

Rethinking the Security Threats of Stale DNS Glue Records

Yunyi Zhang^{†‡*}, Baojun Liu^{‡*}, Haixin Duan^{‡§Φ}, Min Zhang^{†✉}, Xiang Li[‡], Fan Shi[†]
Chengxi Xu[†], and Eihal Alowaisheq[¶]

[†]National University of Defense Technology, [‡]Tsinghua University

[§]Zhongguancun Laboratory, ^ΦQuan Cheng Laboratory, [¶]King Saud University

{zhangyyzyy, zhangmindy, shifan17, xuchengxi}@nudt.edu.cn
{lbj, duanhx}@tsinghua.edu.cn, {x-119}@mails.tsinghua.edu.cn, {ealowaisheq}@ksu.edu.sa

Abstract

The Domain Name System (DNS) fundamentally relies on glue records to provide authoritative nameserver IP addresses, enabling essential in-domain delegation. While previous studies have identified potential security risks associated with glue records, the exploitation of these records, especially in the context of out-domain delegation, remains unclear due to their inherently low trust level and the diverse ways in which resolvers handle them. This paper undertakes the first systematic exploration of the potential threats posed by DNS glue records, uncovering significant real-world security risks. We empirically identify that 23.18% of glue records across 1,096 TLDs are *outdated* yet still served in practice. More concerningly, through reverse engineering 9 mainstream DNS implementations (e.g., BIND 9 and Microsoft DNS), we reveal manipulable behaviors associated with glue records. The convergence of these systemic issues allows us to propose the novel threat model that could enable large-scale domain hijacking and denial-of-service attacks. Furthermore, our analysis determines over 193,558 exploitable records exist, placing more than 6 million domains at risk. Additional measurement studies on global open resolvers demonstrate that 90% of them use unvalidated and outdated glue records, including OpenDNS and Alibaba Cloud DNS. Our responsible disclosure has already prompted mitigation efforts by affected stakeholders. Microsoft DNS, PowerDNS, OpenDNS, and Alibaba Cloud DNS have acknowledged our reported vulnerability. In summary, this work highlights that glue records constitute a forgotten foundation of DNS architecture requiring renewed security prioritization.

1 Introduction

The Domain Name System (DNS) facilitates Internet communications by translating human-readable domain names into machine-friendly IP addresses. The recursive traversal

of a delegation chain from parent to child authoritative nameservers underpins domain name resolution, necessitating that parent domains contain delegation records referencing their subdomains' designated nameservers. However, a paradoxical loop emerges when attempting to resolve a subdomain delegated by its parent domain (e.g. `ns.foo.com` which is delegated by `foo.com`). This inherent recursion is resolved by utilizing *glue records*, which contain nameserver IP addresses stored within the delegating parent's zone file, solely used in referral responses.

The security of delegation records has been widely discussed in the security community. RFC1034 [52] states that authoritative nameserver (NS) records at both parent and child should be “consistent and remain so”, but some researchers found significant inconsistencies [4, 64]. Moreover, researchers have conducted extensive measurement studies on abandoned glue records [33, 63] and have analyzed the potential impact of unsigned glue records on DNSSEC [68]. These inconsistent and inappropriate delegation records provide potential attack vectors, allowing attackers to take over domain names by re-registering expired NS domains [3, 4] or obtaining IP addresses [14, 46].

Although some studies and blog posts [30, 46] have mentioned the potential exploitation of glue records to hijack domains, the exploitation of glue records, especially in the case of out-domain delegation, remains unclear. This is due to the differences in their use, compared to typical authoritative records, which can be directly exploited, as presented in [3, 4, 14]. In this paper, we aim to provide a comprehensive analysis of the management and utilization of glue records, evaluating realistic security threats in the wild.

Our study. Our research conducts a thorough examination of the conditions under which glue records are exploited. We propose a new attack vector called *shadow caching* to exploit the massive number of stale DNS glue records in the domain namespace, even under out-domain delegation. Specifically, through comprehensive analysis, we examine 1,096 TLD zone files alongside 9 predominant DNS implementations. Extensive interrogation of zone file data reveals that over 23.18%

* Both authors contributed equally to this work.

✉ Corresponding author.

of existent glue records are outdated yet still served in real-world DNS infrastructure. More troublingly, our software testing reveals manipulable behaviors in mainstream DNS software where they cache and utilize “unvalidated” glue records during resolution, resulting in a phenomenon we term “shadow caching”. Once attackers obtain the GlueIP, typically from cloud platforms, they can covertly hijack domains using *shadow caching* even under out-domain delegation. Moreover, if the GlueIP is unobtainable, they can still launch a denial-of-service attack against the target. We have demonstrated the widespread nature of this threat, impacting up to 6M domains.

We empirically assessed the real-world exploitability of stale records utilizing our proposed threat models. The analysis identified 193,558 (9.60%) potentially exploitable stale glue records among the 2,016,516 used glue records examined. Alarming, 39,795 (20.56%) of these resided within the Tranco Top 1M domains [57], indicating serious exposure among prominent websites. Overall, we detected 6,398,631 domains vulnerable to hijacking attacks and 784,693 susceptible to denial-of-service, with 496,324 affected by both vectors. Further examination revealed 5,700 impacted domains lie within the Tranco Top 1M, including high-profile companies (e.g., trueconf.net) and ISPs (e.g., i2bnetworks.com), empirically proving the ability to disrupt major providers. Experimental confirmation validated that all tested DNS software and 14 major public resolvers are susceptible, including widely used providers such as BIND [13], Microsoft DNS [51], OpenDNS [15] and Quad9 [59].

We also evaluated the real-world impact of stale records by conducting measurements on open DNS resolvers, selecting only representatives exhibiting stability over two months owing to high churn rates [39]. Over 90% of the tested resolvers exhibited manipulable behaviors with glue records, rendering them susceptible to hijacking exploits. Furthermore, 111,766 resolvers (12.48%) were confirmed to be vulnerable to denial-of-service attacks abusing stale records. Our results provide empirical evidence that the vast majority of DNS resolvers operating in practice are prone to exploitation of stale glue record vulnerabilities. This highlights the seriousness and scale of the threats that have been overlooked by the technical community thus far.

Disclosure and mitigation. We responsibly disclosed the discovered issues to affected TLD registries, resolver vendors, and DNS software. The .info and .org registry confirmed the issue and began remediating impacted records. Also, we reported the issue to ICANN and are collaborating to assist other registries in resolving this problem. Among DNS providers, Microsoft DNS, PowerDNS, OpenDNS, and Alibaba Cloud DNS acknowledged the threat and have either deployed or are implementing mitigation based on our reports. Moreover, we disclosed the threat to affected domains through our national CERT. This responsible disclosure and ongoing collaboration will raise community awareness and drive the adoption of critical defenses against the overlooked risks of




Figure 1: DNS resolution process.

stale glue records.

Contributions. Our contributions are outlined as follows:

Systematic analysis of glue records. A systematic analysis of real-world glue record usage from the lens of domain configurations across 1,096 TLDs and 9 major DNS software.

Novel attack. Exploiting our proposed novel attack vector, *shadow caching*, we propose the new exploitation method for stale glue records, especially under out-domain delegation, enabling domain hijacking and denial-of-service attacks.

Comprehensive evaluation of new attacks. Comprehensive threat evaluation proving over 6 million domains are vulnerable. We empirically demonstrate 90% of stable open resolvers and 14 major public DNS providers are susceptible to the newly proposed attacks.

2 Background

DNS resolution process. The domain name space is a tree structure. Each node and leaf on the tree corresponds to a resource set, with the higher and lower zones called parent and child zones. At the top of the DNS tree is the *root* zone, whose child zones are divided into a collection of Top-Level Domains (TLDs), like .com, and .net. One step down, the child zone of the .com zone is a set of Second-Level Domains (SLDs), such as exam.com. and foo.com.

As shown in Figure 1, the DNS recursive resolver first contacts the *root* zone when it receives a query request from a client (step 1 and 2). Since a DNS zone only contains information about its child zones, the *root* returns a *referral response* that informs the resolver of information of the authoritative nameserver for .com (step 3). Iterating through this process (step 4 and 5), the resolver eventually retrieves the resource records from the authoritative nameserver of exam.com. (step 6) and returns it to the client.

DNS glue record. The DNS resource record (RR) format defined by RFC 1034 [52] is a 5-tuple: $\langle owner, type, class, TTL, RDATA \rangle$. *Owner* describes the domain name that the RR belongs to, *Type* and *RDATA* represent the kind of RR and its value, such as the NS record for the authoritative nameserver and A record for IPv4 IP address.

