

Network Working Group
Internet-Draft
Updates: [4035](#) (if approved)
Intended status: Standards Track
Expires: April 7, 2017

K. Fujiwara
JPRS
A. Kato
Keio/WIDE
W. Kumari
Google
October 4, 2016

**Aggressive use of NSEC/NSEC3
draft-ietf-dnsop-nsec-aggressiveuse-03**

Abstract

The DNS relies upon caching to scale; however, the cache lookup generally requires an exact match. This document specifies the use of NSEC/NSEC3 resource records to allow DNSSEC validating resolvers to generate negative answers within a range. This increases performance / decreases latency, decreases resource utilization on both authoritative and recursive servers, and also increases privacy. It may also help increase resilience to certain DoS attacks in some circumstances.

This document updates [RFC4035](#) by allowing validating resolvers to generate negative answers based upon NSEC/NSEC3 records.

[Ed note: Text inside square brackets ([]) is additional background information, answers to frequently asked questions, general musings, etc. They will be removed before publication. This document is being collaborated on in Github at: <https://github.com/wkumari/draft-ietf-dnsop-nsec-aggressiveuse>. The most recent version of the document, open issues, etc should all be available here. The authors (gratefully) accept pull requests.]

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 <http://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 April 7, 2017.

Copyright Notice

Copyright (c) 2016 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	3
3.	Problem Statement	3
4.	Background	4
5.	Aggressive Negative Caching	5
5.1.	NSEC	5
5.2.	NSEC3	6
5.3.	Consideration on TTL	6
6.	Benefits	6
7.	Update to RFC 4035	7
8.	IANA Considerations	7
9.	Security Considerations	8
10.	Implementation Status	8
11.	Acknowledgments	8
12.	Change History	8
12.1.	Version draft-fujiwara-dnsop-nsec-aggressiveuse-01	10
12.2.	Version draft-fujiwara-dnsop-nsec-aggressiveuse-02	10
12.3.	Version draft-fujiwara-dnsop-nsec-aggressiveuse-03	10
13.	References	11
13.1.	Normative References	11
13.2.	Informative References	11
Appendix A.	Detailed implementation notes	12
	Authors' Addresses	12

[1.](#) Introduction

A DNS negative cache exists, and is used to cache the fact that a name does not exist. This method of negative caching requires exact matching; this leads to unnecessary additional lookups, increases

latency, leads to extra resource utilization on both authoritative and recursive servers, and decreases privacy by leaking queries.

This document updates [RFC 4035](#) to allow recursive resolvers to use NSEC/NSEC3 resource records to synthesize negative answers from the information they have in the cache. This allows validating resolvers to respond with NXDOMAIN immediately if the name in question falls into a range expressed by a NSEC/NSEC3 resource record already in the cache.

Aggressive Negative Caching was first proposed in [Section 6](#) of DNSSEC Lookaside Validation (DLV) [[RFC5074](#)] in order to find covering NSEC records efficiently.

Section 3 of [[I-D.vixie-dnsext-resimprove](#)] "Stopping Downward Cache Search on NXDOMAIN" and [[I-D.ietf-dnsop-nxdomain-cut](#)] proposed another approach to use NXDOMAIN information effectively.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Many of the specialized terms used in this document are defined in DNS Terminology [[RFC7719](#)].

The key words "Closest Encloser" and "Source of Synthesis" in this document are to be interpreted as described in [[RFC4592](#)].

"Closest Encloser" is also defined in NSEC3 [[RFC5155](#)], as is "Next closer name".

3. Problem Statement

The DNS negative cache caches negative (non-existent) information, and requires an exact match in most instances [[RFC2308](#)].

