Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 28, 2012

R. Bellis
Nominet UK

March 27, 2012

Title
draft-bellis-dnsext-multi-qtypes-00

Abstract

This document specifies a method for a DNS client to request additional DNS record types to be delivered alongside the primary record type specified in the question section of a DNS query.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

A commonly requested DNS [RFC1035] feature is the ability to receive multiple related resource records (RRs) in a single DNS response.

For example, it may be desirable to receive both the A and AAAA records for a domain name together, rather than having to issue multiple queries.

The DNS wire protocol in theory supports having multiple questions in a single packet, but in practise this does not work:

- Each question consists of the tuple (QNAME, QTYPE, QCLASS). Since each question has its own QNAME field it would be possible for one name to exist and another to not exist, resulting in an inconsistent response code.
- The idea that only a single question is allowed is sufficiently entrenched that many DNS servers will simply return an error (or fail to response at all) if they receive a query with a question count (QDCOUNT) of more than one.

To resolve both of these issues, this document constraints the problem to those cases where only the QTYPE varies by specifying a new option for the Extension Mechanisms for DNS (EDNS [RFC2671]) that contains an additional list of QTYPE values that the client wishes to receive in addition to that in the primary question.

TODO: why not "ANY" ?

2. Terminology used in this document

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. Description

3.1. Multiple QTYPE EDNS Option Format

The overall format of an EDNS option is shown for reference below, per [RFC2671], followed by the option specific data:

```
       +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
    0: |                          OPTION-CODE                          |
       +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
    2: |                         OPTION-LENGTH                         |
       +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
    4: |                                                               |
       /                          OPTION-DATA                          /
       /                                                               /
       +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

OPTION-CODE: TBD by IANA

OPTION-LENGTH: Size (in octets) of OPTION-DATA.

OPTION-DATA: Option specific, as below:

```
                    +0 (MSB)                            +1 (LSB)
       +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
    0: |QTD|   reserved    |  QTCOUNT  |           QT1 (MSB)           |
       +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
    2: |           QT1 (LSB)           |              ...              |
       +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
       |               ...            ///          QTn (MSB)           |
       +---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
       |           QTn (LSB)           |
       +---+---+---+---+---+---+---+---+
```

QTD: this bit indicates the direction of the packet. It MUST be clear (0) in a query and set (1) in a response.

QTCOUNT: a 3 bit field with range 0 .. 7 specifying the number of QT fields to follow.

QTn: a 2 byte field (MSB first) specifying a DNS RR type. The RR type MUST be for a real resource record, and MUST NOT refer to a pseudo RR type such as "OPT", "IXFR", "TSIG", etc.

3.2. Response Generation

3.2.1. Server Side Processing

A conforming server that receives a Multiple QTYPE Option in a query MUST return a Multiple QTYPE Option in its response.

The QTD bit in that response MUST be set (1) as protection against

servers which simply echo unknown EDNS options verbatim. If the QTD bit in a response is zero the client MUST treat the response as if this option is unsupported.

The server SHOULD attempt to return any resource records known to it that match the additional (QTYPE, QCLASS, QTn) tuples. These records MUST be returned in the Answer Section of the response, but the answer for the primary QTYPE from the Question Section MUST be included first.

For any particular QTn in the query, if the server provides addtional answers, or has knowledge that the RR type type does not exist for that QNAME (a "negative answer"), it must include that QTn value in the Multiple QTYPE Option of its response.

A negative answer is therefore indicated by the combination of the presence of a QTn value in the Multiple QTYPE Option and the absence of a matching record in the Answer Section. This is necessary (in the absence of DNSSEC) to differentiate between absence of the record from the zone and absence of the record from the response.

A server that is authoritative for the specified QNAME on receipt of a Multiple QTYPE Option MUST attempt to return all specified RR types except where that would result in truncation in which case it may omit some (or all) of the records for the additional RR types. Those RR types MUST then also be omitted from the Multiple QTYPE Option in the response.

A caching recursive server receiving a Multiple QTYPE Option SHOULD attempt to fill its positive and negative caches with all of the specified RR types before returning its response to the client.

TODO: is there a case for mandatory answers, i.e. the client saying I _really_ want all these?

3.2.2. Client Side Processing

Recursive resolvers MAY use this method to obtain multiple records from an authoritative server. For the purposes of Section 5.4.1 of [RFC2181] any authoritative answers received MUST be ranked the same as the answer for the primary question.

3.2.3. DNSSEC

If the DNS client sets the "DNSSEC OK" (DO) bit in the query then the server MUST also return the related DNSSEC records that would have been returned in a standalone query for the same QTYPE.

A negative answer from a signed zone MUST contain the appropriate authenticated denial of existence records, per [RFC3403] and [RFC5155].

In a signed zone there is a theoretical risk of valid signatures for one RR type and invalid signatures for another. This is the only case known to the author where the response code for any particular QNAME may be inconsistent across different RR types.

Should a validating resolver produce NOERROR for some RR types and SERVFAIL for others it MUST omit the RR types that failed to validate from its response and from the QTn fields on the Multiple QTYPE option. The client MAY then initiate standalone queries for those RR types.

4. Security Considerations

The method documented here does not change any of the security properties of the DNS protocol itself.

It should however be noted that this method does increase the potential amplification factor when the DNS protocol is used as a vector for a denial of service attack.

5. IANA Considerations

IANA is requested to assign a new value in the DNS EDNS0 Options registry.

6. Acknowledgements

The author wishes to thank the following for their feedback and reviews during the initial development of this document: Michael Graff, Olafur Gudmundsson, Matthijs Mekking, Paul Vixie.

7. Normative References

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS

Specification", RFC 2181, July 1997.

[RFC2671] Vixie, P., "Extension Mechanisms for DNS (EDNS0)", RFC 2671, August 1999.

[RFC3403] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database", RFC 3403, October 2002.

[RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", RFC 5155, March 2008.

Appendix A. Change Log

NB: to be removed by the RFC Editor before publication.

draft-bellis-dnsext-multi-qtypes-00
Initial draft

Author's Address

Ray Bellis
Nominet UK
Edmund Halley Road
Oxford OX4 4DQ
United Kingdom

Phone: +44 1865 332211
Email: ray.bellis@nominet.org.uk
URI: http://www.nominet.org.uk/