

dnsop
Internet-Draft
Updates: [4034](#), [4035](#), [5155](#), [8198](#) (if approved)
Intended status: Standards Track
Expires: 13 August 2021

P. van Dijk
PowerDNS
9 February 2021

NSEC and NSEC3 TTLs and NSEC Aggressive Use
draft-ietf-dnsop-nsec-ttl-03

Abstract

Due to a combination of unfortunate wording in earlier documents, aggressive use of NSEC and NSEC3 records may deny names far beyond the intended lifetime of a denial. This document changes the definition of the NSEC and NSEC3 TTL to correct that situation. This document updates [RFC 4034](#), [RFC 4035](#), [RFC 5155](#), and [RFC 8198](#).

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 13 August 2021.

Copyright Notice

Copyright (c) 2021 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 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.

Internet-Draft

nsec-ttl

February 2021

Table of Contents

1.	Introduction	2
2.	Conventions and Definitions	4
3.	NSEC and NSEC3 TTL changes	4
3.1.	Updates to RFC4034	4
3.2.	Updates to RFC4035	4
3.3.	Updates to RFC5155	5
3.4.	Updates to RFC8198	5
3.5.	A note on incremental signers	6
4.	Zone Operator Considerations	6
4.1.	A Note On Wildcards	6
5.	Security Considerations	6
6.	IANA Considerations	6
7.	Normative References	6
8.	Informative References	7
Appendix A.	Implementation Status	7
Appendix B.	Document history	8
	Acknowledgements	9
	Author's Address	9

[1.](#) Introduction

[RFC editor: please remove this block before publication.]

Earlier notes on this:

- * <https://indico.dns-oarc.net/event/29/sessions/98/#20181013>
(<https://indico.dns-oarc.net/event/29/sessions/98/#20181013>)
- * <https://lists.dns-oarc.net/pipermail/dns-operations/2018-April/thread.html#17420> (<https://lists.dns-oarc.net/pipermail/dns-operations/2018-April/thread.html#17420>)
- * <https://lists.dns-oarc.net/pipermail/dns-operations/2018-March/017416.html> (<https://lists.dns-oarc.net/pipermail/dns-operations/2018-March/017416.html>)

This document lives on GitHub (<https://github.com/PowerDNS/draft-dnsop-nsec-ttl>); proposed text and editorial changes are very much welcomed there, but any functional changes should always first be discussed on the IETF DNSOP WG mailing list.

]

[RFC2308] defines the TTL of the SOA record that must be returned in negative answers (NXDOMAIN or NODATA):

van Dijk

Expires 13 August 2021

[Page 2]

Internet-Draft

nsec-ttl

February 2021

| The TTL of this record is set from the minimum of the MINIMUM
| field of the SOA record and the TTL of the SOA itself, and
| indicates how long a resolver may cache the negative answer.

Thus, if the TTL of the SOA in the zone is lower than the SOA MINIMUM value (the last number in the SOA record), the authoritative server sends that lower value as the TTL of the returned SOA record. The resolver always uses the TTL of the returned SOA record when setting the negative TTL in its cache.

However, [\[RFC4034\] section 4](#) has this unfortunate text:

| The NSEC RR SHOULD have the same TTL value as the SOA minimum TTL
| field. This is in the spirit of negative caching ([\[RFC2308\]](#)).

This text, while referring to [RFC2308](#), can cause NSEC records to have much higher TTLs than the appropriate negative TTL for a zone. [\[RFC5155\]](#) contains equivalent text.

[RFC8198] [section 5.4](#) tries to correct this:

| [Section 5 of \[RFC2308\]](#) also states that a negative cache entry TTL
| is taken from the minimum of the SOA.MINIMUM field and SOA's TTL.
| This can be less than the TTL of an NSEC or NSEC3 record, since
| their TTL is equal to the SOA.MINIMUM field (see [\[RFC4035\]](#),
| [Section 2.3](#) and [\[RFC5155\], Section 3](#)).

| A resolver that supports aggressive use of NSEC and NSEC3 SHOULD
| reduce the TTL of NSEC and NSEC3 records to match the SOA.MINIMUM
| field in the authority section of a negative response, if
| SOA.MINIMUM is smaller.

But the NSEC and NSEC3 RRs should, according to [RFC4034](#) and [RFC5155](#), already be at the value of the MINIMUM field in the SOA. Thus, the advice from [RFC8198](#) would not actually change the TTL used for the NSEC and NSEC3 RRs for authoritative servers that follow the RFCs.

