

DNS Operations
Internet-Draft
Updates: 7871 (if approved)
Intended status: Informational
Expires: 1 September 2026

X.L. Li
Y.Q. Qiu
Nankai University
28 February 2026

Strengthening DNS Query Aggregation against ECS-based Attacks
draft-li-dnsop-ecs-aggregation-fix-01

Abstract

The DNS query aggregation mechanism is a critical defense against DNS cache poisoning attacks that exploit the "Birthday Paradox". However, recent research has revealed that flawed implementations of the EDNS Client Subnet (ECS) option, as specified in RFC 7871, can be exploited to bypass this defense. This allows attackers to force a resolver to issue multiple simultaneous queries for the same domain name by crafting queries with different ECS options. This vulnerability revives the classic DNS Birthday Attack, posing a significant threat to DNS resolvers and the clients they serve.

While Section 11.2 of RFC 7871 notes the general risk of Birthday Attacks and suggests marking whether responding nameservers support ECS, it stops short of providing a concrete normative state machine or explicit mitigation process. This document specifies a stricter and more coherent processing model for the ECS option in DNS resolvers to formalize that detection concept. It introduces a mechanism for resolvers to track the ECS support state of authoritative nameservers ("no-ECS-support" state) and mandates strict query aggregation for zones determined to not support ECS. These changes are designed to close the identified vulnerability and restore the effectiveness of query aggregation as a practical defense mechanism.

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. Requirements Language | 4 |
| 2. The REBIRTHDAY Vulnerability | 4 |
| 3. Updated Resolver Behavior | 5 |
| 3.1. Principle of ECS Coherence | 5 |
| 3.2. Handling Responses Containing an ECS Option | 5 |
| 3.3. Handling Responses Lacking an ECS Option | 5 |
| 3.4. Updated Query Aggregation Logic | 5 |
| 4. Security Considerations | 6 |
| 5. IANA Considerations | 6 |
| 6. References | 6 |
| 6.1. Normative References | 7 |
| 6.2. Informative References | 7 |
| Authors' Addresses | 8 |

1. Introduction

The Domain Name System (DNS) [RFC1034] [RFC1035] is a fundamental component of the Internet, but its security has been a long-standing concern, with DNS cache poisoning [HITCHHIKER] remaining a persistent threat. Following the discovery of widespread vulnerabilities to off-path attacks [KAM_ATTACK], a number of defenses were developed. One of the earliest and most effective of these is the query aggregation mechanism implemented in DNS resolvers. This mechanism ensures that multiple identical incoming queries (matching in qname

and qtype) result in only a single query being sent to an authoritative nameserver. This thwarts attackers seeking to increase their chances of success by forcing many simultaneous outbound queries, a technique known as the DNS Birthday Attack [SACRAMENTO].

The EDNS Client Subnet (ECS) extension [RFC7871] was introduced to provide network location information of the client to authoritative nameservers, enabling them to provide more geographically or topologically relevant answers. While beneficial for performance, the processing logic for ECS has introduced a subtle but critical security flaw.

Research, such as the "REBIRTHDAY Attack" [REBIRTHDAY], has demonstrated that many resolver implementations expand their query aggregation key from <qname, qtype> to <qname, qtype, subnet>. An attacker can exploit this behavior by sending a large number of queries for the same qname and qtype but with different, spoofed client subnets. A vulnerable resolver will fail to aggregate these requests and will instead issue a multitude of distinct queries upstream.

A crucial aspect of this vulnerability is that [RFC7871] specifies that a response from an authoritative server without an ECS option is still a valid reply to a query that contained one. This allows an attacker to inject a single type of forged response (lacking an ECS option) that has a high probability of matching one of the many outstanding queries, thereby poisoning the resolver's cache.

Section 11.2 of [RFC7871] correctly identifies that an attacker "just needs to flood back responses that have no ECS option at all" and suggests that this could be detected "by marking whether the responding nameserver had previously been sending ECS options." However, it does not provide an actionable, specified mechanism to leverage this observation.

This document aims to formalize that observation into a concrete defense mechanism. It rectifies the vulnerability by defining stricter rules for ECS handling. It requires resolvers to maintain explicit state regarding an authoritative server's support for ECS ("no-ECS-support") and to enforce strict aggregation based on this state, effectively neutralizing the attack vector and fulfilling the defensive concept initially touched upon in RFC 7871.

1.1. 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 REBIRTHDAY Vulnerability

The vulnerability stems from the interaction of two factors: a flaw in resolver query aggregation logic and a permissiveness in the ECS specification. The attack proceeds as follows:

1. Query Initiation: The attacker sends a high volume of DNS queries for a target domain (e.g., "victim.example.com A") to a vulnerable resolver. Each query is identical in qname and qtype but contains a unique, forged ECS option (e.g., with different subnets).
2. Bypassing Aggregation: The vulnerable resolver treats each query as unique due to the different ECS options. Instead of sending a single query upstream, it forwards a large number of queries to the authoritative nameserver for "example.com", each with a different source port.
3. Race Condition: The attacker simultaneously floods the resolver with forged DNS responses for "victim.example.com". Crucially, these forged responses do not contain an ECS option. The attacker brute-forces the 16-bit Transaction ID (TxID) and attempts to guess one of the many source ports used by the resolver.
4. Cache Injection: According to [RFC7871], a response without an ECS option is a valid reply. If the attacker's forged response matches the TxID and source port of any of the outstanding queries, the resolver accepts it. The malicious record is then cached, and subsequent legitimate queries for that domain name will receive the poisoned data.

