

Network Working Group
Internet-Draft
Updates: 1034, 1035 (if approved)
Intended status: Standards Track
Expires: February 15, 2019

J. Abley
Afilias
O. Gudmundsson
M. Majkowski
Cloudflare Inc.
E. Hunt
ISC
August 14, 2018

Providing Minimal-Sized Responses to DNS Queries that have QTYPE=ANY
draft-ietf-dnsop-refuse-any-07

Abstract

The Domain Name System (DNS) specifies a query type (QTYPE) "ANY". The operator of an authoritative DNS server might choose not to respond to such queries for reasons of local policy, motivated by security, performance or other reasons.

The DNS specification does not include specific guidance for the behaviour of DNS servers or clients in this situation. This document aims to provide such guidance.

This document updates RFC 1034 and RFC 1035.

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 February 15, 2019.

Copyright Notice

Copyright (c) 2018 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	3
1.1. Terminology	3
2. Motivations for Use of ANY Queries	3
3. General Approach	4
4. Behaviour of DNS Responders	4
4.1. Answer with a Subset of Available RRSets	5
4.2. Answer with a Synthesised HINFO RRSet	5
4.3. Answer with Best Guess as to Intention	6
4.4. Behaviour with TCP Transport	6
5. Behaviour of DNS Initiators	6
6. HINFO Considerations	7
7. Updates to RFC 1034 and RFC 1035	7
8. Implementation Experience	8
9. Security Considerations	8
10. IANA Considerations	8
11. Acknowledgements	8
12. References	9
12.1. Normative References	9
12.2. Informative References	9
12.3. URIs	9
Appendix A. Editorial Notes	10
A.1. Change History	10
A.1.1. draft-ietf-dnsop-refuse-any-07	10
A.1.2. draft-ietf-dnsop-refuse-any-06	10
A.1.3. draft-ietf-dnsop-refuse-any-05	10
A.1.4. draft-ietf-dnsop-refuse-any-04	10
A.1.5. draft-ietf-dnsop-refuse-any-03	10
A.1.6. draft-ietf-dnsop-refuse-any-02	10
A.1.7. draft-ietf-dnsop-refuse-any-01	11
A.1.8. draft-ietf-dnsop-refuse-any-00	11
A.1.9. draft-jabley-dnsop-refuse-any-01	11
A.1.10. draft-jabley-dnsop-refuse-any-00	11
Authors' Addresses	11

1. Introduction

The Domain Name System (DNS) specifies a query type (QTYPE) "ANY". The operator of an authoritative DNS server might choose not to respond to such queries for reasons of local policy, motivated by security, performance or other reasons.

The DNS specification [RFC1034] [RFC1035] does not include specific guidance for the behaviour of DNS servers or clients in this situation. This document aims to provide such guidance.

1.1. Terminology

This document uses terminology specific to the Domain Name System (DNS), descriptions of which can be found in [RFC7719].

In this document, "ANY Query" refers to a DNS meta-query with QTYPE=ANY. An "ANY Response" is a response to such a query.

In this document, "conventional ANY response" means an ANY response that is constructed in accordance with the algorithm documented in section 4.3.2 of [RFC1034] and specifically without implementing any of the mechanisms described in this document.

In an exchange of DNS messages between two hosts, this document refers to the host sending a DNS request as the initiator, and the host sending a DNS response as the responder.

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.

2. Motivations for Use of ANY Queries

ANY queries are legitimately used for debugging and checking the state of a DNS server for a particular name.

ANY queries are sometimes used as a attempt to reduce the number of queries needed to get information, e.g. to obtain MX, A and AAAA RRSets for a mail domain in a single query. There is no documented guidance available for this use case, however, and some implementations have been observed not to function as perhaps their developers expected. Implementers that assume that an ANY query will ultimately be received by an authoritative server and will fetch all existing RRSets, should include a fallback mechanism to use when that does not happen.

