5–7) and NS target (line 11) dependencies for a name d and then includes the parent zone for each target (line 15) and each non-trivial zone in its ancestry (lines 16–21). In Fig. 3 $U'_{www.soccer.com} = \{soccer.com, tennis.com, sports.net\}$.

The sizes of the sets of all, non-trivial, and first-order zones influencing a domain name can be used for a quantitative analysis. However, these metrics are unbounded and may not be entirely representative. It might be the case, for example, that a zone has a large number of influential zones, but 90% of its influence is contained within 5% of those zones. The following metrics use the preceding graph properties to quantify influence with values normalized between 0 and 1, for representative comparison.

First-order ratio: The *first-order ratio* is used to determine the percentage of non-trivial zones that are expressly configured by the administrators of d : $\frac{U'_d}{U_d}$. Values closer to 1 indicate that the administrators are largely in control of the zones comprising the TCB.

Third-party influence: *Third-party influence* (TPI) is used to quantify how much domain name d is influenced by names not explicitly controlled by its administrators, i.e., $I_d(U_d - U'_d)$.

```

1: procedure FIRSTORDERDEPS( $d$ )
2:    $N \leftarrow \text{NONTRIVIALZONES}(d)$ 
3:   /*  $M$  is the set of explicitly configured names for  $d$  */
4:    $M \leftarrow \{d\}$ 
5:   if  $d$  is not a zone then
6:     if  $d$  is an alias then
7:        $M \leftarrow M \cup \{Cname(d)\}$ 
8:      $d \leftarrow Parent(d)$ 
9:   /* Add NS target edges for zone  $d$  to  $M$  */
10:   $M \leftarrow M \cup \{u \in V_d | \exists(d, u) \in A_d, \text{NS target dep.}\}$ 
11:   $D \leftarrow \{d\}$ 
12:  /* Add non-trivial zones in  $M$ 's ancestry to  $D$  */
13:  for all  $u \in M$  do
14:     $v \leftarrow Parent(u)$ 
15:    while  $v \neq r$  do
16:      if  $v \in N$  then
17:         $D \leftarrow D \cup \{v\}$ 
18:       $v \leftarrow Parent(v)$ 
19:  return  $D$ 

```

Fig. 5. The FIRSTORDERDEPS identifies all the non-trivial zone dependencies for domain name d which are explicitly configured by DNS administrators.

Fig. 6 details the THIRDPARTYINF algorithm, which is used to calculate TPI. The TPI of d 's alias, if any (lines 5–7), is combined (lines 15–16) with the TPI of its parent zone (lines 9–10) and that of its collective NS target dependencies (lines 11–14). Two helper algorithms are utilized: the CONTROLLEDALIAS algorithm (lines 17–26) analyzes a name to determine whether or not it aliases (directly or indirectly) another name outside of the set of zones in set D . The THIRDPARTYINFD algorithm (lines 27–42) determines the probability that resolution of u will utilize a name outside the set of zones in set D .

V. DATA COLLECTION AND ANALYSIS

In this section we describe the methodology we employed for collecting data from the DNS infrastructure, and provide analysis of the data collected. We analyze several different characteristics of our data using the metrics presented in Section IV to assess quality of name resolution.

A. Data collection

We populated a database of name dependencies by crawling the namespace of known domain names. A set of over 3,000,000 hostnames was extracted from URLs submitted to the Open Directory Project (ODP) at DMOZ [23] dated April, 2009. These names were combined with over 100,000 names received as queries by the recursive servers at the International Conference for High-performance Computing, Networking, Storage and Analysis (SC08) [24]. The ODP/SC08 seed names represent both names with some intention of being resolved for access to production services (ODP names) and names actually queried by clients (SC08 names).