Assume that the (DNSSEC signed) "example.com" zone contains:

apple.example.com IN A 192.0.2.1

elephant.example.com IN A 192.0.2.2

zebra.example.com IN A 192.0.2.3

If a validating resolver gets a query for cat.example.com, it will query the example.com servers and will get back an NSEC (or NSEC3)

record stating that there are no records between apple and elephant. The resolver then knows that cat.example.com does not exist; however, it does not use the fact that the proof covers a range (apple to elephant) to suppress queries for other labels that fall within this range. This means that if the validating resolver gets a query for ball.example.com (or dog.example.com) it will once again go off and query the example.com servers for these names.

Apart from wasting bandwidth, this also wastes resources on the recursive server (it needs to keep state for outstanding queries), wastes resources on the authoritative server (it has to answer additional questions), increases latency (the end user has to wait longer than necessary to get back an NXDOMAIN answer), can be used by attackers to cause a DoS (see additional resources), and also has privacy implications (e.g: typos leak out further than necessary).

4. Background

DNSSEC [[RFC4035](#)] and [[RFC5155](#)] both provide "authenticated denial of existence"; this is a cryptographic proof that the queried for name does not exist, accomplished by providing a (DNSSEC secured) record containing the names which appear alphabetically before and after the queried for name. In the example above, if the (DNSSEC validating) recursive server were to query for lion.example.com it would receive a (signed) NSEC/NSEC3 record stating that there are no labels between "elephant" and "zebra". This is a signed, cryptographic proof that these names are the ones before and after the queried for label. As lion.example.com falls within this range, the recursive server knows that lion.example.com really does not exist. This document specifies that this NSEC/NSEC3 record should be used to generate negative answers for any queries that the recursive server receives that fall within the range covered by the record (for the TTL for the record).

[Section 4.5 of \[RFC4035\]](#) says:

"In theory, a resolver could use wildcards or NSEC RRs to generate positive and negative responses (respectively) until the TTL or signatures on the records in question expire. However, it seems prudent for resolvers to avoid blocking new authoritative data or synthesizing new data on their own. Resolvers that follow this recommendation will have a more consistent view of the namespace." and "The reason for these recommendations is that, between the initial query and the expiration of the data from the cache, the authoritative data might have been changed (for example, via dynamic update)". In other words, if a resolver generates negative answers from an NSEC record, it will not send any queries for names within that NSEC range (for the TTL). If a new name is added to the zone during this interval the resolver will not know this.

We believe this recommendation can be relaxed because, in the absence of this technique, a lookup for the exact name could have come in during this interval, and so this could already be cached (see [\[RFC2308\]](#) for more background). This means that zone operators should have no expectation that an added name would work immediately. With DNSSEC and Aggressive NSEC, the TTL of the NSEC record is the authoritative statement of how quickly a name can start working within a zone.

5. Aggressive Negative Caching

[Section 4.5 of \[RFC4035\]](#) shows that "In theory, a resolver could use wildcards or NSEC RRs to generate positive and negative responses (respectively) until the TTL or signatures on the records in question expire. However, it seems prudent for resolvers to avoid blocking new authoritative data or synthesizing new data on their own. Resolvers that follow this recommendation will have a more consistent view of the namespace".

This document relaxes this restriction, as follows:

```
+-----+
| Once the records are validated, DNSSEC enabled validating |
| resolvers SHOULD use NSEC/NSEC3 resource records to generate |
| negative responses until their effective TTLs or signatures |
| for those records expire. |
+-----+
```

If the validating resolver's cache has sufficient information to validate the query, the resolver SHOULD use NSEC/NSEC3/wildcard records aggressively. Otherwise, it MUST fall back to send the query to the authoritative DNS servers.

If the query name has the matching NSEC/NSEC3 RR proving the information requested does not exist, the validating resolver may respond with a NODATA (empty) answer.

5.1. NSEC

Implementations which support aggressive use of NSEC SHOULD enable this by default. Implementations MAY provide a configuration switch to disable aggressive use of NSEC and allow it to be enabled or disabled per domain.