ANY queries are frequently used to exploit the amplification potential of DNS servers/resolvers using spoofed source addresses and UDP transport (see [RFC5358]). Having the ability to return small responses to such queries makes DNS servers less attractive amplifiers.

ANY queries are sometimes used to help mine authoritative-only DNS servers for zone data, since they are expected to return all RRSets for a particular query name. If a DNS operator prefers to reduce the potential for information leaks, they might choose not to send large ANY responses.

Some authoritative-only DNS server implementations require additional processing in order to send a conventional ANY response, and avoiding that processing expense might be desirable.

3. General Approach

This proposal provides a mechanism for an authority server to signal that conventional ANY queries are not supported for a particular QNAME, and to do so in such a way that is both compatible with and triggers desirable behaviour by unmodified clients (e.g. DNS resolvers).

Alternative proposals for dealing with ANY queries have been discussed. One approach proposed using a new RCODE to signal that an authoritative server did not answer ANY queries in the standard way. This approach was found to have an undesirable effect on both resolvers and authoritative-only servers; resolvers receiving an unknown RCODE would re-send the same query to all available authoritative servers, rather than suppress future such ANY queries for the same QNAME.

This proposal avoids that outcome by returning a non-empty RRSset in the ANY response, providing resolvers with something to cache and effectively suppressing repeat queries to the same or different authority servers.

4. Behaviour of DNS Responders

Below are the three different modes of behaviour by DNS responders when processing queries with QNAMEs that exist, QCLASS=IN and QTYPE=ANY. Operators/Implementers are free to choose whichever mechanism best suits their environment.

1. A DNS responder can choose to select one or a larger subset of the available RRSets at the QNAME.

2. A DNS responder can return a synthesised HINFO resource record. See Section 6 for discussion of the use of HINFO.
3. Resolver can try to give out the most likely records the requester wants. This is not always possible and the result might well be a large response.

Except as described below in this section, the DNS responder MUST follow the standard algorithms when constructing a response.

4.1. Answer with a Subset of Available RRSets

A DNS responder which receives an ANY query MAY decline to provide a conventional ANY response, or MAY instead send a response with a single RRSset (or a larger subset of available RRSets) in the answer section.

The RRSets returned in the answer section of the response MAY consist of a single RRSset owned by the name specified in the QNAME. Where multiple RRSets exist, the responder SHOULD choose a small subset of those available to reduce the amplification potential of the response.

If the zone is signed, appropriate RRSIG records MUST be included in the answer.

Note that this mechanism does not provide any signalling to indicate to a client that an incomplete subset of the available RRSets has been returned.

4.2. Answer with a Synthesised HINFO RRSset

If there is no CNAME present at the owner name matching the QNAME, the resource record returned in the response MAY instead be synthesised, in which case a single HINFO resource record SHOULD be returned. The CPU field of the HINFO RDATA SHOULD be set to RFCXXXX [note to RFC Editor, replace with RFC number assigned to this document]. The OS field of the HINFO RDATA SHOULD be set to the null string to minimize the size of the response.

The TTL encoded for the synthesised HINFO RR SHOULD be chosen by the operator of the DNS responder to be large enough to suppress frequent subsequent ANY queries from the same initiator with the same QNAME, understanding that a TTL that is too long might make policy changes relating to ANY queries difficult to change in the future. The specific value used is hence a familiar balance when choosing TTL for any RR in any zone, and be specified according to local policy.

If the DNS query includes DO=1 and the QNAME corresponds to a zone that is known by the responder to be signed, a valid RRSIG for the RRSets in the answer (or authority if answer is empty) section MUST be returned. In the case of DO=0, the RRSIG SHOULD be omitted.

A system that receives an HINFO response SHOULD NOT infer that the response was generated according to this specification and apply any special processing of the response, since in general it is not possible to tell with certainty whether the HINFO RRSet received was synthesised. In particular, systems SHOULD NOT rely upon the HINFO RDATA described in this section to distinguish between synthesised and non-synthesised HINFO RRSets.