Glue records are a specific A record in the DNS zone. Theoretically, for the DNS zone, the parent zones can provide all information that is successful in accessing their children’s

(1) in-domain delegation	(2) sibling-domain delegation	(3) out-domain delegation
<pre>;; NS RR test.com. NS ns1.test.com. ;; glue record in .com zone ns1.test.com. A g.l.u.e</pre>	<pre>;; NS RR test.com. NS ns1.foo.com. ;; glue record in .com zone ns1.foo.com. A g.l.u.e</pre>	<pre>;; NS RR test.com. NS ns1.test.net.</pre>
referral response	referral response	referral response
<pre>;; AUTHORITY SECTION: test.com. NS ns1.test.com. ;; ADDITIONAL SECTION: ns1.test.com. A g.l.u.e</pre>	<pre>;; AUTHORITY SECTION: test.com. NS ns1.foo.com. ;; ADDITIONAL SECTION: ns1.foo.com. A g.l.u.e</pre>	<pre>;; AUTHORITY SECTION: test.com. NS ns1.test.net.</pre>

Figure 2: Domain delegation categories.

zones. It is, however, difficult to do this using only the NS RR for children zones. We may encounter a situation where a parent domain is delegated to its child domain (i.e., in-domain delegation), which means that the recursive resolver has to initial the query from the parent domain, causing the DNS resolution loop problem. To fix it, RFC 1034 introduces the glue record, which allows to add non-authoritative data (i.e., the IP address of the NS RR) in the zones [52]. For convenience, we will refer to the domain name of authoritative nameservers in the glue record as GlueFQDN (e.g., `ns1.test.com.` in Figure 2), and the corresponding IP address as GlueIP (e.g., `g.l.u.e` in Figure 2).

In implementation, the provisioning and management of domain names and nameserver delegation records are standardized via the Extensible Provisioning Protocol (EPP) [25–27]. Registrants are not able to interact directly with the EPP servers and can only configure the glue records via the limited interfaces provided by the registrars. Not all registrars support the configuration of glue records; this is determined by the registrars themselves. However, it is worth noting that some famous registrars do provide support for glue records, like GoDaddy [22], and Aliyun [6].

In addition, RFC 1034 requires that glue records are only used as part of a referral response [52]. To do this, depending on the source of the data, different levels of trust are assigned to the RRs, with the glue records having the lowest level of trust [21]. The authoritative data in the `answer` section of the authoritative response is the most trusted, while the trust level of glue records is far below it. DNS software relies on the trust level of resource records to determine which records can be returned to users as `answer` and which records can be cached or updated. Moreover, Li et al. [43] showed that the implementation of record trust levels by different DNS software does not follow the RFC requirements exactly, but the common denominator is a low trust level for glue records.

Categories of domain delegation. The domain administrator delegates their domains to a specific server by configuring the delegation NS RR and glue RR (if needed). According to the relationship of NS RR and the domain name, there are three types of delegation, in-domain, sibling-domain, and out-domain delegation. Below, we describe each category of

delegation and their corresponding referral response in detail.

In-domain delegation configures the in-domain nameservers that are contained in the delegated zone itself [50], where the glue records are necessary; otherwise, resolution loops will occur. The configuration and referral response for `test.com.` with glue records are shown in Figure 2 (1).

Sibling-domain delegation configures nameservers that are not contained in the delegated zone itself, but in another zone delegated from the same parent [50], where glue records are not required, as shown in Figure 2 (2).

Out-domain delegation sets nameservers that cross the TLD, where glue records are not required. The referral response for `test.com.` that delegated to `ns1.test.net.` is shown in Figure 2 (3).

Summary: Under in-domain delegation, glue records are necessary, otherwise the resolution chain will be broken. For sibling domain delegation, glue records are not necessary. However, the referral response usually contains the glue record. Thus, resolvers typically utilize these glue records to improve resolution efficiency. For out-domain delegation, the resolver cannot obtain glue records from the referral response and is required to actively resolve the GlueFQDN.

3 Characterizing the Use of DNS Glue Records

Due to the special nature of DNS glue records, they are not used directly, unlike authoritative records. They are exclusively used in referral responses and are assigned low trust rank. This unique characteristic makes the exploitation of glue records more challenging.

In this section, we investigated the use of glue records in TLD zone files and DNS software to determine whether the problem is caused by incorrect or manipulable usage.

3.1 Stale Glue Records in Zone Files

To evaluate the use of glue records within zone files, we begin by introducing our dataset and describing the methodology employed to identify stale glue records. Then, we analyze the distribution of glue records in the two largest TLD `.com` and `.net` and quantify stale glue records through our approach.




Figure 3: Statistics information of glue records.

Table 1: Overview of glue records.

Total glue ¹	2,827,798	GlueFQDN status ²	2,016,516 (100%)
GlueFQDN	2,518,786	NOERROR	1,438,964 (71.36%)
SLD	975,168	SERVFAIL	385,816 (19.13%)
Used glue	2,283,196	NXDOMAIN	191,160 (9.48%)
Stale glue	529,197	REFUSED	576 (0.03%)

¹: A single GlueFQDN can correspond to multiple IP addresses. Thus, the number of GlueFQDN is less than that of glue records.

²: Only glue records served domains are being considered.

Moreover, we compare differences between legacy and latest IP addresses to reveal the IP changing of stale glue records. Finally, we exhibit the overall delegation distribution.

Dataset. For a comprehensive evaluation of glue records, we downloaded authorized zone files for 1,096 TLDs on August 10, 2023, through ICANN’s CZDS [29]. The zone file of a TLD encompasses delegation information for all its associated SLDs, while concurrently housing glue records, DNSSEC signatures, and other pertinent details. From these zone files, we collected a total of 2,827,798 glue records, represented as $\langle \text{FQDN}, \text{IP} \rangle$ pairs, with 2,518,786 distinct GlueFQDNs, spanning across 975,168 SLDs that correspond to 954,971 GlueIPs. Specifically, 2,283,196 glue records (80.74%) are in use, with the rest abandoned and unused by any domain. Note that we have excluded the abandoned glue records from subsequent analysis. Our active resolution results for GlueFQDNs show that nearly 30% of GlueFQDNs cannot obtain IP addresses, as shown in Table 1. Note that an abnormal resolution status doesn’t necessarily invalidate a GlueFQDN. For instance, if the resolution response for a GlueFQDN is NXDOMAIN, it might be due to the domain name owner forgetting to configure the A record of the GlueFQDN in their nameservers. Nevertheless, the GlueFQDN remains valid for in-domain and sibling-domain delegation. Hence, we need a method to determine the real stale glue records.

Methodology. In this paper, we define *stale glue records* as glue records that persist in the zone file but no longer fulfill their domain resolution function. There are several reasons

for the occurrence of stale glue records, such as when domain owners change their authoritative server’s IP address or when that authoritative server is deprecated. As such, our first step is to identify those glue records where the GlueIP does not align with the A record on their authoritative nameservers. Specifically, we begin with extracting all glue records from the zone files. Subsequently, we resolve these GlueFQDNs actively to collect their latest IP addresses. To enhance the precision of our data collection, we obtain data from three distinct measurement points (Hong Kong, Dubai, and Virginia). At each of these points, data is collected independently, and subsequently, the union of the resource records is derived. Next, we compare the actively obtained IP addresses with those retained GlueIP in the zone files. Glue records exhibiting discrepancies in these results are classified as potential stale glue records.

Then, we check if these glue records continue to provide services for the domains delegated to them. For the glueFQDN with more than 100 domains, we randomly select $\max(1\%, 100)$ as test domains. For those serving less than 100, we test them all. We believe that testing only 1% (100) of the domains when the nameserver is serving a large number of domains provides a representative indicator of its performance while minimizing unnecessary load. This approach also helps reduce the influence of external factors like network instability or potential DNS censorship, ensuring a more accurate assessment of the nameserver’s responsiveness. We only consider the glue record as stale if we did not receive a response for all test domains, suggesting it is inactive.

In summary, we identify stale glues by filtering out active records, abandoned glue records, and records that are consistent with authoritative records. The configuration operation of glue records is complicated, and most require registrants to configure manually at the registrar. Thus, the identified stale glue that does not provide any services, fails to meet the requirements of specific configurations, such as load balancing.

Glue record distribution. The use of glue records is very

prevalent, with 75.27% of NS records in the zone files being configured with glue records. We conducted a statistical analysis on the NS configuration in the two largest TLDs .com and .net, as shown in Figure 3a. We observed that 80% of domains are configured with 2 NS records, yielding an average of 2.42 NS per domain. In addition, we discovered that 95% of the glue records are associated with fewer than 5 IP addresses. This is consistent with common load-balancing strategies, where large enterprises configure multiple primary and secondary servers to provide redundant services, thereby enhancing resolution efficiency and disaster recovery capabilities. However, certain IP addresses are associated with even more than 10,000 glue records, which could be attributed to specific operations by certain registrars [4].

Stale glue records. We found, surprisingly, that 23.18% of glue records are stale. Utilizing our approach, we identified 529,197 stale glue records (504,851 GlueFQDNs) out of 2,283,196 (2,016,516 GlueFQDNs) actively used. Then, we analyzed the resolution status of these stale glues and the number of domains they serve. In Figure 3b, the bar illustrates the stale glue records across various resolution states (primary Y-axis), while the line depicts the count of domains they serve (secondary Y-axis). Stale glue records in a normal resolution status support the majority of affected domains. We deduced that these stale glue records have transitioned to new service IP addresses, yet the historically deprecated IPs remain preserved in zone files. A domain name can be associated with multiple nameservers; if even one operates correctly, the domain can be resolved normally. Subsequently, we quantified the proportion of stale glue records in each domain’s nameservers, as depicted in Figure 3c. The findings indicate that, although some domains depend on both normal and stale records, a significant 83.42% of impacted domains solely rely on stale glue records. We will outline how to exploit them in the following section (Section 4).