Each name was investigated by first surveying each name in its ancestry which had not already been surveyed, beginning with the root. Surveying a domain name consisted of issuing queries to a recursive server to receive an authoritative answer for any matching NS, MX (mail exchange) and CNAME RRs. The relationships between the name and any corresponding targets returned were recorded and subsequently surveyed.

```

1: procedure THIRDPARTYINF( $d$ )
2:    $D \leftarrow \text{FIRSTORDERDEPS}(d)$ 
3:    $P_A \leftarrow 0$ 
4:   if  $d$  is not a zone then
5:     /* If  $d$  is an alias, calculate the TPI of  $Cname(d)$  */
6:     if  $d$  is an alias then
7:        $P_A \leftarrow \text{THIRDPARTYINFD}(Cname(d), D)$ 
8:      $d \leftarrow Parent(d)$ 
9:   /* Calculate the TPI of  $Parent(d)$  */
10:   $P_P \leftarrow \text{THIRDPARTYINFD}(Parent(d), D)$ 
11:  /* Calculate the TPI of each NS target of zone  $d$  */
12:   $P_{NS} \leftarrow 0$ 
13:  for all  $u \in V_d | \exists(d, u) \in A_d, \text{NS target dep.}$  do
14:     $P_{NS} \leftarrow P_{NS} + w(d, u)\text{THIRDPARTYINFD}(u, D)$ 
15:  /* Aggregate the TPI of all name dependencies */
16:  return  $1 - (1 - P_P)(1 - P_A)(1 - P_{NS})$ 

17: procedure CONTROLLEDALIAS( $u, D$ )
18:    $H \leftarrow \{u\}$ 
19:   while  $u$  is an alias do
20:     if  $Parent(Cname(u)) \notin D$  then
21:       return False
22:     else if  $Cname(u) \in H$  then /* Loop detected */
23:       return True
24:      $H \leftarrow H \cup \{u\}$ 
25:      $u \leftarrow Cname(u)$ 
26:   return True

27: procedure THIRDPARTYINFD( $u, D$ )
28:   if  $u$  is not a zone then
29:     /*  $u$  aliases a name outside of  $D$  */
30:     if  $\neg \text{CONTROLLEDALIAS}(u, D)$  then
31:       return 1.0
32:      $u \leftarrow Parent(u)$ 
33:    $P \leftarrow 0$ 
34:   /* Aggregate influence outside  $D$  for  $u$ 's ancestors */
35:   while  $u \neq r$  do
36:      $P_u \leftarrow 0$ 
37:     for all  $v \in V_d | \exists(u, v) \in A_d, \text{NS target dep.}$  do
38:       if  $Parent(v) \notin D$  or  $\neg \text{CONTROLLEDALIAS}(v, D)$  then
39:          $P_u \leftarrow P_u + w(u, v)$ 
40:      $P \leftarrow 1 - (1 - P)(1 - P_u)$ 
41:      $u \leftarrow Parent(u)$ 
42:   return  $P$ 

```

Fig. 6. The THIRDPARTYINF algorithm returns the third-party influence for domain name d . The CONTROLLEDALIAS helper algorithm returns False if domain name u directly or indirectly aliases a name outside of a set of zones, D ; otherwise it returns True. The THIRDPARTYD helper algorithm returns the influence on domain name u by names outside of a set of zones, D .

For each NS RRset, we checked the consistency between parent and child zones by using some extra probing. For zone z we identified the set of servers authoritative for $Parent(z)$ but not z . We issued a query for z to each server in this set, anticipating a referral response without the authoritative answer (AA) flag set. Only if the AA flag was not set could we be sure that the NS RRset returned reflected the delegation records maintained in the $Parent(z)$ zone. A mismatch between the NS RRset returned in the referral and the authoritative NS RRset reflected an inconsistency.

The TTL values of additional address records corresponding to targets of NS RRs in the authority section of server responses were used to identify the presence of glue records in the parent zone. When a referring server returned a non-authoritative response, we issued a second query after a two-second delay (both without the recursion-desired flag set). Since TTL is measured in seconds, the two-second delay

between queries resulted in a decreasing TTL for RRsets in the additional section of the response, if populated from the cache of the referring authoritative server. Since a cached RRset from an authoritative response would take precedence over any glue record, we were unable to determine the existence of a glue record and optimistically assumed one existed, if the NS target was in-bailiwick. If for an NS target there was no corresponding address record provided in the additional section, we inferred that the parent zone was not configured with a glue record.

If the TTL value of an additional record does not vary between the two referral responses, it could indicate one of several things:

- the parent zone is configured with a glue record for the additional record;
- s is (also) authoritative for the zone to which the additional record belongs; or
- s is authoritative for an ancestor of the NS target and has been configured with a glue record for that NS target.

We assume optimistically in this case that if the NS target is in-bailiwick, a glue record is present.