The validating resolver needs to check the existence of an NSEC RR matching/covering the source of synthesis and an NSEC RR covering the query name.

If the validating resolver's cache contains an NSEC RR covering the source of synthesis and the covering NSEC RR of the query name, the validating resolver may respond with NXDOMAIN error immediately.

5.2. NSEC3

NSEC3 aggressive negative caching is more difficult. If the zone is signed with NSEC3, the validating resolver needs to check the existence of non-terminals and wildcards which derive from query names.

If the validating resolver's cache contains an NSEC3 RR matching the closest enclosure, an NSEC3 RR covering the next closer name, and an NSEC3 RR covering the source of synthesis, it is possible for the validating resolver to respond with NXDOMAIN immediately.

If a covering NSEC3 RR has Opt-Out flag, the covering NSEC3 RR does not prove the non-existence of the domain name and the aggressive negative caching is not possible for the domain name.

A validating resolver implementation MAY support aggressive use of NSEC3. If it does aggressive use of NSEC3, it MAY provide a configuration switch to disable aggressive use of NSEC3 and allow it to be enabled or disabled for specific zones.

5.3. Consideration on TTL

The TTL value of negative information is especially important, because newly added domain names cannot be used while the negative information is effective. [Section 5 of RFC 2308](#) states that the maximum number of negative cache TTL value is 3 hours (10800). It is RECOMMENDED that validating resolvers limit the maximum effective TTL value of negative responses (NSEC/NSEC3 RRs) to this same value.

6. Benefits

The techniques described in this document provide a number of benefits, including (in no specific order):

Reduced latency By answering directly from cache, validating resolvers can immediately inform clients that the name they are looking for does not exist, improving the user experience.

Decreased recursive server load By answering negative queries from the cache, validating servers avoid having send a query and wait for a response. In addition to decreasing the bandwidth used, it also means that the server does not need to allocate and maintain state, thereby decreasing memory and CPU load.

Decreased authoritative server load Because recursive servers can answer (negative) queries without asking the authoritative server, the authoritative servers receive less queries. This decreases the authoritative server bandwidth, queries per second and CPU utilization.

The scale of the benefit depends upon multiple factors, including the query distribution. For example, currently around 65% of queries to Root Name servers result in NXDOMAIN responses; this technique will eliminate a sizable quantity of these.

The technique described in this document may also mitigate so-called "random QNAME attacks", in which attackers send many queries for random sub-domains to resolvers. As the resolver will not have the answers cached it has to ask external servers for each random query, leading to a DoS on the authoritative servers (and often resolvers). Aggressive NSEC may help mitigate these attacks by allowing the resolver to answer directly from cache for any random queries which fall within already requested ranges. It will not always work as an effective defense, not least because not many zones are DNSSEC signed at all, but it will still provide an additional layer of defense.

7. Update to [RFC 4035](#)

[Section 4.5 of \[RFC4035\]](#) shows that "In theory, a resolver could use wildcards or NSEC RRs to generate positive and negative responses (respectively) until the TTL or signatures on the records in question expire. However, it seems prudent for resolvers to avoid blocking new authoritative data or synthesizing new data on their own. Resolvers that follow this recommendation will have a more consistent view of the namespace".

The paragraph is updated as follows:

```
+-----+
| Once the records are validated, DNSSEC enabled validating |
| resolvers MAY use wildcards and NSEC/NSEC3 resource records |
| to generate negative responses until their effective TTLs |
| or signatures for those records expire. |
+-----+
```

8. IANA Considerations

This document has no IANA actions.

9. Security Considerations

Newly registered resource records may not be used immediately. However, choosing suitable TTL value and negative cache TTL value (SOA MINIMUM field) will mitigate the delay concern, and it is not a security problem.

It is also suggested to limit the maximum TTL value of NSEC / NSEC3 resource records in the negative cache to, for example, 10800 seconds (3hrs), to mitigate this issue. Implementations which comply with this proposal are recommended to have a configurable maximum value of NSEC RRs in the negative cache.

