

DNS Operations (dnsop)
Internet-Draft
Intended status: Best Current Practice
Expires: 1 September 2026

X.L. Li
Y.Q. Qiu
Nankai University
M.M. Zhang
Zhongguancun Laboratory
28 February 2026

Best Practices for Handling Deeply Nested Domain Delegations in
Recursive Resolvers
draft-li-dnsop-deep-delegation-scrutiny-01

Abstract

The Domain Name System (DNS) relies on caching to ensure its scalability and performance. However, certain behaviors in the DNS delegation and caching mechanisms can be exploited to circumvent domain name revocation. A recently discovered attack, named PHOENIX DOMAIN T2, allows a malicious domain that has been revoked at its parent zone to remain resolvable for an extended period. This is achieved by creating a chain of deeply nested subdomains, effectively keeping the delegation information alive in recursive resolver caches.

This document describes the PHOENIX DOMAIN T2 attack mechanism and proposes a set of operational best practices for DNS recursive resolver operators to mitigate this threat. The primary recommendation is for resolvers to apply additional scrutiny to domain names with an excessive number of labels, which is a key characteristic of this attack.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 1 September 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Background on Domain Name Revocation	2
1.2. The PHOENIX DOMAIN T2 Vulnerability and Related Threats	3
1.3. Requirements Language	4
2. The PHOENIX DOMAIN T2 Attack Mechanism	4
3. Operational Recommendations	5
3.1. Identifying Excessively Deep Domains	5
3.2. Applying Scrutiny	5
3.2.1. Preferential Re-validation	5
3.2.2. Restricting Maximum Cache TTL	6
3.2.3. Monitoring and Logging	6
4. Security Considerations	6
5. IANA Considerations	7
6. References	7
6.1. Normative References	7
6.2. Informative References	7
Authors' Addresses	8

1. Introduction

1.1. Background on Domain Name Revocation

The DNS namespace is structured as a hierarchy, where parent zones delegate authority for sub-zones (children) via NS (Name Server) records [RFC1034] [RFC1035]. Domain name revocation is the process of removing or altering these delegation records in the parent zone to render a domain and its subdomains unresolvable or to redirect them to a sinkhole. This is a critical tool for combating malicious activities such as botnets, phishing, and malware distribution.

However, due to the distributed and heavily cached nature of DNS, revocation is not instantaneous. Resolvers will continue to use their cached delegation information until its Time-to-Live (TTL) expires. The "Ghost Domain" attack first demonstrated that a revoked domain's authoritative server could trick resolvers into refreshing the TTL of cached NS records, keeping the domain alive post-revocation [Jiang2012]. While mitigations for this specific attack were widely deployed, the underlying principles of DNS caching still present opportunities for exploitation.

1.2. The PHOENIX DOMAIN T2 Vulnerability and Related Threats

A more recent study introduced the PHOENIX DOMAIN T2 attack [Li2023], which bypasses existing defenses. This attack does not violate any DNS protocol specifications but rather exploits a de facto standard behavior in how resolvers search their cache. When a resolver needs to resolve a name for which it has no specific cached data (a cache miss), it searches for the "closest known nameserver" by performing a longest suffix match on the queried name against its cache, a behavior clarified in documents such as [RFC2181].

The T2 attack leverages this behavior by creating a chain of single-label subdomains. An attacker can make a resolver cache a new delegation for a deep subdomain (e.g., "s.s.botnet.com") based on the soon-to-expire cached delegation of its parent ("s.botnet.com"). This process can be repeated, potentially up to the maximum label count of 127, allowing a revoked domain to remain effectively resolvable for months or even years. Because this attack exploits a fundamental aspect of cache searching, it affects virtually all mainstream DNS resolver implementations.

The PHOENIX DOMAIN attack is a prime example of vulnerabilities arising from the temporal desynchronization between DNS caching and delegation state. This class of issues is not isolated. A related, widespread problem is that of "dangling domains," where DNS resource records (e.g., CNAME or A records) point to third-party hosting service endpoints that have been deprovisioned but are available for others to claim. As detailed by Zhang et al. [Zhang2023], attackers can take over these dangling domains by registering the same service endpoints, thereby hijacking legitimate traffic for malicious purposes. While the PHOENIX DOMAIN attack exploits the delegation chain, the dangling domain problem exploits the resolution endpoint. Both highlight a critical principle: the security of the DNS ecosystem relies on DNS records accurately reflecting the current, authoritative state of domain control and resources.

This document acknowledges that the diversity of DNS implementations is a strength and not a weakness. The exact mitigations detailed herein are provided as operational guidance and Best Current Practices, rather than rigid Internet Standards. Implementers are encouraged to adapt these mechanisms to suit their specific architectures.

1.3. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. The PHOENIX DOMAIN T2 Attack Mechanism

The attack proceeds in the following steps:

1. Initial State: The attacker controls a domain, e.g., "botnet.com", and its authoritative nameserver. A target recursive resolver is made to query for "botnet.com", thereby caching its NS records with a given TTL (e.g., 24 hours).
2. Revocation: The domain is revoked at the parent registry (e.g., the ".com" TLD). The NS records for "botnet.com" are removed from the ".com" zone. The resolver, however, still holds the original NS records in its cache.
3. Exploitation: Shortly before the 24-hour TTL expires in the resolver's cache, the attacker triggers a query to the resolver for a previously unseen subdomain, e.g., "s.botnet.com".
4. Cache Logic Exploited: The resolver experiences a cache miss for "s.botnet.com". It then performs a longest suffix match and finds the still-valid (though soon-to-expire) NS records for "botnet.com". It sends a query for "s.botnet.com" to the attacker's authoritative nameserver listed in these cached records.
5. Delegation Refresh: The attacker's nameserver responds with a new delegation for "s.botnet.com", pointing to another nameserver (or itself) and providing NS and glue records. This response has a fresh, long TTL (e.g., another 24 hours). The resolver caches this new delegation for "s.botnet.com".