If we were unable to elicit a non-authoritative response with a referral to z from servers authoritative for $Parent(z)$, then we could not determine inconsistencies between NS'_z and NS_z , and their corresponding glue records. However, in practice, if the set of servers authoritative for $Parent(z)$ are a subset of those authoritative for z , then consistency is satisfied implicitly since all servers will send the authoritative records from z over corresponding records from $Parent(z)$ [8]. For all zones in our analysis we assumed a 0.5 likelihood that a resolver had in its cache the authoritative NS RRset for z (i.e., $P_{NS}(z) = 0.5$), such that NS target names in both NS_z and NS'_z were considered for server selection with equal probability.

The results from our analysis are summarized in Table V.

One interesting caveat to our analysis is that the authoritative server names for several top-level domains (TLDs) are not subdomains of the TLDs for which they are authoritative. For example, the authoritative servers for *com* are in the *gld-servers.net* zone. These server names don't result in active influence on affected TLDs because they are in-bailiwick (subdomains of the root zone) and glue records exist. However, these TLDs are subject to passive influence if resolvers have A RRsets in their caches from authoritative responses.

As part of our analysis we evaluated TCB size and third-party influence with different values of $P_{\{s,c\}}(v)$, the likelihood that an authoritative A RRset for any NS target v exists in the cache of the resolver or the referring authoritative server. The third-party influence for subdomains of affected TLDs was always at least the value of $P_{\{s,c\}}(v)$ in such cases. When $P_{\{s,c\}}(v) = 1$, for example, any domain names ending in *com* or *edu* had a third-party influence of 1.0.

While the passive influence of such TLDs is legitimate, its presence in our analysis obscures the level of passive influence incurred due to configuration of subdomains. Because the naming convention of TLD authoritative servers is deliberate, and the effects of passive influence on TLDs are likely negligible, we did not consider their passive influence, with

Measurement	Values
ODP/SC08 hostnames	3,167,594
Total domain names collected	8,439,927
Total zones	2,996,460
NS target dependencies	6,855,379
NS targets requiring glue	3,723,203 (54%)
NS targets missing required glue	901 (0.024%)
Zones for which $NS_z \neq NS'_z$	587,865 (20%)

TABLE V

A SUMMARY OF RESULTS COLLECTED FROM SURVEYING THE DNS NAMESPACE, SEEDED WITH ODP/SC08 HOSTNAMES.

Metric	$P_{\{s,c\}}(v)$	Avg.	Max.
Influential zones	0	5.26	72
Influential zones	> 0	16.53	180
Non-trivial zones	0	2.26	45
Non-trivial zones	> 0	11.65	146
First-order zone ratio	0	0.92	1.0
First-order zone ratio	> 0	0.63	1.0
Third-party zone influence	0	0.08	1.0
Third-party zone influence	0.5	0.38	1.0
Third-party zone influence	1.0	0.55	1.0
Non-trivial organizations	0	1.67	37
Non-trivial organizations	> 0	8.41	113
Third-party organization influence	0	0.04	1.0
Third-party organization influence	0.5	0.34	1.0
Third-party organization influence	1.0	0.49	1.0

TABLE VI

TCB AND INFLUENCE STATISTICS FOR THE ODP/SC08 HOSTNAMES.

the exception of country-code TLDs (e.g., *us*, *fr*), which we felt might be more interesting because influence indicates the potential crossing of political boundaries.

B. Trusted computing base

The raw size of the TCB for hostnames collected in terms of influential zones and non-trivial zones is shown in Fig. 7 as a cumulative density function (CDF), and the statistics are shown in Table VI. Nearly all hostnames have a TCB smaller than 20 zones when $P_{\{s,c\}}(v) = 0$, and the average size of the TCB was 2.26 non-trivial zones and 5.26 total zones—both of

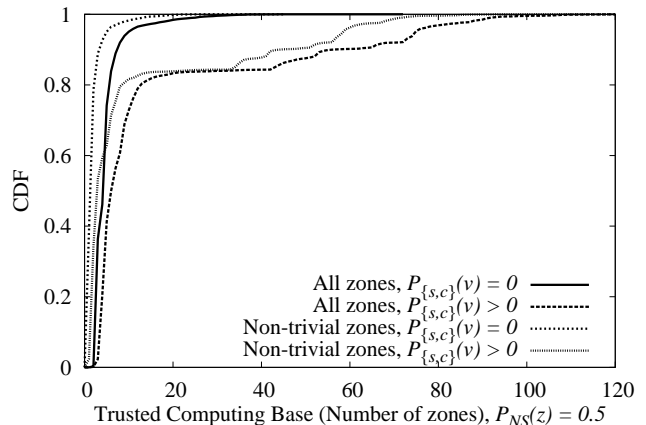


Fig. 7. The CDF for the size of the TCB of ODP/SC08 hostnames. Included are the CDF for the number non-trivial and total zones in the TCB, for $P_{\{s,c\}}(v) = 0$ and $P_{\{s,c\}}(v) > 0$.