The Birthday Paradox dictates that the probability of a successful collision increases dramatically with the number of outstanding queries, making this attack highly practical against vulnerable implementations.

3. Updated Resolver Behavior

3.1. Principle of ECS Coherence

To mitigate this attack, resolvers MUST handle ECS options coherently. Coherence means that a resolver's caching and query aggregation behavior MUST account for whether an authoritative server for a given zone actually uses the ECS information provided to it. A resolver MUST NOT indefinitely treat queries as unique based solely on varying ECS options if there is evidence that the authoritative server does not support or utilize ECS for that zone.

3.2. Handling Responses Containing an ECS Option

When a resolver receives a response that contains an ECS option, it MUST validate the fields in the response against the query that was sent, as described in Section 7.3.1 of [RFC7871]. If there is a mismatch in the Family, Address, or Source Prefix-Length, the entire response SHOULD be discarded. This is existing best practice and is reiterated here for completeness.

3.3. Handling Responses Lacking an ECS Option

This is the critical change to resolver logic.

When a resolver sends a query containing an ECS option to an authoritative nameserver and receives a valid response that lacks an ECS option, the resolver SHOULD treat this as an indication that the authoritative server for this zone does not support or did not use ECS to generate this specific response.

The resolver SHOULD cache this "no-ECS-support" state associated with the queried domain name's zone. The lifetime of this cached state MAY be tied to the TTL of the zone's NS records or a locally configured timer.

3.4. Updated Query Aggregation Logic

Resolver query aggregation logic MUST be updated to incorporate the "no-ECS-support" state.

1. If a resolver receives multiple client queries for the same <qname, qtype> that have different ECS options, and there is no active query or existing "no-ECS-support" state for that zone, the resolver MAY forward one of these queries with an ECS option and queue the others.

2. If the resolver has a cached "no-ECS-support" state for the zone of a given qname, any incoming client query for that qname, regardless of the presence or value of its ECS option, MUST be aggregated with other pending queries for the same <qname, qtype>.
3. When sending an aggregated query for a zone in the "no-ECS-support" state, the resolver SHOULD NOT include an ECS option in the upstream query. This prevents the authoritative server from potentially providing a subnet-specific answer that would be incorrectly served to clients from other subnets whose queries were aggregated.

By implementing this logic, a resolver will quickly learn that a zone does not support ECS and will revert to classic query aggregation based solely on <qname, qtype>, thus defeating the REBIRTHDAY attack.

4. Security Considerations

The procedures outlined in this document are intended to mitigate a specific DNS cache poisoning vulnerability. By enforcing coherent ECS handling and stateful query aggregation, resolvers can prevent attackers from bypassing a fundamental security mechanism. This restores the difficulty of off-path birthday attacks to their pre-ECS levels, which are considered computationally infeasible for a properly configured resolver.

This document does not address all possible DNS cache poisoning vectors. The use of strong entropy for TxID and source port randomization remains essential. Furthermore, the deployment of DNSSEC [RFC4033] provides cryptographic assurance of data integrity and is the most robust defense against cache poisoning attacks. The mechanisms in this document are complementary to, not a replacement for, DNSSEC.

Implementers should be careful about the cache lifetime for the "no-ECS-support" state. A value that is too long could prevent a zone that newly adopts ECS from providing optimized responses. A value that is too short might be cleared by an attacker, re-opening a small window for an attack. Tying the lifetime to the zone's NS record TTL is RECOMMENDED as a reasonable balance.

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>>.
- [RFC7871] Contavalli, C., van der Gaast, W., Lawrence, D., and W. Kumari, "Client Subnet in DNS Queries", RFC 7871, DOI 10.17487/RFC7871, May 2016, <<https://www.rfc-editor.org/info/rfc7871>>.

6.2. Informative References

- [REBIRTHDAY] Li, X., Zhang, M., Xu, Z., Miao, F., Qiu, Y., Liu, B., Zhang, J., Zheng, X., Duan, H., Liu, Z., Zhang, Y., and D. Fan, "RebirthDay Attack: Reviving DNS Cache Poisoning with the Birthday Paradox", In Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS '25) , 13 October 2025.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [SACRAMENTO] Sacramento, V., "Vulnerability in Requests Control of BIND Versions 4 and 8 Allows DNS Spoofing", November 2002, <<https://lists.isc.org/pipermail/bind-users/2002-November/043141.html>>.
- [KAM_ATTACK] Kaminsky, D., "It's the End of the Cache as We Know It", August 2008.

[HITCHHIKER]

Son, S. and V. Shmatikov, "The Hitchhiker's Guide to DNS Cache Poisoning", In Proceedings of the 6th International ICST Conference on Security and Privacy in Communication Systems (SecureComm '10) , September 2010.

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