

Internet Engineering Task Force (IETF)
Request for Comments: 9077
Updates: 4034, 4035, 5155, 8198
Category: Standards Track
ISSN: 2070-1721

P. van Dijk
PowerDNS
July 2021

NSEC and NSEC3: TTLs and Aggressive Use

Abstract

Due to a combination of unfortunate wording in earlier documents, aggressive use of NSEC and NSEC3 records may deny the existence of 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 RFCs 4034, 4035, 5155, and 8198.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9077>.

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
 2. Conventions and Definitions
 3. NSEC and NSEC3 TTL Changes
 - 3.1. Updates to RFC 4034
 - 3.2. Updates to RFC 4035
 - 3.3. Updates to RFC 5155
 - 3.4. Updates to RFC 8198
 4. Zone Operator Considerations
 - 4.1. A Note on Wildcards
 5. Security Considerations
 6. IANA Considerations
 7. Normative References
- Acknowledgements
Author's Address

1. Introduction

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

```
| 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 top-level domain (TLD) named .example with an SOA record like this:

```
example.      900 IN      SOA primary.example. dnsadmin.example. (
                                     1 1800 900 604800 86400 )
```

The SOA record has a 900-second TTL and an 86400-second MINIMUM TTL. Negative responses from this zone have a 900-second TTL, but the NSEC or NSEC3 records in those negative responses have an 86400-second 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

[RFC4034], [RFC4035], and [RFC5155] use the SHOULD requirement level, but they were written prior to the publication of [RFC8198] when [RFC4035] still said:

| However, it seems prudent for resolvers to avoid blocking new
| authoritative data or synthesizing new data on their own.

[RFC8198] updated that text to contain:

| ...DNSSEC-enabled validating resolvers SHOULD use wildcards and
| NSEC/NSEC3 resource records to generate positive and negative
| responses until the effective TTLs or signatures for those records
| expire.

This means that the correctness of NSEC and NSEC3 records and their TTLs has become much more important. Because of that, the updates in this document upgrade the requirement level to MUST.

3.1. Updates to RFC 4034

[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]. Because some signers incrementally update the NSEC
| chain, a transient inconsistency between the observed and expected
| TTL MAY exist.

3.2. Updates to RFC 4035

[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]. Because some signers incrementally update the NSEC
| chain, a transient inconsistency between the observed and expected
| TTL MAY exist.

3.3. Updates to RFC 5155

[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]. Because some signers incrementally update the NSEC3
| chain, a transient inconsistency between the observed and expected
| TTL MAY exist.

Where [RFC5155] says:

```
| * 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:

```
| * 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. Because some signers incrementally update the NSEC3  
| chain, a transient inconsistency between the observed and  
| expected TTL MAY exist.
```

3.4. Updates to RFC 8198

[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.)

4. Zone Operator Considerations

If signers and 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.

Note that some signers might use the SOA TTL or MINIMUM as a default for other values, such as the TTL for DNSKEY records. Operators should consult documentation before changing values.

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 or 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 delay 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 has added 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

- [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>>.
- [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>>.
- [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>>.
- [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>>.

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 the 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