Fig. 8. The CDF for the first-order ratio of the ODP/SC08 hostnames.

which are reasonably small. When $P_{\{s,c\}}(v) > 0$, the average size of the TCB increases several times to 11.65 non-trivial and 16.53 total zones. Only about 80% have fewer than 20 zones; most of the remaining 20% have between 30 and 90 non-trivial and total zones in their TCB. Caching and using NS target names from authoritative sources, rather than glue, can increase the size of the TCB of a domain by several times.

C. Controlled influence

The first-order ratio of the ODP/SC08 hostnames is shown in Fig. 8. The average first-order ratio was 0.92 for $P_{\{s,c\}}(v) = 0$ and 0.63 for $P_{\{s,c\}}(v) > 0$, indicating that control of the TCB is lost as caching of NS target names is introduced. When $P_{\{s,c\}}(v) > 0$, third-party zones comprise more than half of the the non-trivial zones in the TCB of roughly 40% of the hostnames surveyed.

Fig. 9 shows the third-party influence of the ODP/SC08 hostnames. When $P_{\{s,c\}}(v) = 0$, 85% of the hostnames are not influenced at all by third parties. At $P_{\{s,c\}}(v) = 0.5$ only 60% of the hostnames are influenced less than 50% by third parties. When $P_{\{s,c\}}(v) = 1$ nearly half of the hostnames are influenced almost certainly by third parties. Again the behavior of caching preference of NS target names from authoritative sources at the resolver and authoritative servers greatly affects third-party influence of domain names.

D. Alias chains

One behavior affecting the third-party influence of domain names is the practice of chaining aliases. As an example, *www.example.com* is configured as an alias for *www.example.net* which, in turn, is an alias for *www.example.org*. The practice is common among content distribution networks for offloading the burden of addressing from the administrators of the original name (e.g., *www.example.com*). In our survey, we found that 33,873 (1%) of the ODP/SC08 names were affected by alias chains. Using that subset of affected names, we graphed the third-party influence for analysis in Figure 10 using $P_{\{s,c\}}(v) = 0.5$. Only 10% of the names dependent on alias chains exhibit

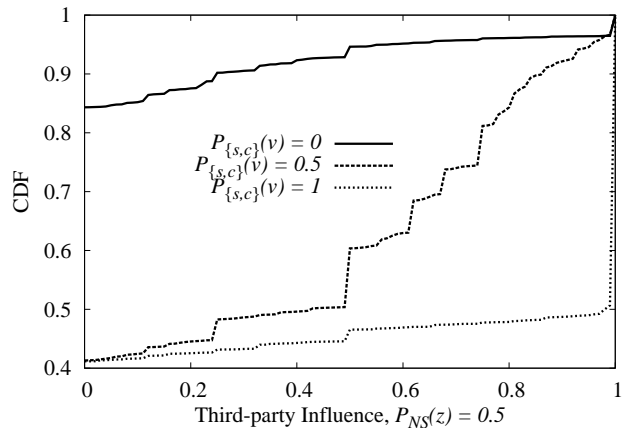


Fig. 9. The CDF for the third-party influence of the ODP/SC08 hostnames, grouped by zone.

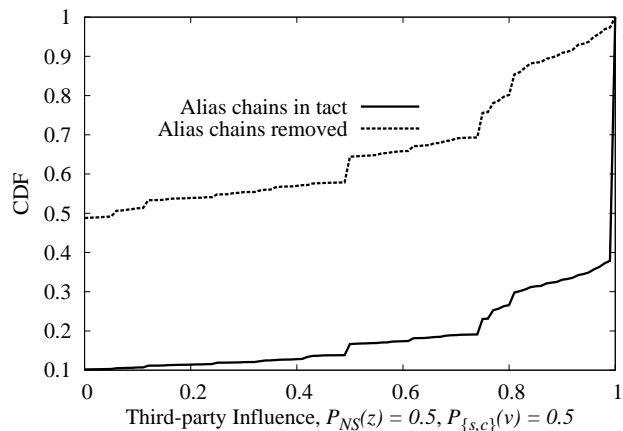


Fig. 10. The CDF for the third-party influence of ODP/SC08 hostnames affected by alias chains, grouped by zone.

no third-party influence, compared to 42% of the entire set of ODP/SC08 names. However, when we removed the effects of the alias chains in our analysis, nearly half of the names avoided third-party influence completely, even with $P_{\{s,c\}}(v) > 0$.