Legacy and latest IP addresses of stale glue records across organizations, these migrations not only directly demonstrate the migration of nameservers but also imply that the GlueIP left in the zone files is forgotten. The authoritative IP addresses of GlueFQDN have been migrated to the new IP address, but the old glueIP remains in the zone file. We selected the legacy and latest IP addresses of the stale glue records with service status `TIMEOUT` to draw the Sankey diagram in Figure 4. The left shows the autonomous system organization of legacy IP addresses for stale glue records, while the right displays the distribution of their current IP addresses. The legacy glue IPs contain many IP addresses from cloud platforms. The results of our analysis confirm the tangible risk that attackers could exploit stale glue records to hijack active domains (Section 5). In addition, compared to previous domain takeover methods, taking over domains via stale glue records is subtler and poses challenges for detection through active scanning.

Delegation category. The use of glue records under sibling-




Figure 4: Legacy and latest IP address changes of stale glue records. The left shows the distribution of legacy IP addresses for stale glue records in the zone file, while the right displays the distribution of their current IP addresses.

domain delegation far exceeds that under in-domain delegation. The TLD’s nameserver responses glue records if the delegation is in-domain or sibling-domain. We analyzed the delegation patterns across 1,096 TLDs and found that a mere 0.29% of domains employ in-domain delegation. Moreover, we discovered that over 50% of domain names are configured with sibling-domain nameservers, which is the main usage scenario for glue records.

3.2 Glue Record Use in DNS Implementations

In this section, we present a systematic analysis of the glue records usage of mainstream DNS implementations.

As glue records in the DNS that are permanently ignored and never returned to users, whether they can be successfully exploited depends on how resolver software uses them. Hence, the aim of our analysis is to answer three questions:

- Q1: Does DNS software validate glue records before use?
- Q2: Whether the software caches unvalidated glue records?
- Q3: How does the software handle it when no response is received from the GlueIP?

Particularly, after reviewing previous work [2, 41, 43–45], we collected 9 DNS software after discarding some that are outdated or not downloadable, as shown in Table 2, including two proprietary software for Windows, Microsoft DNS [51] and Simple DNS Plus [62] and other seven have published source code. For open source software, we downloaded the latest version of the source code from their official website, then compiled and installed them in Docker; for Windows software, we tested them in a local virtual machine. Depending on the nature of delegation, we design specific test scenarios and then observe the resolution behaviors of different DNS software to infer their usage characteristics of glue records. Next, we detail the design of test scenarios and results.

Table 2: DNS operational modes for glue records with different responses in mainstream implementations.

DNS Software		Active Glue			Stale Glue Service Status*			
brand	version	use directly ¹	check actively ²	shadow caching ³	TIMEOUT	SERVFAIL	REFUSED	NXDOMAIN
BIND [13]	9.18.12	✓	✗	✓	SERVFAIL	SERVFAIL	SERVFAIL	NXDOMAIN
PowerDNS Recursor [58]	4.8.4	✓	✗	✓	SERVFAIL	SERVFAIL	SERVFAIL	NXDOMAIN
Unbound [54]	1.17.1	✓	✓	✓	NOERROR	NOERROR	NOERROR	NXDOMAIN
Knot [37]	5.6.0	✓	✗	✓	SERVFAIL	SERVFAIL	SERVFAIL	NXDOMAIN
CoreDNS [19]	1.10.1	✓	✗	✗	NOERROR	NOERROR	NOERROR	SERVFAIL
Technitium [65]	11.1.1	✓	✓	✗	NOERROR	NOERROR	NOERROR	SERVFAIL
MaraDNS [49]	3.5.0036	✓	✗	✗	NOERROR	NOERROR	NOERROR	TIMEOUT
Microsoft DNS [51]	2022	✓	✗	✓	TIMEOUT	SERVFAIL	SERVFAIL	SERVFAIL
Simple DNS Plus [62]	9.1	✓	✗	✓	TIMEOUT	SERVFAIL	SERVFAIL	NXDOMAIN

¹: When receiving the referral response with glue records, whether the software uses the glue record directly.

²: Whether the software updates glue records actively from their authoritative nameservers. ³: Whether software uses shadow caching.

*: The header is the response status of GlueFQDNs, and the table shows the response status returned to the client by different software after receiving different response statuses.

✓: Yes ✗: No

Test scenarios for different delegations. We designed three test scenarios to verify the use of glue records under different delegations. To determine whether DNS software validates glue records before use, we configure two different authoritative nameservers (ip_adns1 and ip_adns2) to return different responses for the test domain, as shown in Figure 5. We use ip_dns1 to represent the stale glue record remaining in the zone file, whereas ip_dns2 refers to the IP address that ns1.example.com actively resolves to. Then, through the responses, we verify how the resolver utilizes glue records for both legacy (i.e., ip_dns1) and latest (i.e., ip_dns2) IPs. For in-domain and sibling-domain delegation, the glue record is attached directly to the additional section of the referral response from the authoritative nameserver of the TLD. However, the out-domain nameserver scenario does not carry glue records. Therefore, we introduce a domain name with a sibling-domain nameserver to inject glue records into the resolver. We do not use domain names with in-domain delegation for glue record injection because they're not always controllable. It's easier to manage and configure registrable domain names with sibling-domain delegation as we prefer.

Manipulable behaviors of glue records. Most DNS software cache and use glue records without validation, neglecting to ensure the GlueIP matches the authoritative response for GlueFQDN, even when the GlueIP is inactive or returns a negative response. While this behavior does not breach protocol specifications, it is susceptible to abuse.

Caching and using glue without validation. All DNS software trusts and uses the glue record that is in-bailiwick when a response with glue records attached is received, as shown in column *use directly* in Table 2. Under in-domain delegation, this behavior is in line with the RFC 1034, but for sibling-domain and out-domain delegation, this is not the behavior expected [52]. There are still some differences in the way they are implemented in practice. Most of the DNS software trusts the glue records completely without actively resolving, including BIND9 [13], PowerDNS Recursor [58], Knot [37],

CoreDNS [19], MaraDNS [49], Microsoft DNS [51] and Simple DNS Plus [62]. However, Unbound [54] and Technitium [65] query the authoritative A records of GlueFQDNs actively to update the GlueIP. On the other hand, in DNS RFCs, the trust level of glue records is very low and is required to be used in the resolution process and cannot be returned as the answer to users. However, RFCs do not specify the scope of the resolution process, whether it is limited to the current resolution only or can extend to future resolution tasks. In the out-domain test scenario, we evaluate the use of shadow caching, which refers to cached unvalidated glue records. We found that CoreDNS [19], Technitium [65], and MaraDNS [49] do not use shadow caching across different resolution tasks, while all other software leverages glue records as shadow caching. Note that since Unbound resolves actively glue records, it only utilizes unvalidated glue records when the authoritative responses are inactivated.

Misplaced trust for unvalidated glue records. Most DNS software seems to place too much trust in the unvalidated glue records, as shown in column *Stale Glue Service Status* in Table 2. When receiving a negative response (e.g., TIMEOUT and NXDOMAIN) from the GlueIP, five out of the nine DNS software adopt a simple retry policy for stale GlueIPs instead of actively resolving GlueFQDNs to obtain authoritative IP addresses, suggesting these software completely trust the unvalidated glue records. Moreover, the NXDOMAIN response for stale GlueIPs will cause all DNS software to abort the recursive resolution, resulting in a NXDOMAIN response to users.

3.3 Summary

The original intention of introducing glue records was to address the resolution loop issue in in-domain delegation. However, our observations suggest that the current usage of glue records lacks uniform standard practices. Firstly, registries and registrars incorrectly handle stale glue records, resulting in an abundance of outdated glue records in the zone files.

```

;; Registrar
example.com NS ns1.example.com
ns1.example.com A ip_adns1 ; glue record

;;ADNS1 [ip_adns1]
example.com NS ns1.example.com
ns1.example.com A ip_adns2
test.example.com A ip_test1

;;ADNS2 [ip_adns2]
example.com NS ns1.example.com
ns1.example.com A ip_adns2
test.example.com A ip_test2

```

(a) In-domain nameserver test configuration. To answer whether or not the resolver will unconditionally trust the attached glue records under the in-domain delegation.

```

;; Registrar
example.com NS ns1.example.com
example-sibling.com NS ns1.example.com
ns1.example.com A ip_adns1 ; glue record

;;ADNS1 [ip_adns1]
example-sibling.com NS ns1.example.com
test.example-sibling.com A ip_test1

;;ADNS2 [ip_adns2]
example-sibling.com NS ns1.example.com
test.example-sibling.com A ip_test2

```

(b) Sibling-domain nameserver test configuration. To answer whether or not the resolver will unconditionally trust the attached glue records under the sibling-domain delegation.

```

;; Registrar
example.net NS ns1.example.com

;; .com zone file
example.com NS ns1.example.com
example-sibling.com NS ns1.example.com
ns1.example.com A ip_adns1 ; glue record

;;ADNS1 [ip_adns1]
example.net NS ns1.example.com
test.example.net A ip_test1

;;ADNS2 [ip_adns2]
example.net NS ns1.example.com
test.example.net A ip_test2

```

(c) Out-domain nameserver test configuration. To answer whether or not the resolver caches and uses the non-validation glue records to resolve the test domain.

