Network Working Group
Internet-Draft
Updates: 1034 (if approved)
Intended status: Standards Track
Expires: December 5, 2020

M. Andrews
ISC

June 3, 2020

Glue In DNS Referral Responses Is Not Optional
draft-ietf-dnsop-glue-is-not-optional-00

Abstract

The DNS uses glue records to allow iterative clients to find the addresses of nameservers that live within the delegated zone. Glue records are expected to be returned as part of a referral and if they cannot be fitted into the UDP response, TC=1 MUST be set to inform the client that the response is incomplete and that TCP SHOULD be used to retrieve the full response.

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 December 5, 2020.

Table of Contents

1. Introduction

The DNS [RFC1034], [RFC1035] uses glue records to allow iterative clients to find the addresses of nameservers that live within the delegated zone. Glue records are expected to be returned as part of a referral and if they cannot be fitted into the UDP response, TC=1 MUST be set to inform the client that the response is incomplete and that TCP SHOULD be used to retrieve the full response.

While not common, real life examples of servers that fail to set TC=1 when glue records are available exist and they do cause resolution failures. The example below shows a case where none of the glue records, present in the zone, fitted into the available space and TC=1 was not set in the response. While this example shows an DNSSEC [RFC4033], [RFC4034], [RFC4035] referral response, this behaviour has also been seen with plain DNS responses as well. The records have been truncated for display purposes.

```
% dig +norec +dnssec +bufsize=512 +ignore @a.gov-servers.net \
        rh202ns2.355.dhhs.gov

; <<>> DiG 9.15.4 <<>> +norec +dnssec +bufsize +ignore \
        @a.gov-servers.net rh202ns2.355.dhhs.gov
; (2 servers found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 8798
;; flags: qr; QUERY: 1, ANSWER: 0, AUTHORITY: 9, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags: do; udp: 4096
;; QUESTION SECTION:
;rh202ns2.355.dhhs.gov.         IN A

;; AUTHORITY SECTION:
dhhs.gov.               86400   IN NS      rh120ns2.368.dhhs.gov.
dhhs.gov.               86400   IN NS      rh202ns2.355.dhhs.gov.
dhhs.gov.               86400   IN NS      rh120ns1.368.dhhs.gov.
dhhs.gov.               86400   IN NS      rh202ns1.355.dhhs.gov.
dhhs.gov.               3600    IN DS      51937 8 1 ...
dhhs.gov.               3600    IN DS      635 8 2 ...
dhhs.gov.               3600    IN DS      51937 8 2 ...
dhhs.gov.               3600    IN DS      635 8 1 ...
dhhs.gov.               3600    IN RRSIG   DS 8 2 3600 ...

;; Query time: 226 msec
;; SERVER: 69.36.157.30#53(69.36.157.30)
;; WHEN: Wed Apr 15 13:34:43 AEST 2020
;; MSG SIZE  rcvd: 500

%
```

This is almost certainly due a wide spread misbelief that all additional section records are optional. This has never been the

case with respect to glue records and later protocol extension have added more cases where records in the additional section are not optional in the response. This includes TSIG [RFC2845], OPT [RFC6891], and SIG(0) [RFC2931].

Glue records are added to the parent zone as part of the delegation process. They are expected to be returned as part of a referral and if they can't fit in a UDP response TC=1 MUST be set to signal to the client to retry over TCP. This document reinforces that expectation.

1.1. Reserved Words

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].

2. Modifications to RFC1034

Replace

"Copy the NS RRs for the subzone into the authority section of the reply. Put whatever addresses are available into the additional section, using glue RRs if the addresses are not available from authoritative data or the cache. Go to step 4."

with

"Copy the NS RRs for the subzone into the authority section of the reply. Put whatever addresses are available into the additional section, using glue RRs if the addresses are not available from authoritative data or the cache. If glue RRs do not fit set TC=1 in the header. Go to step 4."

3. Security Considerations

This document reinforces DNS server behaviour expectations and does not introduce new security considerations.

4. IANA Considerations

There are no actions for IANA.

5. References

5.1. 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>.

5.2. Informative References

[RFC2845] Vixie, P., Gudmundsson, O., Eastlake 3rd, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", RFC 2845, DOI 10.17487/RFC2845, May 2000, <https://www.rfc-editor.org/info/rfc2845>.

[RFC2931] Eastlake 3rd, D., "DNS Request and Transaction Signatures (SIG(0)s)", RFC 2931, DOI 10.17487/RFC2931, September 2000, <https://www.rfc-editor.org/info/rfc2931>.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <https://www.rfc-editor.org/info/rfc4033>.

[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>.

[RFC6891] Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, RFC 6891, DOI 10.17487/RFC6891, April 2013, <https://www.rfc-editor.org/info/rfc6891>.

Author's Address

M. Andrews
Internet Systems Consortium
PO Box 360
Newmarket, NH 03857
US

Email: marka@isc.org