Aggressive use of NSEC / NSEC3 resource records without DNSSEC validation may create serious security issues, and so this technique requires DNSSEC validation.

10. Implementation Status

Unbound currently implements aggressive negative caching, as does Google Public DNS.

11. Acknowledgments

The authors gratefully acknowledge DLV [[RFC5074](#)] author Samuel Weiler and the Unbound developers.

The authors would like to specifically thank Tatuya JINMEI for extensive review and comments, and also Mark Andrews, Stephane Bortzmeyer, Casey Deccio, Alexander Dupuy, Olafur Gudmundsson, Bob Harold, Shumon Huque, John Levine, Pieter Lexis and Matthijs Mekking (who even sent pull requests!).

12. Change History

RFC Editor: Please remove this section prior to publication.

-02 to -03:

- o Integrated a bunch of comments from Matthijs Mekking - details in: <https://github.com/wkumari/draft-ietf-dnsop-nsec-aggressiveuse/pull/1>. I decided to keep "Aggressive Negative Caching" instead of "Aggressive USE OF Negative Caching" for readability.
- o Attempted to address Bob Harold's comment on the readability issues with "But, it will be more effective when both are enabled..." in [Section 5.4](#) - <https://www.ietf.org/mail-archive/web/dnsop/current/msg17997.html>

- o MAYs and SHOULD drifted in the text block. Fixed - thanks to <https://mailarchive.ietf.org/arch/msg/dnsop/2ljmmzxtIMCFMLOZmWcSbTYV0y4>
- o A number of good edits from Stephane in: <https://www.ietf.org/mail-archive/web/dnsop/current/msg18109.html>
- o A bunch more edits from Jinmei, as in: <https://www.ietf.org/mail-archive/web/dnsop/current/msg18206.html>

-01 to -02:

- o Added [Section 6](#) - Benefits (as suggested by Jinmei).
- o Removed [Appendix B](#) (Jinmei)
- o Replaced "full-service" with "validating" (where applicable)
- o Integrated other comments from Jinmei from <https://www.ietf.org/mail-archive/web/dnsop/current/msg17875.html>
- o Integrated comment from co-authors, including re-adding parts of [Appendix B](#), terminology, typos.
- o Tried to explain under what conditions this may actually mitigate attacks.

-00 to -01:

- o Comments from DNSOP meeting in Berlin.
- o Changed intended status to Standards Track (updates [RFC 4035](#))
- o Added a section "Updates to [RFC 4035](#)"
- o Some language clarification / typo / cleanup
- o Cleaned up the TTL section a bit.
- o Removed Effects section, Additional proposal section, and pseudo code.
- o Moved "mitigation of random subdomain attacks" to Appendix.

From [draft-fujiwara-dnsop-nsec-aggressiveuse-03](#) -> [draft-ietf-dnsop-nsec-aggressiveuse](#)

- o Document adopted by DNSOP WG.

- o Adoption comments
- o Changed main purpose to performance
- o Use NSEC3/Wildcard keywords
- o Improved wordings (from good comments)
- o Simplified pseudo code for NSEC3
- o Added Warren as co-author.
- o Reworded much of the problem statement
- o Reworked examples to better explain the problem / solution.

12.1. Version [draft-fujiwara-dnsop-nsec-aggressiveuse-01](#)

- o Added reference to DLV [[RFC5074](#)] and imported some sentences.
- o Added Aggressive Negative Caching Flag idea.
- o Added detailed algorithms.