Figure 5: Test configuration of different delegations.

Furthermore, most DNS software exhibits correct-but-can-be-abused behaviors when handling glue records, even though such behavior does not violate protocol specifications. They cache and use unvalidated glue records, including BIND, PowerDNS Recursor, Knot, CoreDNS, MaraDNS, Microsoft DNS, and Simple DNS Plus. Such behavior provides an opportunity for attackers to exploit stale glue records.

4 Attack Overview

In this section, we first describe the technical concept of *shadow caching*. Then, we introduce the threat models of our domain takeover and DoS attacks, and describe the workflows of the two attacks. Finally, we compare our threat models with the previous work.

4.1 Shadow Caching

According to RFC1034 [52], DNS glue records are only allowed to be used in the referral response of the same TLD, i.e., under in-domain or sibling-domain delegation. Under out-domain delegation, the referral response of the TLD response does not contain glue records, as shown in Figure 2(3). Once glueFQDN (e.g., ns1.test.net) is actively resolved, the resolver will get the authoritative record (the correct record) of the glueFQDN instead of the glue record.

To exploit the glue records under out-domain delegation, we introduce *shadow caching*, a glue record cache resulting from carefully constructed delegation records. According to the analysis in Section 3, mainstream DNS software caches the glue records in the referral response (as shown in Figure 2(1) and (2)) and utilizes them for future resolution. Therefore, we can create sibling-domain delegation to inject specific stale glue records into the target resolver in advance, as shown in the preparation phase in Figure 8. Specifically, we register a domain with the same TLD as the stale glue (i.e., attack.com), configuring its NS to ns1.vulner.com (i.e., the target stale glueFQDN). When querying attack.com,

the target resolver will cache the glue records in the referral response, leading to the *shadow caching*.

In general, an attacker can easily implement shadow caching by registering a domain name with the same TLD as the target glueFQDN and configuring NS records. Moreover, the test results in Section 3 show that 6 mainstream DNS software is affected by this problem.

4.2 Threat Model

Domain takeover. Figure 6 illustrates the threat model of the domain takeover. We assume that the GlueFQDN ns1.vulner.com., on which victim.net. depends, is stale and remain in .com zone file. Moreover, when the recursive resolver receives a referral response from the authoritative nameserver of the TLD, it applies *shadow caching* directly. To complete the domain takeover, the following requirements need to be fulfilled.

(1) Exploitable stale glue records. Most gTLD zone files are publicly accessible and can be acquired from their registries, like .com of Verisign. Moreover, ICANN developed the CZDS to simplify the process of applying for zone data access [29]. Once attackers had access to the zone files, they were able to discover a large number of potentially exploitable stale glue records, as we showed in Section 5.1.

(2) Assignable cloud IPs. The attacker is capable of obtaining IP addresses released by other users on cloud platforms by continuing to allocate and release. Prior studies showed the feasibility of milking the cloud platform IP pool from the perspective of an external attacker [14, 46]. Moreover, a recent study [56] uncovered the cloud squatting problem within a cloud platform, further revealing the severity of cloud IP reuse issues.

DoS attack. When the GlueIP is unobtainable, although attackers cannot directly hijack the target domain, they can still launch a denial-of-service attack against it. Figure 7 illustrates the threat model of the DoS attack. We assume that the stale glue record stops service for the target domain name, and the

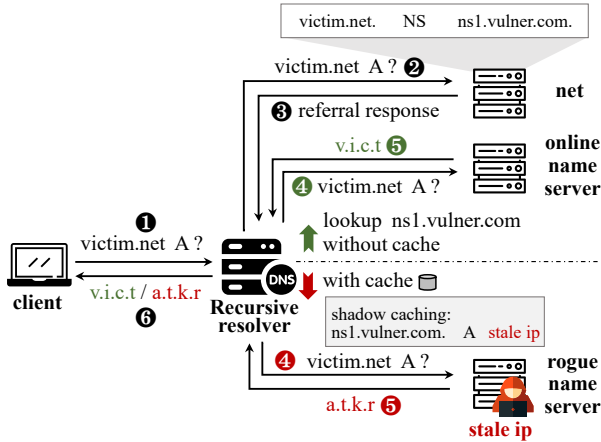


Figure 6: Threat model of domain takeover by stale glue records.

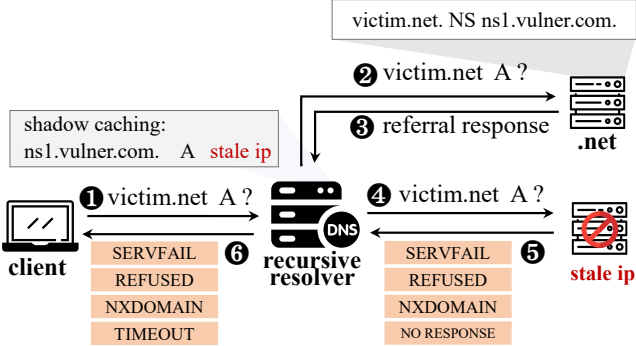


Figure 7: Threat model of DoS attack by stale glue records.

attacker can not get hold of the GlueIP because it is assigned or not a cloud IP. And the target resolvers can be accessed, so the attacker is able to inject the *shadow caching*. Moreover, the target domain configures out-domain nameservers and all GlueFQDNs of nameservers are stale. We found through large-scale experiments that this configuration is very common, making them vulnerable to potential attacks (Section 3). On the other hand, we assume that the recursive resolver uses the *shadow caching* directly.

4.3 Attack Workflow

Figure 8 presents steps of our domain takeover and DoS attack via stale glue records. In the preparation stage (i.e., causing the *shadow caching*), the attacker collects the exploitable glue records (e.g., ns1.vulner.com.), whose GlueIP is assignable (domain takeover) or stops authoritative services (DoS). Next, we require injecting the *shadow caching* into the resolver. To achieve this, we first registered a domain with the same TLD as the stale glue (i.e., attack.com), configuring its NS to ns1.vulner.com. Note that this step does not require

the cooperation of vulner.com; attackers can complete it independently. Subsequently, the attacker queries attack.com to a target resolver, causing it to cache the glue records in the referral response, leading to the *shadow caching*. In the next stage, the attacker can conduct the domain takeover and DoS attack leveraging the above cache.

Domain takeover attack. The part two of Figure 8 illustrates the steps of the domain takeover attack via stale glue records. First, a client queries the domain (e.g., victim.net) to the target resolver (step 1), where victim.net is delegated to ns1.vulner.com. After receiving the referral response from .net nameservers (step 2), the resolver resolves the domain from rogue nameservers controlled by the attacker due to the *shadow caching* is not expired (step 3). Finally, the resolver returns the fake resource records to the client (step 4).

DoS attack. The part three of Figure 8 presents the step of DoS attack by stale glue records. Similar to the domain takeover, when the client queries the domain name to the target resolver, the resolver requests the resource records leveraging the *shadow caching* (step 1 ~ 3). When multiple retries fail to obtain a valid response, the resolver returns a failed response to the client (step 4 and 5).

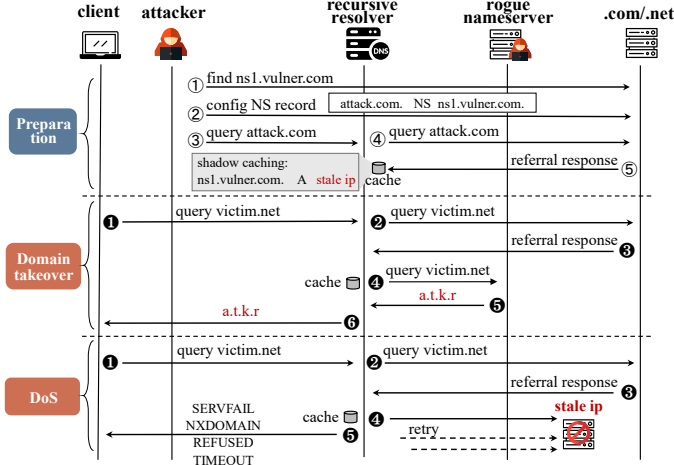


Figure 8: Workflow of domain takeover and DoS attacks.

4.4 Comparison to Prior Work

Similar to previous work on domain takeover [3, 14, 46, 70], our threat model leverages the stale resource records to hijack target domains. However, our threat model differs from these studies by exposing a long-standing overlooked and broader attack surface related to the usage mechanism of glue records, even threatening a large number of active domains, not just expired domains. By analyzing the usage of glue records in DNS software, we introduce the *shadow caching* to address a significant gap in glue record security research, i.e., the exploitation of glue records under out-domain delegation. Compared to prior work [14, 46], the *shadow caching* broadens the range




Figure 9: Identifying potential exploitable stale glue records from zone files.

of impact and increases the flexibility of attacks. Attackers can target any domain name delegated to stale glue records by means of meticulously crafted delegation records, even in out-domain delegation. Furthermore, our findings reveal more affected domains than previous studies. From the 1,096 TLD zone files, we identified 193,558 directly potentially exploitable glue records, impacting 6,687,000 domains.

