

Workgroup: DNSOP
Internet-Draft:
draft-ietf-dnsop-glue-is-not-optional-01
Updates: [1034](#) (if approved)
Published: 12 July 2021
Intended Status: Standards Track
Expires: 13 January 2022
Authors: M. Andrews S. Huque P. Wouters D. Wessels
 ISC Salesforce Aiven Verisign
Glue In DNS Referral Responses Is Not Optional

Abstract

The DNS uses glue records to allow iterative clients to find the addresses of nameservers that are contained within a delegated zone. Servers are expected to return available glue records in referrals. If message size constraints prevent the inclusion of glue records in a UDP response, the server MUST set the TC flag to inform the client that the response is incomplete, and that the client SHOULD use TCP 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 13 January 2022.

Copyright Notice

Copyright (c) 2021 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](#)
 - [1.1. Reserved Words](#)
- [2. Clarifying modifications to RFC1034](#)
- [3. Why glue is required](#)
 - [3.1. Example one: Missing glue](#)
 - [3.2. Example two: Sibling Glue from the same delegating zone](#)
 - [3.3. Example three: Cross Zone Sibling Glue](#)
 - [3.4. Promoted \(or orphaned\) glue](#)
- [4. Security Considerations](#)
- [5. IANA Considerations](#)
- [6. Normative References](#)
- [7. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

The Domain Name System (DNS) [[RFC1034](#)], [[RFC1035](#)] uses glue records to allow iterative clients to find the addresses of nameservers that are contained within a delegated zone. Glue records are added to the parent zone as part of the delegation process. Servers are expected to return available glue records in referrals. If message size constraints prevent the inclusion of glue records in a UDP response, the server MUST set the TC flag to inform the client that the response is incomplete, and that the client SHOULD use TCP to retrieve the full response. This document clarifies 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. Clarifying 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. Why glue is required

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.

3.1. Example one: Missing glue

The example below from June 2020 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. Note that at the time of this writing, this configuration has been corrected and the response correctly sets the TC=1 flag.

```
% 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

%
```

DNS responses sometimes contain optional data in the additional section. Glue records however are not optional. Several other protocol extensions, when used, are also not optional. This includes TSIG [[RFC2845](#)], OPT [[RFC6891](#)], and SIG(0) [[RFC2931](#)].

3.2. Example two: Sibling Glue from the same delegating zone

Sibling glue are glue records that are not contained in the delegating zone itself, but in another delegated zone. In many cases, these are not strictly required for resolution, since the resolver can make follow-on queries to the same zone to resolve the nameserver addresses after following the referral to the sibling zone. However, most nameserver implementations provide them as an optimization to obviate the need for extra traffic.

Here the delegating zone "test" contains 2 delegations for the subzones "bar.test" and "foo.test". The nameservers for "foo.test" consist of sibling glue for "bar.test" (ns1.bar.test and ns2.bar.test).

bar.test.	86400	IN NS	ns1.bar.test.
bar.test.	86400	IN NS	ns2.bar.test.
ns1.bar.test.	86400	IN A	192.0.1.1
ns2.bar.test.	86400	IN A	192.0.1.2
foo.test.	86400	IN NS	ns1.bar.test.
foo.test.	86400	IN NS	ns2.bar.test.

Referral responses from test for foo.test should include the sibling glue:

```
;; QUESTION SECTION:
;www.foo.test.      IN      A

;; AUTHORITY SECTION:
foo.test.          86400      IN      NS      ns1.bar.test.
foo.test.          86400      IN      NS      ns2.bar.test.

;; ADDITIONAL SECTION:
ns1.bar.test.      86400      IN      A      192.0.1.1
ns2.bar.test.      86400      IN      A      192.0.1.2
```

Question: if sibling glue from the same delegating zone does not fit into the response, should we also recommend or require that TC=1 be set?

3.3. Example three: Cross Zone Sibling Glue

Here is a more complex example of sibling glue that lives in another zone, but is required to resolve a circular dependency in the zone configuration.

example.com.	86400	IN NS	ns1.example.net.
example.com.	86400	IN NS	ns2.example.net.
ns1.example.com.	86400	IN A	192.0.1.1
ns2.example.com.	86400	IN A	192.0.1.2
example.net.	86400	IN NS	ns1.example.com.
example.net.	86400	IN NS	ns2.example.com.
ns1.example.net.	86400	IN A	198.51.100.1
ns2.example.net.	86400	IN A	198.51.100.2

3.4. Promoted (or orphaned) glue

When a zone is deleted but the parent notices that its NS glue records are required for other zones, it MAY opt to take these (now orphaned) glue records into its own zone to ensure that other zones depending on this glue are not broken. Technically, these NS records are no longer glue records, but authoritative data of the parent zone, and should be added to the DNS response similarly to regular glue records.

4. Security Considerations

This document clarifies correct DNS server behaviour and does not introduce any changes or new security considerations.

5. IANA Considerations

There are no actions for IANA.

6. 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>>.

7. 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>>.

Authors' Addresses

M. Andrews
ISC

Email: marka@isc.org

Shumon Huque
Salesforce

Email: shuque@gmail.com

Paul Wouters
Aiven

Email: paul.wouters@aiven.io

Duane Wessels
Verisign

Email: dwessels@verisign.com