E. Organizational dependencies

Throughout Sections IV and V we have used zones as the unit of measurement for describing properties of the DNS name dependency graph. In reality the grouping of names or servers for analysis is arbitrary, and the methodology employed depends on certain assumptions and the desired results. Using zones as a unit of measurement may seem reasonable because of the assumption that security and reliability across the servers whose names are in a common zone are consistent. A similar assumption may hold for use of administering organizations as a unit of measurement. Analysis of domain name influence may yield different results when zones are grouped by organization. For example, *sports.com* and *athletics.net* (Table I) are both run by Sports Central, Inc. Although the *athletics.net* zone is a third party to *www.soccer.com*, when

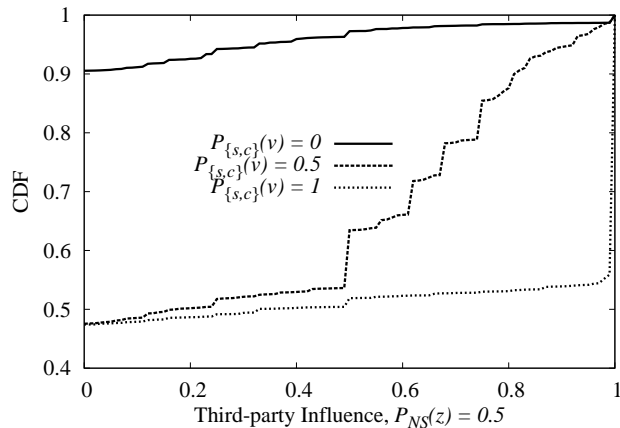


Fig. 11. The CDF for the third-party influence of ODP/SC08 hostnames, grouped by organization.

analyzed by organization it is grouped with *sports.com* under Sports Central, Inc., a first-order dependency.

Every zone is configured with a point-of-contact email in its SOA (start-of-authority) RR. To further our analysis, we used the domain from the email as a method for grouping different zones into a single organization, under the assumption that the email suffix would reflect the organization responsible for maintenance. The results are shown in Figure 11. Using this organizational grouping, we found that the number of domain names with no third-party influence increased from 85% to over 90% when $P_{\{s,c\}}(v) = 0$ and from 42% to 48% when $P_{\{s,c\}}(v) > 0$.

VI. PASSIVE INFLUENCE CONSIDERATIONS

We have observed an increased third-party influence on dependent domain names when there is a higher likelihood of passive influence caused by the existence of an A RRset from an authoritative response in cache for an NS target name for which a glue record exists. This likelihood, $P_{\{s,c\}}(v)$, is comprised of the probability that the A RRset from authoritative source is in the cache of the authoritative server, s , offering the referral or in the cache of the resolver, c , $P_s(v)$ and $P_c(v)$, respectively. We discuss in this section considerations that affect passive influence, as a resource to contain third-party influence with proper configuration.

A. Authoritative servers

The obvious way to minimize $P_s(v)$ is to remove caching functionality from authoritative servers, a DNS best practice [9]. Both NSD and *tinydns* act as authoritative servers without caching. BIND allows dual functionality but provides a configurable option to disallow the inclusion of cached RRsets in the additional section of responses. However, this option is not enabled by default, so if a BIND authoritative server with default settings has records in cache, it will use them to populate the additional section in preference to any glue records in its local configuration.

During our survey of the production DNS space we queried 7,781 distinct servers authoritative for a parent zone in an attempt to detect glue records and discrepancies in the delegation records for a child zone. We found that 1,486 (19%) of these servers returned additional records whose TTL decreased with subsequent requests, indicating that they had come from the cache of the authoritative server. This demonstrates the pervasiveness of caching functionality on production authoritative servers.