5 Evaluating the Impact of Stale Glue Records

In this section, we demonstrate that the risks of stale glue records in the wild are widespread and underplayed. We first evaluate the impact scope of our threat models by analyzing vulnerable glue records in zone files. Subsequently, we conduct comprehensive experiments and measurements to understand the real-world implications of stale glue records, encompassing both open DNS servers and prominent public DNS services.

5.1 Domain Takeover

5.1.1 Methodology

In this part, we propose a novel methodology that finds potential exploitable stale glue records in zone files.

Technical challenges. Glue records can be valid but not resolvable. A glue record works as long as it is submitted to zone files [52]. However, a glue record will only be resolvable if its parent domain has added an A record for the glue record on its authoritative nameservers. Otherwise, the glue record will be invisible in DNS, present only in zone files, and unavailable through active resolve. On the other hand, invalid glue records are not necessarily exploitable, so careful verification is required to confirm whether attackers can acquire the GlueIP. As a result, identifying exploitable stale glue records poses a significant challenge.

Identifying potential exploitable stale glue records from zone files. We extract all glue records from zone files regardless of their resolvability, delineating our analysis scope and reducing complexity. Moreover, the prerequisite that a glue record can be exploited is that it is outdated and assignable, suggesting attackers can obtain it. Therefore, based on the previous studies [14], we develop our approach to identify the potential exploitable stale glue records, as shown in Figure 9.

First, we select GlueFQDNs from the stale glue records found by Section 3 exhibiting a service status of `TIMEOUT`, indicating that the record no longer furnishes authoritative services for any domain. Next, we leverage the ASdb [71], a research dataset that maps public autonomous systems (identified by ASN) to organizations and industry types, to locate the GlueIP belonging to cloud platforms. Finally, we examine the GlueIPs to determine if they are offline and assignable by conducting a port scan on the 33 most commonly used TCP and UDP ports [14], as shown in Table 6 (in Appendix A). When the service status of a GlueFQDN is `TIMEOUT`, and the GlueIP is offline and assignable, we mark it as a potential exploitable stale glue record.

5.1.2 Results

Leveraging our approach to zone files from August 10, 2023, we identified 193,558 exploitable stale glue records mapping to 100,258 cloud IPs. 39,795 (20.56%) of these glue records (5,615 SLDs) are present with Tranco Top 1M, as shown in Table 3. Figure 10a shows the ranking distribution of these stale glue records FQDN and SLD. From a TLD perspective (Figure 10b), the highest number of exploitable stale glue records are found under `.com`, followed by `.net` and `.org`. Moreover, the number of exploitable stale glue records we found far surpassed statistics on orphan and abandoned records in prior research, indicating that the security risks of stale glue records remain largely unaddressed and growing.

These glue records are scattered throughout the DNS space, posing an extensive yet pressing issue requiring resolution. However, the distribution of SLDs partially accounts for stale glue: undocumented registrar operations caused glue records to remain in the zone files. Similar to the operation observed by Akiwate et al. [4], we speculate that some of the stale glue records may also originate from an undocumented operation by registrars. Specifically, when a domain expires, registrars should remove all associated records. However, some instead rename and keep the glue records of expired domains in zone files. One specific example is that for the SLD `directideleteddomain.com`, we identified 1,967 stale glue records, which map to 1,491 IP addresses. Moreover, from the stale glue records, it is possible to identify the domain names before they are renamed. For example, the stale GlueFQDN `zeus1.smmdeyiz.com.directideleteddomain.com` and domain name before renamed `zeus1.smmdeyiz.com`. We successfully traced the historical mapping relationships between the domains and GlueIP using passive DNS data.

Furthermore, we identified 6,398,631 domain names delegated to exploitable stale glue records susceptible to takeover. Of these, 5,947,669 (92.95%) domain names are active. Among these domains, 5,395 ranked in Tranco Top 1M, and Figure 10c shows their ranking distribution. For example, the domain `trueconf.net` (ranked 2,138), a leader in enterprise video conferencing and collaboration in Europe, has




Figure 10: Exploitable stale glue records information.

Table 3: Potential exploitable stale glue record statistics.

	# (in Tranco)	# SLD (in Tranco)
Stale glue record	193,558 (39,795)	99,373 (5,470)
Takeover domain	-	6,398,631 (5,395)
DoS domain	-	784,693 (1,576)

modified the nameserver on their authoritative server but failed to remove the discarded NS RR `ns7.trueconf.net`. Other domain names include the largest U.S. pharmacy chain `walgreens.com`, and even ISP `i2bnetworks.com`.

5.1.3 Discussion

The proportion of stale glue in a domain's nameservers influences the performance of the domain takeover. Two factors must be considered to evaluate the impact of stale glue records on the target domain: the number of NSes configured for the domain and the number of stale GlueIP associated with each NS. We assume that if a domain name is configured with N NS records, then the probability of each NS being selected is $1/N$. Moreover, each NS has M GlueIPs, of which m are stale. Then, the formula to calculate the degree of impact on the domain is as follows:

$$Hijacking\ ratio = 1/N * \sum_{i=0}^N m_i/M_i$$

Where N represents the number of NS for the domain. M_i and m_i respectively represent the number of GlueIPs and stale GlueIPs under the i th NS.

We have calculated the hijacking ratio of affected domains and discovered that half of the traffic for 84% of these domains could be taken over by attackers, and Figure 12 shows the CDF of hijacking ratio (in Appendix B). Subsequently, we examine whether these stale glues are cleaned up in time. We tracked the lifecycle of exploitable stale glue records within

zone files, including whether they are still present in zone files and whether they are still being used by domain names. The results demonstrate that with few glue records removed within a month, attackers have a substantial time window to operate. Figure 11 shows the monthly change in stale glue records identified. The blue dashed line experiences a trough on August 13th, which is attributed to the partial loss of `.net` zone files during our data processing.

Another challenge of hijacking a domain is obtaining cloud IPs. However, a lot of work has been done to prove the feasibility of this [14,46,56]. Liu et al. [46] and Borgolte et al. [14] collected IP pools from different cloud service vendors by applying and releasing, while Pauley et al. [56] examined cloud IP reuse from the internal perspective of a cloud platform, proving that cloud IPs are frequently acquired by other users after release. Following the method of previous work [14,46], we utilize the Amazon Cloud API `allocate_address` and `release_address` to automate the allocation and release of cloud IPs [9] at a very low rate to prove stale GlueIPs could be obtained by attackers. As for attackers, they can apply for a large number of IPs in a short time. Within two weeks, we successfully applied for 27 GlueIPs, costing \$2.3 total. We released IPs immediately after allocation, so there is no actual impact on domains. Moreover, we also reported to domain owners to remind them of the risk of their domains.

5.2 DoS Attack

Attack condition. Analysis of DNS software behavior (Section 3.2) shows that most DNS software misplaces trust in unvalidated glue records. Exploiting stale glue records, attackers can perform a denial-of-service attack on some domains. When domain names meet the following conditions, they may be affected: (1) the domain configures out-domain nameservers; (2) all the GlueFQDNs of nameservers are stale. **Results.** While the attack conditions may seem stringent, we still managed to identify a substantial number of affected domain names. And we exclude domain names that cannot




Figure 11: Number of exploitable stale glue records over time in zone files (August 10-September 10).

be resolved as they hold no value for the attack. As a result, we identified 784,693 active domains meeting the conditions, including 1,576 popular domain names in Tranco Top 1M.

Discussion. After the cache expires, attackers need to reinject stale records into the target resolver again. Here, we are discussing the negative cache TTL, glue record TTL, and the maximum cache TTL of the software. The implementation of negative cache TTL depends on the software itself, but it typically does not exceed the maximum value specified in the standards [10], which is 300 seconds. Moreover, we conducted a statistical analysis on the TTL of 2.8M glue records and found that they have a fixed value of 172,800 (2 days). In addition, Li et al. [41] analyzed the Maximum cache TTL of mainstream DNS software and found that all popular software has a maximum cache TTL of more than one day and 3 of them have a maximum over 6-day TTL by default (BIND9 [13], Simple DNS Plus [62], and Knot [37]).

Table 4: Open DNS resolver statistics.

	# IP	%
DNS resolver on Apr. 03 2023	1,846,535	
DNS resolver alive on Jun. 10 2023	895,674	100%
Domain Takeover*		
in-domain delegation	834,644	93.19%
sibling-domain delegation	830,146	92.68%
out-domain delegation	218,942	24.44%
portion stale glue	201,607	22.51%
Domain Denial-of-Service		
misplaced trust in unvalidated glue records	111,766	12.48%

*: Number of resolvers that use unvalidated glue records directly on different delegations.

5.3 Evaluating DNS Resolvers in the Wild

In this part, we perform extensive experiments to evaluate the impact of stale glue records in the wild. Due to ethical considerations, we evaluate glue record usage under different delegation categories for resolvers to infer the impacts of stale glue records. The primary external factor in our threat models is how resolvers use glue records, so this evaluation is reasonable. Moreover, we conduct DoS attack experiments on controlled domains to identify their impact on open resolvers.

Meanwhile, we also measured open resolvers' non-standard behavior. This behavior of resolvers can expand the attack surface of our threat model.

5.3.1 Collecting DNS Resolver

Our test DNS resolvers include not only public DNS providers but also a large number of stable open resolvers.

