

Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 27, 2016

W. Kumari
Google
G. Huston
APNIC
February 24, 2016

Believing NSEC records in the DNS root.
draft-wkumari-dnsop-cheese-shop-01

Abstract

This document describes a method to generate negative answers from NSEC records for the special case of the DNS root. This improves performance; the resolver can answer immediately, and does not need to query the root. It also cuts down on the so-called "junk" queries.

[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-wkumari-dnsop-cheese-shop>. 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 August 27, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

If I've told you once...

February 2016

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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.	Background	2
2.	Believing NSEC records.	3
2.1.	Requirements notation	3
3.	Generating negatives responses from NSEC	3
4.	IANA Considerations	4
5.	Security Considerations	4
6.	Acknowledgements	4
7.	References	4
7.1.	Normative References	5
7.2.	Informative References	5
Appendix A.	Changes / Author Notes.	5
	Authors' Addresses	5

[1.](#) Background

[This section may be removed before publication... but I'd prefer not, it provides useful context]

If a DNS resolver queries a root zone authoritative name server with the EDNS0 DNSSEC OK option set, for a name that does not exist in the root zone, it gets back an NXDOMAIN response and an NSEC record, which "proves" that the name does not exist. NSEC proves this by providing names (and signatures) for the names which do exist on either side of the queried name. For example, if a nameserver queries for .belkin, it will get back an NXDOMAIN, and an NSEC record showing that nothing exists between (currently) .beer and .bentley [Ed note: There *probably* should be something between a beer and a bentley. :-P]. This means that, if the nameserver subsequently (during the TTL of the NSEC record) gets a query for .beeswax (alphabetically between beer and bentley) it need not attempt to resolve this - it has already been given proof that the name does not

exist.

The title of this draft comes from a famous Monty Python skit - "The Cheese Shop". There are some useful parallels between this problem

and the skit - watching the skit is encouraged to understand the problem - <https://www.youtube.com/watch?v=cWDdd5KKhts>

[2.](#) Believing NSEC records.

This is simply a refinement of [\[I-D.fujiwara-dnsop-nsec-aggressiveuse\]](#), for a limited use case (the root). Full credit to the authors of the aforementioned draft, and this draft does not replace that draft, nor remove the need for the broader consideration of the use of NSEC records as described in [\[I-D.fujiwara-dnsop-nsec-aggressiveuse\]](#).

The scope of this document is limited to the special case of recursive DNSSEC validating resolvers querying the root zone. This is because the root zone has some well known properties which make it a special case - we know it is DNSSEC signed, and uses NSEC, the majority of the queries are "junk" queries, the rate of change is relatively slow, and there are no odd corner cases such as wildcards. See [Section 3](#) for more discussion.

If the (DNSSEC validated) answer to a query to a root server is an NXDOMAIN then the resolver SHOULD cache the NSEC record provided in the response. The resolver SHOULD NOT send further queries for names within the range of the NSEC record for the lifetime of the cached NSEC TTL. Instead, the resolver SHOULD answer these queries directly with NXDOMAIN (and NSEC records if so signalled by EDNS). They SHOULD set the AA bit and AD bits.

[2.1.](#) Requirements notation

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 [\[RFC2119\]](#).

[3.](#) Generating negatives responses from NSEC

[This section is mainly for discussion, and is more informal. It should be deleted before publication.]

[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."

Kumari & Huston

Expires August 27, 2016

[Page 3]

Internet-Draft

If I've told you once...

February 2016

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)."

So, 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.

For the limited use case of this document (the DNS root) we believe that this is an acceptable trade off - the (current) TTL of the "negative cache" (in the SOA) is the same as the NSEC TTL (1 day). This means that, for a new TLD to begin resolving everywhere will require a minimum of a day - and this is true whether or not this is implemented (if someone had queried for the exact name, there would be a negatively cached answer, this simply expands the range of negative caches).

[4.](#) IANA Considerations

This document contains no IANA considerations.

[We MAY want to add something about setting the NSEC TTL appropriately?!]

[5.](#) Security Considerations

The impact of resolver caching is that the resolver will not re-query

an name server for a cached response until the TTL of the cached response expires. This may lead to cases where the resolver responds with outdated information for a period of time for subsequent queries for the name name.

This draft extends the scope of this vulnerability to include queries for all names that fall within the NSEC-defined range.

6. Acknowledgements

The authors wish to thank some folk, including Stephane Bortzmeyer, Bob Harold, Paul Vixie.

7. References

Kumari & Huston	Expires August 27, 2016	[Page 4]
-----------------	-------------------------	----------

Internet-Draft	If I've told you once...	February 2016
----------------	--------------------------	---------------

7.1. 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>>.

7.2. Informative References

[I-D.fujiwara-dnsop-nsec-aggressiveuse]
Fujiwara, K. and A. Kato, "Aggressive use of NSEC/NSEC3", [draft-fujiwara-dnsop-nsec-aggressiveuse-02](#) (work in progress), October 2015.

[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>>.

Appendix A. Changes / Author Notes.

[RFC Editor: Please remove this section before publication]

From -00 to -01.

- o Fairly significant rewrite - no substantive changes, only additional information, explanation and readability.

Authors' Addresses

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

Email: warren@kumari.net

Geoff Huston
APNIC
6 Cordelia St
South Brisbane QLD 4001
AUS

Email: gih@apnic.net