As a theoretical exercise, consider a TLD named ".example" with a SOA record like this:

```
"example. 900 IN SOA primary.example. hostmaster.example. 1 1800 900
604800 86400"
```

The SOA record has a 900 second TTL, and a 86400 MINIMUM TTL. Negative responses from this zone have a 900 second TTL, but the NSEC or NSEC3 records in those negative responses have a 86400 TTL. If a resolver were to use those NSEC or NSEC3 records aggressively, they would be considered valid for a day, instead of the intended 15 minutes.

[2.](#) Conventions and Definitions

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.

[3.](#) NSEC and NSEC3 TTL changes

[3.1.](#) Updates to [RFC4034](#)

Where [[RFC4034](#)] says:

| The NSEC RR SHOULD have the same TTL value as the SOA minimum TTL
| field. This is in the spirit of negative caching ([[RFC2308](#)]).

This is updated to say:

| The TTL of the NSEC RR that is returned MUST be the lesser of the
| MINIMUM field of the SOA record and the TTL of the SOA itself.
| This matches the definition of the TTL for negative responses in

| [[RFC2308](#)].

3.2. Updates to [RFC4035](#)

Where [[RFC4035](#)] says:

| The TTL value for any NSEC RR SHOULD be the same as the minimum
| TTL value field in the zone SOA RR.

This is updated to say:

| The TTL of the NSEC RR that is returned MUST be the lesser of the
| MINIMUM field of the SOA record and the TTL of the SOA itself.
| This matches the definition of the TTL for negative responses in
| [[RFC2308](#)].

3.3. Updates to [RFC5155](#)

Where [[RFC5155](#)] says:

| The NSEC3 RR SHOULD have the same TTL value as the SOA minimum TTL
| field. This is in the spirit of negative caching [[RFC2308](#)].

This is updated to say:

| The TTL of the NSEC3 RR that is returned MUST be the lesser of the
| MINIMUM field of the SOA record and the TTL of the SOA itself.
| This matches the definition of the TTL for negative responses in
| [[RFC2308](#)].

Where [[RFC5155](#)] says:

| o The TTL value for any NSEC3 RR SHOULD be the same as the minimum
| TTL value field in the zone SOA RR.

This is updated to say:

| o The TTL value for each NSEC3 RR MUST be the lesser of the

| MINIMUM field of the zone SOA RR and the TTL of the zone SOA RR
| itself.

3.4. Updates to [RFC8198](#)

[RFC8198] [section 5.4](#) (Consideration on TTL) is completely replaced by the following text:

| 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 \[RFC2308\]](#) suggests a maximum default negative cache
| TTL value of 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.

| A resolver that supports aggressive use of NSEC and NSEC3 MAY
| limit the TTL of NSEC and NSEC3 records to the lesser of the
| SOA.MINIMUM field and the TTL of the SOA in a response, if
| present. It MAY also use a previously cached SOA for a zone to
| find these values.

(The third paragraph of the original is removed, and the fourth paragraph is updated to allow resolvers to also take the lesser of the SOA TTL and SOA MINIMUM.)

3.5. A note on incremental signers

Some DNSSEC signer implementations might not (re-)sign whole zones in one go, instead spreading the work of updating inception/expiration times over some period. Such implementations would not be able to update all NSEC or NSEC3 records in the zone instantly either. To aid these implementations, we additionally specify the following:

| If an implementation cannot update all NSEC or NSEC3 TTLs after a
| SOA change immediately, it MUST still attempt to do so as soon as
| possible during the signing process.

4. Zone Operator Considerations

If signers & DNS servers for a zone cannot immediately be updated to

conform to this document, zone operators are encouraged to consider setting their SOA record TTL and the SOA MINIMUM field to the same value. That way, the TTL used for aggressive NSEC and NSEC3 use matches the SOA TTL for negative responses.

4.1. A Note On Wildcards

Validating resolvers consider an expanded wildcard valid for the wildcard's TTL, capped by the TTLs of the NSEC and NSEC3 proof that shows that the wildcard expansion is legal. Thus, changing the TTL of NSEC or NSEC3 records (explicitly, or by implementation of this document, implicitly) might affect (shorten) the lifetime of wildcards.

5. Security Considerations

An attacker can prevent future records from appearing in a cache by seeding the cache with queries that cause NSEC or NSEC3 responses to be cached, for aggressive use purposes. This document reduces the impact of that attack in cases where the NSEC or NSEC3 TTL is higher than the zone operator intended.