B. Resolvers

Many factors may contribute to an increased probability of RRsets from authoritative sources in a resolver's cache, $P_c(v)$. Some are environment-specific. For example, the rate of queries to the resolver requesting an authoritative answer for the A RRset of an NS target name would certainly affect this probability, but this behavior is local to the resolver's environment. One consideration that has effects in any environment is the dynamics of TTL values for related RRsets.

An ideal configuration is one in which the TTL of a glue record is equal to that of the NS RRset for which it is an NS target. When the TTL of the glue record is less than that of the NS RRset, there is an increased chance that a queries for names within the zone will result in an induced query for an authoritative response for the A RRset.

As a practical example, suppose the TTL of the NS RRset for *tennis.com* (from Table I) is 160 and the TTL for the glue record for *ball.soccer.com* is 80, and suppose that a resolver caches both the NS RRset and the A RRset (from the additional section of the response) at time $t = 0$ (We illustrate this example in Fig. 12). For lookups in *tennis.com* during $0 < t \leq 80$ the resolver will have the address(es) necessary to query authoritative server *ball.soccer.com* using the glue record from the *com* zone. When $80 < t \leq 160$ the *tennis.com* NS RRset is still in the resolver's cache, but now a lookup is required to obtain the address for *ball.soccer.com*, and the answer received from an authoritative source will take the place of the expired glue record. Such is the case with the query at $t = 110$.

Not until $t > 160$, when the *tennis.com* NS RRset expires and it is required to re-query the *com* servers, does the resolver again receive a glue record for *ball.soccer.com* in the additional section of a response. The query at $t = 170$ in Fig. 12 causes this lookup. However, as explained in Section II, records from authoritative sources are considered more trustworthy, so the records received from the additional section do not override the records already in cache. Thus at $t = 260$, when the A RRset from authoritative source has expired, the query induces another authoritative lookup, and passive influence persists until both the NS and the A RRsets expire from cache. After both RRsets have expired, the query at $t = 350$ induces a query by the resolver resulting in a referral response with the NS RRset and the A RRset from glue, which will both be used for subsequent queries.

We observed this pattern of passive influence in both the BIND and unbound resolvers. However, the behavior exhibited by *dnscache* varies because of its cache renewal




Fig. 12. A timeline of queries to a resolver for names within a particular zone and which are not found in the resolver’s cache. The labels on top of the timeline indicate the cache lifetimes of the NS RRset for the zone and the A RRset from the glue and authoritative answers. The ticks on the bottom mark the times at which queries for uncached names in the zone arrive at the resolver, inducing queries for NS or A RRsets, as appropriate. Queries at the circled times represent queries for authoritative A RRsets, and shading along the bottom of the timeline highlights the time in which authoritative A RRsets exist in cache.

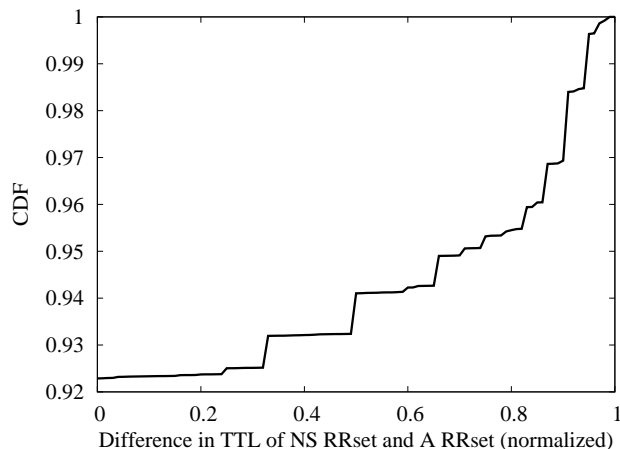


Fig. 13. The CDF showing the normalized difference between NS and A RRsets for NS-type dependencies in the data seeded by the ODP/SC08 names.

policy. BIND and unbound do not renew the lifetime of an NS RRset that has been cached from the authority section of an authoritative response, even when subsequent responses are received with the same contents. However, every time `dnscache` receives a response, it updates its cache with RRsets from all sections. This means that RRsets from the additional section of responses will replace any from the answer section of authoritative responses. It also means that a resolver never needs to obtain a referral from the *com* servers as long as the NS RRsets in its cache are renewed from the authority section of authoritative responses from the *example.com* servers. This two-edged sword means that passive influence may enter less frequently on a `dnscache` resolver, but it may persist longer, depending on query rates.

