Workgroup: dnsop
Internet-Draft: draft-ietf-dnsop-nsec-ttl-02
Updates: 4034, 4035, 5155 (if approved)
Published: 29 January 2021
Intended Status: Standards Track
Expires: 2 August 2021
Authors: P. van Dijk
PowerDNS

NSEC(3) TTLs and NSEC Aggressive Use

Abstract

Due to a combination of unfortunate wording in earlier documents, aggressive use of NSEC(3) records may deny names far beyond the intended lifetime of a denial. This document changes the definition of the NSEC(3) TTL to correct that situation. This document updates RFC 4034, RFC 4035, and RFC 5155.

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 2 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.

Table of Contents

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://lists.dns-oarc.net/pipermail/dns-operations/2018-April/thread.html#17420

* https://lists.dns-oarc.net/pipermail/dns-operations/2018-March/017416.html

This document lives on GitHub; 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 that the SOA TTL to be used in negative answers (NXDOMAIN or NODATA) is

> the minimum of the MINIMUM field of the SOA record and the TTL of the SOA itself

Thus, if the TTL of the SOA in the zone is lower than the SOA MINIMUM value (the last number in a SOA record), the negative TTL for that zone is lower than the SOA MINIMUM value.

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(3) RRs should, per RFC4034, already be at the MINIMUM TTL, which means this advice would never actually change the TTL used for the NSEC(3) RRs.

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(3) records in those negative responses have a 86400 TTL. If a resolver were to use those NSEC(3)s 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(3) 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 NSEC RR SHOULD have the same TTL value as 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 value for any NSEC RR SHOULD be the same TTL value as 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 NSEC3 RR SHOULD have the same TTL value as 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 any NSEC3 RR SHOULD be the same as the lesser of the MINIMUM field of the zone SOA RR and the TTL of the zone SOA RR itself.

3.4. No updates to RFC8198

Instead of updating three documents, it would have been preferable to update one. [RFC8198] says:

> With DNSSEC and aggressive use of DNSSEC-validated cache, the TTL of the NSEC/NSEC3 record and the SOA.MINIMUM field are the authoritative statement of how quickly a name can start working within a zone.

Here, the SOA.MINIMUM field cannot be changed to "the minimum/lesser of the SOA.MINIMUM field and the SOA TTL" because the resolver may not have the SOA RRset in cache. Because of that, this document cannot get away with updating just [RFC8198]. However, if authoritative servers follow the updates from this document, this should not make a difference, as the TTL of the NSEC/NSEC3 record is already set to the minimum value.

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 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(3) proof that shows that the wildcard expansion is legal. Thus, changing the TTL of NSEC(3) 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(3) responses to be cached, for aggressive use purposes. This document reduces the impact of that attack in cases where the NSEC(3) 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

[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>.

[RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", RFC 2308, 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, 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, DOI 10.17487/RFC5155, March 2008, <https://www.rfc-editor.org/info/rfc5155>.

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>.

[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>.

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 .

Implemented in BIND 9.16 and up, to be released early 2021 https://mailarchive.ietf.org/arch/msg/dnsop/ga41J2PPUbmc21--dqf3i7_IY6M https://gitlab.isc.org/isc-projects/bind9/-/merge_requests/4506 .

Implemented in Knot DNS 3.1, to be released in 2021 https://gitlab.nic.cz/knot/knot-dns/-/merge_requests/1219 .

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

Acknowledgements

Ralph Dolmans helpfully pointed out that fixing this in RFC8198 is only possible for negative (NXDOMAIN/NODATA) responses, and not for wildcard responses. Warren Kumari gracefully acknowledged that the current behaviour of RFC8198, in context of the NSEC TTL defined in RFC4034, is not the intended behaviour. Matthijs Mekking provided additional text explaining why this document cannot simply update RFC8198. Vladimir Cunat pointed out that the effect on wildcards should be made explicit. Paul Hoffman, Matt Nordhoff, and Josh Soref provided helpful corrections as native speakers.

Author's Address

Peter van Dijk
PowerDNS
Den Haag
Netherlands

Email: peter.van.dijk@powerdns.com