4.3. Answer with Best Guess as to Intention

In some cases it is possible to guess what the initiator wants in the answer (but not always). Some implementations have implemented the spirit of this document by returning all RRSets of RRTYPE CNAME, MX, A and AAAA that are present at the owner name but suppressing others. This heuristic seems to work well in practice, satisfying the needs of some applications whilst suppressing other RRSets such as TXT and DNSKEY that can often contribute to large responses. Whilst some applications may be satisfied by this behaviour, the resulting responses in the general case are larger than the approaches described in Section 4.1 and Section 4.2.

As before, if the zone is signed and the DO bit is set on the corresponding query, an RRSIG RRSet MUST be included in the response.

4.4. Behaviour with TCP Transport

A DNS responder MAY behave differently when processing ANY queries received over different transport, e.g. by providing a conventional ANY response over TCP whilst using one of the other mechanisms specified in this document in the case where a query was received using UDP.

Implementers SHOULD provide configuration options to allow operators to specify different behaviour over UDP and TCP.

5. Behaviour of DNS Initiators

A DNS initiator which sends a query with QTYPE=ANY and receives a response containing an HINFO resource record or a single RRset, as described in Section 4, MAY cache the response in the normal way. Such cached resource records SHOULD be retained in the cache following normal caching semantics, as it would with any other response received from a DNS responder.

A DNS initiator MAY suppress queries with QTYPE=ANY in the event that the local cache contains a matching HINFO resource record with RDATA.CPU field, as described in Section 4. A DNS initiator MAY instead respond to such queries with the contents of the local cache in the usual way.

6. HINFO Considerations

It is possible that the synthesised HINFO RRSset in an ANY response, once cached by the initiator, might suppress subsequent queries from the same initiator with QTYPE=HINFO. Thus the use of HINFO in this proposal would hence have effectively mask the HINFO RRSset present in the zone.

Authority-server operators who serve zones that rely upon conventional use of the HINFO RRTYPE SHOULD sensibly choose the "single RRSset" method described in this document or select another type.

The HINFO RRTYPE is believed to be rarely used in the DNS at the time of writing, based on observations made at recursive servers, authority servers and in passive DNS.

7. Updates to RFC 1034 and RFC 1035

This document extends the specification for processing ANY queries described in section 4.3.2 of [RFC1034].

It is important to note that returning a subset of available RRSets when processing an ANY query is legitimate and consistent with [RFC1035]; it can be argued that ANY does not always mean ALL, as used in section 3.2.3 of [RFC1035]. The main difference here is that the TC bit SHOULD NOT be set on the response indicating that this is not a complete answer.

This document describes optional behaviour for both DNS initiators and responders, and implementation of the guidance provided by this document is OPTIONAL.

RRSIG queries (i.e. queries with QTYPE=RRSIG) are similar to ANY queries in the sense that they have the potential to generate large responses as well as extra work for the responders that process them, e.g. in the case where signatures are generated on-the-fly. RRSIG RRSets are not usually obtained using such explicit queries, but are rather included in the responses for other RRSets that the RRSIGs cover. This document does not specify appropriate behaviour for RRSIG queries, but note that future such advice might well benefit

from consistency with and experience of the approaches for ANY queries described here.

8. Implementation Experience

In October 2015 Cloudflare Authoritative Name server implementation implemented the HINFO response. A few minor problems were reported and have since been resolved.

An implementation of the subset-mode response to ANY queries was implemented in NSD 4.1 in 2016.

An implementation of a single RRSet response to an ANY query was made for BIND9 by Tony Finch, and that functionality was subsequently made available in production releases starting in BIND 9.11.

9. Security Considerations

Queries with QTYPE=ANY are frequently observed as part of reflection attacks, since a relatively small query can be used to elicit a large response; this is a desirable characteristic if the goal is to maximize the amplification potential of a DNS server as part of a volumetric attack. The ability of a DNS operator to suppress such responses on a particular server makes that server a less useful amplifier.

The optional behaviour described in this document to reduce the size of responses to queries with QTYPE=ANY is compatible with the use of DNSSEC by both initiator and responder.