Public DNS providers. We collect a famous public DNS provider list, whose users are distributed all over the world, and they are used in prior works as well [2, 31, 32, 35, 36, 47, 48, 60], like Google Public DNS, Cloudflare DNS.

Stable and open resolvers. Some studies have demonstrated significant churn of open DNS resolvers [39]. Hence, our measurements focus on stable DNS resolvers. To this end, we scanned the entire IPv4 address space on April 3rd and June 3rd, 2023, using XMAP [42] respectively, then took the intersection of the two results as our measurement subjects (i.e., open resolvers that stably exist for two months), and the detailed statistics are listed in Table 4.

5.3.2 Measurement Setup

This section aims to measure resolvers' usage of unvalidated glue records and assess whether such usage can lead to domain takeover or DoS attacks. To comprehensively evaluate resolver behavior, we employed different configurations, including in-domain, sibling-domain, and out-domain scenarios, as detailed in Section 3.2. Moreover, T1 and T2 were introduced to investigate specific out-domain configurations further. Here, we denote the authoritative A records actively queried by the resolver as QR.

T1: Outage QR, and Online GlueIP. The QR is offline, and the service of GlueIP is normal. We registered four new test domain names to configure different QR and GlueIP. GlueIP provided normal resolution services for these four domain names, while QR returned different error status codes, including SERVFAIL, NXDOMAIN, REFUSED, and without response.

T2: Online QR, and Online and Outage GlueIP. We set up another four new test domain names to avoid the cache effect of open resolvers. One GlueIP and QR provided normal responses to these four domain names, while the other GlueIP returned different error status codes.

Furthermore, special configurations were implemented to assess the impact of DoS attacks. We configure two new domain names, `dos-test.com` and `inject-cache.net`. Among them, all glue records (`ns1.vulner.net` and `ns2.vulner.net`) of `dos-test.com` are stale, making it vulnerable to DoS attack. The `inject-cache.net` is used to cause shadow caching in the target resolver. It is also delegated to `ns1.vulner.net` and `ns2.vulner.net`.

Each experiment is repeated three times to comprehensively and accurately evaluate the behavior of open DNS resolvers. In each experiment, we first inject glue records into resolvers

Table 5: Glue usage of public DNS resolver.

Public DNS	In Domain	Sibling Domain	Out Domain			Vulnerable?	
			Active glue	T1	T2	Takeover	DoS
Cloudflare DNS [18]	●	○	○	○	○	✓	✗
Google Public DNS [23]	●	○	○	○	○	✓	✗
Alibaba Cloud DNS [7]	●	●	●	●	●	✓	✗
114DNS [1]	●	○	○	○	○	✓	✗
OpenDNS [15]	●	○	●	●	●	✓	✓
Level3 DNS [40]	●	○	●	○	○	✓	✗
Quad9 DNS [59]	●	●	●	●	○	✓	✓
Neustar UltraDNS [53]	●	○	○	○	○	✓	✗
Dyn DNS [20]	●	○	○	○	○	✓	✗
CleanBrowsing DNS [17]	●	○	○	○	○	✓	✓
DNSPod Public DNS+ [66]	●	●	●	●	●	✓	✓
Baidu DNS [11]	●	○	○	○	○	✓	✓
Verisign Public DNS [67]	●	○	○	○	○	✓	✗
Yandex.DNS [69]	●	○	○	○	○	✓	✗

●: Full use. ○: Partial use. ○: No use.
 ✓: Yes ✗: No

using the domain name under the same TLD as GlueFQDN and then query different test domain names, respectively.

5.3.3 Measurements and Results

Most of the stale open resolver lacks validation for glue records. As listed in Table 4, under in-domain and sibling-domain delegation, over 90% of resolvers cache and use unvalidated glue records. While for out-domain delegation, 24.44% of resolvers blindly trust shadow caching. On the other hand, 410,406 resolvers resolve actively glue records in each of our tests. Of these, however, 67% resolvers exhibited non-standard behavior. When the service of validated NSes is abnormal, they fall back to using the glue records.

After repetitious experiments, we show that *all 14 public DNS are vulnerable to domain takeover* because they unconditionally cache and use glue records, like Open DNS and Quad9. Although Google Public DNS and Cloudflare DNS only trust the records they resolve actively instead of using glue records under out-domain delegation, they are still affected under in-domain and sibling-domain delegation. Moreover, these public DNS are affected to varying degrees by their complex multi-backend and multi-cache architecture [60], as shown in Table 5. Three Public DNS (Alibaba Cloud DNS, Quad9, and DNSPod) will continue to make full use of glue records under the sibling-domain delegation, but other public DNS are using glue in some of the tests (i.e., partial use). There are two reasons for this phenomenon: 1) the public DNS backend is multi-cache, and our test domain name just happened to miss the glue records we injected. 2) The backend implementations of one public DNS are different, resulting in a difference in their behavior when using glue records. Even so, attackers can still take over part of the target domain name traffic. In addition, we discover that *five public DNS (OpenDNS, Quad9, CleanBrowsing, DNSPod, and Baidu DNS) misplaced trust unvalidated glue records*, who are vulnerable to our DoS threat model.

5.4 Ethical Considerations

Our experiments involve large-scale open DNS resolver scanning, DNS software implementation analysis, and attack effect evaluation against different public DNS and open resolvers. Hence, we have considered many ethical considerations in the experimental design. We strictly follow the existing ethical principles of Menlo Report [34] and best practices of network measurements [38, 39].

First, for large-scale open DNS resolver scanning, to minimize the impact on DNS resolvers, we strictly limit the scan rate. In addition, we limit the number of DNS queries for each resolver in each round of testing.

Second, we perform controlled experiments. We use newly registered experimental domain names to evaluate the effect of attacks. All domain names and authoritative servers are under our control. To avoid the impact of cache and different service statuses, we registered more than 15 new domain names under different TLDs. These new domains have no real-world impact as they are only used in our experiments.

Third, we report vulnerabilities to all relevant domain owners and resolver vendors for responsible disclosure.

6 Discussion and Mitigation

In this section, we discuss the lessons learned from new threats. Then, we propose mitigation measures. Finally, we describe our responsible disclosure.

6.1 Lessons Learned

When RFC 1034 introduced glue records, they aimed to assist in domain name resolution under in-domain delegation and were prescribed for use in referral responses. Nevertheless, as DNS has progressed, such foundational guidelines no longer suffice for the multifaceted applications of DNS. Currently, the management and utilization of glue records are predominantly determined by individual registrars, registries, or resolver software, leading to a noticeable absence of a comprehensive and unified guideline.

Due to a lack of awareness and concern regarding the security risks of glue records, over time, a large number of legacy glue records are left in the TLD zone file, leading to deterioration in the usability and credibility of glue records. On the other hand, the usage of glue records by most DNS software also deviates from their initial design intent. They utilize glue to enhance resolution efficiency, irrespective of whether it's sibling-domain or even out-domain delegation. However, the configuration of glue records is typically distinct from that of conventional A and NS records. Domain owners are required to manipulate specific functions provided by the registrars. This might pose confusion to ordinary users, thereby leading to operational errors or forgetting previous configurations.

Thus, the behavior of DNS software is risky. Our work highlights that glue records constitute a forgotten foundation of DNS architecture. We advocate for thoroughly reassessing security threats for glue records and establishing standardized guidelines for their usage scenarios, with renewing security prioritization.

6.2 Mitigation

The core issue of the stale glue threat lies in managing glue records by registrars and registries. Upon the registrar's comprehensive cleanup of all invalid glue records, attackers will lose their avenue of attack. Therefore, we recommend that registrars and registries establish or refine their strategies to promptly clean up stale glue records. On the other hand, the undocumented and improper operations of registrars on expired domain names have introduced direct attack vectors into TLD zone files. These operations should be standardized, and the incorrectly introduced glue records should be removed.

For resolver software, validating glue records is complex and may reduce efficiency. However, sacrificing security for speed is not advisable. Thus, we recommend that mainstream resolver software actively query the glueFQDN to obtain an authoritative response and use it when encountering a glue record under sibling-domain and out-domain delegation. In the experiments, we observed that Unbound and Technitium have implemented similar strategies. They prioritize glueFQDN's authoritative response over the glueIP.

6.3 Disclosure and Responses

We have reported vulnerabilities to all relevant vendors and DNS software for responsible disclosure and are discussing the mitigation with them.

Regarding resolver software, we understand that given the complexities of domain dependencies, resolver software must consider multiple factors to balance security and performance. Nonetheless, we have submitted reports to mainstream DNS software and public DNS services, recommending the deployment of mitigation measures. As a result, Microsoft DNS confirmed the vulnerability and ranked it as `important`, with plans to release a patch in April 2024, along with a \$1,000 bonus. PowerDNS acknowledged the threat but thought that fixing this issue may impact resolution performance, and we are discussing an optimized solution. Moreover, OpenDNS and Alibaba Cloud DNS have acknowledged the threat and are preparing to implement mitigation based on our reports [16].

Furthermore, we have contacted the respective TLD registries to explore more effective mitigation strategies. `.org` and `.info` have acknowledged the issue and are collaborating with us on a cleanup plan. We have also reported the issue to ICANN and are working together to assist other registries in resolving this problem. Given the impact on 6 million affected domains, it is not feasible for us to individually report to each

one. Thus, we disclosed the information through our national CERT and await their response.