12.2. Version [draft-fujiwara-dnsop-nsec-aggressiveuse-02](#)

- o Added reference to [[I-D.vixie-dnsext-resimprove](#)]
- o Added considerations for the CD bit
- o Updated detailed algorithms.
- o Moved Aggressive Negative Caching Flag idea into Additional Proposals

12.3. Version [draft-fujiwara-dnsop-nsec-aggressiveuse-03](#)

- o Added "Partial implementation"
- o [Section 4](#),5,6 reorganized for better representation
- o Added NODATA answer in [Section 4](#)
- o Trivial updates
- o Updated pseudo code

13. References

13.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", [RFC 2308](#), DOI 10.17487/RFC2308, March 1998, <<http://www.rfc-editor.org/info/rfc2308>>.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", [RFC 4035](#), DOI 10.17487/RFC4035, March 2005, <<http://www.rfc-editor.org/info/rfc4035>>.
- [RFC4592] Lewis, E., "The Role of Wildcards in the Domain Name System", [RFC 4592](#), DOI 10.17487/RFC4592, July 2006, <<http://www.rfc-editor.org/info/rfc4592>>.
- [RFC5074] Weiler, S., "DNSSEC Lookaside Validation (DLV)", [RFC 5074](#), DOI 10.17487/RFC5074, November 2007, <<http://www.rfc-editor.org/info/rfc5074>>.
- [RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", [RFC 5155](#), DOI 10.17487/RFC5155, March 2008, <<http://www.rfc-editor.org/info/rfc5155>>.
- [RFC7719] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", [RFC 7719](#), DOI 10.17487/RFC7719, December 2015, <<http://www.rfc-editor.org/info/rfc7719>>.

13.2. Informative References

- [I-D.ietf-dnsop-nxdomain-cut]
Bortzmeyer, S. and S. Huque, "NXDOMAIN really means there is nothing underneath", [draft-ietf-dnsop-nxdomain-cut-03](#) (work in progress), May 2016.
- [I-D.vixie-dnsexst-resimprove]
Vixie, P., Joffe, R., and F. Neves, "Improvements to DNS Resolvers for Resiliency, Robustness, and Responsiveness", [draft-vixie-dnsexst-resimprove-00](#) (work in progress), June 2010.

Appendix A. Detailed implementation notes

- o Previously, cached negative responses were indexed by QNAME, QCLASS, QTYPE, and the setting of the CD bit (see [RFC 4035, Section 4.7](#)), and only queries matching the index key would be answered from the cache. With aggressive negative caching, the validator, in addition to checking to see if the answer is in its cache before sending a query, checks to see whether any cached and validated NSEC record denies the existence of the sought record(s). Using aggressive negative caching, a validator will not make queries for any name covered by a cached and validated NSEC record. Furthermore, a validator answering queries from clients will synthesize a negative answer whenever it has an applicable validated NSEC in its cache unless the CD bit was set on the incoming query. (Imported from [Section 6 of \[RFC5074\]](#)).
- o Implementing aggressive negative caching suggests that a validator will need to build an ordered data structure of NSEC and NSEC3 records for each signer domain name of NSEC / NSEC3 records in order to efficiently find covering NSEC / NSEC3 records. Call the table as NSEC_TABLE. (Imported from [Section 6.1 of \[RFC5074\]](#) and expanded.)
- o The aggressive negative caching may be inserted at the cache lookup part of the recursive resolvers.
- o If errors happen in aggressive negative caching algorithm, resolvers MUST fall back to resolve the query as usual. "Resolve the query as usual" means that the resolver must process the query as though it does not implement aggressive negative caching.

Authors' Addresses

Kazunori Fujiwara
Japan Registry Services Co., Ltd.
Chiyoda First Bldg. East 13F, 3-8-1 Nishi-Kanda
Chiyoda-ku, Tokyo 101-0065
Japan

Phone: +81 3 5215 8451
Email: fujiwara@jprs.co.jp

Akira Kato
Keio University/WIDE Project
Graduate School of Media Design, 4-1-1 Hiyoshi
Kohoku, Yokohama 223-8526
Japan

Phone: +81 45 564 2490
Email: kato@wide.ad.jp

Warren Kumari
Google
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: warren@kumari.net