We note that the presence of authoritative data for an NS target in cache does not necessarily mean that a resolver selects the respective server for query—only that the source of the address is now a product of passive influence, and that if selected, it exhibits dependencies of that name, following the model described in Section IV. Also, our example assumes for simplicity that only queries for names in *tennis.com* are received by the resolver during the time span shown.

We define the contribution that TTL dynamics have on cache probability at a resolver as follows. Let $v \in NS_z$ be an NS target of zone z , such that v is in-bailiwick, $Parent(v) \neq z$, and a glue record for v exists. We use the difference in TTL value between the NS RRset for z and the A RRset for v as a measure of the impact on cache probability,

$P_c(v)$. This difference represents the window in which a resolver has the NS RRset for z cached, but the A RRset has expired, following an initial query for z . We normalize this value by dividing it by the TTL of the NS RRset. For simplicity, we assume that the TTLs for delegation and glue records in the parent zones are the same as the corresponding RRsets in the child zone.

Fig. 13 shows a CDF of the normalized difference between authoritative NS and glue A RRsets for the 2,713,903 NS target dependencies meeting the criteria for potential passive influence, which comprised 40% of the total NS target dependencies in our data. Over 92% of the dependencies are configured such that there is no difference between the TTLs of the NS and A RRsets. But for nearly 6% of the dependencies, the TTL of the NS RRset is two or more times that of the A RRset for the corresponding target.

VII. CONCLUSION

In this article we have presented a graph model for analysis of name dependencies in DNS, which was based on specification and behavior of deployed DNS servers. We defined the trusted computing base (TCB) of a domain name in terms of zones and organizations. We also derived metrics for assessing the dependency model of a domain name. Among these were the level of influence of influential domain names, and third-party influence—the probability that resolution of a domain name will utilize namespace outside the explicit configuration of domain administrators.

We observed that the TCB of domain names, when measured by influential zones and organizations, is much smaller than previously thought. On average 92% of the non-trivial zones in the TCB of a domain name were explicitly configured by the domain administrators. However, the practice of resolver and authoritative server using the address records for corresponding to NS targets from cache, rather than from additional records in a response or from glue, can increase the size of the TCB and the influence of third-party namespace significantly.

To maximize the reliability of name resolution from the perspective of both resolver and authoritative server, administrators and designers of DNS services should be aware of their server configurations as well as the names and organization comprising the TCBs of their domain names. Administrators should review the role of name servers in their environment to minimize the influence of third parties. The practice of chaining domain name aliases increases the potential for

third-party influence and should be avoided, as suggested in RFC 1912 [10]. Additionally, we recommend that the roles of authoritative server and caching server be kept distinct or that authoritative servers be configured to not include information from their caches, as this creates additional dependence. Also, DNS administrators should coordinate the TTL values of NS RRsets with the TTL values of their NS targets.

A better understanding of DNS dependencies and an application of that understanding will put more control into the hands of DNS administrators and mitigate the risks associated with large and diverse TCBs and high third-party influence.