7 Related Work

DNS Misconfiguration. Although RFC 1034 requires that the delegation data from parent and child zones should be consistent, inconsistent delegation caused by complex DNS configurations is prevalent in practice [12, 61], which is the main reason for the generation of legacy records. Back in 2004, Pappas et al. [55] found that 15% of DNS zones had inconsistent delegation. Kalafut et al. [33] noticed that some glue records remained in the TLD zone even though their parent domains had expired. They referred to these records as orphan records and found that some orphans were evidently used for malicious purposes. In 2020, Sommese et al. [63] reviewed orphan records again. They observed a significant decrease in `.com` and `.net`, suggesting that TLD operators had implemented policies to mitigate this. However, in some new gTLDs, this phenomenon showed an upward trend, such as `.info` and `.mobi`. Moreover, under the common gTLDs (i.e., `.com`, `.net`, and `.org`), Sommese et al. [64] have shown that almost 8% of domain name delegations are inconsistent. They are concerned about how it could affect DNS operations, such as improper load balancing among the nameservers. To mitigate this problem, RFC 7477 [24] proposes a new record type (CSYNC) to synchronize delegation information between parent and child zones. However, actual deployment is rare. These previous works have identified a large number of inconsistent delegation records (NS RR), our work further proposes a novel exploitation method of DNS glue records and systematically evaluates the real-world threat it poses.

Domain Takeover. In essence, the DNS provides a resource mapping service. If the mapping is not purged when the service expires or is terminated, it can be exploited by attackers. Liu et al. [46] proposed four types of exploitable dangling DNS records (Dare), showing that Dare is a real and prevalent threat. Moreover, Borgolte et al. [14] discovered a significant number of stale DNS records pointing to available IP addresses in clouds, introducing a new attack vector. Akiwate et al. [3] utilized comprehensive collections of both active and passive DNS measurements to investigate and quantify the risks of lame delegation and showed that lame delegation in some zone files affected 14% of domains. In addition, Alowaisheq et al. [8] investigated a new category of stale NS records that reside in the domain zone (instead of the TLD zone) for an active domain, which can cause a stealthier hijacking of target domains. Houser et al. [28] have shown that authoritative nameserver deployments in government domains still contain a non-trivial number of configurations that do not comply with RFC requirements and that more than 1,000 domains are vulnerable to hijacking because of these configurations. Unlike previous research, our work introduced the concept of *shadow caching* enabling domain takeover or

DoS under out-domain delegation.

DNS Cache Measurement. DNS depends on extensive caching for good performance, and numerous studies have contributed to DNS cache measurements. Klein et al. [36] proposed methodologies for efficiently discovering and enumerating the caches of the DNS resolution platforms to shed light on architectures and configurations of the caches in DNS resolution platforms. And, Al-Dalky et al. [5] presented a characterization and classification of the multiple recursive resolver pools and showed the pools exhibit a wide range of behaviors. Randall et al. [60] developed the Trufflehunter tool, which models the complex behavior of large multi-layer distributed caching infrastructures, and used it to infer the caching strategies of four popular public DNS resolvers. Moreover, Li et al. [41] measured the caching and transmission mechanism of DNS software and open resolvers. Previous work has only focused on conventional resource records such as the A record, ignoring the glue record hidden in the reference response. As a result, the security community currently lacks a clear understanding of the usage behavior of glue records in the wild. Our work evaluates the caching and usage behavior of glue records by open resolvers and mainstream DNS software, revealing manipulable behavior in mainstream implementations. Exploiting this, we propose new threat models.

8 Conclusion

In this paper, we systematically measure and assess the security risks of glue records, especially under out-domain delegation, proving that they are real and prevalent. We empirically identify that 23.18% of glue records are stale in the zone files of 1,096 TLDs. In addition, through reverse engineering 9 mainstream DNS implementations (e.g., BIND 9 and Microsoft DNS), we reveal the risky behaviors of glue records. Moreover, we identified 193,558 exploitable glue records. These affect over 6M domain names that are vulnerable to takeover or DoS attacks. To evaluate the real-world implications of stale glue records, we perform large-scale measurements toward stable open resolvers and 14 well-known public DNS resolvers. The experimental results show that over 90% of resolvers and all surveyed public DNS cache and use unvalidated glue records, including Quad9, and OpenDNS. To initiate the remediation process, we have notified the affected vendors and actively followed up on the deployment of the mitigation solutions. Our work provides a thorough picture of glue record threats and hopes to inspire the enthusiasm of registrars, registries, and DNS software for the security threats of glue records.

Acknowledgments

We sincerely thank all anonymous reviewers and our shepherd for their valuable comments to improve the paper. We also

thank Ruixuan Li, Zhijie Xie, and Weijing Liang for their help with the paper. This work is in part supported by the National Key Research and Development Program of China (No. 2023YFB3105600), the National Natural Science Foundation of China (62102218), CCF-Tencent Rhino-Bird Young Faculty Open Research Fund (CCF-Tencent RAGR20230116). Haixin Duan is supported by the Taishan Scholars Program. Min Zhang is the corresponding author.

References

- [1] 114DNS. 114DNS. <https://114dns.com/>, 2023.
- [2] Yehuda Afek, Anat Bremler-Barr, and Lior Shafir. Nxn-sattack: Recursive DNS inefficiencies and vulnerabilities. In *29th USENIX Security Symposium*, 2020.
- [3] Gautam Akiwate, Mattijs Jonker, Raffaele Sommese, Ian Foster, Geoffrey M Voelker, Stefan Savage, and KC Claffy. Unresolved issues: Prevalence, persistence, and perils of lame delegations. In *Proceedings of the ACM Internet Measurement Conference*, 2020.
- [4] Gautam Akiwate, Stefan Savage, Geoffrey M. Voelker, and Kimberly C. Claffy. Risky bizness: risks derived from registrar name management. In *IMC '21: ACM Internet Measurement Conference*, 2021.
- [5] Rami Al-Dalky and Kyle Schomp. Characterization of collaborative resolution in recursive DNS resolvers. In *Passive and Active Measurement - 19th International Conference, PAM*, 2018.
- [6] Alibaba Cloud. Alibaba Cloud: Cloud Computing Services. <https://www.alibabacloud.com/>, 2023.
- [7] AliDNS. AliDNS. <https://alidns.com/>, 2023.
- [8] Eihal Alowaisheq, Siyuan Tang, Zhihao Wang, Fatemah Alharbi, Xiaojing Liao, and XiaoFeng Wang. Zombie awakening: Stealthy hijacking of active domains through dns hosting referral. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 2020.
- [9] Amazon EC2. EC2. <https://boto3.amazonaws.com/v1/documentation/api/latest/reference/services/ec2.html>, 2024.
- [10] Mark P. Andrews. Negative caching of DNS queries (DNS NCACHE). *RFC*, 1998.
- [11] Baidu. BaiduDNS. <https://dudns.baidu.com/>, 2023.
- [12] David Barr. Common DNS operational and configuration errors. *RFC*, 1996.

