

Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 14, 2018

K. Fujiwara
JPRS
January 10, 2018

Returning additional answers in DNS responses
draft-fujiwara-dnsop-additional-answers-01

Abstract

This document proposes to document the ability to provide multiple answers in single DNS response. For example, authoritative servers may add a NSEC resource record or A/AAAA resource records of the query name. This is especially useful as, in many cases, the entity making the request has no a priori knowledge of what other questions it will need to ask. It is already possible (an authoritative server MAY already send what it wants in the additional section). This document does not propose any protocol changes, just explanations of an already acceptable practice.

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 July 14, 2018.

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 (<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. Introduction	2
2. Background	3
3. Terminology	3
4. Returning multiple answers	4
5. Possible additional answers	4
6. Stub-Resolver Considerations	5
7. Use of Additional information	5
8. IANA Considerations	5
9. Security Considerations	5
10. Acknowledgments	6
11. Change History	6
11.1. 00 to 01	6
12. References	6
12.1. Normative References	6
12.2. Informative References	7
Appendix A. Comparisons of multiple response proposals	7
A.1. draft-wkumari-dnsop-multiple-responses	7
A.2. draft-fujiwara-dnsop-additional-answers	8
A.3. draft-bellis-dnsexp-multi-qtypes	8
A.4. draft-yao-dnsop-accompanying-questions	8
A.5. draft-vavrusa-dnsop-aaaa-for-free	8
A.6. QDCOUNT>1 idea	8
A.7. Comparison chart	9
Author's Address	9

1. Introduction

[I-D.wkumari-dnsop-multiple-responses] proposes pseudo resource record that controls resource records added into additional section. It offers any combinations of owner names and record types that are added into additional section.

In many cases, combinations are limited and DNS software developers knows well. This document proposes that DNS server software developers choose the combination of additional data.

By providing multiple answers in single response, authoritative name servers can assist full-service resolvers in pre-populating their cache before stub resolvers or other clients ask for the subsequent queries. Apart from decreasing the latency for end users [RFC6555],

this also decreases the total number of queries that full-service resolvers need to send and authoritative servers need to answer.

By providing NSEC/NSEC3 resource record that matches a query name, validating resolvers can generate NODATA or NXDOMAIN responses with Aggressive Use of DNSSEC-validated cache [RFC8198].

Developers of DNS servers know end users' query patterns or full-service resolvers' query patterns well. Authoritative DNS servers may add any authoritative data in the additional section. For example, QTYPE MX queries are followed by mail exchange hosts A/AAAA queries. When an authoritative server receives a QTYPE MX query, some implementations add mail exchange hosts A/AAAA resource records in additional section if the authoritative server have authoritative data of mail exchange hosts.

Other typical examples are A and AAAA, SRV and Target A/AAAA, TLSA RR and corresponding server addresses.

This technique, described in this document, is purely an optimization and enables authoritative servers to distribute some other related answers that the client is likely to need along with an answer to the original request. Users get a better experience, full-service resolvers need to send less queries, authoritative servers have to answer fewer queries, etc.

2. Background

The DNS specifications ([RFC1034], for instance section 4.3.2) allow for supplemental information to be included in the "additional" section of the DNS response, but in order to defeat cache poisoning attacks most implementations either ignore or don't trust additional records they didn't ask for. For more background, see [RFC2181].

Some implementations add mail exchange A/AAAA resource records in MX responses (an actual example is given in section 3.7.1 of [RFC1034]). Some implementations add Target A/AAAA resource records in SRV responses.

3. Terminology

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 RFC 2119 [RFC2119].

Many of the specialized terms used in this specification are defined in DNS Terminology [RFC7719] and [I-D.ietf-dnsop-terminology-bis].

Additional records: Additional records are records that the authoritative nameserver has included in the Additional section.

4. Returning multiple answers

An authoritative nameserver MAY include any additional records that help name resolution. These additional records are appended to the additional section of the response.