6. Iteration: The attacker can repeat this process. Shortly before the TTL for "s.botnet.com" expires, they can trigger a query for "s.s.botnet.com", and so on. Each step adds a label and refreshes the cache lifetime, effectively keeping subdomains of the revoked "botnet.com" continuously resolvable.

3. Operational Recommendations

To disrupt the chain of delegation renewals central to the PHOENIX DOMAIN T2 attack, recursive resolver operators SHOULD implement mechanisms to identify and apply additional scrutiny to DNS queries for a domain name with a high number of labels.

3.1. Identifying Excessively Deep Domains

Research indicates that the vast majority of legitimate domains used in practice are not deeply nested [Li2023]. A study of DNS responses for Alexa Top 1 Million domains found that over 99% of observed names have 10 or fewer labels.

Therefore, it is RECOMMENDED that resolvers use a threshold of 10 labels as a heuristic to identify potentially suspicious domain names that warrant further scrutiny. This threshold MAY be configurable by the operator.

3.2. Applying Scrutiny

When a resolver receives a query for a domain name that exceeds the configured label-count threshold, it SHOULD perform one or more of the following actions:

3.2.1. Preferential Re-validation

Instead of immediately trusting the delegation information provided by a child zone's nameserver (especially when that information is already in cache), the resolver SHOULD perform a re-validation query to the parent zone. For a query to "label10...label11.example.com", the resolver should explicitly query the nameservers for "example.com" to validate the delegation for "label11.example.com". This action directly confirms whether the delegation is still valid from the parent's perspective, effectively stopping the T2 attack.

3.2.2. Restricting Maximum Cache TTL

If re-validation is not performed, resolvers SHOULD enforce a significantly reduced maximum TTL on any resource records received for an excessively deep domain. For example, an operator could configure that any records for a domain with more than 10 labels, regardless of the TTL provided by the authoritative server, will not be cached for more than a short duration (e.g., 1 hour). This would not prevent the attack but would drastically reduce its longevity and increase the cost for the attacker.

3.2.3. Monitoring and Logging

Queries for domain names with an unusually high number of labels SHOULD be logged with a distinct severity level. This provides network administrators with visibility into potential T2 attack activity within their networks, enabling further investigation or blocking.

4. Security Considerations

This entire document is focused on improving the security and resilience of the DNS ecosystem against cache-based attacks that undermine domain name revocation.

The recommendations provided herein directly mitigate the PHOENIX DOMAIN T2 attack by disrupting its core mechanism of iterative cache refreshing for deep subdomains. By applying scrutiny based on label count, resolvers can more effectively enforce the intent of domain revocation actions taken at parent zones.

A potential risk is the impact on legitimate services that rely on deeply nested domain names. However, as supported by measurement studies [Li2023], such domains are exceptionally rare in common usage. The recommended actions, particularly preferential re-validation, would not break resolution for legitimate deep domains but would simply ensure their delegation is valid. The ability for operators to configure the label threshold provides a mechanism to balance security with any specific local operational needs.

It is important to note that the operational recommendations in this document are specifically designed to mitigate the PHOENIX DOMAIN T2 attack vector. However, resolver operators and domain owners must recognize that robust DNS security requires a multi-faceted approach. Other significant threats, such as the hosting-based domain takeovers described in [Zhang2023], depend not on deep delegations but on diligent DNS record management and proper domain ownership validation by hosting providers. While the scrutiny of deep domains improves

resolver resilience against delegation-based attacks, it does not replace the fundamental need for domain owners to purge stale DNS records pointing to deprovisioned third-party services.

5. IANA Considerations

This document has no IANA actions.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

6.2. Informative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/rfc/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/rfc/rfc1035>>.
- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, DOI 10.17487/RFC2181, July 1997, <<https://www.rfc-editor.org/rfc/rfc2181>>.
- [Jiang2012] Jiang, J., Liang, J., Li, K., Li, J., Duan, H., and J. Wu, "Ghost Domain Names: Revoked Yet Still Resolvable", in Proceedings of the 19th Annual Network and Distributed System Security Symposium (NDSS '12) , February 2012.
- [Li2023] Li, X., Liu, B., Bai, X., Zhang, M., Zhang, Q., Li, Z., Duan, H., and Q. Li, "Ghost Domain Reloaded: Vulnerable Links in Domain Name Delegation and Revocation", in Proceedings of the Network and Distributed Systems Security (NDSS) Symposium 2023 , February 2023, <<https://dx.doi.org/10.14722/ndss.2023.23005>>.

[Zhang2023]

Zhang, M., Li, X., Liu, B., Lu, J., Zhang, Y., Chen, J., Duan, H., Hao, S., and X. Zheng, "Detecting and Measuring Security Risks of Hosting-Based Dangling Domains", Proc. ACM Meas. Anal. Comput. Syst. Vol. 7, No. 1, Article 9, March 2023, <<https://doi.org/10.1145/3579440>>.

Authors' Addresses

Xiang Li
Nankai University
38 Tongyan Road
Tianjin
Tianjin, 300355
China
Email: lixiang@nankai.edu.cn

Yuqi Qiu
Nankai University
38 Tongyan Road
Tianjin
Tianjin, 300355
China
Email: qiuyuqi@mail.nankai.edu.cn

Mingming Zhang
Zhongguancun Laboratory
Beijing
Beijing, 100081
China
Email: zhangmm@mail.zgclab.edu.cn