10. IANA Considerations

The IANA is requested to update the Resource Record (RR) TYPEs Registry [1] entry as follows:

Type	Value	Meaning	Reference
*	255	A request for some or all records the server has available	[RFC1035][RFC6895] [This Document]

11. Acknowledgements

David Lawrence provided valuable observations and concrete suggestions. Jeremy Laidman helped make the document better. Tony Finch realized that this document was valuable and implemented it

while under attack. Richard Gibson identified areas where more detail and accuracy was useful. A large number of other people also provided comments and suggestions we thank them all for the feedback.

12. References

12.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [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>>.
- [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>>.

12.2. Informative References

- [RFC5358] Damas, J. and F. Neves, "Preventing Use of Recursive Nameservers in Reflector Attacks", BCP 140, RFC 5358, DOI 10.17487/RFC5358, October 2008, <<https://www.rfc-editor.org/info/rfc5358>>.
- [RFC6895] Eastlake 3rd, D., "Domain Name System (DNS) IANA Considerations", BCP 42, RFC 6895, DOI 10.17487/RFC6895, April 2013, <<https://www.rfc-editor.org/info/rfc6895>>.
- [RFC7719] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", RFC 7719, DOI 10.17487/RFC7719, December 2015, <<https://www.rfc-editor.org/info/rfc7719>>.

12.3. URIs

- [1] <http://www.iana.org/assignments/dns-parameters/dns-parameters.xhtml#dns-parameters-4>

Appendix A. Editorial Notes

This section (and sub-sections) to be removed prior to publication.

A.1. Change History

A.1.1. draft-ietf-dnsop-refuse-any-07

Address AD's concerns: more colour to describe updates to 1034/1035 in the abstract; don't rely upon HINFO RDATA formatting; language cleanup around guess intent. Add Evan as author (originator of the "choose one record" response idea).

A.1.2. draft-ietf-dnsop-refuse-any-06

Update RFC 1034 as well as RFC 1035; define the term "conventional ANY response"; soften and qualify ANY does not mean ALL; note that the subset mode response lacks signalling.

A.1.3. draft-ietf-dnsop-refuse-any-05

Minor editorial changes. Soften advice on RRSIG queries. Version bump.

A.1.4. draft-ietf-dnsop-refuse-any-04

These are the changes requested during WGLC. The title has been updated for readability The behavior section now contains description of three different approaches in order of preference. Text added on behavior over TCP. The document is clear in how it updates from RFC1035. Minor adjustments for readability and remove redundancy.

A.1.5. draft-ietf-dnsop-refuse-any-03

Change section name to "Updates to RFC1034", few minor grammar changes suggested by Matthew Pounsett and Tony Finch.

Text clarifications, reflecting experience, added implementation experience.

A.1.6. draft-ietf-dnsop-refuse-any-02

Added suggestion to call out RRSIG is optional when DO=0.

Number of text suggestions from Jeremy Laidman.

A.1.7. draft-ietf-dnsop-refuse-any-01

Add IANA Considerations

A.1.8. draft-ietf-dnsop-refuse-any-00

Re-submitted with a different name following adoption at the dnsop WG meeting convened at IETF 94.

A.1.9. draft-jabley-dnsop-refuse-any-01

Make signing of RRSets in answers from signed zones mandatory.

Document the option of returning an existing RRSset in place of a synthesised one.

A.1.10. draft-jabley-dnsop-refuse-any-00

Initial draft circulated for comment.

Authors' Addresses

Joe Abley
Afilias
300-184 York Street
London, ON N6A 1B5
Canada

Phone: +1 519 670 9327
Email: jabley@afilias.info

Olafur Gudmundsson
Cloudflare Inc.

Email: olafur+ietf@cloudflare.com

Marek Majkowski
Cloudflare Inc.

Email: marek@cloudflare.com

Evan Hunt
ISC
950 Charter St
Redwood City, CA 94063
USA

Email: each@isc.org