To increase the probability that these extra data will actually be useful for the resolver, it is suggested to send them only if:

- o The query has DNSSEC OK bit set.
- o The authoritative server is authoritative for the additional records, and the records to be returned are DNSSEC signed. The additional records contain RRSIGs.
- o To prove the non-existence of the resource record type, additional records may be NSEC/NSEC3 resource records for the query name and some other query names (for example, TLSA owner name). Validating resolvers can generate negative NODATA/NXDOMAIN response with Aggressive Use of DNSSEC-validated cache [RFC8198].
- o Responses with additional records fit in the required response size.

Additional records may be controlled by server configuration. "enable additional a/aaaa" or "enable additional nsec*" options are possible.

5. Possible additional answers

Possible query and additional records pairs are:

- o NAME A : NAME AAAA (or NAME NSEC/NSEC3)
- o NAME AAAA : NAME A (or NAME NSEC/NSEC3)
- o NAME MX : mail exchange A/AAAA (and/or mail exchange NSEC/NSEC3)
- o NAME SRV : Target host A/AAAA (and/or Target host NSEC/NSEC3)
- o NAME A/AAAA : _443._tcp.NAME TLSA (and/or NAME NSEC/NSEC3)
- o _443._tcp.NAME TLSA : NAME A/AAAA (and/or NAME NSEC/NSEC3)

TLSA / MX / SRV pairs have different query names.

6. Stub-Resolver Considerations

No modifications need to be made to stub-resolvers to get the predominate benefit of this protocol, since the majority of the speed gain will take place between the validating recursive resolver and the authoritative name server. However, stub resolvers and full-service resolvers may use this technique if stub-resolvers are validating stub resolvers.

7. Use of Additional information

When deciding to use additional records in the additional section, a resolver should follow certain rules:

- o Additional records are validated before being used.
- o Additional records SHOULD have lower priority in the cache than answers received because they were requested. This is to help evict Additional records from the cache first (to help prevent cache filling attacks).
- o Recursive resolvers MAY choose to ignore Additional records for any reason, including CPU or cache space concerns, phase of the moon, etc. It may choose to accept all, some or none of the Additional record sets.
- o Recursive resolvers SHOULD support "Aggressive use of DNSSEC-validated cache" [RFC8198].

These rules are derived from [RFC2181] and DNSSEC RFCs.

8. IANA Considerations

This document has no IANA actions.

9. Security Considerations

The use of DNSSEC guarantees that these additional records will be accepted and cached by the resolver only if they can be proved genuine.

The technique described in this document makes DNS response size large. If DNS response size exceeds path MTU, the response will be fragmented and the fragmentation may cause problems. Authoritative DNS server software developers and operators need to choose suitable response size limit.

10. Acknowledgments

The author acknowledges authors of [I-D.wkumari-dnsop-multiple-responses] because many part of idea and texts are copied from the draft.

The author would like to specifically thank Stephane Bortzmeyer for extensive review and comments.

11. Change History

11.1. 00 to 01

Sync with IETF 100 presentation

- o Added system wide configuration that controls additional records
- o Added "draft-vavrusa-dnsop-aaaa-for-free"
- o Updated comparison table

12. References

12.1. Normative References

- [I-D.ietf-dnsop-terminology-bis]
Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", draft-ietf-dnsop-terminology-bis-08 (work in progress), November 2017.
- [I-D.wkumari-dnsop-multiple-responses]
Kumari, W., Yan, Z., Hardaker, W., and D. Lawrence, "Returning extra answers in DNS responses.", draft-wkumari-dnsop-multiple-responses-05 (work in progress), July 2017.
- [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>>.

- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, DOI 10.17487/RFC2181, July 1997, <<https://www.rfc-editor.org/info/rfc2181>>.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, DOI 10.17487/RFC6555, April 2012, <<https://www.rfc-editor.org/info/rfc6555>>.
- [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>>.
- [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>>.