- [13] BIND. BIND 9. <https://www.isc.org/bind/>, 2023.
- [14] Kevin Borgolte, Tobias Fiebig, Shuang Hao, Christopher Kruegel, and Giovanni Vigna. Cloud strife: Mitigating the security risks of domain-validated certificates. In *25th Annual Network and Distributed System Security Symposium*, 2018.
- [15] Cisco. OpenDNS. <https://www.opendns.com/>, 2023.
- [16] Cisco PSIRT. Add validation of DNS glue records. <https://bst.cloudapps.cisco.com/bugsearch/bug/CSCwf99517>, 2024.
- [17] CleanBrowsing. CB DNS. <https://cleanbrowsing.org/>, 2023.
- [18] CloudFlare. CloudFlare DNS. <https://1.1.1.1/dns/>, 2023.
- [19] CoreDNS. CoreDNS. <https://coredns.io/>, 2023.
- [20] Dyn. Dyn DNS. <https://help.dyn.com/internet-guide-setup/>, 2023.
- [21] Robert Elz and Randy Bush. Clarifications to the DNS specification. *RFC*, 1997.
- [22] Godaddy. Domain Names, Websites, Hosting & Online Marketing Tools. <https://www.godaddy.com>, 2023.
- [23] Google. Google Public DNS. <https://dns.google/>, 2023.
- [24] Wes Hardaker. Child-to-parent synchronization in DNS. *RFC*, 2015.
- [25] Scott Hollenbeck. Extensible provisioning protocol (EPP). *RFC*, 5730:1–67, 2009.
- [26] Scott Hollenbeck. Extensible provisioning protocol (EPP) domain name mapping. *RFC*, 2009.
- [27] Scott Hollenbeck. Extensible provisioning protocol (EPP) host mapping. *RFC*, 2009.
- [28] Rebekah Houser, Shuai Hao, Chase Cotton, and Haining Wang. A comprehensive, longitudinal study of government DNS deployment at global scale. In *52nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, 2022.
- [29] ICANN. Centralized Zone Data Service. <https://czds.icann.org/>, 2023.
- [30] Ivan Ristić. Subdomain Takeover Prevention and Dangling DNS Detection. <https://www.hardenize.com/blog/dangling-dns-detection-and-subdomain-takeover-prevention>, 2024.
- [31] Philipp Jeitner and Haya Shulman. Injection attacks reloaded: Tunnelling malicious payloads over DNS. In *30th USENIX Security Symposium*, 2021.
- [32] Jian Jiang, Jinjin Liang, Kang Li, Jun Li, Hai-Xin Duan, and Jianping Wu. Ghost domain names: Revoked yet still resolvable. In *19th Annual Network and Distributed System Security Symposium, NDSS*, 2012.
- [33] Andrew J Kalafut, Minaxi Gupta, Christopher A Cole, Lei Chen, and Nathan E Myers. An empirical study of orphan dns servers in the internet. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, 2010.
- [34] Erin Kenneally and David Dittrich. The menlo report: Ethical principles guiding information and communication technology research. Available at SSRN 2445102, 2012.
- [35] Amit Klein. Cross layer attacks and how to use them (for DNS cache poisoning, device tracking and more). In *42nd IEEE Symposium on Security and Privacy*, 2021.
- [36] Amit Klein, Haya Shulman, and Michael Waidner. Counting in the dark: DNS caches discovery and enumeration in the internet. In *47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, 2017.
- [37] Knot Resolver. Knot Resolver. <https://www.knot-resolver.cz/>, 2023.
- [38] Maciej Korczyński, Michał Król, and Michel Van Eeten. Zone poisoning: The how and where of non-secure dns dynamic updates. In *Proceedings of the 2016 Internet Measurement Conference*, 2016.
- [39] Marc Kühner, Thomas Hupperich, Jonas Bushart, Christian Rossow, and Thorsten Holz. Going wild: Large-scale classification of open dns resolvers. In *Proceedings of the 2015 Internet Measurement Conference*, 2015.
- [40] Level3. DNS. <https://www.publicdns.xyz/public/level3.html>, 2023.
- [41] Xiang Li, Baojun Liu, Xuesong Bai, Mingming Zhang, Qifan Zhang, Zhou Li, Haixin Duan, and Qi Li. Ghost Domain Reloaded: Vulnerable Links in Domain Name Delegation and Revocation. In *Proceedings of the 30th Annual Network and Distributed System Security Symposium*, 2023.
- [42] Xiang Li, Baojun Liu, Xiaofeng Zheng, Haixin Duan, Qi Li, and Youjun Huang. Fast IPv6 Network Periphery Discovery and Security Implications. In *Proceedings of the 2021 IEEE/IFIP International Conference on Dependable Systems and Networks*, 2021.

- [43] Xiang Li, Chaoyi Lu, Baojun Liu, Qifan Zhang, Zhou Li, Haixin Duan, and Qi Li. The Maginot Line: Attacking the Boundary of DNS Caching Protection. In *Proceedings of the 32nd USENIX Security Symposium*, 2023.
- [44] Xiang Li, Dashuai Wu, Haixin Duan, and Qi Li. DNS-Bomb: A New Practical-and-Powerful Pulsing DoS Attack Exploiting DNS Queries-and-Responses. In *Proceedings of 2024 IEEE Symposium on Security and Privacy (S&P '24)*, 2024.
- [45] Xiang Li, Wei Xu, Baojun Liu, Mingming Zhang, Zhou Li, Jia Zhang, Deliang Chang, Xiaofeng Zheng, Chuhan Wang, Jianjun Chen, Haixin Duan, and Qi Li. TuDoor Attack: Systematically Exploring and Exploiting Logic Vulnerabilities in DNS Response Pre-processing with Malformed Packets. In *Proceedings of 2024 IEEE Symposium on Security and Privacy (S&P '24)*, 2024.
- [46] Daiping Liu, Shuai Hao, and Haining Wang. All your dns records point to us: Understanding the security threats of dangling dns records. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016.
- [47] Keyu Man, Zhiyun Qian, Zhongjie Wang, Xiaofeng Zheng, Youjun Huang, and Haixin Duan. DNS cache poisoning attack reloaded: Revolutions with side channels. In *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security*, 2020.
- [48] Keyu Man, Xin'an Zhou, and Zhiyun Qian. DNS cache poisoning attack: Resurrections with side channels. In *CCS '21: 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021.
- [49] MaraDNS. MaraDNS. <https://maradns.samiam.org/>, 2023.
- [50] Mark P. Andrews, Shumon Huque, Paul Wouters, Duane Wessels. DNS Glue Requirements in Referral Responses. <https://www.ietf.org/archive/id/draft-ietf-dnsop-glue-is-not-optional-09.html>, 2023.
- [51] Microsoft DNS. Domain Name System (DNS) Docs. <https://docs.microsoft.com/en-us/windows-server/networking/dns/dns-top>, 2023.
- [52] Paul V. Mockapetris. Domain names - concepts and facilities. *RFC*, 1987.
- [53] Neustar. UltraDNS Public. <https://www.publicdns.neustar/>, 2023.
- [54] NLnet Labs. UNBOUND. <https://nlnetlabs.nl/projects/unbound/about/>, 2023.
- [55] Vasileios Pappas, Duane Wessels, Daniel Massey, Songwu Lu, Andreas Terzis, and Lixia Zhang. Impact of configuration errors on DNS robustness. *IEEE J. Sel. Areas Commun.*, 2009.
- [56] Eric Pauley, Ryan Sheatsley, Blaine Hoak, Quinn Burke, Yohan Beugin, and Patrick McDaniel. Measuring and mitigating the risk of ip reuse on public clouds. In *2022 IEEE Symposium on Security and Privacy (SP)*, 2022.
- [57] Victor Le Pochat, Tom van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczynski, and Wouter Joosen. Tranco: A research-oriented top sites ranking hardened against manipulation. In *26th Annual Network and Distributed System Security Symposium*, 2019.
- [58] PowerDNS. PowerDNS. <https://doc.powerdns.com/recursor/>, 2023.
- [59] Quad9 DNS. Quad9 DNS. <https://doc.powerdns.com/recursor/>, 2023.
- [60] Audrey Randall, Enze Liu, Gautam Akiwate, Ramakrishna Padmanabhan, Geoffrey M. Voelker, Stefan Savage, and Aaron Schulman. Trufflehunter: Cache snooping rare domains at large public DNS resolvers. In *IMC '20: ACM Internet Measurement Conference*, 2020.
- [61] Artur Romão. Tools for DNS debugging. *RFC*, 1994.
- [62] Simple DNS Plus. DNS. <https://simplifiedns.plus/download>, 2023.
- [63] Raffaele Sommese, Mattijs Jonker, Roland van Rijswijk-Deij, Alberto Dainotti, Kimberly C Claffy, and Anna Sperotto. The forgotten side of dns: Orphan and abandoned records. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, 2020.
- [64] Raffaele Sommese, Giovane C. M. Moura, Mattijs Jonker, Roland van Rijswijk-Deij, Alberto Dainotti, Kimberly C. Claffy, and Anna Sperotto. When parents and children disagree: Diving into DNS delegation inconsistency. In *Passive and Active Measurement - 21st International Conference, PAM*, 2020.
- [65] Technitium. Technitium DNS. <https://technitium.com/dns/>, 2023.
- [66] Tencent. DNSPod Public DNS+. <https://www.dnspod.com/>, 2023.
- [67] Verisign. Verisign Public DNS. <https://www.publicdns.xyz/public/verisign.html>, 2023.
- [68] Zheng Wang. The availability and security implications of glue in the domain name system. *CoRR*, abs/1605.01394, 2016.

- [69] Yandex. Yandex.DNS. <https://dns.yandex.com/>, 2023.
- [70] Mingming Zhang, Xiang Li, Baojun Liu, Jianyu Lu, Yiming Zhang, Jianjun Chen, Haixin Duan, Shuang Hao, and Xiaofeng Zheng. Detecting and measuring security risks of hosting-based dangling domains. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2023.
- [71] Maya Ziv, Liz Izhikevich, Kimberly Ruth, Katherine Izhikevich, and Zakir Durumeric. Asdb: a system for classifying owners of autonomous systems. In *Proceedings of the 21st ACM Internet Measurement Conference*, 2021.

A Common TCP and UDP ports

We selected 13 UDP and 20 TCP ports associated with common protocols, as shown in Table 6. These ports have been utilized in previous studies for the purpose of verifying whether an IP address is currently in use [14].

Table 6: Most frequently used TCP and UDP ports [14]

Protocol (port)	
UDP	FTP (21), SSH (22, 2222, 22022), Telnet (23), SMTP (25, 587), DNS (53), Kerberos (88), POP3 (110), IMAP (143), LDAP (389), MYSQL (3306)
TCP	FTP (21), SSH (22, 2222, 22022), Telnet (23), SMTP (25, 587), WHOIS (43), DNS (53), HTTP (80,8000,8080), Kerberos (88), POP3 (110), IMAP (143), LDAP (389), HTTPS (443, 8443), MS SQL (1433), MYSQL (3306)

B Hijacking scope of affected domains

Figure 12 displays the cumulative distribution of the hijack ratio for affected domains. For the majority of affected domains, attackers can capture half of their traffic.




Figure 12: Hijacking ratio for each domain. Half of the traffic for 84% of affected domains can be taken over by attackers.