REFERENCES

- [1] P. Mockapetris, "RFC 1034: Domain names - concepts and facilities," 1987. [Online]. Available: <http://tools.ietf.org/html/rfc1034>
- [2] —, "RFC 1035: domain names - implementation and specification," 1987. [Online]. Available: <http://tools.ietf.org/html/rfc1035>
- [3] P. Mockapetris and K. J. Dunlap, "Development of the domain name system," vol. 18, no. 4. New York, NY, USA: ACM, 1988, pp. 123–133.
- [4] ISC BIND. [Online]. Available: <http://www.isc.org/products/BIND/>
- [5] Unbound. [Online]. Available: <http://www.unbound.net/>
- [6] djbdns. [Online]. Available: <http://cr.yt.to/djbdns.html>
- [7] NSD. [Online]. Available: <http://www.nlnetlabs.nl/projects/nsd/>
- [8] R. Elz and R. Bush, "RFC 2181 - clarifications to the DNS specification," 1997. [Online]. Available: <http://tools.ietf.org/html/rfc2181>
- [9] R. Chandramouli and S. Rose, "Secure domain name system (DNS) deployment guide," National Institute of Standards and Technology. [Online]. Available: <http://csrc.nist.gov/publications/nistpubs/800-81r1/sp-800-81r1.pdf>
- [10] D. Barr, "RFC 1912 - common DNS operational and configuration errors," 1996. [Online]. Available: <http://tools.ietf.org/html/rfc1912>
- [11] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "RFC 4034: Resource records for the DNS security extensions," 2005. [Online]. Available: <http://tools.ietf.org/html/rfc4034>
- [12] V. Ramasubramanian and E. G. Sirer, "Perils of transitive trust in the domain name system," in *IMC '05: Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, USENIX Association, Berkeley, CA, USA: USENIX Association, 2005, pp. 379–384.
- [13] V. Pappas, Z. Xu, S. Lu, D. Massey, A. Terzis, and L. Zhang, "Impact of configuration errors on DNS robustness," in *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, ACM, New York, NY, USA: ACM, 2004, pp. 319–330.
- [14] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "RFC 4033: DNS security introduction and requirements," 2005. [Online]. Available: <http://tools.ietf.org/html/rfc4033>
- [15] IKS, "DNSSEC." [Online]. Available: <https://www.iks-jena.de/leistungen/dnssec.php>
- [16] E. Osterweil, D. Massey, and L. Zhang, "Deploying and monitoring DNS security (DNSSEC)," in *25th Annual Computer Security Applications Conference (ACSAC '09)*, December 2009.
- [17] "SecSpider." [Online]. Available: <http://secspider.cs.ucla.edu/>
- [18] E. Osterweil, M. Ryan, D. Massey, and L. Zhang, "Quantifying the operational status of the DNSSEC deployment," in *Proceedings of the 6th ACM/USENIX Internet Measurement Conference (IMC'08)*, October 2008.
- [19] D. Dagon, M. Antonakakis, P. Vixie, T. Jinmei, and W. Lee, "Increased DNS forgery resistance through 0x20-bit encoding: security via leet queries," in *Proceedings of the 15th ACM conference on Computer and communications security*. ACM, 2008, pp. 211 – 222.
- [20] R. Perdisci, M. Antonakakis, X. Luo, and W. Lee, "WSEC DNS: Protecting recursive DNS resolvers from poisoning attack," in *Dependable Systems & Networks, 2009. DSN '09. IEEE/IFIP International Conference on*. IEEE, July 2009, pp. 3 – 12.
- [21] L. Yuan, K. Kant, P. Mohapatra, and C.-N. Chuah, "DoX: A peer-to-peer antidote for DNS cache poisoning attacks," in *ICC '06: IEEE International Conference on Communications*.
- [22] V. Ramasubramanian and E. G. Sirer, "The design and implementation of a next generation name service for the internet," in *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*. New York, NY, USA: ACM, 2004, pp. 331–342.
- [23] Open Directory Project. [Online]. Available: <http://www.dmoz.org/>
- [24] SC08: The International Conference for High-performance Computing, Networking, Storage and Analysis. [Online]. Available: <http://sc08.supercomputing.org/>



Dr. Casey Deccio is a Senior Member of Technical Staff at Sandia National Laboratories in Livermore, CA. He joined Sandia in 2004 after receiving his BS and MS degrees in Computer Science from Brigham Young University, and he earned his doctoral degree from the University of California, Davis, in 2010. Dr. Deccio's research interests lie primarily in modeling and availability analysis of DNS and DNSSEC.



Jeff Sedayao is an enterprise architect in Intel's IT Research Group. He focuses on applying distributed systems technologies such as cloud computing to enterprise IT problems. Jeff has participated in IETF working groups, published papers on policy, network measurement, network and system administration, and authored the O'Reilly and Associates book, *Cisco IOS Access Lists*.



Dr. Krishna Kant



Dr. Prasant Mohapatra is currently the Tim Bucher Family Endowed Chair Professor and the Chairman of the Department of Computer Science at the University of California, Davis. In the past, he has been on the faculty at Iowa State University and Michigan State University. He has also held Visiting Scientist positions at Intel Corporation, Panasonic Technologies, Institute of Infocomm Research (I2R), Singapore, and National ICT Australia (NICTA). He was/is on the editorial board of the IEEE Transactions on Computers, IEEE Transactions on Mobile Computing, IEEE Transaction on Parallel and Distributed Systems, ACM WINET, and Ad Hoc Networks. He has been on the program/organizational committees of several international conferences.

Dr. Mohapatra received his doctoral degree from Penn State University in 1993, and received an Outstanding Engineering Alumni Award in 2008. He is a Fellow of the IEEE.

Dr. Mohapatras research interests are in the areas of wireless networks, sensor networks, Internet protocols, and QoS.