12.2. Informative References

- [I-D.bellis-dnsexst-multi-qtypes] Bellis, R., "DNS Multiple QTYPEs", draft-bellis-dnsexst-multi-qtypes-05 (work in progress), January 2018.
- [I-D.vavrusa-dnsop-aaaa-for-free] marek@vavrusa.com, m. and O. Gu[eth]mundsson, "Providing AAAA records for free with QTYPE=A", draft-vavrusa-dnsop-aaaa-for-free-00 (work in progress), March 2016.
- [I-D.yao-dnsop-accompanying-questions] Yao, J., Vixie, P., Kong, N., and X. Lee, "A DNS Query including A Main Question with Accompanying Questions", draft-yao-dnsop-accompanying-questions-04 (work in progress), September 2017.

Appendix A. Comparisons of multiple response proposals

A.1. draft-wkumari-dnsop-multiple-responses

[I-D.wkumari-dnsop-multiple-responses] proposes pseudo resource record that controls resource records added into additional section.

No protocol changes between authoritative servers and full-service resolvers. New authoritative server software required. Zone operators need to configure. Supports different owner names and types. Answer size becomes large if the query matches operators configuration. Requires DNSSEC.

A.2. draft-fujiwara-dnsop-additional-answers

draft-fujiwara-dnsop-additional-answers proposes that authoritative servers add well used additional records and NSEC/NSEC3 resource records in additional section.

No protocol changes between authoritative servers and full-service resolvers. New authoritative server software required. No configuration. Supports different owner names and types. Answer size becomes large (always). Requires DNSSEC and [RFC8198].

A.3. draft-bellis-dnsexp-multi-qtypes

[I-D.bellis-dnsexp-multi-qtypes] proposes new EDNS options that carry additional query types.

New authoritative server software required. New full-service resolver software required. No configuration. No support of different owner names.

A.4. draft-yao-dnsop-accompanying-questions

[I-D.yao-dnsop-accompanying-questions] proposes new EDNS option that carry additional query names, query types and rcodes.

New authoritative server software required. New full-service resolver software required. No configuration.

A.5. draft-vavrusa-dnsop-aaaa-for-free

[I-D.vavrusa-dnsop-aaaa-for-free] proposes additional AAAA resource records in answer section. New authoritative server software required. New full-service resolver software required because existing full-service resolvers ignore additional AAAA resource records. No configuration.

A.6. QDCOUNT>1 idea

No drafts. QDCOUNT is not limited to 1 in [RFC1035].

No protocol changes between authoritative servers and full-service resolvers, however, some implementations (For example, BIND 9, NSD, Unbound) treats QDCOUNT>1 as FORMERR. New authoritative server software required. New full-service resolver software required. Supports different owner names and types, however, it cannot answer different rcodes. No configuration. A database that each IP address support QDCOUNT>1 is required in full-service resolvers.

A.7. Comparison chart

Draft	additional answers	multiple responses	aaaa for free	multi qtypes	accompanying querstions
Protocol change	No	No	Yes?	Yes	Yes
Code size	little	some	little	large?	large?
Resolver modification	No	No	Yes?	Yes	Yes
Config complexity	No	Yes	No	No	No
Multiple names	Yes	Yes	No	No	Yes
Multiple types	Yes	Yes	AAAA	Yes	Yes
Multiple rcodes	(NSEC*)	---	---	---	Yes
Negative response	Yes	No	No	Yes	Yes
Fat response if	always	config	always	query	query
Stub support ?	No	No	?	possible	possible
Deployment	easy	easy	gradual	gradual	gradual
Require DNSSEC	(Yes)	(Yes)	No	No	No
IP addr Database	No	No	No	EDNS	EDNS

Author's Address

Kazunori Fujiwara
 Japan Registry Services Co., Ltd.
 Chiyoda First Bldg. East 13F, 3-8-1 Nishi-Kanda
 Chiyoda-ku, Tokyo 101-0065
 Japan

Phone: +81 3 5215 8451
 Email: fujiwara@jprs.co.jp