6. IANA Considerations

IANA is requested to add a reference to this document in the "Resource Record (RR) TYPEs" subregistry of the "Domain Name System (DNS) Parameters" registry, for the NSEC and NSEC3 types.

7. Normative References

- [RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", [RFC 2308](https://www.rfc-editor.org/info/rfc2308), DOI 10.17487/RFC2308, March 1998, <<https://www.rfc-editor.org/info/rfc2308>>.

- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", [RFC 4034](https://www.rfc-editor.org/info/rfc4034), DOI 10.17487/RFC4034, March 2005, <<https://www.rfc-editor.org/info/rfc4034>>.

- [RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", [RFC 5155](https://www.rfc-editor.org/info/rfc5155), DOI 10.17487/RFC5155, March 2008,

<<https://www.rfc-editor.org/info/rfc5155>>.

- [RFC8198] Fujiwara, K., Kato, A., and W. Kumari, "Aggressive Use of DNSSEC-Validated Cache", [RFC 8198](#), DOI 10.17487/RFC8198, July 2017, <<https://www.rfc-editor.org/info/rfc8198>>.
- [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, <<https://www.rfc-editor.org/info/rfc4035>>.
- [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>>.

8. Informative References

- [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>>.

Appendix A. Implementation Status

[RFC Editor: please remove this section before publication]

Implemented in PowerDNS Authoritative Server 4.3.0
<https://doc.powerdns.com/authoritative/dnssec/operational.html?highlight=ttl#some-notes-on-ttl-usage>
(<https://doc.powerdns.com/authoritative/dnssec/operational.html?highlight=ttl#some-notes-on-ttl-usage>) .

Implemented in BIND 9.16 and up, to be released early 2021
https://mailarchive.ietf.org/arch/msg/dnsop/ga41J2PPUbmC21--dqf3i7_IY6M (https://mailarchive.ietf.org/arch/msg/dnsop/ga41J2PPUbmC21--dqf3i7_IY6M) https://gitlab.isc.org/isc-projects/bind9/-/merge_requests/4506 (https://gitlab.isc.org/isc-projects/bind9/-/merge_requests/4506) .

https://gitlab.nic.cz/knot/knot-dns/-/merge_requests/1219
(https://gitlab.nic.cz/knot/knot-dns/-/merge_requests/1219) .

Implemented in ldns, patch under review
<https://github.com/NLnetLabs/ldns/pull/118>
(<https://github.com/NLnetLabs/ldns/pull/118>)

Implementation status is tracked at
<https://trac.ietf.org/trac/dnsop/wiki/draft-ietf-dnsop-nsec-ttl>
(<https://trac.ietf.org/trac/dnsop/wiki/draft-ietf-dnsop-nsec-ttl>)

[Appendix B](#). Document history

[RFC editor: please remove this section before publication.]

From [draft-vandijk-dnsop-nsec-ttl-00](#) to [draft-ietf-dnsop-nsec-ttl-00](#):

- * document was adopted
- * various minor editorial changes
- * now also updates 4035
- * use .example instead of .com for the example
- * more words on 8198
- * a note on wildcards

From [draft-ietf-dnsop-nsec-ttl-00](#) to [draft-ietf-dnsop-nsec-ttl-01](#):

- * various wording improvements
- * added Implementation note from Knot, expanded the BIND one with the GitLab MR URL
- * reduced requirement level from MUST to SHOULD, like the original texts

From [draft-ietf-dnsop-nsec-ttl-01](#) to [draft-ietf-dnsop-nsec-ttl-02](#):

- * updated the second bit of wrong text in 5155

From [draft-ietf-dnsop-nsec-ttl-02](#) to [draft-ietf-dnsop-nsec-ttl-03](#):

- * document now updates resolver behaviour in 8198

- * lots of extra text to clarify what behaviour goes where (thanks Paul Hoffman)
- * replace 'any' with 'each' (thanks Duane)
- * upgraded requirement level to MUST, plus a note on incremental signers

Acknowledgements

This document was made possible with the help of the following people:

- * Ralph Dolmans
- * Warren Kumari
- * Matthijs Mekking
- * Vladimir Cunat
- * Matt Nordhoff
- * Josh Soref
- * Tim Wicinski

The author would like to explicitly thank Paul Hoffman for extensive reviews, text contributions, and help in navigating WG comments.

Author's Address

Peter van Dijk
PowerDNS
Den Haag
Netherlands

Email: peter.van.dijk@powerdns.com

